

Please note: This is an AI generated transcription - there may be slight grammatical errors, spelling errors and/or misinterpretation of words.

Revamping Retirement Episode 88

Intro: Covering the ever evolving retirement plan landscape to help identify the biggest opportunities for plan sponsors, CAPTRUST presents Revamping Retirement.

Matt Patrick: Hello, and welcome to this month's episode of Revamping Retirement. I am Matt Patrick, and I am joined today by Jennifer Doss. Jennifer, great to be co-hosting with you again.

Jennifer Doss: Yeah, absolutely

Matt Patrick: We hope, everyone listening is staying cool out there as we're in the peak of what is a very warm summer on the East Coast. So I know, everyone on here has, been enjoying that so far. But, today's episode, we are gonna be discussing artificial intelligence. and feels like a topic that comes up on pretty much every episode that we do just because it's so top of mind for everything that we're talking about.

We actually haven't done a dedicated AI episode on this show in over a year, which given the pace of AI development, is an eternity. So, wanted to hit that one again. And part of what brought this up is, a couple months ago, I attended a conference, and there was a session there, from Lucas Erb, who's the founder of a group called AI Experts.

and the main topic of the session was talking around prompting for AI and how to get better at that. But woven in there throughout was just some background on how we got to where we are with AI and some good thoughts for developing an AI strategy at a corporate level. And it really started getting us thinking around, for one, a lot of plan sponsors and people on committees are business owners, business leaders who are thinking about AI strategy for their firm.

But also we expect in the near future that both plan sponsors and retirement plan participants will interact a lot more with either AI-generated content or directly with AI tools that are coming from the various vendors that they work with. So wanted to dive into that topic, and to do that with us today, we are joined by Lucas Erb.

Lucas, welcome to the show.

Lucas Erb: Matt, thank you for having me. I'm thrilled to be here and, to be the first AI podcast end-to-end that we're talking about some of the things that are happening in the industry, it's a big pressure for me, right? There's a lot of things that we could cover, but I'm more than thrilled to be the person bringing some of these conversations, especially as you mentioned, the last year of change has just been nothing short of exponential and exponential off of, what we thought was in previous years already exponential growth.

So you guys know charts and, everything is logarithmic in the space, and it's really fun to chat.

Jennifer Doss: Thank you, Lucas, maybe just to start out, we talked a little bit about what you do. Matt was describing it, but what's your background?

And tell us a little bit more about AI Experts.

Lucas Erb: Absolutely. I started my career, over 10 years ago in an IT test lab at HP. So I've been in software engineering and information for a long time. in, undergrad, I studied computer science, specialized in big data, and AI at the time, it was something that I was, experimenting with and building in personal projects and things.

But it was this technology which, was near impossible to use effectively, right? I had professors that were very well-intentioned that said, "Hey, AI is what you pursue and study if you wanna become a PhD-level researcher. And if that's not on the horizon for you, practical AI is just simply not affordable for most businesses today.

If you wanna go and work in enterprise, you develop software, right? You build end-to-end software systems and leave the AI out of it." That was, not that long ago, and less than a decade ago, that was the conversation around AI. It was this pie-in-the-sky technology. certainly business useful.

I don't wanna make it sound, too impossible because there were various forms of natural language learning and algorithms on machine learning and vision, for example, that were becoming very useful in industry. so my career then is more recently acknowledging as a part of my job at Deloitte, I was an emerging tech researcher, and in that research process, acknowledging that AI is something that very quickly from companies like OpenAI and then famously Google, they put the T into GPT.

They invented the transformer, which then allowed this whole process that we're seeing with LLMs today. we started to research these technologies and realize, hold on, there's this exponential law behind them that if you just can put more data inside these models, more information, if you can give them more to train on, they're actually gonna continue to get better and more competent.

And, seeing those charts, maybe similar to, Jeff Bezos' early inspiration on Amazon. He said, "I saw sort of the exponential growth in the internet, and I had to leave everything to jump into that space." similarly for me, I couldn't resist, the potential to learn everything about a brand-new technology that was, truly going to impact each and every one of our lives and to put my thumb on the scale.

