Business Better (Season 4, Episode 8): Cyber Advisor – A Comparison of AI Regulatory Frameworks

Speakers: Greg Szewczyk and Paolo Sbuttoni

Steve Burkhart:

Welcome to Business Better, a podcast designed to help businesses navigate the new normal. I'm your host, Steve Burkhart. After a long career at Global Consumer Products Company, BIC, where I served as Vice President of Administration, General Counsel and Secretary, I'm now special counsel in the litigation department at Ballard Spahr, a law firm with clients across industries and throughout the country.

This episode is part of our Cyber Advisor series where we discuss emerging issues in the world of privacy and data security. Today we're joined by Paolo Sbuttoni, a partner at Foot Anstey with years of experience specializing in technology and data. We compare the AI regulatory landscape in the European Union, the United Kingdom and the United States. We provide insight on the scope of recent Acts that have been enacted in the EU and UK, and those that will be enacted on a state by state basis in the US. Leading the discussion is my Ballard Spahr colleague, Greg Szewczyk, leader of Ballard's Privacy and Data Security Group. So now let's turn the episode over to Greg.

Greg Szewczyk:

Welcome everyone to another episode of the Cyber Advisor webcast. And for those listening to the podcast, today we're going to be taking a look at the always topical issue of AI regulation. And I think we're going to have a particularly interesting conversation today because joining me is Paolo Sbuttoni, a partner at the UK firm of Foot Anstey LP.

Paulo is going to give us his analysis and perspective on the difference between the EU and UK approaches, and then I'll follow up with an analysis of the approaches from Colorado and California. So Paulo, thank you for being here today.

Paolo Sbuttoni:

Thanks very much, Greg. Really good to be here.

Greg Szewczyk:

Before we start, I want to remind everybody that we are only discussing a couple potential sources of obligations when it comes to AI. As you can see on the screen, there are numerous sources that could place requirements relating to AI use on a company, whether the company's a developer or just using third party tools.

And we've done hour long webinars on most of these individual topics, and that's not really the goal for today. Instead, we just want to give a high level overview of the emerging approaches, how they may differ, how they may overlap and a couple of international regulations and a couple of state approaches. But when a company's assessing its own compliance and governance stance, it needs to consider all of the factors that may be relevant to it.

And so with that caveat on the limited scope of today, I'm going to turn it over to Paulo to start us off.

Paolo Sbuttoni:

Thanks a lot, Greg. So over the next sort of 10, 15 minutes or so, I'm going to give you all a high level overview of the current position in the European Union and the United Kingdom on artificial intelligence regulation. So we do have currently very different approaches and stages of development as we'll discuss.

In the EU we have legal certainty with the AI Act that's about to come into force on the 1st of August 2024. But in the UK there is potentially a new and quite different direction on AI regulation as we've just recently had a government change to the Labor Party and a new prime minister in Keir Starmer and AI is a hot topic for the new government as well.

But practically speaking for clients, whether the laws currently impact your business or whether you're looking down the line at the changes, there are a lot of steps that we think you can be taking now to make sure that you're well-prepared for rules as they develop in the area, which is also what we'll cover.

So first of all, starting with the EU AI Act. Now this is the world's first attempt at a comprehensive piece of legislation governing the development, marketing and use of AI.

It was approved in May this year, and as I mentioned, we'll come into force on the 1st of August. After that a transition period begins where certain parts of the Act will come into force at different stages, and I'll share highlights of some of the key timelines.

Firstly, I wanted to mention the Act's overall aim, and that's really to ensure that AI technologies are used safely and ethically while also promoting innovation and competition. It's part of the EUs broader digital strategy to become a leader in trustworthy AI. Looking at balancing technology advancement with a core principle of European Union fundamental rights. The Act quite usefully includes an updated definition of what artificial intelligence is to clarify exactly what's being regulated.

The definition distinguishes AI systems from simpler traditional software and technology and general data processing systems. So we've got the definition to be regulated. The AI system has to operate with some autonomy and must adapt after being put into use and a key aspect is the ability to learn from the data it receives to generate outputs.

The AI Act will apply to different players across the AI distribution chain. And this is really quite wide-ranging. And a key point to take away is that essentially if a business is even simply using AI in its activities, it can potentially be caught by the Act.

