Barclay Damon Live Presents *Cyber Sip*™ Season 3, Episode 9: "Developing a Robust AI Governance Plan," With Jessica Copeland

Host: **Kevin Szczepanski**, Barclay Damon

[Kevin Szczepanski]: Welcome back to *Cyber Sip*. It is the first *Cyber Sip*—and it took us till season three to get there—that I am wearing glasses, but I need them to see, so you'll never get that few minutes back. But I am really excited, on a serious note, to have our very special guest, Jessica Copeland, with us today, for those of you that don't know Jessica, and there may still be one or two…

[Jessica Copeland]: [Laughs] Several, several, several…

[Kevin]: Jessica Copeland is a partner with Bond, Schoeneck & King, our neighbors to the south here in the Avant Building here in Buffalo, New York, where we are today. Jessica is the chair of Bond's Cybersecurity and Data Privacy Practice. She is chair of the firm's AI Committee, and she is a certified information privacy professional for the United States through the ANSI Accredited International Association of Privacy Professionals. She joins us here today to talk about AI governance and use. Jessica, thank you so much for coming on *Cyber Sip*.

[Jessica]: Well, thank you for inviting me, Kevin. I look for every opportunity to talk to you about the techie things in the world, including AI. So I appreciate the opportunity.

[Kevin]: Yes, we had a nice conversation last week off camera. We're going to pick it up now on camera. So before we talk about AI governance, which is such an important and hot topic, we should probably start a few steps back and talk about AI. What is it? And particularly, what is generative AI and maybe you could also tell us a little bit about what are the benefits and risks of generative AI.

[Jessica]: Sure. So many people are talking about AI like it's this new technological development, but I always like to bring the audience to the baseline of we have been utilizing the technology of artificial intelligence for decades. The term "artificial intelligence" was actually born in 1956 through a Dartmouth conference where several professors in the mathematics department endeavored a program to identify a way for computer systems to simulate the human intelligence. And that study was looking into the neurological synopses of the brain function and a way in which they could overlay that in a computer code. Fast forward to 2022, 2023. We had the identification of generative AI, but so many steps along the way provided that capability. So artificial intelligence, as we all know and use the audience can think of it as Amazon. You shop on Amazon and you buy one product and then surprisingly you have an advertisement for a product that's similar and that is the use of artificial intelligence underlaid in the Amazon shopping experience that knows that this user in particular is going to look for this brand of sneaker and this size and this color. And so it'll promote that. So that's one example. Another example is if you use Netflix and you are a rom-com fan, well, you're not going to be surprised that you'll get recommendations for rom-com sitcoms, you know, if you've watched movies, as opposed to thrillers. So that's just a really simple end user experience identification of an example of what artificial intelligence is. It uses machine learning, so it just understands structured datasets; large volumes of datasets to predict an outcome. Generative AI is taking that technology and creating a natural language exchange so that you can enter a prompt and the model—

which you might have heard of LLM, large language models, will process that prompt, look through and assess in an incredibly fast pace large volumes of data to generate a response such that you feel like you're talking to a person. And so there's the distinction that you... it's the large language model analysis of a query that allows for natural language response.

[Kevin]: And really this particular iteration of generative AI came to us in in grand public fashion in November 2022, where we all just discovered what OpenAI's Chat GPT was, but now we're into several iterations of Chat GPT, Jessica. And as you say, it's a large language model, it's predictive, it's literally using an algorithm to predict the next word in the sentence. But in order to do that and be so good and we all assume most of our audience has experimented with it, it's got to have or train on enormous amounts of data. Tell us where does that data come from? Where, how does Chat GPT get all the data that it has?

[Jessica]: Well, Open AI has utilized data scraped through any crevice of the world wide web, of the internet that you've posted. So it can be on public facing social media sites, websites, content provider sites where, you know, certain publications or articles are freely available and it's collected all of that data to create the algorithm... or to have a repository of information to be able to generate responses in a tone that sounds human, and that's OpenAI. Meta also has generative AI resources and continues to invest in and develop their own technology. And Meta has the benefit of using all of its social media platforms. Google has the benefit of everyone that uses G Suite... that I don't know if you recall a few years ago, Google was in a little bit of trouble because of the way that they were handling or retaining content of email inboxes from Gmail accounts. And so if you think of every aspect of your day that relies on a device, whether you're online shopping, you're on social media, you're sending an email from Google or Yahoo, it is possible that that information is accessible by these companies that are developing these LLMs.