And so that's the birth of AI Experts, aiexperts.com is, now a community of consultants and contractors that are delivering the very best across AI automation for mostly mid-size financial services firms. our standard sales process is investigating first how the business is working,

Doing really good discovery that's personal and human to understand where AI is at currently for their business, and then figuring out what the best first and second and third steps are for them to take in the space. it's very much a, a consultative approach. We don't offer a product or a platform, though we do have a productized service offering, which is, delivering the best value for some of our clients today.

Jennifer Doss: That's great. And, you just mentioned a little bit about this, that things are this logarithmic growth, right? that we see. but tell us about the progress of AI, so far to date. Like, where are we?

Lucas Erb: it is, incredible, right? Incredible mostly because, the skeptics were, and in many cases in their lens still are right about AI's limitations today, That, a year ago when we were sampling these large language models, most of the language explaining what was going on under the hood was these are next word prediction machines,

You feed in a sentence like, "I like to eat blank," and this model is just smart enough to insert pizza at the end 'cause pizza's the most common answer to that sentence, But then beyond that, it's not thinking, it's not clever. the model can't do anything real. what are you gonna make, a really nice poem?

Okay, great. Thank you for showing me that. Fun demo, but what am I gonna do for my business with this? That was just about a year ago, I would say,

maybe two years ago if you were really, looking at the cutting edge of the space. And more recently, so quickly, we've been able to grow the capabilities of these models to eliminate things that were really common in high percentage, like hallucinations.

We've been able to give these models tools access to the various things that you already do in your work day to day, right? Things like making a PowerPoint file, which were once just there's no way a large language model could do that. Now they're just a tool away that the AI can call to create a PowerPoint.

And then we've also been able to give these models, an enhanced level of reasoning where they think on a problem for a very long time. So we're spending tons and tons of these tokens. What we're doing is we're actually simulating the thought process that might happen almost instantaneously inside of a human brain.

We're able to simulate that and then allow us to reproduce real human workflows and work, and that's the big differentiator here is we're moving out of this, text-based thing. Actually, I call this compound AI. Every time you chat with your AI chat window, you're actually talking to a whole bunch of different AI algorithms and systems using a router that routes your request to the appropriate model and the appropriate response.

So you're really getting a mixture of experts' models that are replying to your question based off of what you asked, doing search, onto Google if you have a question that's related to search or something that happened yesterday as opposed to when the model was trained weeks back.

And that's why we're getting now this, you might call it PhD level intelligence in your pocket

Matt Patrick: Just you talking about the speed of improvement and referencing PowerPoints in particular. I do feel like, Jennifer, we were working with some tools and six months ago we were like, "it's just not any good at that." And now we're like, "Well, they're, looking pretty good now as we, go and try and design stuff using the tools."

Soto your point, it's like you try it once, you give it a couple months, and it's tackled that challenge.

you mentioned hallucinations in the model. Could you just quickly define what that is in case someone's not familiar with that term?

Lucas Erb: A hallucination is any time the model responds with something that is confidently correct, but it's actually incorrect when you check under the hood, AI has a tendency to, gaslight the user into believing what it's saying. And, you may have experienced this personally, there was a viral internet phenomenon about a year ago where they were testing the AI's ability to measure the number of Rs in the word strawberry,

And so the AI, with an inability to count character by character, would respond wrong to that question almost every time, there's confidently two Rs in the word strawberry." that's just one example, It's the AI being wrong about something. It also occurs when you give it large files, for example.

Let's just say you have a PDF that's 180 pages. PDFs are, to the computer's eyes, harder to read, Because you're talking about essentially a solidified image as opposed to just a text file where you can read the words right off the screen.

And so the PDF file, oftentimes we'd see the AI cite some source that was in the document, like 32% for some statistic, and then you look in the PDF, 32% doesn't exist anywhere. So where did the AI get that information? So we've actually been able to, over time, the industry I should say, not me personally, but we've improved these models to the point where those are much less likely to occur,

These multi-step workflows now are becoming more accurate because the model's able to reason and think about and cite its sources correctly as it goes. That's huge, because then you can trust it to do something that takes 30 minutes for the model that would've taken a human 10 hours to do, and you can just check the end result and make sure that it's accurate at the end of the task.

