Proskauer

AI Bias Audits

The Proskauer Brief on **July 31, 2025**

In this episode of The Proskauer Brief partner Guy Brenner, who leads Proskauer's D.C. Labor & Employment practice and is head of the Government Contractor Compliance Group, and Jonathan Slowik, senior counsel, Labor & Employment, in the firm's Los Angeles office, discuss laws requiring employers who use artificial intelligence (AI) to conduct bias audits and impact assessments to root out potential algorithmic discrimination. Guy and Jonathan discuss two high profile laws applicable to employers in New York City and Colorado, what employers need to do to ensure compliance, and the possible proliferation of bias audit requirements in other jurisdictions.

Guy Brenner: Welcome to The Proskauer Brief: Hot Topics in Labor and Employment Law. I'm Guy Brenner, a partner in our employment litigation and counseling group based in Washington, DC, and I'm joined by my colleague Jonathan Slowik, senior counsel in our practice group based in Los Angeles. Today's episode is about laws requiring employers to conduct bias audit impact assessments to ensure their artificial intelligence tools or AI tools are not resulting in illegal discrimination. Jonathan, thank you for joining me today.

Jonathan Slowik: Very happy to be here.

Guy Brenner: So, Jonathan, can you begin by giving us a quick primer on the laws we're talking about and the AI tools at issue?

Jonathan Slowik: Sure. So New York City's Local Law 144, which took effect in 2023, and Colorado's Artificial Intelligence Act, which will take effect in February 2026, target AI tools used by employers to make or assist in making decisions in the employment context. The focus of these laws is on technologies that assist employers with conducting tasks that are otherwise intensive or costly, like resume screening and performance management. These tools have become quite popular, as employers have been inundated with resumes for job openings that applicants can submit on, and various tools, especially online, that facilitate the application process. Given the volume of these applications, and in many cases, there's just simply no way for employers to review all the resumes. And a lot of those resumes are from unqualified candidates. So AI tools can quickly review all of those resumes or applications and identify promising candidates. And this process can save companies significant time and resources.

Guy Brenner: Yeah, that efficiency you mentioned is precisely why these tools are so attractive for employers. So why would New York City and Colorado's laws want to target this technology?

Jonathan Slowik: So it boils down to what is often referred to as algorithmic bias. We've discussed this in previous episodes of this podcast, and for a full explanation of what algorithmic bias is in the eye context specifically, I encourage listeners to tune in to those episodes, but at a high level. Basically, these AI tools are not infallible. In most cases, the tools are only as good as the data that was used to train them. And if that data is problematic in some way, for example, if it's not representative of the population applying for the jobs, if it contains biases, or if there's simply not enough data, it's possible for the tool to disfavor certain groups of applicants or employees unintentionally. So, in a worst-case scenario, the tool might have been unintentionally trained in a way that produces biased outcomes on the basis of protected characteristics like race, sex, or religion.

Guy Brenner: Right. And as we've talked about in our prior episodes, what you put into AI, you know, determines what you get out of it. And if the data used to train the AI reflects biases, even unintentionally, you can run into some trouble. So, what do these laws require employers to do to address this bias issue?

Jonathan Slowik: So in each, in their own bespoke way, these laws basically require employers to test the tools for bias. These requirements to test for bias apply when tools are used for certain purposes or deployed in certain ways. The New York City law refers to this this this testing process as conducting a bias audit. The Colorado law uses slightly different terminology or refers to them as impact assessments.

Jonathan Slowik: But, these are really referring to practically the same thing. The idea is that employers need to test these tools to ensure that they're operating as intended, and not creating unintentionally biased or discriminatory outcomes.

Guy Brenner: All right. So let's take a look at New York City's law. What do the bias audits look like there? What's required?

Jonathan Slowik: So, New York City requires bias audits of what's referred to in the ordinance as “automated employment decision tools.” That term, automated employment decision tools is defined to include AI or any other algorithmic tool that quote, “issues simplified output, including a score, classification or recommendation that is used to substantially assist or replace discretionary decision making or making employment decisions that impact natural persons.” Now I space that language substantially. “Assist” might seem really broad and cover a wide range of use cases, but there's an implementing rule that narrows the scope of that that language significantly under this rule, to substantially assist or replace discretionary decision making means either one to rely solely on a simplified output, with no other factors considered to use a simplified output as one of a set of criteria where the simplified output is weighted more than any other criterion in the set, or three to use the simplified output to overrule conclusions derived from other factors, including human decision making.