[Kevin]: Yeah, for sure. So we're sitting here, we're going to talk about why governance is so important. But before we get to that, Jessica, can you tell us a little bit about why an organization sitting here—and I know we're going to have another episode to talk about law firms in particular. So putting law firms aside, if you're an organization out there today, why might you be thinking about using either an open sourced AI or a proprietary AI source in order to do business?

[Jessica]: So the fundamental benefit of any gen AI tool is efficiency. Because of the speed at which the algorithm can process such large volumes of data, it really does permit the pace at which we conduct business, whatever business sector you're in to stay with the times, right? You have a customer that needs a response and it's 8:30 p.m. Eastern Standard Time and no one's there to answer, but you can have a chat bot on your online e-commerce site and it will allow you to have a conversation that's efficient and responsive if that LLM is trained on information related to your e-commerce. The other very easy example is I want to create a deck to promote this new product. How do... where do I start? Well, if you have Copilot through Microsoft, you can simply put a few prompts in that says, I want ten slides explaining this particular product, what the price point is and with the click of a button you will have a deck to start from. So it really is a time savings, but with every benefit comes a risk. And I know we want to talk about risk. So risks are...is the outputs, is the product that's generated accurate and is it without bias? Because you know bias... if you've heard of any other criticisms of AI, it's accuracy and bias.

[Kevin]: Right.

[Jessica]: The bias comes from what data the models trained on. So there's a great example early-stage development in Amazon, and they were feeding a tool with resumes. All of the resumes were of men between the age range of 25 to 40. So when it came to selecting a resume based on that model for interview, no women were selected. And so that's a really easy example of how bias can exist just based on the dataset that the LLM is trained on. And the accuracy part comes in where the LLM might not want to give a response. And it's odd to think of a computer that is "thinking," but there is a sense of the job of the computer is to provide

an answer, and it's the job of the computer is to provide an answer of "I don't know," is not preferable. Right. And so it leads the algorithm to create an answer that's called a hallucination. So any time you use any of these tools—again, whatever industry or sector you're in—you need to be mindful of reviewing the output for accuracy, checking sources and making sure that it is refined to be in the tone of you, you individually and your company.

[Kevin]: And confidentiality is also an issue, too. I suppose that in theory, many of us can't wait to rush to generative AI by because it is so fast and efficient. So let's say I'm in charge of writing my company's business plan, and I use an open AI source like ChatGPT to do it. If in theory, I have created Acme Corporation's development plan or business plan for a particular year and somebody else knows to look for that, or somebody else is creating a similar plan based on this large language model, whatever I have entered into this open source LLM is free to anyone who wants to or who happens to discover it, right? So there's this tremendous issue with protection of not only sensitive business information, but also potentially personally identifiable information or protected health information. Anything you put into that system is fair game for anyone else who is entering a prompt to search it.

[Jessica]: It is. So it's fair game. It results in a risk not only of the open-source generative AI tool, utilizing those prompts to train the model, which is where you were going with this, which is I might find my business plan as an output or pieces of it as an output. It also creates a scenario where if you have confidential and proprietary information, if your, you know, your R&D engineers are entering sensitive company confidential, potentially trade secret information that is no longer considered confidential because you've entered it into essentially the public domain and we'll table the legal concepts of that for the next episode. But for all intents and purposes, you should leave any confidential information entered into that chat now in the public domain. The other risk is if that entity (which OpenAI actually has had a data breach previously) has a data breach, then that company sensitive information is also available to threat actors.

[Kevin]: Mm hmm. So we're sitting here and we know we want to use it because it has such great potential for efficiency, game-changing efficiency. We know there are risks of hallucinations and bias and potential loss of confidentiality, but we're still sitting here as an organization, and we know that we want to think about whether and how best to implement it. And we've got to have some guardrails in place. So let's talk about AI governance. Jessica, what are the key components of a robust AI governance plan that every organization should be thinking about?

