

REGULATORY OVERSIGHT s04E11: AI LEGISLATION: THE STATEWIDE SPOTLIGHT

Stephen Piepgrass: Welcome to another episode of Regulatory Oversight, a podcast dedicated to delivering expert analysis on the latest developments shaping the regulatory landscape. I'm one of the hosts of the podcast, Stephen Piepgrass, and the leader of our firm's regulatory investigation, strategy, and enforcement, or RISE Practice group.

This podcast highlights insights from members of our practice group, including its nationally ranked state attorney's general practice, as well as guest commentary from industry leaders, regulatory specialists, and government officials. Our team's committed to bringing new, valuable perspectives, in-depth analysis, and practical advice from some of the foremost authorities in the regulatory field today.

Today I am pleased to share an episode from one of our firm's other podcasts, the Consumer Finance Podcast. Our State Attorney General Practice and RISE team frequently work with the Consumer Financial Services Group and the Privacy and Cyber Group. And I thought our listeners might find this topic particularly interesting because it includes a crossover with members from each of those groups.

In this episode, Chris Willis, Kim Phan and Gene Fishel delve into the evolving world of state AI legislation. They focus on Colorado's comprehensive AI law and its implications for the financial services industry. In particular, they highlight the challenges and considerations surrounding algorithmic discrimination, privacy, and cybersecurity. And finally, give their thoughts on whether other states may adopt similar legislation.

I hope you enjoy it.

[INTRO]

Chris Willis: Welcome to the Consumer Finance Podcast. I'm Chris Willis, Co-Leader of Troutman Pepper Locke's Consumer Financial Services Regulatory Practice. Today, we're going to be talking about state AI legislation.

But before we jump into that topic, let me remind you to visit and subscribe to our blogs, troutmanfinancialservices.com and consumerfinancialserviceslawmonitor.com. And don't forget about all of our other podcasts, FCRA Focus, Unauthorized Access, The Crypto Exchange, Payments Pros, and Moving the Metal. All of those are available on all popular podcast platforms. Speaking of those platforms, if you like this podcast, let us know. Leave us a review on your podcast platform of choice and let us know how we're doing.

Now, as I said, we're going to be talking today about state AI legislation. We had one of those pieces of legislation passed about a year ago and there have been other moves in that direction by a number of other states. So, it's becoming a very important emerging issue, I think, for the financial services industry.

Joining me to talk about that are two of my partners, Kim Phan, who's a partner in our privacy and cyber practice group, and Gene Fishel, who's a partner in our RISE group, which stands for Regulatory Investigation Strategy and Enforcement, which is the group that houses our nationally renowned state attorney general group.

[EPISODE]

Chris Willis: So, Kim, Gene, thanks for being on the podcast with me today.

Kim Phan: Thank you for having us.

Gene Fishel: Thank you, Chris.

Chris Willis: So, let's set the stage here because so far, we have one piece of enacted AI legislation in the country, and that's Colorado. So, would y'all talk a little bit to the audience about what Colorado's law says and requires and kind of where we stand with implementation of that law.

Kim Phan: If I may clarify, Chris, we have one comprehensive AI bill that hasn't been enacted into law. We have lots of little AI bills that have been enacted in other states, ones that are narrowly focused on, say, human resources and employee interviews or others that North Dakota has a fun one. It merely says that artificial intelligence is not a person, and so it does not have natural rights for individuals.

Chris Willis: So, I can murder Alexa and not go to jail for it in North Dakota, I suppose?

Kim Phan: Correct. We have other narrower, like Utah enacted at the end of last year, an AI disclosure law, and California has some AI disclosure laws. But Colorado really is, let's say the big kahuna, the one comprehensive artificial intelligence law that we have here in the United States. While some might compare it to what the EU has done with the EU AI Act, it is a little bit different. And the big question is, will it be copied by other states this year or coming years?

Chris Willis: Thanks for that clarification, Kim. Why don't we kind of jump into a discussion of the broad-brush status of the Colorado law because it was enacted in May of last year, but it's not effective yet. What's going on with it?

