Consumer and Worker Protection

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

Dear Sir / Madam –

On October 24, 2022, the Department of Consumer and Worker Protection of the City of New York will hold a public hearing to discuss and hear comment on the aforementioned ordinance. **We request that this ordinance be written to more clearly articulate its true spirit, namely, to reduce or eliminate the use of artificial intelligence and machine learning in selection processes.** The following paragraphs share our rationale for this request.

Criteria Corp. is a leader in pre-employment assessment (e.g., online tests, structured interviewing), with more than 4,500 clients across more than 60 countries. Since our inception in 2006, we have delivered more than 40 million assessments to help thousands of organizations make better hiring decisions, enhancing diversity among their workforce, and improving bottom line results. Criteria was founded on the belief that all job candidates should have an equal opportunity to demonstrate their capabilities in a fair and unbiased manner, regardless of their background or socioeconomic status. We use assessments to achieve this goal.

Criteria does NOT use artificial intelligence or machine learning in our assessments. All of Criteria's assessments are based in science, have a logical rationale for their use in selection for a given role, and are extensively proven to predict job performance. For example, we know that there are more than 22,000 studies across more than 5 million participants, demonstrating that smarter employees tend to perform their jobs better, so Criteria has assessments to measure candidate intelligence. There are countless studies demonstrating that employees who work harder tend to perform their jobs better, so Criteria has assessments that measure candidates' propensity to try their best and work hard on a given job. There are many studies demonstrating that specific characteristics, such as sales ability, customer service orientation, and empathy, contribute to success in sales, service, and leadership jobs, respectively, so Criteria has assessments that measure these characteristics.

In addition to the published literature, Criteria has conducted hundreds of studies across our client base demonstrating that our assessments predict job performance, and do so in a fair and unbiased manner. In essence, our assessments empower job candidates to demonstrate their abilities in a manner that transcends their backgrounds and socioeconomic status, as well as their gender, ethnicity, age, nation of origin, and other characteristics that are irrelevant to job performance. The use of our assessments creates a fairer selection process by removing bias.

Given that Criteria does NOT use artificial intelligence or machine learning in our assessments, we understand that we are not impacted by this upcoming New York City law. However, the law as it is currently written is very brief, and subsequently very broad, such that one could misconstrue our assessments to be impacted.

For example, given the brevity and subsequent breadth of the current law, if a company in New York City wanted to hire an Accountant, and they used a math test as part of the hiring process, they could potentially be fined based on this law. If a company in NYC wanted to hire an HR Manager who is tolerant and empathetic, and they assessed job candidates for that position with an online test that measured tolerance and empathy, they could potentially be fined under this law. If a company in NYC wanted to hire a leader who is strong in integrity, they could potentially be fined if they used an assessment to hire that person.

In each of these examples, there is an a-priori rationale for the use of the test content (math for accounting; tolerance and empathy for HR; integrity for leadership). However, given that all online tests measuring these characteristics use an algorithm (even an algorithm as simple as just adding the total number of math questions answered correctly), they could be construed by some as falling under this law.

We therefore ask that the law be written to more clearly articulate its true spirit, namely, to reduce or eliminate the use of artificial intelligence and machine learning in selection processes. Note that the key difference between a test using artificial intelligence and machine learning, versus a test that does not is based in whether the test uses an a-priori rationale for what it is measuring. If the test has an a-priori rationale for what it is measuring, then it is not using artificial intelligence and machine learning. We therefore ask that any assessment using an a-priori rationale be excluded from this law.

Please also note that, given its efficiency, objectivity, and accuracy, testing cannot simply go away. There are far too many job applicants for most companies to interview all of them, so testing needs to exist as an efficient, accurate, and unbiased means for identifying candidates who are most likely to succeed on the job. The law as currently written could possibly, and unintentionally, include virtually any test. It would be far better to target the subset of tests that lack an a-priori rationale (e.g., artificial intelligence and machine learning based assessments only) with this law.

Thank you for taking the time to read our request. We look forward to the upcoming public hearing.

Respectfully yours,
Brad Schneider, Ph.D.

Brad Schneider, Ph.D.
Vice President, Strategic Consulting
mobile +1 (747) 800-1657

The following comments pertain to the proposed automated employment decision tools rules.

1. 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?
2. The proposed bias audit is based solely on the generation of impact ratios, and thus equates bias with score or categorization differences by group. As discussed in Chapter 3 of the 2014 Standards for Educational and Psychological Testing, measurement bias, and more broadly fairness, are complex concepts. These concepts cannot be boiled down to a simple impact ratio. Further, the validity of the scores or categorizations must be considered to evaluate bias. The proposed bias audit process is therefore likely to yield oversimplified, inaccurate and misleading results.
3. How is it envisioned that the bias audit would be handled when a new automated decision tool is initially implemented? If it is a new tool, obviously the user would have no data to compute impact ratios. Would the expectation be that the bias audit as defined would be performed after 1 year of use?

Joseph D. Abraham, Ph.D.
Vice President, Assessment Solutions

TALOGY
PREVIOUSLY PSI TALENT MANAGEMENT

O: 01 818 847 6180, Ext. 7306 | **M:** 01 918 404 0050
Time zone: Central Time

Talogy.com | LinkedIn

October 24, 2022

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 Clarification 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 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 January 1, 2023.

I. Overview of Proposed Rules for Local Law 144

The Institute applauds the DCWP's efforts to clarify Local Law 144. Enhanced clarity established by clearly defining terms and requirements allow employers to ensure they are compliant with this law when it goes into effect. Our members are encouraged by various clarifications in the proposed rules so far.

However, there are some areas where the proposed rules fall short of providing clarifying guidance and/or deviate from standards and accepted practices related to evaluating disparate impact. Additionally, there are areas in which inconsistencies raise questions that require further clarification to allow employers and auditors to understand how to properly comply with Local Law 144. The remainder of these comments will focus on these areas in which we encourage the DCWP to consider further revising and clarifying the provided definitions and requirements.

II. Who is eligible to conduct a bias audit?

In §5-300, the DCWP defines an independent auditor as: "… a person or group that is not involved in using or developing an AEDT that is responsible for conducting a bias audit of such AEDT." This definition suggests that neither an employer using an AEDT nor a vendor developing an AEDT will suffice as an independent auditor. However, in the first example of § 5-301, the statement is made that:

> The employer asks the vendor for a bias audit. The vendor uses historical data it has collected from employers on applicants selected for each category to conduct a bias audit…

Understanding that this example is describing a situation in which the employer is considering the use of an AEDT (as opposed to an employer having already used an AEDT) such that no employer-specific data are available and historical vendor data may be required, the phrasing suggests that a vendor is conducting the bias audit of their own AEDT for an employer.

It is important to distinguish between who is conducting the bias audit and who is providing the data upon which the bias audit is based. There will most certainly be use cases, such as the one outlined in the DCWP's provided example, wherein an AEDT has not yet been deployed for an employer, and a vendor may need to provide relevant data to an independent auditor to be used in conducting the bias audit analyses. However, a vendor providing data for a bias audit and a vendor being eligible to serve as the independent auditor of their own AEDT and conducting the bias audit are quite different. As such, we request the DCWP clarify this issue in the proposed rules to reflect their intent. If the definition of independent auditor precludes vendors from performing bias audits, then the example should clarify that the vendor is providing data for the bias audit to an independent auditor, but not performing audits of their own tools.

III. What data should a bias audit be based on for different use cases?

The proposed rules do not make it clear whether employers are required to use their specific data or whether cross-client, aggregate data from a vendor might suffice. We recognize that this answer may depend on the specific circumstances. In some situations, an employer will have their specific data available to use (*e.g.*, post deployment of an AEDT, pilot studies with an AEDT), while other situations may occur where no employer-specific data are available (*e.g.*, prior to deployment of an AEDT). In situations where employer-specific data are not available, there will be some circumstances in which historical, archival data from other sources may be available, and some in which no historical, archival data are available (*e.g.*, custom-built AEDT pre-deployment) to estimate the potential bias of an AEDT.

We request the DCWP provide clarification for three typical use cases regarding the data that should be used for the purpose of conducting bias audit analyses. Specifically, we recommend that the DCWP address the following use cases and provide guidance in the proposed rules.

1. Implementing an AEDT where sufficient employer data are available: This scenario aligns well to the Uniform Guidelines and should most closely mirror the reality of operational use of the AEDT. When should employer data be analyzed for the bias audit?

2. Implementing an AEDT where sufficient employer data are not available, but sufficient cross-employer vendor data are available: This is similar to a longstanding practice by vendors to use aggregated data to estimate how an assessment would be expected to function, absent employer-specific data. When should cross-client, aggregate vendor data be analyzed for the bias audit?

3. Implementing an AEDT where no existing data are available: How should employers proceed when seeking to deploy a custom-built AEDT, for which no prior data exists? Are there options available for auditors to evaluate bias via a proxy method (*e.g.*, another similar type of AEDT, research literature related to a type of AEDT), or will the City grant a grace period for this particular use case?

Providing clarification on data eligible for analysis in these situations would be beneficial to employers and auditors to ensure that the bias analyses conducted are compliant with Local Law 144.

Additionally, DCWP should seek feedback as to whether the specific language used in the examples in the proposed rules are consistent with the remainder of the information contained in the proposed rules and with the intention of the City.

IV. Are there parameters that define data to be used in bias audit analyses?

After determining the source of data that are eligible for use in different situations, it is important to consider relevant parameters to conducting the bias audit. To build consistency in what employers are expecting, what auditors are delivering, and what the City is evaluating for compliance, the City would be well served by directly addressing the following parameters:

1. Candidate location: Should bias audit analyses restrict the candidate pool to include only NYC residents, or may a larger candidate pool be considered (*e.g.*, state, region, national)?

2. Relevance of data when employer-specific data cannot be used: In some situations, AEDTs may be used across a very wide range of jobs (*e.g.*, hourly workers, sales professionals, managers, executives) or in very different settings (*e.g.*, finance, manufacturing, customer service). Should bias audit analyses be restricted those jobs and settings that more closely represent the jobs for which the employer intends to use the AEDT?

V. The definition of impact ratio

We applaud the DCWP for providing more clarity around the bias audit analyses that are expected. The impact ratio is defined in § 5-300 as:

> [E]ither (1) the selection rate for a category divided by the selection rate of the most selected category or (2) the average score of all individuals in a category divided by the average score of individuals in the highest scoring category.

Impact Ratio = Selection rate for a category
~~Selection rate of the most selected category~~

OR

Impact Ratio = Average score of individuals in a category
~~Average score of individuals in the highest scoring categ~~ory

The first equation is one of the well-recognized, practical significance metrics for evaluating selection rate differences between two different demographic groups. This metric is consistent with the "four-fifths rule" as described in the Uniform Guidelines on Employee Selection Procedures[1] ("Uniform Guidelines" Section 4D) and their Interpretation and Clarification (Questions and Answers) document[2] ("Q&A" Question 12).

The second equation, however, is not a metric that is used in practice based on our experience. It contains a fundamental flaw that does not allow for meaningful conclusions about the difference between scores on an AEDT for two different demographic groups because it does not account for variability in scores. This metric, as defined, may lead to purposefully or

inadvertently illusory, misleading results. In simple terms, if a child comes home from school and tells you that they scored an 85 on a test at school, your first question might be, "Is that good?" From the information provided, you cannot tell—you lack relevant context. To know if this is a good score on the test, you need to know how others did perform (or typically perform) on the test—you need to understand the distribution (variability) of scores on the test.

That said, we very much appreciate that the DCWP has recognized the need for a bias indicator to handle the situation wherein an AEDT is not being used with a strict threshold (*e.g.*, cut-score, band, categorization) that allows for computing a selection rate, but instead scores are generated and are being directly used as a factor in determining which candidates proceed in a hiring process. This is a very common use case, and we agree that it would be best if a consistent method were used across auditors.

To that end, we submit that there are multiple, professionally-accepted, practical significance measures that may be used for effectively evaluating the difference between two "average scores" on an AEDT, and that the DCWP consider adopting one of these. For example, one such statistic that is commonly used is the d-statistic.[3][4][5] In simple terms, this metric allows one to describe the difference in magnitude of two average scores, regardless of the distribution of scores, in a consistent, meaningful, and interpretable way.

[1] *Uniform guidelines on employee selection procedures*, Equal Employment Opportunity Commission, Office of Personnel Management, Department of Justice, Department of Labor, and Department of Treasury, 43 *Fed. Reg*. 38290-38315 (August 25, 1978). *See* 29 CFR Part 1607 (1978) available at https://www.ecfr.gov/current/title-29/subtitle-B/chapter-XIV/part-1607 and 42 CFR Part 60-3 (1978) available at https://www.ecfr.gov/current/title-41/subtitle-B/chapter-60/part-60-3?toc=1.

[2] *Adoption of Questions and Answers to Clarify and Provide a Common Interpretation of the Uniform Guidelines on Employee Selection Procedures*; Equal Employment Opportunity Commission, Office of Personnel Management, Department of Justice, Department of Labor, and Department of Treasury (1979). 44 *Fed.Reg*. 11996-12009 (March 2, 1979); available at https://www.eeoc.gov/laws/guidance/questions-and-answers-clarify-and-provide-common-interpretation-uniform-guidelines.

[3] Cohen J. (1988). Statistical Power Analysis for the Behavioral Sciences. New York, NY: Routledge Academic.

[4] The d-statistic is sometimes referred to as an effect size difference or a standardized mean difference.

[5] This represents a pooled standard deviation which is commonly used for this statistic. This provides a more accurate representation of the overall standard deviation for the pairwise comparison of group 1 and group 2 averages.

To compare this to the second impact ratio metric described in the proposed rules, consider an example of three different AEDTs where the average score of two groups is being compared. Then consider how the interpretation of results is much different once you have information about the distribution of scores (*i.e.*, the variance or standard deviation). The second impact ratio defined in the proposed regulations is represented by the "Averages Ratio" column, the variability of scores in the distribution is represented by the "Standard Deviation" column, and the result of calculating the standardized difference in average scores is represented by the "d-statistic" column.

Scenario	Average Group 1	Average Group 2	Averages Ratio	Standard Deviation[5]	d-statistic
AEDT 1	50	44	.88	24.0	0.25
AEDT 2	50	44	.88	12.0	0.50
AEDT 3	50	44	.88	6.0	1.00

In all three cases, the Averages Ratio appears to be the same—one could draw the same conclusion about the difference in demographic group scores. However, when considering the distribution of scores with a statistic such as the d-statistic, it is clear that the same conclusion should not be drawn. With AEDT 1, the difference is relatively small because there is a large

amount of variation in scores; with AEDT3, the difference is large as there is little variation in scores. The interpretation of bias for these two AEDTs would be quite different.
It should be clear from this example that the currently-defined second impact ratio is not very meaningful and will be unlikely to yield interpretable results. Given this, we request that the DCWP define a practical significance metric commonly used or professionally-accepted that will be more appropriate for use in comparing the average score of two groups, of which the dstatistic is but one example.