So I wanted to give you a brief overview of who it applies to. Well, firstly, there's a category called AI providers. These are businesses who actively develop AI systems, so going to be mainly technology businesses and they're developers of AI systems or general purpose AI models, or even if you have one contracted and developed for you and then place it into the EU market or into service brand name.

The second category that the AI Act applies to are so called deployers, and this is one of the really broad categories, and that's those who use AI systems except if you're using them for personal use.

Third category is AI importers, and that's those who are located in the EU and put an AI system into the market. And these are businesses that from outside of the EU. So the fourth category is AI distributors. So you've got those in the supply chain other than the provider or the importer that makes AI systems available on the market.

So depending on the category that a business falls into, the obligations differ quite broadly for each. So generally you've got very strict and wide range of requirements on the providers of AI, which makes sense given that they're the ones developing the technology. And I'll go into some examples of these different categories on a later slide. I wanted to mention that there are lots of similarities between the approach of the AI Act, and I know because both Greg and I specialize broadly in technology and data protection.

There's similarities between AI Act and the GDPR and one key area where there's a similarity is the extraterritorial effects. And this will be of note for many of Ballard Spahr's US client base and any businesses that are outside the EU. So the Act. Applies to providers located outside the EU where they make an AI system or a general purpose AI model available on the EU market.

The Act also applies even where only the output generated by the AI system is used in the EU. So it's really very broad application and an area for non EU businesses to pay close attention at this early stage to see if they're caught by the Act. Also, another point to note for US businesses is that non EU providers of general purpose AI systems and higher risk AI systems, which I'll talk about shortly, are required to appoint an AI representative in the EU to act as a contact point for the EU regulators.

If you're a GDPR or data protection practitioner, this will be familiar to you because it's quite similar to the designated representative approach of the GDPR.

Now, a key and very specific part of the AI Act is the risk-based approach. Now, the EUs approach to this, and the reasons behind it is that of course, technology like AI can be beneficial or it can be quite risky or challenging depending on its use. So the approach that's taken in the law is by placing risk central to the law the legislators have a way of regulating AI that does not regulate the specific technology as such, but regulates what is done with that technology through its use.

So the principle is that the higher the risk of harm to society through the use of AI, then the stricter the rules. So the Act establishes four categories of AI systems that are based on the probability of an occurrence of harm and then the severity of that harm.

And the rules around categorizing risk are quite detailed in the Act and are, I would say, the most sophisticated example of EU risk-based approach legislation. So I'm just going to introduce them at quite a high level here in the limited of time that we have.

So the broad categories are first prohibited AI systems. Now, there are some systems that are completely banned for use in the EU under the Act. And examples of these are social scoring, compiling facial recognition databases, and real time biometric identification in public the accessible spaces, but of course subject to certain exemptions that are in the interest of the public.

The second category is higher risk AI systems. And these are those AI systems that have a high potential to cause significant harm or infringe rights. They're not prohibited, but they require strict regulation and oversight to mitigate any risks. So they include a use of AI in biometric systems or used in critical infrastructure education and employment and border control management, for example.

The third category is limited risk AI systems. Here you still need to adhere to certain safeguards, but the regulatory requirements are less stringent. So an example of a limited risk system might be a AI powered customer service chatbot used to provide automated responses to customer questions.

And the next category is mineral risk AI systems. The AI systems in this category are subject to much lighter regulatory burdens, and one example might be something that's been around for many years in basic email filters that classify messages as spam. So there's a low likelihood of negative impact to individual drives.

Now I just briefly wanted to mention penalties because that's always of interest and particular concern for businesses where the AI Act imposes quite significant fines for non-compliance. And it does have a tiered structure for the fines and the maximum fines are up to 35 million euros or 7% of total worldwide turnover, whichever's higher, for non-compliance with the strictest provisions on prohibited AI practices.

But at the other end of the scale, for smaller and medium-sized enterprises, the AI Act allows for lower scale of penalties to be applied and requires the interests of SMEs and their economic viability are taken into account when imposing fines. So quite a broad scale there, but headline figures are obviously quite concerning. Similar to when the GDPR came into force.