Matt Patrick: So I feel like we're dancing around this, but I guess in general, could you summarize with where we are today and the capabilities as we sit here? maybe we'll timestamp it as, we're here late July of 2026 since, these things change so quickly. but what would you say in general AI tools that are available today, what do they do well for the most part?

And then maybe where are some areas where there's still room for improvement?

Lucas Erb: Many, many places where things can be improved. But, I think briefly I'll reference a slide that we often use with clients. thinking practically about how you can use AI in your business, what are the things that, OpenAI

and Anthropic and, the labs at Google, what are they gonna improve on their models?

And then what are we as Fortune 100, Fortune 500 enterprises that are focused on using AI to deliver value for our customers without breaking the working systems we have, Without shipping, bad stuff out to the general public,

How do we build a system that's reliable using AI? So really quick, OpenAI or Anthropic, their labs are gonna continue to optimize for these, standard operating procedure thinking machines. So it's not that they're gonna know your SOP, your standard operating procedure in your business, but they're gonna optimize for, what you call alignment, the system being able to do what the human asks it to do, and being able to follow the tasks.

and if you say, "Step one, do this. Step two, do this. Step three, do that," it should be able to follow what you ask it to do every single time. So they're gonna increase the predictability. They're gonna increase the intelligence. Just naturally, they have these, benchmarks of the intelligence of the model, and now one of the, biggest ones in this space is called Humanity's Last Exam.

It's such a difficult exam, the idea being that a team of PhD researchers have gotten together to create questions that stump the AI, and, there's no chance that any single human could solve for this test. It's too complicated. It's in multiple different languages. a lot of different difficult, elements to this.

But the AI models, they're increasing their ability to actually solve these complex things. And so if you can think about the models getting smarter and better at following the job or the task that you list out to it, now on our side, what do we have to do as businesses to capitalize on that?

we have to create a lot more clear documentation about what good looks like and what needs to be done. What are the steps? What are the workflows, What can you teach this AI to do? If right now, you just ask it to do some-something simple with a one-sentence prompt, every time you're gonna end up at the same result as some other competitor of yours that asked it to do something simple with a one-sentence prompt.

It'll be the mean. It'll be the norm. But if you really clearly articulate what your business does and why you need it to do particular way, if you give it a template or an example of past experience that you've done that's relevant to your business, relevant to your clients, whatever it might be, that's where you then end up with this PhD-level thinking machine that really follows the job well.

And the more of those workflows you can begin to process and build the AI around, that's where you end up with, an engineering exercise fit for maybe the next 10-year evolution in your business using AI.

Jennifer Doss: Yeah, that's a good segue into, my next question, if you could think about where we're going. And, I think originally, you ask people like, "Oh, what's gonna happen over the next five to 10 years?" I get asked that all the time in our business, what's gonna happen with the retirement industry in the next five to 10 years?

It feels like a really tough question to ask you, Lucas. I'm not gonna lie, you're like, what's gonna happen in the next five months? just tell me what's gonna happen in the next five months. what do you think? so I'm, gonna say five to 10 years, and I have no expectation that you're gonna be able to go out that far

Lucas Erb: I briefly worked at Deloitte. I shouldn't say briefly. this was, three years of my career there was, was focused in specifically on emerging tech research with the CTO and the chief futurist of Deloitte. And, we used to do these exercises, where we'd predict roughly three to five years out.

Ten years is anyone's game, I would say, personally, right? you end up guessing on things. But we also used to joke, that futurists are closet historians, that really, if you look at the past, you look at what has happened with technology, and you map that forward into the future, it's a little bit easier to say, actually, the technology, we can see where it's headed because it's just headed in the same trajectory of all these other innovations we've already seen occur over the past five years.