Gene Fishel: Right now, the Colorado law is scheduled to take effect in January of 2026, and pursuant to the laws and the process out there in Colorado, the attorney general is now undertaking a rulemaking process. So, really, that office is going to clarify some of the provisions and set some standards by which companies can comply and report to the attorney general's office when necessary.

Now, we don't know when exactly they are going to promulgate these rules or when these rules will come out, but that's the process we're in right now in Colorado.

Chris Willis: Okay, so I know the attorney general has been taking public comment in connection with that rulemaking, Gene. But we don't, as you said, have any timeline for when some proposed rules might come out or have any idea about what they might say when they do, right?

Gene Fishel: That's correct. I think it's important to note what Kim referenced earlier that Colorado is significant because other states are looking at what they passed. And potentially considering passing legislation, there's legislation. I think as we are recording this, 19 states have some sort of AI legislation. Several of those have legislation that's modeled on Colorado. Virginia introduced a legislation in this past General Assembly session here in 2025 that made it through both the House and the Senate and went to the governor, and it was substantially similar to Colorado's, in that it's focused on preventing algorithmic discrimination, which is the main focus of Colorado.

Virginia, the fate of the Virginia legislation was a lot different. It received virtually no Republican support in the House and Senate here in Virginia. It was carried mostly by Democrats, who control both houses, of course, in Virginia. Here we have a Republican governor and he vetoed the legislation and that it's not surprising given what the Trump administration and other Republican governors have come out and said about AI and regulating AI. They want to take more of a hands-off approach to letting AI developing companies innovate in this area.

So, not surprising it was vetoed here in Virginia, but I bring this up because companies should be aware that a Colorado AI law could be passed soon in one of these other states that are considering it.

Kim Phan: I will point out, though, that when the governor of Virginia vetoed the Virginia bill, he raised, and this was Governor Glenn Youngkin, who, as you know, is a Republican. He raised many of the same issues that the governor of Colorado, Jared Polis at the time, raised when he signed Colorado's bill into law, that he thought that there were lots of very onerous obligations being posed on businesses, potentially set up barriers to innovation for companies that might be looking for ways to utilize AI to increase efficiency and other types of benefits to not only the economy, but directly to consumers.

So, I think there is an open question. Do other states follow in Colorado's footsteps or do they look at some of this criticism and find opportunities to enhance and improve some of this legislation as it moves forward. I know that Polis specifically had hoped that through the regulatory rulemaking process or a legislative amendment process that some of the issues that he had cited with regard to the Colorado bill would get fixed over time. Now, we haven't seen that come to pass yet, but we can certainly hope.

Chris Willis: Kim, one of the states that sort of always on the top of my list to enact something like this and do it early is California. But California has had its own episode with comprehensive AI legislation. Can you tell us what happened there?

Kim Phan: Yes. Similarly, there was a comprehensive AI piece of legislation that, like the Virginia Bill, made it through the California Assembly, made it through the California Senate, and was vetoed by Governor Gavin Newsom. More targeted and narrowly tailored legislation addressing AI was moved through. He did sign into law a few new laws, one with regard to, if there is AI content, having some sort of disclosure so that consumers who are viewing something that was created by Generative AI, can identify that this was not created by a person, this was created by AI. And also, disclosures with regard to the training data that is being used to train AI with regard to its output.

So, more narrow bills have moved through, but he had the same hesitation with regard to some of the requirements that have been bandied about with regard to comprehensive AI regulation.

Gene Fishel: Kim points out the two most significant California bills that were passed, that the training using data sets, personal identifying information, and disclosure of AI content. And actually, California passed nine other AI-related laws, but they are very, very near dealing with medical situations, human resources, those sorts of things. But it is interesting, given California's history, particularly with privacy and cybersecurity, when they've always kind of been the first out there to pass these comprehensive laws. And to have that hesitation regarding AI kind of highlights really the complexity of this issue within the states.

Kim Phan: Yes. Even if it found the legislature, we're seeing efforts among the states to look at AI and specifically for the financial services industry as a counterparty to California, we also see New York being frequently front and center of some of these issues. And we've seen New York's Department of Financial Services issue a number of bulletins and other advisory pronouncements over the past year addressing how financial institutions should be thinking about AI.

