



Comments Received by the Department of
Consumer and Worker Protection on

Proposed Rules related to Automated Employment Decision Tools

IMPORTANT: The information in this document is made available solely to inform the public about comments submitted to the agency during a rulemaking proceeding and is not intended to be used for any other purpose



January 18, 2023

Via email: Rulecomments@dcwp.nyc.gov

Commissioner Vilda Vera Mayuga
Department of Consumer and Worker Protection
Consumer Services Division
42 Broadway, 9th Floor New
York, NY 10004

Re: Future of Privacy Forum Comment on Local Law 144 or 2021 Proposed Rules

Dear Commissioner Mayuga and Members of the Department of Consumer and Worker Protection:

The Future of Privacy Forum welcomes this opportunity to provide feedback on the New York City Department of Consumer and Worker Protection’s (DCWP) draft rules to implement Local Law 144 of 2021 (LL 144) concerning the use of automated employment decision tools (AEDT).¹ FPF recommends that the Department seek to clarify audit requirements as they are required under the law, establish standards for auditors, and pay specific heed to the needs of all marginalized and multi-marginalized populations.

The Future of Privacy Forum (FPF) is a non-profit organization dedicated to advancing privacy leadership, scholarship, and principled data practices in support of emerging technologies in the United States and globally. FPF brings together industry, academics, consumer advocates, and other thought leaders to explore the challenges posed by technological innovation and develop privacy protections, ethical norms, and workable business practices. In 2018, FPF launched a working group to examine the implications of artificial intelligence on privacy and data use.² Through this work stream, we have produced several important resources examining the AI landscape, the scope of harms implicated by AI tools, and the risks of AI and Machine Learning to privacy and security, among other things.³

¹ New York City Department of Consumer and Worker Protection, “Notice of Public Hearing and Opportunity to Comment on Proposed Rule” (Dec. 15, 2022), <https://rules.cityofnewyork.us/wp-content/uploads/2022/12/DCWP-NOH-AEDTs-1.pdf>.

² Brenda Leong, “FPF Launches AI and Machine Learning Working Group and Releases New AI Resource Guides,” Future of Privacy Forum (August 21, 2018)

<https://fpf.org/blog/fpf-launches-ai-and-machine-learning-working-group-and-release-new-ai-resource-guides/>

³ The opinions expressed herein do not necessarily reflect the views of FPF’s supporters or Advisory Board.

Future of Privacy Forum Comments on LL 144 AEDT Rulemaking

Emerging technologies such as artificial intelligence and machine learning systems have clear potential to improve performance and efficiency across a variety of domains.¹ However, without necessary safeguards and review processes, these tools can also create or reinforce discriminatory impacts on individuals, which is particularly harmful when those impacts limit the availability of important personal and professional growth opportunities.² New York City's LL 144 is a landmark framework designed to support the transparent and accountable use of automated tools in employment decisions.³ In finalizing rules to implement the legislation and clarify requirements for the use of AEDT within New York City, FPF recommends that the Department consider the following principles for conducting bias audits of AI systems.

1. Recommendation #1 The Department should clarify what constitutes an adequate audit under LL 144

The Department should clarify what constitutes an adequate audit under the law. The law defines impact ratio, but does not define the kind of auditing procedures that would determine whether the audit was done correctly or adequately.⁴ Many different kinds of auditing procedures are available for AEDT systems.⁵ The Department needs to promulgate clear rules so that companies and individuals know what kinds of auditing tools meet LL 144's requirements.⁶ Clear rules would help companies better identify, implement, and use compliant auditing tools; a rule would be particularly beneficial if it focused on factors beyond impact ratios.

2. Recommendation #2 The Department should set clear professional standards for auditors

¹ "Key Ways Artificial Intelligence Can Improve Recruiting In The Hiring Process," Forbes (August 27, 2021) <https://www.forbes.com/sites/forbescoachescouncil/2021/08/27/key-ways-artificial-intelligence-can-improve-recruiting-in-the-hiring-process/?sh=3d1cd9882c52>

² For example, in 2018, a Reuters report alleged that Amazon suspended the development of an automated employment screening system that showed disproportionate bias against women's resumes. Jeffrey Dastin "Amazon scraps secret AI recruiting tool that showed bias against women," Reuters (Oct. 10, 2018), <https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G>

³ For example, LL 144 is already informing legislative efforts to establish rights and protections for the use of Artificial Intelligence systems in employment decisions in other jurisdictions including New York State and New Jersey (New York State Assembly, A567 available at https://nyassembly.gov/leg/?default_fld=&leg_video=&bn=A00567&term=2023&Summary=Y&Actions=Y&Text=Y; New Jersey Assembly, A4909 available at <https://www.njleg.state.nj.us/bill-search/2022/A4909>.)

⁴ New York City Department of Consumer and Worker Protection, "Notice of Public Hearing and Opportunity to Comment on Proposed Rules" (September 19, 2022) <https://rules.cityofnewyork.us/wp-content/uploads/2022/09/DCWP-NOH-AEDTs-1.pdf>

⁵ For example, some auditing systems examine a "snapshot" of how an AEDT system performs at a specific moment in time. Other audits allow both the deployer of the AEDT system and the auditor to get feedback in real time.

⁶ Ellen P. Goodman and Julia Trehu, "AI Audit-Washing and Accountability," German Marshall Fund (November 15, 2022) <https://www.gmfus.org/news/ai-audit-washing-and-accountability>

The Department should specify the factors that determine which firms qualify as “independent auditors” under LL 144, and what professional standards apply to auditors’ work. There are Future of Privacy Forum

Comments on LL 144 AEDT Rulemaking

questions as to what independent auditors must do within the context of an audit, as well as the standards that should apply from both an internal and external perspective. Affirmatively determining standards for auditors ensures that vendors and employers can certify they are working with the appropriate auditing firms. A failure to elaborate on what is required under the law in terms of professional standards could hamper the goals of the legislation and create risks for employers, candidates, employees, and vendors.

3. Recommendation #3 The Department should take into account the data equities of marginalized communities

The Department, while addressing hiring bias at scale, should keep in mind the data equities of marginalized communities. The issue of scale is not only one that the employees and job candidates struggle with but both vendors and employers. While there are fairly accurate approximations of race and gender that are easily accessible, there are other protected categories where data is not so clear. Disability, sexual orientation, and religious affiliation can be difficult to approximate based on limited information, and people who are LGBTQIA+, have a disability, or have different religious practices may not feel comfortable sharing that information. If the Department seeks to expand the classes of people to be considered in the audit beyond race and gender, the Department should consider the different privacy and data equities of the communities.

Thank you for this opportunity to provide input on New York City’s AEDT rulemaking. We look forward to participating in the upcoming public hearing and future opportunities to provide resources or information to assist in this important effort. If you have any questions regarding these comments, please contact Bertram Lee at blee@fpf.org.

Sincerely,

Bertram Lee
Senior Policy Counsel, Data Decision Making, and Artificial Intelligence

Amber Ezell
Policy Counsel

Jan 18, 2023



January 20, 2023

VIA E-MAIL: Rulecomments@dcwp.nyc.gov

City of New York
Department of Consumer and Worker Protection
42 Broadway
Manhattan, New York 10004

Re: Letter of Comment on the DCWP Updated Proposed Rules for NYC Local Law 144

Dear Sir/Madam:

The Institute for Workplace Equality (“IWE” or “The Institute”) submits the following comments in response to the New York City (“NYC” or the “City”) Department of Consumer and Worker Protection’s (“DCWP” or the “Department”) invitation. The Department’s Notice of Proposed Rules is seeking to clarify the requirements set forth by NYC’s Local Law 144 that will regulate the use of automated employment decision tools (“AEDT”) wherein hiring or promotion decisions are made or substantially assisted by algorithmically-driven mechanisms.

Background on The Institute for Workplace Equality

The Institute is a national, non-profit employer association based in Washington, D.C. The Institute’s mission includes the education of federal contractors as to their affirmative action, diversity, and equal employment opportunity responsibilities. Members of The Institute are senior corporate leaders in EEO compliance, compensation, legal, and staffing functions representing many of the nation’s largest and most sophisticated federal contractors.

The Institute recognizes the responsibility of all employers to create a nondiscriminatory workplace. To that end, NYC’s DCWP has an important role in enforcement efforts related to Local Law 144 and additional rules to clarify the requirements for compliance. This is critical to ensuring that employers understand their requirements and can effectively comply with the law beginning April 15, 2023.

Comments on Proposed Rules for Local Law 144

The Institute appreciates the efforts taken by the DCWP to update the proposed rules for Local Law 144 by incorporating some of the written and verbal comments previously provided by the Institute as well as other individuals and organizations. The clarifications and additions to the definition section, the clarifications for the bias audit, published results, and notices, and the addition of the data requirements section all enhance the ability for employers and auditors to understand and meet the requirements of Local Law 144. Our members found these updates to be very useful. However, the Institute would ask that the DCWP consider adding some important additional clarifications to the new section, §5-302 Data Requirements.

As context for the points on which we seek clarification, we suggest that there are four (4) use cases that will define the vast majority of bias audits required by Local Law 144. These use cases are as follows:

1. Implementing an AEDT where sufficient employer (historical) data are available, and vendor (test) data are not available.
2. Implementing an AEDT where sufficient employer (historical) data are not available, but sufficient cross-employer vendor (test) data are available.
3. Implementing an AEDT where sufficient employer (historical) data are available, and cross-employer vendor (test) data are also available.
4. Implementing an AEDT where there are no data available.

Section §5-302 Data Requirements contains three components. The first two components (a & b, listed below) suggest that an audit should be comprised of an employer's historical data where sufficient, and test data where historical data are insufficient—and where test data are used, an explanation of why and how should be provided.

“§ 5-302 Data Requirements.

- (a) A bias audit conducted pursuant to section 5-301 of this Chapter must use historical data of the AEDT. If insufficient historical data is available to conduct a statistically significant bias audit, test data may be used instead.
- (b) If a bias audit uses test data, the summary of results of the bias audit must explain why historical data was not used and describe how the test data used was generated and obtained. ...”

If one considers these two components of the Data Requirements in isolation, it suggests auditors do the following for each use case:

- Use Case #1 – auditors should use employer, historical data.
- Use Case #2 – auditors should use vendor, cross-employer, test data.
- Use Case #3 – auditors should use employer, historical data.
- Use Case #4 – the data to be used are less clear here, however, an auditor might work with the employer (and/or vendor) to consider alternative sources of test data (e.g., data from a validation study, data from a pilot study, simulated data) to serve as the basis for the bias audit, until such time as sufficient employer historical data became available (presumably in the following year’s audit).

The third component listed in Section §5-302 (c, listed below), however, introduces ambiguity related to the data to be used in Use Case #3.

“§ 5-302 Data Requirements.

... (c) A bias audit of an AEDT used by multiple employers or employment agencies may use the historical data of any employers or employment agencies that use the AEDT. However, an employer or employment agency may rely on a bias audit of an AEDT that uses the historical data of other employers or employment agencies only if it provided historical data from its use of the AEDT to the independent auditor for the bias audit or if it has never used the AEDT.”

This impacts the use case (#3) for which there is both sufficient employer (historical) data and sufficient cross-employer vendor data where the same AEDT is being used. The third component allows the bias audit to either, 1) be based solely on an analysis of the employer’s (historical) data, or 2) be based on an analysis of a vendor’s, cross-employer dataset, so long as the employer’s data comprise part of the dataset.

This may have been intentionally introduced to provide employers with multiple options for having an auditor complete their bias audit. However, if this is unintentional, then the specific parameters of this section should be revised accordingly.

There are four questions that the Institute sees as important to clarify to ensure that auditors, employers, and vendors are properly interpreting the intention of the city as it relates to the newly added section (§ 5-302 Data Requirements) in the proposed rules:

Question 1: Will the city provide guidance on what constitutes sufficient (versus insufficient) historical data, or will this be left to the judgment of the independent auditor?

Question 2: In Use Case #3 (where both employer historical data and vendor crossemployer data are available and sufficient), must the bias audit be conducted solely on the employer historical data, or may a cross-employer dataset provided by a

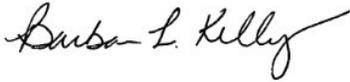
vendor, which includes data from the specific employer, serve as the basis for the bias audit instead?

Question 3: Will the city provide parameters for identifying appropriate test data for Use Case #4 (where no historical or vendor cross-employer data are available), or will the city defer to the judgment of the independent auditor and employer?

Question 4: May a bias audit for Use Case #4 use a simulated set of test data?

Thank you in advance for your consideration of The Institute's comments. We are happy to provide any additional information you may need, or to answer any questions you may have.

Best wishes,

A handwritten signature in cursive script that reads "Barbara L. Kelly".

Barbara L. Kelly
The Institute for Workplace Equality Director



Sanford Heisler Sharp, LLP
1350 Avenue of the Americas, Floor 31
New York, NY 10019
Telephone: (646) 402-5650
Fax: (646) 402-5651
www.sanfordheisler.com

David Tracey, Partner (646)
402-5667
dtracey@sanfordheisler.com

New York | Washington, DC | San Francisco | Palo Alto | Atlanta | Baltimore | Nashville | San Diego

January 22, 2023

VIA EMAIL

Rulecomments@dcwp.nyc.gov

City of New York, Department of Consumer and Worker Protection

42 Broadway #5

New York, NY 10004

Re: Comment on the Proposed Rule Amendments to NYC Local Law 144

To the Department of Consumer and Worker Protection:

Sanford Heisler Sharp, LLP commends New York City for passing Local Law 144 of 2021, which represents an important first step in curtailing discriminatory practices in artificial intelligence-based hiring. We are, however, deeply concerned that the Department of Consumer and Worker Protection's (the "Department") Proposed Rule Amendments ("Amendments") are inconsistent with language of the statute and will diminish its usefulness as tool to combat employment discrimination. As one of the largest worker-side employment law firms in the country, Sanford Heisler Sharp, LLP submits this comment to raise concerns about how the Amendments threaten to curtail Local Law 144's impact for its intended beneficiaries—applicants and employees.

One of our many concerns about the Amendments stems from the proposed definition of Automated Employment Decision Tool ("AEDT").¹ Local Law 144 defines AEDTs to include automated tools that "substantially assist" discretionary decision making in employment decisions. N.Y.C. Admin. Code § 20-870. The Department's original proposed rule significantly narrowed this definition by limiting it to tools that had a dispositive or predominating impact on employment decisions. Specifically, the tool's output had to be (i) the sole factor considered, (ii) "weighted more than any other criterion," or (iii) used to "overrule or modify conclusions." Now, the proposed rule further limits the definition by striking "or modify," so the third factor applies only to outputs that "overrule conclusions" about employment decisions.

Under the proposed definition, the exceptions threaten to swallow Local Law 144. Employers will undoubtedly characterize any use of an algorithmic tool as part of a holistic inquiry attendant to each employment decision, in which no one factor is dispositive or weighted more than any other.

Department of Consumer and Worker Protection
January 22, 2023
Page 2 of 2

In so doing, their tools will evade the important notice, audit, and publication requirements of Local Law 144. To be sure, employers who disingenuously characterize their tools will risk regulatory enforcement and penalties. But uncovering potential violations may prove difficult and establishing a violation even more so.

A definition more in line with the purpose of the statute will reduce the possibility of evasion. Such a definition should focus on all manners in which automated tools may “substantially assist” an employment decision, whether or not they are dispositive or predominating factors in such decisions. Specifically, the Department should strike from the second definition the phrase “where the simplified output is weighted more than any other criterion in the set.” Additionally, the Department should reinsert the phrase “or modify” into the third definition. These simple edits will provide increased protection to applicants and employees and honor the statute’s language and intent.

Thank you for considering these comments.

Sincerely,

A handwritten signature in black ink, appearing to read 'D. Tracey', is written over a horizontal line.

David Tracey

¹ Organizations, such as the New York Civil Liberties Union, have previously offered other important feedback that we encourage the Department to continue carefully considering.



U.S. Chamber of Commerce

1615 H Street, NW
Washington, DC 20062-2000
uschamber.com

January 23, 2023

Commissioner Vilda Vera Mayuga
NYC Department of Consumer and Worker Protection
Consumer Services Division
42 Broadway, 9th Floor
New York, NY 10004

Re: Proposed Rules, NYC Department of Consumer and Worker Protection; Automated Employment Decision Tools (Updated); (January 23, 2023)

Dear Commissioner Mayuga:

The U.S. Chamber of Commerce's Chamber of Technology Engagement Centers (C_TEC) appreciates the opportunity to provide further feedback to the Department of Consumer and Worker Protection (DCWP) on the proposed rules for "Requirement for the use of Automated Employment Decision Tools." As we stated in our first filing, we believe that the use of A.I. in the hiring and promoting process has been essential in helping streamline the review, outreach, vetting, and onboarding process of potential employees, and A.I. has become an essential tool for employers to use to avoid their own unconscious bias in the hiring process.

C_TEC has long recognized that "fostering public trust and trustworthiness in A.I. technologies is necessary to advance its responsible development, deployment, and use." We believe it is essential that DCWP make sure that the rules implementing Int. 1894-2020 in relation to "automated employment decision tools," are made in a considerate and balanced manner to ensure that the deployment of such tools benefits the employer/employment agency, employee, and/or independent contractor to streamline the process.

We are concerned that the regulations, as currently drafted, would impede the ability of businesses to find and hire qualified candidates in New York City, by reducing the number of candidates that may be considered for an open position. Such a scenario, during our current challenging time of an acute labor shortage, deprives businesses of the tools that would allow for a review of a larger volume of resumes. Our comments seek to highlight many of the areas in which regulations can be better tailored to help the business community as well as candidates for employment.

Definition of Independent Auditor: The current definition of "independent auditor" is overly restrictive, resulting in higher costs and potential obstacles in hiring. While outside

assistance should never be prohibited, it should be noted that there is a well-documented risk¹ of engaging third-party auditors. There currently are no universal standards and certifications regarding third-party audits. This means there is no guarantee that auditors can deliver verifiable measurement methods that are valid, reliable, safe, secure, and accountable. For this reason, we would encourage the definition to be changed to how it was initially proposed to mean “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.”

Definition of Automated Employment Decision Tool: We ask for the following changes to the definition of “automated employment decision tool.” We would ask the following to be stripped from the current definition “(ii) 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 (iii) to use a simplified output to overrule conclusions derived from other factors including human decision-making.” Furthermore, we would ask that you add the following sentence: “Automated employment decision tool,’ or ‘AEDT,’ does not include the automated searching of resumes to identify candidate qualifications, including relevant skills or experience.”

Bias Audit: The examples provided in subsections (b) and (c) of 5-301 are both prescriptive in who bears responsibility for the bias audit (i.e., the employer/deployer or the vendor/developer) without accounting for the range of possible scenarios. For this reason, we prefer that the examples be made clear that they aren't necessarily exhaustive of all scenarios and remove the specificity of responsibility in each of the two examples, allowing for flexibility to account for the range of scenarios.

We request the following changes to subsection (b):

- In the example, strike “provides historical data” and all that follows and replace it with “uses test data to conduct a bias audit as follows:”

We request the following changes to subsection (c):

- In the example, strike the word “planned” from the phrase “planned use of the AEDT.”
- Also, in the example, strike “provides historical data” and replace it with “uses test data.”

Finally, both examples suggest that the bias audit should compare selection rates of not just gender and race/ethnicity – the usual categories required to be compared under the Uniform Guidelines of Employee Selection Procedures – but also on the intersectional categories of gender and race/ethnicity (e.g., Hispanic Males, NonHispanic Female Whites, etc.). Data on these intersectional categories, however, typically is not collected by employers or vendors, as applicants and employees are given the opportunity to separately self-identify their gender and

¹ <https://www.gmfus.org/sites/default/files/2022-11/Goodman%20%26%20Trehu%20-%20Algorithmic%20Auditing%20-%20paper.pdf>

their race/ethnicity. Furthermore, many employers and vendors do not collect any gender or race/ethnicity data on their applicants; please clarify how such employers and vendors should conduct a bias audit in the circumstance in which they do not have any or complete demographic data.

In section 5-303, we also suggest striking the phrase in subsection (a)(1) "the selection rates and impact ratios for all categories," and replacing it with "a statement on adverse impact." Further, the current language in the proposed rule is inconsistent as the definition of "impact ratio" includes either selection rate "or" scoring rate (whereas the wording in the publication requirement mistakenly requires publishing the impact ratio "and" the selection rate).

Data Requirements: C_TEC would also like to highlight our concerns with section 5-302 on data requirements. Specifically, we would ask for clarification on how historical data is made available. There are many practical implementation challenges with using historical data. For instance, the data may reside with multiple entities, making it impossible to compile and share. Furthermore, the sharing of such data introduces privacy concerns and could be used by vendor competitors to create their own Automated Employment Decision System. For this reason, we would encourage that for any vendor-initiated audits, the test data should be the default instead of historical data.

Vendor Audits: The proposed rules contain an example in section 5-301(b) that strongly implies that employers can rely upon bias audits commissioned by vendors using historic applicant data collected by the vendor and not the employer's own data. We ask that the rule explicitly state that this is permissible and satisfies the "bias audit" requirement. It should also make clear that "historical data" may not be available and "test data" would be sufficient.

Lookback Period: While we appreciate that the DCWP provided a temporary delay in the enforcement of the AEDT law until April 15, 2023, in order to provide final rules, the U.S. Chamber strongly encourages the Department to provide a lookback period from April 15, 2023 of at least twelve (12) months from such date to businesses and organizations as they look to implement the final rule.

Conclusion:

We appreciate the opportunity to comment on the implementing rules. We believe it is essential that we get these regulations correct so that New York City does not impose overly broad requirements, which in turn could create significant uncertainty regarding the use of automated employment decision tools in hiring.

Any potential limitation of the use of technology for hiring purposes for businesses could lead to unnecessary barriers to finding qualified candidates for a job. This is not an appropriate policy choice, as we are in a historically tight labor market. Accordingly, we believe that businesses should be able to use tools to identify as wide a pool of applicants as possible. The

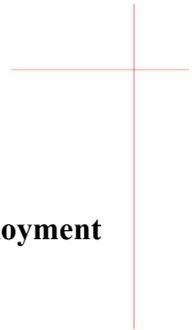
current draft regulations deprive businesses of the ability to do so. This is also harmful to those who are seeking employment as well.

Automated employment decision tools are essential in helping streamline the hiring and promotion process, so we ask that you make the following above-proposed changes to give the business community the necessary certainty they will need. If you have any questions, do not hesitate to contact Michael Richards at mrichards@uschamber.com.

Sincerely,

A handwritten signature in black ink, appearing to read 'TK' followed by a long, sweeping horizontal stroke.

Tom Quaadman
Executive Vice President
Chamber Technology Engagement Center
U.S. Chamber of Commerce



ADP Comments on the DWCP Revised Proposed Rules on Automated Employment Decision Tools

January 23, 2023

ADP appreciates the opportunity to provide comments on the Department of Consumer and Worker Protection's (DCWP) Revised Proposed Rules Regarding Automated Employment Decision Tools (AEDTs). ADP provides a range of administrative solutions to over one million employers worldwide, enabling employers of all types and sizes to manage their employment responsibilities from recruitment to retirement, including payroll services, employment tax administration, human resource management, benefits administration, time and attendance, retirement plans, and talent management. ADP has been a leader in AI ethics, including through publication of a set of AI ethics principles and establishing an AI Data & Ethics Committee comprised of internal and external experts.

As with the original Proposed Rules, the Revised Proposed Rules provide helpful clarification of how Local Law No. 144 operates and give companies developing and deploying AEDTs greater certainty regarding how to meet the law's obligations. At the same time, we ask that DCWP revert to the provision in the first draft of the Proposed Rules allowing companies to use internal resources to conduct independent audits, so long as those resources were not involved in the development of the AEDT. We also ask that DCWP delay enforcement of the law to give companies the opportunity to implement it in light of the Revised Proposed Rules.

- **Independent Auditor.** In the original Proposed Rules, DCWP defined an independent auditor to include persons or groups that might be part of the same company but were not involved in the development or use of the AEDT. Oftentimes, others in the company will be in the best position to conduct an audit, given their expertise in the systems the company uses/develops and the particulars of the machine learning the company employs. While third parties are increasingly entering the AI audit space, this industry is still nascent, so enabling companies to rely on internal experts who were not involved in the development or use of the AEDT helps ensure that the bias audit is effective.
- **Delayed Enforcement Date.** Given the revision to the Proposed Rules and the new hearing, DCWP has already delayed the enforcement date of the Act to mid-April. We ask that DCWP further delay enforcement of the law until at least 180 days after adoption of the Revised Proposed Rules. Given the planned timeline for the adoption of the Revised

Proposed Rules, a short enforcement delay would give companies time to fully implement the law as clarified by the regulations.

We also want to call out an additional positive aspect of the Revised Proposed Rules. Several commenters to the Proposed Rules asked that bias audits be performed by each employer, even when a tool, such as those we provide, are used across many employers. This would have resulted in much duplicative effort by those employers, imposing a substantial burden on small and medium-sized businesses and reducing uptake of AEDT's, which offer substantial benefits to employers and candidates alike. The Revised Proposed Rules provide that employers can rely on a single bias audit of an AEDT used across multiple employers, so long as they contribute data to the audit, thereby addressing concerns about bias resulting from the AEDT itself.

We have also attached our comments on the original Proposed Rules. We ask that DCWP continue to consider this feedback, particularly the need for additional clarity regarding to which roles the law applies.

* * *

ADP appreciates the opportunity to submit these comments on the Proposed Rules and DCWP's consideration of them. If you have any questions or would like additional information, please do not hesitate to contact Jason Albert, Global Chief Privacy Officer, ADP at jason.albert@adp.com.

ADP Comments on the DCWP Proposed Rules on Automated Employment Decision Tools

October 24, 2022

ADP appreciates the opportunity to comment on the Department of Consumer and Worker Protection's (DCWP) Proposed Rules Regarding Automated Employment Decision Tools (AEDTs). The Proposed Rules provide helpful clarification regarding how Local Law No. 144 operates and give companies developing and deploying AEDTs greater certainty regarding how to meet the law's obligations. At the same time, we ask that DCWP consider further defining to which jobs the law applies and, given the timing of finalization of the Proposed Rules, delay enforcement of the law to give companies the opportunity to implement it in light of the provisions of the Proposed Rules.

ADP provides a range of administrative solutions to over 990,000 employers worldwide, enabling employers of all types and sizes to manage their employment responsibilities from recruitment to retirement, including payroll services, employment tax administration, human resource management, benefits administration, time and attendance, retirement plans, and talent management. ADP has been a leader in AI ethics, including through publication of a set of AI ethics principles and establishing an AI Data & Ethics Committee comprised of internal and external experts.

The Proposed Rules provide greater clarity as to how the law operates and by doing so will help companies more effectively meet the requirements and objectives of the law. Specifically:

- ADP appreciates the added precision in the Proposed Rules on what constitutes an AEDT subject to the law's requirements. By specifying that the system must be the sole factor, or outweigh any other factor, or overrule or modify human decision-making, to be a covered AEDT, the Proposed Rules ensure that the law applies to instances where the AEDT is the primary factor in the decision whether to hire or promote an individual. Importantly, this helps ensure that the law doesn't inadvertently impinge on supplemental uses of machine learning technology in the hiring/promotion process.
- ADP further appreciates that DCWP defined an independent auditor to include persons or groups that might be part of the same company but were not involved in the development or use of the AEDT. Oftentimes, others in the company will be in the best position to conduct an audit, given their expertise in the systems the company uses/develops and the particulars of the machine learning the company employs. While third parties are increasingly entering the AI audit space, this industry is still nascent, so enabling companies to rely on internal experts who were not involved in the development or use of the AEDT helps ensure that the bias audit is effective.
- ADP also appreciates the added specificity regarding the minimum requirements for a bias assessment under the law and what information must be published in the summary. The minimum requirements set forth in the Proposed Rules are clear and achievable, while leaving room for companies to do more sophisticated analyses based on market demand.
- In addition, ADP appreciates the approach to providing notice to New York City residents set out in the Proposed Rules. By enabling companies to provide notice on their websites prior to posting positions, it ensures that the recruiting process can run smoothly and quickly, and NYC residents are not disadvantaged relative to applicants from other jurisdictions by delayed consideration of their applications.

While applauding the aspects of the Proposed Rules mentioned above, we ask that DCWP take two additional steps to help ensure that New York City residents are not disadvantaged relative to other job applicants by the law and to enable companies to implement the law in light of the helpful guidance provided by the Proposed Rules. Specifically:

- Additional clarity regarding to which roles the law applies would be helpful. The law states that notice and opt-out must be provided to New York City residents, but otherwise merely says that the law applies "in the city." These leaves unclear exactly what roles fall within the law's scope. To avoid concerns about long-arm jurisdiction and ensure clarity as to the law's scope, ADP asks DCWP to clarify that the law applies to posted jobs where the role will be physically located in New York City. If the scope of the law's applicability remains

unclear, potential employers in other states might avoid considering New York City residents for their open roles.

- ADP also asks that DCWP delay enforcement of the law until at least 180 days after adoption of the Proposed Rules. Given the short timing between when the Proposed Rules were proposed, much less adopted, and the effective date of the law, a short enforcement delay would give companies time to fully implement the law as clarified by the regulations.

* * *

ADP appreciates the opportunity to submit these comments on the Proposed Rules and DCWP's consideration of them. If you have any questions or would like additional information, please do not hesitate to contact Jason Albert, Global Chief Privacy Officer, ADP at jason.albert@adp.com.



From: **Shea Brown, Ph.D.**
Chief Executive Officer
BABL AI Inc.
sheabrown@babl.ai

BABL AI INC.
The Algorithmic Bias Lab
630 Fairchild Street
Iowa City, Iowa 52245
<https://babl.ai>

To: **NYC Department of Consumer and Worker Protection**
Rulecomments@dca.nyc.gov

Re: **Public Comments on the Proposed Rules for Local Law 144 of 2021**

January 23, 2023

To Whom It May Concern:

On behalf of the team at BABL AI, I thank the department for the opportunity to provide public comments on the proposed rules for enforcing Local Law 144 requiring annual bias audits for automated employment decision tools (AEDTs).

We commend the department for providing guidance and clarity on many aspects of the proposed rules. In particular, we are pleased that the new proposed rules

1. Tightened the requirements for independence,
2. Required disclosure of intersectional analysis,
3. Established data requirements for auditing, and
4. Limited the generalizability of audit results.

As a company that audits algorithms for ethical risk, effective governance, bias, and disparate impact, BABL AI believes that the spirit of this law furthers our mission to ensure safe and fair algorithms that prioritize human flourishing.

Although the new rules directly address many of our previous concerns, we would like to comment on three specific areas:

(1) Definition and scope of AEDT: The current definition of AEDT is overly stringent. In particular, the bar an automated tool must meet to qualify as being able to “substantially assist or replace discretionary decision making” is excessively high.

In our experience it is extremely rare that an automated system is deployed in such a manner that would (1) provide users with no other factors besides one simplified output, (2) have the simplified output be the primary factor in user decision making, or especially, (3) overrule conclusions by human decision-making.

Quite the contrary, the majority of the automated systems both involving and not involving AI/ML are intended to *assist* users—e.g., recruiters and hiring managers. They do so often by providing a plethora of information about candidates, including perhaps a “simplified output” by the AI/ML component, among other non-AI/ML-based data regarding the candidates such as their inputted profile or assessment scores. For example, AI interview tools not only display results from their AI/ML components such as their computed personality trait scores¹, but also give immediate access to the candidate’s non-AI/ML-based data—available on the same screen or a click away. Employers and vendors of these tools can thus claim exemption from a bias audit simply because their tools’ outputs (1) are not the only factor being considered by the recruiters—evidenced by there being other pieces of information they can examine, (2) can be argued not to be the primary influencing factor, and (3) do not overrule recruiter decision-making—by design. **Under this current AEDT definition, such paradigmatic automated tools whose harm of bias and discrimination is intended to be prevented by the proposed rules would risk being exempt from a bias audit.**

We strongly urge the Department to remove the expanded definition of “substantially assist or replace discretionary decision making.” We believe that employers procure these tools because they believe they will substantially assist in their employment decisions, and further attempts to clarify these words will unnecessarily narrow the scope.

(2) Definition of “machine learning, statistical modeling, data analytics, or artificial intelligence”: As the current scope of AEDT tends to focus on AI- and ML-based methods, the department should be cautious that **an automated tool does not need AI or ML to result in bias or discrimination.** The current definition of “machine learning, statistical modeling, data analytics, or artificial intelligence,” despite being precise, focuses primarily on methods which are based on training and optimization of parameters. We believe the overemphasis on this technical aspect of ML would exempt many simple algorithms and automated systems that nonetheless likely exacerbate bias. We provide more details in this video² where we illustrate the ways simple algorithms can amplify discrimination and give a concrete example of how one can design a simple algorithm around the requirements of the proposed rules that do not rely on training and parameter-optimization.

We suggest that the Department remove the new definition of “machine learning, statistical modeling, data analytics, or artificial intelligence.” What is currently written does not cover the extent of what these words mean, nor what the spirit of the law intended.

(3) Number of applicants for the scoring rate method: The examples provided in the proposed rules only show the number of applicants for the selection rate method but not for the scoring rate method. The number of applicants is important because it gives readers of public summaries, regardless of level of technicality—some information about the uncertainty and generalizability of the impact ratios. Omission

¹ <https://www.technologyreview.com/2021/07/07/1027916/we-tested-ai-interview-tools/>

² <https://www.youtube.com/watch?v=3VAYGnMLLS8>

of such numbers creates an information asymmetry, and lowers the level of transparency for one of the methods.

We encourage the department to include disclosure of the number of applicants in the example for the scoring rate method. This should equalize the requirements for disclosure for the two methods of impact ratio calculation and provide greater transparency for readers of the resulting summaries.

Again, we would like to thank the NYC Department of Consumer and Worker Protection for providing us the opportunity to comment on the proposed rules of Local Law 144 of 2021, and we would be happy to provide further clarification on any of the above comments.

Contact

A handwritten signature in black ink, appearing to read 'Shea Brown', written in a cursive style.

Shea Brown, Ph.D.
CEO & Founder
BABL AI Inc.
sheabrown@babl.ai



January 23, 2023

To: New York City Department of Consumer and Worker Protection
42 Broadway
New York, NY 10004

Re: Revised Proposed Rules to Implement Local Law 144 of 2021 on Automated Employment Decision Tools

The Center for Democracy & Technology (CDT) respectfully submits these comments on the Department of Consumer and Worker Protection’s (“DCWP”) revised proposed rules to implement Local Law 144 of 2021 (“LL 144”) relating to the auditing and notice requirements for employers’ use of automated employment decision tools (“AEDTs”). CDT is a nonprofit, nonpartisan 501(c)(3) organization that advocates for stronger civil rights protections in the digital age. CDT’s projects include advocating standards and safeguards to ensure that algorithm-driven systems do not interfere with workers’ access to employment.

As CDT previously explained, LL 144’s bias auditing and notice requirements would not ensure accountability in the use of AEDTs.¹ We submitted comments in October 2022 in which we described how the initial proposed rules would weaken LL 144’s effectiveness further and identified changes to mitigate these concerns.²

We are encouraged to see some of our recommendations from our previous comments reflected in the revisions. Namely:

- The definition of “independent auditor” has been expanded to minimize auditors’ potential conflicts of interest. This definition can be strengthened further by stating that an independent auditor is “a person or group that exercises objective and impartial judgment on all issues” within the scope of an AEDT’s bias audit – instead of someone

¹ Matthew Scherer and Ridhi Shetty, *NY City Council Rams Through Once-Promising but Deeply Flawed Bill on AI Hiring Tools*, Center for Democracy & Technology (Nov. 12, 2021), <https://cdt.org/insights/ny-city-council-rams-through-once-promising-but-deeply-flawed-bill-on-ai-hiring-tools/>.