It's actually going towards the same end. at the risk of answering this question too in-depth, I think the history I would look at is information broadly. And with information broadly, we graduated in simplicity primarily. That's the thing that's made AI useful today, is that it's simpler to use, that you don't need a PhD and a lab coat to use AI, that now you can just prompt it with natural language, with conversational English, whatever it might be for your native dialect, and you get an intelligent response back.

And what we saw, for example, with the spreadsheet, the advent of Excel, was that it didn't replace accounting as a profession. If anything, it made accounting a more relevant and approachable profession for the masses. It made that job much more in demand. We saw accounting becoming this thing where now everybody ubiquitously uses Excel because it just makes the job simpler, more intuitive to do.

And I think AI, we're gonna see something similar across professions where for the next five to 10 years at least, AI has increased the simplicity of access to knowledge. We have this knowledge and information abundance, and our economy is based off information. And so now, with all this abundance, with this technology at our fingertips, with this ability to not just ask questions, but also to tell the AI to go and do some task, now I think the, the human problem is, how do we level up our ability to manage these machines, to figure out with agency what to do next, to use this now simpler interaction and interface, this direct, immediate access we have to knowledge and information to add more value for the business?

I believe that much with the reference I made to Excel, that ultimately this will only lead to more abundance for human beings. But I think the short term is you absolutely gotta be using it and learning how to use it, because if not, somebody else is gonna come along, and they will.

they'll staff 10 agents to do some job, and they'll use that as a competitive differentiator for their business and leave other firms behind if they don't quickly pick this up and learn how to use it.

Jennifer Doss: we talk to a lot of businesses in our line of work, and they're all doing it, but it's not readily apparent from the outside. So a lot of people are, thinking about remapping processes and operations and all those things internally. So the way that these businesses run internally may look very different in the next five years, but maybe externally they're still providing, relatively the same services.

I mean, maybe they've, thought about delivery as well, but I don't know. I found that we're at just the first stage of this implementation. it tends to be under the covers or under the hood,

Lucas Erb: Absolutely. and there's a real engineering effort. the IT desk is already busy enough managing the myriad of complicated technologies we've introduced over previous years in the internet, migrations to the cloud are still just as relevant today as they were, some years back.

But I think now they've had this new technology ushered upon them as a mandate from the business. now you have to be an AI person too, and it's a brand new thing. And how to engineer it to work specifically for your business and not, again, generally just any business is the hard part. it's the part that's taking some firms to the stratosphere because they found some way to create sort of a self-learning loop of AI within their business, where it's taking their

data and it's taking their customers' responses to what they give them and optimizing for it, responding and learning and then building some more useful resources for their business, creating code from scratch,

Engineering new products and solutions constantly on this feedback loop where it's almost autonomous as an organization. That's sort of what we're tracking towards. Maybe that's the 10 year or the, the 15 years. We'll probably see a lot more autonomous organizations where they're learning from the data that's coming in and then using that to directly apply AI-executed work.

maybe we'll see more businesses selling the end result or the end labor of AI as opposed to just selling like, "Hey, we deliver services X, Y, or Z thing on a monthly subscription basis." Rather, "We can deliver to you a more efficient and a cheaper end result or end product because we're using AI under the hood."

frankly, it's all up in the air at this point, and so I don't wanna sound like I know too much, but, it's incredibly exciting and it is a, a whole workflow reinvention.

Matt Patrick: I really like the framing of the advancement really is in the simplicity of how you use it. I think you see a lot of that and is in line with even something from the retirement industry of like, if you go back, 20, 30, 40 years retirement plans, there was a real focus on choice. Choice in number of investments, choice of options and you'd see these really large lists of investment funds.

You're like, let me just make as many available as possible 'cause people like choice." And what you found is that was, overwhelming. And so there's been this consolidation down, and there's a direct tie to increases in participating in retirement plans with when you simplify the actual decision.

as I think about the way AI gets talked about now, where it's like it's gonna do more and more over time. if you're someone who's intimidated by where it is right now and you're like, "It's gonna do more and more, I'm just gonna get more behind and more lost." But when you're like, actually as it advances, it will become simpler to use, and that's where the power in it is gonna come, is it's, to your point, a few years ago it wasn't in natural language.