VI. Parameters to clarify related to calculations of an impact ratio

After defining the impact ratio, there are practical parameters to a bias audit analysis that are important to consider. These parameters are related to the stability of results and the confidence one might have in drawing conclusions based on the analyses conducted. We recommend the DCWP consider addressing some key parameters in the proposed rules or indicate that there is allowance for reasonable professional judgment by the independent auditor.

Note that some of these parameters may be informed by the Uniform Guidelines and their Q&A document since the proposed rules for Local Law 144 state that:

> These calculations are consistent with Section 1607.4 of the EEOC Uniform Guidelines on Employee Selection Procedures. *See* 29 CFR §1607.4.

1. <u>Sample size required for a pairwise comparison</u>: Small sample sizes are especially problematic for impact ratio analyses. Evidence based on small numbers is considered in the Uniform Guidelines (4D) and their Q&A document (e.g., Question 20, Question 21). For example, in Question 20 they provide an example whereby an impact ratio based on a pool of 20 males and 10 females, where 3 males and 1 female are selected, yielded a conclusion that "the number of selections is too small to warrant a

determination of adverse impact." In Question 21, they provide an example in which impact ratio results would change appreciably with the shift of a single individual, noting that this indicated results that were unstable. Given the nature of impact ratios, this instability at small sample sizes is to be expected—the influence of a decision made about a single individual may unduly effect the statistical conclusions.

We request the DCWP directly acknowledge the relevance of the federal guidance, directly specify guidance on minimum sample sizes to warrant analyses, or state that this is subject to independent auditor discretion.

2. Sample size required for a demographic group to be evaluated: We would appreciate the DCWP confirming that the guidance set forth by the Uniform Guidelines Section 15A-2[6] and Question 16 of their Q&A document, indicating that groups which

[6] 29 CFR §1615(a)(2); 42 CFR § 60-3.15(A)(2).

comprise less than 2% of the pool should be excluded from analysis, should be applied when conducting Local Law 144 bias audits.

3. Interpretive guidelines for impact ratios: Interpretation of an impact ratio has become more nuanced over the years. Impact ratio results can be interpreted in a binary framework based on the four-fifths rule, or with a multi-categorical interpretation (*e.g.*, OFCCP Q&A on Practical Significance, January 15, 2021[6]), or with more of a continuous consideration taking the statistic as a way to gauge the potential for a tool to yield bias. Will the DCWP provide interpretation guidance, or will this be left to the discretion of the independent auditors?

4. Which impact ratio to calculate: §5-301-b indicates that the averages ratio must be calculated when candidates are scored. However, in the vast majority of cases, any candidate that has been assigned to a classification or band has been given a score upon which the assignment is based. As such, the proposed rules read that a classification or banding approach will require both be calculated.

   Given the definition of impact ratio in §5-300, which clearly indicates (via the "OR") that the intent is to calculate one metric (not both), it would behoove all for the DCWP to align language in the proposed rules to reflect that only one metric should be calculated and that this should be the one that most closely mirrors the operational implementation of the AEDT.

VII. Which pairwise comparisons are required for analysis?

The tables in the two examples in the proposed rules imply intersectional analyses need to be conducted and reported (*i.e.*, the combination of race/ethnicity and sex). While intersectional analyses can be conducted, in most cases the primary interest is in evaluating race/ethnicity and sex separately, with intersectional analyses considered as supplemental at most.

This is another situation in which it would behoove all involved for some clarification on the pairwise comparisons that are intended to be evaluated in a bias audit for Local Law 144. Should auditors analyze race/ethnicity and sex separately only, analyze intersectional combinations only, or do both (yielding three tables per analysis strata: 1 race/ethnicity, 1 sex, 1 intersectional).

[6] https://www.dol.gov/agencies/ofccp/faqs/practical-significance

VIII. Effective date for Local Law 144

NYC's Local Law 144 is slated to go live on January 1, 2023. However, there is still significant confusion related to the specifics for employers coming into compliance with this law. Further, as noted here, there are some additional clarifications that are necessary and will impact the ability for employers and auditors to comply with what is required. Finally, with the rescheduling of the October 24 public hearing to a later date, both the feedback to the DCWP, as well as any revisions to the proposed rules based on this feedback, will be further delayed. We respectfully request that the City immediately consider a six month delay to the effective date. In the alternative, the City could proceed with a soft launch in which there is a grace period for employers to comply, so long as an independent auditor is identified and the work is scheduled by January 1, with a requirement to have the work completed before the year concludes.

Conclusions

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,

Barbara L. Kelly

Barbara L. Kelly
The Institute for Workplace Equality Director

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

October 24, 2022

BSA | The Software Alliance appreciates the opportunity to submit comments on the proposed regulations implementing Local Law 144 of 2021—New York City's ordinance on automated employment decision tools ("Ordinance"). BSA is the leading advocate for the global software industry before governments and in the international marketplace. Our members are enterprise software companies that create the technology products and services that power other businesses.[7] They 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.

BSA members are on the leading edge of providing businesses — in every sector of the economy — with innovative software services, including Artificial Intelligence (AI). 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.

While the adoption of AI can unquestionably be a force for good, we also recognize the significant risks to society if this technology is not developed and deployed responsibly. BSA agrees that when AI is used in ways that could unlawfully discriminate, the public should be assured that such systems have been thoroughly vetted to identify and mitigate risks associated with unintended bias.

A. BSA's Framework to Build Trust in AI

This has been an area of particular focus for BSA and our member companies over the last several years. In June 2021, BSA released Confronting Bias: BSA's Framework to Build Trust in AI, which sets forth a risk management approach for confronting concerns about bias.[2] The Framework was informed by leading AI developers and outlines a lifecycle-based approach for performing impact assessments to identify risks of AI bias and corresponding best practices for mitigating those risks. AI is used in so many different contexts that only a flexible, risk management approach will be successful. There are a variety of AI development and deployment models, and the Framework recognizes that the appropriate allocation of risk management responsibilities will

[7] BSA's members include: Adobe, Alteryx, Atlassian, Autodesk, Bentley Systems, Box, Cisco, CNC/Mastercam, CrowdStrike, Databricks, DocuSign, Dropbox, Graphisoft, IBM, Informatica, Intel, 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.

[2] Confronting Bias: BSA's Framework to Build Trust in AI, *available at* https://ai.bsa.org/confrontingbias-bsas-framework-to-build-trust-in-ai.

vary depending on the type of system, including who develops the algorithm, trains the model, and ultimately deploys the system.

BSA supports the overarching goal of the proposed regulations, which is to ensure high-risk uses of AI are subject to safeguards. One crucial safeguard that promotes responsible uses of AI systems is ensuring that companies that develop or use high-risk AI systems establish a comprehensive approach for performing impact assessments. Impact assessments are widely used in a range of other fields—from environmental protection to data protection—as an accountability mechanism that promotes trust by demonstrating that a system has been designed in a manner that accounts for the potential risks it may pose to the public. The purpose of an impact assessment is to establish organizational processes to guide the development and use of high-risk systems by requiring internal stakeholders to identify the risks that a system may pose, quantify the degree of harm the system could generate, and document any steps that have been taken to mitigate those risks to an acceptable level. By establishing a process for personnel to document key design choices and their underlying rationale, impact assessments are an important transparency and accountability mechanism.

B. The Department of Consumer and Worker Protection's Proposed Regulations

The proposed regulations address the need for safeguards that promote responsible uses of AI by focusing on how companies are to implement bias audits under the Ordinance. As the proposed regulations recognize, one critical aspect of operationalizing such audits is identifying a set of individuals who can conduct them. Importantly, the proposed regulations appropriately recognize that internal stakeholders who are not involved in the development or use of an automated employment decision tool are competent to conduct a bias audit. Acknowledging the independence of internal stakeholders will incentivize companies to implement multiple layers of independent review, which will enhance trust in the use of these systems and create safeguards that function in practice.
We have recommendations, however, about three aspects of the proposed regulations:

First, we suggest deleting "or modify" in the definition of "automated employment decision tool" in Section 5-300. The definition refers to relying "solely" on simplified outputs "with no other factors considered" and use of simplified outputs to "overrule" conclusions derived from other factors including human decision-making. As illustrated by these descriptions, the intent of this definition appears to be identifying circumstances that could "substantially assist or replace discretionary decision making." This is a relatively high threshold. Outputs that do not rise to this level, such as merely contributing to or modifying, should be excluded from the scope of this defined term.

Second, we suggest adding a notation explaining that the examples in Section 5-301 for the bias audit are not exhaustive of all scenarios. Along the same lines, we recommend revising the two examples in this section to recognize a broader range of possible scenarios regarding the appropriate entity to conduct a bias audit. There are myriad development and deployment models, and multiple stakeholders may play varying roles depending on the circumstances. Specifically, we recommend offering an illustrative example without inadvertently suggesting that the types of entities described in the example will always have the same roles.
Further, in the examples for both subsections (a) and (b) of this section, we recommend replacing "historical data" with "test data" because there may not always be historical data if the organization has not previously used the technology. In addition, historical data may not be available in many scenarios—such as when that data is subject to privacy or security safeguards that prevent its dissemination. In addition, we recommend deleting "planned" from the example in subsection (b), as it inserts ambiguity regarding the timing of the use of the automated employment decision tool.

Third, we recommend omitting the requirement to publish the selection rates and impact ratios for all categories and instead require a summary statement on adverse impact. Although Section 20-871(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 contemplated 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 striking “the selection rates and impact ratios for all categories” in Section 5302 and replacing it with “a statement on adverse impact.” We further recommend aligning the categories for the selection rates and impact ratios with the EEOC’s approach for disparate impact testing.

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

October 24, 2022

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

Re: 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 proposed rules relating to Local Law 144 of 2021 ("LL 144"), the ordinance that the New York City Council passed last fall relating to employers' notice and audit requirements for the use of Automated Employment Decision Tools (AEDTs). CDT is a nonprofit, nonpartisan 501(c)(3) organization based in D.C. that advocates for stronger civil rights protections in the digital age, including algorithmic tools that can affect workers' access to employment.

CDT supports laws and regulations that ensure that AEDTs are nondiscriminatory and job-related, include proper candidate notice, and are subjected to adequate auditing, oversight, and accountability. CDT appreciates the Council's efforts to establish some level of accountability for AEDTs through the passage of LL 144, and likewise appreciates the Department's efforts to bring clarity to that ordinance through its proposed rules. Nevertheless, as CDT explained after the ordinance passed, we do not believe the ordinance goes far enough in ensuring that AEDTs are tested for potential discriminatory effects, that candidates receive proper notice of the manner in which AEDTs assess them, or that adequate enforcement and remedies are available.[8]

Unfortunately, several of the proposed rules would further diminish the ordinance's effectiveness by narrowing its scope and application in ways that are not supported by the ordinance's text. Our comments discuss the ways in which we believe the draft rules can be strengthened to ensure that the City's workers are protected to the maximum degree possible under the ordinance.

Proposed interpretation of "Automated Employment Decision Tool" is narrower than the statutory definition

The statutory definition of "automated employment decision tool" covers all automated processes that "substantially assist or replace discretionary decision making" in employment decisions.[9] The proposed rules interpret this term as requiring that the output of an AEDT be the "sole[]" factor in a decision, be "one of a set of criteria where the output is *weighted more than any other criterion in the set*," or "overrule or modify conclusions derived from other factors." This interpretation is unjustifiably narrower than the statutory definition, which encompasses any system that "substantially assist[s]" human decision-making. The term "substantial" means "of ample or considerable amount, quantity, size, etc"[10] as well as items that are "important" or "essential."[4] Therefore, the plain meaning of the statutory definition would designate tools used in a broad range of contexts as AEDTs.

By narrowing the definition to tools that are given sole or primary weight, the proposed rules create a loophole companies could exploit. The rules do not require employers to provide meaningful transparency about how the outputs of AEDTs are weighted and about other factors considered to assess the same candidate. For an AEDT whose output is "one of a set of criteria," an employer could easily design a selection process that circumvents the

[8] Matthew Scherer and Ridhi Shetty, *NY City Council Rams Through Once-Promising but Deeply Flawed Bill on AI Hiring Tools*, Center for Democracy & Tech. (Nov. 12, 2021), https://cdt.org/insights/ny-city-council-rams-through-once-promising-but-deeply-flawed-bill-on-ai-hiring-tools/.
[9] N.Y.C. Admin. Code §20-870.
[10] Dictionary.com, *substantial*, based on the Random House Unabridged Dictionary, https://www.dictionary.com/browse/substantial (accessed October 11, 2022). [4] *Id.*

ordinance or argue that the AEDT does not meet the rules' narrow definition of "substantially assist." For instance:

- An employer could rely on multiple distinct AEDTs and give all of their outputs the same weight; under the rules' definition, none would then qualify as AEDTs.
- An employer could rely on multiple AEDTs whose outputs are weighted differently, so only the one given the most weight would be subject to the ordinance's requirements.
- An employer could give an AEDT's output the same weight as the outputs of non-automated tools or methods used in the overall process, and such an AEDT would fall outside the audit requirements under the rules.

The proposed rules could also lead to inconsistent results--an AEDT that is given 40% weight in a hiring decision might not qualify (if another factor is weighted the same or slightly more), while one that is only given 15% weight might qualify (if there are numerous other factors with lesser weight also considered).
We urge the Department to either leave the term undefined, as the word "substantial" is commonly understood and can be given its plain and ordinary meaning, or else to define it as including tools whose output is "an important or significant factor in an employment decision."

"Candidate for employment" should include all workers subjected to an AEDT
The ordinance defines "employment decision" to mean "to screen candidates for employment or employees for promotion in the city." This definition is already too narrow as it only covers decisions to hire or promote.[11] This excludes decisions regarding "terms, conditions, and privileges of employment" that were previously covered in the legislative text when it was originally introduced.[12]

The proposed rules would unduly restrict this definition even further by limiting what it means to "screen" a "candidate for employment." The proposed rules specify that to "screen" is to determine "whether someone should be selected or advanced in the hiring or promotion process" and that a candidate is only someone "who has applied for a specific employment position." Those definitions are inconsistent with the text of the ordinance; while the ordinance's *notice* requirement only applies to "a candidate *who has applied for a position* for an employment decision," the bias audit requirement applies more broadly to tools used by employers to "screen a candidate for an employment decision." The rules' definition thus inserts language that is not consistent with the term's use in the text of the ordinance.
By narrowing the ordinance's scope to the screening of candidates who actively apply for a job, the proposed rules exclude several categories of automated tools, such as those that target job advertisements or search resume databases or social media to recruit potential candidates. This would exclude a significant number of workers from the ordinance's protection because such tools increasingly determine the candidate pool for jobs in the city.[13]
The definition of "screen" is also seemingly narrower than the definition of discriminatory practices under the current City Administrative Code, which states that it is unlawful to "represent that any employment or[14]position is not available when in fact it is available." N.Y.C. Admin. Code § 8-107(a)(1).
The rules instead should align more closely with the City's civil rights laws by clarifying that to "screen" means to "make a determination about whether someone should be selected or advanced in the hiring or promotion process, or whether to represent that any employment or job position is available, on the basis of characteristics protected under N.Y.C. Admin. Code § 8-107." The rules should define "candidate for employment" as "a person

[11] N.Y.C. Admin. Code §20-870.
[12] Int. No. 1894.
[13] Note that even among workers who do submit an application, the proposed rules' requirement that a worker submit information "in the format required" by the employer could give employers an opportunity to exclude workers from the bill's notice requirements (or from the adverse impact calculations in a bias audit) on a technicality.
[14] RCNY §5-249 (emphasis added).

who applies for a position or who is subjected to an employment decision." This will ensure that all people screened or assessed by an AEDT--and thus all people whose employment prospects the AEDT affects--are covered by the ordinance's bias audit requirements.