² Center for Democracy & Technology, Comments on New York City Department of Consumer and Worker Protection’s Proposed Rules Implementing Local Law 144 (Oct. 24, 2022), <https://cdt.org/wp-content/uploads/2022/10/CDT-Comments-on-NYC-AEDT-proposed-rules.pdf> [hereinafter “October 2022 Comments to NYC DCWP”].

“that is capable of” doing so – to clarify that the auditor is obligated to actually exercise objective and impartial judgment.

- The information that employers and employment agencies must make publicly available must include the source and explanation of data used to conduct the bias audit.
- Bias audits must also calculate the AEDT’s impact on intersections of sex with race or ethnicity.
- Notice to workers must include information on an AEDT’s data retention policy, the type of data collected for the AEDT, and the source of the data.

However, these positive changes are outweighed by revisions to the proposed rules that further dilute the protections that LL 144 is supposed to secure. We urge DCWP to make further changes to better protect workers.

New proposed definitions of “automated employment decision tool” and “screen” limit the law’s scope even further

The initial proposed rules recognized that LL 144 limits the definition of AEDTs to tools that “substantially assist or replace discretionary decision making,” and defined this phrase to mean that the tool either relies “solely on a simplified output,” uses a set of criteria in which the simplified output is the criterion given the greatest weight, or uses the simplified output to overrule or modify conclusions derived from other factors. Our previous comments explained that this would enable employers to use AEDTs that still substantially influence the overall employment decision but are not the sole or primary criterion.¹

The revised proposed rules make only one change to the definition of AEDTs: the phrase “or modify” is removed from the final clause of the definition, conveying that a tool is an AEDT *only* if it uses a simplified output that *overrules* conclusions derived from other factors, but not if the tool uses a simplified output that modifies such conclusions. This one change even further restricts the types of automated tools that would be covered, leaving more workers unprotected.

This change does not ensure the definition is “focused,” as the statement of basis and purpose suggests; it instead creates a loophole that could swallow the law. Employers may be able to evade the requirements of LL 144 simply by casting AEDT’s outputs as “recommendations” that human decision-makers either rubber-stamp or hesitate to contradict. We again urge the Department to adopt a definition that restores the full scope of the statutory text, either by

¹ *Id.* at 2-3.

leaving the term “substantially assist” with its plain and ordinary meaning, or by ensuring that the definition includes tools whose output is “an important or significant factor in an employment decision.”

Similarly, the definition of “screen” previously applied to a determination about “whether *someone* should be selected or advanced in the hiring or promotion process” (emphasis added). Our previous comments noted that this definition would exclude targeted job advertising and targeted recruiting.¹ We advised that this definition should be broadened to apply to determinations, based on protected characteristics, about whether to represent that any employment or job position is available.² Instead, the revised proposed rules further narrow the definition of “screen” to only apply to a determination about whether a “candidate for employment or employee being considered for promotion” should be selected or advanced. This more explicitly excludes determinations about how job advertisements and recruiting efforts are targeted, enabling the use of automated tools to limit which workers learn of job opportunities.

Other critical definitions remain inappropriately narrow

Despite our previous suggestions, the definitions of “candidate for employment,” “employment decision,” and “employment agency” are left unrevised. As a result, the rules still exclude certain automated tools and actors from coverage that are part of the hiring or promotion process and contribute to discriminatory outcomes throughout the employment cycle. For instance, workers would not be protected from discriminatory targeted job advertising, as they would not be considered candidates and the advertising practice would not be recognized as an employment decision. Platforms that direct job advertisements or recruiters to specific people on behalf of employers would not be recognized as employment agencies despite performing the functions of an employment agency as defined under New York City’s anti-discrimination laws and New York State law.³ We urge the Department to broaden these proposed definitions, in accordance with our corresponding comments on the original proposed rules, which we incorporate herein by reference.⁴

Revisions to the bias audit requirements create new ambiguity

A new addition to the rules’ section on bias audits states that the selection rate and impact ratio required for bias audits must separately calculate the AEDT’s impact on sex categories, race/ethnicity categories, and intersectional categories of sex, ethnicity, and race. The unrevised language immediately preceding the addition already establishes this requirement with respect

¹ October 2022 Comments to NYC DCWP, *supra* note 2, at 3.

² *Id.* at 4.

³ N.Y.C. Admin. Code § 8-102; N.Y. Gen. Bus. Law § 171.

⁴ October 2022 Comments to NYC DCWP, *supra* note 2, at 3-4.

to sex categories and race/ethnicity categories: it states that “a bias audit must, at a minimum” calculate the selection rate for each category and calculate the impact ratio for each category. In its statement of basis and purpose of proposed rule, DCWP states that the revisions clarify that the impact ratio must be calculated separately to compare an AEDT’s impact on each category. To better provide this clarification, Section 5-301(b)(3) should state that the calculations required in Section 5-301(b)(1)-(2) must be used to *compare* the impacts of the AEDT on each individual category and on intersectional categories of sex, ethnicity, and race.

Section 5-301(b)(4) states that when the AEDT classifies candidates or employees into “groups,” the AEDT’s selection rate, impact ratio, and impact on each category must be separately calculated for each group. Based on the parenthetical provided in the provision, the term “groups” appears to be referring to job qualifications and characteristics, but the term is vague in this context. Instead, the provision could refer to an AEDT “assigning” or “generating a classification” for candidates or employees being considered for promotion. This would align with the rules’ definitions of “selection rate” and other terms. Another alternative would be for the section to refer to an AEDT classifying candidates or employees based on how job qualifications or characteristics are demonstrated. This would make it all the more important to define “job qualifications and characteristics” as well, which as our previous comments mention, neither the law nor the proposed rules define. The Department should include such a definition in its final rules.

The proposed rules state how the impact ratio can be calculated for an AEDT that selects or classifies candidates or employees and for an AEDT that scores candidates or employees. For the latter type of AEDT, the revised rules replace the previous formula, which was based on a comparison of the average score of people in a given category with the average score of people in the highest scoring category. In the revised rules, the formula is based instead on a comparison of scoring rate between these two categories. The revised rules also add a definition for “scoring rate”: “the rate at which individuals in a category receive a score above the sample’s median score, where the score has been calculated by an AEDT.” To provide the necessary context to reliably determine whether a comparison of scoring rates reflects disparities, the results should also include the number of total applicants and selected applicants included in the sample.

New data requirements for bias audits enable employers to evade scrutiny

The revised proposed rules add a new Section 5-302 to establish that a bias audit must use historical data of the AEDT. If there is insufficient historical data for the bias audit, this section states that test data “may be used instead.” If test data is indeed used instead, the bias audit must explain why historical data was not used for the audit, and it must explain how the test data was generated and obtained. Instead of stating that test data “may be used” if there is

insufficient historical data, the section should state that test data “must be used instead,” to make explicit that employers’ obligations do not cease if they have insufficient historical data.

Section 5-300 adds definitions for “test data” and “historical data.” “Historical data” is defined as “data collected during an employer or employment agency’s use of an AEDT” to assess candidates or employees. This definition allows employers the discretion to decide the type and extent of data that must be used in the bias audit. “Test data” is defined simply as “data used to conduct a bias audit that is not historical data.” Here, too, without any further parameters regarding the detail or extent of such data, any non-historical data could be sufficient for employers to consider test data for the purpose of bias audits.

Further, the new Section 5-302 states that a bias audit can use historical data for any employer or employment agency that uses the AEDT when that AEDT is used by multiple employers or employment agencies. This section adds that an employer or employment agency may only rely on a bias audit that uses another’s historical data if that employer or employment agency also provides the independent auditor with its own historical data. The revised proposed rules do not address what the independent auditor is obligated to do with an employer’s own historical data if the employer relies on a bias audit that uses another employer or employment agency’s historical data. In addition, if an employer relies on a bias audit that uses a different employer or employment agency’s historical data, that bias audit is even less likely to accurately capture the AEDT’s impact on race, sex, and intersectional categories that would be reflected in that employer’s own historical data.

To address these issues, we urge the Department to modify Section 5-302 to require an employer’s bias audit to use *all* of their own historical data covering the period since the last bias audit was conducted or, if it is the first bias audit for a tool, all available historical data from the employer’s own use of the tool. The Department should also clarify that a bias audit can only use data from other employers if all the following conditions are met:

- Other employers’ historical data is only used to *supplement* the employer’s own historical data;
- Other employers’ historical data is used only to the extent needed to draw meaningful conclusions about the tool’s discriminatory impact;
- The bias audit distinguishes the employer’s own historical data from the data it is using from other employers.

These changes would ensure that employers and vendors cannot manipulate what data the bias audit includes, and ensure that all relevant historical data for each employer is included in each bias audit.

Notice requirements still keep workers at a disadvantage

As mentioned above, the revised proposed rules take a positive step by providing for candidates or employees to receive information about the AEDT data retention policy and the type and source of data collected. Our previous comments urged the DCWP to ensure that employers must proactively provide these details and describe all criteria that will be evaluated and the purpose for using such criteria to workers.¹ We added that all of this information should be provided to workers prior to an AEDT's use and through multiple channels to make notice more accessible to workers with different needs. Otherwise, the onus is left on workers to try to find these details on employers' websites or submit a written request for these details.

However, the revised proposed rules do not require these details to be provided prior to an AEDT's use. Section 5-304(d) continues to allow for employers to wait until they receive a written request to provide information about the AEDT data retention policy and type and source of the collected data, in which case employers are required to provide instructions in a "clear and conspicuous manner" on how to make this request. The revisions clarify that employers must comply with the written request within thirty days, but they do not address whether this notice would be provided before or after an AEDT's use. A fundamental part of all data policies is *prior* notice explaining what data is collected, where it comes from, and how it will be treated – it should be required, not optional, to inform workers about how their data will be handled *before* they provide their data. Therefore, we urge the Department to issue rules explicitly stating that, where practicable, these forms of notice should come before the AEDT's use and should be provided through multiple channels.

Meanwhile, the only information that Section 5-304(b)-(c) requires employers to provide ten days prior to an AEDT's use is (1) that an AEDT will be used and (2) the "job qualifications and characteristics" that it will evaluate. The term "job qualifications and characteristics" remains undefined, so this ten-day notice period will not ensure that candidates and employees receive the necessary details about how they will be assessed early enough in the selection process to make an informed decision about whether they may need to request accommodations or alternative selection processes. As recommended above, the Department should add a definition for "job qualifications and characteristics" that includes the criteria an AEDT will assess and the purpose of using such criteria.

The revised proposed rules also remove in-person notice from Section 5-304(b)-(c) as a way for employers to provide prior notice, so employers would only be required to provide notice by U.S. mail or email ten days prior to using an AEDT. Unless notice by U.S. mail reaches the worker ten days prior to the AEDT's use, notice by U.S. mail may not provide for timely notice compared to email, and both methods do not account for workers who may rely on in-person notice

¹ *Id.* at 6-7.

because they may not have a stable mailing address or internet access. Because this revised definition puts already-marginalized workers at a further disadvantage, the Department should restore the in-person notice option.

Conclusion

We appreciate that some of our recommendations are reflected in the revised proposed rules, but other new changes to the rule will further exclude certain types of discriminatory automated assessment methods from transparency requirements and weaken the auditing and notice obligations under LL 144. DCWP should structure the final rules to establish clearer and more effective auditing obligations for the field of automated assessment methods and to make sure workers receive a meaningful opportunity to access accommodations or alternative methods that provide a fairer selection process.

Respectfully submitted,

Ridhi Shetty
Policy Counsel, Privacy & Data Project
rshetty@cdt.org

Matthew Scherer
Sr. Policy Counsel, Privacy & Data Project
mscherer@cdt.org



January 20, 2023

Honorable Eric Adams
New York City Mayor
City Hall
New York, NY 10007

Dear Mayor Adams,

We write to you as a local coalition of civil rights and community organizations who believe that the time for meaningful progress on fair employment is long overdue. We appreciate your leadership on this issue, as demonstrated at the end of last year when you signed legislation to help promote diversity in our fire department.

Today, we want to draw your attention to another new law that has the potential to make New York City a beacon of equality and opportunity: Local Law 144 of 2021, commonly known as the “bias audit” law for automated employment decision tools (AEDTs). The Department of Consumer and Worker Protections (DCWP) released its second version of draft rules for LL-144 on December 23, 2022. Unfortunately, the latest proposed guidelines will fall short in actualizing the spirit of the law.

As the rulemaking process is still ongoing, we strongly encourage DCWP to adopt final rules for LL-144 that will push NYC businesses to prioritize equity and fairness in the hiring process. This law represents a major opportunity for New York City to lead the nation on an important aspect of workforce diversity.

AEDTs have massive consequences for racial and socioeconomic equity because employers rely on them to conduct standardized screenings of job candidates. The nature of these tools can vary from personality tests administered on computers to sophisticated AI platforms. AEDTs are often purchased from third-party vendors, though employers sometimes create them internally.

Since the civil rights era, we have known that certain employment screening tools can perpetuate discrimination. Today, many of the most common assessments employers use to screen job candidates are biased against Black and Brown people. Just as lower scores on the SAT have barred candidates from diverse backgrounds access to higher education, lower performance on hiring assessments has restricted access to economic opportunities.

You may wonder: Why don't employers, especially those who have repeatedly emphasized their commitments to workforce diversity, simply stop using biased AEDTs? The response: Some do not even know if the tools they are using are problematic. The unfortunate reality is that clear information on the extent of bias caused by AEDTs is almost never available to employers, job candidates, or the public. And



Advocates Holding Employers Accountable on Diversity

employers who suspect an issue are hardly incentivized to go looking for evidence of possible discrimination.

The opacity regarding bias in AEDTs is the problem LL-144 seeks to change. If this law is properly implemented, when an employer uses an AEDT to evaluate New Yorkers as job candidates, they will be required to post a document to their career website known as a bias audit. A bias audit will include the results obtained when a third-party auditor tests the AEDT for discriminatory effects, also known as “disparate impact.”

Bias audits are about transparency, as they simply let the public know whether a company’s employment tools are likely to disadvantage particular demographic groups.

As you well know, over the past few years, employers have been making bold statements about their desire to provide historically disadvantaged communities with job opportunities. They have told civil rights advocates that they want to be held accountable to these commitments, especially in the wake of George Floyd’s death in the summer of 2020.

And yet, many of the very same organizations who insisted that Black Lives Matter are now balking at the prospect of rooting out biased hiring tools. The response from Big Business to LL-144 has been combative and reactionary.

Over the past several months, pro-business interests have characterized the law as complex, burdensome and vague. Further, they have encouraged regulators to draft an extremely narrow definition of AEDT in the proposed rules that will render almost all common hiring tools exempt from the statute’s provisions.

The motivations of employers who are attempting to weaken this law are very simple: they are well aware that bias audits will make it abundantly clear when an organization’s stated commitments to workforce diversity ring hollow. For these businesses, criticizing the transparency requirements outlined by this law and calling for large loopholes is a more appealing strategy than abandoning the biased AEDTs they have been using for decades.

As strong supporters of this law and its goal of promoting meaningful transparency, our message to you is that the influence of Big Business on LL-144 is a major concern. We urge your administration to recognize their sentiments as what they are: thinly veiled attempts to render this precedent-setting law meaningless before it goes into effect.

As your administration continues to do the hard work of implementing this groundbreaking legislation in the near term, you will undoubtedly continue to hear complaints from businesses who are resistant to change. We strongly urge you to stand firm in insisting that Local Law 144 will not be diluted during the rulemaking process or at any time after the law goes into effect. There can be no loopholes created or



Advocates Holding Employers Accountable on Diversity

penalties suspended. There can be no allowances made for business leaders who make vague claims about the law being confusing on points they know to be quite clear. If the only thing employers must do to comply is the *transparency* requirements outlined by the law, the question to these organizations should be simple: What are you so eager to hide?

Thank you for your consideration of our perspective. We are proud that New York City is paving the way for equity and fairness in employment, and we look forward to seeing LL-144 implemented in the near term.

Sincerely,

AHEAD

Advocates Holding Employers Accountable for Diversity

Bertha Lewis, The Black Institute

Kirsten John Foy

Bedford Stuyvesant Restoration Corporation Bridge Street Development Corporation Myrtle

Avenue Brooklyn Partnership Freelancers Union

National Black MBA Association (NBMBA) Digital Girl, Inc

East NY Press

Andrew Freedman House

Black Gotham Experience

Universe City NYC

Mastermind Connect

BMS Family Health and Wellness Centers NY for Seniors

Purelements - An Evolution In Dance AfroLatino Fest NYC

Billy Council of CouncilHim

The Andrew Freedman Home

Youth Action Youth Build



Dear Chair and Members of the DCWP,

Recognizing that the final rulemaking for this Law has been an iterative process, as noted in our [previous](#) comments, Credo AI welcomes NYC Law No. 144 (LL-144) as an important step forward to address bias and reduce discrimination in the hiring process when using automated employment decision tools (AEDTs). Credo AI has focused these comments on the Department of Consumer and Worker Protection (DCWP)'s latest clarification rules.

[Credo AI](#) is a private venture backed company that empowers organizations to develop, procure and use AI responsibly (Responsible AI: AI systems held to the highest ethical standards). Our Responsible AI Governance SaaS platform helps companies align their specific AI use case contexts with applicable Responsible AI governance requirements - including LL-144 - to test and evaluate AI systems against those requirements to ensure AI applications are compliant, fair, safe and auditable. We work with a number of companies in the human resources sector that are procuring, building and deploying AEDTs in New York and elsewhere.

Regarding the changes addressed in the latest rulemaking, Credo AI suggests that the DCWP consider the following recommendations:

- **The revised AEDT definition is limiting, and will exclude a number of tools that can propagate historical bias.** The law states that an AEDT is “automated” if it is used “to substantially assist or replace discretionary decision making²²,” and the rulemaking clarifies that this means “(i) to rely solely on a simplified output (score, tag, classification, ranking, etc.), with no other factors considered; (ii) 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 (iii) to use a simplified output to overrule conclusions derived from other factors including human decision-making.” **Credo AI believes that this emphasis on the “weight” of the AEDT decision versus human input will limit the number of AEDTs that actually fall under the limits of the law, since very few organizations actually measure or weigh the amount which an AEDT output contributes to a human decision compared to other factors.** It is also unclear how to validate whether human input in decision-making actually occurred. Ostensibly, a user of a tool could claim that human input was the primary input to the tool's use, rendering the tool not "automated." We agree with suggestions that have been made by other entities that including use case examples (via an appendix or otherwise) which are covered by LL-144 would be very

²² NYC AC § 20-870

useful, since many vendors are unclear whether their tools fall under the jurisdiction of the law. **It is our view that the new DCWP rulemaking has overly-focused on a narrow form of machine learning that will exclude a number of tools that can propagate historical bias.** For instance, a system that does keyword matches between “resume” and “job description,” and uses this as a score of fit, would not fit under the current AEDT definition because it is not trained or tested on any data. LL-144 is a precedent setting regulation that can bring fairness and equity to hiring, and we are concerned that the revised AEDT definition will not affect change.

- **Clarification on the distinction between evaluation of AEDTs (on their own), and hiring processes (incorporating both AEDT output and human decision-making based on the AEDT) for disparate impact would be helpful.** Most AEDTs are used in conjunction with human decision making, so in evaluating the disparate impact of the tool alone, employers may fail to identify disparate impacts caused by a combination of bias in the AEDT and bias in the rest of the hiring process (or compounded bias caused by use of multiple slightly biased tools). For example, an AEDT and human decision-makers might serve different roles in a hiring process. It is possible that the two may yield selection rates that are considered non-discriminatory (they could each individually pass the “4/5ths rule” with an impact ratio of 81%), but when combined to obtain an overall impact ratio, the *system* is considered discriminatory (e.g. with an overall impact ratio of 66%). The reverse is also possible, where a human may compensate for an AEDT with known weak-spots and biases leading to a more effective and unbiased system. The DCWP should clarify whether employers should assess the AEDT alone, the AEDT’s use in a larger decision-making system, or both.
- **The DCWP should consider customizing job categories per AEDT.** Credo AI contends that the current EEO-1 categories are too high level, and not granular enough for most AEDT applications, which may lead to false positives or false negatives detecting disparate impact. Credo AI suggests that the DCWP consider allowing each employer to select the job categories on which their AEDT is evaluated, and then justify those categories as well as any reasons for exclusion.
- **The DCWP should clarify the 10 day notice period for “instant use” AEDTs.** If an AEDT is used to evaluate a candidate’s resume right when it is submitted to a job posting, it is unclear what qualifies as a sufficient notice. Clarification for both employers and vendors would be extremely useful here.
- **The DCWP should clarify the source of demographic data:** While the test versus historical data clarification in the new rule is helpful, further clarification is needed regarding whether using inferred data is acceptable (and what types of inference are preferable); many employers do not collect a sufficient amount of self-reported data across demographic groups in order to avoid bias in the test data set.
- **More can be included in disparate impact calculation:**
 - We appreciate that the impact ratio is defined. We have two suggestions to lend further clarity and accountability. First, the “scoring rate” version of the impact

ratio is potentially overly permissive. Downstream users of candidate's scores rarely "select" people who are above the median score. More often higher thresholds are used, for instance only selecting candidates in the top 90th percentile. A system that is unbiased at the 50th percentile may show bias at the 90th. Given this, we suggest requiring the "scoring" version of the impact ratio to be calculated at a number of quantiles, with special attention paid to the top-end which have the highest chance of affecting downstream decisions. For instance, report impact ratio at 50th, 75th, 90th, and 95th percentiles. Adding this additional information does not increase regulatory burden, as it is easy to calculate these additional statistics using the same data.

- Second, further clarification is needed around how impact ratios should be interpreted when calculated separately for multiple job categories. If an AEDT is employed for many job categories, there will naturally be a distribution of impact ratios. Even when a system is "fair", some job categories may appear "biased" just by chance. We suggest that the law provides some guidance on how multiple impact ratios should be interpreted holistically.

We commend your administration for its hard work bringing this groundbreaking Law to fruition in order to ensure fair employment, and prioritize equity in the hiring process. Credo AI strongly urges the DCWP to consider the above recommendations; we view these clarifications and considerations as critical to enabling the DCWP and the City of New York to deliver on the impact that this Law aspires to have.

In particular, it is our view that the new DCWP rulemaking has overly-focused on a narrow form of machine learning that will exclude a number of tools that can propagate historical bias. LL-144 is a precedent setting Law that can bring fairness and equity to hiring, and we are concerned that the revised AEDT definition will not effect change. We would welcome the opportunity to discuss these recommendations further and share our work.

Thank you in advance for your consideration.

Navrina Singh

Founder & CEO of Credo AI www.credo.ai

Email: navrina@credo.ai



January 23, 2023

To the New York City Department of Consumer and Worker Protection Commissioners:

Thank you for the opportunity to comment on the proposed rules for implementation and enforcement of Local Law 144, regarding the use of automated employment decision tools (AEDTs) in hiring and promotion processes.

Data & Society is an independent, nonprofit research institute studying the social implications of data-centric technologies and automation. We produce empirical research that challenges the power asymmetries created and amplified by technology in society, and work to help ensure that artificial intelligence (AI) systems are accountable to the communities within which they are applied.

Local Law 144 is one of the first laws in the world to mandate an independent audit of any algorithmic systems for bias, and therefore this law and rule-making process has important implications beyond the jurisdiction of the DCWP. Not only are such systems used in employment contexts, they are increasingly used across the economy and government in sensitive domains, such as distribution of social welfare, educational opportunity, housing, and access to financial resources. As is well-documented in scholarly literature, government reports, and investigative journalism, the use of machine learning to train these computational systems is prone to bias against vulnerable and historically disadvantaged groups of people. At their core, these systems learn to replicate the past decisions and behaviors recorded in their training data—these systems predict how we *would have* acted in similar contexts, leaving little room for adjusting how we *should have* acted. Regardless of the efficiency that machine learning systems promise to those who use them, *society has an obligation to ensure that such efficiency is not gained on the backs of vulnerable populations.*

Local Law 144 addresses that obligation by requiring those who deploy these systems in an employment context to transparently account for how their systems behave toward the actual population of job seekers, notifying the public and applicants of their use, and giving applicants the right to request alternative methods. The City of New York is right to pry open this black box for job seekers, and should continue to do so for other domains in the future.

However, as an organization of scholars and policy experts in the social consequences of data technologies, we are concerned that some details of these proposed rules will dramatically blunt the effectiveness of this law and subvert the intent of the New York City Council. We note that these rules may unnecessarily limit the scope of these auditing obligations in three ways:

- 1. Narrowly defining Automated Employment Decision Tools;**
- 2. Misunderstanding how machine learning bias is propagated; and**
- 3. Restricting bias audits to gender and race/ethnicity.**

Defining Automated Employment Decision Tools

The proposed rules define automated decision tools in employment to mean a system substantially assists or replaces discretionary decision making:

- i. to rely solely on a simplified output (score, tag, classification, ranking, etc.), with no other factors considered;
- ii. 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
- iii. to use a simplified output to overrule conclusions derived from other factors including human decision-making.

We are concerned that this definition is so narrow as to exclude the majority of AEDT applications on the market, and misses the core motivations behind LL144.

This definition appears to assume that the biased outcomes that result from AEDTs derive only from *automated* decisions. However, the economic rationale of AEDTs for most employers is not to render a final hiring or promotion decision on the basis of machine learning outputs alone, and the market for such a tool therefore has few (if any) options that operate in a fully automated fashion. Some very large employers, most notably Amazon warehouses, utilize what appear to be fully-automated recruitment software that is developed in-house (though there are few public accounts of how these hiring processes actually work). While those systems certainly deserve scrutiny, and Amazon's many potential warehouse employees deserve the protections offered by LL144, it does not appear that the City Council intended to limit the scope of the law to a very small number of large employers.

Rather, most employers use these systems to create time and economic efficiencies in decision-making by the humans tasked with final hiring or promotion decisions. The practical reality in most cases is that humans still make the final decision from a pool that has been filtered, sorted, scored, and/or narrowed by prior computation. However, this definition restricts the scope of the rules to systems where an algorithmically-generated score is either the sole or predominant factor in the hiring or promotion decision.

Our strong suspicion is that very few employers who one could judiciously say currently deploy an AEDT have a system that would meet reasonable interpretations of this proposed definition. *Most job seekers who are algorithmically scored would not be protected under these rules, which we do not believe was the intent of the City Council in passing LL144.*

At the very least, this proposed definition leaves open the door to subsequent legal challenges that may gut the intent of LL144. In particular, the language in subphrase (ii) raises the question of how stakeholders might measure “weighted more than any other factor”—the ambiguity here leaves open specious legal interpretations that nearly every employer could adopt to avoid conducting audits. Does “weighted more” mean that if the algorithmic score is weighted 49% then it is not open to the scrutiny of an audit? Does it mean that if the automated score is weighted at 20% and eight other factors are weighted individually at 10% and collectively 80%, then the system is subject to the rules? How would the DCWP ask employers to reliably account for the weights that are used?

Furthermore, there are abundant studies examining the highly complex and fraught relationship between human discretion and algorithmic scores. The story that emerges is that a multitude of factors determine to what extent humans utilize discretion when presented with algorithmic scores. Even when prompted to use discretion, humans in organizational contexts where they are otherwise incentivized to trust the computer will follow the algorithmic predictions the vast majority of the time. Given the many different corporate structures, incentives, and internal information systems at private employers over which DCWP has no insight or control, it would seem that the spirit of LL144 requires assuming that *any* system which uses algorithmic scores is potentially a source of algorithmic bias and therefore subject to audit.

We also note that the only change in the wording of this definition between the prior proposed rules (considered for public comment in October, 2022) is removal of the word “modify” from subphrase (iii), replacing “overrule or modify conclusions” with simply “overrule conclusions.” This change again significantly reduces the number of AEDTs that would fall under scope, reducing the practical reach of these rules.

We suggest to the DCWP that the definition of AEDTs used in the rule-making process should hew more closely to that plainly stated in the text of LL144:

“The term “automated employment decision tool” means any computational process, derived from machine learning, statistical modeling, data analytics, or artificial intelligence, that issues simplified output, including a score, classification, or

recommendation, that is used to **substantially assist or replace** discretionary decision making for making employment decisions that impact natural persons.” [emphasis added]

The definition offered in this proposal accounts in practice only for those systems that *replace* discretionary decisions, not for those that *assist* discretionary decisions, and therefore is too dependent on the methods used by the developer and/or the intent of the employer. The simplest and most direct route to providing the protections of job seekers that LL144 plainly intends is to *subject all algorithmic scoring systems to independent audits*. Audit obligations should apply without consideration of the degree to which the employer *intends* to weight those scores.

Misunderstanding of how bias is propagated in machine learning systems

The second significant flaw in these proposed rules is that the definition of “Machine learning, statistical modeling, data analytics, or artificial intelligence” is overly narrow, and as a consequence misunderstands how bias operates in machine learning. Potential routes to biased hiring and promotion decisions could be technically excluded.

The proposed definition is as follows:

Machine learning, statistical modeling, data analytics, or artificial intelligence. “Machine learning, statistical modeling, data analytics, or artificial intelligence” means a group of mathematical, computer based techniques:

- i. that generate a prediction, meaning an expected outcome for an observation, such as an assessment of a candidate’s fit or likelihood of success, or that generate a classification, meaning an assignment of an observation to a group, such as categorizations based on skill sets or aptitude; and
- 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.

The specific error here is found in (ii): “for which a computer at least in part identifies the inputs.” In machine learning, algorithms are used to find patterns in a collection of features (categories of data, such as educational level, degree, years of experience, previous job titles, previous employers, etc.) that statistically indicate a certain outcome is likely to occur (such as a

candidate is likely to be successful in a job role). The pattern that indicates success is then structured as a *model*, a set of mathematical instructions for an application to predict future outcomes based on live inputs (such as the content of new applicant resumes). The efficiency of machine learning is finding the *optimal* arrangements (weights) of features to predict success in the objective function (such as finding a good candidate).

In *some* cases of machine learning, the computer chooses which features/inputs to utilize in building the model. Those methods are often known as “deep learning” wherein a very large unstructured dataset of features—many of which may be facially irrelevant in human judgment—is analyzed by the algorithms to generate an optimized model. Data scientists are rightfully concerned about how deep learning can unintentionally and inscrutably propagate historical biases embedded in their training data. However, deep learning is only likely to comprise a small proportion of the types of machine learning utilized to construct AEDTs because the data available in a resume is already highly structured, labeled, and pre-determined by the expectations of job-seekers and hiring managers. Before machine learning has entered the picture the inputs are already chosen simply because resumes are largely standardized, which means that many AEDTs could be technically excluded by this definition.

Additionally, developers of machine learning systems are often substantially engaged in crafting the models—despite the marketing rhetoric around automation, there is nearly always significant human input and artfulness that goes into shaping these applications and services. It is likely a rare occurrence for the computer to choose relevant features, optimal weights, and parameters alone. The developer's choice of statistical techniques may also introduce opportunities for bias such that even the most rudimentary forms of machine learning result in bias. Similarly, these systems are often customizable by the employer. A hiring manager may manually choose certain weights (defined here as “the relative importance placed on those inputs”) that still drive algorithmically-biased consequences. In other words, especially in machine learning systems meant to intervene in human social processes like AEDTs do, human discretion in the construction of the model is just as likely to introduce bias as deep learning techniques.

Therefore, it is possible that AEDTs which use fairly simple (and very common) machine learning techniques to evaluate candidates on the basis of their resumes would evade this definition.

Of course, some ambitious AEDTs may use features/inputs that require deep learning methods, such as intelligence tests, personality tests, or biometrics. Such applications may require the machine learning system to model the relevant features at a fine granularity, such as the pattern of mouse movement to complete a task or the structure of a person's face when smiling in a

video interview. However, even in those applications, at a gross scale humans are still choosing the relevant inputs, such as efficiency and emotional state. Using the rules as currently proposed leaves the DCWP open to legalistic objections and evasions on this question.

In short, only the actual measurement of bias really matters here—exactly how the system is constructed is largely irrelevant to the question of whether bias may be present.

Simply striking point (ii) in this definition would resolve this error and still leave DCWP with a defensible and adequately capacious definition. Alternatively, the conjunction “and” could be replaced by the disjunction “or” in (ii) to clarify that *any* of those techniques is classified as machine learning.

Restricting bias testing to race and gender

The proposed rules only require bias auditing of gender and race/ethnicity features as defined by the US Federal Equal Opportunity Commission (EEOC). There is good reason to use these standardized categories from EEOC rules, insofar as they are commonly understood, nearly universally collected, and do not generate conflict with other statutes. We affirm that the DCWP is correct to use these categories in the bias audits.

However, we note that LL144 does not specify any requirement to limit bias audits to race and gender features, and therefore leaves the door open to auditing against a more expansive list. The EEOC categories should be a floor, not a ceiling; all AEDTs should be audited for bias along these features, but other biases should be audited for if the system implicates relevant features.

For example, multiple commercial AEDT products analyze audiovisual content of video interviews to predict personal characteristics such as personality or affect. In one audit²³, a group of journalists found that a commercially-available personality profiling AEDT generated significantly different results if candidates wore glasses, changed their background to include a bookshelf, put on a headscarf, or changed their lighting conditions. Obviously, none of these characteristics are correlated with stable personality features predictive of job performance, and thus the product is itself dubious. However, such products can also introduce unexpected biases:

²³ <https://interaktiv.br.de/ki-bewerbung/en/>

affect can be associated with gender and sexual identity, headwear can be correlated with religion, and eyeglasses are a prosthetic to correct for a disability that doesn't affect job performance. Similarly, text-based tests or cognitive tests may be swayed by neurodivergence or

cognitive disability unrelated to job performance and/or amenable to reasonable accommodation as required by the ADA. None of these known, well-demonstrated, and illegal types of bias in AEDTs would be accounted for in the proposed rules.

Our recommendation to the DCWP is that AEDTs should be audited according to the type of bias they are likely to propagate based on the inputs chosen by the developers and deployers.

Developers and deployers of AEDTs are responsible for choosing the features used in these systems. Multiple ethical algorithm design resources for tracking these risks are publicly-available, including resources developed by the National Institute of Science and Technology. The independent auditors mandated by LL144 are capable of identifying such features and the bias risks associated with their use, and responsible AEDT developers already do so. A more expansive audit would not pose an undue burden.

Beyond transparently accounting for bias, this would also promote the desirable consequence of weeding out “algorithmic snake oil” offerings in the AEDT marketplace. Algorithmic systems will always excel at making measurements and offering predictions in a manner that appears useful and economically valuable. But whether those predictions are relevant to the objective function (e.g., job performance), or desirable by society at large (e.g., fair opportunity), is often unanswered. Algorithmic snake oil is a common mode for injecting unfairness into a system because irrelevant measurements can be disguised as objective mathematical judgment.

But there is a simple solution to this problem: if the consequences of including a particular feature cannot be included in a bias audit, then that feature need not be used. Transparent and independent bias audits are one mechanism to force a developer to justify their choice to include certain features and prove that it does not create illegal bias, *but only if the relevant types of bias are accounted for in the audit.*

There is no justifiable reason to treat the EEOC gender and race/ethnicity categories as a ceiling rather than a floor and permanently exclude other, known types of bias. If the audits are restricted to EEOC gender and race/ethnicity categories for now, there should be explicitly stated intent to include more categories in the future as audit methodologies and data collection are improved over time and become routine.

Conclusion

We believe that LL144 is a much-needed and ground-breaking intervention in a market that has been harmful and largely neglected by regulators. However, the proposed rules are too narrow to provide the protections intended by the City Council.

The DCWP should pursue a simple principle in the next round of rule-making: developers and deployers of AEDTs are responsible to measure and transparently report how their systems behave when imposed upon the job-seeking public regardless of how the systems were constructed. This version relies too heavily on the presumption that bias is introduced *by* machines that replace human judgment, when in fact *all algorithmic bias is introduced and/or mitigated only by the choices of the developer and deployer* to engage machine learning for these tasks.