Like you'd have to be a little more advanced, and now you can interact directly with it. Or I think about, even the way like Google's embedded it where you can do a Google search the same way that you have done it since Google started, but now you'll get the AI summary. you didn't have to even change behavior to start interacting with the AI tool, And at least for me, like I always end up reading it

'cause it's right there and it gave me more than I was looking for. And those are the types of things where you're like, I didn't even think about let me use AI or didn't open up ChatGPT today to answer this question, but I've searched it the way I always have at a muscle memory and there was AI answering the question for me.

So I think the framing it as like simplicity is the future of it makes a lot of sense in terms of the practical nature of using it.

Lucas Erb: it puts you on the right side of natural, flow of human engineering and design, where if you just think about is the consumer going to do, in this case, to target it towards this audience, the participant, But then also probably the sponsors who are working, across their, i- individual company to make wise decisions.

What is the natural thing that they're gonna be doing in five years in interaction with the computer? They're probably gonna be going to their internal chatbot or maybe their, ChatGPT license or their Copilot license that they have. They're probably gonna go and ask that question, assess these plans?

Who should be our provider? What are their offerings?" "What is the most value that they can deliver?" so just from a, workflow reinvention standpoint, there's a number of things that are gonna occur almost automatically because the user's gonna gravitate towards the simplest step forward for them to solve some problem that they face or some question that they face.

And to that end, there's a number of companies doing AI engine optimization as opposed to search engine, where Google would be the search engine. How do you optimize for your company appearing at the top of a ranking on somebody searching on ChatGPT, for example? totally different problem.

Similar, in what you're doing. You're measuring whether or not you appear on these various search terms. but you're shifting where the power sits and where the, ultimate, delivery lies for what is fundamentally just a huge part of the bottom line of the business, is finding the right customers and generating, warm leads from that.

Jennifer Doss: Yeah, I think we've already seen that there's been a couple of examples where clients reached out to us or we were in a prospecting situation and they reached out and said, "Oh, I asked ChatGPT whether you were better than this person and what value you could bring. And I put in all your documents that you sent me, and I asked it for its advice."

And so I think we're already seeing that, it's a tool to help people. Now, what it's saying, I don't know. But, they became clients, so it must have been fine.

Lucas Erb: Yeah, absolutely. if you're an honest company with truly the best offering, the best value to add to the consumer, my hope would be that it allows for more meritocracy, Where information is abundant, people can make wise decisions

Jennifer Doss: I've seen plan sponsors using it as, "Here's what you told me. I put it in here, and it gave me a couple questions to ask you."

So, like from our perspective, we're seeing people ask better questions. they're asking deeper questions, because I think they feel empowered to, do that. and I hope that they're learning something from that So, I think that gets to your, hey, it's simpler to use, and so information is more ubiquitous than it was before, and so we can have a more impactful relationship because, we've got that back and forth there, and we really understand each other.

So I-- that's interesting what we're seeing there.

Matt Patrick: Lucas,

if you were putting yourself in the head space of like, I'm offering a plan, but I need to be thinking about employees on my team that will be using the retirement plan, that might be engaging with tools from third parties, what would you think about in terms of developing a strategy for either education or thinking about how do you roll that out, put proper risk controls in place?

how would you think about that when you're offering something to people on your team?

Lucas Erb: I think the engineering problem of making, like a, public-facing chatbot work well is a much harder, engineering problem, and it should be. one of the earliest examples we saw in the news of a company leveraging AI chat, to try to drive more business, more sales, it was an auto dealership, who had published a public-facing chatbot.