[Jessica]: Sure. So we talked about open source generative AI tools, but we haven't talked about closed source. We haven't talked about the companies that are developing this... these algorithms or licensing the use of the algorithm from Open AI and overlaying it with their data, contractually preventing Open AI from accessing that data, so it creates a more protected environment. And then what you're referring to, Kevin, is developing a policy for your company that looks to, how do we select these tools? Who makes the call, who's analyzing the security features? Who's analyzing the effectiveness, the accuracy? So depending on the sector that you're in, certainly the health sector, banking industry, very different attorney-client privilege concerns in the legal industry. So it varies. It's not like it's a one-size-fits-all, but every AI policy should have an elevation of, okay, I have a tool I'm interested in using. Who do I go to? Do I go to that... to the CISO? Do I go to the CIO? Do I go to the CFO, because we need to purchase these tools. And so within the policy, each organization has a different set of decision makers and operational decision makers and financial decision makers. And you need to know who within your business is going to make that business decision. But you also need to have the security posture of the tool assessed. And I view it as looking at the security stack of these companies the same way that you would any software vendor, frankly, any vendor that your company utilizes that's going to touch your data, you need to know how are they protecting it, how are they processing it, who has access to it, and when did they destroy it? When do they return it?

[Kevin]: And you don't... I always say to clients, you have to assume we hope for the best, but we plan for the worst. So if something goes wrong and there's a breach and you're talking to a regulator, you don't want to have to say to the regulator, gee, we really don't know how this product works. We don't understand its safeguards. And we didn't really understand the risks that we undertook. So. All right. So we start with it sounds to me like the threshold part of AI governance is whether we're going to use AI, what product we're going to use, how we're going to use it. It may be for certain discrete functions in the organization. Who's going to be responsible and what policies are you going to have, what guardrails you're going to have to face? Let's break those last two down. We're talking about who's going to be responsible. Is this something that we can I suppose it depends on the size of the organization, but let's suppose you have a small to medium-sized business. They're thinking about who in the organization is responsible. Do you recommend this is something we can lay off on the information technology folks should this plan with the CISO. Should there be a committee? How... what's your recommendation for best practice?

[Jessica]: It would need to be either a director that or CISO that is familiar with cybersecurity concerns because the information technology department might be aware of security risks, but they might not know exactly how deep to look for the security stack. So I would suggest that it not be the IT support. It would be more if you've had risk assessments perform, you'll say you're a small company, you don't have in-house CISO, you have a VCISO. So I would have the VCISO assess it the way again, you should have a vendor, any vendor that is providing a software solution and hosting your data or....

[Kevin]: Let me just jump in when you say VCISO you're referring to what we call a virtual CISO. So we should do a separate episode on this. But a virtual CISO is a great option for a small company or even a mid-sized company. You can't... you want the services of a chief information security officer, but you don't think that you can financially support that position. Going to a virtual CISO firm can be very helpful, but you're going to get in theory all the benefits of an actual CISO either organization-wide or for discrete parts of your organization without the burden of having to keep that person on the payroll.

[Jessica]: Correct. So ...and many organizations really don't need an FTE in that role because of the size of company, the type of data processed or collected and frankly, the financial resources. And so it allows for you to have the best of both worlds.

[Kevin]: Yeah, I agree. I think that's right. So let's say we've got the right people involved. We've got the organization has a CISO and the CISO is going to be responsible. Maybe they have an AI committee. So what's the next step? Do we need to have a policy or, or a set of procedures in place to make sure that we're using AI as designed and as intended?

[Jessica]: Yes. So and one other addition to that selection of who should approve it. If your organization has in-house counsel that is looking at general legal compliance, you certainly want to have Legal review it because there are several states where there are AI laws percolating at the bill stage. And you want to make sure that you're not running afoul of any laws that are percolating.

[Kevin]: Right? No. Good point. All right. So let's say we've done that. We have an AI committee. The CISO is the chair of the committee, and now the CISO turns to in-house counsel and says, okay, what do I have to have in place? Do I need an AI policy?

[Jessica]: Yes. And the policy. Sorry, Kevin. So the policy gets to the acceptable use of AI tools. So you might have an acceptable use policy. You should have an AI use policy that will provide examples of what you can use the tool for if approved, and what should never be used, what you should never use the tool for. You should also in the policy require training both on the policy and the way that you can develop the right prompts to elicit the most accurate and effective output from the tool.

[Kevin]: And what about… let's say you have…well, what I really want to ask you about is that's just the reality I find in organizations where I think the younger people are using AI, They have a facility with it, but they may or may not be using AI strictly in accordance with the organization's policy. How—and maybe training is the answer, but how do you how does an organization get its arms around those the young folks that want to run 100 miles an hour towards AI while at the same time guarding against the risks that's maybe running a bridge too far a step too far, and that it could put the organization at risk.