Guy Brenner: All right. So, it sounds like the test is does the AI essentially make the decision or is a significant factor in in making the decision. And so, in those circumstances, bias audits are triggered under New York City's law. So, Jonathan, all three of those circumstances that trigger the bias audit requirements in New York City, there are situations where the employer has meaningfully delegated its discretion to the AI tool. Is that right?

Jonathan Slowik: Yeah. These are all situations where the tool is more or less making the decision or it's the tool is so important that human decision makers are largely deferring to its output when they're making decisions. On the other hand, if the tool is merely providing one data point among several that a human decision maker is taking into account, then these bias audit rules don't apply under the New York City law.

Guy Brenner: Jonathan, surely there are some employers who are using automated tools that trigger the bias audit requirement under the New York City ordinance. What do rules require in cases where they apply?

Jonathan Slowik: There's a detailed set of requirements that set forth in the law. But in summary, the employer has to conduct a bias audit annually and publish the results of that audit on their website. Those are that's the main kind of the main summary of what the employer needs to do. The bias audit needs to be performed by an independent auditor, and it needs to examine what the law refers to as a selection rate. So this is the rate at which applicants are selected for an employment decision, like hire a promotion. Or if there's no, you know, selection rate, so to speak. It must examine what’s called the scoring rate. So the rate at which an applicant or an employee receives a score above the samples, median score. And these rates are calculated and examined for individual based on their race, ethnicity and sex. And then also there's an intersectionality component. So also various combinations of those characteristics. Once those rates are calculated they're compared to the rates for the highest selected or highest scoring candidates. And that comparison creates what the law calls an impact ratio.

Guy Brenner: So that was like a pretty detailed set of requirements to look under the hood of these tools and see how they're operating. So let's shift gears a bit. Let's talk about what's required under Colorado's law.

Jonathan Slowik: There are some important similarities between Colorado and New York City. The Colorado Artificial Intelligence Act, like the New York City law, also requires impact assessments. To be completed on an annual basis. There's an additional trigger, though. It also requires these assessments to be completed within 90 days of any intentional or substantial modification to the AI tool. So, like New York City, the nuts and bolts of this process are laid out in a lot of detail. In Colorado, the impacts assessment needs to include a number of statements, metrics and analyzes. Some of the key ones are one, a clear statement disclosing the purpose intended use, and benefits of the AI tool to and analysis of whether the AI tool poses any known or reasonably foreseeable risks of algorithmic discrimination. The nature of that risk, and the steps taken to mitigate that risk. And a description of the post-deployment monitoring and user safeguards provided, including the oversight, use and learning process the employer will use to address any issues that arise.

Guy Brenner: These bias audits and impact assessments sound pretty rigorous, might sound more rigorous than the actual benefit you're supposed to gather from using the tool as a practical matter, right? What if the employer just doesn't have the historical data needed that appears to be fueling these bias audits?

Jonathan Slowik: The New York City law speaks to this explicitly. It allows first time users of the tools to rely on data from other employers who use the same tool or on test data. If there's insufficient data about the employer's own track record. Using the tool. If the even if it's an employer, is not a first-time user. So, they've used the tool in the past, there's still a provision that allows employers to rely on data that's not limited only to that organization if certain requirements are met, the most important one of which is that the employer needs to at least also make its data available for that purpose.

Guy Brenner: So, Jonathan, what happens if after conducting one of these analyzes, it appears that there potentially was some sort of disparate impact occurring as it relates to a protected group?

Jonathan Slowik: So the New York City law doesn't actually say so. There's no explicit requirement under the New York City law that the employer do anything necessarily. Now, that being said, of course, there are existing laws against discrimination in employment that are generally broad enough to encompass decisions made or insisted by an AI tool. So it's no defense to say that the I did it. The practical effect of that is if there is a bias audit revealing evidence of algorithmic discrimination, that's going to be a strong incentive for the employer to take some kind of action, whether that's changing how they use the tool or just discontinuing use of the tool completely. So interestingly enough, although California doesn't require bias audits, regulators there have taken a similar approach. There are regulations that are expected to become final this summer that state that it's, quote, relevant to a discrimination claim or to the employer's defense if the employer has taken action in response to anti-bias testing. On the other hand, Colorado's law does have some explicit requirements placed on the employer. If an impact assessment does uncover evidence of a disparate impact.

So if an impact assessment reveals algorithmic discrimination that a Colorado employer needs to notify the state attorney general within 90 days, if the Colorado Attorney General brings an enforcement action against the employer, the employer will have a rebuttable presumption that it used a reasonable care. If it has a risk management policy that complies with certain statutory requirements, and if it has made that timely disclosure to the Attorney General.