Definition of "employment agency" relies on narrow language and should be clarified

The proposed rules use the definition of "employment agency" given in § 5-249 of the Department's regulations, which only says that an "employment agency" includes "all persons who, for a fee, render vocational guidance or counseling services, *and* who directly or indirectly represent" that they do one of the enumerated functions that can help people apply for or secure jobs.[8] The same provision states that "vocational guidance or counseling services" means "services which consist of one or more oral presentations" *and* perform other enumerated functions for people seeking employment.[15] Since the proposed rules only apply to people who have applied for specific job positions, the proposed rules' reference to § 5-249 is only relevant to entities that provide oral presentations to jobseekers *and* that represent that they can secure job positions for the jobseekers who have already applied to those positions.

Given the increasing prevalence of *employer* service providers in the labor market, the definition of "employment agency" should be broad enough to subject any person using AEDTs to try to procure employees for employers or job opportunities for workers to the bill's notice and bias audit requirements. New York State's full definition of "employment agency" captures employer service providers because it applies to "any person who, for a fee, procures or attempts to procure (i) employment or engagements for someone seeking employment or engagements, or (ii) *employees for employers seeking the services of employees*."[16] Similarly, NYC's anti-discrimination laws define the term as "any person undertaking to procure employees or opportunities to work."[17] Both of these definitions would apply to entities that the proposed rules aim to hold accountable. Therefore, the final rules should more directly refer to the broader definition in the State's law and the City's anti-discrimination laws.

The requirements for "independent auditors" should be strengthened to ensure truly independent audits

The proposed definition of "independent auditor" would ensure some semblance of independence by requiring an auditor be "a person or group that is not involved in using or developing" the AEDT for which they perform a bias audit. This does not go far enough, however, in ensuring true independence. An independent auditor should be required to be free of financial or personal conflicts of interest. In addition, the auditor should have a duty to accurately report the results of all tests and analyses conducted during the course of the audit; otherwise, a vendor may be able to pressure an auditor to withhold or modify audit results that are unfavorable to the vendor.

Consequently, we recommend that "independent auditor" be defined as "a person or group that is responsible for conducting a bias audit of an AEDT and that:

- Is not involved in using or developing such AEDT;
- Certifies in the audit report that all results from the audit have been reported and are complete and accurate; and

[15] *Id.*
[16] N.Y. Gen. Bus. Law § 171 (emphasis added).
[17] N.Y.C. Admin. Code § 8-102.

- Has no financial, personal, or other conflict of interest that would compromise or call into question the auditor's duty or ability to design and conduct the audit to measure impact effectively and to completely and accurately report the audit results."

Notice should be communicated directly to candidates

The ordinance requires an employer to give candidates notice that an AEDT will be used to assess them "no less than ten business days before such use and allow a candidate to request an alternative selection process or accommodation."[18] The rules permit an employer to satisfy this notice requirement simply by posting such a notice "on the careers or jobs section of its website in a clear and conspicuous manner." But the mere posting of such information does not satisfy the notice requirement contemplated by the ordinance.

As the ordinance text indicates, a key purpose of this requirement is to allow candidates who may need to request an alternative assessment method or accommodation an opportunity to do so.[19] The information required by § 20-871(b)(2) regarding the "job qualifications and characteristics" the AEDT will measure is, relatedly, essential for a candidate to determine whether they may need to make such a request. However, these requirements to notify candidates about the use of an AEDT and qualifications and characteristics it will measure are not enough to ensure candidates have the opportunity to request accommodations or alternative assessment methods.

First, neither the ordinance nor the proposed rules define "job qualifications and characteristics." This leaves employers with the discretion to provide vague or limited information about the characteristics or qualifications that an AEDT will assess. Posting such information on a website does not address the information gap candidates may encounter. Rather, notice should describe to candidates all criteria that the AEDT will evaluate, the purpose of evaluating that criteria, and the data used to measure those criteria.

Second, by allowing employers to satisfy their "notice" obligation by posting the information on a website, the ordinance effectively places the onus on workers to find out that they are being subject to an AEDT and how the AEDT will be used. This exacerbates a flaw in the ordinance itself, which does not require employers to proactively provide candidates with notice about the type of data an AEDT will collect, the sources of such data, or the employer's data retention policy. Rather, it requires employers who do not include this information on their websites to provide it to candidates *upon written request*. The proposed rules simply affirm this, only adding a requirement to post instructions on employers' websites on how to make such written requests--again, placing the onus on candidates to find the necessary information. This is not "notice" in any meaningful sense.

Third, neither the ordinance nor the proposed rules require employers to ensure that websites where notices are provided are easily accessible to candidates with disabilities or candidates without reliable Internet access. Further, the proposed rules do not clarify that employers have a duty to fulfill requests for reasonable accommodations or alternative assessment methods.

We recommend that the final rules:

- Affirm employers' duty to provide accommodations, recognizing that the use of an AEDT does not relieve employers of this responsibility.
- Require employers to proactively notify candidates that they will be assessed by an AEDT, with a description of the qualifications and characteristics the AEDT will measure, types of data it will use, the sources of the data, and employers' data retention policies.
- Require employers to provide this notice through multiple channels that make the notice accessible to candidates with different needs, rather than allowing employers to choose a single method of notice that does not effectively give candidates the information they need to decide whether they should request accommodations.

[18] N.Y.C. Admin. Code § 20-871(b)(1).
[19] *Id.*

Conclusion

While LL 144 recognizes the need for accountability in the use of AEDTs, the law's requirements fall short by merely clarifying employers' existing obligations to report race-, ethnicity-, or sex-based disparate impact and offering notice requirements that employers could easily circumvent. The proposed rules do not effectively implement--much less strengthen--the law's obligations. Instead, the proposed rules limit the law's protections further with narrowed definitions that would exclude several tools from the auditing and transparency requirements. We urge the Department to ensure that final rules maximize the law's requirements to more effectively hold employers accountable for preventing discriminatory outcomes when they deploy AEDTs.

Respectfully submitted,

Matt Scherer
Senior Policy Counsel for Workers' Rights
Privacy & Data Project

mscherer@cdt.org

Ridhi Shetty
Policy Counsel
Privacy & Data Project

rshetty@cdt.org

U.S. Chamber of Commerce

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

October 24, 2022

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

Re: Comment on Proposed Rules

Dear Commissioner Mayuga:

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. The recent data from the Labor Statistics for the New York City Region indicates that "The city's seasonally adjusted unemployment rate was 6.6 percent in August 2022, up 0.6 percent from July."[20] This increase in the unemployment rate is due to a "rise in the number of people entering the labor force." In a time when the region is seeing growth in the labor market, and people are looking to enter the workforce, we feel it is essential that the City make sure that the rules implementing Int. 1894-2020 in relation to "automated employment decision tools," are made in a thoughtful and balanced manner, which will allow for the deployment of such tools to benefit the employer, employee, and/or independent contractor to help streamline the process.

The U.S. Chamber of Commerce has long recognized that "fostering public trust and trustworthiness in A.I. technologies is necessary to advance its responsible development, deployment, and use." For this reason, we at the Chamber appreciate the opportunity to provide the following comment on the New York City Automated Employment Decision Tools Regulation, which we believe will help provide more certainty to the framework.

<u>Definition of Automated Employment Decision Tool</u> , We urge you to revise the definition of "automated employment decision tool" and, more specifically, the definition of the phrase "to substantially assist or replace discretionary decision making." First, we would ask that you strike the last phrase "or to use a simplified output to overrule or modify conclusions derived
from other factors including human decision making." In its place, we 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."

Furthermore, most employers do not use AEDT as the sole factor for determination of whom to employ but use it in a more holistic approach as one of many factors in evaluating a candidate. Although an employer may deploy an algorithm on every potential candidate, that does not mean that the output is always used by the employer. While an AEDT may review

[20] https://dol.ny.gov/labor-statistics-new-york-city-region

and/or score each candidate, an employer may still empower the hiring manager with the discretion to determine whom to interview and the amount of reliance and weight they put on the tool's output. With hiring managers potentially weighing their expertise and analysis more than the AEDT output, we ask for clarification on category (2) ("to use a simplified output as one of a set of criteria where the output is weighted more than any other criterion in the set").

Finally, if you choose not to delete category (3) ("to use a simplified output to overrule or modify conclusions derived from other factors including human decision making"), we ask that "or modify" be struck or clarification be provided on what it means to "modify" a conclusion.

Bias Audit: The examples provided in subsections (a) and (b) 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 (a):

- Revise the initial phrase to read: "Where an AEDT is used to screen a candidate or employee for an employment decision, a bias audit required by § 20-871 of the Code must, at a minimum"
- In the example, strike "historical data" and replace it with "test data."

We request the following changes to subsection (b):

- In the example, strike the word "planned" from the phrase "planned use of the AEDT."
- Also, in the example, strike "historical data" and replace it with "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, Non-Hispanic 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 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 demographic data.

Published Results: We ask for the addition of the following italicized phrase in section (a) "Prior to the use of an AEDT to screen a candidate or employee for an employment decision, employers and employment agencies in the city must make the following publicly available on the careers or jobs section of their website in a clear and conspicuous manner:"

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."

Definition of Screen in relation to Employment Decisions: The law states that "employment decision" means "to screen candidates for employment or employees for promotion within the city." The proposed rule defines "screen" as "to make a determination about whether someone should be selected or advanced in the hiring or promotion process." We ask for clarity on how the use of some A.I. tools would fit within this definition.

Vendor Audits: The proposed rules contain an example in §5-301(a) 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.

Frequency of Audits: The law states that an AEDT cannot be used unless a bias audit was "conducted no more than one year prior to the use of such tool." It is not clear whether that language requires yearly bias audits of the tool or if conducting one audit on a tool is sufficient unless or until the tool is replaced or materially modified. The Statement of Basis and Purpose of the Proposed Rule states that a bias audit is required "within one year of use of the tool" which implies that the audit may take place within the 12 months which follow implementation of the tool. We understand this is not the intent, so this should be clarified.

Grace Period: While we understand the legislation requires the rules and regulations to go into effect by January 2023, the U.S. Chamber strongly encourages the Department to provide a grace period of at least twelve (12) months to businesses and organizations as they prepare to implement the final rule.

Conclusion: We appreciate the opportunity to comment on the implementing rules. It is essential that these regulations are implemented by New York City in a manner that does not impose overly broad requirements, which in turn could create significant uncertainty regarding the use of automated tools in hiring. Potential limitations of the use of technology for hiring purposes for businesses could lead to unnecessary barriers to finding qualified candidates for a job; this is particularly challenging during periods when we see both labor shortages and increases in the labor market, as businesses are put in a position where they receive more resumes/applications than they have the capability to review, which inhibits their ability to identify potential candidates. The use of automated employment decision tools is essential in helping streamline the hiring and promotion process. Thank you for considering the aboveproposed changes to give the business community the necessary certainty they will need. If you have questions, do not hesitate to contact Michael Richards at mrichards@uschamber.com.

Sincerely,

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

Indeed's Comments on Proposed Regulations to NYC Local Law No. 144
New York City Department of Consumer and Worker Protection
October 24, 2022

Indeed's mission is to help people get jobs. We are committed to increasing fair hiring and removing barriers for job seekers. In the New York City area, *5.5M job seekers* use Indeed to search for work and *124,000 businesses* use Indeed to find talent (Indeed Hiring Lab).
In the next 10 years, Indeed is committed to helping 30M job seekers facing barriers get hired. In order to reach our goal, we must ensure that algorithms increase fairness and inclusiveness in hiring. At Indeed, we monitor our algorithms for unintended bias by establishing checkpoints set out by Indeed's AI Ethics Committee, a team within ESG composed of data scientists, software engineers, product managers, and more.

In partnership with the government, we leverage our 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 committing to host a series of hiring events to help public schools hire teachers, administrators, counselors, and other school staff.

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

Feedback on Proposed Regulation

We support the Proposed Regulations to LL No. 144 for providing necessary clarity and we'd like to raise a few points for further clarification.
Under section 5-301 "Bias Audit", the "Example" discloses 1) demographic categories that intersect (both race *and* gender) and 2) sample size of candidates who applied and were selected.

1) Intersectional Demographic Categories: Reporting on multiple demographic categories within one group could allow individuals to be identified (particularly with small sample sizes) and may raise privacy concerns. Alternatively, data would be more anonymous if audits were conducted for race and gender separately. **We'd like clarification as to whether it is required to conduct selection rates and impact ratio for each intersectional group.**
2) Sample Size: Sample size metrics (number of applications, number of selections) can correlate to meaningful financial information about company performance. If this information is disclosed to clients and competitors, it could create financial concerns. **We'd like to confirm that it is *not* required to disclose sample size as part of the bias audits.**

Thank you for leading on this new issue and we'd love to be a resource as you finalize your regulations. I look forward to the opportunity to discuss our feedback and any other input you receive. Please feel free to reach out to Alison Klein (aklein@indeed.com) with any further questions.

October 24, 2022

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

By electronic delivery: LinkedIn Comments on the New York City Department of Consumer and Worker Protection's Proposed Rules Implementing Local Law 144 of 2021 Regarding Automated Employment Decision Tools

Dear Commissioner Mayuga:

LinkedIn, a New York City (NYC) employer with over 1,700 NYC employees, appreciates the opportunity to submit comments on the proposed rules issued by the NYC Department of Consumer and Worker Protection (Department or DCWP) to implement Local Law 144 of 2021 regarding Automated Employment Decision Tools.

LinkedIn is the world's largest professional network, whose vision is to provide economic opportunity for every member of the global workforce. LinkedIn provides an online platform that allows our 850 million members worldwide to create professional connections, access employment opportunities, and learn new skills. The platform also provides tools for employers and recruiters to search for prospective job candidates.

As a company whose mission to connect workers lies at the intersection of talent and technology, LinkedIn not only recognizes the opportunities that come with the use of automated technologies in the employment sector, including those using artificial intelligence and/or machine learning, but also fully recognizes the importance of thoughtful and responsible use of these technologies. In carrying out our mission to create economic opportunity for every member of the global workforce, delivering fair and equitable experiences for each of our 850M members is paramount to this work. As a result, we design our systems in ways that help us provide the right protections, mitigate unintended consequences, and ultimately better serve our members, customers, and society.

