

NYC's AI Hiring Law Is Now Final and Effective July 5, 2023

April 12, 2023

The New York City Department of Consumer and Worker Protection (the “DCWP”) has adopted final rules (the “Final Rules”) regulating the use of artificial intelligence for hiring practices. The DCWP’s [Automated Employment Decision Tool Law](#) (the “AEDT Law” or the “Law”) requires covered employers to conduct annual independent bias audits and to post public summaries of those results. To recap, the DCWP released an [initial set of proposed rules](#) on September 23, 2022, and held a public hearing on November 4, 2022. Due to the high volume of comments expressing concern over the Law’s lack of clarity, the DCWP issued a [revised set of proposed rules](#) on December 23, 2022, and held a second public hearing on January 23, 2023. After issuing the Final Rules, the DCWP delayed enforcement of the Law for the second time from April 15, 2023 to July 5, 2023.

The Final Rules largely adopt the December proposal with a few notable changes addressing concerns raised during the second public hearing. Most significantly, the Final Rules clarify that use of an AEDT to screen resumes requires a bias audit, even when the employer is not using the AEDT to make the final hiring decision. It also provides examples clarifying when employers can use test data instead of historical data for their bias audits. In this Debevoise Data Blog post, we discuss the current state of the AEDT Law and highlight how the final changes impact employers’ compliance obligations.

Defining an Automated Decision Tool

Under the AEDT Law, an “automated employment decision tool” (“AEDT”) includes (i) any computational process derived from machine learning, statistical modeling, data analytics or artificial intelligence; (ii) that issues simplified output, including a score, classification or recommendation; and (iii) that is used to substantially assist or replace discretionary decision-making for making employment decisions that impact natural persons. Examples of tools outside this definition include junk email filters, antivirus software, calculators, spreadsheets, databases, and other compilations of data. The Final Rules have modified the application of this definition in two ways.

Clarification on the Scope of an Employment Decision

Before using an AEDT for an “employment decision,” employers must ensure the AEDT has undergone a bias audit, conducted no more than one year prior to the use of the tool. The term “employment decision” means “to screen candidates for employment or employees for promotion within the city.” The term “screen” means “to make a determination about whether a candidate for employment or employee being considered for promotion should be selected or advanced in the hiring or promotion process.”

The Final Rules reiterate that the AEDT Law applies to employment decisions (i) based solely on the AEDT’s output, (ii) based on a number of factors, but the AEDT’s output is weighed more than any other criterion, or (iii) based on the AEDT’s output overruling conclusions derived from other factors, including human decision-making. Thus, under the Final Rules, if an AEDT’s output is one of many factors leading to a decision, the tool is only within the Law’s scope if the output outweighed all other factors or overruled a human decision.

The Final Rules clarify that an “employment decision” does not need to be a final hiring decision, and provide an example where an employer wants to use an AEDT to screen resumes and schedule interviews for a job posting. The example notes that the employer must conduct a bias audit before the planned use even though the tool is only being used to screen at an early point in the application process. With this clarification, some employers who may have believed that they were exempt may now conclude that they likely fall within the scope of the Law.

Broader Application to Complex Models

The Final Rules also broaden the Law’s application to complex models. The Final Rules redefine “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, if applicable, other parameters for the models in order to improve the accuracy of the prediction or classification. The Final Rules eliminate a previous limitation that required inputs and parameters to be refined through cross-validation or training and testing data. Many simpler AI tools may now fall within the revised definition.

Bias Audit Requirements

Protected Categories

The AEDT Law requires employers to conduct “bias audits” for AEDT tools, which include but are 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 Final 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 cover 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). Employers must consider intersectionality, and conduct the bias audit for each EEO-1 Category jointly. The Final Rules maintain that bias audit calculations must also be separately conducted for standalone sex and ethnicity/race categories (e.g., male, female, Hispanic, Black, Asian, White, etc.).

Clarity on Bias Audit Data Requirements

To conduct a bias audit, employers must use “historical data,” which is defined as “data collected during an employer’s or an employment agency’s use of an AEDT to assess candidates for employment or employees for promotion.” If there is not sufficient historical data available for a statistically significant bias audit, “test data” may be used, but the public summary of the bias audit results must then explain why historical data was not used and describe how the test data was generated. “Test data” is defined as any data other than historical data. The Final Rules answer a question that we raised in our last blog regarding how employers can conduct bias audits when they lack sufficient historical data. Specifically, the Final Rules clarify when employers may rely on other employers’ historical data or on test data:

- An employer may always rely on other employers’ “historical data” so long as the employer provides any personally collected “historical data” to an independent auditor for use and consideration.
- An employer may rely on “test data” when (i) using an AEDT for the first time or (ii) lacking sufficient “historical data” for a statistically significant bias audit. If relying on test data, the public summary must then explain why “historical data” was not used and describe how the “test data” was generated.

Calculating Selection Rates and Impact Ratios

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. However, as in the previous proposal, the Final Rules describe 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 scoring rate of all individuals in a particular EEO-1 Category divided by the scoring rate 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 Final Rules maintain the 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.

- Selection 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 specific groups (e.g., leadership styles), steps (i) and (ii) must be taken with respect to each such classification.