We have a flow chart of some of the key issues to think about at the outset and how the EU AI Acts could apply to you and your business. And I'm just briefly going to touch on some of these.

Firstly, really a key step is to work out exactly what AI systems you are using as a business currently to decide whether the Act applies first and foremost, and then whether you fall into one of the categories of a developer, a provider, or simply a deployer or user of AI.

So what we're seeing at the moment with clients currently and the kind of questions that we are getting asked is quite similar to when GDPR came into force. Businesses are looking at carrying out an AI audit similar to a data audit, working out exactly what AI they're using, creating an inventory of current AI systems and models. And this is really good practice because unless you do this AI audit, you are then not in a position to be able to actually see how the Act applies. So that's a really key step to take even if you're not a hundred percent sure which rules apply to you and would apply across any jurisdictions as well as in the EU.

So in this slide, a few key points mentioned here are what you might then take into consideration once you've done the AI audit is whether there's any exemptions apply, but really you're looking at, well, what are some of the AI systems you're using and whether they're prohibited or high risk as priority.

Again, the similarities here with when we engage with GDPR requirements are around whether certain safeguards are needed according to level of risk, according with the specific activities. And one tool that we've seen obviously used extensively in data protection is privacy impact assessments. And we're seeing clients adopting similar AI governance tools for AI impact assessments to assess risks. So looking at AI systems that are borderline high risk and documenting those and whether you are able to use them or what safeguards should be in place is going to be key for compliance.

We wanted to just highlight some of the obligations that are triggered depending on the role that you play once you've done that analysis around how the Acts could apply to you. Obviously, it's very important to work out this as the obligations greatly. So for example, providers of technology along with the GP AI providers are subject to some of the strictest requirements including preparing declarations of conformity for each high risk system.

The broader categories of user, the deployers of AI are also quite far ranging. So for example, there's a requirement to ensure input data is relevant and sufficiently representative to the extent that the user exercise some control over it. You also have obligations around monitoring the operation of high risk AI systems and reporting incidents to the provider and relevant national supervisory authorities. And a key governance point is to complete fundamental rights impact assessments before using high risk AI systems.

So again, the point there around AI impact assessments is going to be really important. Providers and deployers of certain AI systems and GP AI models are also subject to general transparency obligations. For example, use of chatbots and other automated systems where you currently can't tell if AI is being used. We'll see this being made clearer as there's a requirement to ensure that users are aware that they're interacting with AI and inform users when emotion recognition and biometric categorization systems are being used.

And then there's a requirement to label that AI generated content as such. So we'll see a lot more transparency there as this develops. So as you'll see, the law comes into force on the 1st of August, but its provisions come into effect in stages. So the fact that the laws in force 1st of August is only really the beginning or the continuation of quite a long and winding road to navigate AI in Europe and around enforcement.

A key point to remember is that it's going to be fully applicable 24 months after coming into force for new systems introduced since the law came into force. But those AI systems already on the market, which there are many, have a longer compliance deadline with some even up until 2028. So the existing systems key milestones are longer, so you have longer to comply. I wanted to flag particularly that certain parts of the law come into effect earlier.

The earliest times would be aware of is bans on prohibited AI. The strictest requirements will apply six months after the entry into four states, so in February 2025. So that's the key priority now, and that's really what we are seeing with clients is some of the bigger users of AI, the bigger technology providers are looking at that, making that kind of assessment. But that's not going to really impact the vast majority of businesses who are using more limited risk or even borderline risk AI. The compliance obligations will kick in later on, but it is really important obviously to reiterate the point to work out what AI you're using, how you're using it, and then look to categorize that before the rules do come into force, particularly the key milestone of August of next year for the vast majority of obligations.

So that was really just a whistle-stop tour through some of the key areas in the EU AI Act.

So I wanted to talk briefly about the UK approach, given it's obviously not subject to EU laws directly.

So up until now we've had the conservative government under Rishi Sunak who took a pro-innovation approach to AI and said that they wouldn't bring in an all-encompassing AI law. Instead, the approach was really around developing a non-binding cross-sector principle based framework to enable existing regulators, existing sector focused regulators. For example, the data protection regulator, the information commissioner and the financial conduct authority regulates financial services providers to set their own rules.