[Jessica]: But, you know, I think that you need to look at it as a wholesale restriction is never going to be effective because the first thing that those new employees are going to want to do is go on their mobile device. And if you're restricted on your Wi-Fi, they'll go off of wi fi. And so, you know, it's almost like saying telling someone not to look to the left and they're immediately going to look to the left, you cannot police it that well. But in organizations that are intelligently evaluating and responsibly looking at adoption of these tools and sharing that with your workforce, I think it will go a long way because it will empower those that are trained to want to effectively use the tools. And it will show them that you are staying with the times and that you're willing to progress, but in a responsible way as opposed to a full restriction. Or frankly, you know, a wholesale adoption, because even those young employees might be concerned that they're putting information on this open source because they understand the difference between open source and closed source and they understand risks maybe with these tools that the more senior folks don't even identify.

[Kevin]: Yeah. No, I think that's right. So before we get to some best practices, Jessica, I want to ask you, so having talked about all this, all of this makes perfect sense. I'm sure everyone listening saying, yes, this works. We should decide what we're going to do and who's going to be in charge and how we're going to train our employees. Where are we? As a practical matter, here in …we're recording in May 2024, where are we really? Are organizations running to folks like you and me to say, hey, we really want to develop a cogent AI policy? We actually just did one… finished one up this morning. I know some people are, but is that where most people are or do we need to do a better job, you and me and our colleagues out there in the cyber world to educate organizations on what they really should be doing yesterday when it comes to AI governance.

[Jessica]: Yeah, I think the more education, the more likely we'll see an uptick in adoption. I certainly have coached and strategized with clients on what they're really looking for. You know, a lot of clients say, what is this Gen AI and what do I need to do about it? And I ask, how do you want to use it? Because there are so many different solutions. Are you just talking about having a chat bot or do you want a content creator or document creator or PowerPoint creator? Do you want to create audio visual aids or do you want to use audio visual props? And so it's not necessarily that… it's funny because I'm a lawyer and I love lawyering, but I also think I'm just a business consultant when it comes to these things. And in understanding where your clients are coming from is critical. So I, I don't see hundreds and troves of clients running to adopt. I think there's a fair amount in different industries that are ahead of the curve and then there are many that are still reluctant.

[Kevin]: How much does an organization really have to understand about how that AI technology works, Jessica? I take it we don't have to all become PhDs in engineering to understand it, but we have to go beyond, well, I paid the $5.99. Now I got AI, if something goes wrong. I don't know.

[Jessica]: Well, as we talked about earlier, understanding where your data goes, once you enter a prompt, how long the company is retaining that information, identifying you, your company as the owner of that information, as opposed to, you know, providing ownership or licensing over or transferring over ownership to the entity that you're entering the prompt is one of the things that you need to know. You need to also know what is the percentage of accuracy of these outputs. That's a great question to ask some of these companies… and the smaller companies I would say you should focus on, because the larger companies, Microsoft, OpenAI, you know, they're …they'll be able to give you a great statistic and they'll be willing to do so. If you ask a startup company that's still in development phase, that might be a good way of kind of learning, early

stage, where is the accuracy of your model? But you don't need to be a PhD in computer engineering, and nor would that help you because most of the algorithms are proprietary, so you won't be able to look that deeply under the hood. But what you should do is understand how your data is being handled and processed, stored, retained, and also contractually having the company represent everything that they're saying in the sales pitch. Because, as you know, Kevin, what is said in a proposal is not always finally in the contract.

[Kevin]: All right, So Jessica, one last question because I know we're running short on time, if you had to... It sounds like we've talked about understanding the nature and, and of, of generative AI how it's going to work for your organization privacy, security, privacy of the data that you input, the security of the AI system, and also mentioning risks shifting through the contracts that you have with your AI provider. Are there any other best practices or anything else that you think we should know about AI that we don't know? We haven't discussed yet.

[Jessica]: I know we mentioned training. I don't know that I can emphasize that enough. I put this in with my cybersecurity hat. Training your workforce to understand the risks associated with these tools will be the best way to mitigate those risks.

[Kevin]: Yeah, I think that's absolutely right. Let's leave it there. Let's talk more about this on future episodes. But Jessica, thank you so much for joining us. I'm so glad you could.

[Jessica]: Same here. Thank you so much.

[Kevin]: And thanks to all of you will be back soon with another episode.

[Kevin]: The *Cyber Sip* podcast is available on barclaydamon.com, YouTube, LinkedIn, Apple Podcasts, Spotify, and Google Podcasts. Like, follow, share, and continue to listen.