We share the goals of the NYC City Council in seeking to ensure that workers are protected throughout the recruiting and hiring process and that the right skilled candidates have an equitable opportunity to thrive in the workforce. As such, we applaud the Department's efforts in the proposed regulations to provide additional clarity with regard to the scope of tools that are covered by Local Law 144, along with the Department's recognition that there are various legal and other considerations that are applicable during the early stages of recruiting and sourcing of candidates. It is at this very point in the hiring process—searching for prospective candidates—where we believe automated tools can many times serve an important function of *increasing* access to the job searching process for potential candidates of all backgrounds.

We appreciate the Department's attention to these important issues and the opportunity to provide feedback on the proposed rules. Our goal is to ensure that any final regulations are clear, harmonized with existing employment laws where appropriate, and avoid creating any unintended barriers to implementation. Please find our comments below, including three general themes related to the proposed Definitions, the Bias Audit Process, and the Timing of Notice requirements.

DEFINITIONS

"Automated Employment Decision Tool": LinkedIn appreciates the improved clarity that the proposed regulations provide with regard to the scope of tools regulated by Local Law 144. However, given the general definition of automated employment decision tool (AEDT) in the proposed regulations focuses on a tool that *replaces or substantially assists* human decision making, we are concerned that the third alternative definition is unclear in that it refers simply to use of a simplified output to "modify" a human decision. In practice, when simplified outputs from tools are used as data points as part of a holistic human review, it will be difficult for employers to identify when and how this language may apply, resulting in inconsistent employer and candidate experiences that do not align with the apparent intent of the law.

To aid in clarifying the definition of "automated employment decision tool" in the proposal, we would recommend striking the phrase "or modify" from the final clause of the definition.

"Screen": Ensuring internal consistency across definitions within the regulation, particularly where important distinctions are made, is vital to achieve the regulation's intended goals. We appreciate the Department's recognition, in its proposed definition of "candidate," that there are different laws and policy considerations that come into play during the sourcing/candidate discovery phase of recruiting. Notably, there is no "screening out" of individuals happening at that early stage; rather, this is the stage at which recruiters and employers search for prospective candidates to *invite to submit* applications to a publicly available position. In our experience, this is an area of great opportunity for modern technologies to enable hiring companies and other employers to conduct such searches with the specific goal of developing a diverse candidate pool. This may help employers advance affirmative action obligations consistent with federal and state laws, where they exist, or simply create a broader and more diversified pool of candidates for a role.

To increase clarity and consistency across the draft regulation's definitions and advance the goals of the proposal, we would recommend replacing "someone" with "a candidate or employee" in the definition of "screen."

BIAS AUDIT PROCESS

Classifications: The way in which this section speaks to "classifying individuals into groups" may be confusing to employers, causing different employers to apply the regulations differently. Typically, in the context of employment decision-making audits, using a selection rate/impact ratio calculation suggests that there is a positive outcome (an individual is hired or moves forward in process) and a negative outcome (an individual is rejected or does not move forward) to be compared for the purpose of the calculations. If a tool is used in some way to classify individuals into groups, it is entirely possible that there is no "advantaged"—or positive—group but simply differences among groups.

For example, one could envision a tool that classifies resumes with significant sales experience as primarily sales related and resumes with significant engineering experience as primarily engineering related, but without creating any preference or advantage in the hiring or promotion process. Calculating selection rates and impact ratios across such classifications would not then yield meaningful outcomes about any sort of screening or selection and may create confusion and unnecessary burden.

We would recommend amending the language to clarify that the regulation applies where an AEDT "classifies those individuals into groups in a way that prefers a group to move forward in a hiring or promotion process." In addition, we recommend replacing "individuals" with "candidates or employees" for the same reasons related to providing clarity and consistency across the draft regulation identified in our comments on the definition of the term "screen."

Additionally, it would be helpful to insert language that clarifies that the included examples are only that—illustrative examples of acceptable methods of bias audit—and not prescriptive of exactly who is responsible for conducting audits or the specific data required to be audited in every case. In our reading of the proposed regulations, this appears to be the intended meaning, but providing express acknowledgment of this intent would improve understanding and enable better employer compliance.

<u>Timing of the Audit</u>: The plain language of the regulation, and the proposed regulations' reference to the distribution date, appear to contemplate that a bias audit must be conducted within a year of the AEDT being implemented/placed into use. Nonetheless, questions remain about whether the language creates an obligation for employers using an AEDT to perform an annual bias audit. We believe this is *not* the intent of Law 144 based on 1) the plain language of the law; 2) the fact that the alternative would imply some sort of rolling year calculation where every single individual use of the AEDT would trigger the bias audit calculation (an untenable and infeasible requirement); and 3) the fact that once an AEDT is in use as a screening procedure, it is already subject to existing employment laws such as the Uniform Guidelines on Employee Selection Procedures (UGESP) and general adverse impact discrimination law, which create the need for monitoring of ongoing impact. The proposed regulations do appear to recognize this in various provisions of the draft (*e.g.*, Example 2 in the Bias Audit section and the incorporation of the distribution date definition into the disclosure section). For avoidance of doubt, we suggest clarifying that the audit requirement is intended to be conducted in the year before placing an AEDT into use.

To facilitate this clarification, we recommend modifying section 5-302(a) to clarify that the bias audit is required "prior to *implementing* use of an AEDT" and section 5-302(c) to clarify that the posting is required "for the duration of the time the AEDT is used" and for 6 months after that. Also, in the Bias Audit section, we recommend changing Example #1 to mirror the existing language in Example #2 and clearly state that the audit is required "within a year from the *planned* use of the AEDT."

<u>Publication of Audit Results</u>: While well-intended, we believe a requirement to publish the specific selection rate and impact ratio information may in practice undermine the goals of the regulation given the existing legal landscape. Specifically, depending on the type of audit and data reviewed, it may require publication of information that could appear to constitute adverse impact under existing employment law guidance. Notably, existing anti-discrimination law acknowledges that the existence of adverse impact based on numerical calculations alone is not unlawful under all circumstances. Thus, a summary statement requirement better aligns with the letter and spirit of existing employment laws prohibiting discrimination, while also accomplishing the goals of this ordinance.

To better achieve the desired outcomes of the regulation, we recommend striking "the selection rates and impact ratios for all categories" in Section 5-302(a)(1) and replacing it with "a statement on adverse impact." Alternatively, we ask the Commission to consider leveraging language from the recently released White House framing of algorithmic discrimination protections, and replacing the final clause of 503-2(a)(1) with "which may include confirmation that if negative impacts were observed on any category, the employer or employment agency has retained the relevant materials and justification for any continued use of the AEDT."

TIMING OF NOTICE REQUIREMENT

Notifying Candidates: The language of Law 144 and the proposed regulation contemplate notice to candidates and employees 10 days prior to use of an AEDT. It is not clear whether this applies to every separate individual review of candidate/employee materials by an AEDT or whether it applies to the fact that an AEDT is in use in the process, though the option to effect notice via policy or career website suggests the latter. If interpreted as requiring 10 days-notice to each individual candidate before an AEDT could be used with respect to that candidate information, the consequence would be significant for employer hiring processes, slowing down the ability to hire by potentially placing candidates or employees in a 10 day "holding pen" before a resume could even be reviewed. This would have a negative impact on hiring employers but also would have a significant impact on job seekers.

To aid in clarifying this issue, we would recommend modifying the language of Section 5-303 to clarify that the notice is required 10 days prior to placing an AEDT into use for the job in question, rather than 10 days prior to each and every independent review by the AEDT.

CONCLUSION

LinkedIn appreciates the opportunity to provide comments on the proposed regulations implementing Local Law 144 governing automated employment decision tools. We look forward to working with the Department to implement these rules in a way that supports the employers and jobseekers of NYC.

Sincerely,

Henry Fong
VP, Legal
LinkedIn Corporation

Enclosure

Enclosure accompanying: **LinkedIn Comments on the New York City Department of Consumer and Worker Protection's Proposed Rules Implementing Local Law 144 of 2021 Regarding Automated Employment Decision Tools**

Proposed Rule Amendments

Section 1. Chapter 5 of Title 6 of the Rules of the City of New York is amended to add Subchapter T as follows:

Subchapter T: Automated Employment Decision Tools §

5-300. Definitions.

As used in this subchapter, the following terms have the following meanings:

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 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.

Bias Audit. “Bias audit” means “Bias audit” as defined by § 20-870 of the Code.

Candidate for Employment. “Candidate for employment” means a person who has applied for a specific employment position by submitting the necessary information and/or items in the format required by the employer or employment agency.

Category: “Category” means any component 1 category required to be reported by employers pursuant to subsection (c) of section 2000e-8 of title 42 of the United States code as specified in part 1602.7 of title 29 of the code of federal regulations, as designated on the Equal Employment Opportunity Commission Employer Information Report EEO-1.

Distribution Date. “Distribution date” means the date the employer or employment agency began using a specific AEDT.

Employment Decision. “Employment decision” means “Employment decision” as defined by § 20-870 of the Code.

Employment Agency. “Employment agency” means “Employment agency” as defined by 6 RCNY § 5-249.

Independent Auditor. “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.

Impact Ratio. “Impact ratio” means either (1) the selection rate for a category divided by the selection rate of the most selected category or (2) the average score of all individuals in a category divided by the average score of individuals in the highest scoring category.

$$\text{Impact Ratio} = \frac{\text{selection rate for a category}}{\text{selection rate of the most selected category}}$$

OR

$$\text{Impact Ratio} = \frac{\text{average score of individuals in a category}}{\text{average score of individuals in the highest scoring category}}$$

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.

Screen. "Screen" means to make a determination about whether ~~someone~~ a candidate or employee should be selected or advanced in the hiring or promotion process.

Selection Rate. "Selection rate" means the rate at which individuals in a category are either selected to move forward in the hiring process or assigned a classification by an AEDT. Such rate may be calculated by dividing the number of individuals in the category moving forward or assigned a classification by the total number of individuals in the category who applied for a position or were considered for promotion.

Example. If 100 Hispanic women apply for a position and 40 are selected for an interview after use of an AEDT, the selection rate for Hispanic women is 40/100 or 40%.

Simplified output. "Simplified output" means a prediction or classification as specified in the definition for "machine learning, statistical modelling, data analytics, or artificial intelligence". A simplified output may take the form of a score (e.g., rating a candidate's estimated technical skills), tag or categorization (e.g., categorizing a candidate's resume based on key words, assigning a skill or trait to a candidate), recommendation (e.g., whether a candidate should be given an interview), or ranking (e.g., arranging a list of candidates based on how well their cover letters match the job description). It 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.

§ 5-301 Bias Audit

(a) Where an AEDT selects ~~individuals~~candidates or employees to move forward in the hiring process or classifies those individuals into groups [in a way that prefers one group to move forward in a hiring or promotion process], a bias audit required by § 20-871 of the Code must, at a minimum:

(1) Calculate the selection rate for each category;

(2) Calculate the impact ratio for each category; and
(3) Where an AEDT classifies individuals into groups as described above, the calculations in paragraphs (1) and (2) of this subdivision must be performed for each such classification.

Example: 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 within a year from the planned use of the AEDT. The employer asks the vendor for a bias audit. The vendor uses historical data it has collected from employers on applicants selected for each category to conduct a bias audit as follows:

			# of Applicants	# Selected	Selection Rate	Impact Ratio
Hispanic or Latino		Male	205	90	43.9%	0.841
		Female	190	82	43.2%	0.827
Non/Hispanic or Latino	Male	White	412		52.2%	1.000
		Black or African American	226	215 95	42.0%	0.806
		Native Hawaiian or Pacific Islander	87	37	42.5%	0.815
		Asian	321	167	52.0%	0.997
		Native American or Alaska Native	24	11	45.8%	0.878
		Two or More Races	115	52 197	45.2%	0.866
	Female	White	385		51.2%	0.981
		Black or African American	164	75	45.7%	0.876
		Native Hawaiian or Pacific Islander	32	15	46.9%	0.898
		Asian	295	135	45.8%	0.877
		Native American or Alaska Native	17	7	41.2%	0.789
		Two or More Races	98	44	44.9%	0.860

(b) Where an AEDT scores applicants or candidates, a bias audit required by § 20-871 of the Code must, at a minimum:
(1) Calculate the average score for individuals in each category;
(2) Calculate the impact ratio for each category;

Example: An employer uses an AEDT to score applicants for "culture fit". To do so, the employer must ensure that a bias audit of the AEDT was conducted within a year from the planned use of the AEDT. The employer uses historical data on average "culture fit" score of applicants for each category to conduct a bias audit as follows:

-	-	-	Average Score	Impact Ratio
Hispanic or Latino		Male	76	0.884
		Female	47	0.547
Non/Hispanic or Latino	Male	White	86	1.000
		Black or African American	76	0.884
		Native Hawaiian or Pacific Islander	65	0.756
		Asian	71	0.826
		Native American or Alaska Native	69	0.802
		Two or More Races	84	0.977
	Female	White	68	0.791
		Black or African American	56	0.651
		Native Hawaiian or Pacific Islander	74	0.860
		Asian	67	0.779
		Native American or Alaska Native	52	0.605
		Two or More Races	86	1.000

§ 5-302 Published Results

(a) Prior to ~~the~~implementing use of an AEDT, employers and employment agencies in the city must make the following publicly available on the careers or jobs section of their website in a clear and conspicuous manner:

(1) the date of the most recent bias audit of such AEDT and a summary of the results, [including a statement on any adverse impact] OR [which may include confirmation that if negative impacts were observed on any category the employer or employment agency has retained the relevant materials and justification for any continued use of the AEDT]~~ which shall include the selection rates and impact ratios for all categories~~ ; and,

(2) the distribution date of the AEDT to which such bias audit applies.

(b) The requirements of subdivision (a) of this section may be met with an active hyperlink to a website containing the required summary of results and distribution date, provided that such link is clearly identified as a link to results of the bias audit.

(c) Employers and employment agencies must keep the summary of results and distribution date posted for the duration of the time that they use the AEDT and for at least 6 months after last using such AEDT for an employment decision.

§ 5-303 Notice to Candidates and Employees

(a) For the purpose of compliance with § 20-871(b)(1) and (2) of the Code, an employer or employment agency may provide notice to a candidate for employment who resides in the city by:
   (1) Including notice on the careers or jobs section of its website in a clear and conspicuous manner at least 10 business days prior to ~~use of an~~placing an AEDT into use;
   (2) Including notice in a job posting at least 10 business days prior to placing ~~use of~~ an AEDT into use for the relevant job; or,
   (3) Providing notice to candidates for employment via U.S. mail or e-mail at least 10 business days prior to ~~use of an~~ placing an AEDT into use for the relevant job.