This approach prioritize remaining agile as new technologies emerged, but it arguably failed to deliver the regulatory certainty needed by UK firms and investors. So we saw a lot of clients taking in the UK a wait and see approach to exactly see what the sector focus regulators would do and then see how that would impact their businesses.

But now we've had a bit of a shift, as you'll know as of last month, we have a new labor government and prime minister in Keir Starmer, and a question has been, well, how will a new labor government regulate AI?

Recently we had the King speech in which the King shares priorities for the new government and lists out the types of laws that the government has promised or will look to implement over their term in power. There are no details as yet, and there wasn't actually an announcement of an AI bill, which some commentators were expecting, but the King speech did indicate a shift from the sector led approach of the previous government. So there was a statement in the speech that indicates Labor's

priority is to regulate developers of the most powerful artificial intelligence models. So it does indicate a shift from the approach of the broad application previously of sector specific focused regulation.

So there are also indications as well that you can look from labor's pre-election promises that there will be a statutory code requiring AI developers to share safety test data with the government. And labor previously also said that they would create regulatory innovation office to encourage greater speed and adaptability to new technology from regulators. So whilst there's no concrete details as yet, it is quite likely that on the horizon we'll see some form of overarching law that will bring similar legal certainty along the lines of the EU AI Acts, but with potentially more of a focus on the developers rather than the broad application to deployers and users.

So that's my coverage of the UK and I hope that's been helpful to contrast it with the approach in the EU with the EU AI Act. And I'll like to pass back to Greg. Thanks very much.

Greg Szewczyk:

Thanks Paul. Moving on to the US model. As of now, we're seeing something very similar to what we saw in the privacy world in the US, which is a state by state approach as opposed to concerted effort at the federal level.

And currently there are really two states that are driving the conversation, Colorado and California, and we're going to start with Colorado. Colorado currently has the only dedicated AI law in the United States, and while privacy laws regulate the use of AI when it's processing personal data in certain contexts, the Colorado AI Act applies to its use more generally.

So it's obviously going to be an expansion, but it's not a blanket expansion to all AI uses. Instead, it follows the same kind of risk-based model that Paolo was just talking about in the EU AI Act, albeit slightly different. Under the Colorado model, it only applies to certain defined high risk AI systems.

And this is a good chance to point out that in AI regulation we're likely going to see something similar to what we've seen in privacy regulation, which is that the same or similar terms are going to have slightly different meanings across different laws making just pure discussion of compliance for international companies a little bit difficult and tricky and requiring some specificity and attempts to really figure out where you could have overlapping compliance or where you might need to have different regimes.

The Colorado AI Act defines high risk AI systems to mean any AI system that makes or supports a substantial factor in making a decision involving education, employment, financial or lending, healthcare, housing, insurance or legal services. And this scope was a big focus of the legislative history. The original bill applied to a much broader set of activities, but compromises were necessary to get it passed.

And that's something that we are likely going to see be an issue as more states in the United States pass these types of laws is how broadly they apply. Do they only apply to high risk systems? Do they only apply to very high risk systems? The different treatment at these different data sets is going to be a big issue.

And even in Colorado, the scope may still be changing. The law doesn't go into effect until February of 2026. And when Governor Pollis signed the law, he did so with reservations noting that he did not want this law to hurt innovation in the state. And after the law was passed, Governor Pollis, the Attorney General and the Bill Spotter wrote a letter to the legislature asking them to consider amending the bill next year. So it's very possible that the scope of the law is going to change before it goes into effect, but in any event, the Attorney General will have rulemaking authority.

And like with the Colorado Privacy Act regulations, the Attorney General has made clear that he will be really listening to input from all stakeholders ranging from consumer advocates to business interests. So regardless of whether the legislature amends the bill, we may also see rules that change how we interpret areas of the law where there's discretion or flexibility.

But as of now, the law will go into effect as written. And the law applies to AI developers and deployers. It also does not have a volume threshold like US privacy laws do. So it could apply to companies that do business in Colorado not currently subject to the CPA. There are several exemptions in the AI law that contain a lot of nuances, but none of those are entity level exemptions. So unlike with many state privacy laws, federally regulated entities will still have to consider their compliance obligations under this law.