When those choices result in bias, they alone are accountable in each case.

Thank you for the opportunity to include our remarks on this critically important public policy. We hope that the work DPWC is doing on this topic can shape other efforts in the future.

Sincerely,

Jacob Metcalf, PhD
AI on the Ground Initiative, Program Director



VIA E-MAIL

Commissioner Vilda Vera Mayuga
New York City Department of Consumer and Worker Protection
42 Broadway, 8th Floor
New York, NY 10004

**Re: Continued Comments Regarding The Proposed Rules to the Use of Automated
Employment Decision Tools Under Local Law 144 of 2021**

Dear Commissioner Mayuga:

We submit these comments on January 23, 2023 in response to the second public hearing on New York City Local Law 144. After reviewing the revised version of proposed rules issued by the Department of Consumer and Worker Protection regarding automated employment decision tools (AEDT) following the November public hearing, it's apparent that many of the previously submitted elements were integrated, resulting in a more clear explanation of the law and its parameters. In the interest of continuing to refine such language and define the use of AI audits, our additional feedback is shared below.

Working with a broad array of employers, leveraging AI technology for talent acquisition and talent management solutions, retrain.ai has gained a deep knowledge of AI, Responsible AI, and the legal implications of designing, developing and applying sophisticated algorithmic technologies in the workplace. Likewise, retrain.ai's membership in the World Economic Forum includes working together with public- and private-sector leaders to help define and develop the standards for responsible AI.

Our expert data science groups work on Responsible AI requirements across technologies, including the use of artificial intelligence and machine learning algorithms to help with key processes including sourcing, hiring, retention, workforce planning, employee management, talent development and diversity, equity and inclusion (DEI) initiatives.

1. The law needs to better define what constitutes an adequate data sample size, and include an understanding of what procedural changes will be instituted if a sample size

is too small, as with certain ethnicities that have a small representation in the population for example. Without a robust enough data set for analysis, accurate detection of bias is likely impossible. Hence, clearer guidelines for sample size of data are required.

2. The law needs to clarify parameters in the absence of sufficient historical data. The current wording in section 5-302(c) on Data Requirements reads:

“A bias audit of an AEDT used by multiple employers or employment agencies may use the historical data of any employers or employment agencies that use the AEDT. However, an employer or employment agency may rely on a bias audit of an AEDT that uses the historical data of other employers or employment agencies only if it provided historical data from its use of the AEDT to the independent auditor for the bias audit or if it has never used the AEDT.”

But while the law states that a vendor may use historical data of multiple employers or employment agencies that use the AEDT, it doesn't address the challenge presented when other employers or employment agencies won't provide their consent to share private data. At that point, what happens if there is no data from the employer, nor the other employers, due to either a lack of data or a refusal of employers or agencies to share private data? This could be addressed as a situation when no historical data exists.

The law's guidelines suggest that if there is no historical data, an auditor can use test data; but the definition of test data is vague and doesn't include parameters for what is appropriate:

“‘Test data’ means data used to conduct a bias audit that is not historical data.”

“(a) A bias audit conducted pursuant to section 5-301 of this Chapter must use historical data of the AEDT. If insufficient historical data is available to conduct a statistically significant bias audit, test data may be used instead.

(b) If a bias audit uses test data, the summary of results of the bias audit must explain why historical data was not used and describe how the test data used was generated and obtained.”

Data is the basis of the entire audit, therefore ‘test data’ necessitates a much more detailed definition. Without one, there remains a risk of using an inappropriate dataset during a bias audit. This can affect the ability to facilitate the audit properly in order to accurately expose the actual performance of the AEDT.

We at retrain.ai understand that when used responsibly, AI can empower employers to greatly enhance unbiased hiring practices that lead to the proven benefits of a diverse, inclusive workforce. We look forward to the further refinement of Local Law #144 for the betterment of hiring practices not just in New York City, but also beyond our city limits, as countless national and international companies are linked to NYC through business operations based here, many of which require hiring of personnel within the city’s five boroughs. Thank you for including an array of voices in the conversation.

Should the Council have questions or comments about this submission letter, retrain.ai is happy to answer and share our perspective on this important topic.

Sincerely,

retrain.ai Inc.

Comments on the DCWP Proposed Rules implementing LL144 related to Automated Employment Decision Tools

Anupam Datta, Co-Founder & Chief Scientist, TruEra
Shayak Sen, Co-Founder & CTO, TruEra
Will Uppington, Co-Founder & CEO, TruEra

January 22, 2022

Please find below a few comments on the proposed rules for implementing LL144.

- 1. The expanded definition of how an AEDT is used “to substantially assist or replace discretionary decision making” is too restrictive. This definition makes it unlikely that bias audits will be required in very common hiring situations, while leaving those situations susceptible to algorithmic bias.**

Note this definition from the proposed rules: “Automated Employment Decision Tool. “Automated employment decision tool” or “AEDT” means “Automated employment decision tool” as defined by § 20-870 of the Code where the phrase “to substantially assist or replace discretionary decision making” means (i) to rely solely on a simplified output (score, tag, classification, ranking, etc.), with no other factors considered; (ii) 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 (iii) to use a simplified output to overrule conclusions derived from other factors including human decision-making.”

AEDTs are typically used as *one* factor in the hiring process alongside other process steps (e.g., manual review of resumes, or interviews). The final decision is often made by a hiring manager taking into account all of these factors, without necessarily explicitly weighting them. For example, an AEDT that processes a large set of resumes and outputs a smaller list of the top candidates for recruiters and hiring managers to use as a starting point can introduce bias in the decision making process even though humans are making the final decisions taking into account other factors as well.

In this typical setting, the criteria (i), (iii) do not apply; whether (ii) applies or not may be unclear. Thus, the revision in the definition risks ruling out from the bias audit the common case of how AEDTs assist or replace discretionary decision making, and through that introduce potential fairness harms.

The language of the proposed rules in December 2022 was more expansive in that the third criterion was “to use a simplified output to overrule or modify conclusions derived from other factors including human decision-making”, which would bring the typical setting more clearly within the scope of the law.

We recommend that the revised definition of “substantially assist or replace discretionary decision making” be reverted back to the language of the proposed rules in December 2022.

2. List out a non-exhaustive list of use cases that are covered by the law.

An alternative way of clarifying the AEDTs covered and not covered by the law is to include examples in the proposed rules. For example, an AEDT that processes a large set of resumes and outputs a smaller list of the top candidates should be within the scope of a bias audit even if there are other factors that go into the final decision. In contrast, an automated tool that simply organizes and tracks resumes of applicants through the hiring process and enables humans to search through them is likely not covered by the law.

3. Updated definition of impact ratio and scoring rate.

The updated definition of the impact ratio for models that output scores now make use of a scoring rate.

$$\text{Impact Ratio} = \frac{\text{scoring rate for a category}}{\text{scoring rate of the highest scoring category}}$$

“Scoring Rate. “Scoring Rate” means the rate at which individuals in a category receive a score above the sample’s median score, where the score has been calculated by an AEDT.”

This definition is in line with our recommendation submitted through comments on October 21, 2022. It addresses the challenge of scaling with the previous definition based on average scores. We hope to see this language persist in the final draft.



FORHUMANITY

January 23rd, 2023

Re: Local Law #144 rule-making

From Ryan Carrier - Executive Director of ForHumanity
On behalf of ForHumanity's civil-society community of volunteers

ForHumanity²⁴ is a US 501(c)(3) tax-exempt public charity and our mission is to *examine and analyze downside risk associated with the ubiquitous advance of AI, algorithmic and autonomous systems and where possible to engage in risk mitigation to maximize the benefits of these systems... ForHumanity*

ForHumanity created its #employment-and-hiring team in April of 2021. Our intent is to identify the best way to serve the people of New York with the tools that ForHumanity has been developing to facilitate governance, accountability, oversight and trust in AI, algorithmic and autonomous systems.

Independence continues to be a word that the DCWP is defining. The definition is improved, but continued deviations from the numerous tried and tested definitions of “Independence” opens the law up to challenge. ForHumanity advises strict adoptions of existing definitions some of which are listed below. SEC, PCAOB and Sarbanes-Oxley are harmonized. However, it is important to indicate that DCWP updates to “Independence” are welcomed and improved over previous versions.

²⁴ ForHumanity (<https://forhumanity.center/>) is a 501(c)(3) nonprofit entity dedicated to addressing the Ethics, Bias, Privacy, Trust, and Cybersecurity in artificial intelligence and autonomous systems. ForHumanity uses an open and transparent process that draws from a pool of over 1100+ international contributors to construct audit criteria, certification schemes, and educational programs for legal and compliance professionals, educators, Data Auditors, developers, and legislators to mitigate bias, enhance ethics, protect privacy, build trust, improve cybersecurity, and drive accountability and transparency in AI and autonomous systems. ForHumanity works to make AI safe for all people and makes itself available to support government agencies and instrumentalities to manage risk associated with AI and autonomous systems.

This definition is at odds with legal and regulatory definitions from all of the following:

- 1) Sarbanes-Oxley Act of 2002
- 2) EU Digital Services Act of 2022.
- 3) The Public Company Accounting Oversight Board (PCAOB)
- 4) The Securities and Exchange Commission

In DCWP's response to the October hearing, its first point said "to produce an AEDT definition that is focused." ForHumanity believes that DCWP has erred in its interpretation. Respondent after respondent cautioned DCWP NOT to narrow the focus of the law, and to adhere to the spirit of the council's lawmaking by ensuring bias audits were conducted on automated decision making tools impacting hiring and employment. Let us reflect, that Local Law #144 refers to "Automated" employment decision tools. Here is the definition of "Automated" from the Blueprint for an AI Bill of Rights from the Office of Science and Technology Policy at the White House. Any deviation from this definition by DCWP should be viewed as a reckless challenge to Federal level experts.

AUTOMATED SYSTEM: An "automated system" is any system, software, or process that uses computation as whole or part of a system to determine outcomes, make or aid decisions, inform policy implementation, collect data or observations, or otherwise interact with individuals and/or communities. Automated systems include, but are not limited to, systems derived from machine learning, statistics, or other data processing or artificial intelligence techniques, and exclude passive computing infrastructure. "Passive computing infrastructure" is any intermediary technology that does not influence or determine the outcome of decision, make or aid in decisions, inform policy implementation, or collect data or observations, including web hosting, domain registration, networking, caching, data storage, or cybersecurity. Throughout this framework, automated systems that are considered in scope are only those that have the potential to meaningfully impact individuals' or communities' rights, opportunities, or access.

Subsequent efforts to further define AEDTs or the AI, algorithmic or autonomous components such as machine learning have only served to weaken the law considerably.

In this context, ForHumanity has established AEDT Comprehensive Provider Certification Scheme and defined the following 13 categories of AEDTs with a subsequent 58 use cases²⁵. DCWP should adopt a list and definition based approach, and use ForHumanity's classification scheme

1. Job Attraction
2. Recruiting & Hiring
3. Pay & Benefits
4. Onboarding
5. Task Allocation & Scheduling
6. Monitoring & Surveillance
7. Productivity Tools
8. Learning & Development
9. Rewards & Recognition
10. Disciplinary Action
11. Reviews & Coaching
12. Career Progression
13. Separation

The 58 use cases cross-referenced to these 13 categories are posted at the end of this letter in Appendix A.

Definition of "Bias Audit"

The requirements of a bias audit continue to focus unfortunately on "outcomes-only". Bias exists in many forms and throughout the algorithmic lifecycle. We would encourage DCWP to avoid a narrow focus on bias. We have identified three critical areas where Bias can and should be mitigated:

- 1) In Data
- 2) In Architectural inputs to AI, Algorithmic and Autonomous Systems
- 3) In Outcomes

²⁵ Miller, C. L. (2023). 2023 Ultimate automated employment decision system use case reference guide. The Center for Inclusive Change, Lewes, DE. <https://www.InclusiveChange.org/>

Further, if DCWP is interested in mitigating bias to New Yorkers in AEDTs, then the definition of bias should be all encompassing, covering the following sources of Bias²⁶:

1. Statistical Bias
2. Cognitive Bias
3. Non-Response or Technology Barrier Bias

If rules-making continues to narrow the focus of bias type, then New Yorkers will remain largely exposed to bias in AEDTs, and Rule 144 will quickly become meaningless.

DCWP's definition of "Machine learning, statistical modelling, data analytics, or artificial intelligence"

"Machine learning, statistical modelling, data analytics, or artificial intelligence" means a group of mathematical, computer-based techniques:

i. that generate a prediction, meaning an expected outcome for an observation, such as an assessment of a candidate's fit or likelihood of success, or that generate a classification, meaning an assignment of an observation to a group, such as categorizations based on skill sets or aptitude; and

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.

We disagree that the definition of machine learning, statistical modeling, data analytics and artificial intelligence needed further clarification. We also would argue that this definition has no other basis in law or practice in any country, notably it is inconsistent with the EU Artificial Intelligence Act and the recently published white paper from the Office of Science and Technology Policy regarding the AI Bill of Rights.

The new revised definition encourages AEDT providers to reduce the sophistication of their tool and produce lower applications with weaker data analytics. Our research indicates that vendors of AEDTs are already leveraging this language to avoid the law. **Therefore, ForHumanity argues that this further definition is excessive and unnecessary and we strongly recommend striking all references beyond the original definition.**

²⁶ Brown, Carrier, Hickok, Smith - June 2021 Bias Mitigation in DataSets

The field of artificial intelligence and algorithmic systems is broad and defined in a variety of ways. Therefore excessive precision defining machine learning, statistical modeling, data analytics or artificial intelligence will be exploited. Narrowing the definition compromises the intended scope of the law from the City Council - which we believe targeted that most AEDTs shall undergo audits for bias, because the City Council decided (which we agreed with that decision) that many of these tools have multiple forms of bias embedded in them. Nothing about this definition results in mitigating bias and the inclusion of this definition means a lot more tools will go without a bias audit and remain largely unregulated. These changes in definitions will allow the very same harm and damages to New Yorkers to continue, that Rule 144 seeks to prohibit.

In support of Local Law #144.
 Ryan Carrier Executive Director,
 ForHumanity
 ryan@forhumanity.center
 Appendix A

Attract	AI-generated Job descriptions (gender-ized language)
	Hiring ad placements (job boards, social media platforms)
Recruit & Hire	Resume keyword scoring
	Resume scraping / predictive hiring analysis (ML)
	Applicant tracking systems
	Applicant screening chat bot
	Video interviews (one-way, applicant to AI, AI-determined questions, predictive scoring)
	Video interviews (two-way, applicant to human, AI-assisted questions, AI-analyzed, predictive scoring)
	Pre-employment job function tests (non-personality-based)
	Pre-employment personality tests for job fit and longevity predictions
	Video game testing
	Applicant verification / background checks [issues: name spelling, similar name, inaccurate data, out of date data]
	Applicant verification / social media activity analysis
Pay & Benefits	Recommender system (for pay/salary offers, raises, bonuses)
	Automated payroll processing (local, national, global)
	On-demand pay
	Facial recognition clock in/out
	AI-assisted shift swapping
	Automatic pay docking (for exceeding break time)
	Automatic pay docking (for computer inactivity)
	Crowdwork platforms - demand-based pay variability (e.g., Mterk, Deliveroo, Uber)
	Algorithmic pay raise, bonus & incentive determinations (gig workers, customer satisfaction-based pay incentives)

	Nudge technologies (targeted to HR staff - e.g., employees with no raise/promotion in 2 years, no training in 6 mo, no response to applicant in 10 days, excessive employee leave accumulation)
	Nudge technologies (targeted to employees - e.g., flu shots available, flex account total and time left, improved adption of wellness programs, benefits selection support, PTO reminders, company events)
	Mental health chat bots (pulse check on mood, stress level, burn out)

	Reward-based chat bots (thank you gift likes/dislikes)
	Benefits chat bots (answer questions, time off requests, self service changes: address, beneficiary, marital status)
	Anonymous policy violation reporting (discrimination and harassment)
Onboarding & Security / Access	AI-powered onboarding (automated materials distribution, tech provisioning, FAQ chat bot)
	Facial recognition-based security access
Work Design, Task Allocation & Scheduling	Automated decision support tools (managerial support to balance/reallocate workload)
	Automated job design (goal setting, work methods, task significance, job complexity, demands)
	Automated scheduling (e.g., service industry - retail, food service, hospitality, etc.)
	Scheduling nudges (gig-workers)
	AI-assisted team creation based on skills alignment
	Customer self-service tools (employee reductionism)
Performance Monitoring & Surveillance	Automated decision support tools
	Technical supervision
	Behavioral supervision
	File access monitoring (NDA, trade secrets)
	Electronic trackers (wearable tech: warehouses scanners; GPS trackers: delivery drivers, gig workers, public transportation)
	Electronic driver performance monitoring (e.g., speed, breaking, backing up, road conditions, etc.)
	Keyboard and mouse movement trackers (productivity detection)
	Automatic screenshots logging (productivity detection)

	Automatic webcam shots (presence detection)
	Chair sensors (detecting break durations)
	Wristbands to guide workers' hands (warehouse pickers)
	Tracking eye movements to evaluate tiredness
	Bluetooth employee badges (with microphones)
	Nudge technologies (to increase productivity)
	Metabolism monitors & biological measuring devices (high-risk jobs - e.g., high-voltage repairs, harness-based work)
	Social media monitoring (to evaluate code of conduct compliance)
Productivity-Enhancing Tools	Automated decision support tools (e.g., customer-assisted technologies)
	Internal communications tools (when using NLP for employee listening)
	Communications apps used for business purposes on personal devices (WhatsApp, FB Messenger, GoogleMyBusiness)
	Video conferencing (mandatory web cam usage)
	Emotion/engagement analysis (voice and video-based)
	"Que" notifications (call analysis - identifying upselling opportunities, que for emotions, scoring sales agents)
	CRM / algorithmic lead distribution
	Networking badges with visual "shared interest signaling" (badges light up when near-field detects person with similar interest - post event dashboard available to manager)
	Writing assistants (e.g., Grammerly, Textio)
	DEI-based word flagging in Slack/Teams
	Text/story generating tools (e.g., marketing/web copy, internal communications copy) using GPT-3
Learning & Development	Automated decision support tools (recommender systems for training needs)
	Talent intelligence - skill gap analysis and recommendations
	Technology adoption chat bots (FAQ's about new tech implimentations)
	Institutional knowledge question routing through Teams/Slack (NLP-driven, skill-based, routed to "internal experts")
Reward & Recognition (Engagement)	NPL of engagement survey feedback
	Auto-response chat bot used in making suggestions to CEO
	Management nudges (reminders for in-the-moment recognition for staff)

Disciplinary Action	Algorithmic replacing (low-skilled workers/2nd tier of reserve workers; gig-workers/slow replacement)
Reviews & Coaching	Listening tools (intent to quit, intent to unionize, burnout detection)
	Annual review scoring
	AI-assisted mentor matching
	Virtual coaching (for presenters - e.g., training staff)
	Virtual coaching (for managers)
	Sales call listening
	Real-time performance feedback
	Feedback coaching (call center - caller needs, caller emotions, operator emotions)
Career Progression	Recommender system (for promotions based on algorithmically calculated productivity or keywords on resumes)
	Personalized career development tools
Separation	Automated termination decisions (or announcements) (missing productivity targets; gig-workers based on customer satisfaction ratings)
	Automated deviant behavior monitoring and alerts (flight risk chatter, performance, fraud)

Source: Miller, C. L. (2023). 2023 Ultimate automated employment decision system use case reference guide. The Center for Inclusive Change, Lewes, DE. <https://www.InclusiveChange.org/>

Testimony of Julia Stoyanovich before the New York City Department of Consumer and Worker Protection regarding Local Law 144 of 2021 in Relation to Automated Employment Decision Tools (AEDTs)

January 23, 2023

Dear Chair and members of the Department:

My name is Julia Stoyanovich. I hold a Ph.D. in Computer Science from Columbia University in the City of New York. I am an Associate Professor of Computer Science and Engineering at the Tandon School of Engineering, an Associate Professor of Data Science at the Center for Data Science, and the founding Director of the Center for Responsible AI at New York University. In my research and public engagement activities, I focus on incorporating legal requirements and ethical norms, including fairness, accountability, transparency, and data protection, into data-driven algorithmic decision making.²⁷ I teach responsible data science courses to graduate and undergraduate students at NYU.²⁸ In fact, I am speaking to you today from an NYU classroom, and am accompanied by about 90 students who are taking my Responsible Data Science course.

I would like to commend the Department of Consumer and Worker Protection on their continued efforts to make AEDT regulation a reality. And I would like to underscore the exceptional competence and dedication of Irene Byhovsky, Legislative Counsel to the Committee on Technology, who has been an incredible advocate to all New Yorkers in her work on this law from its inception.

For background: I actively participated in the deliberations leading up to the adoption of Local Law 144 of 2021^{29,30} and have carried out several public engagement activities around this law when it was proposed³¹. Informed by my research and by opinions of members of the public, I have written extensively on the auditing and disclosure requirements of this Law, including an

²⁷ See <https://dataresponsibly.github.io/> for information about this work, funded by the National Science Foundation through NSF Awards #1926250, 1934464, and 1922658.

²⁸ All course materials are publicly available at <https://dataresponsibly.github.io/courses/>

²⁹ Testimony of Julia Stoyanovich before New York City Council Committee on Technology regarding Int 30 -2020, November 12, 2020, available at

https://dataresponsibly.github.io/documents/Stoyanovich_Int1894Testimony.pdf

³¹ Public engagement showreel, Int 1894, NYU Center for Responsible AI, December 15, 2022 available at <https://dataresponsibly.github.io/documents/Bill1894Showreel.pdf>

opinion article in the New York Times³² and an article in the Wall Street Journal³³. I have also been teaching members of the public about the impacts of AI and about its use in hiring, most recently by offering a free in-person course at the Queens Public Library called “We are AI”³⁴. (Course materials are available online³⁵.) Based on my background and experience, I would like to make 4 recommendations regarding the enforcement of Local Law 144.

Recommendation 1: Clarify and provide explicit guidance on the “Notice to Candidates and Employees” portion of the law. Disclose information about job qualifications and characteristics for which the AEDT screens in a manner that is comprehensive, specific, understandable, and actionable for job seekers and employees.

Rule making on Local Law 144 has thus far focused almost solely on the bias audit provisions. Bias audits are an important part of the law, but they are by far not the only or, in my opinion, the most important. Notice to Candidates and Employees, if implemented as the law intends, will help get at the validity of predictions made by AEDT. Without sufficient information about job qualifications and characteristics for which the AEDT screens, we risk continuing to legitimize tools that lead to arbitrary and capricious decisions. Further, without sufficient information, these decisions will remain uncontestable by job seekers.

There is evidence to suggest that recommendations of many AEDT are inconsistent and arbitrary³⁶. AEDTs that don’t work hurt job seekers and employees, subjecting them to arbitrary decision-making with no recourse. AEDTs that don’t work also hurt employers, they waste money paying for software that doesn't work, and miss out on many well-qualified candidates based on a self-fulfilling prophecy delivered by a tool.

I recommend showing job seekers and employees simple, standardized labels that list the factors that go into the AEDT’s decision both before they are screened and after a decision is

³² We need laws to take on racism and sexism in hiring technology, Alexandra Reeve Givens, Hilke Schellmann and Julia Stoyanovich, The New York Times, March 17, 2021, available at <https://www.nytimes.com/2021/03/17/opinion/ai-employment-bias-nyc.html>

³³ Hiring and AI: Let job candidates know why they were rejected, Julia Stoyanovich, The Wall Street Journal Reports: Leadership, September 22, 2021, available at <https://www.wsj.com/articles/hiring-job-candidates-ai-11632244313>

³⁴ “We are AI” series by NYU Tandon Center for Responsible AI and Queens Public Library helps citizens take control of tech, March 14 2022, available at <https://engineering.nyu.edu/news/we-are-ai-series-nyu-tdandon-center-responsible-ai-queens-public-library>

³⁵ “We are AI: Taking control of technology”, NYU Center for Responsible AI, available <https://dataresponsibly.github.io/we-are-ai/>

³⁶ “Resume Format, LinkedIn URLs and Other Unexpected Influences on AI Personality Prediction in Hiring: Results of an Audit,” Rhea et al., AAAI/ACM AIES 2022, available at <https://dl.acm.org/doi/10.1145/3514094.3534189>

made. Job seekers, employees, and their representatives should be directly involved in the design and testing of such labels.

Figure 1 gives an example of a possible “posting label” with a short and clear summary of the screening process. (See my recent Wall Street Journal article for details⁶.) This label is presented to a job seeker before they apply, supporting informed consent, allowing them to opt out of components of the process, or to request accommodations. Giving an opportunity to request accommodations is particularly important in light of the recent guidance by the Equal Employment Opportunity Commission on the Americans with Disabilities Act and the use of AI to assess job applicants and employees³⁷.

ACCOUNTANT	
Acme Partners	
<hr/>	
Qualifications:	BS in accounting, GPA >3.0, Knowledge of financial and accounting systems and applications
Personal data to be analyzed:	An AI program could be used to review and analyze the applicant's personal data online, including LinkedIn profile, social media accounts and credit score.
Additional assessment:	AI-assisted personality scoring
<hr/>	
ALERT: Applicants for this position DO NOT have the option to selectively decline use of AI analysis for any of their personal data or to review and challenge the results of such analysis.	

Figure 1: A posting label is a short, simple, and clear summary of the screening process. This label is presented to a job seeker before they apply, supporting informed consent, allowing them to opt out of components of the process or to request accommodations.

Recommendation 2: Expand the scope of auditing beyond bias to also interrogate whether the AEDTs work, based on input from all key stakeholders, including job seekers, employees, and their representatives.

In my own work, done in collaboration with an interdisciplinary team that included several data scientists, a sociologist, an industrial-organizational (I-O) psychologist, and an investigative journalist, I evaluated the validity of two algorithmic personality tests: AEDTs that are used for pre-employment assessment⁹, Humantic AI and Crystal. Importantly, rather than challenging or

³⁷ The Americans with Disabilities Act and the use of software, algorithms, and AI to assess job applicants and employees, US Equal Employment Opportunity Commission, 2022, <https://www.eeoc.gov/laws/guidance/americans-disabilities-act-and-use-software-algorithms-and-artificialintelligence>

affirming the assumptions made in psychometric testing — that personality traits are meaningful and measurable constructs, and that they are indicative of future success on the job— we framed our methodology around testing the assumptions made by the vendors themselves.

We found that both systems show substantial instability on key facets of measurement, and so cannot be considered valid testing instruments. For example, Crystal frequently computes different personality profiles if the same resume is given in PDF vs. in raw text, while Humantic AI gives different personality profiles on a LinkedIn profile vs. a resume of the same job seeker, violating the assumption that the output of a personality test is stable across job-irrelevant input variations. Results are summarized in Table 1. Such tools cannot be allowed to proliferate, and Local Law 144 should help protect candidates and employers from their use!

Facet	Crystal	Humantic
Resume file format	✗	✓
LinkedIn URL in resume	?	✗
Source context	✗	✗
Algorithm-time / immediate	✓	✓
Algorithm-time / 31 days	✓	✗
Participant-time / LinkedIn	✗	✗
Participant-time / Twitter	N/A	✓

Table 1: Summary of stability results for Crystal and Humantic AI: ✓ indicates sufficient rank-order stability in all traits, while ✗ indicates insufficient rank-order stability or significant locational instability in at least one trait, and N/A indicates the facet was not tested in our audit. Results are detailed in <https://dl.acm.org/doi/10.1145/3514094.3534189>.

Recommendation 3: Involve job seekers, employees, and their representatives in defining standards for AEDT audits and notices.

Local Law 144 is an incredible opportunity for New York City to lead by example, but only if this law is enacted in a way that is responsive to the needs of all key stakeholders. The conversation thus far has been dominated by the voices of commercial entities, especially by AEDT vendors and organizations that represent them, but also by employers who use AEDT, and by commercial entities wishing to conduct AEDT audits. However, as is evident from the fact that we are testifying in front of the Department of Consumer and Worker Protection, the main stakeholder group Local Law 144 aims to protect – from unlawful discrimination, and arbitrary and capricious decision-making – are job candidates and employees. And yet, their voices haven’t been heard prominently in the conversation!

New York City must ensure active participation of a diverse group of job seekers, employees and their representatives in both rule making and enactment of Local Law 144. The NYU Center

for Responsible AI (R/AI) conducted numerous public engagement activities under my leadership, both broadly on AI and automated decision making, and specifically on AEDTs, and we see substantial interest from members of the public. R/AI will be happy to assist the City in convening diverse groups of stakeholders.

Recommendation 4: Expand the scope of auditing for bias beyond disparate impact to include other dimensions of discrimination, based on input from all key stakeholders, including job seekers, employees, and their representatives.

For example, the most prominent thread in readers' comments on a New York Times opinion piece I co-authored in March 2021, entitled "We need laws to take on racism and sexism in hiring technology" concerned age-based discrimination in hiring. Local Law 144 does not currently include any provisions to audit for this type of discrimination. This is problematic! It is also problematic that intersectional discrimination is not considered in scope. And that inclusion of individuals with disabilities has received no attention during rule making.

To conclude, I would like to keep my testimony today brief. I am enclosing a copy of the testimony I entered on June 6, 2022 and a copy of the testimony I entered on November 4, 2022, for additional background on Automated hiring systems, and for details regarding my recommendations on rules for auditing and notice (disclosure) requirements of Local Law 144.

Testimony of Julia Stoyanovich before the New York City Department of Consumer and Worker Protection regarding Local Law 144 of 2021 in Relation to Automated Employment Decision Tools (AEDTs)

November 4, 2022

Dear Chair and members of the Department:

My name is Julia Stoyanovich. I hold a Ph.D. in Computer Science from Columbia University in the City of New York. I am an Associate Professor of Computer Science and Engineering at the Tandon School of Engineering, an Associate Professor of Data Science at the Center for Data Science, and the founding Director of the Center for Responsible AI at New York University. In my research and public engagement activities, I focus on incorporating legal requirements and ethical norms, including fairness, accountability, transparency, and data protection, into data-driven algorithmic decision making.³⁸ I teach responsible data science courses to graduate and undergraduate students at NYU.³⁹ Most importantly, I am a devoted and proud New Yorker.

I would like to commend New York City on taking on the ambitious task of overseeing the use of automated decision systems in hiring. I see Local Law 144 as an incredible opportunity for the City to lead by example, but only if this law is enacted in a way that is responsive to the needs of all key stakeholders. The conversation thus far has been dominated by the voices of commercial entities, especially by AEDT vendors and organizations that represent them, but also by employers who use AEDT, and by commercial entities wishing to conduct AEDT audits. However, as is evident from the fact that we are testifying in front of the Department of Consumer and Worker Protection, the main stakeholder group Local Law 144 aims to protect – from unlawful discrimination, and arbitrary and capricious decision-making – are job candidates and employees. And yet, their voices haven't been heard prominently in the conversation!

As an academic and an individual with no commercial interests in AEDT development, use, or auditing, I am making my best effort to speak today to represent the interests of the job candidates, employees, and the broader public. However, I cannot speak on behalf of this diverse group alone. Therefore, my main recommendation today is that New York City must ensure active participation of a diverse group of job seekers, employees and their representatives in both rule making and enactment of Local Law 144.

³⁸ See <https://dataresponsibly.github.io/> for information about this work, funded by the National Science Foundation through NSF Awards #1926250, 1934464, and 1922658.

³⁹ All course materials are publicly available at <https://dataresponsibly.github.io/courses/>

For background: I actively participated in the deliberations leading up to the adoption of Local Law 144 of 2021⁴⁰⁴¹ and have carried out several public engagement activities around this law when it was proposed⁴². Informed by my research and by opinions of members of the public, I have written extensively on the auditing and disclosure requirements of this Law, including an opinion article in the New York Times⁴³ and an article in the Wall Street Journal⁴⁴. I have also been teaching members of the public about the impacts of AI and about its use in hiring, most recently by offering a free in-person course at the Queens Public Library called “We are AI”⁴⁵. (Course materials are available online⁴⁶.) Based on my background and experience, I would like to make 4 recommendations regarding the enforcement of Local Law 144.

Recommendation 1: Involve job seekers, employees, and their representatives in defining standards for AEDT audits and notices.

The NYU Center for Responsible AI (R/AI) conducted numerous public engagement activities under my leadership, both broadly on AI and automated decision making, and specifically on AEDTs, and we see substantial interest from members of the public. R/AI will be happy to assist the City in convening diverse groups of stakeholders.

Recommendation 2: My second recommendation is about the extremely important component of Local Law 144 - the bias audit requirement. I recommend that we expand the scope of auditing for bias beyond disparate impact to include other dimensions of discrimination, based on input from

⁴⁰ Testimony of Julia Stoyanovich before New York City Council Committee on Technology regarding Int

⁴¹ -2020, November 12, 2020, available at

https://dataresponsibly.github.io/documents/Stoyanovich_Int1894Testimony.pdf

⁴² Public engagement showreel, Int 1894, NYU Center for Responsible AI, December 15, 2022 available at

<https://dataresponsibly.github.io/documents/Bill1894Showreel.pdf>

⁴³ We need laws to take on racism and sexism in hiring technology, Alexandra Reeve Givens, Hilke

Schellmann and Julia Stoyanovich, The New York Times, March 17, 2021, available at

<https://www.nytimes.com/2021/03/17/opinion/ai-employment-bias-nyc.html>

⁴⁴ Hiring and AI: Let job candidates know why they were rejected, Julia Stoyanovich, The Wall Street Journal

Reports: Leadership, September 22, 2021, available at <https://www.wsj.com/articles/hiring-job-candidates-ai-11632244313>

⁴⁵ “We are AI” series by NYU Tandon Center for Responsible AI and Queens Public Library helps citizens take control of tech, March 14 2022, available at

<https://engineering.nyu.edu/news/we-are-ai-series-nyu-tdandon-center-responsible-ai-queens-public-library>

⁴⁶ “We are AI: Taking control of technology”, NYU Center for Responsible AI, available

<https://dataresponsibly.github.io/we-are-ai/>

all key stakeholders, including job seekers, employees, and their representatives.

For example, the most prominent thread in readers' comments on a New York Times opinion piece I co-authored in March 2021, entitled "We need laws to take on racism and sexism in hiring technology" concerned age-based discrimination in hiring. Local Law 144 does not currently control include any provisions to audit for this type of discrimination. This is problematic!

Recommendation 3: Expand the scope of auditing beyond bias to also interrogate whether the AEDTs work, based on input from all key stakeholders, including job seekers, employees, and their representatives.

There is evidence to suggest that recommendations of many of these tools are inconsistent and arbitrary⁴⁷. AEDTs that don't work hurt job seekers and employees, subjecting them to arbitrary decision-making with no recourse. AEDTs that don't work also hurt employers, they waste money paying for software that doesn't work, and miss out on many well-qualified candidates based on a self-fulfilling prophecy delivered by a tool.

In my own work, done in collaboration with an interdisciplinary team that included several data scientists, a sociologist, an industrial-organizational (I-O) psychologist, and an investigative journalist, I evaluated the validity of two algorithmic personality tests: AEDTs that are used for pre-employment assessment⁹, Humantic AI and Crystal. Importantly, rather than challenging or affirming the assumptions made in psychometric testing — that personality traits are meaningful and measurable constructs, and that they are indicative of future success on the job— we framed our methodology around testing the assumptions made by the vendors themselves.

We found that both systems show substantial instability on key facets of measurement, and so cannot be considered valid testing instruments. For example, Crystal frequently computes different personality profiles if the same resume is given in PDF vs. in raw text, while Humantic AI gives different personality profiles on a LinkedIn profile vs. a resume of the same job seeker, violating the assumption that the output of a personality test is stable across job-irrelevant input variations. Results are summarized in Table 1. Such tools cannot be allowed to proliferate, and Local Law 144 should help protect candidates and employers from their use!