Chris Willis: Yes. And didn't California DFPI put out something like that too, Kim?

Kim Phan: They did. So, we're seeing it on all fronts. I mean, I think the government, whether on the federal or state level, I think appreciates that this is not an industry-specific phenomenon. Right? There's going to be efforts to utilize AI in all aspects of various industries and all aspects within an individual company, whether or not it is taking meeting notes at your board meetings, whether or not it is during the HR hiring process, during implementation, during your data security checkups, these are all aspects where AI could find really useful solutions and tools for companies and I will flag that across all of these bills that we're seeing.

Financial services is absolutely in the mix. When they talk about AI, when they talk about the potential high-risk implications of AI. Financial services and decisioning based on AI is pretty much always included. I can't think of a bill yet that hasn't flagged the financial consequences of utilizing AI's potential risk.

Chris Willis: Sure. Let's talk about the content of these laws. The Colorado one and the more general ones that are modeled after it. What are the primary hot-button issues that those laws are trying to address in the implementation of AI?

Gene Fishel: Right off the top of the list has to be what's termed algorithmic discrimination. That's the gist of the Colorado law. That is what a lot of the legislation you're seeing through the state wants to address. And that's simply trying to prevent an AI system from engaging in bias or really producing results that are biased or unfair or discriminate in some sense. So, this type of legislation, just like Colorado, it requires that companies implement risk management

systems that they conduct impact assessments that they're monitoring the inputs going into the system so that these inputs aren't biased in some way.

I would also just point out that bias and discrimination beyond AI-specific laws are really on the radars of state attorneys general around the country. And over the past year, a handful of state attorneys general have come out and said that utilizing AI could potentially violate anti-discrimination laws if the outputs are biased, if the impacts even are unfair, regardless of the intent of the company deploying the AI system or developing the AI system.

So, bison discrimination seems to be one of the primary considerations. There are others, but I would raise that as really the top issue. And that's based on what the Colorado law and what Virginia tried to pass here also. Kim, do you have any other thoughts on that?

Kim Phan: Yes. Bison discrimination is certainly a hot button topic. The intellectual property rights with regard to the underlying data, you also hear that as a frequent topic of discussion where I come in at the part of the privacy and cyber group is with regard to the privacy and cybersecurity considerations related to AI. There's heavy concerns, especially in the financial services industry, and this intersects as well with some of the new state privacy laws that we're seeing that impose obligations on companies to narrowly utilize data only for the purposes for which it was originally collected.

So, if you collected consumer information without intending to use it to train your artificial intelligence systems, but now you want to use previously collected information to do that trainees, that permissible under some of these state laws, do you need to provide opportunity for consumers to either consent or opt out of those uses? So, some of these privacy struggles really hit up against the reality that these are all very, very new state laws. So, not only the privacy laws, but the AI laws.

Then you think about security. I think everyone appreciates the reality that artificial intelligence, whether or not being utilized by companies affirmatively, have to be thinking about using artificial intelligence defensively, because the bad guys 100% are using artificial intelligence to engage in sophisticated phishing, vishing, voice, quishing with QR codes, and other types of AI-

powered attacks against companies. So, the reality is the best way to fight AI is with AI. So, companies need to be thinking about that as well.

It's a hard position for companies to be thinking about because the reality is when you're thinking about bias and discrimination, what if you get a false positive, right? Chris Willis is potentially a fraudster, and he's a security risk, so we're not going to allow him access to our financial platform. Is that a problem, right? And how do you test and monitor for that?

Chris Willis: Gene and Kim, here's the interesting thing to me. When I see this AI legislation and using, again, Colorado as an example, it imposes this duty on the users of an AI model, which is basically anything that uses a computer, honestly, to take reasonable steps to prevent algorithmic discrimination. It doesn't say how to do that. But more importantly, how would a state regulator know whether somebody did that or not? What is the path for enforcement for an obligation like that under a state AI law, like Colorado or some that might pass in another state?

Kim Phan: Yes, I will say Colorado has added in this particular law a lot of process with regard to who has to notify who and what information they have to provide. It is on multiple parties. So, the developers of AI have obligations to provide certain information to the deployers of AI about how their AI was developed and tested. They have to provide publicly available information to consumers with regard to that AI and they have to provide notice to the State Attorney General about their AI tools and solutions.