But on the plus side, it does not have a private right of action. It's enforced by the Attorney General with fines being under Colorado's unfair and Deceptive Acts Law, which means fines could be up to $20,000 per violation as to how those fines will be calculated. We're going to have to wait and see a bit.

Looking to the current obligations for developers, developers have a general duty of care and there are then certain specific requirements. And if a developer meets those requirements, it creates a rebuttable presumption that they used reasonable care, that those specific requirements are owed to different groups. With respect to obligations to other developers and deployers developers must make available a general statement describing the reasonably foreseeable uses and known harmful or inappropriate uses of subject AI systems.

They also must make available documentation disclosing things such as a high level summary of the type of data used to train the AI and known or foreseeable limitations of the system.

We also have to make available documentation describing things such as how the system was evaluated for performance and mitigation of algorithmic discrimination and intended outputs of the system.

And then the last two bullets that we have up here are actually very important and have already started to change the flow of information during contractual diligence. And that's making documentation reasonably available to assist the deployer in understanding the outputs and monitor the performance of the system for algorithmic discrimination and documentation necessary for the deployer to complete their own risk assessment.

And as I said, this has already started changing a little bit how we have approached contractual negotiations and what kind of information about the underlying AI system needs to be made available when as a deployer you are purchasing a third party tool from a developer. Developers also have to make available on their website a summary of its systems and how the developer manages known and foreseeable risks.

And finally, if the developer learns that its high risk AI system has been deployed and has caused or is reasonably likely to cause algorithmic discrimination, it must disclose that to the Attorney General and all knowing deployers or other developers within 90 days of the discovery.

And so in essence, this new requirement would create similar to a breach notification type of a law, but in the context of AI, and you can imagine how broadly some of those notices could be for large scale developers.

Moving on to deployer obligations, deploy similarly have a general obligation to use reasonable care to protect consumers from foreseeable risks, and they also get a rebuttable presumption if they meet the specified requirements.

Deployer requirements include implementing a risk management policy and program to govern their deployment of the subject AI systems, and there are specific requirements for that risk management program defined in the law.

They also would need to complete an impact assessment for subject systems, the specifics of which are again outlined. There needs to be notification to consumers about how the AI system is being used and if applicable, information regarding the right to opt out of profiling under the Colorado Privacy Act.

Deployers need to inform consumers about how to appeal adverse decisions made by subject systems, and that appeal must include human review if technically feasible.

And finally, there are some website disclosures that need to be made summarizing information such as the types of subject AI systems that are currently deployed and how the deployer manages known or foreseeable risks.

There are more important provisions including relating to disclosures made to consumers who are interacting with AI systems, and there are a lot of nuances in this law. But for purposes of today, I think the point is to try to highlight the approach and potential compliance steps.

Turning to California. California does not have a dedicated AI law, and in fact, it doesn't even have final regulations regarding automated decision-making technology involving personal data under the CPA. Those regulations are still in draft form and based on recent California Privacy Protection Agency meetings, I think it's safe to assume that even the current draft regulations will be changing as formal drafts are release and comments are received.

But I do think it's worth noting that the basic framework seems to be forming in a way that follows both the Colorado AI Act to a degree and the EU AI Act Paolo discussed a minute ago. And that basic framework involves notice opt-out rights and performing risk assessments on subject systems.

I also think it's worth pointing out that definitional differences between the CCPA regulations and the Colorado AI Act could have important impacts on their scope. Differences between automated decision-making technology and high-risk AI systems can have a big impact on whether certain types of activities are captured.

And those differences may have a particularly big impact in the employment context because as the CCPA regulations are currently framed, a lot of companies that are not in tech still use common third party tools related to employee monitoring, whether for efficiency or for safety or for timekeeping. And so it's very important to keep an eye on how the regulations are shaping up because they could have big operational impacts for companies that are going to be subject to the CCPA.

So, where does that leave us? We have different approaches, but we have some overlaps that can help drive a general compliance plan. I'll put Paolo on the spot and ask if you have a client operating across multiple jurisdictions, how do you go about advising on getting a compliance plan together at high level?