⁴⁷ "Resume Format, LinkedIn URLs and Other Unexpected Influences on AI Personality Prediction in Hiring: Results of an Audit," Rhea et al., AAAI/ACM AIES 2022, available at <https://dl.acm.org/doi/10.1145/3514094.3534189>

Facet	Crystal	Humantic
Resume file format	X	✓
LinkedIn URL in resume	?	X
Source context	X	X
Algorithm-time / immediate	✓	✓
Algorithm-time / 31 days	✓	X
Participant-time / LinkedIn	X	X
Participant-time / Twitter	N/A	✓

Table 1: Summary of stability results for Crystal and Humantic AI: ✓ indicates sufficient rank-order stability in all traits, while X indicates insufficient rank-order stability or significant locational instability in at least one trait, and N/A indicates the facet was not tested in our audit. Results are detailed in <https://dl.acm.org/doi/10.1145/3514094.3534189>.

Recommendation 4: Disclose information about job qualifications and characteristics for which the AEDT screens in a manner that is comprehensive, specific, understandable, and actionable for job seekers and employees.

I recommend showing job seekers and employees simple, standardized labels that list the factors that go into the AEDT’s decision both before they are screened and after a decision is made. Job seekers, employees, and their representatives should be directly involved in the design and testing of such labels.

Figure 1 gives an example of a possible “posting label” with a short and clear summary of the screening process. (See my recent Wall Street Journal article for details⁶.) This label is presented to a job seeker before they apply, supporting informed consent, allowing them to opt out of components of the process, or to request accommodations. Giving an opportunity to request accommodations is particularly important in light of the recent guidance by the Equal Employment Opportunity Commission on the Americans with Disabilities Act and the use of AI to assess job applicants and employees⁴⁸.

⁴⁸ The Americans with Disabilities Act and the use of software, algorithms, and AI to assess job applicants and employees, US Equal Employment Opportunity Commission, 2022, <https://www.eeoc.gov/laws/guidance/americans-disabilities-act-and-use-software-algorithms-and-artificialintelligence>

ACCOUNTANT

Acme Partners

Qualifications: BS in accounting, GPA >3.0, Knowledge of financial and accounting systems and applications

Personal data to be analyzed: An AI program could be used to review and analyze the applicant's personal data online, including LinkedIn profile, social media accounts and credit score.

Additional assessment: AI-assisted personality scoring

ALERT: Applicants for this position DO NOT have the option to selectively decline use of AI analysis for any of their personal data or to review and challenge the results of such analysis.

Figure 1: *A posting label is a short, simple, and clear summary of the screening process. This label is presented to a job seeker before they apply, supporting informed consent, allowing them to opt out of components of the process or to request accommodations.*

I would like to keep my testimony today brief. I am enclosing a copy of the testimony I entered on June 6, 2022, for additional background on Automated hiring systems, and for details regarding my recommendations on rules for auditing and notice (disclosure) requirements of Local Law 144.

Testimony of Julia Stoyanovich before the New York City Department of Consumer and Worker Protection regarding Local Law 144 of 2021 in Relation to Automated Employment Decision Tools

June 6, 2022

Dear Chair and members of the Department:

My name is Julia Stoyanovich. I hold a Ph.D. in Computer Science from Columbia University. I am an Associate Professor of Computer Science and Engineering at the Tandon School of Engineering, and an Associate Professor of Data Science at the Center for Data Science, and the founding Director of the Center for Responsible AI at New York University. In my research and public engagement activities, I focus on incorporating legal requirements and ethical norms, including fairness, accountability, transparency, and data protection, into data-driven algorithmic decision making.⁴⁹ I teach responsible data science courses to graduate and undergraduate students at NYU.⁵⁰ Most importantly, I am a devoted and proud New Yorker.

I actively participated in the deliberations leading up to the adoption of Local Law 144 of 2021^{51,52} and have carried out several public engagement activities around this law when it was proposed⁵³. Informed by my research and by opinions of members of the public, I have written extensively on the auditing and disclosure requirements of this Law, including an opinion article in the New York Times⁵⁴ and an article in the Wall Street Journal⁵⁵. I have also been teaching members of the public about the impacts of AI and about its use in hiring, most recently by

⁴⁹ See <https://dataresponsibly.github.io/> for information about this work, funded by the National Science Foundation through NSF Awards #1926250, 1934464, and 1922658.

⁵⁰ All course materials are publicly available at <https://dataresponsibly.github.io/courses/>

⁵¹ Testimony of Julia Stoyanovich before New York City Council Committee on Technology regarding Int

⁵² -2020, November 12, 2020, available at

https://dataresponsibly.github.io/documents/Stoyanovich_Int1894Testimony.pdf

⁵³ Public engagement showreel, Int 1894, NYU Center for Responsible AI, December 15, 2022 available at

<https://dataresponsibly.github.io/documents/Bill1894Showreel.pdf>

⁵⁴ We need laws to take on racism and sexism in hiring technology, Alexandra Reeve Givens, Hilke Schellmann and Julia Stoyanovich, The New York Times, March 17, 2021, available at

<https://www.nytimes.com/2021/03/17/opinion/ai-employment-bias-nyc.html>

⁵⁵ Hiring and AI: Let job candidates know why they were rejected, Julia Stoyanovich, The Wall Street Journal Reports: Leadership, September 22, 2021, available at <https://www.wsj.com/articles/hiring-job-candidates-ai-11632244313>

offering a free in-person course at the Queens Public Library called “We are AI”⁵⁶. Course materials are available online⁵⁷.

In my statement today I would like to make three recommendations regarding the enforcement of Local Law 144 of 2021:

1. **Auditing:** The scope of auditing for bias should be expanded beyond disparate impact to include other dimensions of discrimination, and also contain information about a tool’s effectiveness - about whether a tool works. Audits should be based on a set of uniform publicly available criteria.
2. **Disclosure:** Information about job qualifications or characteristics for which the tool screens the job seeker should be disclosed to them in a manner that is comprehensible and actionable. Specifically, job seekers should see simple, standardized labels that show the factors that go into the AI’s decision both before they apply and after a decision on their application is made.
3. **An informed public:** To be truly effective, this law requires an informed public. I recommend that New York City invests resources into informing members of the public about data, algorithms, and automated decision making, using hiring ADS as a concrete and important example.

In what follows, I will give some background on automated hiring systems, and will then expand on each of my recommendations.

Automated hiring systems

Since the 1990s, and increasingly so in the last decade, commercial tools are being used by companies large and small to hire more efficiently: source and screen candidates faster and with less paperwork, and successfully select candidates who will perform well on the job. These tools are also meant to improve efficiency for the job applicants, matching them with relevant positions, allowing them to apply with a click of a button, and facilitating the interview process.

⁵⁶ “We are AI” series by NYU Tandon Center for Responsible AI and Queens Public Library helps citizens take control of tech, March 14 2022, available at <https://engineering.nyu.edu/news/we-are-ai-series-nyu-tandon-center-responsible-ai-queens-public-library>

⁵⁷ “We are AI: Taking control of technology”, NYU Center for Responsible AI, available <https://dataresponsibly.github.io/we-are-ai/>

In their 2018 report, Bogen and Rieke⁵⁸ describe the hiring process from the point of view of an employer as a series of decisions that form a funnel: “Employers start by *sourcing* candidates, attracting potential candidates to apply for open positions through advertisements, job postings, and individual outreach. Next, during the *screening* stage, employers assess candidates—both before and after those candidates apply—by analyzing their experience, skills, and characteristics. Through *interviewing* applicants, employers continue their assessment in a more direct, individualized fashion. During the *selection* step, employers make final hiring and compensation determinations.” Importantly, while a comprehensive survey of the space lacks, we have reason to believe that automated hiring tools are in broad use in all stages of the hiring process.

Despite their potential to improve efficiency for both employers and job applicants, hiring ADS are also raising concerns. I will recount two well-known examples here.

Sourcing: One of the earliest indications that there is cause for concern came in 2015, with the results of the AdFisher study out of Carnegie Mellon University⁵⁹ that was broadly circulated by the press⁶⁰. Researchers ran an experiment, in which they created two sets of synthetic profiles of Web users who were the same in every respect — in terms of their demographics, stated interests, and browsing patterns — with a single exception: their stated gender, male or female. In one experiment, the AdFisher tool stimulated an interest in jobs in both groups, and showed that Google displays ads for a career coaching service for high-paying executive jobs far more frequently to the male group (1,852 times) than to the female group (318 times). This brings back memories of the time when it was legal to advertise jobs by gender in newspapers. This practice was outlawed in the US 1964, but it persists in the online ad environment.

Screening: In late 2018 it was reported that Amazon’s AI resume screening tool, developed with the stated goal of increasing workforce diversity, in fact did the opposite thing: the system taught itself that male candidates were preferable to female candidates.⁶¹ It penalized resumes

⁵⁸ Bogen and Rieke, “*Help Wanted: An Examination of Hiring Algorithms, Equity, and Bias*”, Upturn, (2018) <https://www.upturn.org/static/reports/2018/hiring-algorithms/files/Upturn%20-%20Help%20Wanted%20An%20Exploration%20of%20Hiring%20Algorithms,%20Equity%20and%20Bias.pdf>

⁵⁹ Datta, Tschantz, Datta, “*Automated experiments on ad privacy settings*”, Proceedings of Privacy Enhancing Technology (2015) <https://content.sciendo.com/view/journals/popets/2015/1/article-p92.xml>

⁶⁰ Gibbs, “*Women less likely to be shown ads for high-paid jobs on Google, study shows*”, The Guardian (2015) <https://www.theguardian.com/technology/2015/jul/08/women-less-likely-ads-high-paid-jobs-google-study>

⁶¹ Dastin, “*Amazon scraps secret AI recruiting tool that showed bias against women*”, Reuters (2018)

<https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scrap-secret-ai-recruiti>

that included the word “women’s,” as in “women’s chess club captain,” and downgraded graduates of two all-women’s colleges. These results aligned with, and reinforced, a stark gender imbalance in the workforce at Amazon and other platforms, particularly when it comes to technical roles.

Numerous other cases of discrimination based on gender, race, and disability status during screening, interviewing, and selection stages have been documented in recent reports⁶²⁶³.

These and other examples show that, if left unchecked, automated hiring tools will replicate, amplify, and normalize results of historical discrimination.

Recommendation 1: Expanding the scope of auditing

Bias audits should take a broader view, going beyond disparate impact when considering fairness of outcomes. Others surely spoke to this point, and I will not dwell on it here. Instead, I will focus on another important dimension of due process that is closely linked to discrimination — substantiating the use of particular features in decision-making.

Regarding the use of predictive analytics to screen candidates, Jenny Yang states: “Algorithmic screens do not fit neatly within our existing laws because algorithmic models aim to identify statistical relationships among variables in the data whether or not they are understood or job related.[...] Although algorithms can uncover job-related characteristics with strong predictive power, they can also identify correlations arising from statistical noise or undetected bias in the training data. Many of these models do not attempt to establish cause-and-effect relationships, creating a risk that employers may hire based on arbitrary and potentially biased correlations.”⁶⁴

In other words, identifying what features are impacting a decision is important, but it is insufficient to alleviate due process and discrimination concerns. I recommend that an audit of an automated hiring tool should also include information about the job relevance of these features.

A subtle but important point is that even features that can legitimately be used for hiring may capture information differently for different population groups. For example, it has been

[ng-tool-that-showed-bias-against-women-idUSKCN1MK08G](#)

⁶² Emerging Technology from the arXiv, “*Racism is Poisoning Online Ad Delivery, Says Harvard Professor*”, MIT Technology Review (2013) <https://www.technologyreview.com/s/510646/racism-is-poisoning-online-ad-delivery-says-harvard-professor/>

⁶³ Stains, “*Are Workplace Personality Tests Fair?*”, Wall Street Journal (2014) <http://www.wsj.com/articles/are-workplace-personality-tests-fair-1412044257>

⁶⁴ Yang, “*Ensuring a Future that Advances Equity in Algorithmic Employment Decisions*”, Urban Institute (2020) <https://www.urban.org/research/publication/ensuring-future-advances-equity-algorithmic-employment-decisions>

documented that the mean score of the math section of the SAT (Scholastic Assessment Test) differs across racial groups, as does the shape of the score distribution.⁶⁵ These disparities are often attributed to racial and class inequalities encountered early in life, and are thought to present persistent obstacles to upward mobility and opportunity.

Some automated hiring tools used today claim to predict job performance by analyzing an interview video for body language and speech patterns. Arvind Narayanan refers to tools of this kind as “fundamentally dubious” and places them in the category of AI snake oil.⁶⁶ The premise of such tools, that (a) it is possible to predict social outcomes based on a person's appearance or demeanor and (b) it is ethically defensible to try, reeks of scientific racism and is at best an elaborate random number generator.

The AI snake oil example brings up a related point: that an audit should also evaluate the effectiveness of the tool. Does the tool work? Is it able to identify promising job candidates better than a random coin flip? What were the specific criteria for the evaluation, and what evaluation methodology was used? Was the tool's performance evaluated on a population with demographic and other characteristics that are similar to the New York City population on which it will be used? Without information about the statistical properties of the population on which the tool was trained (in the case of machine learning) and validated, we cannot know whether the tool will have similar performance when deployed.⁶⁷

In my own work, I recently evaluated the validity of two algorithmic personality tests that are used by employers for pre-employment assessment⁶⁸. This work was done by a large interdisciplinary team that included several data scientists, a sociologist, an industrial-organizational (I-O) psychologist, and an investigative journalist. My colleagues and I developed a methodology for an external audit of stability of algorithmic personality tests, and used it to audit two systems, Humantic AI and Crystal. Importantly, rather than challenging or affirming the assumptions made in psychometric testing — that personality traits are meaningful and measurable constructs, and that they are indicative of future success on the job— we framed our methodology around testing the underlying assumptions made by the vendors of the algorithmic personality tests themselves.

⁶⁵ Reeves and Halikias “Race gaps in SAT scores highlight inequality and hinder upward mobility”, Brookings (2017)

<https://www.brookings.edu/research/race-gaps-in-sat-scores-highlight-inequality-and-hinder-upward-mobility>

⁶⁶ Narayanan, “How to recognize AI snakeoil” (2019)

<https://www.cs.princeton.edu/~arvindn/talks/MIT-STS-AI-snakeoil.pdf>

⁶⁷ Stoyanovich and Howe, “Follow the data: Algorithmic transparency starts with data transparency”

(2019) <https://ai.shorensteincenter.org/ideas/2018/11/26/follow-the-data-algorithmic-transparency-starts-with-data-a-transparency>

⁶⁸ An external stability audit of framework to test the validity of personality prediction in AI hiring, Rhea et al., 2022, available at <https://arxiv.org/abs/2201.09151>

In our audits of Humantic AI and Crystal, we found that both systems show substantial instability on key facets of measurement, and so cannot be considered valid testing instruments. For example, Crystal frequently computes different personality scores if the same resume is given in PDF vs. in raw text, while Humantic AI gives different personality scores on a LinkedIn profile vs. a resume of the same job seeker. This violated the assumption that the output of a personality test is stable across job-irrelevant input variations. Among other notable findings is evidence of persistent — and often incorrect — data linkage by Humantic AI. A summary of our results are presented in **Table 1**.

Facet	Crystal	Humantic
Resume file format	X	✓
LinkedIn URL in resume	?	X
Source context	X	X
Algorithm-time / immediate	✓	✓
Algorithm-time / 31 days	✓	X
Participant-time / LinkedIn	X	X
Participant-time / Twitter	N/A	✓

Table 1: Summary of stability results for Crystal and Humantic AI, with respect to facets of measurement: ✓ indicates sufficient rank-order stability in all traits, while X indicates insufficient rank-order stability or significant locational instability in at least one trait, and N/A indicates the facet was not tested in our audit. Results are detailed in <https://arxiv.org/abs/2201.09151>.

In summary, I recommend that the scope of auditing for bias should be expanded beyond disparate impact to include other dimensions of discrimination, and also contain information about a tool’s effectiveness. To support compliance and enable a comparison between tools during procurement, these audits should be based on a set of uniform criteria. To enable public input and deliberation, these criteria should be made publicly available.

Recommendation 2: Explaining decisions to the job applicant

Information about job qualifications or characteristics that the tool uses for screening should be provided in a manner that allows the job applicant to understand, and, if necessary, correct and contest the information. As I argued in Recommendation 1, it is also important to disclose why these specific qualifications and characteristics are considered job relevant.

I recommend that explanations for job seekers are built around the popular nutritional label metaphor, drawing an analogy to the food industry, where simple, standardized labels convey information about the ingredients and production processes.²⁰

²⁰ Stoyanovich and Howe, “*Nutritional labels for data and models*”, IEEE Data Engineering Bulletin 42(3): 13-23 (2019) <http://sites.computer.org/debull/A19sept/p13.pdf>

An applicant-facing nutritional label for an automated hiring system should be comprehensible: short, simple, and clear. It should be consultative, providing actionable information. Based on such information, a job applicant may, for example, take a certification exam to improve their chances of being hired for this or similar position in the future. Labels should also be comparable: allowing a job applicant to easily compare their standing across vendors and positions, and thus implying a standard.

Nutritional labels are a promising metaphor for other types of disclosure, and can be used to represent the process or the result of an automated hiring system for auditors, technologists, or employers.⁶⁹

ACCOUNTANT	
Acme Partners	
Qualifications:	BS in accounting, GPA >3.0, Knowledge of financial and accounting systems and applications
Personal data to be analyzed:	An AI program could be used to review and analyze the applicant's personal data online, including LinkedIn profile, social media accounts and credit score.
Additional assessment:	AI-assisted personality scoring
ALERT: Applicants for this position DO NOT have the option to selectively decline use of AI analysis for any of their personal data or to review and challenge the results of such analysis.	

⁶⁹ Stoyanovich, Howe, Jagadish, “*Responsible Data Management*”, PVLDB 13(12): 3474-3489 (2020) <https://dataresponsibly.github.io/documents/mirror.pdf>

Figure 1: *A posting label is a short, simple, and clear summary of the screening process. This label is presented to a job seeker before they apply, supporting informed consent, allowing them to opt out of components of the process or to request accommodations.*

Figure 1 shows a posting label, a short and clear summary of the screening process. This label is presented to a job seeker before they apply, supporting informed consent, allowing them to opt out of components of the process or to request accommodations. Giving job seekers an opportunity to request accommodations is particularly important in light of the recent guidance by the Equal Employment Opportunity Commission (EEOC) on the Americans with Disabilities Act and the use of AI to assess job applicants and employees ⁷⁰.

If a job seeker applies for the job but isn't selected, then he or she would receive a "decision label" along with the decision. This label would show how the applicant's qualifications measured up to the job requirements; how the applicant compared with other job seekers; and how information about these qualifications was extracted.

Recommendation 3: Creating an informed public

My final recommendation will be brief. To be truly effective, this law requires an informed public. Individual job applicants should be able to understand and act on the information disclosed to them. In Recommendation 1, I spoke about the need to make auditing criteria for fairness and effectiveness publicly available. Empowering members of the public to weigh in on these standards will strengthen the accountability structures and help build public trust in the use of ADS in hiring and beyond. In Recommendation 2, I spoke about nutritional labels as a disclosure method. We should help job seekers, and the public at large, to understand and act upon information about data and ADS.

I recommend that New York City invests resources into informing members of the public about data, algorithms, and automated decision making, using hiring ADS as a concrete and important example. I already started this work, having developed "We are AI", a free public education course on AI and its impacts in society. This course is accompanied by a comic book series, available in English and Spanish.

⁷⁰ The Americans with Disabilities Act and the use of software, algorithms, and AI to assess job applicants and employees, US Equal Employment Opportunity Commission, 2022, <https://www.eeoc.gov/laws/guidance/americans-disabilities-act-and-use-software-algorithms-and-artificialintelligence>



Conclusion

In conclusion, I would like to quote from the recently released position statement by IEEE-USA, titled “Artificial Intelligence: Accelerating Inclusive Innovation by Building Trust”.⁷¹ IEEE is the largest professional organization of engineers in the world; I have the pleasure of serving on their AI/AS (Artificial Intelligence and Autonomous Systems) Policy Committee.

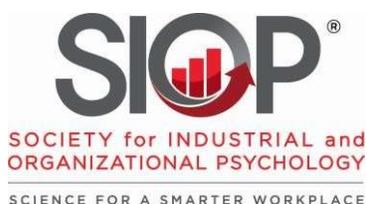
“We now stand at an important juncture that pertains less to what new levels of efficiency AI/AS can enable, and more to whether these technologies can become a force for good in ways that go beyond efficiency. We have a critical opportunity to use AI/AS to help make society more equitable, inclusive, and just; make government operations more transparent and

⁷¹ IEEE-USA, “*Artificial Intelligence: Accelerating Inclusive Innovation by Building Trust*” (2020)
<https://ieeusa.org/wp-content/uploads/2020/10/AITrust0720.pdf>



accountable; and encourage public participation and increase the public's trust in government. When used according to these objectives, AI/AS can help reaffirm our democratic values.

If, instead, we miss the opportunity to use these technologies to further human values and ensure trustworthiness, and uphold the status quo, we risk reinforcing disparities in access to goods and services, discouraging public participation in civic life, and eroding the public's trust in government. Put another way: Responsible development and use of AI/AS to further human values and ensure trustworthiness is the only kind that can lead to a sustainable ecosystem of innovation. It is the only kind that our society will tolerate.”



The Society for Industrial and Organizational Psychology (SIOP) is submitting these comments in response to a request for public comment from the City of New York Department of Consumer and Worker Protection regarding proposed rule amendments to new legislation related to the use of automated employment decision tools (AEDTs).

Industrial and organizational (I-O) psychology addresses workplace issues at the individual and organizational level. I-O psychologists apply research that improves the well-being and performance of people and the organizations that employ them. SIOP is the professional organization representing a community of over 10,000 members, including academics, consultants, and students of I-O psychology, working to promote evidence-based policy and practice in the workplace. Many I-O psychologists specialize in topics related to employee selection and the design and implementation of employment decision tools. SIOP's *Principles for the Validation and Use of Personnel Selection Procedures* (2018) is the authoritative document on how to develop and evaluate employment decision tools. Thus, this expertise is well suited to address the issues at hand in the proposed rule amendments.

SIOP is supportive of efforts to implement a regulatory framework around the use of automated employment decision tools (AEDTs). Although the proposed rule amendments help to clarify the original legislation, we continue to have several concerns about the definitions and implementation of key aspects of the proposed rules.

Specifically, one of the proposed methods of calculating the adverse impact ratio (computing the ratio of selection rates) is commonly accepted but the other method described (computing the ratio of individuals scoring above the median) is not. The second approach is not mentioned in authoritative texts on calculating adverse impact (see Outtz, 2010; Dunleavy, Howard, and Morris, 2015; Dunleavy and Morris, 2016), nor is it a common metric in case law. This proposed new way of considering impact should be carefully scrutinized as it does not consider any cutoff scores that are used to make hiring decisions or the actual selection rates of various groups. Thus, it is possible to have equivalent rates of majority and minority group members scoring above the median (i.e., the "Scoring Rate") and disproportionate hiring rates for one group compared to the other. Therefore, the ratio based on the Scoring Rate does not provide meaningful insights into whether adverse impact does or does not exist for a test used with a particular cutoff score.

The use of historical data from multiple employers using the same AEDT also raises several concerns. The proposed amendments do not clarify when the AEDT used for one employer is sufficiently similar for comparison with data collected from another employer. Moreover, the proposed amendments do not provide any guidance about acceptable time frames within which the historical data should be collected. By their very nature, AEDTs that are based on machine learning and other forms of statistical modeling or artificial intelligence can be updated frequently as new data become available. As a result, the algorithm that underlies the AEDT used by one employer may not be the same as the underlying algorithm used by another employer. This issue will be exacerbated over longer periods of time and is particularly salient in the example of "culture fit" used in the proposed rule amendments. Given that culture is inherently unique to a particular organization, it is unclear that the algorithm used to predict culture fit in one organization

will be the same as the algorithm used to predict culture fit in another organization. In addition, the exact same AEDT used by multiple employers on different applicant pools may result in different levels of adverse impact. For example, an AEDT used for selection into managerial roles may show different levels of adverse impact compared to the same AEDT used for selection into retail sales positions.

The concerns with the use of the “Scoring Rate” described previously are also important when considering the relevance of historical data to a specific employment context. If the historical data were collected from samples that are not representative of the current applicant pools encountered by the employer, the scoring rates calculated using the historical data may not accurately represent the rates that will be observed in future applications. Moreover, the relevance of the historical data to the job(s) the AEDT will be used for is also critically important for establishing scores on the AEDT as a business necessity, as defined by the *Uniform Guidelines on Employee Selection Procedures* (UGESP).

Our final concern is with the requirement to calculate impact ratios for the intersectional categories of sex, ethnicity, and race. Requiring these comparisons without any consideration of the number of people in each of these intersectional categories can provide misleading results that are difficult to interpret. When small samples are used to calculate score differences and/or impact ratios, the results will be unstable and potentially inaccurate. These inaccurate results could lead employers or vendors to conclude that there are no differences in the selection rates for some intersectional categories, only to find out later that substantial differences exist in larger samples that are more representative of the applicant pool (or vice versa).

Expert Contacts

SIOP welcomes the opportunity to submit these comments and provide further expertise and insight. Please reach out to the following SIOP issue experts with additional questions:

Dr. Eric Dunleavy
Director of Employment & Litigation Services
DCI Consulting edunleavy@decisoncult.com

Dr. Christopher Nye
Associate Professor of IO Psychology, Michigan State University
nyechris@msu.edu

Field Code Changed

Dr. Ann Marie Ryan
Professor of IO Psychology, Michigan State University
ryanan@msu.edu

Dr. Nancy Tippins
Principal
The Tippins Group
Nancy@tippinsgroup.com

References

Dunleavy, E. and Morris, S.B. (2016). *Adverse impact analysis: Understanding data, statistics, and risk*. Routledge.

Dunleavy, E., Morris, S., & Howard, E. (2015). Measuring adverse impact in employee selection decisions. In C. Hanvey & K. Sady (Eds.), *Practitioner's guide to legal issues in organizations* (pp. 1–26). Springer International Publishing AG. https://doi.org/10.1007/978-3319-11143-8_1

Ottz, J. L. (Ed.). (2010). *Adverse impact: Implications for organizational staffing and high stakes selection*. Routledge/Taylor & Francis Group.

Society for Industrial and Organizational Psychology (2018). [Principles for the Validation and Use of Personnel Selection Procedures](https://doi.org/10.1017/iop.2018.195), *Industrial and Organizational Psychology: Perspectives on Science and Practice*, 11(Supl 1), 2–97. <https://doi.org/10.1017/iop.2018.195>



**BSA | The Software Alliance Comments on the New York City Department of
Consumer and Worker Protection’s Revised Proposed Rules Implementing Local
Law 144 of 2021 Regarding Automated Employment Decision Tools**

January 23, 2023

BSA | The Software Alliance appreciates the opportunity to submit comments on the revised proposed regulations implementing Local Law 144 of 2021 — New York City’s ordinance on automated employment decision tools (“Ordinance”).¹ BSA is the leading advocate for the global software industry. Our members are enterprise software companies that are on the leading edge of providing businesses — in every sector of the economy — with innovative technology products and services, including those powered by artificial intelligence (AI).² Our members provide trusted tools that help other businesses innovate and grow, including cloud storage services, customer relationship management software, human resources management programs, identity management services, and collaboration software. As leaders in the development of enterprise AI systems, BSA members have unique insights into the technology’s tremendous potential to spur digital transformation and the policies that can best support the responsible use of AI.

The success of AI products and services will be based on public trust and confidence in these technologies. To earn that trust, organizations that develop AI, and those that use AI, must do so responsibly and in a manner that accounts for the unique opportunities and risks the technology poses. Policymakers can enhance public confidence and trust by establishing a legal and regulatory environment that supports responsible innovation, including by creating safeguards that help prevent unlawful discrimination. BSA supports the overarching goal of the proposed regulations to ensure that AI systems have been thoroughly vetted to identify and mitigate risks associated with unintended bias. However, we have recommendations on four aspects of the proposed regulations — the independent auditor requirement, publication of bias audit results, the definition of automated employment decision tool, and the use of historical data — to better achieve this objective.

¹ BSA’s comments on the initial set of proposed rules is available at <https://www.bsa.org/files/policyfilings/10242022nyctools.pdf>.

² BSA’s members include: Adobe, Alteryx, Atlassian, Autodesk, Bentley Systems, Box, Cisco, CNC/Mastercam, CrowdStrike, Databricks, DocuSign, Dropbox, Graphisoft, IBM, Informatica, Juniper

Networks, Kyndryl, MathWorks, Microsoft, Okta, Oracle, Prokon, PTC, Salesforce, SAP, ServiceNow, Shopify Inc., Siemens Industry Software Inc., Splunk, Trend Micro, Trimble Solutions Corporation, TriNet, Twilio, Unity Technologies, Inc., Workday, Zendesk, and Zoom Video Communications, Inc.

A. Independent Auditor

As the proposed regulations recognize, one critical aspect of operationalizing bias audits is identifying a set of individuals who can conduct them. Importantly, the earlier version of the proposed regulations appeared to recognize that internal personnel who are not involved in the development or use of an automated employment decision tool are competent to responsibly conduct a bias audit. Acknowledging the independence of internal stakeholders would incentivize companies to implement multiple layers of independent review as they develop and test their products, which would enhance trust in the use of these systems and create safeguards that function in practice.

The revisions to the proposed rules eliminate this flexibility, clarifying that no employees are considered independent and, therefore, may not serve as auditors. We support the aim of the bias audit but respectfully recommend an alternative approach for two reasons.

First, there is currently no consensus on AI auditing standards. Unlike other areas, such as privacy and cybersecurity, where standards underpin notable certifications, AI lacks this essential element to create trust and accountability in the marketplace. Standards bodies, such as the International Organization for Standardization, are only in the early phases of their AI projects. In addition, recognizing the lack of auditing tools to detect bias and discrimination, Stanford University recently launched an AI audit challenge to help with this effort. A Stanford professor, who is also the President and CEO of Sagewood Global

Strategies, a technology policy and risk advisory firm, and the former Senior Director for Cyber Policy on the White House National Security Council, recently examined this issue and acknowledged the lack of auditable criteria, concluding that “[t]he AI audit ecosystem is immature, at best.”⁷²

Although the proposed rules specify how selection rates and impact ratios should be calculated, this guidance does not replace the need for broader universal, consensus-based standards developed in a multistakeholder process. Nor does it address the lack of available bias detection tools. Without common standards, companies can shop around for auditors based on their preferences for particular methods, criteria, and scope. In addition, the quality of audits will vary significantly and are likely to correlate with price, undermining efforts to establish common objective benchmarks.

Second, there are no professional organizations to govern or train third-party auditors for AI systems. Auditors typically have professional bodies that create baseline criteria and maintain ethical guidelines. SOC audits, for example, are conducted by CPAs and governed by the American Institute of Certified Public Accountants. In addition, educational bodies are in place in other fields — such as privacy — to train professionals. No such body exists for AI auditors.

In short, the AI auditing landscape is nascent, lacking common developed standards and professional oversight bodies. For these reasons, requiring third-party audits is not a feasible or optimal approach. Instead, we recommend that the Department of Consumer and Worker Protection (DCWP) reinstate the earlier definition of “independent auditor,” which would have

allowed internal personnel to conduct the audit so long as they were not involved in the development of the AI system.

⁷² Andrew Grotto, *Sagewood Global Strategies LLC, Audit of AI Systems: Overview, Current Status, and Future Prospects*, 5 (Nov. 2022), available at <https://www.regulations.gov/comment/FTC-20220053-0906>.

B. Bias Audit

As we highlighted in our initial comments, we also recommend omitting the requirement to publish the selection rates and impact ratios for all categories and instead require a summary statement on adverse impact. As an initial matter, the definition of impact ratio includes either selection rates or scoring rates, so there is no need to publish both the impact ratio and the selection rate. In addition, although Section 20-871(a)(2) of the Ordinance requires a “summary of the results” of the bias audit to be published, it does not call for the level of specificity contemplated by the proposed regulations. Publishing the specific information required by the proposed regulations could inadvertently undermine the goals of the Ordinance. For example, it may discourage applicants from groups that are selected less frequently from applying to an organization at all, hampering efforts to attract a diverse workforce. Moreover, requiring the public disclosure of such specific information could disincentivize companies from conducting thorough audits to avoid possible results that may not be optimal. Accordingly, we recommend a more flexible approach to strike “the selection rates and impact ratios for all categories” in Section 5-303 and replace it with “a statement on adverse impact.”

We also note that it is unclear how employers should account for applicants who do not disclose their race, ethnicity, or gender. For example, if an employer has 1,000 applicants, but only 750 of them disclose those fields, the result of the bias audit will be statistically meaningless. Because the selection rate calculation requires an evaluation of the total number of applicants, the disparate impact result from the 750 people who report those fields will be inaccurate and incomplete. A selection rate calculation that excludes the additional 250 people who did not report those demographic fields would also render inaccurate results, as it doesn't reflect the racial and gender diversity that could exist in 25% of the applicant pool. As a result, a requirement that employers publish the results of these calculations on its website could lead to misleading or inaccurate information.

We further recommend aligning the categories for the selection rates and impact ratios with the EEOC's approach for disparate impact testing. The regulations contemplate that race, ethnicity, and sex will be evaluated both separately and intersectionally. However, the US Equal Employment Opportunity Commission (EEOC) does not require intersectional testing. Accordingly, we recommend that you revise the proposed regulations to require that these demographic categories only be tested separately, consistent with the well-established EEOC approach.

C. Definition of Automated Employment Decision Tool

We welcome the proposed regulations' revision to the definition of automated employment decision tool, which strikes “or modify.” This change appropriately focuses the definition on circumstances that overrule human decision making.

The proposed regulations would benefit from clarification on one additional point relating to this definition. If an automated employment decision tool produces a simplified output, but there is functionality in the tool to override or disregard that score (e.g., the user has the ability to change the score or select candidates independent of the score), it is unclear whether the company needs to produce artifacts to show that its automated employment decision tool is not in scope because users are not solely relying on the simplified output. We interpret the rules' silence on this issue as not requiring the production of such artifacts, but we would appreciate further clarifications in the proposed regulations.

D. Use of Historical Data to Conduct Bias Audits

The requirement in Section 5-302 to use historical data to conduct bias audits poses many practical challenges and raises additional concerns regarding the data privacy of employees and job candidates. First, it is unclear what circumstances render historical data unavailable, triggering the option to use test data. Second, the historical data may reside with multiple employers, and vendors may be legally and contractually prevented from collecting or accessing that data. Third, as written, employers would be required to provide sensitive employee and candidate information to third-party auditors. Yet DCWP's rules do not impose limitations on how that data is handled by these third-party firms. Indeed, third parties

may conceivably reuse this sensitive personal information for commercial uses or even sell it to other third parties. For these reasons, we recommend that: (1) employers should not be required to use historical data to conduct bias audits; (2) the use of test data for vendorinitiated audits should be the default rule; and (3) employers should not be required to provide sensitive personal information of their employees and job candidates to third parties.

Finally, we note that the current enforcement date is April 15, 2023. However, without final regulations, organizations cannot undertake all the efforts necessary to comply. We urge you to allow sufficient time for compliance once the regulations have been finalized and, if necessary, postpone the enforcement date.

We thank you for the opportunity to provide comments and look forward to serving as a resource as you finalize the proposed regulations.

