

# New York Releases Proposed Rules for Its AI Employment Law

September 27, 2022

In July, we wrote about New York City's [Automated Employment Decision Tool Law \(the "AEDT Law" or the "Law"\)](#), which requires employers to conduct an independent bias audit of their AI employment tools by January 1, 2023. On September 23, 2022, New York City's Department of Consumer and Worker Protection ("DCWP") released [proposed rules](#) (the "Proposed Rules") that would implement the Law, and which clear up some, but not all, of [the Law's ambiguities](#). The Proposed Rules are now subject to a comment period ending on the day of DCWP's public hearing, which is October 24, 2022. DCWP will then publish the final rules at some point thereafter, hopefully before the AEDT Law comes into effect.

---

## Limiting the Scope of an Automated Decision Tool

Under the AEDT Law, an "automated employment decision tool" ("AEDT") includes:

- any computational process derived from machine learning, statistical modeling, data analytics or artificial intelligence;
- that issues simplified output, including a score, classification or recommendation; and
- that substantially assists employment decisions being made by humans.

The Proposed Rules appear to narrow this definition in two ways.

First, the Proposed Rules would limit the scope of the Law to instances where the AEDT's output is the most important factor in the employment decision. The Proposed Rules provide that the phrase "to substantially assist or replace discretionary decision making" means "to rely solely on a simplified output (score, tag, classification, ranking, etc.), *with no other factors considered*, or to use a simplified output as one of a set of criteria where the output is *weighted more than any other criterion in the set*, or to use a simplified output to *overrule or modify conclusions* derived from other factors including

human decision-making” (emphasis added). So, if an AEDT creates an output that is just one data point among a number of factors leading to a decision, under the Proposed Rules, that tool would be outside the scope of the AEDT, unless that output was weighed more than any other criterion or used to overrule a human decision.

Second, the Proposed Rules may limit the Law’s application to complex models. The Proposed Rules define “machine learning, statistical modeling, data analytics, or artificial intelligence” as a group of mathematical, computer-based techniques (i) that generate a prediction or a classification; (ii) for which a computer at least in part identifies the inputs, the relative importance placed on those inputs, and other parameters for the models in order to improve the accuracy of the prediction or classification; and (iii) for which the inputs and parameters are refined through cross-validation or by using training and testing data. This definition seems to apply only to sophisticated models, and may exclude simpler tools that only conduct a linear analysis of inputs to reach a defined output. But exactly how the New York City Corporation Counsel (which is responsible for the AEDT Law’s enforcement) will interpret the provision remains to be seen.

---

## The Requirement for the Bias Audit

### Protected Categories

The AEDT Law requires “bias audits” to be conducted for AEDT tools, which includes but is not limited to assessing the tool’s disparate impact on “persons of any component 1 category to be reported by employers pursuant to subsection (c) of section 2000e-8 of title 42 of the United States Code as specified in C.F.R. Title 29, part 1602.7.” The Proposed Rules specify that this consists of the categories designated on the Equal Employment Opportunity Commission Employer Information Report EEO-1 (the “EEO-1 Categories”), which covers binary gender (male or female), ethnicity (Hispanic or Latino or non-Hispanic or Latino), and race (Black or African American, Native Hawaiian or Other Pacific Islander, Asian, Native Hawaiian or Other Pacific Islander, White, or Two or More Races).

The Proposed Rules, through their examples, suggest that it is not sufficient to assess disparate impact of an AEDT for each EEO-1 Category in isolation. Rather, it appears that employers must consider intersectionality, and conduct the bias audit for each EEO-1 Category jointly as follows:

<b>Hispanic or Latino</b>	<b>Male</b>	
	<b>Female</b>	
<b>Non-Hispanic or Latino</b>	<b>Male</b>	<b>White</b>
		<b>Black or African American</b>
		<b>Native Hawaiian or Pacific Islander</b>
		<b>Asian</b>
		<b>Native American or Alaska Native</b>
		<b>Two or more races</b>
	<b>Female</b>	<b>White</b>
		<b>Black or African American</b>
		<b>Native Hawaiian or Pacific Islander</b>
		<b>Asian</b>
		<b>Native American or Alaska Native</b>
		<b>Two or more races</b>

Importantly, the Proposed Rules do not address the question of what companies should do if they do not have a complete set of data about their applicants’ gender, race, and ethnicity that would be needed to conduct this kind of bias audit, as EEO1 Category data is often provided by candidates on a voluntary basis. To the extent that employers will need to infer any missing EEO-1 Category data from applicants’ names and addresses, that will likely inject a degree of error into the audit that may create questions about the data’s validity. The potential challenge is compounded by the fact that it is typically not advisable to ask applicants about their race or ethnicity during the application process.