Paolo Sbuttoni:

Thanks, Greg. I think there's definitely a lot to be learned from past practices around data governance that you can use for AI governance.

And I would say going back to first principles there, first step is knowing what you're dealing with. And so I touched on this when I was talking about the EU and UK approaches, but I would say that first and foremost is really working on building an inventory of what you've got in terms of current AI systems and models and also doing an audit of IT processes to establish how AI is being used as part of these, and then working through classifying any AI models and systems based on risk level that we talked about. And then work to then design an appropriate compliance program accordingly in order to manage each of those respective risk levels. So that would be the first approach I would say. And I would also say don't be afraid to use what you've already got in place. So you've got current data governance or other risk and governance frameworks.

Look at how you can use some of those, particularly around data protection impact assessments to implement similar and AI governance framework and how you can build on that to help manage AI obligations.

I guess the other point then is well look at some technical issues and perhaps look at maybe what's out there currently in the market that can help you. I know there are new standards for security around ISO, so there's a new ISO standard for AI management systems. So looking at that and looking at the technical guidance offered by that ISO will help with managing AI risks, although it doesn't obviously guarantee that you'll automatically be compliant. It will be a good practice and help with fleshing out your governance programs.

And the other point I thought was quite interesting from Greg's discussion was around the other providers and what people are doing in terms of technology and disclosures.

One thing you can start doing actually, and we've helped some clients with is look at how some of your service providers in the market might be changing their terms and looking at their platforms and consider whether their terms or conditions are being updated to comply with EU rules currently. And you can start by asking the questions, and we've been helping some businesses that very likely use AI to communicate that with their customers and discuss how they're addressing these issues to show that they are going to be fully in compliance.

So whether you are a user of that kind of service, that's a step towards good governance to start working on those kinds of discussions with third parties that are utilizing AI that you are contracting with and help that with your overarching or contractual governance framework. So those are some of the key tips that I would think about, but that's just scratching the surface.

Greg Szewczyk:

Yeah, and I mean I completely agree with everything you just said. Leveraging that privacy compliance program is just such a key component of making sure you have that inquiry, making sure you're doing those risk assessments. Where I think in the

US things are going to be a little different or at least more difficult is unlike for the GDPR, state privacy laws have applicability thresholds. So larger companies that are already subject to laws like the Colorado Privacy Act or the CCPA or Connecticut's Privacy Law or any number of these state privacy laws, we now have 18 of them passed, are going to be able to leverage that and probably have the framework in place to know what they're doing.

But smaller companies that aren't subject to the [inaudible 00:35:25] and don't deal with as much personal data so that they just aren't subject to them. May be having to build this from a little bit of the ground up.

And I think that's going to be a challenge, but I think that the point of the Colorado AI Act doesn't go into effect until February of 2026. Taking this as an opportunity to start building that general framework, even if you know it's going to be changing on the specifics, is something that especially these companies who are not subject to privacy law should start doing because it's going to take some organizational differences, but at the end of the day, it's very much what you said of inventory, risk assessments, and making sure that you're really managing your vendors.

So again, like you said, we're scratching the surface here. That's kind of the best we can do with the amount of time we have, but it's definitely an interesting area where there's a lot of moving pieces on the regulation and an area where companies need to be paying attention.

With that, I'll thank everybody for watching or listening today. If anybody has any questions, please don't hesitate to reach out to Paulo or me. If you're watching, we have the contact information up on the screen. If you are not watching, I'd advise you, please go to the Cyber Advisor blog. You'll be able to get both of our contact information from there.

We're happy to speak with anyone on these topics or on privacy topics or just feel free to reach out. Paulo, thank you again for joining us today. Thank you everybody for watching and listening, and we'll see you next time.

Steve Burkhart:

Thanks again to Paolo Sbuttoni and Greg Szewczyk. Make sure to visit our website, www.ballardspahr.com, where you can find the latest news and guidance from our attorneys. Subscribe to the show in Apple Podcasts, Spotify, YouTube, or your favorite podcast platform. If you have any questions or suggestions for the show, please email podcast@ballardspahr.com. Stay tuned for a new episode coming soon. Thank you for listening.