**Comments of the Partnership for New York City to the Department of Consumer and
Worker Protection
Proposed Rules Implementing Local Law 144 of 2021
January 23, 2023**

Thank you for the opportunity to comment on the proposed rules implementing Local Law 144 of 2021 (LL144) regulating automated employment decision tools (AEDT). The Partnership for New York City represents private sector employers of more than one million New Yorkers. We work together with government, labor, and the nonprofit sector to maintain the city's position as the preeminent global center of commerce, innovation, and economic opportunity.

The Partnership is grateful to the Department of Consumer and Worker Protection for pausing the enforcement of LL144 and giving impacted parties time to comply. The Partnership also appreciates the Department's consideration of our previous comments and suggestions. Below are recommended changes and questions on the updated proposed rules:

§ 5-300 – Definition of Independent Auditor

- The previous version of the proposed rules allowed for some flexibility as to what type of entity could be considered an independent auditor. We are disappointed that the Department has chosen to remove this flexibility and require the hiring of an external third party. Since auditing of this type of technology is an emerging practice and widely accepted standards do not exist, it is unclear how employers would determine the competence of third-party auditors or the reliability of their results. The only thing that is certain is that this requirement will be a boon to individuals holding themselves out as auditors and will add substantial expense, time, and process for employers.
 - Recommendation: Allow a wider variety of actors to serve as an independent auditor, including units within an employer that are not involved in the development or use of the AEDT.
- It is unclear whether a third-party consultant would be considered an independent auditor (*i.e.*, does not have a “direct financial interest or a material indirect financial interest”) if the consultant 1) has a pre-existing consulting relationship with an employer or vendor or 2) has previously audited the same or a different AEDT for the employer or vendor.
 - Recommendation: Clarify 1) that an independent auditor can be a pre-existing consultant and 2) that the same independent auditor can repeatedly audit the same AEDT and can audit multiple AEDTs for the same employer or vendor.

§ 5-301 – Bias Audit

- The examples provided in §§5-301(a) and (b) are useful, but it would be helpful to explicitly state that the examples do not show all of the possible options for the audit process.
- Employer's commonly use “score-banding” with AEDTs, a practice in which all applicants who receive a score within a range are grouped together. It is unclear how employers would perform the calculation required in (c) if they use score-banding.

- o Recommendation: Clarify the process for the bias audit where score-banding is used.
- The Partnership appreciates the Department’s efforts to address the potential for employers to have incomplete demographic data on job applicants, often due to applicants’ unwillingness to voluntarily disclose such data. In addition to the option to use test data, it would be helpful if employers had the option to use their existing data more accurately.
 - o Recommendation: Allow employers to use “unknown” as a category for demographic information that is not voluntarily disclosed by a job applicant.

§5-302 – Data Requirements

- It is helpful to have the option to use test data for a bias audit, but it is not clear under what circumstances there would be “insufficient historical data [] available.” For example, could an employer use test data where a significant number of applicants chose not to provide demographic information or where the employer’s AEDT does not store such data?
 - o Recommendation: Clarify the circumstances in which sufficient historical data would not be available.
- There is also confusion over what test data an audit would be allowed to use.
 - o Recommendation: Provide guidance and examples of what test data would be acceptable.

§5-303 – Published Results

- Since an audit may not be required to calculate both the selection rate and the impact ratio, it appears that the text in (a)(1) should read “The date of the most recent bias audit of the AEDT and a summary of the results, which shall include the source and explanation of the data used to conduct the bias audit and the selection rates and impact ratios, as required, for all categories.

Thank you for considering our comments.



Rafael Espinal
Executive Director
Freelancers Union Testimony
RE: LL-144 (Bias Audit Local Law)

Good morning Commissioner Mayuga and DCWP,

First and foremost thank you Commissioner Mayuga and the team at DCWP for convening this hearing. My name is Rafael Espinal and I am the Executive Director of the Freelancers Union. My organization represents the interests of over half a million workers nationwide operating in diverse industries. Previously, I had the honor of serving as the Council Member for New York City's 37th district.

I have spent much of my career looking for and creating opportunities to disrupt systemic inequality. I believe that we are currently facing such an opportunity with Local Law 144. We are at a crossroads. The decisions this administration makes in the coming months about how to execute this law will speak volumes about our city's priorities.

If New York City is to prioritize equality of opportunity, Local Law 144 will unapologetically push businesses to take a hard look at the way they evaluate job candidates. Upon taking a look, many of these organizations will have to come to terms with the fact that their hiring tools have played a role in perpetuating historical disadvantage. Employers, regulators, and the public will be armed with data. This data will make it possible to distinguish between employers who are genuinely working to promote workforce diversity and those who are not.

If New York City is to prioritize maintaining the status quo, Local Law 144 will be implemented in a manner that panders to traditionalist business interests. Regulators will publish rules that make bias audits an extremely rare requirement that can easily be circumvented by savvy corporate legal teams. Information regarding which types of hiring tools disadvantage minorities will remain hidden, as it has been since the civil rights era. Employers who use biased AEDTs will face no incentives to change.

Like many voices who have spoken out about this issue to date, I sincerely hope New York City



takes the former approach regarding Local Law 144. As a society, we are becoming less and less tolerant of opacity regarding how important life decisions are made, and for good reason. Transparency is the future when it comes to understanding bias in contexts like hiring, education, healthcare, and housing.

In terms of next steps for Local Law 144, DCWP has already put considerable thought into the proper form and structure of bias audits. Their work to date has several strengths. For example, I am pleased that the agency's proposed rules specify that an auditor must be genuinely independent from the auditee organization. Additionally, DCWP has been wise to shape bias audits around the well-established concept of disparate impact analysis enshrined by Title VII of the Civil Rights Act.

While I am appreciative of the agency's efforts, there is also a major weakness in the proposed rules that needs to be amended before the guidance is finalized. When our City Council passed this legislation, the language was very clear in stating that bias audits would be required for all kinds of automated hiring tools, regardless of their technical nuances. Whether inadvertently or otherwise, DCWP's proposed rules now call for a dramatically limited scope. I believe the scope is so limited that it would make Local Law 144 practically useless.

I would like to leave the administration with this message: Opportunities for disrupting systemic inequality, particularly on issues related to racial equity, are rare. Opportunities to lead the nation in this disruption are even less common. It is my sincere hope that we seize this moment in history and be bold in our fight against bias in hiring.



212-607-3300 www.nyclu.org

Legislative Affairs
One Whitehall Street
New York, NY 10004

Testimony of Daniel Schwarz
On Behalf of the New York Civil Liberties Union
Before the New York City Department of Consumer and Worker Protection Regarding the
New Proposed Rules to Implement Local Law 144 of 2021

January 13, 2023

The New York Civil Liberties Union (“NYCLU”) respectfully submits the following testimony regarding the proposed rules to implement Local Law 144 of 2021. The NYCLU, the New York affiliate of the American Civil Liberties Union, is a not-for-profit, non-partisan organization with eight offices throughout the state and more than 180,000 members and supporters. The NYCLU’s mission is to defend and promote the fundamental principles, rights, and values embodied in the Bill of Rights, the U.S. Constitution, and the Constitution of the State of New York. The NYCLU works to expand the right to privacy, increase the control individuals have over their personal information, and ensure civil liberties are enhanced rather than compromised by technological innovation.

The New York City Council enacted Local Law 144 of 2021 (“LL 144”) which laudably attempts to tackle bias in automated employment decision tools (“AEDT”). AEDT, similar to automated decision systems in other areas, are in urgent need of transparency, oversight, and regulation.⁷³ These technologies all too often replicate and amplify bias, discrimination, and harm towards populations who have been and continue to be disproportionately impacted by bias and discrimination: women, Black, Indigenous, and all people of color, religious and ethnic minorities, LGBTQIA people, people living in poverty, people with disabilities, people who are or have been incarcerated, and other marginalized communities. And the use of AEDT is often accompanied by an acute power imbalance between those deploying these systems and those affected by them, particularly given that AEDT operate without transparency or even the most basic legal protections.

Unfortunately, LL 144 falls far short of providing comprehensive protections for candidates and workers. Worse, the rules proposed by the New York City Department of

⁷³ See: Testimony on Oversight and Regulation of Automated Decision Systems, NEW YORK CIVIL LIBERTIES UNION (2020), <https://www.nyclu.org/en/publications/testimony-oversight-and-regulationautomated-decision-systems>.

Consumer and Worker Protection (“DCWP” or “Department”) would stymie the law’s mandate and intent further by limiting its scope and effect. The NYCLU submitted comments in response to the first proposed rules by the Department on October 24, 2022. Disappointingly, the subsequent update to the proposed rules did not ameliorate the many shortcomings we identified. We therefore resubmit our comments with minor revisions to respond to these changes.

The DCWP must strengthen the proposed rules to ensure broad coverage of AEDT, expand the bias audit requirements, and provide transparency and meaningful notice to affected people in order to ensure that AEDT do not operate to digitally circumvent New York City’s laws against discrimination. Candidates and workers should not need to worry about being screened by a discriminatory algorithm.

Re: § 5-300. Definitions

LL 144 defines AEDT as tools that “substantially assist” in decision making. The proposed rules by DCWP further narrow this definition to “one of a set of criteria where the output is weighted more than any other criterion in the set.” This definition goes beyond the law’s intent and meaning, risking coverage only over certain scenarios and a subset of AEDT. In the most absurd case, an employer could deploy two different AEDT, weighted equally, and neither would be subject to this regulation. More problematically, an employer could employ an AEDT in a substantial way that doesn’t meet this threshold, while still having significant impact on the candidates or workers.⁷⁴ The Department should revise this definition to be consistent with the statute.

The proposed definition of “simplified output” would exclude “output from analytical tools that translate or transcribe existing text, e.g., convert a resume from a PDF or transcribe a video or audio interview.” However, existing transcription tools are known to have racial bias,³ and their outputs could very well be used as inputs to other AEDT systems, resulting in biased results.

Re: § 5-301 Bias Audit

The definition for bias the audit in LL 144, § 20-870, explicitly lists disparate impact calculation as a component but not the sole component (“include but not be limited to”). The examples given in section 5-301 of the proposed rules do not account for an AEDT’s impact on age and disability, or other forms of discrimination.

⁷⁴ Aaron Rieke & Miranda Bogen, *Help Wanted*, UPTURN (2018), <https://upturn.org/work/help-wanted/>. ³ Allison Koenecke et al., *Racial disparities in automated speech recognition*, 117 PNAS 7684 (2020).

At a minimum, in addition to an evaluation of disparate impact of the AEDT, any evaluation that could properly qualify as a bias audit would need to include an assessment of:

- the risks of discriminatory outcomes that an employer should be aware of and control for with the specific AEDT, including risks that may arise in the implementation and use of the AEDT;⁷⁵
- the sources of any training/modeling data, and the steps taken to ensure that the training data and samples are accurate and representative in light of the position's candidate pool;
- the attributes on which the AEDT relies and whether it engages in disparate treatment by relying on any protected attribute or any proxy for a protected attribute;
- what less discriminatory alternative inputs were considered and which were adopted;
- the essential functions for each position for which the AEDT will be used to evaluate candidates, whether the traits or characteristics that the AEDT measures are necessary for the essential functions, and whether the methods used by the AEDT are a scientifically valid means of measuring people's ability to perform essential job functions.

Similar essential components are outlined in the federal EEOC guidance, which recommends including “information about which traits or characteristics the tool is designed to measure, the methods by which those traits or characteristics are to be measured, and the disabilities, if any, that might potentially lower the assessment results or cause screen out.”⁷⁶

The bias audit should clearly state the origin of the data used for the statistics reported. This includes where the data was gathered from, by who, when, and how it was processed. It should also provide justification for why the source of the data for the bias audit model population is believed to be relevant to this specific deployment of the AEDT.

The proposed rules for the ratio calculations also make no mention of appropriate cutoffs when a specific candidate category (per EEO-1 Component 1) has a small or absent membership that could result in unrepresentative statistics.

⁷⁵ See: Algorithmic Discrimination Protections, THE WHITE HOUSE OSTP (2022), <https://www.whitehouse.gov/ostp/ai-bill-of-rights/algorithmic-discrimination-protections-2/>.

⁷⁶ The Americans with Disabilities Act and the Use of Software, Algorithms, and Artificial Intelligence to Assess Job Applicants and Employees, US EEOC (2022), <https://www.eeoc.gov/laws/guidance/americans-disabilities-act-and-use-software-algorithms-andartificial-intelligence>.

Re: § 5-303 Published Results

The disclosure of the bias audit on employers' and employment agencies' websites should not be limited to the selection rates and impact rates results described in §5-301. It should include all the elements mentioned in our comments on the bias audit. The summary should describe the AEDT appropriately and include information on traits the tool is intended to assess, the methods used for this, the source and types of data collected on the candidate or employee, and any other variables and factors that impact the output of the AEDT. It should state whether any disabilities may impact the output of the AEDT.

Additionally, the published results should list the vendor of the AEDT, the specific version(s) of product(s) used, and the independent auditor that conducted the bias audit. The DCWP should provide examples that include such information.

The "distribution date" indicated in the proposed rules for the published results should also describe which particular part of the employment or promotion process the AEDT is used for on this date. It is insufficient to note "an AEDT with the bias audit described above will be deployed on 2023-05-21" unless there are already clear, public indicators that describe which specific employment or promotion decision-making process happened on that date. Any examples should be updated to include a reasonable deployment/distribution description.

Published results should include clear indicators about the parameters of the AEDT as audited and testing conditions, and the regulations should clarify that employers may not use the AEDT in a manner that materially differs from the manner in which the bias audit was conducted. This includes how input data is gathered from candidates or employees compared to how the comparable input data was gathered from the model population used for the bias audit. For example, if testing of an AEDT used a specific cutoff or weighting scheme, the cutoff or weighting scheme used in the actual deployment should match it as closely as possible, and the publication should indicate any divergence and the reason for it. A tool may not show a disparate impact when cut offs or rankings are set at one level and show a disparate impact with other levels. Likewise, if one input variable is hours worked per week, the model population for the bias audit derives those figures from internal payroll data, but candidate data will come from self-reporting, then the publication should indicate that divergence and provide commentary on the reason for the divergence and an assessment of the impact the divergence is likely to have on the relevance of the bias audit.

Lastly, the rules should clarify that the published results must be disclosed in machine readable and ADA compliant formats in order to be accessible to people with various assistive technologies.

Re: § 5-304 Notice to Candidates and Employees

Section 20-871(b)(2) of LL 144 requires the disclosure of the job qualifications and characteristics that the AEDT will use in the assessment. The rules should clarify that candidates or employees should be provided with as much information as possible to meaningfully assess the impact the AEDT has on them and whether they need to request an alternative selection process or accommodation.⁷⁷

The law also requires that the employer “allow a candidate to request an alternative selection process or accommodation.” The regulations should provide employers with parameters of how to provide alternative selection processes or accommodations, including what processes may be used to give equal and timely consideration to candidates that are assessed with accommodations or through alternative processes. By merely stating in the regulations that “Nothing in this subchapter requires an employer or employment agency to provide an alternative selection process,” the regulations suggest that the law provides an empty protection to candidates to be solely allowed to make a request for an alternative without any obligation on the part of the employer to in any way consider or honor that request.

Conclusion

In conclusion, the NYCLU thanks the Department of Consumer and Worker Protection for the opportunity to provide comments to the new proposed rules. The Department’s rulemaking is instrumental in ensuring a productive implementation of Local Law 144 and making clear that discriminatory technology has no place in New York. We strongly urge the Department to amend and strengthen the proposed rules to deliver on the law’s promise to mitigate bias and provide people with the protections and information they need.

⁷⁷ See: Algorithms, Artificial Intelligence, and Disability Discrimination in Hiring, BETA.ADA.GOV (2022), <https://beta.ada.gov/resources/ai-guidance/>.



January 23, 2023

City of New York
Department of Consumer and Worker Protection
42 Broadway
New York, New York 10004

Re: **Sapia.ai Comments on New York City Local Law 144 of 2021 (Subchapter T)**

To Whom It May Concern:

Sapia.ai welcomes the opportunity to submit comments to and questions about New York City Local Law 144 of 2021 and the proposed Rules (“the Law”), regarding the use of automated employment decision tools (“AEDT”). Our comments are specifically focused on the proposed amendments included in Subchapter T.

Sapia.ai is a global software technology provider that offers a text chat-based candidate screening tool for high volume recruitment. We have conducted more than 2 million interviews, with close to 1 billion words shared by candidates in conversations from around the world. Our Chat Interview (Ci) product offers a blind and untimed form of a structured interview conducted via an asynchronous text chat. It combines the reliability of a structured interview, which is well established in organizational psychology literature, with an experience that candidates love—evidenced by the consistent positive feedback we receive.

We strongly believe Local Law 144 is an important step in the right direction toward the responsible use of AEDT, and we are highly encouraged by the City’s efforts to promote transparency with the goal of increasing fairness when companies use AEDT. Mitigating bias is a core founding principle of Sapia.ai. To that end, we have made specific design decisions to mitigate the risk of bias when building and implementing our tools. For example, presently, we do not use video- or audio-based AI due to known risks in those approaches to induce bias.

As a responsible AI provider, we have launched our own initiatives to define and publish a framework for upholding fairness when using AI in recruitment called the Fair AI for Recruitment (FAIR) framework. FAIR identifies four dimensions—bias, validity, inclusivity, and explainability—with clearly defined metrics under each for demonstrating the tool is not causing adverse impact. (A copy of the FAIR white paper can be obtained here:

<https://sapia.ai/sapia-labs/fair-ai-sapia-ai-recruitment-software/>.) Similar to the proposed law, FAIR uses the impact ratio (IR) to test for bias in its models, but additionally uses effect size tests and error rates (where applicable) when determining adverse impact. FAIR goes further by testing for inclusivity using metrics such as candidate satisfaction rates and dropout rates across groups to make sure the experience of the tool does not cause adverse impacts. It is important to note that bias tests such as IRs fail to capture this aspect of AEDT. For example, significantly more female candidates might fail to complete a screening test than male candidates, but a bias test that

considers only the completed candidates will fail to capture this aspect of experience related impacts. We have also adopted the concept of a “model card” that provides transparency to our customers by presenting all model related information such as descriptions of training data used, feature weights, partial dependency graphs (model behaviour data) and adverse impact test outcomes for every model built. We have also published multiple research papers in peer reviewed journals and at conferences, expanding wider transparency of our technology. We offer real-time access to bias testing outcomes and candidate experience stats, including unfiltered candidate feedback to our customers for all live models via insights dashboards. This means customers can check at any point if a live model is inadvertently causing adverse impact, regardless of the model being bias free at the deployment stage.

With this experience, we respectfully submit the following questions and comments for your consideration:

1. **Real time access to impact ratio stats:** It is our experience and well documented in machine learning literature that changes in input data distributions (known as “data drift”) can occur between training and inference data. This means a machine learning model that has been tested for bias at one point in time may still lead to biased outcomes if a data drift has occurred. Capturing data drifts requires real-time monitoring of models.

In order to reduce any adverse impacts due to data drifts, it is important to offer more frequent or real-time bias test outcomes on live models, in addition to a point-in-time bias audit, as required by the law. Hence, we strongly believe that it's part of good governance related to responsible use of AEDT that employers have access to this information in the form of a report or live dashboard to allow for timely interventions if such impacts are discovered. Thus, we recommend that tools be required, or at least encouraged, to provide this real-time information to employers. As proposed, the formal audit conducted by an independent auditor can remain to be a yearly event.

2. **Minimum sample size requirements:** We welcome the inclusion of intersectional categories in the bias audit. However, this requirement further increases the need for clarity around the minimum sample size to be used in the bias audit (i.e. total number of candidates in the case of selection, regardless of the distribution of categories) or the minimum sample size under each category (i.e. total number of candidates or the percentage from the total in a given intersectional category (e.g. Black or African American Females). For instance, the latter can be 30, typically regarded as the minimum sample size for statistical testing. It can also be defined as a percentage from the total sample used in the bias audit (e.g. each intersectional category should contain at least 2% of the total sample, 2% being arbitrary here). It is important that the law provides clarity around the minimum sample size given the bias audit is centered around the calculation of ratios.
 - a. A related secondary point is, how to handle categories that do not meet the recommended minimum sample size. Can those categories be excluded from the impact ratio calculation or still reported but excluded from any interpretations?

- b. This is particularly important when the most selected or highest scoring category is a group with a small sample size. Results may be misleading when the most selected category or highest scoring category has a small sample size.
 - c. Again, it is helpful to have clear guidelines on how to handle these scenarios.
3. **Definitions of test data and historical data:** Based on the definitions given under “5300. Definitions,” for both historical data and test data, it is not clear what data can be considered valid under both categories. It remains vague. For example:
- a. When no historical data is available (e.g. an employer is using AEDT with a custom built model for the first time or first time for a given role type), can historical data from a similar employer from a different geography (i.e. within or outside of the US) be used for the bias audit?
 - b. If a dataset from outside of the US is used, it may not conform to standard EEO-1 categories. In such an instance, can the auditors make a reasonable mapping from provided categories to EEO-1 categories?
 - c. Can data collected from survey platforms such as Amazon Mechanical Turk or other open surveys be used as test data?

We find that the overarching challenge related to the above examples is the lack of clarity around the representativeness of the dataset that can be used to conduct a bias audit. It is also important to highlight that not all employers collect demographic information, and thus, to comply with the law, employers and/or vendors may need to infer sex and race/ethnicity using methods such as BISG that can lead to inaccurate outcomes.

- a. Would the law accept the use of inferred sex and race/ethnicity when this data is not available in an existing historical dataset?
- b. If not above, would the Department provide further time for employers or vendors to collect self-reported sex and race/ethnicity from candidates to formulate a more accurate historical dataset for the bias audit?

Further clarity around the definitions of the datasets that can be used in the bias audit can lead to less confusion and more relevant audit outcomes.

We thank you again, for the opportunity to provide comments on the proposed law. We are happy to share our learnings from building responsible AI solutions in hiring, should the Department require further input.

Sincerely,

_____/s/ Barb Hyman_____
Barb Hyman
Founder & CEO

_____/s/ Buddhi Jayatilleke_____
Buddhi Jayatilleke, PhD
Chief Data Scientist



January 23, 2023

Commissioner Vilda Vera Mayuga
New York City Department of Consumer and Worker Protection
42 Broadway, 8th Floor New
York, New York 10004
<http://rules.cityofnewyork.us>
Rulecomments@dcwp.nyc.gov

Re: Comments regarding the revised rule on the Requirement for the Use of Automated Employment Decisionmaking Tools (AEDT)

Dear Commissioner Mayuga:

For over 75 years, [SHRM](#), the Society for Human Resource Management, has been the foremost global thought leader and convener on all issues concerning work, workers, and the workplace. We appreciate the New York City (NYC) Department of Consumer and Worker Protection (DCWP or Department) efforts to revise and clarify the rules governing the use of AEDTs, employer obligations, and bias audit requirements; however, the revisions still leave several concerns, and SHRM respectfully submits the following additional comments.

SHRM previously submitted comments on the original rule proposal on October 24, 2022, commending the Department's efforts to address ambiguities and provide specificity around the requirements and implications of New York City Local Law 144 of 2021 (LL 144). Clear guidance is still needed regarding the scope and coverage of the law. The revised definition of AEDT remains broad and still leaves employers at a disadvantage in guessing what automated processes and systems are encapsulated in the law. SHRM urges the Department to expressly exclude human-directed assessments, such as scheduling tools, skill-screening technologies, and personality or leadership-style assessments from the tools covered by LL 144.

AEDT Definition Clarity

SHRM appreciates the Department's amendments to the definition of AEDT in an effort to provide additional clarity regarding which tools are covered by LL 144. Unfortunately, the proposed definition remains broad and can be interpreted as capturing tools that sort candidates based on demonstrated ability to perform a certain type of work at a given time and place, but where the tool does not perform an analysis or calculation beyond scheduling availability (e.g., unavailable days; unavailable times of day; etc.) and geographic availability (e.g., the ability to work in one particular location versus another). SHRM **recommends** that the proposed rule specifically and clearly define precisely what constitutes an AEDT. Consistent with our previous

comment letter, the definition still captures many tools and technologies commonly used for operational efficiencies, optimization, and professional development.

Bias Audits

SHRM commends DCWP's effort to provide employers and AEDT vendors with at least two alternative methods for conducting the bias audit. However, the second standard – *scoring rate for a category/scoring rate of the highest scoring category* – is not a proven or tested method. The first standard is consistent with the U.S. Department of Labor Office of Federal Contract Compliance Programs (OFFCP), but the second proposed standard is an untested and unproven means to calculate the impact ratio. SHRM **recommends** that DCWP revise the second impact ratio to a calculation consistent with one or more methodologies employed by statisticians, labor economists, or industrial-organizational psychologists in the employment context. Calculating adverse impacts on specific communities should be supported by metrics found in interpretative guidance or peer-reviewed literature. Moreover, SHRM submits that employers should be afforded the flexibility needed to select the specific bias audit methodologies that are most relevant when assessing the impact ratio of a given assessment.

With regard to the definition of an independent auditor, SHRM **recommends** that the Department maintain the flexibility included in the previously proposed regulation related to the selection of individuals to perform such an audit. Limiting the ability of an employer or vendor to conduct its own bias audit will only serve to increase costs associated with using AEDTs. Many practitioners are not necessarily “involved in using or developing” an AEDT that can be leveraged to conduct the required review, given their expertise in determining the potential impact and validity of a particular assessment.

Notice to Candidates and Employees

SHRM **recommends** that DCWP amend the rule to allow the job postings to satisfy notification requirements to candidates for employment or promotion. The “10 business days prior” requirement is infeasible because it means AEDT-use notifications would have to be made public before the job posting is available. The requirement negatively impacts innovation and prolongs candidate searches and the time it takes to fill job vacancies. AEDTs can assist employers to save time and create efficiencies to identify the top talent and skills to help their organization thrive.

Extension of Enforcement Period

SHRM **recommends** a further extension in the delayed enforcement period that is currently scheduled to end on April 15, 2023. Since the proposed rules are still under consideration, employers will need additional time to implement the requirements needed to comply with LL 144. Conducting the required bias audits will be a significant undertaking, and employers and vendors simply do not have the requisite clarifications to comply by April 15, 2023.

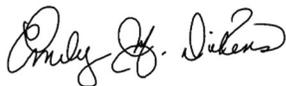
Conclusion

In light of the key benefits that AEDTs offer, SHRM reiterates its prior comments that the requirements and obligations contained in the proposed rules should be viewed through the lens of minimizing limitations on the growth and advancements that could benefit everyone. SHRM champions the creation of better economic opportunities for overlooked and untapped talent pools. As shared in our initial comments, the availability of AEDTs to better recognize the knowledge, skills, and abilities of these workers provides an important tool for organizations seeking to build a more equitable and inclusive workplace.

SHRM promotes the value of untapped talent as more than a matter of social responsibility or goodwill; these groups of workers are proven to show high returns on investment and skills for employers. SHRM recognizes that appropriate safeguards must be balanced against heavyhanded regulatory restrictions that will set key HR functions back and impede the ability to create and identify broader, more inclusive talent pipelines.

SHRM commends the Department's efforts to get this first-of-its-kind law right and calibrated to best serve the NYC business community. SHRM and our over 5,500 NYC-based members are always ready to work with city policymakers to develop public policies that create better workplaces and opportunities for all New Yorkers. We appreciate the opportunity to comment on the revised rule.

Sincerely,

A handwritten signature in cursive script, appearing to read "Emily M. Dickens".

Emily M. Dickens
Chief of Staff & Head of Public Affairs



DISTRICT OFFICE
244 UNION AVENUE
BROOKLYN, NY 11211
TEL: (718) 963-3141
FAX: (347) 228-4347

CITY HALL OFFICE
250 BROADWAY, SUITE 1740
NEW YORK, NY 10007
TEL: (212) 788-7095
E-MAIL: jgutierrez@council.nyc.gov

THE COUNCIL
OF
THE CITY OF NEW YORK

JENNIFER GUTIÉRREZ

COUNCIL MEMBER, 34th DISTRICT, BROOKLYN

CHAIR
TECHNOLOGY

COMMITTEE
ECONOMIC DEVELOPMENT
EDUCATION
ENVIRONMENTAL PROTECTION
HOSPITALS
STATE AND FEDERAL LEGISLATION
WOMEN AND GENDER EQUITY

Dept. of Consumer and Worker Protection
42 Broadway, 8th Floor
New York, NY 10004

Dear Commissioner Mayuga and Agency Sta ,

As you know, Local Law 144 was passed by an overwhelming majority of the City Council in 2021. Over the last few months, your agency has been engaged in the challenging task of drafting rules to shape this precedent-setting initiative. Your efforts to date are greatly appreciated. However, as Chair of the Council's Committee on Technology, I feel compelled to share some concerns with you regarding the latest version of your proposed rules.

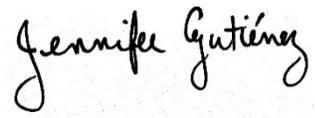
LL144 was passed with the intention of removing bias in the tools that large companies use to screen and hire employees. In my role as Technology Chair, it is my priority to ensure that technology is used as a tool to empower all communities, especially those of color and women, not to serve as another barrier to entry. We must be focusing on how these tools affect the lives of real people, not on how the systems are built. The latest version of draft rules from DCWP is based on the awed assumption that only the most modern and technologically advanced forms of employment tools are worthy of scrutiny. I fear we are gutting much of the original intent of this law.

Some of the rules make it almost virtually impossible to apply to all hiring tools, exempting most of the companies that this law intended to regulate. The notion that this can only apply to companies that solely use this technology – with no human input - is absurd. Humans are involved at every step of the hiring process, and technology is merely an aide, over which there must be supervision. Automated decision-making tools have been in existence for decades, and this law is not just looking at the future, but righting historical wrongs.

These technologies are vast and overreaching, they range from gathering data to making assessments on applicants' emotional intelligence - and if we are not demanding transparency from a wide variety of employment tools, we are all complicit in contributing to systematic inequality and biases in hiring.

I hope that DCWP is approaching this nal stage of its rule-making with the full weight of both the discriminatory history and our futures in mind. If you keep this context in mind, I am confident that the nal rules for Local Law 144 will bring about meaningful change in how New York City prioritizes equity when it comes to employment technology.

Sincerely,

A handwritten signature in black ink that reads "Jennifer Gutiérrez". The signature is written in a cursive style with a large initial "J" and a long, sweeping underline.

Jennifer Gutiérrez, Council Member (34th District)

Gerald T. Hathaway
Partner
gerald.hathaway@faegredrinker.com +1
212 248 3252 direct

Faegre Drinker Biddle & Reath LLP
1177 Avenue of the Americas, 41st Floor
New York, New York 10036
+1 212 248 3140 main
+1 212 248 3141 fax

January 23, 2023

SUBMITTED BY EMAIL: RULECOMMENTS@DCWP.NYC.GOV

Commissioner Vilda Vera Mayuga
New York City Department of Consumer and Worker Protection
42 Broadway, 8th Floor
New York, New York 10004

Re: Proposed Rules for Implementing Local Law 144 of 2021

Dear Commissioner Mayuga and Department of Consumer and Worker Protection:

Faegre Drinker Biddle & Reath LLP is a large law firm with offices throughout the United States, and in London, England and Shanghai, China. Faegre Drinker has an interdisciplinary team, called the AI-X team, that focuses on artificial intelligence (“AI”), algorithmic decision-making & big data. We combine the regulatory and litigation experience of attorneys who have advised clients on these issues for years with the data scientists at Tritura, Faegre Drinker’s AI subsidiary that employs many data scientists, who understand the innerworkings of these technologies and can test clients’ algorithms for unintended discrimination. With an interdisciplinary approach, we help clients leverage the power of AI and algorithms in a safe, legal and ethical manner, understand the risks related to automated decision-making, respond to emerging laws and regulations, and keep up with industry best practices. We represent and advise employers that use or are planning to use AI in the employee selection process. Some of our clients receive thousands of applications for a single job opening, and it would be impossible to process those thousands of applications without the use of AI.

We have the following comments regarding the revised, proposed regulations published on December 23, 2022 relating to Local Law 144 of 2021 (“LL 144”), which regulates the use of automated employment decision tools (“AEDT”).

Limiting the bias audit to qualified applicants. An employer may have very specific qualifications for applicants, and the absence of the qualifications will eliminate the candidacy. For example, a position available in New York City can require an applicant to have a Commercial Driver’s License, or CDL, a requirement of both [federal](#) and [state](#) law for driving certain types of vehicles. Even with the required qualification clearly stated, the employer may well receive many applications from individuals not having CDLs. If an AEDT excludes from consideration those unqualified applicants, what is to be gained by assessing and reporting on the unqualified applicants, and combining the data for the unqualified applicants with the data for the qualified applicants who are not hired? The picture created by the combined data is not helpful, and any reporting on the numbers of those selected/not-selected, by race and sex, could well be distorted. The bias audit should be limited to assessing successful/not successful candidates by looking at only qualified candidates.

Remote Workers Based in New York City: It appears from the Proposed Rules that LL 144 applies only to positions for employment located in New York City, and that the notices required by LL 144 are to be given only to residents of New York City. XYZ Company has no office anywhere in New York City, but does have employees all over the world, and for many positions, XYZ Company allows employees to work remotely from any place in the world, with the work location being entirely of the employee's own choosing. If XYZ Company solicits applications for a new or vacant position that can be performed anywhere in the world, would LL 144 apply to the New York City residents (by choice) who apply for such a position? We suggest that for purposes of LL 144, the location of the employer should be the location from which the work of the New York City remotely working employee is directed, and that the non-city employer not be regarded as being located within the city if it engages remote workers within the city. The required bias audit is an expensive undertaking, and employers engaging a few employees within New York City will not be inclined to make the expenditure, and so may, unless the regulations are clarified, advise potential applicants, either, "New York City residents are not eligible for employment," or, "No remote work can be performed within New York City." Either way, we do not think that employers should be chased out of New York City, nor do we think that job opportunities should be denied to New York City residents, but this law may do just that.

Very truly yours,



Gerald T. Hathaway

GTH/jq



TECHNET
THE VOICE OF THE
INNOVATION ECONOMY

25th
ANNIVERSARY

TechNet Northeast | Telephone 774.230.6685
One Beacon Street, Suite 16300, Boston, MA
02108
www.technet.org | @TechNetNE

January 23, 2023

New York City Department of Consumer and Worker Protection

Re: Revised Proposed Rules Relative to the Use of Automated Employment Decision Tools

Thank you again for the opportunity to submit updated comments regarding DCWP's proposed rules implementing Local Law 144 of 2021, regarding automated employment decision tools (AEDT).

TechNet is the national, bipartisan network of technology CEOs and senior executives that promotes the growth of the innovation economy by advocating a targeted policy agenda at the federal and 50-state level. TechNet's diverse membership includes dynamic American businesses ranging from startups to the most iconic companies on the planet and represents over five million employees and countless customers in the fields of information technology, e-commerce, the sharing and gig economies, advanced energy, cybersecurity, venture capital, and finance. TechNet has offices in Austin, Boston, Chicago, Denver, Harrisburg, Olympia, Sacramento, Silicon Valley, and Washington, D.C.

TechNet's membership includes both employers in NYC and companies that make some of the tools that may be regulated under these rules.

Bias Audit

Our prior testimony sought clarity in the roles and responsibilities of each party and while additional details are provided in the latest draft, we still find inconsistencies in the roles of vendors of AEDT technologies and employers who use them. In the first example of a bias audit in Section 5-301, it appears the AEDT vendor would be required to provide historical data to the independent auditor, while in the following example it puts the onus on the employer seeking to use the AEDT.

In the former case, we remain concerned that vendors will not have adequate historical demographic information to conduct a bias audit as they are not required under the federal Equal Opportunity Employment Act (EOEA) to collect it. Moreover, an employer's use of AEDT for any downstream process may be entirely employer – or even job – specific. Each employer may utilize the tool in a way the vendor has little to no insight into.

In the latter example, the concern remains that under the EOEA, owing to small sample sizes or the fact that employers are not allowed to compel answers to demographic questions, data they provide to auditors may be misleading or statistically insignificant.