This data may also paint a misleading picture of the company’s racial, ethnic, and gender diversity. As noted above, the gender options are restricted to male and female, and do not include any consideration of anyone identifying as transgender or gender non-conforming. Similarly, people who are “white (not Hispanic or Latino)” are defined as “persons having their origin in Europe, the Middle East, or North Africa,” which means that people of Middle Eastern or North African descent are not included as a separate ethnic category. Sexual orientation is also not included in any of these EEO-1 Categories. Thus, companies conducting a bias audit focused solely on the EEO-1 Categories may appear to be less diverse than is actually the case.

**Calculating “Selection Rate” and “Impact Ratio”**

Other than requiring a disparate impact assessment for the EEO-1 Categories, the AEDT Law does not provide a method for conducting the required bias audit. The Proposed Rules, however, introduce two metrics that must be ascertained as part of the audit: a “selection rate” and an “impact ratio”:

- Selection Rate is the rate at which individuals in an EEO-1 Category are either selected to move forward in the hiring process or assigned a classification by an AEDT—for example, how many Asian women a resume screening tool recommends for an interview. The Selection Rate metric is calculated by dividing the number of individuals in the EEO-1 Category moving forward or assigned a classification by the total number of individuals in that EEO-1 Category who applied for a position or were considered for promotion. So, if 100 people applied for a nursing position, and 10 of those applicants were Asian women, and three of those Asian women were selected for an interview, the Selection Rate for Asian women by the AEDT for that position would be 0.3 or 30%.
- Impact Ratio is the ratio of either (i) the selection rate for a particular EEO-1 Category divided by the selection rate of the *most selected* EEO-1 Category or (ii) the average score of all individuals in a particular EEO-1 Category divided by the average score of individuals in the *highest scoring* EEO-1 Category. So, continuing with the above example, if 10 of the 100 applicants were white men, and five of those men were selected for an interview, the Selection Rate for white men would be 0.5 or 50%. Assuming white men had the highest Selection Rate of any EEO-1 Category, the Impact Ratio for white men would be 0.5/0.5 or 1.0. The Impact Ratio for Asian women would be 0.3/0.5 or 0.6.

### Selection vs. Scoring

The Proposed Rules set forth requirements for conducting the bias audit in two different scenarios: (i) where an AEDT selects individuals to move forward in the hiring process or classifies individuals into groups (e.g., those that will receive an interview or be considered for a promotion); and (ii) where an AEDT provides applicants or candidates with scores that effectively rank them.

For a Selecting or Classifying AEDT, the employer’s bias audit must (i) calculate the selection rate for each EEO-1 Category and then (ii) calculate the impact ratio for each EEO-1 Category. In addition, where an AEDT classifies individuals into groups, steps (i) and (ii) must be taken with respect to each such classification. To illustrate this scenario using an example from the Proposed Rules, where an employer is using an AEDT to screen resumes and schedule interviews for a job posting, an independent auditor conducting the bias audit on behalf of an employer would look at the data on selected applicants and perform the hypothetical calculations as follows:

		# of Applicants	# Selected	Selection Rate	Impact Ratio
Hispanic or Latino	Male	205	90	43.9%	0.841
	Female	190	82	43.2%	0.827
Male	White	412	215	52.2%	1.000
	Black or African American	226	95	42.0%	0.806

Non-Hispanic or Latino		Native Hawaiian or Pacific Islander	87	37	42.5%	0.815
		Asian	321	167	52.0%	0.997
		Native American or Alaska Native	24	11	45.8%	0.878
		Two or More Races	115	52	45.2%	0.866
	Female	White	385	197	51.2%	0.981
		Black or African American	164	75	45.7%	0.876
		Native Hawaiian or Pacific Islander	32	15	46.9%	0.898
		Asian	295	135	45.8%	0.877
		Native American or Alaska Native	17	7	41.2%	0.789
		Two or More Races	98	44	44.9%	0.860

For a Scoring AEDT, the employer’s bias audit must (i) calculate the average score for individuals in each EEO-1 Category and (ii) calculate the impact ratio for each EEO-1 Category. For example, using an illustration from the Proposed Rules where an employer uses an AEDT to score applicants for “culture fit,” an independent auditor conducting the bias audit on behalf of an employer would look at the data on the average “culture fit” score of applicants for each category and perform the hypothetical calculations as follows:

		Average Score	Impact Ratio	
Hispanic or Latino	Male	76	0.884	
	Female	47	0.547	
Non-Hispanic or Latino	Male	White	86	1.000
		Black or African American	76	0.884
		Native Hawaiian or Pacific Islander	65	0.756
		Asian	71	0.826
		Native American or Alaska Native	69	0.802
		Two or More Races	84	0.977
		Female	White	68
	Black or African American	56	0.651	
	Native Hawaiian or Pacific Islander	74	0.860	
	Asian	67	0.779	
	Native American or Alaska Native	52	0.605	
	Two or More Races	86	1.000	

---

## Who Can Serve as the Independent Auditor

The Proposed Rules clarify the degree of independence required for the bias audit. The auditor must be “a person or group that is not involved in using or developing an AEDT that is responsible for conducting a bias audit of such AEDT.” This suggests that an employer using a developer’s tool cannot simply rely on the developer’s internal bias audit. However, the extent to which an employer can rely on a developer’s third-party audit remains unclear. It is also unclear whether the third-party auditor can separately provide consulting services to the employer or vendor with respect to their development and use of the tool.

---

## Publication of Audit Results and Notice Provisions

The Proposed Rules set forth prescriptive requirements for employers to comply with the AEDT Law’s requirement that employers make publicly available a summary of their bias audit results and the distribution date of the AEDT to which the audit applies. Additionally, employers must provide notices to candidates: (a) that an AEDT will be used in connection with the assessment or evaluation, and allow a candidate to request an alternative selection process or accommodation; (b) the job qualifications or characteristics that the AEDT will use in connection with the assessment; and (c) the types of data collected for the AEDT, the source of the data, and the employer’s retention policy. The Proposed Rules add several clarifications to these disclosure and notice obligations:

**Published Results.** The Proposed Rules clarify that the publication of the bias audit summary must (a) be either posted on the careers or jobs section of the employer’s website or linked to an external website (provided the link clearly identifies that it points to the results of the bias audit); (b) include the selection rates and impact ratios for all EEO-1 Categories; (c) note the distribution date of the AEDT, which is defined under the Proposed Rules as the date the employer began using a specific AEDT; and (d) remain posted for at least six months after the employer last used the AEDT to make an employment decision.

- **Notices to Candidates.** The Proposed Rules set out a number of options for how employers can provide required notices to candidates or employees under the AEDT Law, though they vary slightly depending on whether the recipient of the notice is a candidate or an employee. Employers can provide notice to candidates through the career or jobs section of its website in a clear and conspicuous manner, and employers can provide notice to employees through a written policy or procedure provided to the employees. For both categories, however, employers can also satisfy

the notice provisions by either including the notice in a job posting or providing notice in person, via U.S. mail, or email.

- **Notice Regarding Requests for Accommodations or Alternative Selection Processes.** The Proposed Rules mandate that the candidate notices must include instructions for how the candidate can request an alternative selection process or accommodation. Notably, however, the Proposed Rules also make clear that the AEDT Law does *not* require an employer to actually provide an alternative selection process. This creates a strange situation in which employers must provide the ability to request an opt-out, but need not grant any request or actually have an alternative available, unless it is for an accommodation required under the Americans with Disabilities Act or other applicable laws.

---

## Which Employers Are Covered by the AEDT Law?

One question that the Proposed Rules do not answer is which employers are covered by the AEDT Law. The AEDT Law is clear that it applies to companies located in New York City, that are hiring or promoting City residents, for jobs that are located in the City using a covered AEDT. The AEDT Law is not clear, however, as to whether it applies when a company located outside of the City is hiring New York City residents or when a company based in the City is hiring applicants from outside the City, and the Proposed Rules do not help resolve that ambiguity.

Some or all of the remaining ambiguities we describe above will likely be flagged by commenters or at the upcoming hearing, and will hopefully be addressed in the final rules.

To subscribe to our Data Blog, please click [here](#).


\* \* \*

Please do not hesitate to contact us with any questions.

**NEW YORK**



Avi Gesser  
agesser@debevoise.com



Jyotin Hamid  
jhamid@debevoise.com



Tricia Bozyk Sherno  
tbsherno@debevoise.com



Scott M. Caravello  
smcaravello@debevoise.com



Anna R. Gressel  
argressel@debevoise.com