(b) For the purpose of compliance with § 20-871(b)(1) and (2) of the Code, an employer or employment agency may provide notice to an employee who resides in the city by:
   (1) Including notice in a written policy or procedure that is provided to employees at least 10 business days prior to ~~use of~~ placing an AEDT into use;
   (2) Including notice in a job posting at least 10 business days prior to use of an AEDT for the relevant job; or,
   (3) Providing written notice in person, via U.S. mail or e-mail at least 10 business days prior to use of an AEDT.

(c) The notice required by § 20-871(b)(1) of the Code must include instructions for how to request an alternative selection process or accommodation.

(d) For the purpose of compliance with § 20-871(b)(3) of the Code, an employer or employment agency must:
   (1) provide notice to an employee or candidate who resides in the city by:
      i. Including notice on the careers or jobs section of its website in a clear and conspicuous manner; or,
      ii. Providing written notice in person, via U.S. mail or e-mail within 30 days of receipt of a written request for such information.
   (2) where notice is not included on its website, post instructions for how to make a written request for such information on the careers or jobs section of the employer or employment agency's website in a clear and conspicuous manner.
   (3) Provide an explanation to a candidate or employee for refusing to provide information pursuant to § 20-871(b)(3) of the Code where disclosure of information required by the notice would violate local, state, or federal law, or interfere with a law enforcement investigation.

(e) Nothing in this subchapter requires an employer or employment agency to provide an alternative selection process

Consumer Data Industry Association 1090 Vermont Ave., NW,
Suite 200 Washington, D.C. 20005-4905

P 202 371 0910 Writer's direct dial: +1 (202) 408-7407

CDIAONLINE.ORG

October 24, 2022

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

Via Email: Rulecomments@dcwp.nyc.gov

Re: Proposed rules relating to the use of automated employment decision tools

Dear Commissioner Maygua:

I write on behalf of the Consumer Data Industry Association (CDIA) to offer our comments to the Department of Consumer and Worker Protection ("DCWP" or "Department") on proposed rules relating to the use of automated employment decision tools. The proposed rules are being issued under Local Law 144 (2021) ("LL 144").[21]

An employer's use of augmented intelligence is rarely a binary decision. When used appropriately by employers, a combination of human and technology tools can help reduce bias and get people to work more quickly in ways that are fair, legal, and replete with second chances. Technology in hiring can reduce unintentional discrimination

[21] The Consumer Data Industry Association ("CDIA") is the voice of the consumer reporting industry, representing consumer reporting agencies including the nationwide credit bureaus, regional and specialized credit bureaus, background check and residential screening companies, and others. Founded in 1906, CDIA promotes the responsible use of consumer data to help consumers achieve their financial goals, and to help businesses, governments and volunteer organizations avoid fraud and manage risk. Through data and analytics, CDIA members empower economic opportunity all over the world, helping ensure fair and safe transactions for consumers, facilitating competition and expanding consumers' access to financial and other products suited to their unique needs

in employment and reduce the chances that human error can create a barrier to employment. Technology in hiring meets the urgent need to hire quickly, hire smartly, and hire remotely. Any rule from DCWP must be carefully balanced to meet the needs of consumers as asserted in the passage of LL 144, with the needs of employers and employees to higher quickly. The proposed rule has the potential to slow down the hiring process across the City.

People should be able to have confidence that automated decision tool are lawfully compliant, and are used responsibly in their development, deployment, and use. Given the speed of technology, trust in automated decision tools is a shared responsibility of government and the private sector. Governance frameworks should be flexible and driven by a voluntary, and multistakeholder process. With those considerations in mind, CDIA offers comments on the proposed rule.

Definition of an automated decision tool (§ 5-300)

CDIA respectfully requests three changes to this definition. First, we request that the Department strike "or to use a simplified output to overrule or modify conclusions derived from other factors including human decision making" and substitute in its place the following "'Automated employment decision tool,' or 'AEDT,' does not include (a) the automated searching of resumes to identify candidate qualifications, including relevant skills or experience, or (b) tools that merely covert a manual process to an automated process without the introduction of machine learning or artificial intelligence, and does not include tools that do not result in a final employment decision."

Most employers do not use AEDT as the only factor in determining who to employ. Employers take a more holistic approach. For example, just because an employer may use an algorithm on every potential candidate, it does not mean that that tool is *always* used. While an AEDT may review each candidate, the discretion may still rest with a hiring manager to determine whom to interview, as well as the potential reliance and weight placed on the use of the tool.

Bias audit (§ 5-301)

We respectfully request that the examples offered in the proposed rule clarified so that, among other things, they should not necessarily represent an exhaustive list of all the possible scenarios that might be available to employers. We ask the Department to remove the specificity of responsibility in each of the two examples of a bias audit to bring flexibility to account for a myriad of employment scenarios.

The examples provided in Subsections (a) and (b) are prescriptive in noting who bears responsibility for the bias audit (i.e., the employer/deployer or the vendor/developer). Yet, the examples fair to cover a range of possible hiring scenarios. For this reason, the examples fall short of guidance for employers.

We request a change to section (a): "Where an AEDT is used to *Screen a candidate or employee for an employment decision*, a bias audit required by § 20-871 of the Code must, at a minimum:" We also request that the Department change Subsection (a) in the example, strike "historical data" and replace it with "test data." In Subsection (b) we ask that the Department strike "historical data" and replace it with "test data." In the model we request that the Department strike "planned."

Published results (§ 5-302)

We respectfully ask for two changes. First, an addition of the following in section (a) "Prior to the use of an AEDT to screen a candidate or employee for an employment decision, employers and employment agencies in the city must make the following publicly available on the careers or jobs section of their website in a clear and conspicuous manner:" Second, we request that the Department strike "the selection rates and impact ratios for all categories" and replace with "a statement on adverse impact."

Definition of screen in relation to employment decisions (§ 5-300)

Local Law 144 defines an "employment decision" as a "screen [of] candidates for employment or employees for promotion within the city." § 20-870. We request clarity on how the use of some A.I. tools would fit within this definition. For example, we do not believe that if a tool makes a recommendation about an individual, but that recommendation is subject to human review, that that recommendation qualifies as a "determination." Further, since LL 144 does not extend beyond the limits of New York City, we encourage the Department to clarify the proposed rule to state that "within the city" means positions located in the City of New York, regardless of where the employee resides.

Independent auditor (§ 5-301)

The proposed rule defines an independent auditor as "a person or group that is not involved in using or developing an AEDT." We request that the proposed rule confirm the implication that an employer or vendor can conduct a bias audit with internal employees, provided that the group conducting the audit is not involved in using or developing the specific tool.

Obligation to provide an alternative selection method (§ 5-303)

Local Law 144 says that required notices to candidates must allow them "to request an alternative selection process or accommodation." § 20-871.b.1. But, the proposed rule states that nothing in the rules "requires an employer or employment agency to provide an alternative selection process." We read that as saying employers must provide *an opportunity* to request an alternative selection process in their notices. This requirement is a challenge because employers are never obligated to provide the alternative. The rule should confirm this logical interpretation.

Frequency of audits

Local Law 144 says that an AEDT cannot be used unless a bias audit was "conducted no more than one year prior to the use of such tool." § 20-871.b.1. It is not clear whether that language requires yearly bias audits of the tool or if conducting one audit on a tool is sufficient unless or until the tool is replaced or materially modified. The Statement of Basis and Purpose of the Proposed Rule states that a bias audit is required "within one year of use of the tool." This clause implies that the audit may take place within the 12 months which follow implementation of the tool. Since this requirement was not the intent of the proposal, it should be clarified.

Grace period

While we understand the legislation requires the rules and regulations to go into effect by January 2023. We strongly encourage the Department to provide a grace period to businesses and organizations as they look to implement the final rule.

Summary

CDIA appreciates the opportunity to comment on the implementing rules. An employer's use of augmented intelligence is rarely a binary decision. When used appropriately by employers, a combination of human and technology tools can help reduce bias and get people to work more quickly in ways that are fair, legal, and replete with second chances. Any rule from DCWP must be carefully balanced to meet the needs of consumers as asserted in the passage of LL 144, with the needs of employers and employees to higher quickly. The proposed rule has the potential to slow down the hiring process across the City.

Sincerely,

Eric J. Ellman
Senior Vice President, Public Policy & Legal Affairs

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

October 24, 2022

New York City Department of Consumer and Worker Protection

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

Thank you for the opportunity to submit 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. We are interested in ensuring clarity in the roles and responsibilities of each party, and whether the test of applicability is the location of the posted position or posting company, or the location in which the tool was used. We further hope to focus the proposed rules on the kinds of tools and use cases that may actually have a meaningful, measurable impact on final employment-related decisions. Additionally, we seek clarity on the interoperability of the proposed rules with existing state and federal laws and standards, including the Equal Employment Opportunity Act.

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 search.

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 employment position; (2) the automated searching of resumes or other relevant documents to identify candidate qualifications including relevant skills or experience; or (3) filtering or automated labeling of candidates to help identify qualifications

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.

Bias Audit

When it is determined that an AEDT has been used for screening purposes, it is unclear to TechNet exactly which party – the prospective employer or the maker of the AEDT – is responsible for performing the required bias audit. We are concerned that both parties may experience significant operational difficulties in complying with the rules as currently proposed.

Under existing federal standards, the demographic information required in the proposed bias audit may be unavailable or uninstructive depending on which party is deemed responsible. Under the Equal Employment Opportunity Act, the provider of the AEDT is not required to collect demographic information, while most (but not all) prospective employers are.

TECHNET
THE VOICE OF THE
INNOVATION ECONOMY

Furthermore, employers cannot require applicants to provide that information, generally offering an option like "prefer not to say". In this case, in order to comply with the proposed NYC rule, is the Department proposing that relevant parties will infer the missing information or compel applicants to provide information that is not currently required under existing standards?

If certain groups opt out of providing demographic information at higher rates than others, the resulting audits would be inaccurate and misleading. This problem would be particularly acute in smaller sample sizes, and is compounded by the requirements to conduct intersectional analysis, which further diverges from existing US standards. If companies have a small pool of relevant job postings or applicants, aberrations that are not otherwise statistically significant could paint a picture of an employer that does not accurately reflect reality. Would the Department consider a trigger that ensures these audits would be required only in cases where the volume or quality of data available would produce useful and actionable audits that provide statistically significant information?

Enforcement

In the interest of addressing the many open questions, the potential difficulties of operationalizing this new rule, and the length of time that it may take to conduct a bias audit compliant with the proposed rules, we ask that the Department consider a delay in enforcement for at least 6 months after final rules are promulgated.

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

Legislative Affairs
One Whitehall Street
New York, NY 10004
212-607-3300 www.nyclu.org

NYCLU
ACLU of New York

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 Proposed Rules to Implement Local Law 144 of 2021

October 24, 2022

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.[22] 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 Consumer and Worker Protection ("DCWP" or "Department") would stymie the law's mandate and intent further by limiting its scope and effect.

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

[22] *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.

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.[23] 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"). follows. 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.

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;[24]
- 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 where 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

[23] 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).

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

traits or characteristics are to be measured, and the disabilities, if any, that might potentially lower the assessment results or cause screen out."[25]

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.

LL 144 mandates that any "tool has been the subject of a bias audit conducted no more than one year prior to the use of such tool." The DCWP's rules should clarify that this requires annual auditing for continued AEDT use and provide an example similar to the rate calculation scenarios.

Re: § 5-302 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.

[25] 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.

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: § 5-303 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.[26]

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 testimony. 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 information they need.

[26] *See:* Algorithms, Artificial Intelligence, and Disability Discrimination in Hiring, BETA.ADA.GOV (2022), https://beta.ada.gov/resources/ai-guidance/.

Comments of the Partnership for New York City to the Department of Consumer and Worker Protection
Proposed Rules Implementing Local Law 144 of 2021
October 24, 2022

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 appreciates the work the Department of Consumer and Worker Protection has done to craft these proposed rules. Below are recommended changes and questions:

§ 5-300 – Definition of Automated Employment Decision Tool

- Issue: Applying the law to tools that simply "modify" conclusions goes beyond what is required by LL144. It also leads to substantial confusion because it is unclear whether the use of a tool's results as one factor in a human's decision would be considered to "modify conclusions derived from other factors including human decision-making." o For example, does a screening tool count as an AEDT when used as follows? The tool compares an applicant's resume and application to the posted job description and provides a score reflecting how closely the applicant matches the job description. The score is used to prioritize applicants for review by a recruiter, however the recruiter reviews the resumes and applications for all applicants. The recruiter determines whether the applicant is qualified to move forward in the hiring process.

o Recommendation: Strike "or modify".

§ 5-300 – Definition of Independent Auditor

- It is not clear whether an independent auditor must be external to the employer.

o Recommendation: Clarify that a group within a company that does not use the AEDT may serve as an independent auditor for the AEDT used by another group in the company. This is especially important for companies that develop and use AEDTs.

§ 5-301 – Bias Audit

- Employers who use any kind of tests in employment are already required by the federal Equal Employment Opportunity Commission (EEOC) to conduct an adverse (a.k.a. disparate) impact analysis on the test. The analysis requires employers to make comparisons between a minority group and a majority group (*e.g.*, women vs. men) but does not require the analysis to include intersectional subgroups (e.g., Black men vs. Caucasian men). Including intersectional groups requires substantially more work and may result in meaningless output where specific population sizes are too small.

o Recommendation: The rules should mirror the EEOC's adverse impact analysis requirements.

- The proposed rule appears to require employers or vendors to use actual or historical data to conduct the bias audit. This has two problems:
  - Employers may request job applicants disclose demographic data voluntarily. If applicants do not volunteer this information, the bias audit results may be inaccurate. This would be particularly true if certain categories of applicants refuse to provide the information at higher rates.
  - Requiring the use of historical data (as in the example in §5-301(a)) implies that the AEDT has previously been used. This would preclude the use of new tools.
  - Recommendations: Allow the bias audit to be conducted with "test data" instead of actual or historical data.
    - If the final rules do not allow the use of test data, they should clarify whether the data used in the bias audit should be from all applicants or New York City residents only.
- It is unclear whether an employer must conduct the §5-301(a) analysis or the §5-301(b) analysis if the employer uses an AEDT to create a score that selects individuals to move forward in the hiring process.
  - Recommendation: Clarify which test the employer is required to conduct.
- The examples provided in §§5-301(a) and (b) are specific about who performs the bias audit and do not account for all of the possible actors who could perform the audit in each example. Although these are listed as examples, this may still be confusing to employers.
  - Recommendation: Remove the reference to the actor performing the audit ("the vendor" in (a) and "the employer" in (b)). It would also be helpful to include a statement that the examples do not show all of the possible options for the audit process.
- The example in §5-301(a) suggests that the employer can ask the vendor providing the AEDT to conduct the bias audit. Since it is common for the AEDT vendor to also be the developer, allowing the vendor to conduct the bias audit appears to conflict with the definition of "Independent auditor" (*i.e.*, "a person or group not involved in using or developing the AEDT").
  - Recommendation: Clarify whether a vendor who is also the developer can be considered an independent auditor.
- In the example in §5-301(b), the word "planned" is confusing. LL144 requires an AEDT to have been audited within one year of use, not within one year of planned use.
  - Recommendation: Delete the word "planned."