The new proposal further expands the definition of independent auditor to disqualify anyone who has an employment relationship with an employer seeking to use an AEDT or has a direct or indirect financial stake in an employer seeking to use an AEDT or the vendor thereof. In order to conduct a meaningful, accurate audit, the auditor must have an in-depth understanding of the system they are auditing. Restrictions on internal experts or those with even a superficial connection to the employer or tool they are auditing will reduce the quality of audits. Further, a thirdparty audit could take up to twice as long as in-house, given the amount of time and effort required to onboard. This may lead to a significant backlog of bias audits, resulting in companies being forced to turn off these important tools in the City simply because the auditing industry lacks the capacity to provide timely audits for all affected vendors and employers.

We recommend the Department revert to its original proposed definition.

“‘Independent auditor’ means 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.”

Definition of AEDT

TechNet remains concerned that the definition of AEDT in LL144 is drafted so broadly as to encompass simple organizational tools like keyword searching, filtering, and labeling. Particularly when a posting attracts a large pool of applicants, these basic tools can be crucial in scanning cover letters and resumes for requisite core competencies listed in the job description.

To help address that concern, we ask that the Department consider further tightening the proposed definition of AEDT as follows:

Automated Employment Decision Tool. “Automated employment decision tool” or “AEDT” means “Automated employment decision tool” as defined by § 20-870 of the Code where 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 substantially 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. “Automated employment decision tool” or “AEDT” does not include (1) tools to help improve the overall quality of candidates for employment or the efficiency of the recruiting process, prior to a natural person applying for a specific

Definition of Screening

Ultimately, the proposed definition of AEDT in the proposed rule hinges more on the human use of the tool than the essence of the tool itself. In order to align the rest of the rules with the intent signaled in that definition, we propose an adjustment to the definition of “screen” to focus the rules on the processes that may have concrete effects on decision making:

Screen. “Screen” means to make a determination that a candidate for employment should not be hired or that an employee should not be granted a promotion. ~~about whether someone should be selected or advanced in the hiring or promotion process.~~

We further recommend clarifications throughout the rules to specify that the bias audit requirement is triggered when an employer uses an AEDT for screening purposes.

Thank you for your consideration. TechNet’s members are eager to be of service as the Department considers these complex issues. Please do not hesitate to contact me if I can provide any additional information.

Sincerely,



Christopher Gilrein
Executive Director, Northeast
TechNet
cgilrein@technet.org



Comments on New York City LL 144 Proposed Rules
NYC Department of Consumers and Worker Protection
January 23, 2023

I write on behalf of [Indeed](#), the world's leading job site whose mission is to help people get jobs. In New York City, over 800,000 people use Indeed to search for work and over 15,000 businesses use Indeed to find talent.

In order to [increase equity in hiring](#) and [remove barriers](#) for job seekers, Indeed works toward building algorithms that promote fairness. Indeed's AI Ethics Committee - a diverse team of data scientists and subject matter experts - consult on the development of algorithms, as well as the monitoring and auditing for unintended bias.

In partnership with the government, Indeed leverages its tools and solutions to address the most pressing hiring challenges in our community. In a recent effort, [Indeed joined the White House call-to-action to address the current shortage of education workers](#) by providing our tools to help public schools across the country hire teachers, administrators, counselors, and other school staff.

Our goal is to help New Yorkers find jobs and the right regulation of AEDTs can minimize harm, and maximize fast, successful hires. We appreciate the City's leadership on this important issue.

Below is Indeed's feedback on the Proposed Rules to LL 144. Thank you in advance for your consideration and please don't hesitate to reach out if I can provide any additional information.

Sincerely,

Alison Klein
Government Relations and Public Policy, State & Local
Indeed aklein@indeed.com

Definition of Independent Auditor

Definition:

"Independent auditor" means 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. An auditor is not an independent auditor of an AEDT if the auditor:

- 1. is or was involved in using, developing, or distributing the AEDT;*
- 2. 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
- 3. 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 continue to use the AEDT or in a vendor that developed or distributed the AEDT.*

Feedback:

This definition of "independent auditor" raises several concerns.

1. Quality

To conduct a meaningful, accurate bias audit, an auditor must have an in-depth understanding of the system it is auditing. Qualified auditors are those with strong institutional knowledge of the system and compared to an internal business expert, a third party lacks the knowledge that is required.

2. Partiality

Objectivity in bias audits are determined by the design of the audit itself, not by who is conducting them. At present, there are no standards for what is considered a sufficient bias audit and auditors, both first-party and third-party, are able to set their own parameters. Without standardization in the auditing process, third party auditors are no more "objective and impartial" than first-party auditors.

3. Feasibility

A carefully considered and rigorous internal audit can take months to complete, with open channels of communication between auditors and product teams. External audits conducted with the same rigor and care will take much longer. Given the current capacity of the bias auditing industry (estimated about 200 qualified professionals in the U.S. at present), third party auditors would not be able to meet

demand for NYC employers that need to audit their tools. This would create a tremendous audit backlog, potentially requiring businesses to turn off their AEDTs in the city in order to not violate the law. As a result, NYC workers will not be able to access or benefit from the essential services that allow them to find jobs.

Recommendation:

To help ensure that all audits are conducted with the same level of quality and objectivity, regardless of who it is conducted by, we recommend that the City create a clear set of standards for what is considered a sufficient bias audit.

To make auditing feasible at scale, we propose expanding the pool of “independent auditors” by reverting to the original definition proposed by the Department:

“Independent auditor” means 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.”

By creating a set of bias audit standards and expanding the pool of auditors, quality, objective bias audits will be able to be conducted at scale.

Definition of Scoring Rate and Impact Ratio

Definition:

“Scoring Rate” means the rate at which individuals in a category receive a score above the sample’s median score, where the score has been calculated by an AEDT.

$$\text{Impact Ratio} = \frac{\text{Scoring rate for a category}}{\text{scoring rate of the highest scoring category}}$$

Feedback:

The new definition of "scoring rate" uses the median as the threshold for use of a continuous score. The median scores of two groups are compared to create a measure of demographic parity. While this is a straightforward metric of fairness, it is not recommended. Many employment decision tools produce continuous scores at

some level; however, practical applications of scores almost always lead to a threshold for selection. For example, one employer may have a university GPA cutoff of 3.0, while another may use a threshold of 3.6, and for both the median GPA may be 2.8. It is statistically more likely for a continuous score to be fairer as the threshold lowers and more candidates pass (see Figure 1).

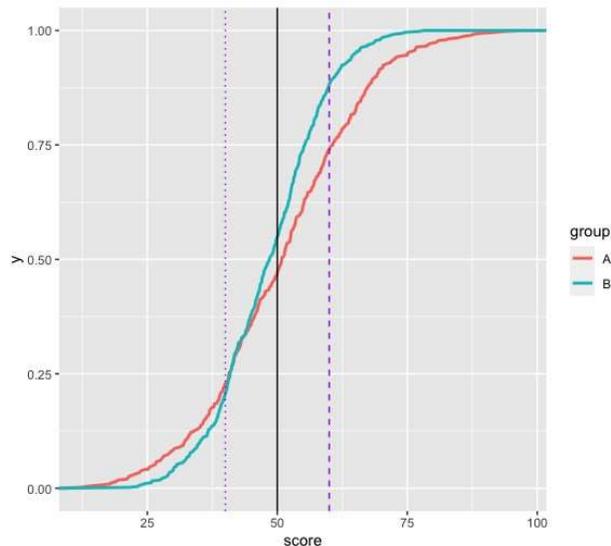


Figure 1. The median score (black line) is considered fair (solid black line; impact ratio of .85). However, the score is unfair to group A if employers used a stricter threshold of 60 (dashed purple line; impact ratio of .45), and fair but skewed in favor of group A if employers used a more permissive threshold of 40 (dotted purple line, impact ratio of 1.028).

Using the median as the *de facto* threshold for all applications of a score would penalize employers using a looser criterion for success (e.g., a GPA of 2.5 when the median is 2.8). More importantly, using the median as a threshold would more likely fail to find bias when employers use a stricter criterion for success (e.g., a GPA of 3.6 when the median is 2.8). Figure 2 illustrates a comparison of two groups with statistically equal mean scores of 50, but with different variance, resulting in a fair result when using a median cutoff of 50, but an unfair when using a higher threshold of 70.

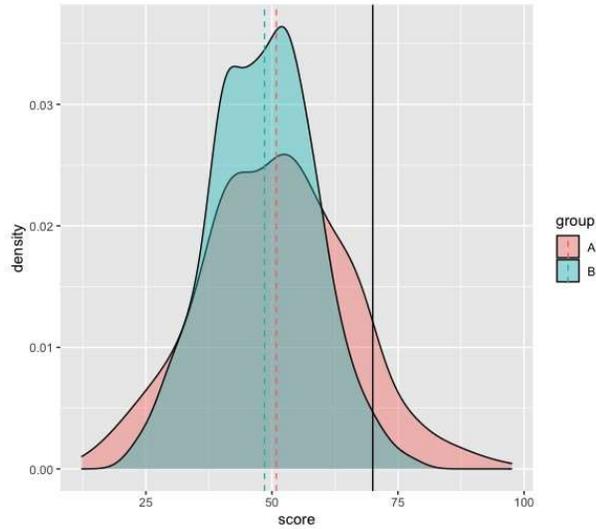


Figure 2. Illustration of impact using the median as a de facto threshold for comparing continuous scores. Groups A and B have statistically equal medians (50.8 and 48.5 respectively), but different variance. This results in a determination of fairness when comparing medians (impact ratio of .955), when the use of a strict selection threshold at a score of 70 would be considered unfair (impact ratio of .214).

Recommendation:

It is recommended that AEDT's be audited in a way that accounts for the thresholds recommended for use by the product guidelines and/or the thresholds actually used to select candidates.

VIA EMAIL

January 23, 2023

Hon. Vilda Vera Mayuga
Commissioner
New York City Department of Consumer and Worker Protection
42 Broadway #5
New York, NY 10004

RE: Automated employment decision tools

Dear Commissioner Mayuga:

Thank you for the opportunity to comment on proposed rules related to Local Law 144 of 2021. The local law established that it was “unlawful for an employer or an employment agency to use an automated employment decision tool to screen a candidate or employee for an employment decision.”

The Retail Council of New York State was founded in 1931 and now represents 5,000 stores statewide. Our member stores range in size from the nation’s largest and best-known brands to the smallest Main Street entrepreneurs that fuel local economies. According to [New York State Comptroller Thomas DiNapoli](#), there were 311,200 retail jobs in New York City in November 2022, making the industry the city’s second largest private-sector employer behind the “office sector.”

We respectfully submit the following comments for your consideration:

- 1.) First, we request that part (ii) of the definition of “Automated Employment Decision Tool” be omitted or clarified to further define what the Department intends with their use of the word “weighted.” With the current definition, it is difficult for employers to determine if the Department is requesting a mathematical calculation or an objective set of criteria for which candidates are reviewed.
- 2.) Similarly, we request that examples be included in the definition for “machine learning, statistical modelling, data analytics, or artificial intelligence” to increase clarity for employers.
- 3.) We greatly appreciate the additional clarity on what constitutes an independent auditor. Unfortunately, some ambiguity remains as to whether a third-party consultant, not otherwise involved in the development / distribution of the AEDT but who has a pre-existing consulting relationship with an employer, has a "material indirect financial interest" and is therefore excluded.
- 4.) There are several points of concern as it relates to the bias audits:

- a. The proposed rule explicitly states employers must conduct a bias audit and publish the standalone results for race/ethnicity and sex, and for the intersection of those categories.
The former is inconsistent with Uniform Guidelines on Employee Selection Procedures (UGESP) requirements. Requiring an intersectional analysis could result in small populations being compared. This is neither useful for public consumption nor indicative of discrimination.
- b. Additionally, it is our belief that the annual requirement to conduct a bias audit will create a “cottage industry” of independent auditors for whom universal standards are still being developed.
- c. It is our understanding that there are exceptions under employment law for the need to consider certain populations that fall below a critical threshold. However, there is no exception when the data available for a subpopulation within the pool of employees or candidates is too small to measure the ratios for that subcategory, as required in §5301(b).
- d. We respectfully request clarity related to the “test data” an employer might use if they do not have sufficient historical data. Examples for this situation would be helpful.
- e. The bias testing examples do not provide for an "unknown" category for race or gender, which could be a significant number. There will likely be a considerable number of employers who do not collect this information from applicants and may not have useful data to perform the audit.
- f. The requirements for automated employment decision tools that score candidates do not translate to common employer practice of score banding. Employers would need additional clarity on what is required in these situations.

To be clear, we agree that there is no place for discrimination of any kind during the hiring process. Our concerns are related to unclear definitions and the need for clarification.

Thank you, again, for the opportunity to provide comment on the proposed rule. We will remain constructive throughout the regulatory process and are available if you have any questions.

Sincerely,

A handwritten signature in black ink that reads "Kelsey Prado". The signature is written in a cursive, flowing style.

Director of State and Local Government Relations
Retail Council of New York State



SAP America, Inc.
3999 West Chester Pike
Newton Square, PA 19073
T 610.661.1000
F 610.661.1000
info@sap.com

January 23, 2023

SENT VIA ELECTRONIC MAIL

City of New York
Department of Consumer and Worker Protection
42 Broadway
New York, New York 10004

RE: SAP Comments on the Proposed Rule for Automated Employment Decision Tools

To Whom it May Concern:

SAP appreciates the opportunity to submit comments in response to the proposed rule on Automated Employment Decision Tools (AEDT). Overall, SAP is encouraged by the agency's continued engagement with the public and private sectors in the development of this rule.

SAP is a globally recognized leader in the field of information technology. We are the market leader in enterprise application software, helping organizations of all sizes and in all sectors run at their best. Our customers generate 87% of total global commerce (\$46 trillion). Additionally, 99 out of the 100 largest companies in the world are SAP customers. We operate in over 150 countries and have over 110,000 team members worldwide. We have over 110,000 employees, including over 500 employees in New York City alone, and operate in over 150 countries. SAP is also a proud resident of New York City with our Hudson Yards office.

We share the view that responsible developers and deployers of Artificial Intelligence (AI)-enabled tools, including AEDT, need clear guardrails to make the protections in LL 144 workable and avoid unintended disruptions to New York City's economic prosperity. Recognizing the significant impact of AI on people, our customers, and society, SAP designed guiding principles to steer the deployment of our AI software to help the world run

better and improve people's lives. *SAP's Guiding Principles for Artificial Intelligence*⁷⁸ requires our technical teams to gain a deep understanding of the business problems they are trying to solve and the data quality this demands. We seek to increase the diversity and interdisciplinarity of our teams, and we are investigating new technical methods for mitigating biases.

We believe SAP is uniquely suited to provide the Department of Consumer and Worker Protection (DCWP) with insights into the opportunities and challenges in the proposed rule on AEDT. This comment letter serves as SAP's response to published amendments to the proposed rule. Overall, we offer the following considerations:

1. Strengthen on-ramps for employers to comply with LL 144.

While a growing number of organizations are integrating AI-enabled tools into their processes, the utilization of AI is still an emerging technology. Many employers do not fully understand what is required of them under LL 144 or the steps they need to take to comply. Employees and job candidates do not know what to expect from employers or what recourse they may have under LL 144. We recommend providing guidance, a question and answer document, and practical tips on how employers can comply with LL 144. Additionally, the agency should clarify in broad communication to prospective job candidates in New York City how this rule impacts them and what changes they may observe from New York City employers.

2. Harmonize the AEDT definition with domestic and international statutes and policies.

AI and automated decision systems have become the most prominent terms used for legal and regulatory purposes. By giving clear meaning to the term 'automated employment decision tools', the legal definition can resolve ambiguity and communicate to various audiences that interact and relate to LL 144 differently (e.g., lawyers, judges, civil servants, corporations, and the public). The proposed rule would benefit from aligning with U.S. and international statutes and policies. The technical assistance document from the U.S. Equal Employment Opportunity Commission on the Americans with Disabilities Act (ADA) and AI⁷⁹, the European Union General Data Protection Regulation (GDPR)⁸⁰, and the Government of Canada Directive on Automated Decision-Making⁸¹ provide harmonized definitions around automated decision systems to include employment tools. We recommend the agency amend the AEDT definition to make it clear that the law is intended to apply to tools that operate with limited human intervention for the purpose of making employment decisions.

3. Avoid boxing in the City of New York to bias mitigation methodologies that may become obsolete in the future.

⁷⁸ [SAP's Guiding Principles for Artificial Intelligence](#)

⁷⁹ [The Americans with Disabilities Act and the Use of Software, Algorithms, and Artificial Intelligence to Assess Job Applicants and Employees – U.S. Equal Employment Opportunity Commission](#)

⁸⁰ [Article. 22 GDPR – Automated individual decision-making including profiling – European Union](#)

⁸¹ [Government of Canada - Directive on Automated Decision Making](#)

SAP believes that harnessing the transformative potential of AI requires a commitment to actively develop and use it responsibly. A critical aspect of the responsible development of AI is the focus on identifying and mitigating bias. In recent years, SAP AI Research⁸² has shared findings and made contributions in the development of tools that support bias mitigation. The advancements being made in the AI ethics space are rapidly evolving with new bias mitigation methodologies and techniques on the horizon. We recommend that the agency consider adopting less prescriptive bias auditing techniques to ensure the rule can adapt to new (and potentially superior) methods for the mitigation of bias in AEDT. Additionally, the agency should review the National Institute for Standards and Technology (NIST) *AI Risk Management Framework (AI RMF)*⁸³ expected to be published on January 26, 2023 to ensure the rule aligns with the consensus-driven process to address AI risks throughout the lifecycle of AI products, services, and systems, to include AEDT.

4. Requirement to use historical data presents a significant challenge.

§5-302 establishes requirements in which historical data must be used to conduct bias audits, and if “insufficient historical data” is not available, test data may be used as an alternative. It is unclear what threshold must be met to deem historical data as “insufficient”, and thus triggering the option to use test data. Additionally, the unintended consequence of the data requirement is that it introduces bias in favor of employers or employment agencies that are able to conduct audits using historical data. The agency’s proposed rule creates a competitive advantage rather than a neutral position. Therefore, we recommend removing the requirement to use historical data, and only allow the use of test data to meet the bias auditing requirements described in §5301.

5. Permit independent auditing flexibilities.

LL 144 requires that AEDTs be subjected to an impartial evaluation by an independent auditor. The proposed amendments to the rule would remove the flexibility of employers or employment agencies to utilize internal AI experts to meet the definition of an independent auditor. Additionally, the amendment should take into consideration that internal auditing, documentation, and governance of AI products, services, and systems are a core element to building trustworthy AI under the NIST AI RMF. Lastly, the AI third-party auditing landscape is relatively new, lacks standards, ethics rules, and an oversight body. Therefore, we propose incorporating a model of accountability used in federal financial services regulation (e.g., Generally Accepted Accounting Principles or GAAP⁸⁴) which offers employers the choice to establish an

⁸² [SAP AI Research](#)

⁸³ [NIST AI Risk Management Framework Home Page](#)

⁸⁴ [Financial Accounting Standards Board – GAAP](#)

internal audit function, consistent with auditing best practices, or contract with a third-party firm to conduct the testing and review of the AEDTs being deployed.

In closing, SAP encourages DCWP to continue engagement and gaining feedback as the agency develops this regulation by conducting another comment period. We hope that SAP's recommendations support the advancement of positive change that leads to the responsible development and use of AI. If well-crafted this rule will be a foundation to provide organizations with an opportunity to enhance New Yorkers confidence and trust in AI. Thank you again for the opportunity to engage on the development of this regulation. We look forward to collaborating more with the City of New York. For any questions, please contact Mr. Brian Logan in SAP Government Affairs at Brian.Logan@sap.com.

Respectfully Submitted,

Brian Logan
Senior Manager, U.S. Government Affairs

January 23, 2023

New York City Department of Consumer and Worker
Protection 42 Broadway, 8th Floor New York, New
York 10004 <http://rules.cityofnewyork.us>
Rulecomments@dcwp.nyc.gov

Re: Comments Regarding Proposed Rules for Implementing Local Law 144 of 2021

To the Department of Consumer and Worker Protection:

Gibson, Dunn & Crutcher, LLP (“Gibson Dunn”) is a go-to global law firm for novel and high-stakes matters, providing full-service representation through the most challenging labor and employment issues. Gibson Dunn’s prominent Labor & Employment and Artificial Intelligence & Automated Systems Practice Groups—among others—routinely work to provide cross-disciplinary expertise and counsel to employers regarding federal, state, and local requirements, including as they relate to the use of automated decision-making tools.

Gibson Dunn submits on behalf of a client these comments to the proposed rules issued by the Department of Consumer and Worker Protection (the “Department”) on December 23, 2022 that seek to implement New York City Law 144 of 2021 (“Local Law 144”) regarding the use of automated employment decision tools (“AEDT”) by employers in hiring and promotion decision making (the “Proposed Rules”). Our client urges the Department to further clarify the scope and applicability of the Department’s most recent definition of AEDT under Local Law 144, as refined by the Proposed Rules.

I. The Revised Definition Of AEDT Should Be Narrowly Tailored So Algorithmic Tools That Merely Assist Employers In Sorting, Filtering, Or Reviewing Candidates Are Not Unintentionally Included.

Our client encourages the Department to revise the definition of AEDT to more clearly capture the intended scope: to exclude automated organizational tools that are used to help employers simplify their review of candidates for employment through the use of recommendations or suggestions, or the ability to sort, filter, organize or review such candidates, but whose outputs are not, with respect to making employment decisions, used

as or intended to be (i) the sole factors considered, (ii) weighted more heavily than any other criteria, or (iii) overruling conclusions from other factors.

We believe the Department’s inclusion of “test data” in the Proposed Rules further supports this narrowed definition and interpretation, as tools excluded by this definition of AEDT likely cannot be reliably audited using test data in the manner Local Law 144 requires.

A. Revisions To The Definition Of AEDT In The Proposed Rules Help To Define The Scope Of Included Tools, But Additional Clarification Should Be Provided.

AEDT are defined within Section 20-870 of Local Law 144 as “any computational process, derived from machine learning, statistical modeling, data analytics, or artificial intelligence, that issues simplified output, including a score, classification, or recommendation, that is used to *substantially assist* or *replace* discretionary decision making for making employment decisions that impact natural persons” (emphases added).

In Section 5-300 of the Proposed Rules, the phrase “to substantially assist or replace discretionary decision making” is further refined as follows (changes **bolded** as compared to the September 23, 2022 version):

(i) to rely solely on a simplified output (score, tag, classification, ranking, etc.), with no other factors considered; (ii) 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 (iii) to use a simplified output to overrule ~~or~~ **modify** conclusions derived from other factors including human decision-making.

Our client appreciates the Department’s efforts to modify the definition of AEDT in the Proposed Rules to make it more focused.⁸⁵ Removal of “or modify” in the third prong of the revised definition is a helpful and necessary clarification that *excludes* simplified outputs that would merely “modify” conclusions derived from other factors, including human decision-making. For further confirmation, the Department should revise the second prong (prong (ii)) of the aforementioned definition to specify that a covered AEDT would only “substantially assist or replace discretionary decision making” if it uses “a simplified output as one of a set of criteria where the simplified output is weighted **substantially** more than any other criterion in the set.” By including “substantially,” the definition appropriately tracks the terminology of the original definition, which requires that the tool either wholly replace, or “**substantially** assist” the employer’s decision-making. Importantly, this would more clearly narrow the scope to tools in which the

⁸⁵ See Statement of Basis and Purpose of Proposed Rule (“Various issues raised in the comments have resulted in changes to the proposed rules. These changes include . . . [m]odifying the definition of AEDT to ensure it is focused.”).

simplified output is meaningfully being used in an employer’s *decision-making*—not tools which merely sort or organize a simplified output for the employer to then engage in its own decision-making process.

Local Law 144’s definition recognizes as much, and identifies certain tools that may be used to assist review, but do not “substantially assist or replace discretionary decision-making processes,” such as “a junk email filter,” “calculator,” “spreadsheet,” “or other compilation of data.” In view of this, the Department should additionally confirm in the Proposed Rules the list of examples that covered AEDT do *not* include, such as: (1) tools performing automated searching of resumes or other relevant documents to identify qualifications about candidates for employment, including relevant skills or experience; or (2) any automated organizational tools that are used to assist in the searching, identifying, sorting, filtering, labeling, analyzing, or organizing of candidates for employment for further human review.

The Department’s confirmation of the types of tools included within the scope of Local Law 144 is necessary to ensure that automated tools that merely *assist* employers in hiring and promotion review process, but do not overrule the *end decision-making*, are excluded from the law, as these are not the tools the Department sought to govern.

B. Organizational Tools Assist Employers In Identifying Qualified Candidates, But Are Not The Sole Factor In The Decision, Do Not Substantially Outweigh Any Other Criterion, And Do Not Overrule Human Decision-Making.

As additional background on why we believe the Department reasonably excluded these automated organizational tools, and why the additional suggested revisions provide important context: many automated tools seek to provide recommendations and suggestions for employers about candidates for employment, or identify and organize candidates for further review. These tools may help play a vital role in simplifying an employer’s review of such candidates, but they are in no way intended to “overrule” an employer’s conclusion, be the sole or most heavily weighted criteria or factor in the hiring or promotion process, or even to be involved in the *decision* at all.

Rather, these tools are used for the mere searching, identifying, sorting, filtering, analyzing, organizing, or suggesting of candidates as a way to simplify an employer’s review, with humans not only overseeing the use of such processing, but also making the actual decision-making downstream of these processes.

For example, certain types of tools may engage in algorithmic matching by identifying applicants who are potential high-match candidates for a given position. This matching is based on a limited universe of factors concerning the job posting (such as a job’s title, location, and description) and the candidate (such as the candidate’s general employment background, location, skills, and/or certifications). Similarly, a “spreadsheet”

(which is excluded from the primary definition of an AEDT) may be used by the employer to automatically filter or search based on keywords or categories. These types of automated organizational tools are just that, organizational, and intended to streamline and improve the efficiency of an employer’s recruiting process.⁸⁶

Such tools do not remove candidates for employment from the employer’s applicant pool, do not replace, overrule, or otherwise serve as the sole or primary decision-making criterion, and do not overrule conclusions derived from other factors. In practice, how automated organizational tools are used, and the *subsequent* employment *decision*-making, are incredibly employer-specific (for example, the definitions clearly recognize the issue with a spreadsheet that can be filtered, but does not explicitly account for various other tools that present similar complexities). Accordingly, we believe the Department intended to exclude such automated organizational tools from the definition of an AEDT because the ultimate employment decision is largely driven by the employer’s human decision-making.

C. The Definition Of “Test Data” Further Supports The Clarified Scope Of AEDT.

Automated organizational tools should also be outside the scope of this law because such tools lack the very type of inputs and algorithms Local Law 144 seeks to govern, and therefore cannot be reliably audited for bias in the manner proposed, even as most recently edited.

Section 5-302 of the Department’s revised Proposed Rules specifies that bias audits may use test data if “insufficient historical data is available to conduct a statistically significant bias audit.” Section 5-300 defines “historical data” as “data collected during an employer or employment agency’s use of an AEDT to assess candidates for employment or employees for promotion,” and “test data” as “data used to conduct a bias audit that is not historical data.” The permissibility of test data to conduct a bias audit as a theoretical concept is welcome as it will ensure that the innovation of new tools necessarily lacking historical data is not inhibited.

Nevertheless, it is difficult to identify in practice what “test data” may be used for automated organizational tools, and the limitations would make conducting such a bias audit confusing or inapplicable, if not impossible. For example, where a “tool” is neither intended to collect sex, race/ethnicity, or intersectional categories in the first instance, nor uses any of that information as part of its matching or organizational algorithms, it is not clear what (if any) test data could (let alone *should*) be used to calculate potential disparate impact. Indeed, even if a tool endeavored to request such information, applicants have the

⁸⁶ On average, each corporate opening attracts 250 resumes. See Glassdoor, *50 HR & Recruiting Stats That Make You Think* (2015), <https://www.glassdoor.com/employers/blog/50-hr-recruiting-stats-make-think/>.

ability to decline to provide it, which may skew the findings of any bias audit given the inconsistency of such information being provided.

Moreover, where an employer’s use of the tool for any downstream review process is largely employer- (or even job-) specific, a tool developer has little or no insight into such process. This user-specificity renders it infeasible (if not impossible) to test every permutation of potential employer use and the outcomes of downstream and unknowable employment decisions—similar to the concerns presented by a spreadsheet tool with filtering. Each of the practical considerations that make the definition of “test data” difficult to apply in these instances support these tools being appropriately classified as outside the scope of an AEDT.

II. Conclusion

The Proposed Rules should further clarify that Local Law 144 intends to exclude automated organizational tools that solely provide recommendations or suggestions to employers, or the ability to search, identify, sort, filter, label, analyze, or organize candidates for employment but are not (i) the sole factor considered, (ii) weighted substantially more than any other criterion, or (iii) overruling conclusions from other factors. Further, the types of data that the law seeks information about is unavailable or inapplicable for these tools. As such, we urge the Department to more clearly identify the scope of tools reasonably covered—and, more importantly, not covered—by the definition of an AEDT.

* * * *

Please do not hesitate to contact us if you believe additional information would be of assistance. Thank you for your time and attention to our submission.



Commissioner Vilda Vera Mayuga
NYC Department of Consumer and Worker Protection
Consumer Services Division
42 Broadway, 9th Floor
New York, NY 10004

Re: New York Staffing Association Comments on Proposed Rules

Dear Commissioner Mayuga,

On behalf of the New York Staffing Association (“NYSA”), we respectfully submit the below comments to the Department of Consumer and Worker Protection regarding proposed rules for New York City’s Local Law 144, in relation to automated employment decision tools (“AEDTs”).

Introduction

NYSA represents temporary and contract staffing firms in New York. The staffing industry plays a critical role in finding jobs for New Yorkers, employing over 700,000 temporary and contract workers annually, thus contributing significantly to employment and the State’s economic vitality and growth.

Employees work for staffing firms to supplement their incomes and/or while looking for more traditional work arrangements. Approximately 70% of employees working for staffing firms ultimately find permanent jobs either with staffing firm clients or through the skills they received while temping, making staffing firms excellent job generators.

Comments on Local Law 144

We thank the Department for its thoughtful consideration of the rules governing the use of AEDTs. These tools have become vital to employers who, with the advent of the internet, are now inundated with applications in response to job advertisements- a result of the comparative ease of submission. Without the use of automated tools of some form, a large number of potential applicants would never have their submissions reviewed, as employers simply do not have the manpower to review every inquiry received. The use of AEDTs allows for the evaluation of a much greater number of applicants and in turn promotes fairness. Additionally, these tools are commonly used by employers to avoid the unconscious bias and discrimination that may result from human decision makers.

Data Requirements

After reviewing the most recent draft rules, we would like to address a few concerns with respect to the data requirements presented to employers. We understand that an auditor must utilize some form of

data in order to perform a test on a given system's biases. But while larger employers may collect the information that would be required in such a dataset— which would include data on intersectional categories of race, gender, and ethnicity—most small employers and staffing firms do not collect any information of this sort. Staffing firms in particular do not collect this data because they do not see their employees on a regular basis and would have difficulty complying. The EEOC has long recognized this fact and has exempted staffing firms from having to accumulate this data and provide it for their employees on the EEO-1 form. As a result, staffing firms (and small employers who are exempt from EEO-1 reporting because of their size) do not have any historical data available to provide to an auditor.

We worry, however, that because many employers do not obtain any of this information from applicants, they might be excluded from a multi-employer audit scenario as described in part “(c)” of the Data Requirements section. This section states that a bias audit may use the “historical data” of any employer or employment agency that makes use of the subject AEDT. It goes on to say that employers or employment agencies may rely on a bias audit of an AEDT that uses the historical data of other employers or agencies *only* if these parties have provided their own historical data for use in the audit.

While we understand the spirit of fairness in this point, many employers simply do not have backlogs of selection data centered upon gender, race, and ethnicity to provide to an auditor. In the absence of this capability, we ask the Department to allow *all* employers who make use of a tool to receive access to the bias audit. By only allowing those who provide historical data to rely on the multi-employer bias audit, the Department effectually excludes small employers and staffing firms.

We also encourage the Department to allow all such audits to make use of test data, since “historical data” is largely nonexistent except through large employers. Indeed, we contend that test data is more than capable of providing an accurate measure of the implicit biases of a given tool.

Responsibility of the Audit

In our previously submitted comments, we noted the ambiguity around who is responsible for performing an audit of the AEDT technology. We believe that the intent of the bill was for the creators and vendors of AEDTs to perform, or hire specialists to perform, bias

audits of their own proprietary technology. While we support performing audits to ensure hiring software using AI is not biased, it is important that the creators of this technology be responsible for the audits, rather than employers who do not have the ability or the means to assess the performance of these platforms. We ask the Department to explicitly state that the AEDT provider will be responsible for effectuating and paying for the bias audit.

Definition of AEDT

We seek additional clarity under the Department's definition of AEDT. Specifically, we would like to inquire as to whether some job boards would qualify as AEDTs under this section. We do not believe that these mechanisms for accessing resumes should qualify as "substantially" assisting or replacing discretion in making employment decisions. However, we ask for clear delineation for the following scenarios:

- A firm pulls resumes off of a job board where a candidate did not respond to a particular job post but instead added their resume to a database (like Monster.com) and were identified by users of the software putting in search terms using AI. Since the individuals uncovered through

this method did not reply to a specific job opening, we believe this type of software is not considered AEDT under the legislation but want to confirm.

- A job board forwards applications from candidates (along with their resumes) who respond to job postings. Would the Department agree that in cases where all the responses are simply forwarded on to the utilizer that such use would not be considered an AEDT? If the job board did rank the applications for fit etc. would that be considered an AEDT?

In addition, we understand it is the position of a number of software providers that their software is not an AEDT and therefore they do not need to perform an audit at all. It is unclear in the legislation whether a utilizer of hiring software is responsible for making the decision as to whether or not it is an AEDT. We would like to ask the Department to explicitly state that a utilizer of hiring software can rely upon the determination of the supplier of the software as to whether or not it is an AEDT.

Timing of Implementation

Lastly, we feel the Department should clarify that the "one year" look-back period described in § 20871(a)(1) commences on January 1, 2024. This clarification is consistent with the Council's legislative intent as it will provide employers with the envisioned one-year period (starting in January 1, 2023, and ending December 31, 2023) in order to conduct the required bias audits. The clarification is also consistent with the Department's regulatory cadence, since currently employers do not have final guidelines for performing the bias audits.

Again, we appreciate your careful consideration of this matter, thank you for the opportunity to comment in this rulemaking process.

Sincerely,

Douglas W. Klares

Doug Klares
President, New York Staffing Association



January 23, 2023

New York City Department of
Consumer and Worker Protection
42 Broadway New York, NY 10004

Submitted online to <https://rules.cityofnewyork.us/>

**Re: HireVue Comments on Rules Proposed on Dec. 23, 2022
New York City Local Law 144 of 2021**

To Whom It May Concern:

HireVue is a video interviewing and assessment platform. We support both the candidate and employer interview experience in a broad range of industries for customers around the globe.

We share NYC's interest in protecting job seekers by notifying them when AI is being used in the hiring process and ensuring these technologies are audited before making decisions. HireVue previously testified in November relating to New York City's Local Law 144 and we are pleased to see much of the feedback reflected in the revised and proposed rules.

Specifically, we applaud the addition of data requirements in Section 5-302 in recognition of well-established industry best practices. We assume that auditors will maintain professional discretion over the data included as to matters of statistical significance, identification, and/or non-representativeness.

HireVue seeks to offer a fairer, more inclusive, and equitable hiring experience using science-backed methods and industry best practices. Our additional comments today are rooted in our expertise and commitment to the responsible development of our AI technology to make the hiring process better for candidates and employers.