That auto dealership then had a customer chatting with it that requested a free car. The car was granted for free by the chatbot based off of a little bit of persuasion, and then the courts ruled that, yes, because you quoted publicly on your website that the consumer would get this free car,

You're held to this. You have to actually grant the car to the consumer. and that's just one example, it's a silly example, but you can think about the myriad different ways that, that somebody with the wrong intentions or maybe with the best of intentions just being misled on your public-facing website, based off of what the chatbot is doing.

there's a lot of things that we can do on the engineering side with AI, guardrails, harnesses, things that make evaluating the model evals, things that make it behave better from a risk standpoint before you go and do something like this. I would discourage anybody from thinking that they can just, strap a ChatGPT API directly into their front end of their website and call it a done deal, because that's a huge factor for risk.

Let's just, like, talk practically now. From, a practical standpoint, what can you be doing? Well, absolutely these things are valuable, and so the first step we see most people taking is in our back office, in our, sales team, which is directly interfacing with consumers, how do we empower people to have all the right information at their fingertips to do more?

How do we empower them to do some very long, maybe 10-hour, maybe 15-hour, maybe week-long process they were doing of reviewing two files? How do we make that something that happens almost instantaneously or automatically with an AI for them? for example, just simplifying here, we know that AI is really good at summarization.

You can feed it a massive PDF file and then say, "Extract just the most relevant pieces of information from this. Cite and source them. show me exactly which page this information came from. Give me just a summarization, an abstract, if you will." If AI is good at that, then it's also too capable of taking, say, like a contract you're considering signing with a, target customer and then, a couple of emails that you sent back and forth where they're negotiating terms.

It's capable of reviewing and comparing the two and deciding on a quick sort of level, high level, here's the information that the human needs to know before they make this decision. Here's the risk that it presents to the firm legally, financially, whatever it might be, And that's just one example.

And so we're empowering the team then to make quicker decisions. We're empowering the review to happen almost instantaneously so that then you can get back to doing what humans do best, right? Signing the actual deal, maybe, swinging a few golf clubs on the course that actually allow for the human relationship to occur and the sale to occur.

All of that should be empowered by AI agents that are doing the boring and the mundane and the rote memorization part of the job.

Jennifer Doss: Yeah. I think we have to be very careful in our industry for a number of reasons, right? We're in financial services, and so it's a highly regulated industry for very good reason. And, AI famously will not give you investment advice, but anybody who has used one of those tools knows that it's - feels like it's giving you investment advice,

and a lot of times it's so agreeable with you. It's like whatever you give it, it's like, "That's a great idea, Jennifer." It loves all of my ideas. and maybe that is the best, maybe it's not. I think that there's particularly sensitivity around the tools that participants would interact with and how far it would go and how good the questions are that it's gonna ask.

It's like, I didn't tell you this one thing about me that actually changes the whole picture, and it's like, oh, well, given that, this is what I would say, you know, and it's a completely different answer. So I think you just have to be really thoughtful about the difference between education and advice specifically, in our industry and what goes out to participants, to your point about being very careful about offering people free cars,

Like we just gotta be really thoughtful about what we put in front of participants. but I do wanna ask you about just cybersecurity, which I know is a boring topic, but a really important topic I think for AI. So, as our clients are thinking about, the cybersecurity of AI and these systems that they're using, what should they be thinking about in that regard?

What should they be asking?

Lucas Erb: So one of the first things people are thinking about, is how do we make sure we're not giving away our IP to the AI companies? and there are terms contractually if you sign an enterprise deal with ChatGPT or with, Anthropic, you'll sign terms that basically say for almost all of the model, I don't wanna say broadly speaking this applies to everything.

We should consider the individual equation you're dealing with here. But for almost all the models, you sign a contract that says, "We're not gonna train on your data. We're not gonna use your data at all. We're not gonna store it for an extended period of time," Where you're essentially just hitting the data center much like you would, borrow, say, data centers at AWS or at Azure

some of these, cloud providers, You're just hitting the data center with a question, then the AI will do what's called inference and respond with the answer, and that's it. They're not storing your data. That's important. There's exceptions to that that are really important for people to know too,

So if you have, we call this shadow IT in the industry, if you have people using, say, ChatGPT with their own personal account and you haven't authorized the users to do that, and they're going off on their own and they're pasting in some sort of pre-release financial information or something about the company into ChatGPT, by default, the consumer experience says that they're allowed to retain that information and use it to improve model behavior.