§5-302 – Published Results

- Publishing the impact ratio and/or selection rate does not provide the applicant with a way to assess the existence of bias. Moreover, LL144 only requires a summary of results, not the publication of the calculations used to conduct a bias audit.

o Recommendation: Require a "statement on adverse impact" which is more akin to a summary than specific rates.

§5-303 – Notice to Candidates and Employees

- Although §§5-303(a)(3) and (b)(3) allow an employer to provide individual notices only to job applicants who live in New York City, this is not administratively feasible (especially if the employer is required to wait 10 days to use the AEDT after providing notice).

○ Questions:

- Is the intent to require employers to provide the notice to all applicants for all jobs, regardless of where they are located?
- Is the employer obligated to wait 10 days between notifying the applicant and administration of the AEDT?
- Could an applicant be provided a notification that includes a choice to take an assessment that is an AEDT within the 10-day period?

Effective Date

- The final rules seem likely to be released shortly before the January 1, 2023 effective date for LL144. This will not leave employers with sufficient time to implement the rules.

○ Recommendation: The Department should provide an extended time for employers to implement the rules before enforcing the law.

Thank you for considering our comments and questions.

From: Davies, Rachael <Rachael.Davies@claconnect.com>
Sent: Tuesday, October 11, 2022 5:26 PM
To: Rulecomments
Subject: [EXTERNAL] Clarifications on Rules set forth for upcoming AI legislation

Some people who received this message don't often get email from rachael.davies@claconnect.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).

Hello,

To whom it may concern I have the following questions I hope may be addressed at the upcoming comment session on Oct. 24th.

1) In regard to the independent bias audit
    a. Does the audit need to be done on candidate data that is specific to that employer's use of the tool, or can the audit be done on the tool as a whole. Ie the AI vendor's data for its use across multiple employers?
    b. If the bias audit needs to be run on candidate data from the specific employer, how would one acquire said candidate data if the tool is not yet in use due to the requirement to run a bias audit ahead of its implementation?
2) In regard to the 10 day notice requirement
    a. Does that refer to 10 business days prior to the tool's general use? Or is it 10 business days' notice prior to the use for each specific role? Meaning a requisition would need to remain open and posted at minimum 10 business days prior to reviewing any of the candidate's since they are entitled to 10 days notice before the use of the tool?

Thanks in advance for the clarification

Rachael Davies (she/her/hers)

Employment Practices Senior People Solutions
Direct 626-844-2035
CLA (CliftonLarsonAllen LLP)
rachael.davies@CLAconnect.com

October 24, 2022

Hon. Vilda Vera Mayuga, Esq.
Commissioner
Department of Consumer and Worker Protection
42 Broadway, 9th Floor
New York, NY 10004
Rulecomments@dcwp.nyc.gov (By Email Only)

RE: Comments to Proposed Rules for Local Law 144 of 2021

Dear Commissioner Mayuga:

Resolution Economics, LLC, an international consulting firm with offices in New York, Los Angeles, Chicago, Washington, DC, and Charlotte, NC, 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 regarding the Proposed Rules implementing the Automated Employment Decision Tools law ("AEDT law") that goes into effect on January 1, 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 and employment and financial services matters, including in litigation settings, across every industry. Our professionals include highly trained and technical team members with PhDs, JDs, MAs, MBAs, CPAs, CFEs, and other qualifying expertise. Resolution Economics has been and is currently advising employers on how to evaluate the impact and navigate compliance obligations around automated employment decision tools.

Based on our review of the AEDT law and proposed regulations and our experience in this area, Resolution Economics submits the following questions regarding the Proposed Rules:

1. <u>What type of analysis is required in order to be in compliance with the bias audit requirement and not subject to penalties under the AEDT?</u>

Section 5-301 of the Proposed Rules raises several issues regarding the type of analysis required for a "bias audit."

a. That Section indicates that a bias audit must "calculate the impact ratio" for certain categories. (5 5-301(a) and (b)). The definition of "bias audit" in Local Law 2021/144, however, specifies that the bias audit must include testing "to assess the tool's disparate impact on persons" within those categories. (AEDT law 20-870).

The concept of "disparate impact" encompasses a much broader analysis than simple impact ratio calculations. How are these two different types of analyses to be reconciled? More specifically: If an impact ratio is outside the 0.8 and 1.25 range, is use of the tool automatically a violation of the law that subjects the employer to penalties pursuant to 20-872 of the AEDT law?

Or does the employer have the opportunity to show that there are business or job related legitimate reasons for the observed differences, as would occur under disparate impact analysis?

b. The impact ratio tables (S 5-301(a) or (b)) given as an example in the Proposed Rules calculate the impact ratios for the intersection of race and gender categories. Does this mean that race-only and gender-only impact ratios are not required?

2. What level of specificity is required for a bias audit?

Generally, an analysis to determine whether a selection tool has a particular impact must be performed on a use case basis — that is, be based on data that is specific to a certain population and time period and that is specific to the employer's particular application of the tool. For example, if an employer were using a selection tool in relation to filling two positions at different times — say, one for a front desk clerk in February and another for a warehouse worker six months later separate analyses would need to be performed for each of these different applications of the tool.

It is not clear if the Proposed Rules require this type of analysis for AEDTs or instead permit some sort of single "generic" study that applies to all uses of a given tool. The example given under 5-301 (a) says "The employer asks the vendor for a bias audit. The vendor uses historical data it has collected from employers on applicants selected for each category to conduct a bias audit...". Does this mean that the vendor may use historical data collected from other employers that used this AEDT to carry out the bias audit? Furthermore, may different employers that plan to use the AEDT from the same vendor use the same bias audit conducted by the vendor using historical data from other employers to comply with the proposed rule? Or is it required that the data used for a bias audit be specific to a particular employer and to the specific application of the AEDT at issue?

The language of 5-302(a) (Published Results) raises this same issue. It states that employers must make the results of an AEDT bias audit publicly available prior to use of the AEDT. It would seem, however, that an employer would need to first implement an AEDT for some period of time to collect the data necessary as an input to the bias audit. Is it the case, therefore, that an employer needs to conduct some sort of pilot study with the AEDT prior to relying upon it to make employment decisions? Or may the employer (a) rely on a generic audit done by the vendor; and/or (b) use synthetic data — that is data that is generated partially or completely artificially, rather than being measured or extracted from real-world events — to conduct a bias audit prior to actual use of the AEDT?

3. What level of autonomy is required for the person/group performing the bias audit?

The AEDT Law states that the bias audit must be done by "an independent auditor." (5 20-870). Although the AEDT Law does not provide a definition of "independent," Merriam Webster defines independent as "(1): not subject to control by others; self-governing; (2): not affiliated with a larger controlling unit." In the context of financial statement audits by public accounting firms, employees of an audited entity are not generally considered to qualify as "independent auditors." See, e.g., 17 C.F.R. 210.2-01(c)(2) ("An accountant is not independent if, at any point during the audit and professional engagement period, the accountant has an employment relationship with an audit client").

The Proposed Rules define "independent auditor" as "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."

Where an employer is required to conduct a bias audit, does the Department intend this definition to allow the bias audit to be performed by an employee or group of employees of the employer?

Thank you for considering these questions.

Paul J. White

Paul White, PhD
Partner
Resolution Economics, LLC
Washington, DC

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, CA

October 24, 2022

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 Proposed Rules to the Use of Automated Employment Decision Tools Under Local Law 144 of 2021

Dear Commissioner Mayuga:

The Society for Human Resources Management ("SHRM") is the nation's preeminent HR trade association. More than 95% of Fortune 500 companies rely on SHRM as their go-to resource for all things work and their business partner in creating next-generation workplaces. With more than 300,000 HR and business executive members in 165 countries, SHRM enhances the lives of more than 115 million workers and families globally. SHRM creates better workplaces where businesses and workers thrive together. As the voice of all things work, workers, and the workplace, SHRM is the foremost expert, convener, and thought leader on issues impacting today's evolving workplaces. The SHRM vision is to build a world of work that works for all. Employers, like any other group, often look to technology as a path to innovation and progress. Technology may offer advanced tools that foster improved operations, create efficiencies, optimize interpersonal relations, and provide novel solutions to complex issues that arise within the workplace. During the past several decades, New Yorkers saw the workplace evolve as the use of fax machines, email, cloud computing, cellular phones, WiFi, GPS tracking, and other technologies changed and innovated how New Yorkers and the world worked. We now stand witness to the advent of artificial intelligence in the workplace.

SHRM submits these comments in response to the proposed rules issued by the Department of Consumer and Worker Protection ("DCWP" or "Department") on September 23, 2022, that stand to implement New York City Local Law 144 of 2021 ("LL 144") regarding automated employment decision tools ("AEDT"). We offer these comments with the insight of human resources professionals so as to make the proposed rule more effective.

SHRM applauds the Department's efforts to clarify the obligations of employers and employment agencies under LL 144. SHRM agrees that clearly defined terms and clarity regarding the bias audit requirements, the notice requirements, and other obligations under LL 144 serve to promote compliance and reduce confusion about the nuances of this first-of-its-kind law.

SHRM also commends the Department for its focus on addressing ambiguities and providing specificity around the requirements and obligations of LL 144 while also refraining from expanding the scope of the rules to require obligations beyond those that already exist. *See* Americans with Disabilities Act of 1990 (ADA), 42 U.S.C. §§ 12101-12213 (2018); NYC Human Rights Law, N.Y.C. Admin. Code § 8-107. In particular, we recognize the Department's efforts to clarify that while employers have certain obligations to reasonably accommodate workers on the basis of a disability, there is no obligation to provide alternative selection or evaluation processes. *See* §§5-303 (c) and (e).

The Equal Employment Opportunity Commission has recently provided valuable guidance to employers on the issue of automated employment decision tools and the impacts on individuals with disabilities. *See* U.S. Equal Employment Opportunity Commission, The Americans with Disabilities Act and the Use of Software, Algorithms, and Artificial Intelligence to Assess Job Applicants and Employees, 29 CFR Part 1630 & app. (May 12, 2022), *available at* https://www.eeoc.gov/laws/guidance/americans-disabilities-act-and-use-software-algorithmsand-artificial-intelligence. SHRM submits that all NYC employers and employees will stand to benefit from that guidance even as NYC moves forward with the requirements set forth in LL 144.

While recognizing that the issues are complex and require thoughtful implementation, the proposed rules require additional clarification. In the sections that follow, we summarize the key areas where further clarification could assist employers as they work towards implementing LL 144.

A. THE DEFINITION OF AUTOMATED EMPLOYMENT DECISION TOOL IS OVERLY BROAD AND RISKS UNINTENDED APPLICATION

SHRM commends the Department on its effort to further clarify the definition of AEDT. As it stands, LL144 § 20-870 defines AEDT 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. The term "automated employment decision tool" does not include a tool that does not automate, support, substantially assist or replace discretionary decisionmaking processes and that does not materially impact natural persons, including, but not limited to, a junk email filter, firewall, antivirus software, calculator, spreadsheet, database, data set, or other compilation of data.

The Department's proposed rules, § 5-300, add that "the phrase 'to substantially assist or replace discretionary decision making'" as used in LL 144 § 20-870, means to: (1) "rely solely on a simplified output (score, tag, classification, ranking, etc.), with no other factors considered," *or* (2) "to use a simplified output as one of a set of criteria where the output is weighed 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."

Further, "machine learning, statistical modeling, data analytics, or artificial intelligence," as used in LL 144 § 20-870, as defined in the proposed rules, 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.

Further, "simplified output," as used in LL 144 § 20-870, means a prediction or classification as specified in the definition for "machine learning, statistical modeling, data analytics, or artificial intelligence." A simplified output may take the form of a score (e.g., rating a candidate's estimated technical skills), tag or categorization (e.g., categorizing a candidate's resume based on keywords, assigning a skill or trait to a candidate), recommendation (e.g., whether a candidate should be given an interview), or ranking (e.g., arranging a list of candidates based on how well their cover letters match the job description). It 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.

The purpose of this law will be served best where employers have clear guidance as to when the law is triggered (i.e., which tools, processes, or systems are covered) and what, specifically, employers need to do to comply. While these further definitions seek to create clarity, SHRM observes that they nevertheless cast a very wide net as to the employer tools, processes, or systems that might fall into the definition of AEDT. This ambiguity is likely to lead to employers being left to guess as to what might or might not be covered by this definition. It is helpful that the definition excludes a "junk email filter, firewall, antivirus software, calculator, spreadsheet, database, data set, or other compilation of data," however the broad and conceptual phrasing of these definitions insufficiently identifies specific, concrete, reliable criteria for determining what tools, processes, or systems are covered. We encourage the final regulations to explicitly exclude automated tools that carry out human-directed assessments.

For example, some employers might use a scheduling tool that captures employee availability for purposes of both shift scheduling and candidate evaluation. That tool may "generate a prediction" where "a computer … identifie[d] the inputs" and decided the "relative importance placed on those inputs … in order to improve the accuracy of the prediction or classification," which involved "inputs and parameters" that were "refined through crossvalidation or by using training and testing data." And the purpose of this tool may be "to screen candidates for employment" by virtue of determining and calculating how to best schedule individuals to work. Moreover, with large candidate pools, employers rely on automation to screen candidates based on core job-related decision points such as educational attainment or relevant licensure. Presumably, these types of tools are not intended to be covered by this law, as subjecting this type of tool to the requirements of LL 144 would likely create a severe burden upon employers, and the continued use of such tools would be untenable in the face of the requirements and penalties under LL 144.

Furthermore, nearly 1 in 4 organizations use automation or AI to support HR-related activities, with 79% of the organizations utilizing AI in recruitment and hiring. Mostly extralarge organizations with 5,000+ employees (42%) and large employers with 500-4,999 employees (26%) are utilizing AI to support HR. In the next five years, 1 in 4 organizations plan to start using or increase their use of AI in the recruitment and hiring process. These tools will become more essential to communicating with candidates, reviewing resumes, and identifying skills. These tools will provide more opportunities and employment for untapped talent pools – veterans, opportunity youth, older workers, individuals living with disabilities, and people with criminal records.

Clarity and specificity will foster compliance with LL 144, which in turn will promote the goals and objectives of LL 144. Thus, SHRM recommends further specificity and clarity to assist New York City employers in compliance with this new law.

B. THE BIAS AUDIT PROVISIONS REQUIRE FURTHER CLARITY

SHRM commends the Department on its effort to provide clarity regarding the criteria for a "bias audit." However, additional clarification on the bias audit requirements is needed.

First, the scope of the candidate pool required to be tested is unclear. Specifically, it is not clear whether employers are required, or conversely, permitted, to conduct a bias audit on all candidates globally (i.e., in and outside of New York City in the case of multijurisdictional employers), only New York City candidates, or a sample set of either category. Presumably, the jurisdictional reach of LL 144 does not extend outside of New York City. Otherwise, multijurisdictional employers with significant workforces outside of New York City could be unreasonably burdened by the requirements and obligations of LL 144.