- Scoring AEDT. The employer's bias audit must (i) calculate the median score for the full sample of applicants; (ii) calculate the rate at which individuals in each EEO-1 Category receive a score above the median score (the "scoring rate"); and (iii) calculate the impact ratio for each EEO-1 Category by dividing the scoring rate of each group by the scoring rate of the highest-scoring group.

Who Can Serve as the Independent Auditor?

Bias audits must be conducted by an independent auditor, meaning "a person or group that is capable of exercising objective and impartial judgment on all issues within the scope of a bias audit of an AEDT." Like the previous proposal, the Final Rules outline three circumstances in which an auditor would not be considered independent:

- If an auditor is or was involved in using, developing, or distributing the AEDT;
- If an auditor, at any point during the bias audit, has an employment relationship with an employer or employment agency that seeks to use or continue to use the AEDT or with a vendor that developed or distributes the AEDT; or
- If an auditor, at any point during the bias audit, has a direct financial interest or a material indirect financial interest in an employer or employment agency that seeks to use or to continue to use the AEDT or in a vendor that developed or distributed the AEDT.

Since an auditor is expected to be paid for their work, a "disqualifying financial interest" must be more than that. The Final Rules provide no examples explaining what kinds of financial interests would disqualify an auditor. It remains unclear whether an auditor would be prohibited from having any other ongoing business relationships with the employer seeking to rely on the audit.

Publication of Audit Results and Notice Provisions

Like the previous proposal, the Final 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) of the job qualifications or characteristics that the AEDT will use in connection with the

assessment; and (c) of the types of data collected for the AEDT, the source of the data, and the employer's retention policy.

Published Results

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 number of applicants or candidates, selection rates, scoring rates (if applicable) and impact ratios for all EEO-1 Categories; (c) note the distribution date of the AEDT, which is defined under the Final 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.

The Final Rules introduce two circumstances where the public summary must include additional details:

- If an AEDT assessed individuals that were excluded from the impact ratio calculations because they fall into an unknown sex or ethnicity/race category, the number of assessed individuals must be included in the public summary; and
- Independent auditors may now exclude categories from the impact ratio calculations if the category comprises less than 2% of the data used for the bias audit. If such an exclusion is made, the auditor must include their reasons for the exclusion, number of applicants, and scoring or selection rate for the excluded category in the public summary.

Enhanced Obligations for Providing Notice to Candidates and Employees

Notice to Candidates and Employees

The Final Rules maintain the requirement that employers must provide notice to candidates or employees under the AEDT Law. Specifically, employers must notify candidates or employees residing in the city about both the use of an AEDT in their assessment or evaluation and the job qualifications and characteristics used by the AEDT. Notice must be provided at least ten business days before use of the AEDT. Both candidates and employees may be notified through a job posting or via U.S mail or email. Employers may also provide notice for candidates on the employment section of the website, and may provide employees notice through written policies or procedures.

Notice about AEDT Data Collection

Employers must also provide notice about (i) the type of data collected for the automated employment decision tool; (ii) the source of such data; and (iii) the employer or employment agency's data retention policy. Previously, employers could provide the

required notice in one of three ways. Now, the Final Rules require employers to use all three options. Employers must:

- provide information on the employment section of its website in a clear and conspicuous manner about its AEDT data retention policy, the type of data collected for the AEDT, and the source of the data;
- post instructions on the employment section of its website in a clear and conspicuous manner for how to make a written request for such information, and if a written request is received, provide such information within 30 days; and
- provide an explanation to a candidate for employment or employee being considered for a promotion why disclosure of such information would violate local, state, or federal law, or interfere with a law enforcement investigation.

Instructions for Accommodations or Alternative Selection Processes

The Final Rules maintain the requirement that notices to candidates must include instructions for how candidates can request an alternative selection process or accommodation. Notably, however, 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.

Penalties for AEDT Violations

Employers found to have violated the AEDT Law will face civil penalties of up to \$500 for a first violation and each additional violation occurring on the same day as the first violation. Each subsequent violation will incur a penalty between \$500 and \$1,500. Each day the AEDT is used in violation of the law gives rise to a separate violation, and failure to provide notice is also a separate violation.

Which Employers Are Covered by the AEDT Law?

The Final Rules leave the question of which employers are covered by the AEDT Law unanswered. 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. Notably, employers are only obligated to notify employees or candidates *residing in the city*. This would suggest that employers have different obligations depending on where potential applicants live, but the Final Rules do not help resolve these ambiguities.

Next Steps. As mentioned above, enforcement of the Law begins in under three months (July 5, 2023). The Final Rules address many of the issues raised during the comment period but also increase the compliance burden for employers. If employers use an AEDT for any part of their hiring process, the use very likely triggers compliance with the Law, and employers must ensure compliance with all notice requirements.

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

The *Debevoise Artificial Intelligence Regulatory Tracker* (“DART”) is now available for clients to help them quickly assess and comply with their current and anticipated AI-related legal obligations, including municipal, state, federal, and international requirements.

* * *

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



Anna R. Gressel
argressel@debevoise.com



Basil Fawaz
bfawaz@debevoise.com