We would like to call attention to a few items for consideration that may have unintended effects based on the recent proposed rules:

- 1) The revised definition of Machine Learning requires three components to meet its definition: the first two components are related to what Machine

Learning is, but the third component is an optional design step in Machine Learning development. Cross-validation is an important step in testing and deploying a *good* machine learning model, but it is not a necessary step.

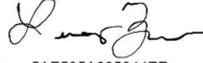
AEDT providers releasing machine learning models without cross-validation would avoid being subject to the Law as it is written.

- 2) Second, some Machine Learning tools use natural language processing to infer user meaning. For example, a chatbot asking minimum qualification screening questions such as “Are you over 18 years old?” would not use Machine Learning on an expected user answer of “yes”, but it would use Machine Learning to infer that a response of “yep” or “sure” also means a “yes” answer. We ask for clarification on whether the law is intended to apply to these simplistic uses of machine learning.
- 3) Finally, although not part of the proposed revisions, we continue to be concerned about the impact of the 10-day notice requirement. The standard hiring process for many jobs has been streamlined through use of technology down to a few days or even a few hours. This waiting period will inevitably disadvantage NYC candidates when candidates from other areas of the region or remote candidates can be hired quicker. This period may also adversely affect NYC employers, who now need to wait to make business critical hiring decisions in industries already strained by labor shortages.

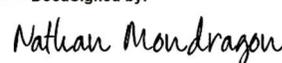
HireVue continues to appreciate the thoughtful revisions reflected in the most recent proposed rules. As always, ongoing dialogue is the key to creating legislation that protects candidates, companies, and innovation. HireVue encourages transparency and supports laws that promote openness and enhance the fairness and efficiency of the hiring process for all individuals.

Thank you for your time and consideration.

Sincerely,

DocuSigned by:

5AE535A605244EF...

Lindsey Zuloaga
Chief Data Scientist

DocuSigned by:

489CA36B1E1D486...

Nathan Mondragon
Chief I/O Psychologist

CHAIR
TRANSPORTATION AND INFRASTRUCTURE

COMMITTEES
FINANCE
HOSPITALS
LAND USE
RULES, PRIVILEGES, AND ELECTIONS
SMALL BUSINESS

SUB-COMMITTEES
COVID RECOVERY AND RESILIENCY



THE COUNCIL
OF
THE CITY OF NEW YORK
SELVENA N. BROOKS-POWERS
MAJORITY WHIP
31ST DISTRICT, QUEENS

ROCKAWAY OFFICE
1931 MOTT AVENUE, SUITE 410
FAR ROCKAWAY, NY 11691
TEL: (718) 471-7014
FAX: (718) 327-4794

LAURELTON OFFICE
222-02 MERRICK BLVD.
LAURELTON, NY 11413
TEL: (718) 527-4356
FAX: (718) 527-4402

CITY HALL OFFICE
250 BROADWAY, SUITE 1865
NEW YORK, NY 10007
TEL: (212) 788-7216
FAX: (212) 227-1210

DATE: January 23, 2023

TO: NYC Dept. of Consumer and Worker Protections

FROM: Council Member Selvena Brooks-Powers

SUBJECT: DCWP rulemaking for Intro 1894-2020 / Local Law 144-2021

My name is Selvena Brooks-Powers, Majority Whip within the New York City Council and I represent the 31st Council district covering parts of Southeast Queens and the Rockaways. I am proud to be one of the 38 lawmakers who voted “yes” on Local Law 144 in December of 2021.

During prior testimony on this issue, I spoke about the massive opportunity offered by Local Law 144 to bring about real progress on racial equity in hiring. I believe this law represents a necessary shift in the City’s approach to bias in the hiring process, which has for too long resulted in disparate treatment of Black and Brown people. I am committed to ensuring the implementation of LL 144 reflects the goals behind its passage: to make the hiring process less biased and more transparent.

Today, I want to discuss concerns I have with the updated rules released last week by DCWP - particularly regarding the proposed definition of Automated Employment Decision Tool, or AEDT. If regulators do not revisit this language, the reach of Local Law 144 may be diminished, and the law will be less effective at reducing racial bias in hiring.

In DCWP’s most recent draft rules, AEDTs are more narrowly defined only as those types of technologies that fully replace, or “overrule,” human decision-making in the hiring process. The problem, however, is that this isn’t how hiring works. Human-decision making always has some role to play in the hiring process. If interpreted strictly, this definition may provide employers a loophole that allows them to evade the requirements of this law.

When my colleagues and I originally voted “yes” on Local Law 144, we voted for transparency in the hiring process to support racial equity. We voted to make it clear that certain types of hiring tools systematically disadvantage Black and Brown people.

We voted for employers to stop hiding behind vague claims about their commitments to workforce diversity.

The intentions of New York City's elected legislators will not be realized if businesses are granted a broad invitation to keep hiding the nature of their hiring practices. Decades of evidence have already demonstrated what types of employment tools employers will use if they are not held accountable for their choices.

In my prior testimony, I also mentioned that many employers and HR tech vendors already collect important data about the racial consequences of their hiring tools and have done so since the civil rights era. As such, Local Law 144 is only revealing what such organizations have known about systematic bias for decades. I strongly urge DCWP to keep the long history of disparate impact reporting in mind as they revisit the appropriate breadth of tools to subject to bias audits.

Finally, I want to re-emphasize my recommendation from my testimony last year: DCWP must recognize the need for expertise in the enforcement of this law. Each rule tweak has implications on the impact of the law and employers' compliance responsibilities. The department should ensure they are staffed with experts to ensure these regulations hold bad actors accountable, while avoiding - as I mentioned previously - hampering legitimate, inclusive AEDTs that seek to expand the hiring pool in the city.

It is my sincere hope that this administration follows through on creating rules that actualize the spirit of this law. Doing so represents an incredible opportunity to disrupt bias in hiring. It would be a shame to squander that opportunity. Thank you.

Sincerely,



Selvena N. Brooks-Powers
New York City Council Majority Whip
District 31, Queens
Chair, Committee on Transportation and Infrastructure



Department of Consumer and Worker Protection
Deputy General Counsel
42 Broadway, 8th Floor
New York, NY 10004

January 23, 2023

Dear NYC Department of Consumer and Worker Protections,

My name is Andrew Hamilton and I am the former president of the New York Metro chapter of the National Black MBA Association.

Over two years ago, I testified before the City Council's Technology Committee in support of Local Law 144. This past November, I reiterated my support during DCWP's hearing regarding the rulemaking process. Today, I continue to believe in the massive potential offered by this legislation, but only if the agency heeds some important feedback on the latest draft rules.

During my November testimony, I encouraged DCWP to be realistic in predicting how employers will feel about the transparency requirements outlined by this law. Specifically, I urged regulators to implement this legislation with a definition of Automated Employment Decision Tools, or AEDTs, that was broad and inclusive of many types of technologies. I argued that limiting the requirement of bias audits to only very modern technologies would perversely incentivize employers to avoid innovation.

My message to the administration today is simple. You cannot implement Local Law 144 with the proposed definition of AEDT unless your goal is to ensure that this will have zero impact. The draft language published by DCWP only requires bias audits in circumstances where an employer has *completely* handed over control of hiring to computers with no human oversight whatsoever. But very few, if any, employers hire in this manner. It would be a waste to limit Local Law 144 to only the most extreme cases of automation, because these situations are not the ones affecting New York City job candidates.

To repeat another point from my November testimony, I urge DCWP to recognize that automated decision-making tools have been around for *many* years; they are not limited to recent advancements like AI. Just think about any standardized test that is scored by a computer and then the scores are used to sort people into "yes" and "no" piles. Black people in the U.S. have a long history of being evaluated by biased technology in contexts like lending, housing, and hiring. This history is not over. The question for New York City is whether we are going to allow it to continue behind closed doors.

The implementation of Local Law 144 represents an incredible opportunity to differentiate between decision-making tools that were built with racial equity in mind and those that were not. As someone who



has been awaiting the enactment of this law for years, it is my sincere hope that the City Council's important work is not stamped out during this final push. Thank you for your time.

Regards,

A handwritten signature in black ink, appearing to read "Andrew Hamilton". The signature is fluid and cursive, with the first name "Andrew" written in a larger, more prominent script than the last name "Hamilton".

Andrew Hamilton

Immediate Past President of New York Metro Chapter
National Black MBA Association



NYC Dept. of Consumer and Worker Protections
42 Broadway #5
New York, NY 10004

January 23, 2023

Good morning,

Thank you to the Department of Consumer and Worker Protections for holding this important convening regarding the NYC Bias Audit Law. I am Billy Council, founder of the CouncilHim Foundation. My organization is committed to helping young men of color realize their educational and professional potential.

The nature of my work has allowed me to witness the realities of bias in the hiring process far too frequently. I know that I am not alone in my frustrations with how often systemic forms of discrimination lock people of color out of opportunities.

When the City Council passed Local Law 144, our elected officials made it clear that New York City workers are owed a fair hiring process. More than that, they sent a message that the public is entitled to know which employers have fair hiring processes and which don't.

To be clear, a bias audit is simply a piece of paper that includes metrics for how likely a hiring tool is to disadvantage minorities. The metrics are generated during a third-party evaluation to ensure they are accurate and objective. An employer who wants to use this hiring tool is simply required to post this document to their hiring website, so that the public understands the diversity implications.

Many advocates across New York City did not even realize how significant Local Law 144 was at the time it was first passed, particularly against the chaotic backdrop of COVID-19. But in recent months, the importance has become clear. We know this because this statute has struck a major nerve with Big Business.

Traditional business interests have responded to Local Law 144 by attempting to co-opt DCWP's rulemaking process for the sake of self-preservation. These well-resourced groups have scrambled to lobby for loopholes to be added that will dramatically weaken the law. They have attempted to characterize transparency as dangerous for economic prosperity. Moreover, they have claimed that the morale of Black and Brown job candidates would be lowered if the truth about most hiring tools came out.

Unfortunately for these employers, the time to be blindly trusting of an employer's intentions regarding workforce diversity is over. The goal for the administration in the coming months must



be to hold firm in this message. If Local Law 144 is executed as our elected officials intended, it will be a way to separate the genuine champions of equity from the woke-washers.

We all know that Big Business has always had a strong voice in New York City, and evidence of this influence can be seen in the most recent draft rules released by DCWP. Under the proposed guidelines, it appears that the vast majority of hiring technologies would not actually be subject to a bias audit. This cannot stand.

My message today is simple: this administration must not allow the whims of large corporations with homogenous workforces to guide decisions about this law. Local Law 144 can only make a difference for New York City workers if it disrupts the secrecy surrounding how these organizations select candidates. You cannot let their fears about reputational risks and preferences for the status quo distract you from the mission of this initiative.

Please do everything in your power to ensure that the transparency regime established by this legislation is a meaningful one. Bias audits should be the rule for most hiring technologies, not the exception. If we can make transparent reporting about the extent of bias in hiring tools universal, behavior change from employers will inevitably follow.

Thank you for your time.

Sincerely,

Billy Council
Founder
COUNCILHIM Foundation Inc.

January 23, 2023

Dear Commissioner Mayuga and DCWP Team,

Please accept our comments on the draft rules for Local Law 144 that the agency published on December 23, 2022.

Our overall position is that Local Law 144 should be implemented with a very inclusive definition of AEDTs. The definition need not rely on the details of how an automated hiring tool is built or used by the employer. Our position on this issue has not changed since providing feedback on the draft rules your team published in September of last year.

As you may recall, in our prior comments, we expressed concerns that DCWP's efforts to limit the requirement of bias audits would create perverse incentives for employers. For example, if the rules only define AEDTs as those technologies reliant on machine learning and/or artificial intelligence, employers could elect to use older hiring assessments as a strategy to avoid being subject to this law.

Upon reviewing the latest version of the proposed rules, we were disappointed to see that the minor edits made to the definition of AEDT only served to further limit the reach of this legislation. We would like to be clear here: if bias audits are only required when a hiring tool is completely replacing or overriding human reviews, there will be no bias audits conducted in New York City. The reality is that almost no employer relies on the output of a single hiring tool to evaluate job candidates and hiring tools are almost never designed to fully override or replace human decision-making.

In terms of what types of technologies *should* be subject to bias audits, we firmly believe that the agency does not need to reinvent the wheel. As we have previously discussed, in contexts like job applicant screening, "automated decision making" is not a new term. Ideally, the scope of Local Law 144 will simply align with the breadth of technologies that would be subject to disparate impact audits under EEOC and OFCCP regulations. Alignment to federal regulation would be particularly appropriate given that bias audit reports are already meant to include disparate impact analyses conducted in accordance with Title VII of the Civil Rights Act.

We further recommend that DCWP should consider providing the public with a list of types of hiring tools that fall under the definition of AEDT and which do not. Please see the attached sample FAQ document for details.

Finally, we do not believe that bias audits should include disparate impact analyses conducted on intersectional demographic groups. Such analyses will almost certainly generate spurious metrics that will only undermine this law's goal of providing reliable and

transparent information about bias testing to the public. We encourage DCWP to think carefully about incorporating standards such as minimum sample sizes to avoid the generation of unreliable bias audit reports.

Thank you for your work thus far. We appreciate the opportunity to contribute our perspective.

Best,

Frida Polli
Founder and Former CEO
pymetrics

Sara Kassir
Former Policy and Research Principal
pymetrics

SAMPLE FAQs for IMPLEMENTATION of LL-144

What is the definition of an automated employment decision tool (AEDT)?

Automated decision-making tools have been around for decades; the term AEDT is not intended to only refer to the most modern technologies, such as those built with artificial intelligence or machine learning.

An AEDT is defined by having the capacity to:

1. Collect data input from or about job candidates
2. Apply some computational process to the data
3. Render an output that is intended to serve as an inference about the candidate's suitability for a job

What are some examples of AEDTs?

The following is a non-exhaustive list of tools that, if administered and scored by a computer, would be considered AEDTs by this law:

- Assessments of cognitive ability or aptitude
- Assessments of personality or behavioral traits or attributes
- Assessments of integrity, judgment, honesty and/or emotional intelligence
- Assessments of "fit" to a role, company, or team culture
- Assessments of technical/professional competencies, such as coding, programming, writing, etc.
- Technologies that automatically analyze responses to case studies, role play, or other simulations
- Technologies that automatically analyze resumes, CVs, or online profiles and produce a score, rank, or other rating
- Technologies that automatically analyze interview recordings and produce a score, rank, or other rating
- Technologies used to verify employment history and/or applicant credibility

Tools that are explicitly excluded from the definition of AEDT include:

- Assessments of language proficiency
- Criminal background checks
- Licensing exams, agency certification, and/or professional certification tests
- Assessments of highly specialized knowledge
- Assessments where responses can *only* be reviewed by a human. (e.g., A written case study is submitted to an HR profession via email who then reviews it manually; this assessment would not be subject to a bias audit. However, if the candidate's submission is also subject to an automated screening algorithm, that process would be subject to a bias audit.)
- Technologies that simply collect information from a candidate without applying analysis or scoring (e.g., a raw data is summarized/presented to HR personnel)

From: Cathy O'Neil <cathy.oneil@gmail.com>
Sent: Friday, January 20, 2023 10:35 AM
To: Rulecomments
Cc: Jacob Appel; Thomas Adams
Subject: [EXTERNAL] Comments on AEDT(Updated) Rules from ORCAA

Some people who received this message don't often get email from cathy.oneil@gmail.com. [Learn why this is important](#)

CAUTION: This email originated from outside of the organization. Do not click links or open attachments unless you recognize the sender and know the content is safe. Forward suspect email to phish@oti.nyc.gov as an attachment (Click the More button, then forward as attachment).

Dear DCWP Rules Committee,

Here are two suggestions on the proposed rules for Automated Employment Decision Tools (Updated).

1. We would suggest that the new definition of AEDT is too focused on the relative influence of the AEDT versus other factors in employment decisions. Many companies using AEDTs use them alongside other hiring methods (e.g. in-person interviews or manual review of resumes), and hiring managers consider all the factors simultaneously in making a decision, sometimes without an explicit weighting scheme. Such use cases should require a bias audit of the AEDT, but employers could argue they are not covered by the current definition since the AEDT is not the sole factor considered, is not weighted more heavily than every other factor, and does not formally overrule a conclusion derived from the other factors. A thumb on the scale could still make things really unfair, depending on how heavy the thumb is.

2. We would also suggest that, to flesh out the definition of AEDT and to clarify its requirements, you could write down specific use cases and say they are (or are not) covered by the law. We think such a list of use cases should include at least (1) a company uses a resume filter to produce a "short list" of candidates from a large set of applicants (this is very common e.g. for call centers); and (2) a scenario like the above, where an AEDT output is one of multiple factors considered in a hiring decision, in order to clarify when it is necessary to have such a system audited.

Thanks very much,
Cathy O'Neil

CEO, ORCAA
orcaarisk.com

Online comments: 19

- **Chris**

Dear Team,

While most comments about the law and its definition have already been made; I'd like to add another perspective on the role of additional players in this new markets, which ought to be regulated (i.e., auditors).

As the attempt is to regulate new technologies in markets (e.g., AI in hiring) questions arise with the alongside development of new services that themselves can become an industry (i.e., auditing in AI).

As AI-auditing itself becomes an industry, which is not an NGO, it can be postulated that they follow economic interests themselves. Thus, as it is necessary to regulate industries (e.g., pharma, finance, technology, etc.) and auditors in other industries have to follow rules, the same principles should apply here.

Again, as auditors are part of an economic industry and thus follow economic interests, those who are being audited cannot ensure that auditors themselves are independent (e.g., favoring large clients) if they do not have to follow rules.

This should be considered when implementing those rules because those who need and want to be audited will a) either have to do it themselves or b) rely on services in an unregulated market itself (i.e., AI-auditing).

I'd like to bring this point up, because companies who deploy AI technology in the market are seen as the only player who needs to be regulated in a new market, while in other industries (e.g., finance) not only the one's who are offering financial products are regulated but also those who are auditing them.

It's unquestionable that new technologies need to be regulated and this never been questioned.

Comment added January 2, 2023 1:16pm

- **Chris**

Dear Team,

While most comments about the law and its definition have already been made; I'd like to add another perspective on the role of auditors. As the attempt is to regulate new technologies in markets (e.g., AI in hiring) questions arise with the alongside development of new services that themselves can become an industry (i.e., auditing in AI).

As AI-auditing itself becomes an industry, which is not an NGO, it can be postulated that they follow economic interests themselves. Thus, as it is necessary to regulate industries (e.g., pharma, finance, technology, etc.) and auditors in other industries have to follow rules, the same principles should apply here.

Again, as auditors are part of an economic industry and thus follow economic interests, those who are being audited cannot ensure that auditors themselves are independent (e.g., favoring large clients) if they do not have to follow rules.

This should be considered when implementing those rules because those who need and want to be audited will a) either have to do it themselves or b) rely on services in an unregulated market itself (i.e., AI-auditing).

I'd like to bring this point up, because companies who deploy AI technology in the market are seen as the only player who needs to be regulated in a new market, while in other industries (e.g., finance) not only the one's who are offering financial products are regulated but also those who are auditing them.

It's unquestionable that new technologies need to be regulated and this never been questioned.

Comment added January 2, 2023 1:17pm

- **Tim**

My main comment would be to make sure you understand the status quo in hiring, and how utterly dreadful it is. Traditional hiring is fundamentally flawed from a bias perspective. Any job application typically starts with a candidate submitting a CV, containing innumerate pieces of noise that reveal their gender, race, religion, socioeconomic status and many other things that are irrelevant to figuring out 'is this the best person for the job'?

Once the CV is received, a human then – if the candidate is lucky – scans the CV for 10 seconds and makes a yes or no call based on their gut. This process has very low accuracy (the best candidate often misses out), is fundamentally biased, costs a lot of money, slows down the hiring process, and leaves 0 candidates with any meaningful feedback – they just get the 'sorry, not sorry' email.

By using products that actually measure candidates' on things that are scientifically proven to predict job performance, such as their skills, personality and intelligence, is clearly a fairer way, that some random person glancing at CV and deciding they don't like where the candidate went to school.

There is 0 current auditability of current practices. No company in the world could explain why they rejected a candidate based on their CV. How could they? The logic of the decision – made in 10 seconds let's not forget – is not recorded anywhere. And how could it be? What would they record? 'Oh I didn't like their name', 'Their formatting was bad', 'They had a spelling error' or any number of ridiculous reasons – I really suggest you research how this is currently done in practice to realise how bad it is.

At least with products that actually measure things, you'll always be able to go back and say 'Ok, this job required Python skills and this candidate scored 10% on a Python test, so that's why they got rejected'. No matter how imperfectly the Python test is constructed, surely that's a more legitimate reason for rejection than just someone's whim?

And this is just the screening stage. The rest of the process of traditional hiring is also filled to the brim with bias, like unstructured interviews where all decisions are based on gutfeel and something approaching astrology.

Please, understand the current market conditions and how incredibly unfair traditional hiring is before implementing this law and throwing the baby out with the bathwater.

Comment added January 12, 2023 4:41am

- **Mike Fetzer**

I applaud and fully support the efforts of the City of New York and the Department of Consumer and Worker Protection (DCWP) for being at the forefront of addressing the potential for bias in the use of artificial intelligence and machine learning applications to make automated employment decisions. In order to best facilitate the efficient and intended impact of this legislation, I would respectfully recommend that the DCWP adopt the standard definitions of artificial intelligence and machine learning that were established by Congress in the National Artificial Intelligence Act of 2020 at sections 5002(3) and 5002(11) respectively:

(3) ARTIFICIAL INTELLIGENCE. The term "artificial intelligence" means a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations or decisions influencing real or virtual environments. Artificial intelligence systems use machine and human-based inputs to—

- (A) perceive real and virtual environments;
- (B) abstract such perceptions into models through analysis in an automated manner; and
- (C) use model inference to formulate options for information or action.

(11) MACHINE LEARNING. The term “machine learning” means an application of artificial intelligence that is characterized by providing systems the ability to automatically learn and improve on the basis of data or experience, without being explicitly programmed.

These definitions have already been adopted by the U.S. Equal Employment Opportunity Commission (EEOC), the U.S. Department of Justice (DOJ), and numerous other Federal and State agencies. Using commonly accepted definitions of artificial intelligence and machine learning will enable the DCWP to focus the definition of automated employment decision tools (AEDTs), draw upon the precedence set at the state and federal levels, and avoid any confusion and/or conflicts that might arise from the use of non-standard definitions. Further, I would recommend striking the terms “statistical modelling” and “data analytics” from LL 144 to maintain focus on artificial intelligence and machine learning.

Alternatively, if the DCWP decides to retain or revise the current definition of machine learning, statistical modelling, data analytics, or artificial intelligence, I would recommend the following be added to part (iii) of the definition:

Cross-validation is a statistical method of evaluating and comparing machine learning algorithms by dividing a single, identical data set into two parts: training and testing data. The training data can be further segmented into training set used to fit the parameters of the model and validation set used to optimize the model parameters. The testing

data is then used to provide an unbiased evaluation of a final model fit on the training data set.

Comment added January 19, 2023 3:35pm

- **Holistic AI Team**

Please see the attached comments relating to issues surrounding small sample sizes and concerns about how the metric for regression systems can be fooled by bimodal distributions.

[Comment attachment](#)

DCWP-Comment-on-Updates-.docx

Comment added January 20, 2023 11:44am

Holistic AI questions regarding the DCWP's updates to Proposed Rules to implement Local Law 144 of 2021 (Automated Employment Decision Tools)

January 20, 2023

Department of Consumer and Worker Protection
DCWP Commissioner Vilda Vera Mayuga
42 Broadway, New York
NY, 10004

RE: Proposed Rules to implement Local Law 144, Automated Employment Decision Tools

Dear Commissioner of the New York City Department of Consumer and Worker Protection (DCWP),

Thank you for the opportunity to provide questions on this important matter.

1. About Holistic AI

Holistic AI is an AI Governance Risk and Compliance (GRC) company, with a mission to empower enterprises to adopt and scale AI with confidence. Holistic AI has a multidisciplinary team of AI and machine learning engineers, data scientists, ethicists, business psychologists, and law and policy experts.

We have deep practical experience auditing AI systems, having assured over 100 enterprise AI projects covering more than 20,000 different algorithms. Our clients and partners include Fortune 500 corporations, SMEs, governments, and regulators. We work with several companies to conduct independent bias audits, including in preparation for Local Law 144.

2. Key questions

We would like to clarify some points regarding updates to the Proposed Rules and their implementation.

2a. Calculating Impact Ratios with Small Sample Size

The Proposed Rules specify that the ethnicity/race categories that should be examined are Hispanic or Latino, White, Black or African American, Native Hawaiian or Pacific Islander, Asian, Native American or Alaska Native, or two or more races. However, there are likely multiple categories with small samples, particularly for the Native Hawaiian or Pacific Islander, Native American or Alaska Native, and two or more races categories. However, the DCWP does not provide any clarification on what is considered an adequate sample size for analysis to be meaningful. Notably, examples provided by the updated Proposed Rules keep the categories separate, with one category listed representing less than 1.5% of the workforce, but the EEOC's clarifications on the Uniform Guidelines⁸⁷ specify that adverse impact analysis should only be carried out for groups who represent at least 2% of the labor force, meaning that the examples provided conflict with this

⁸⁷ <https://www.eeoc.gov/laws/guidance/questions-and-answers-clarify-and-provide-common-interpretation-uniform-guidelines>

guidance. Additionally, analyses based on samples representing less than 2% of the population are unlikely to be meaningful.

One approach to increase sample sizes is to combine the Native Hawaiian or Pacific Islander or Native American or Alaska Native categories into one broader “Other” category⁸⁸. However, there is no guarantee that this will increase the sample to a sufficient size for a robust analysis and could mean that it is harder to identify and mitigate bias for particular subgroups if they do not have their own category.

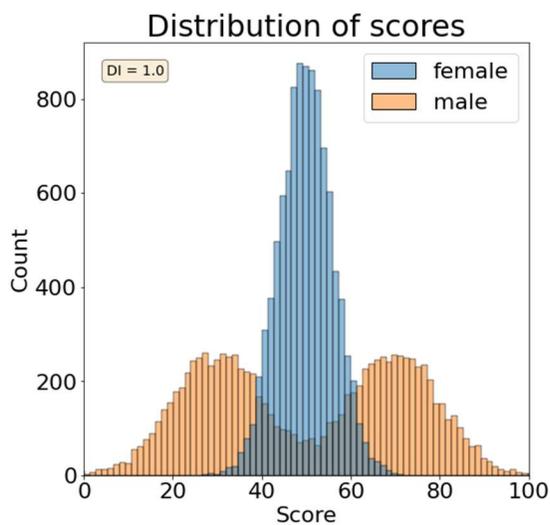
We therefore request that the DCWP clarify what is considered an adequate sample size for analysis to be meaningful and to release guidance on calculating impact ratios when sample sizes are small or propose an alternative metric that is more suited to smaller samples. Such guidance would be particularly useful for intersectional analyses, where sample sizes are often small.

2b. Calculating Impact Ratios for Regression Systems

While the revised metric for calculating the impact ratio for regression systems using the median is a notable improvement over the initially proposed metric using the average score, this metric still has some concerning limitations. Most notably, the impact ratio metric is not always adequate for detecting bias in regression data.

Example 1: Bimodal Distribution vs Unimodal Distribution

Let’s consider the case where male and female candidates are scored from 0 to 100. Whereas male candidates consistently get scores around 50 (unimodal distribution), female candidates seem to be scored either approximately 25 or 75 (bimodal distribution), as seen in the figure below. The median of the full dataset (across males and females) is 50, because half the data falls below it and the other half falls above it.

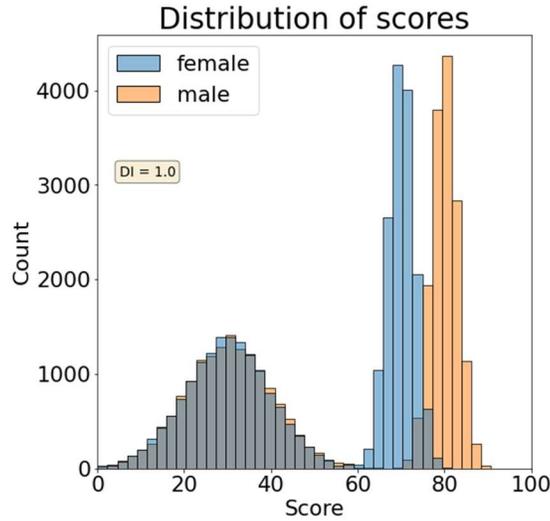


If all candidates scoring above the median value of 50 are hired, then the data will be perfectly fair. However, if only the top 20% of candidates are hired, then almost all chosen candidates will be female. Given that the median is rarely used as a cutoff score, the system is likely to result in biased outcomes even if the audit does not find any evidence of bias using the median metric.

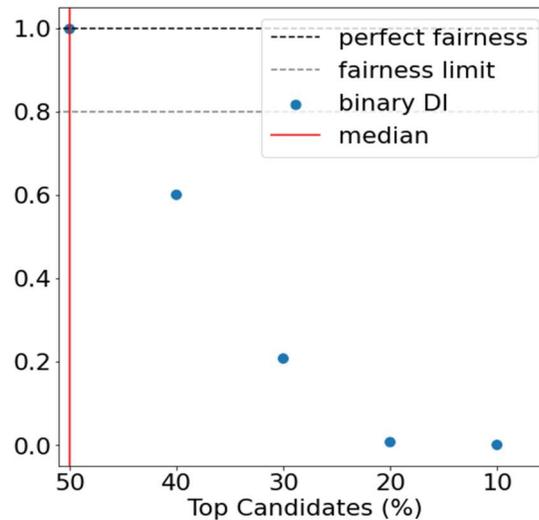
⁸⁸ <https://www.eeoc.gov/data/introduction-race-and-ethnic-hispanic-origin-data-census-2000-special-eeo-file>

Example 2: Two Bimodal distributions

In this example, male candidates are scored either approximately 30 or 70, whereas female candidates are scored either approximately 30 or 80. In this example, seen in the figure below, there are two peaks for each, with the lower one being consistent for both male and females while the higher one is slightly different. Like in the above example, the median of the full dataset is approximately 50, because half the data falls below it and the other half falls above it.



If all candidates scoring above 50 are hired, then the data will be perfectly fair. This is because we are essentially reducing our data to a binary pass/fail classification, and the difference in the higher scores for male and females will be masked. However, as soon as our notion of success changes, big differences are revealed. To better observe this phenomenon, we can calculate how the impact ratio varies when the candidates hired are respectively the top 50% (which is equivalent to the median), 40%, 30%, 20% and 10%.



As is seen from the figure above, the computed binary disparate impact will greatly depend on the threshold we use. At the median value, we obtain perfect fairness. For any values above the median, the fairness rapidly decreases due to the distribution of the data.

Given these concerns, we encourage the DCWP to consider alternative metrics that would be better suited for measuring regression bias. For alternative metrics, see Holistic AI's open-source library.⁸⁹

3. Holistic AI resources

In lieu of the fact that the field of algorithm audit and assessment is relatively new, below we link some resources and references to our open source and academic research.

- [Holistic AI Open Source](#)
- [The New York City Bias Audit Law: Regulating AI and automation in HR](#)
- [Towards Algorithm Auditing: A Survey on Managing Legal, Ethical and Technological Risks of AI, ML and Associated Algorithms](#)
- [Systematizing Audit in Algorithmic Recruitment](#)
- [Perceived Fairness of Algorithmic Recruitment Tools](#)
- [Overcoming Small Sample Sizes When Identifying Bias](#)

4. Concluding statement

Holistic AI welcomes the opportunity to provide comments on this important matter. We appreciate the open, transparent and collaborative approach taken by the DCWP.

We support the important objectives of Local Law 144. We stand ready to support the DCWP, the New York City Council or other public authorities involved in the implementation and enforcement of this important law.

Please contact we@holisticai.com for any further information or follow-up on this submission.

Sincerely,
Holistic AI
<https://www.holisticai.com/>

⁸⁹ <https://www.holisticai.com/open-source>

- **Barbara Kelly**

The Institute for Workplace Equality (“IWE” or “The Institute”) submits the attached comments in response to the New York City (“NYC” or the “City”) Department of Consumer and Worker Protection’s (“DCWP” or the “Department”) invitation. The Department’s Notice of Proposed Rules is seeking to clarify the requirements set forth by NYC’s Local Law 144 that will regulate the use of automated employment decision tools (“AEDT”) wherein hiring or promotion decisions are made or substantially assisted by algorithmically-driven mechanisms.

[Comment attachment](#)

2023.01.20-IWE-NYC-Local-Law-144-Letter-of-Comment.pdf

Comment added January 20, 2023 12:29pm

- **Joseph Abraham**

Thank you for the opportunity to provide comment. Note that these comments reflect my personal observations and recommendations, and not necessarily the position of my employer.

[Comment attachment](#)

proposed-NYC-rule-comments-v2.docx

Comment added January 20, 2023 2:06pm

1. Most importantly, the proposed rule amendments could benefit from further clarifications to state that human guided processes (e.g., traditional psychometric measurement techniques) are not intended to fall under the definition. Otherwise, employers may fall back into relying upon inherent human bias (who you know, similar-to-me) at the expense of job-related assessment tools (such as job knowledge exams or self-report personality inventories). To that end, I would recommend the following (blue highlighted) language be added to clarify the definition of “machine learning, statistical modeling, data analytics, or artificial intelligence.”

Machine learning, statistical modelling, data analytics, or artificial intelligence. “Machine learning, statistical modelling, data analytics, or artificial intelligence” means a group of mathematical, computer- based techniques:

- i. that generate a prediction, meaning an expected outcome for an observation, such as an assessment of a candidate’s fit or likelihood of success, or that generate a classification, meaning an assignment of an observation to a group, such as categorizations based on skill sets or aptitude; and
- 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.

Techniques wherein human judgment determines the inputs, the relative importance placed on those inputs, and other parameters, do NOT constitute “machine learning, statistical modelling, data analytics, or artificial intelligence,” even if statistical information is used to help inform human judgment.

2. There remains confusion regarding numerous scenarios posed during the previous comment period. I would recommend that the agency issue a questions-and-answers document to clarify outstanding questions posed during both comment periods. An example follows which would benefit from such a questions-and-answers document to guide employers.

“The proposed definition for “automated employment decision tool” could benefit from further explanation and illustrations, particularly with regard to the phrase “or to use a simplified output as one of a set of criteria where the output is weighted more than any other criteria in the set.” What about a situation where a candidate score is generated that is based 90% on structured panel interview ratings (numerical ratings averaged) and 10% on a tool that uses machine learning? Or a score that is based 90% on a job knowledge exam (multiple choice test score) and 10% on a tool that uses machine learning?”

3. The bias audit section lacks clarity, particularly around what statistics must be reported in varying scenarios. For example, the distinction between the two statements below is not clear.

— —

“Where an AEDT selects candidates for employment or employees being considered for promotion to move forward in the hiring process or classifies them into groups, a bias audit must, at a minimum:”

Versus

“Where an AEDT scores candidates for employment or employees being considered for promotion, a bias audit must, at a minimum:”

Is the distinction here “selects” and “classifies” versus “scores?” If so, that distinction is unclear and would benefit from elaboration and/or further examples.