And so now you've suddenly just leaked a piece of your company's private data right out to the open market. Who knows, Maybe there's an engineer sitting over there at OpenAI somewhere that says, "I can find some alpha here, and I can trade pre-release market information when somebody pastes financials."

and so that's a real risk as well, and you gotta think about what people are using. a lot of our clients, for example, they trust the most Microsoft because they have what they call, it's marketing speak, for all the things they have under the hood, but enterprise data protection.

And that enterprise data protection says that everything that you're using with AI lives within the walls of your Microsoft 365 ecosystem. And so for a lot of financial service-related firms, that's kind of the idea, is like, "Can we just keep everything in our account management system and protect it from outside walls?"

Matt Patrick: a lot of plan sponsors have been digging more into cybersecurity because there was some guidance related to like, you can't just hire a plan provider and just call it a day. You need to have some, understanding of what their cybersecurity procedures are.

But I think your point around where's your data being used, there's already some of those themes in the industry today, and maybe it's just amplified there's gonna be more tools out there. There's gonna be more tools asking for potentially your participants' data or information about them.

When you're picking someone, when you're monitoring them ongoing, like you need to be staying up to date on not just what is your own internal firm's policy as it relates to it, but these are groups that are interacting with your plan, they're interacting with your data, they're interacting with your employees in the plan,

and you need to understand how they're using all of it and maybe that's a part of ongoing service agreements that's gonna be a lot more spelled out in the future as that becomes more and more pervasive that people are using AI tools.

Lucas Erb: Yeah, and the IP is even more valuable in an age of AI, Leaking out these things like the workflows and the processes that your business does best, the core or the pillar of how you're providing value to the consumer or towards your customer. it's absolutely something you don't want the AI model providers to have.

We've seen already Anthropic releasing a number of products and services across legal, across finance, offering AI agents that are up for the job of certain things. So the more you can protect that and you can use AI, but not just hand that over, uncaringly to various providers, that's very important.

Jennifer Doss: a lot of people listening are clients that are also managing businesses, That are not in financial services or, completely unrelated. and they're all thinking about how to incorporate AI into their companies and what kind of policies and education, resources to offer.

And, we have an AI policy, and so I think this gets back to good governance as well, which is you should be asking if there are some AI policies that for the, vendors that you're using and, again, what they're doing with your data and all those things. But I guess what best practices would you offer, business leaders in terms of thinking about rolling it out to their employees and having a policy and how to use it?

And, You probably didn't have to do that when you rolled out Excel, but this is a little different, and you probably do have to think about how you wanna put that in front of your employees.

Lucas Erb: it's a hard conversation to have, right? governance policies could be, you know, back to 180-page PDF. There could be a lot in there that needs distilling down. but I think just from broader 30,000-foot view perspective on where all of this is headed, you absolutely have to give the end user on your team access to AI.

You absolutely have to give them the ability to do now work at the speed and the efficiency that the rest of the market is at. And so if you hinder them by saying, "Ah, I'm sorry, but, our AI chat, you're only able to ask rudimentary third-grade level questions to and that's it. We're stopping there for security purposes."

yes, the most secure system is the system which allows no users into it, but that is not a winning strategy for business by any means. and so I think, our process at AI Experts, we do two things. We do superhumans and super tools. Superhumans is where we train people to use AI better.

Super tools is where we build custom workflow automations that engineer the standard operating procedures, the files, the systems, the templates of your business into how the AI behaves back for the humans that I just described in the superhumans process, The idea is how do you take your people and you empower them to use the tool better so they can get to the end result of their job, so they can do the risk review faster and better than they did before?

And, frankly, I think that's, a very difficult problem to solve because it's just human change management now that we're talking about. And from a security standpoint, it's how do you coach your team to use, say, Copilot, the tool that we've decided and chosen is the one that we're going with, as opposed to, again, their private ChatGPT license or whatever they're beginning to be familiar with.

How do we then teach them to use it, specifically with this set of workflows and prompts that we've developed, that we've tested, that are secure, that are accessing our data endpoints, specifically through our secure firewall and not some other route or direction that they go.