Employers should be expressly permitted to use broad data for the required bias audit purposes. For instance, § 5-302 requires that "[p]rior to the use of an AEDT, employers and employment agencies in the city must make … publicly available" "the date of the most recent bias audit of such AEDT and a summary of the results," as well as "the distribution date of the AEDT to which such bias audit applies." This requirement appears to create a "chicken-or-theegg" problem. An NYC employer looking to implement a new AEDT (or an updated version of existing AEDT, as discussed below) would have no prior data upon which to perform a bias audit, at least not in the first year, and thus could not begin to use the AEDT to obtain such data. The effect appears to be a de facto prohibition on implementing new AEDT tools. This issue would be easily addressed if the proposed rules were revised to clarify that NYC employers with operations outside the jurisdiction of NYC are permitted to include candidate data from other jurisdictions or alternately to use sample test data to conduct a bias audit of the AEDT to comply with LL 144.

Second, the proposed rules should clarify the timing implications with regard to the cadence of any required bias audits. At least one potential issue arises for employers who update or otherwise improve an AEDT. There should be a good faith standard to assess compliance in those instances. So long as a bias audit has been conducted on the tool itself, employers should not be required to conduct a new audit whenever enhancements are introduced. This is particularly true given that AEDT tools are continuously evolving.

Given that AEDTs are being used in the context of rapidly changing technological advancements, as well as the potential for penalties even in spite of employers' good faith and best efforts, SHRM recommends that the Department include a safe-harbor provision that permits employers to be deemed to have complied with LL 144 based on good faith efforts toward compliance. The Department should further consider providing that this safe harbor permits an AEDT to remain in use for the remainder of the time period during which the last bias audit was conducted.

Third, SHRM respectfully submits that the Impact Ratio metric that relies simply on "average scores" is a flawed methodology that does not adequately consider the effect size or the magnitude of the difference in scores. It appears to be a novel approach that, while seemingly simple to calculate, does not appear to be readily used by statisticians, labor economists, or industrial-organizational psychologists in the employment context. Indeed, we have been unable to find any interpretative guidance or peer-reviewed literature to support the use of the proposed "average score" methodology to determine an impact ratio result.

The use of the "average scores" impact ratio analysis is particularly problematic because it appears that it will be required in virtually all AEDTs. While two impact ratio methodologies are defined in the proposed rules and are seemingly designed to be used in the alternative,[1] it is difficult to conceive of an AEDT that "selects individuals in the hiring process or classifies individuals into groups" as set forth in § 5-301(a), without scoring applicants or candidates in some way as described in § 5-301 (b). Accordingly, under the proposed rules, NYC employers will be required to conduct a bias audit based on a flawed methodology that lacks an adequate benchmark for assessment purposes.

Fourth, while SHRM generally supports transparency efforts, there is some concern that the required notices could risk disclosure of candidates' personal identifiable information ("PII") or

[1] Impact Ratio. "Impact ratio" means either (1) the selection rate for a category divided by the selection rate of the most selected category *or* (2) the average score of all individuals in a category divided by the average score of individuals in the highest scoring category.

other information that is otherwise protected from disclosure. It is presumed that neither LL 144 nor the proposed rules require or otherwise seek disclosure of such information. SHRM thus recommends that the Department consider providing clarification in § 5-302 that "Nothing in this subchapter requires an employer or employment agency to disclose any personal identifiable information or any other protected or confidential information of any candidate or employee." Finally, SHRM commends the Department on the flexibility it has offered employers in determining who is considered to be an "independent auditor" of AEDTs. As NYC employers move quickly to come into compliance, being able to rely on a broad range of individuals or companies to implement the appropriate audit mechanisms will be of critical importance.

C. THE DEPARTMENT SHOULD GUARD AGAINST THE DANGERS OF STIFLING INNOVATION

SHRM commends the Department on its efforts to propose rules to create protections for candidates while providing clear, workable obligations for employers in this developing area where technology and law intersect.

The development and enhancement of advanced technologies in the workplace require flexibility to ensure that optimal processes benefit both the employer community, and candidates and employees in New York City. The war for talent, particularly after the pandemic, has created unique challenges for employers, including increased complexities due to the ever-evolving remote workforce.

As employers continue to navigate these real challenges, employers are exploring the use of AEDTs in myriad ways, including quickly sourcing and expanding candidate pools, identifying skills gaps, enhancing talent pipelines, reducing turnover, increasing productivity, and bolstering diversity equity and inclusion goals. Nearly 80% of employers report that candidates that score highly on a pre-employment assessment but do not meet the minimum experience requirement are likely to advance in the application and hiring process. Nearly 70% report the same for candidates that fail to meet the minimum education requirements for a position and are more likely to advance with high skills assessments. When carefully developed and implemented, employers and workers alike experience the positive benefits of AEDTs. Employers experience improved efficiencies which in turn benefits workers by optimizing the selection of candidates to ultimately place them in jobs that provide increased fulfillment and opportunities for success while enhancing career progression and movement in the workforce.

In light of the key benefits that AEDTs offer, SHRM submits 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, including veterans, older workers, individuals living with disabilities, opportunity youth, and people with a criminal record. 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. Appropriate safeguards must be balanced against heavy-handed regulatory restrictions that will set key HR functions back and impede the ability to create and identify broader, more inclusive talent pipelines. To do otherwise will only dampen innovation and stifle creativity within the confines of employment for NYC workers.

SHRM commends the Department on its efforts to propose rules that pioneer the realm of artificial intelligence and machine learning, making these rules among the first of their kind. We appreciate the opportunity to present our views on the proposed rules and would welcome the opportunity to assist the Department as it continues its rule-making process. If you should have any questions or require any additional information, please feel free to contact James Redstone at James.Redstone@shrm.org or C. Mitch Taylor at Mitch.Taylor@shrm.org.

Sincerely,

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

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)

October 23, 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.[1] I teach responsible data science courses to graduate and undergraduate students at NYU.[2] 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 and the voices of their representatives are conspicuously missing from 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.

For background: I actively participated in the deliberations leading up to the adoption of Local Law 144 of 2021[3][4] and have carried out several public engagement activities around this law

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

[2] All course materials are publicly available at https://dataresponsibly.github.io/courses/

[3] Testimony of Julia Stoyanovich before New York City Council Committee on Technology regarding Int

[4] -2020, November 12, 2020, available at https://dataresponsibly.github.io/documents/Stoyanovich_Int1894Testimony.pdf

when it was proposed [1]. 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 Times[2] and an article in the Wall Street Journal[3]. 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"[4]. (Course materials are available online[5].) 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: 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 concerned age-based discrimination in hiring. Local Law 144 does not currently control include any provisions to audit for this type of discrimination.

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 arbitrary[6]. 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

---

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

[2] 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

[3] 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

[4] "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

[5] "We are AI: Taking control of technology", NYU Center for Responsible AI, available https://dataresponsibly.github.io/we-are-ai/

[6] "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

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 assessment[9], 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**.

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 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 details[6].) 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 employees[1].

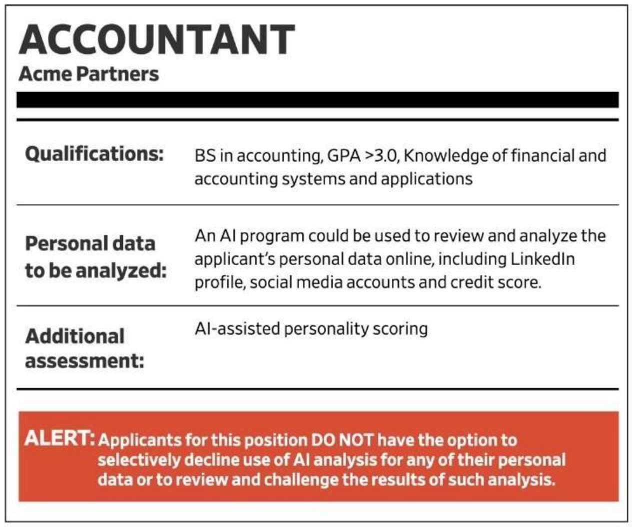

Figure 1: *A posting label is a short, simple, and clear summary of the screening process. This label is presented to a job seeker before they apply, supporting informed consent, allowing them to opt out of components of the process or to request accommodations.*

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

[1] 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

ai/

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.[1] I teach responsible data science courses to graduate and undergraduate students at NYU.[2] 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[3][4]and have carried out several public engagement activities around this law when it was proposed [5]. 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 Times[6] and an article in the Wall Street Journal[7]. 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"[8]. Course materials are available online[9].

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.

---

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

[2] All course materials are publicly available at https://dataresponsibly.github.io/courses/

[3] Testimony of Julia Stoyanovich before New York City Council Committee on Technology regarding Int

[4] -2020, November 12, 2020, available at https://dataresponsibly.github.io/documents/Stoyanovich_Int1894Testimony.pdf

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

[6] 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

[7] 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

[8] "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

[9] "We are AI: Taking control of technology", NYU Center for Responsible AI, available https://dataresponsibly.github.io/we-are-ai/

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. In their 2018 report, Bogen and Rieke[1] 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 University[2] that was broadly circulated by the press[3]. 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

---

[1] 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%20%20An%20Exploration%20of%20Hiring%20Algorithms,%20Equity%20and%20Bias.pdf

[2] 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

[3] 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

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.[1] It penalized resumes 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 reports[2][3]. 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."[4]

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.

[1] 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-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK08G

[2] 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-profess or/

[3] Stains, *"Are Workplace Personality Tests Fair?*", Wall Street Journal (2014) http://www.wsj.com/articles/are-workplace-personality-tests-fair-1412044257

[4] 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-deci sions

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 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.[1] 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.[2] 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.[3]

In my own work, I recently evaluated the validity of two algorithmic personality tests that are used by employers for pre-employment assessment[4]. 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.

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

[1] 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-mobil ity

[2] Narayanan, "*How to recognize AI snakeoil*" (2019) https://www.cs.princeton.edu/~arvindn/talks/MIT-STS-AI-snakeoil.pdf

[3] 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-dat a-transparency

[4] 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

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	✗	✓
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, with respect to facets of measurement: ✔ 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://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.[1]

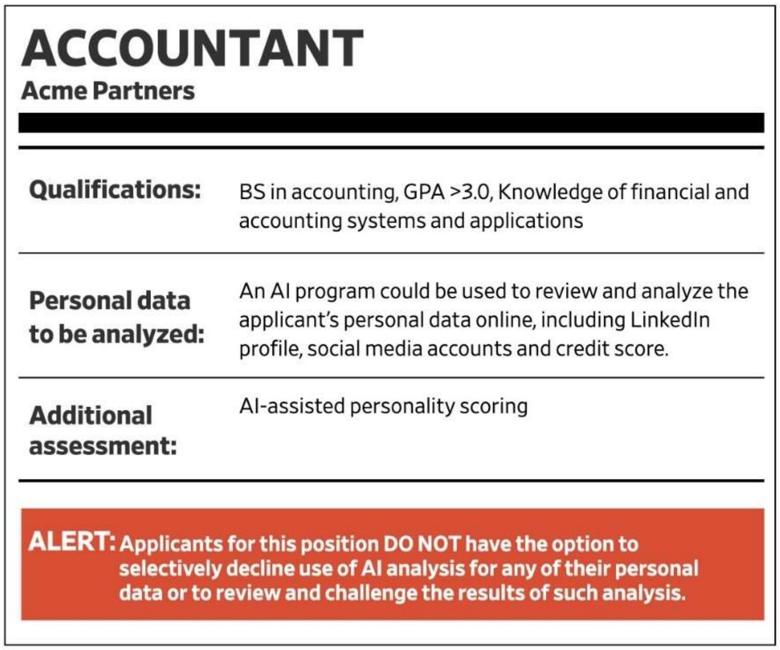

Figure 1: *A posting label is a short, simple, and clear summary of the screening process. This label is presented to a job seeker before they apply, supporting informed consent, allowing them to opt out of components of the process or to request accommodations.*

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 [2].

[1] Stoyanovich, Howe, Jagadish, *"Responsible Data Management"*, PVLDB 13(12): 3474-3489 (2020) https://dataresponsibly.github.io/documents/mirror.pdf
[2] 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,

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.

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".[1] 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 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."

---

https://www.eeoc.gov/laws/guidance/americans-disabilities-act-and-use-software-algorithms-and-artificialintelligence

[1] IEEE-USA, "*Artificial Intelligence: Accelerating Inclusive Innovation by Building Trust*" (2020) https://ieeeusa.org/wp-content/uploads/2020/10/AITrust0720.pdf

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

October 21, 2022

Please find below a few comments on the proposed rules for implementing LL144. In our opinion, providing some additional clarity in these areas will help with implementation at a technical and organizational level.

Score ratios are underspecified

The proposal for impact ratio for scores is the following metric:

$$\text{Impact Ratio} = \frac{\text{average score of individuals in a category}}{\text{average score of individuals in the highest scoring category}}$$

This kind of a metric can be underspecified since they depend on the scale in which the scores are measured.

For example, consider a probability scale and a numeric scale between 500-800 (like a FICO score) computed from the probability score as follows

Scaled score = 500 + 300* probability score

The table below shows how average scores from the probability scale translate into the numeric scale

	Average Score (prob)	Average Score (scaled)
Men	0.9	770
Women	0.5	650

If we calculate the **impact ratio in the probability scale, the impact ratio is 0.56** (=0.5/.09)

If we calculate the **impact ratio in the scaled score, the impact ratio is 0.84** (=650/770)

Thus, applying the four-fifths rule would lead to different conclusions depending on how the score is scaled.

We propose the following adjustment to clarify how the impact ratio is calculated:

Calculate impact ratio with selection rates at a given percentile threshold, where the percentile can depend on the use case.

The rule could also propose specific settings for the threshold.

In the example above, if we assume that the 50%ile threshold in the probability scores is 0.75, the corresponding threshold in the scaled score is 725. If we assume the following selection rates for men and women:

	Above 50%ile (0.75) (prob)	Above 50%ile (725) (scaled)
Men	0.5	0.5
Women	0.4	0.4

For both the probability and scaled scores, the impact ratios are equal to 0.8.

This approach avoids the issue described above because percentiles don't depend on the scaling.

Populations are underspecified

It will be useful for the proposed rules to specify which population the impact ratios have to be computed on.

Some of the options are

A. All job applicants that an AEDT tech system from a vendor scores.
B. All job applicants per employer that an AEDT system scores.
C. All job applicants per job family (above a certain number of applicants) per employer that an AEDT system scores.

Expectations from independent auditors

A. Beyond independence, what additional criteria are independent auditors expected to meet (for example: specific skills or experience)?
B. Does the auditor have to necessarily calculate impact ratios separately? Or can they focus on reviewing and validating the robustness of the process, and the completeness of the data and models/ systems, used by the auditee to calculate the ratio?