The deployers of AI solutions also have similar obligations with regard to notifying developers of issues that they encounter when deploying AI. They also have notification obligations to both consumers and the Attorney General. Well, so we're seeing, specifically in Colorado, a path whereby companies let the state enforcers, like the Attorney General know that they are deploying these tools so that the state AG, to the extent that they want to, can keep an eye on that activity.

But we've seen, alternatively, in some states, legislation that takes a different path, where instead of just providing notice of the use of these AI tools and solutions, it actually requires the provision of the underlying testing data, evidence that their algorithmic machine learning models

are not actually resulting in bias or discrimination. So, we're seeing a range of onerous requirements that start with Colorado could go up from there.

Chris Willis: Yes. And that latter formulation, Kim, I remember seeing in one piece of proposed legislation a year or so ago, the idea that the users of an AI system would have to do the discrimination testing and then provide the results of that to a state regulator seems to me to be kind of the nightmare scenario for industry and that it creates a whole lot of work for industry and makes sort of targeted enforcement by the state regulators very, very easy. It would be my thinking.

Kim Phan: I agree with that. To a certain extent, it is argued that they wouldn't have the ability to have that insight unless they were basically being spoon-fed this data by the developers and deployers of this information. But regulators that have been able to operate for years, Attorney General, intake different concerns from consumers in a lot of different ways, right? Whether or not it's consumer complaints followed with the AG's office or other vehicles by which they can identify potential problems that are arising. So, is that extra layer of reporting, of notice, how burdensome is that, and how much will it slow the ability of companies to really be able to effectively deploy these tools in a timely manner?

Chris Willis: Yes. And Gene, to Kim's point, that AGs have always been able to operate without this kind of reporting, there has actually already been some AI-related state AG enforcement activity, hasn't there?

Gene Fishel: There has, notably from Texas last year. And Texas, of course, does not have an AI-specific law, but Texas proceeded against a medical services company, Pieces Technology, under the Texas, basically their Consumer Protection Act. And what allegedly Pieces was doing, Pieces provides medical charts and data to physicians and nurses in medical facilities, Pieces apparently was advertising the fact that their AI system had an extremely low hallucination rate and of course, hallucination rate are basically false results, right? The incidence of false results that come out of using AI system.

Texas did some traditional investigation using subpoenas, court orders under their Consumer Protection Act, and launched an investigation and inquiry into those acts, reached a settlement with Pieces, technology under the Consumer Protection Act, alleging that their advertisement, that their system was so accurate, was actually false and misleading under the consumer protection statute. So, they reached a settlement where now Pieces has to actually report in the future to the AG's office when they utilize an AI system in certain ways.

But this is an important case because beyond the Texas action, which was the first action of its kind relative to a Generative AI system. Several other states, including Massachusetts, New Jersey, Oregon, have all warned over the past year they've issued formal guidance warning that the use of AI systems could potentially violate consumer protection laws, particularly when a company may misrepresent how the AI systems being used, how data is being used within the system, what kind of data is used to train a model, or also, as we saw in Texas, the accuracy of the AI system.

Not only consumer protection, as Kim referenced earlier, state AGs are saying that companies deploying AI may run afoul of our privacy laws. Of course, now we have 19 states that have passed comprehensive consumer privacy laws. Companies need to be providing adequate notice of how they're using consumer personal identifying information within AI systems. And also, they need to be effectuating because under these Comprehensive Consumer Privacy Laws, consumers have more control over their data. They can ask that data be deleted. They can ask that it be corrected or various different rights now that are given to consumers.

How is the company able to effectuate these consumer requests and particularly what concerns me having looked at this in the past is specifically deletion requests of data. If data is being entered into an AI system or an AI system is touching consumer data, there are issues with particular systems in forgetting things. I mean, part of the magic of Generative AI as it consumes all this information. It's coalescing information and producing results. Well, how is your company removing consumer data from that AI system once a consumer says, "I don't want you to use my data in that." That's a concern for state regulators and something companies need to watch out for.