Guy Brenner: Jonathan, as we discussed and you mentioned earlier, not all applicants are created equal. And part of the reason these tools are so attractive is because there are so many applicants applying who are not qualified for the job. Okay, we recognized that because not all applicants are really equally qualified or even qualified at all. There should be a series of standards to determine who and should not be included as an applicant for purposes of the analysis. For example, many candidates, as we discussed, don't meet the basic qualifications for the position. Under this guidance, such candidates would be removed from the analysis entirely because there was no way that individual would have been hired or could have been hired, irrespective of their race or sex.

Similarly, if a candidate withdrew, they should be removed from the analysis as well because there's no way that you could hire them. They withdrew themselves, they said they're not interested anyway, so that helped make the analysis more accurate, because you were only considering those who were qualified and wanted the job. I don't know if New York City or Colorado thought about these issues, and I'm unaware of any guidance they provided to employers. So employers in those jurisdictions need to think about how they're going to conduct these audits before employing the AI tools.

For example, do they need to collect race and sex information from their applicants to have data to conduct the analysis? How will they safeguard that information? Do they need an applicant tracking system like government contractors used to have? That includes codes that identify candidates that can be removed for audit purposes. Who needs to be trained on how to properly code those candidates? I mean, these and other questions, these questions and implementing any desired sort of outcomes requires planning, time, resources, and training. This isn't something that employers can just do on the fly. They really need to think about it.

Jonathan Slowik: Yeah, those are really good points and it's not clear. I think to me at least, that the states, considering these kinds of bias audit laws or New York City, have solely thought those questions through, I think moving forward, it's worth keeping an eye on Colorado. Considering that law does not even take effect until February of 2026 and maybe amended before then. When he signed a bill into law in 2024, the Colorado governor, Jared Polis, issued a really remarkable signing statement. He signed the legislation, the law, but expressed reservations about how broad reaching it was. And he urged stakeholders to amend it to I think he said fine tune the law before it takes effect. And then more recently, in May of this year, the governor, Senator Michael Bennet and some other prominent Colorado elected representatives, as you will, issued a letter urging the legislature to delay implementation of law even further until January of 2027, to give all the stakeholders time to make those reforms that that Governor Polis, it called for when he initially signed the bill.

Guy Brenner: Well, it sounds like people recognize that, that these laws have significant burdens that need to be thought through and prepared for. And we'll see if the Colorado legislature listens. I think it's also important for employers to consider is the publication requirement, particularly under New York City's laws. Certainly, no employer wants to publish the results of bias on it, reflecting that there may have been a problem with their area tool, or that there was some kind of bias in the hiring process. That's likely to spur some unfavorable results, potentially litigation. Also, I worry that just because data suggests an adverse impact, that doesn't mean that discrimination occurred. Sometimes the race or sex of a candidate chosen doesn't reflect the predominant groups in the applicant pool. That won't be reflected in the bias analysis and will likely lead to confusion and unnecessary disputes.

Again, this is something I don't think New York City or Colorado have thought about. The last thing I'll say is that, as other jurisdictions look to regulate AI tools in the employment context, I fear that they're likely going to look at New York City and Colorado as models and likely incorporate some form of audit process into their laws. So, if you're an employer outside of Colorado or New York City, this may not impact you yet, but it may soon.

Jonathan Slowik: Yeah, I think a lot of jurisdictions may be looking to you, since the Colorado law was a statewide law to require this kind of process. And since it's still a work in progress itself, I think maybe other jurisdictions may be looking to Colorado as a bellwether in this regard, but other jurisdictions have not been sitting on the sidelines. The New York State Senate has considered a bill that would require developers and deploys of so-called high risk AI systems, which could include employers to require these developers and employers to conduct audits as well to evaluate these kinds of risks without algorithmic discrimination. As lawmakers continue to grapple with these tools available to employers, we can assume that, you know, if Congress doesn't say they can't, that jurisdictions will continue to regulate in this space more and more as, you know, as the tools become more widespread. And so, these kinds of mandatory bias audits could be coming to more places. And we'll just need to wait and see.

Guy Brenner: Very interesting, Jonathan and all the more reason for employers to remain vigilant and to keep a close eye on these developments. This was a great conversation. And Jonathan, I look forward to having many more of these conversations and podcasts in the future. Thank you for joining us on The Proskauer Brief today. Also, be sure to follow us on Apple Podcasts, YouTube Music, and Spotify so you can stay on top of the latest hot topics in labor and employment law.

Related Professionals

- **Guy Brenner**
  Partner
- **Jonathan P. Slowik**
  Partner