If you don't get on top of teaching them how to do it your way, then they're gonna get this information from some guy on YouTube who has no concern for your governance risk policy at all, and now I sound like I'm selling our service. Maybe that's the impetus behind the whole industry of AI consultation and transformation is you do have to consider this carefully and teach people well.

Matt Patrick: No, I think that makes sense. it's something we see and something we're certainly in the phase of grappling with is we have, again, Jennifer mentioned, we have an AI policy and we've embraced it, but there's still as you get further and further down, like how do you make sure everyone is, aware of what tools you have access to, what they're good at, how do you think about training?

but also like to your point, inspiring them to use it so they do get better over time. And it's hard when people are used to doing their job a certain way and that's easy and they know how to do it that way. And, maybe at first it feels different or less efficient to learn the new tool, but you gotta sell them on the idea of like, but if you get in there and learn it, you can, really add value above

and beyond what you're doing today because some of those repetitive procedures like that's where it's gonna help you.

And then you're left to do the more innovative work, and more, high value work that's there. So, I think you're hitting in all the right spots. Lucas, appreciate you running through all that. I think we covered a lot of really good ground and I think great takeaways for, anybody that's thinking about either third party AI tools that they're gonna impact their plan, but also like from a business perspective how to, think about the future of AI and how to manage that.

So appreciate you running through all that. Getting personal though, we wrap up with this question with everyone. we've talked about retirement some. We haven't heard you talk about retirement. how do you think about retirement for you personally, and what does that look like for you?

Lucas Erb: I'm still a long ways off of personal retirement. I think there's something interesting to me in, the nature of a lot of successful people, which is that they never fully retire, right? You're always finding a way to continue serving your community well, adding value in some way, working, but maybe it doesn't feel like work because you're doing it with your friends and with your family and in an industry or in a topic area that you're just obsessed with and love.

And for me, that's, the way I think about it broadly. From a financial perspective, yes, I need to be investing in a retirement account and building for the future so that when my body inevitably starts to break down, I'm still able to afford food and rent. But, from a purely philosophical standpoint, man, I never wanna retire,

it's good to continue learning and growing and, enjoying the experience of, serving other people and being needed and wanted in business. It's great.

Jennifer Doss: Well, the good news is it seems like you already do what you love, so check. And, a lot of people when they retire, they say they're gonna be a consultant. So check. You're already there, you're doing great,

Lucas Erb: Oh man, too kind. Thank you

Jennifer Doss: thank you so much for spending all this time with us. We really appreciate it, and lending us your expertise. again, this is about retirement. This is Revamping Retirement, but we do realize that AI has its, tentacles in everything, so we wanted to provide kind of a high-level overview.

So, for all of you, listeners, if you would like and subscribe if you like our podcast, and we will catch up with you next time. Thank you.

Speaker:

The discussions and opinions expressed in this podcast are those of the speaker and are subject to change without notice. This podcast is intended to be informational only. Nothing in this podcast constitutes a solicitation, investment advice, or recommendation to invest in any securities. CAPTRUST Financial Advisors is an investment advisor registered under the Investment Advisors Act of 1940.

This presentation does not contain legal, investment, or tax advice. Views expressed are those of the speakers and interviewees and not necessarily the views of CAPTRUST. Opinions are subject to change without notice.

“CAPTRUST” or “CAPTRUST Financial Advisors”) is an Investment Adviser registered under the Investment Advisers Act of 1940. However, CAPTRUST video presentations are designed to be educational and do not include individual investment advice. Opinions expressed in this video are subject to change without notice. Statistics and data have come from sources believed to be reliable but are not guaranteed to be accurate or complete. This is not a solicitation to invest in any legal, medical, tax or accounting advice. If you require such advice, you should contact the appropriate legal, accounting, or tax advisor. All publication rights reserved. None of the material in this publication may be reproduced in any form without the express written permission of CAPTRUST: 919.870.6822 © 2026 CAPTRUST Financial Advisors