Kim Phan: Yes. I think many other areas that we're seeing, the states are more than happy to step up and fill the perceived void that's being left from the federal government, and we've already seen that President Trump has issued his executive order on AI, essentially rescinding all of President Biden's prior executive order on AI. We've seen other areas of the federal government that sort of pull back on some of their AI. DHS no longer has their artificial intelligence advisory board. Some of the other agencies are also pulling back on some of their AI regulatory initiatives.

So, I think we can only expect that there will be increased state activity in this area in the year to come.

Chris Willis: So, given that fact, Kim, I mean, obviously companies in the industry need to start preparing themselves for what they might need to do in light of laws like this, like the Colorado one and others that may be to come. So, what are some best practice suggestions that you and Gene would suggest for companies to be ready for this evolution in the regulatory landscape?

Kim Phan: I would suggest that companies really just think about AI the same way they would think about any other new technology, the same way when personal assistance came out, when email first came out. Most companies will have something called an acceptable use policy in which they are determining what are the specific use cases for a new type of technology. They should be thinking about artificial intelligence, just a new type of technology. What are the effective governance structures we need around this new technology? What type of training and testing do we need amongst our employees before we deploy live some of these tools into our production environment?

Once it is deployed in a post-deployment world, how are we measuring the success of these new technologies, whether that's basic artificial intelligence or Generative AI? How are we monitoring for potential harms? And what corrective actions do we need or should we be taking to address things like ethical implications of the use of AI and other types of potential harms, whether or not they're direct harms like a consumer being denied a loan, or whether or not they are more theoretical harms, which are some of the concerns that I think some of the consumer advocates have raised.

But thinking about it the same way any financial institution would think about their compliance management system. They should be thinking about how they should structure something similarly around the deployment of AI.

Gene Fishel: Yes, and just to add on to Kim's thoughts there, I completely agree. I think maybe of primary importance is the governance structure, and who in your organization is touching the AI system? Are you the one in control? You have a third party in control? Who in that third party has access to the data? Is it properly segmented? Are you using some sort of open-source AI system, which opens up a lot of problems? But conducting impact assessments, having a risk management program, effectuating if you're dealing with consumers, effectuating data rights requests, all of these are going to go a long way towards compliance and will go a long way when, if ever, your company comes under regulatory scrutiny, you can point to these policies and these procedures regarding AI.

I think there's a lot of unknown out there and there's some fear among consumers, and even regulators who don't fully know the capacity of these AI systems. But Kim's advice about treating it like any other new technology and taking these thorough steps should help go a long way.

Kim Phan: And needless to say, as evidenced by this conversation, they should have some functionality to monitor for these changing legal and regulatory expectations with regard to AI.

Gene Fishel: Indeed, consulting competent counsel, outside counsel to help with these because this is an ever-changing landscape, really, on a month-by-month basis at this point.

Chris Willis: Well, Kim, Gene, thank you very much for this discussion and I feel so lucky to have colleagues like you who are so closely monitoring the progress of this state AI legislation as well as the related state attorney general and other regulator efforts toward AI as they affect our financial institution clients. So, thank you both for being on the podcast today.

And thanks to our audience for listening as well. Don't forget to visit and subscribe to our blogs, troutmanfinancialservices.com and consumerfinancialserviceslawmonitor.com. And while you're at it, why not visit us on the web at troutman.com. You can add yourself to our Consumer Financial Services email list and be notified of the alerts and advisories that we put out, as well as get invitations to our industry-only webinars that we host from time to time. And of course, stay tuned for a great new episode of this podcast every Thursday afternoon. Thank you all for listening.

Copyright, Troutman Pepper Locke LLP. These recorded materials are designed for educational purposes only. This podcast is not legal advice and does not create an attorney-client relationship. The views and opinions expressed in this podcast are solely those of the individual participants. Troutman does not make any representations or warranties, express or implied, regarding the contents of this podcast. Information on previous case results does not guarantee a similar future result. Users of this podcast may save and use the podcast only for personal or other non-commercial, educational purposes. No other use, including, without limitation, reproduction, retransmission or editing of this podcast may be made without the prior written permission of Troutman Pepper Locke. If you have any questions, please contact us at troutman.com.