4. The use of intersectional category reporting for impact ratios is not called for by the federal Uniform Guidelines on Employee Selection Procedures and therefore adds administrative burden to employers. Further, sample sizes can dwindle with intersectional analysis and result in unstable results. Similarly, the published, peer-reviewed literature on assessments rarely reports intersectional group difference results, leading to difficulty with interpreting findings in the context of existing scientific knowledge. I would strongly advise against requiring intersectional reporting, which adds burden with little associated benefit.
5. “Scoring rate” is not a term commonly used in assessment science, nor is the interpretation of “scoring rate” clear. There also remains a serious concern over equating “bias” with differential “scoring rates” or selection rates, without consideration of the validity of the underlying assessment scores.
6. A bias audit using selection rates from multiple employers is unlikely to be practical or meaningful in many, if not most situations. Selection rates are a product of both (a) the selection tool itself and (b) how employers use the tool, which can vary substantially from employer to employer. For example, employers vary widely in the extent to which they place weight on selection tools and the cutoff scores they use, if any, based on selection tool scores.

- **BABL AI Team**

Please consider the attached comments on the new proposed rules, submitted on behalf of the team at BABL AI.

[Comment attachment](#)

BABL-AI-LL144-Comments-III.pdf

Comment added January 22, 2023 4:42pm

- **Merve Hickok (Alethicist.org)**

Thank you for the opportunity to provide further comments as DCWP continues to clarify this pioneering law. Please find attached recommendations and feedback submitted on behalf of Alethicist.org.

[Comment attachment](#)

DCWP_NYC-Public-Comment_Merve-Hickok_Jan2023.pdf

Comment added January 22, 2023 7:33pm



Written Comments from Merve Hickok
Regarding Proposed Rules on NYC Local Law 144 of 2021 in relation to Automated Employment
Decision Tools

TO: New York City Department of Consumer and Worker Protection (DCWP)
Rulecomments@dca.nyc.gov

01/22/2023

Dear Chair and Members of DCWP,

As the Founder of Alethicist.org and Lighthouse Career Consulting LLC, I welcome another round of public comment opportunity for the rules proposed by DCWP⁹⁰ on implementation of Local Law 144 of 2021⁹¹, regulating automated employment decision tools (AEDT).

I submitted written comments to DCWP's previous rulemakings – in June⁹² and October 2022⁹³ I also published on this local law and its impact extensively.^{94 95}

My work is focused on Artificial Intelligence (AI) ethics and AI policy and regulation globally. I am also a certified human resource professional with almost two decades of experience across Fortune 100 companies. As the founder of Alethicist.org, I provide research, training, and consulting on how to develop, use and govern algorithmic systems in a responsible way. I am also the Research Director at Center for AI & Digital Policy and lecturer on data science ethics at University of Michigan.

I commend DCWP on providing further clarification on several concepts related to this very important and impactful law. In particular, the below are clear and reflect the spirit of the law.

⁹⁰ New York City Department of Consumer and Worker Protection (December 2022) Notice of Public Hearing and Opportunity to Comment on Proposed Rules to implement new legislation.

<https://rules.cityofnewyork.us/wpcontent/uploads/2022/12/DCWP-NOH-AEDTs-1.pdf>

⁹¹ Local Law 144 of year 2021: <http://nyc.legistar1.com/nyc/attachments/c5b7616e-2b3d-41e0-a723cc25bca3c653.pdf>

⁹² Merve Hickok (June 6, 2022). Written Comments regarding NYC legislation on Automated Employment Decision Tools (Local Law 144), submitted to New York City Department of Consumer & Worker Protection.

<https://rules.cityofnewyork.us/rule/force-fed-products-open-captioning-in-motion-picture-theaters-and-automatedemployment-decision-tools/>

⁹³ Merve Hickok (October 23, 2022). Written Comments regarding NYC legislation on Automated Employment Decision Tools (Local Law 144), submitted to New York City Department of Consumer & Worker Protection.

<https://www.nyc.gov/assets/dca/downloads/pdf/about/PublicComments-Proposed-Rules-Related-to-Automated-Employment-Decision-Tools.pdf>

⁹⁴ Center for AI and Digital Policy (August 11, 2022). Policy Brief – NYC Bias Audit Law.

<https://www.caidp.org/app/download/8407232163/Alethicist-NYC-Bias-Audit-Law-08112022.pdf>

⁹⁵ Center for AI and Digital Policy (September 2, 2022). Policy Brief - State of AI Policy and Regulations in Employment Decisions. <https://www.caidp.org/app/download/8410870963/Alethicist-HumanResourcesRecentDevelopments-09022022.pdf>

Published results of bias audit: Ensuring the selection rates and impact ratios for all categories are transparent is extremely critical for the success and impact of Local Law 144. A bias audit which only calculates these rates and ratios, but then is not transparent is not beneficial to anyone. Transparency creates accountability.

Definition of “independence”: I congratulate the Department on this clarification and the confirmation that auditor cannot be internal to either the vendor or the employer.

Similarly, it is equally important a lawyer-client privileged relationship is not used for audit purposes. Such a relationship cannot be independent. It can also have negative consequences for DCWP and impacted parties and their access to adequate audit materials on AEDTs if/when needed.

There is still need for further clarification in the proposed rules – to reflect Subchapter 25 of the Code and the intended scope of NYC Council’s decision. Also concerning are attempts to

- narrow the scope of the law by diluting the definition of AEDTs, and
- remove the requirement to make quantitative results of the bias audit public.

Definition of “automated employment decision tool”: Proposal has multiple qualifiers in definition of output, and tools which significantly narrow the scope and intent. Qualifiers such as “to substantially assist or replace discretionary decision making”, “rely solely on a simplified output”, “with no other factors considered” or “weighted more than any other criterion in the set” compromises the scope, and intent of the law to protect candidates in NYC.

Employers use multiple methods during the hiring process and invest in AEDTs mainly for efficiency, speed, and cost-cutting reasons. Employers usually do not elaborate weighing mechanisms for all the criteria used in a decision. Therefore, subjective assessments and qualifiers create loopholes.

- Well-informed employers who do not want transparency may use these loopholes to escape scrutiny and transparency.
- Uninformed employers may not understand the extent of bias in AEDTs⁷, hence not feel the necessity to examine further.
- Candidates, in their individual capacity and without protection of this law, never have a way to find out what tools they are subjected to, or whether the tool was biased.

I strongly recommend DCWP not include any qualifiers which weaken the scope and enforcement of the law.

The definition should simply be: ‘automated employment decision tool’ means any computational process, derived from machine learning, statistical modeling, data analytics, or artificial intelligence, that issues simplified output (including prediction, classification, score, tag, categorization, recommendation, ranking, or similar results) used to automate or support employment decision-making that impact natural persons.’”

Definition of “Machine learning, statistical modelling, data analytics, or artificial intelligence”:

- The word “and” at the end of each point makes the definition extremely narrow, focusing on mainly on ML techniques. Not all AEDTs are AI/ML-based systems.
- A model which does not identify the inputs, or the relative importance of the inputs/other parameters can still be biased. In other words, inputs may be determined by human developers and still reflect bias of developers or historical biases in the dataset. Therefore, should be subject to bias audit.

I recommend DCWP not to try to define these techniques. This would also allow the law to be flexible for future technological innovation and new methods.

Definition of “bias”: The limitations of dataset and model, and design/model decisions made by vendor and employers must be included as audit criteria. For example, different cut-off thresholds may change the results significantly across different groups.

⁷⁷ Merve Hickok (July 2020) Why was your job application rejected: Bias in Recruitment Algorithms? Medium. <https://medium.com/@MerveHickok/why-was-your-job-application-rejected-bias-in-recruitment-algorithms-part-1-4ab24573c384>

§ 5-302 Data Requirements for analysis of Impact Ratio: Proposal suggests “an employer may rely on a bias audit of an AEDT that uses the historical data of other employers only if it provided historical data from its use of the AEDT to the independent auditor for the bias audit or if it has never used the AEDT” If a vendor uses a general model and does not tailor / train models with client’s own data, an aggregate dataset made up of all clients’ historical data may suffice for a bias audit & report.

- However, if the vendor tailors its model and/or trains model with client employer’s data, there is a possibility for the outcomes to vary across different employers. If an employer is liable for the outcomes of an AEDT model trained with its own data, then running an audit on the aggregate data of many employers may not correctly reflect impact ratios for different employers. This necessitates separate audits for each model in use. In this case, a possible solution would be to conduct 2 separate selection rate and impact ratios analyses: 1) for all candidates in the system, regardless of clients, AND 2) a separate client-based analysis.

Impact Ratio analysis, as defined by Local Law 144, has been around for decades and should have been utilized by both vendors and employers as good practice, this should not be an extra burden.

Proposal defines Bias Audit as a simple Impact Ratio analysis. Vendors should already be monitoring client models for quality and liability purposes. Employers should already be monitoring their hiring pipelines. This approach would also give more confidence to employers and provide evidence of good will and practice.

- § 20-871(b)(2) of the Code requires employers to notify on “job qualifications and characteristics that such AEDT will use in the assessment of such candidate or employee”. This would mean that for each audit, the data should also be analyzed for different ‘job categories.’ For example, AEDT outcomes might be biased for females in administrative jobs with traditionally female hires, or vice versa for executive or technical jobs with traditionally male hires. If multiple job categories (which require different skills, traits) are combined for a single analysis, the aggregate dataset may not correctly reflect the possible biases in hiring decisions. Using job categories would also allows vendors and employers to compare their results to existing labor market demographics.

Notice to Candidates: § 20-871(b)(2) of the Code also requires the information about specific job qualifications and characteristics the tool will use in the assessment of candidate so that candidates can request an alternative selection process or accommodation. To be able decide if accommodation is needed and to request an appropriate accommodation or alternative, a candidate needs to know if the use or assessment of the AEDT require any physical, cognitive, or motor skills, or mental, emotional, character capabilities or competencies. DCWP Proposal omits this last part in the notice requirement. Ideally, the information should be both included in the candidate notice, and in published audit results.

We need vendors to create products responsibly and employers not to create disparate impact. Transparent audit results give the businesses a chance to walk their talk about equity, diversity and inclusion. Scope of law and transparency requirements will define if NYC law can be a blueprint for future local, state and national jurisdictions around the world; and prioritize diversity, equity, and civil rights.

Thank you for your consideration of my views. I would welcome the opportunity to discuss further about these recommendations.

A handwritten signature in black ink, appearing to read 'Merve Hickok', is positioned above the typed name.

Merve Hickok, SHRM-SCP

Founder, AIethicist.org, and Lighthouse Career Consulting LLC merve@lighthousecareerconsulting.com

- **retrain.ai Team**

Please see the attached comments relating to questions concerning adequate data sample sizes, sufficient historical data, and the definition of test data.

[Comment attachment](#)

JANUARY-retrain.ai-Comments-on-NYC-Local-Law-144.pdf

Comment added January 22, 2023 11:37pm

- **Fred Oswald**

Thank you for this opportunity to provide further input on this important law. My comments are attached, in hopes they are helpful.

[Comment attachment](#)

DCWP-NYC-AEDT-comment-Fred-Oswald-230123.pdf

Comment added January 23, 2023 8:52am



January 23, 2022

Department of Consumer and Worker Protection
Commissioner Vilda Vera Mayuga
42 Broadway, 9th Floor
New York NY, 10004

RE: Proposed Rules to implement Local Law 144, Automated Employment Decision Tools

Dear Commissioner of the New York City Department of Consumer and Worker Protection (DCWP),

Thank you for the opportunity to share a set of interrelated concerns with the DCWP's Proposed Rules to implement Local Law 144, Automated Employment Decision (updated).

Definitional concerns

1. "Machine learning, statistical modeling, data analytics, or artificial intelligence" - Selection tools would be exempted from this definition if they did not* use cross-validation (required in part iii of the definition), and yet cross-validation is a recommended practice to help ensure, all other things being equal, that selection models are not overfitting data and overstating their benefits.

2. "Simplified output"

(2a) The definition depends on the definition of "machine learning, statistical modeling, data analytics, or artificial intelligence" which itself is problematic (see above);

(2b) The end of the definition states that "[i]t does not refer to the output from analytical tools that translate or transcribe existing text, e.g., convert a resume from a PDF or transcribe a video or audio interview," and yet this output very often produces or contributes meaningfully to a "score...tag or categorization...or ranking." Therefore, it is very unclear, and concerning, why these forms of processed applicant data -- involving some of the most machine-learning-heavy processes -- would be excluded from simplified output.

3. "Scoring rate" - Although the scoring rate is used in machine learning, reporting subgroup scores above the overall median in the selection context is an unorthodox and (at best) indirect way to indicate the performance of an AEDT. For instance, the scoring rate ignores

underlying subgroup score distributions and proportions, and the scoring rate does not necessarily reflect the actual effect of selection on the overall distribution of scores.

Additional concerns

4. Sampling distortions - The independent bias audit is based on "historical data regarding applicant selection that the vendor has collected from multiple employers." Thus, information from the audit report is a gross aggregate: e.g., across wide-ranging industries, jobs, and specific applicant populations. The selection rates and impact ratios in employer settings may also vary greatly, and yet they will be averaged away in the vendor audit report
5. Statistical inaccuracies - Setting aside issue #4 above, small sample sizes are associated with greater inaccuracies in the selection rates and impact ratios reported, yet the audit report does not come with any indication of these inaccuracies (e.g., 95% confidence intervals). Note that because intersectional samples refine the overall sample, they can be more important yet will be even less accurate than higher-level aggregates (a potential tradeoff).

Biographical information

My professional background involves 23+ years of psychometrically developing and evaluating selection measures (e.g., job knowledge and skills, personality, interests) used within organizational, military, and educational settings. In the last several years, I have engaged in a range of activities relevant to the use of AI assessments in employment settings: e.g., teaching machine learning workshops and seminars, publishing peer-reviewed research articles, and serving on technical advisories (e.g., with the Society for Industrial-Organizational Psychology, and the Institute for Workplace Equality). Current roles of relevance include serving as Chair of the Board on Human-Systems Integration (BOHSI) at the National Academies (which recently produced a report on human-AI teams) and member of the National AI Advisory Committee (NAIAC), which advises the Secretary of Commerce and the President. Comments expressed herein are solely my own, and do not represent the views of these aforementioned groups.

With appreciation for these continued important deliberations, please reach out any time to discuss further.

Sincerely,



Frederick L. Oswald
Professor and Herbert S. Autrey Chair in Social Sciences
Department of Psychological Sciences
Rice University

- **Julia Stoyanovich**

See attachment

[Comment attachment](#)

Stoyanovich_144_Jan23_2023.pdf

Comment added January 23, 2023 9:18am

- **Daniel Schwarz**

Comments attached

[Comment attachment](#)

NYCLU-Testimony-DCWP-Employment-ADS-20230113.pdf

Comment added January 23, 2023 11:21am

- **Mitch C. Taylor**

Attached please find the public comment from SHRM, the Society for Human Resource Management.

[Comment attachment](#)

SHRM-NYC-AEDT-Revision-Comment-1.23.2023.pdf

Comment added January 23, 2023 12:31pm

- **Rose Mesina**

Thank you for this opportunity to provide further input. Please see attached for comments.

[Comment attachment](#)

Second-Set-of-Comments-Questions-on-the-Proposed-Rules-January-23-2023.pdf

Comment added January 23, 2023 2:52pm

January 23, 2023

Hon. Vilda Vera Mayuga, Esq.
Commissioner
Department of Consumer and Worker Protection
42 Broadway
New York, N.Y. 10004

RE: Comments on New Proposed Rules to Implement
Automated Employment Decision Tools Law (Local Law
144)

Dear Commissioner Mayuga:

Resolution Economics, LLC, an international consulting firm with offices in New York, Los Angeles, Chicago, Washington, D.C, Charlotte, N.C., and London, makes this submission in response to the Notice of Public Hearing and Opportunity to Comment that was issued by the New York Department of Consumer and Worker Protection on December 15, 2022, regarding the revised Proposed Rules implementing the Automated Employment Decision Tools law (“AEDT law”), which the Department has announced it will begin enforcing on April 15, 2023.

Resolution Economics provides economic and statistical analysis, investigations and advisory services, tailored technology, and analytical solutions as well as expert testimony to law firms, companies, and government agencies. We specialize in global labor, employment, and litigation-related matters across every industry. Our professionals include highly trained and technical team members with PhDs, MAs, MBAs, CPAs, CFEs, and other qualifying expertise. Resolution Economics has been and is currently advising employers how to evaluate the impact and navigate compliance obligations around automated employment decision tools.

Based on our experience in this area, we submitted questions in October 2022 regarding the previous iteration of the Proposed Rule. We are pleased to see that the Department has addressed some of our concerns in the revised Proposed Rules. Improvements include:

- (a) A revised definition of “independent auditor” that makes clear that a bias audit must be conducted by a person or entity that is truly autonomous from and not subject to control by the employer and that plainly states that such audits are not permitted to be conducted by an employer’s employees; and
- (b) Clarification that separate impact ratios must be provided for race, for gender, and for intersectional categories.

However, the revised Proposed Rules leave some questions unanswered. Moreover, the revisions themselves raise additional issues and concerns about how the bias audits required by the AEDT law are to be conducted. We address a number of these issues and concerns below.

1. Use of Test/Synthetic Data

Section 5-302(a) of the new Proposed Rules states:

(a) A bias audit conducted pursuant to section 5-301 of this Chapter must use historical data of the AEDT. If insufficient historical data is available to conduct a statistically significant bias audit, test data may be used instead.

(b) If a bias audit uses test data, the summary of results of the bias audit must explain why historical data was not used and describe how the test data used was generated and obtained.

This Proposed Rule leaves unaddressed several key issues regarding the use of test (also known as “synthetic”) data for bias audits.

- a. Can test/synthetic data be used not only during the first year of an AEDT’s use but also later if the AEDT or the employer’s use of the AEDT changes?

The statement that test data may be used “if insufficient historical data is available” could be read as permitting the use of test data only during the first year – that is, only where the tool has not been used at all previously (or only very sporadically). However, test/synthetic data may be useful not only during the first year of an AEDT’s use but also throughout an AEDT’s use.

These tools will be evolving and their use likely will expand within a business or organization over time. Rather than waiting to determine *after the fact* that a change – either in the tool itself or in the way the tool is used by the employer – results in unintentional bias, it would be wise to use test/synthetic data to test the changes first in a controlled environment. The Proposed Rule should make clear that such use of test data is permitted.

- b. The Proposed Rules should clarify what types of test/synthetic data are permitted to be used for a bias audit

The new Proposed Rules provide a very rudimentary definition of “test data,” defining it only as “data used to conduct a bias audit that is not historical data.” (Section 530). This barebones definition may enable unscrupulous employers or vendors to skirt the very purpose of the AEDT law: ensuring a valid assessment of whether an AEDT produces biased outputs.

That is because test data can be fabricated to artificially produce “unbiased” results, thus making an AEDT appear to be unbiased even if it is actually biased. For example, if certain employment predictors used by the AEDT

are simulated in the test data based on a different distribution than they have in reality, the AEDT would generate “unbiased” employment outcomes between protected and unprotected groups in the bias audit while such unbiasedness would not be achieved if the AEDT were applied to real historical data.

Given this reality, the Department may want to consider providing more specific requirements for the creation of test data that are to be used in a bias audit. Such requirements might include the number and the nature of simulated predictors that should be used in the audit test data and the criterion the test data need to satisfy to ensure the data’s similarity to real historical data. Specifically, the Department might consider requiring that all potential predictors used in the AEDT need to be in the test data and the test data need to be calibrated to external benchmarks to ensure the similarity of test data and real historical data.

2. Missing race and gender data

Not all AEDTs seek data regarding race, ethnicity or gender. And even where the AEDT does ask for such information, an increasing number of individuals choose not to disclose their race, ethnicity and/or gender. The new Proposed Rules do not provide any guidance as to how a bias audit should take account of such situations.

Is imputation allowed? Should individuals who choose not to identify race, ethnicity or gender be excluded from the respective race or gender analyses?

3. Allowing multiple employers to rely on the same bias audit

Section 502(c) of the new Proposed Rules states that

A bias audit of an AEDT used by multiple employers or employment agencies may use the historical data of any employers or employment agencies that use the AEDT. However, an employer or employment agency may rely on a bias audit of an AEDT that uses the historical data of other employers or employment agencies only if it provided historical data from its use of the AEDT to the independent auditor for the bias audit or if it has never used the AEDT.

The new Proposed Rules’ Statement of Purpose indicates that this language is intended to “clarify[] that multiple employers using the same AEDT may rely upon the same bias audit so long as they provide historical data, if available, for the independent auditor to consider in such bias audit.” But this new Proposed Rule raises a host of questions and potential problems.

These problems can be illustrated by looking at one of the examples provided in the new Proposed Rule itself. That example (appended to Section 5-301(b)), states:

An employer wants to use an AEDT to screen resumes and schedule interviews for a job posting. To do so, the employer must ensure that a bias audit of the AEDT was conducted no more than a year prior to the planned use of the AEDT. The employer asks the vendor for a bias audit.

The vendor provides historical data regarding applicant selection that the vendor has collected from multiple employers to an independent auditor who will conduct a bias audit

Imagine that the resume screening tool in this example has a bias audit in place (done less than one year prior) that relied on historical data from two employers:

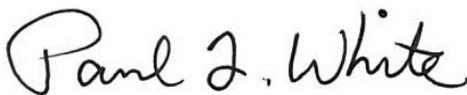
- + Employer A's data comes from the process that screens Registered Nurse resumes (a mostly female occupation)
- + Employer B's data comes from the process that screens air traffic controller resumes (a mostly male occupation)

Given the clear distinctions between the requirements for these two jobs, the AEDT will – or at least certainly *should* - have completely different algorithms in place, one to be used for one job and another for the other job. (In addition, of course, the labor market for these two occupations do not overlap at all.)

Now, imagine that Employer C wants to use the same AEDT to hire truck drivers. Because Employer C has not yet used the AEDT, it does not provide any of its own data to the independent auditor. Is Employer C even aware that the bias audit on which it is relying – and which it will post on its website as proof that the AEDT it is using is not biased – relies on RNs and air traffic controllers? Does such a bias audit – conducted using data wholly unrelated to data that is germane to the job being filled by Employer C – even tell us anything about whether the AEDT is or is not biased as regards selections for truck drivers?

We urge the Department to address these matters before seeking to enforce the new AEDT law.

Sincerely,



Paul White,
Partner
Resolution Economics, LLC

Cc: Kevin Bandoian, CPA, Resolution Economics, New York, NY
Tricia Etzold, CPA, Resolution Economics, New York, NY
Rick Holt, PhD, Resolution Economics, Washington, DC
Victoria A. Lipnic, JD, Resolution Economics, Washington, DC
Ali Saad, PhD, Resolution Economics, Los Angeles, C

- **Jiahao Chen, PhD**

Please see attachment for further comments.

[Comment attachment](#)

2023-01-23-NYC-AEDT.pdf

Comment added January 23, 2023 3:50pm



Jiahao Chen, PhD
Owner, [Responsible Artificial Intelligence LLC](#)
450 Lexington Ave, Unit 609
Manhattan, New York, 10163-9619
jiahao@responsibleai.tech

January 23, 2023

Department of Consumer and Worker Protection
New York City
42 Broadway
Manhattan, New York 10004-1617 Email:
rulecomments@dcwp.nyc.gov

Dear committee members:

I am pleased to submit comments for *Requirement for Use of Automated Employment Decision-making Tools* (Ref.

No. DCWP-21, dated December 15, 2022; “The Updated Rules”). Having previously submitted comments on October 23, 2022 (many of which are still relevant to the current version of the Rules), I would like to focus my comments here on just the new aspects of the Updated Rules.

1. Independent Auditor. The current definition of Independent Auditor could still use some clarification.

On one hand, the current definition of “financial interest” may be too restrictive - an auditor hired by a solutions vendor or employer and paid a fee for the audit could ostensibly be considered having a financial interest in the auditee, since the auditor has a vested interest in sustaining a business relationship in anticipation of future audit opportunities, and hence sustained revenue from audit fees. If such an interpretation were to be enforced so strictly, there would be no business opportunity for for-profit algorithmic auditors. Therefore, clarification would be useful on what constitutes “financial interest” when it comes to compensation for performing audit services, or for a sustaining agreement for such future services.

Conversely, an independent auditor meeting the current definitions in the Updated Rules may still be prevented from producing an objective and impartial audit. For example, a contract for auditing may restrict the tests the auditor can perform and report, and may require pre-clearance for publication to assess reputational risks or other disadvantageous disclosures, or even that the audit be conducted partially or fully under attorney-client privilege, which would undermine the goal of transparency in such audits. If auditees are allowed to exercise control over what is reported from an audit, they may cherry-pick the reported tests beyond the minimum requirements to show only those that portray the auditee positively, and suppress other test results that may be derogatory. There is therefore a need to clarify what restrictions may or may not be placed on auditors and their auditing contracts to prevent the gaming of the audit requirements, and the potential for ethics-washing by hiring auditors that are independent in name only.



2. Risks from reporting the statistics of small numbers. The example provided in the Updated Rules highlights clear privacy and statistical validity risks that could result from the audit, which I had described as a potential problem in my previous comments. The intersectional category of Female / Native American or Alaska Native shows that 7 of 17 candidates were selected, with a corresponding impact ratio of 0.789. A statistically naive auditor may be tempted to report this ratio as falling below the commonly-used threshold of 0.8 and hence be a potential concern for discriminatory risk. However, it could be argued that with just 17 people in this category, the computed ratio is not statistically robust.

There is also a privacy risk inherent in reporting results from small numbers. In the extreme limit, an audit may report that 0 out of 1 applicants in some demographic category were hired. This means that if we happen to know this applicant, we can immediately know from the audit that the application was not hired. In this way, the audit may represent an unintended way to leak information about an employment decision. Such issues have been well studied, particularly in the compilation of census data, where such concerns are of high priority. Current best practices for mitigating such privacy risks include differential privacy and omitting the reporting of very small categories.

In summary, tests of statistical validity, such as reporting the level of significance in hypothesis tests, should be included in the minimum requirements for reporting; otherwise, there should be provisions to omit reporting for very rare categories that lack sufficient representation to provide meaningful statistics.

3. Construct validity of the AEDT. As described in my previous comments, it is vital to test that systems to make employment decisions are made on reasonable criteria, as already required under U.S. Federal employment laws like the Equal Employment Opportunity Act. An employment examination that does not test skills necessary to perform a particular job would not be considered valid for use in hiring decisions for that job. Such requirements ought to extend to data-driven decisions too. Under current regulations for fair lending compliance, a bank would not be permitted to approve or deny credit based on data that have no bearing on creditworthiness, such as the color of the applicant's car or the applicant's taste in music. By analogy, a hypothetical employer may simply perform a coin toss to decide who to hire. Such a decision-making process is trivially unbiased since there is no risk of preferential discrimination on a prohibited basis, but is arguably unfair since no qualifying characteristics of the applicant are considered. Such risks are endemic to AI systems trained on bad data, even in the presence of cross-validation or other validation techniques. Therefore, all the discrimination testing in the world can still obscure the fact that an AEDT may simply be executing a very expensive coin flip, disguised by good performance on cherry-picked test data. Therefore, there is ample precedent for my recommendations, which are:

- A. to broaden the definition to AEDT to any data-driven decision-making system (whether or not it is executed on a computer or by hand; the medium of computation should be irrelevant); and



Jiahao Chen, PhD
Owner, [Responsible Artificial Intelligence LLC](#)
450 Lexington Ave, Unit 609
Manhattan, New York, 10163-9619
jiahao@responsibleai.tech

- B. to have algorithmic audits of AEDTs include tests for whether the data used for its decision making has sufficient predictive signal to enable useful decision making.

4.Data collection and retention requirements for discrimination testing. As described in my previous comments, collecting accurate demographic labels of applicants and employees will be vital for proper testing. However, the Updated Rules do not define minimum requirements for collecting and keeping such records on prohibited bases. Without such specifications, an employer has no incentive to improve their collection of demographic information, which may have downstream privacy and security risks on collected data, and obscuring the extent of discrimination on candidates and employees whose racial/sexual/etc. identities are unknown to the employer. Furthermore, if such records were not properly kept, it would be difficult to validate an audit if its results were significant enough to trigger further enforcement action from a regulator. Therefore, the Rules ought to specify:

- A. the minimum efforts employers are required to spend on collecting information on race, sex, age, etc.;
- B. acceptable usage of such records, such as restricting their use solely to auditing; and
- C. how records should be kept and retained for any further investigation.

Once again, I would like to congratulate the City for its progressive innovation toward eradicating employment discrimination in the age of data-driven decision-making, and encourage the City to consider these comments (in addition to those previously submitted) to maximize the intended effect of the Law toward building a more equitable employment market in the City. Please do not hesitate to reach out if I may be able to provide further clarifications on these comments.

Yours sincerely,

Jiahao Chen, Ph.D.

Owner

Responsible Artificial Intelligence LLC

- **Workday**

Please see our comments on DCWP's proposed amended rules attached.

[Comment attachment](#)

WDAY-Comments_NYC-LL-144-Rules_Second-Version_1.23.23.pdf

Comment added January 23, 2023 4:22pm



Public Comments on Proposed Rule Amendments to Implement New York City Local Law 144

January 23, 2023

Dear Commissioner Mayuga:

Thank you for the opportunity to comment on the New York City (NYC) Department of Consumer and Worker Protection's (DCWP or Department) proposed rule amendments implementing *Local Law 144 of 2021* (LL 144).

Workday is a leading provider of enterprise cloud applications for finance and human resources, helping customers adapt and thrive in a changing world. Workday applications for financial management, human resources, planning, spend management, and analytics have been adopted by thousands of organizations around the world and across industries—from medium-sized businesses to more than 50% of the Fortune 500. Workday advocates for thoughtful regulation to build trust in AI and has engaged with federal, state, and local governments, as well as governments abroad, on AI policy and regulatory best practices.

Workday supports LL 144's goal of addressing public concerns about unlawful discrimination in hiring and promotion. We believe that the successful implementation of LL 144 requires clear rules that offer deployers and developers of automated employment decision tools (AEDTs) a workable path to compliance.

We applaud DCWP's improvements to its proposed rules, including to its definition of an AEDT, which ensures that the LL 144 is targeted in scope. On the whole, however, the Department's amendments represent a step backwards from its original proposal due to its restrictive approach to independent evaluations, its continued misalignment with federal guidance on testing, and the privacy risks it poses to NYC's workers.

With the successful implementation of LL 144 in mind, we offer the following comments and recommendations.

Definition of an Automated Employment Decision Tool

When establishing a framework for regulating AEDTs, precise definitions are critical. We commend DCWP for refining the definition of AEDTs, as this will assist in bringing certainty to deployers and developers of AEDTs in New York.

Recommendation: DCWP should retain its amended definition of an AEDT.

Independent Auditors

LL 144 calls for an “impartial evaluation” of AEDTs by an “independent auditor.” Unfortunately, DCWP’s proposed amendments adopt a restrictive approach that bars employers from conducting internal independent evaluations.

We note that employers conducting internal audits have strong incentives to ensure that AEDTs are not used in a discriminatory manner, as these practices would result in significant legal, financial, and reputational repercussions. By contrast, third-party AI auditors do not have a respected independent professional body to establish baseline auditing criteria or police unethical practices among auditors. *Absent such professional standards, the advantages of the Department mandating audits be carried out by third parties is debatable and may be outweighed by the practical challenges and unintended consequences of a restrictive approach.*

We urge DCWP to recognize the immature state of the AI auditing field and retain the flexibility of its original proposed rules regarding independence.

Recommendation: We recommend the Department adopt the changes to its proposed amendments below.

“Independent Auditor. “Independent auditor” means 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.

An auditor is not an independent auditor of an AEDT if the auditor:

- i. is or was involved in using, developing, or distributing the AEDT;
- ii. ~~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~~
- iii. ~~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 continue to use the AEDT or in a vendor that developed or distributed the AEDT.”~~

Data Requirements

Workday urges DCWP to reconsider its proposed amendments in light of the following privacy and data protection considerations.

First, we caution DCWP against assuming that vendors have access to employers' historical data. Cloud software providers such as Workday are legally and contractually limited in how and when they can access, use, and disclose the data of enterprise customers. These safeguards are in place to protect the privacy and security of personal information and underpin both modern privacy laws and the enterprise cloud software market.

Second, NYC employers seeking to comply with DCWP's amended rules would have to give sensitive personal information of NYC employees and job candidates to third-party auditors. *We note, however, that DCWP's proposed amendments do not prevent thirdparty auditors from reusing New Yorkers' personal information for commercial purposes or from selling it to other third-parties.* Without such restrictions in place, DCWP's proposed amendments pose privacy risks to NYC's workers and job candidates.

Recommendation: The Department should convene stakeholders to discuss how unintended privacy risks posed by its proposed amendments may be avoided.

Bias Audit

DCWP's proposed rules for conducting a disparate impact analysis should align with federal guidelines.

Aligning with existing federal requirements provides two advantages. First, it allows employers to comply with a single, unified standard at the federal, state, and city level. Second, it does not commit the Department to an approach that may conflict with forthcoming federal guidance. We note that the U.S. Equal Employment Opportunity Commission (EEOC) is holding a hearing on January 31, 2023, to explore AEDTs in the context of federal law. Last year, the EEOC and the Department of Justice issued rules with respect to AEDTs and the Americans with Disability Act and is considering additional guidance.

Recommendation: The Department should align its amendments with federal guidelines, specifically, the Uniform Guidelines on Employee Selection Procedures, by removing intersectionality from its disparate impact testing requirement.

Public Disclosure Requirements

DCWP's proposed amendments require employers to publish a summary of the bias audit of an AEDT, including the selection rates and impact ratios for all categories. Releasing such raw data without context would create a situation ripe for misinterpretation. The risk of misinterpretation may drive employers to seek testing models that produce less candid results, undermining the aims of LL 144. Publishing an independent auditor's summary of the impact of the AEDT accomplishes the same objective.

Recommendation: The Department should revise the requirement to publish the selection rates and impact ratios for all categories and replace it with a requirement to publish a summary statement on any adverse impact identified by the audit.

We appreciate your consideration of our comments. Please contact Michelle Richardson, Senior Director, U.S. Public Policy, at michelle.richardson@workday.com, if Workday can provide further information as the Department finalizes these regulations.

- **Hannah Wade**

NYU Langone Health Comments on Proposed Rule Amendment,
Subchapter T: Automated Employment Decision Tools

RULE TITLE: Requirement for Use of Automated Employment
Decisionmaking Tools

REFERENCE NUMBER: 2022 RG 061

RULEMAKING AGENCY: Department of Consumer and Worker
Protection

On behalf of NYU Langone Health, please accept our comments on the proposed rules to implement Local Law 144 of 2021 related to the use of automated employment decision tools, or AEDTs. We appreciate the Department of Consumer and Worker Protection (DCWP) for the opportunity to comment again.

As the healthcare system strives to recover from the COVID-19 pandemic, New York City hospitals are facing significant workforce challenges. The DCWP should consider providing an exemption to the healthcare field due to ongoing public health crises. These crises, including recovery from the COVID-19 epidemic, the monkeypox outbreak, and the recent influx of asylum seekers, have put significant stress on all New York City hospitals. We are deeply troubled by any additional measures that prevent us from fulfilling our mission to provide safe, quality care for our patients.

At NYU Langone Health, we are opposed to any additional barriers to fill urgently needed positions including nursing, allied health, clinical support and other support services. In particular, we have concerns about the potential hiring delays presented by the requirement (Section 5-304) to provide notice to candidates and employees 10 business days prior to the use of an automated employment decision tool, or AEDT. This requirement presents an unnecessary waiting period

that will prolong staffing shortages and negatively impact patients in New York City.

During our Fiscal Year 2022, we received 368,536 applications for 12,796 posted positions which require the use of data analytics to effectively process. Delays of 10 business days in processing time would pose an undue hardship on our healthcare system as we work to recruit and employ talent to best serve our patients.

Once again, thank you for the opportunity to comment. Please reach out to us with any questions or for additional information.

Comment added January 23, 2023 5:00pm

- **Gibson, Dunn & Crutcher, LLP**
Please consider the attached comments.

[Comment attachment](#)

Comments-Regarding-Proposed-Rules-for-Implementing-Local-Law-144-of-2021.pdf