Public Comment on Rules to Implement Local Law 144 of 2021
New York City Department of Consumer and Worker Protection

October 24, 2022

Dear Commissioner Mayuga:

Thank you for the opportunity to comment on the Department of Consumer and Worker Protection's (DCWP or Department) draft rules implementing Local Law 144 of 2021 (LL 144). We appreciate that the Department's proposal reflects a serious consideration of the concerns raised by business and technology leaders regarding the provisions of LL 144. These draft rules attempt to establish a workable framework that ensures that New Yorkers can trust automated employment decision tools (AEDTs) will be deployed with protections against harmful discrimination.

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 supports thoughtful regulation to build trust in AI and has engaged with policymakers on AI policy at the federal, state, local, and global level.

We believe that developers and deployers of AEDTs in New York City (NYC or City) need clear rules to successfully meet LL 144's requirements. These rules should (1) ensure the independent review of AEDTs for disparate impact are in line with best practices from the field, (2) provide guidance to employers on their responsibilities when conducting a disparate impact audit, (3) inform job candidates and employers of their duties and rights (4) be harmonized with national and global frameworks, and (5) allow employers the time to comply. We offer the following comments and recommendations to highlight where DCWP's proposal is aligned with the goals outlined above and propose modest technical clarifications to ensure there is a clear path to compliance for employers:

1. **The proposed definition of independence with respect to AEDT requirements offers a positive path forward and aligns with best practices.** The Department's draft rules rightly provide employers with flexibility needed to meet LL 144's independence requirement. In doing so, the Department is aligning with best practices in auditing independence while also recognizing the nascent state of the AI auditing field.

DCWP's proposal follows best practices for independence by requiring that audits are conducted by a person or team of persons removed from developing the AEDT. An auditor's independence comes from the freedom to plan an inquiry, conduct a review, and report their objective judgments without pressure. Successful models for independent internal auditors include the financial services industry. Similarly, government agencies often rely on internal inspectors general to conduct audits, inspections, and investigations.

The technical standards and professional codes of acceptable practice needed to effectively audit AEDTs are still in their infancy. New York City employers deploying AEDTs will have differing levels of sophistication and technical capabilities. The draft rules therefore wisely provide employers with multiple paths to meet their obligation for independent reviews. The Department's proposed definition of independence also enables employers to meet LL 144's requirement in a manner that changes with the maturing AI governance landscape.

For the above reasons, we applaud the agency for their proposed definition and their commitment to ensuring broad compliance in a field where stringent audit standards remain in their infancy.

2. **Developers and deployers of AEDTs will need a ramp up period to build the necessary controls to comply with this first-in-the-nation law and novel regulatory requirements**. The Department should be sensitive to the needs of employers making good faith efforts to build the systems needed to implement these rules.

   DCWP is issuing these rules with only weeks between the time when they would be adopted and when they would take effect. Even the most forward-thinking employers will need time to ensure that their controls are both compliant and functional. The City has taken this approach in other instances when adopting broad, novel regulations. Similarly, other jurisdictions pursuing groundbreaking laws on automation and privacy have provided a year or more for companies to adjust their practices before pursuing any enforcement action.[1]

Recommendation: Establish a period for companies operating in good faith to implement LL 144's regulations without the threat of enforcement. The following language may be adopted:

> The corporation counsel shall not bring any enforcement action or proceeding for any violation of Subchapter 25 of Chapter 5 of Title 20 of the New York City Administrative Code until 270 days after the publication of these rules by the Department of Consumer and Worker Protection.

3. **Guidance on conducting a compliant disparate impact audit should follow the example set by federal guidelines.** The Department's proposal provides meaningful guidance to employers for auditing AEDTs, demonstrating that it understands that

[1] For example, California's Privacy Rights Act included a two-year ramp-up period for businesses to comply with the new and revised obligations.

establishing baseline normative standards is necessary for employers because LL 144 has charted entirely new ground.

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 is expected to issue guidance on AEDTs and employment decisions. It has already done so with respect to AEDTs and the *Americans with Disability Act*. **Recommendation**: Include language that makes clear that the examples provided in the regulation are illustrative guidance but not exhaustive and clarify that employers are compliant with LL 144 if they meet federal guidelines.

4. **When establishing a framework for regulating AEDTs, it is critical to use precise definitions that explicitly list the AI-powered tools regulated by LL 144 and omits those it does not.** The Department has offered a sound proposal that defines AEDTs with workable and practical limits. However, the Department's final rules should more precisely capture the tools LL 144 intends to regulate and, as importantly, clearly exclude other tools that were not intended to be regulated from their scope.

   Recommendation: Work with applied AI experts and industry leaders to further refine the definition of AEDTs.
5. **The Department's proposal requiring notice and public disclosure will allow consumers to make informed decisions on whether to engage with employers using AEDTs.** The draft rules require: (1) employers to provide notice to candidates and employees and (2) public disclosure of the bias audit results. DCWP should consider the following expansion and revisions to these provisions:

   - **Notices to candidates and employees in an alternative selection process.** The Department's draft rules do not require an employer to provide an alternative selection process, but, if offered, requires the employer to "include instructions for how to request an alternative selection process or accommodation."

     The Department should expressly allow employers to request from candidates or employees any information necessary for the purposes of identifying the candidate or employee and the associated employment opportunity. Without this, employers may be unable to accommodate job seekers opting for a manual review of their applications.

   - **Published results of the bias audit.** The draft rules 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: 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. The following language may be adopted to allows employers to facilitate a manual review:

> As a part of any alternative selection or accommodation request process, employers or employment agencies may reasonably request from candidates or employees information for the purposes of identifying the candidate or employee and the associated employment opportunity for which the candidate or employee has expressed interest. In the event that an employee or candidate does not provide such information, the employer or employment agency shall not be obligated to provide an alternative selection process or accommodation or other such information pursuant to §20-871(b) of Subchapter 25 of Chapter 5 of Title 20 of the New York City Administrative Code.

****

We appreciate your consideration of our comments. Please contact Jarrell Cook (jarrell.cook@workday.com), Senior Manager for State and Local Government Affairs, if Workday can provide further information as the Department finalizes these regulations.

Thank you for the opportunity to comment on recent guidance for this landmark law. As vendors in the model audit space, we provide tools, technology, and services to employers and HR vendors who offer AEDTs to their clients. From this perspective, alongside our work with other regulating bodies, we hope to share some insight from the front lines of our work with clients, policy-makers, and academia alike.

Setting an (Insufficient) Standard

As activists and academics in this space, our team has been working on issues of AI fairness, explainability, accountability, and model accuracy for many years. Our clients come to us because they know we possess deep knowledge of the many ways in which models can go wrong. In the past year, we have performed a number of audits prior to the release of this guidance that have sought to dig far deeper than the guidance requires. As has been previously stated by many in our community, the scope of the law and its proposed guidance, while desperately needed, may miss its intended effects of reducing discriminatory algorithmic behavior. Algorithms typically discriminate in a number of ways, from age to disability to race and gender, with many intersectional combinations thereof. By limiting the scope of the guidance as it currently stands, our prospective clients may be less motivated to pursue thorough examinations of their technology, leading to the shaping of our field's form of practice. The remainder of our comments will seek to illuminate many of the complexities and conflicts within the practice of algorithmic auditing, in order to display a few of the ways in which the guidance may fail to uncover discrimination where it may continue to exist.

Vendors vs. Employers

The proposed guidance contains a premise that may in many cases be untrue: that by modeling reporting from EEOC guidance, the commercial lift on industry will be reduced. However, employers are seldom the ones responsible for the development of AEDT tools, and employers far more frequently rely on vendors to provide the AEDT for a hiring process. This provides significant tension between the types of data needed for analysis, and the data that each entity possesses. Questions remain as to whether employers may themselves be allowed to simply repurpose the highest level reporting from the AEDT vendors, or whether these employers must also themselves perform an audit on their own limited candidate supply. One of the many ways that AEDTs can go wrong is due to generalizability errors. This is to say that vendor suitability may vary by employer (i.e. a tool that works well in healthcare may perform terribly and/or discriminate on engineers). The guidance, assuming it should take place at the global AEDT level, fails to address this kind of error, providing one of many ways that algorithms can be biased, but which will not be discovered in due course of a model audit seeking to replicate this draft for public reporting. One of the major conflicts between employer and vendor analysis for compliance with this law, deals with the question of whether and which entity may possess demographic data.

Demographic Data

Because the guidance assumes it will be produced by employers, who may not have access to match scores at scale provided by their vendor, many of them have requested copies of this report from their vendors for their own websites. This presents a catch-22 for the vendors, in that they largely intentionally limit their collection practices for demographic candidate data, in order to comply with conflicting laws that require data minimization and candidate privacy like the

GDPR. In the absence of self-reported data at application time, demographic data is difficult (and costly) to obtain. Candidates are often reluctant to provide this data post-hoc, as is well known in the financial industry who are permitted to collect this data only after applying to credit applications for compliance with fair lending laws, resulting in very low survey response rates. Even the simple act of retaining this data can generate enhanced risk for AEDT vendors, in that the leakage or theft of this data can result in fines or other legal penalties.

Availability of Intersectionality Supply

We applaud the guidance in that it seeks to apply concepts of intersectionality in dividing the report into combinations of race and gender. However, when one requires that we investigate bias by multiple co occurring protected categories, we are making intersectional comparisons that split each of the broader categories into smaller ones, significantly creating large differences in who is represented in the audit. These more rare intersectional identities can often bear the majority of bias and prejudice throughout their careers and experience a great deal of gatekeeping. There have already been experiences wherein our investigations have resulted in intersectional slices for minority categories for which the amount of data in each category is so limited that the results are misleading, in that they do not represent statistically significant segments. Employers and AEDT vendors in this situation are faced with a significant problem, in that releasing the numbers for these categories may imply discrimination when the evidence of discrimination is simply not there. In some cases, there may even be categories for which no data exists. In this situation, we often see vendors turn to synthetic data or globally representative datasets which may not be sufficiently connected to the candidates truly subjected to the AEDT. The science on this notion is fairly limited, yielding great uncertainty in the audit community, and paving the way for insufficient audits that fail to represent reality.

Conflicting Laws

Some good practices can conflict with each other. For example, data minimization and anonymization can protect our privacy, but also hide the very information we need to conduct an audit. Some industries (like the financial sector) use gender and race inference, or synthetic/global datasets to avoid the privacy issue, but this adds a thick layer of uncertainty to the audit that could cause some audits to under or overestimate bias, which we'll elaborate on later.

Confounding Factors (i.e. Human Selection Bias)

In defining "ground truth" it may be tempting to use some signal of approval (e.g. a candidate was hired, or a candidate was moved forward in the hiring process). However, these signals are human in nature, and therefore even more full of potential for discriminatory behavior (or the lack thereof) to obscure the behavior of the algorithm itself. An exceedingly DEI-consious hiring manager's decision-making may cancel out a highly discriminatory algorithm, and of course the inverse is also true. A deeper examination of the tool's training data, predictions, methods, structure, and behavior at scale in the context of the UX of the system can indeed illuminate this bias, but today's guidance requires none of this type of reporting, and will therefore miss opportunities for improvement.

Mathification of Subjectivity

When assessing algorithmic discrimination, it is vital to have a definition of "ground truth". In the case of hiring, this notion is quite subjective, where the definition of "good candidate to job fit" can differ from organization to organization, and even among the hiring managers within that organization. This makes the challenge of a model audit an inconsistent one, where these definitions will vary significantly by audit vendor. In short, it is entirely possible to "game the system", allowing vendors to provide audits that reflect a lack of bias where bias truly exists. The guidance in its current form does make way for one method to avoid assessing the human factor, by allowing for analysis of adverse impact by match score alone. However, later on in these comments we will detail just a few scenarios in which this simplified reporting may miss many forms of bias that remain, despite "passing" metrics. In order to assess algorithmic discrimination, a combination of quantitative and qualitative analysis is required, in order to contextualize and fully situate the impact to candidates amid the totality of the system. Candidate positioning, ranking, and display qualities matter a great deal to a candidate's likelihood of receiving an offer. In addition, there are many standardization practices that AEDT vendors can undertake to limit discrimination that can only be uncovered through an assessment of their risk and control practices. By neglecting the qualitative elements of the field of algorithmic impact, the city paves the way for these reports to be misleading, and ultimately to fail to reflect real-world discrimination where it exists.

Demographic Inference

As we've previously stated, employers may possess demographic data for their hired candidates, but the vendors who provide this technology often make active effort not to collect this vital information. As a result, these AEDT vendors often turn to methods like BISG to infer race and gender characteristics. BISG, as the most prevalent of these methods, was developed in healthcare research, and has been employed at great scale within the financial sector. However, besides concerns around accuracy, the methods themselves pose structural inequity. Race itself is a subjective attribute, and one which many have claimed can never be truly inferred. These methods also only allow for analysis on a gender binary, obscuring discrimination which may occur against others along the gender spectrum. An unintended consequence of this guidance may be the proliferation of these techniques, which have received deep scrutiny and criticism for their lack of inclusivity, and propensity for error. In fact, these error rates may in many cases be high enough to further obscure discrimination or lack thereof. If a set of candidates are improperly associated to the incorrect protected group, this may result in low enough accuracy to make the report incorrect, and therefore misleading. Additionally, common inference methods like BISG can only be effective in regions where we can assume that redlining, white flight, and gentrification have homogenized the racial makeup of the area. This seems broadly inadequate for a city as diverse as New York where there may be just as many Black John Smiths in the same zip code as there are white John Smiths. In our field, the vast consensus is that the only proper way to use demographic data in analysis is when it is volunteered from the candidates themselves. We recommend to our AEDT vendor clients that they engage in post-hoc surveys, despite our expectations that response rates will be low, because it will yield the greatest accuracy. These surveys take time, however, and in many cases the clients who have only begun this analysis in the second half of 2022 will not have adequate time to complete this initiative sufficiently prior to the release of their public reports.

AEDT Snake Oil

At Parity, we are lucky enough to have the principle and luxury of refusing to perform audits for technology that entirely fails to accurately represent accurate matching capabilities to employers. We have been approached by AEDT vendors who seek to provide technology that attempts to apply pseudo-scientific principles to the concept of job fit, and these vendors may be able to display reports that make them seem somehow fairer or more accurate than tools that avoid these forms of algorithmic phrenology. Stated another way, a tool that is always incorrectly guessing a candidate's fit for the job, or that may approve or deny every group equally along vectors of race and gender, might appear fair. However, this fairness is only trivial – an algorithm that denies everyone for every job may be fair, but it is also not useful. These pseudo-scientific algorithms present far greater danger for candidates with disabilities, but the law in its current form will fail to capture this discrimination entirely.

Inaccruate Reporting of Discrimination/Lack of Discrimination

Finally, that this guidance is so high-level allows for many opportunities to obscure discrimination, or to report discrimination where it does not exist, leading to undeserved risk and scrutiny to the vendors/employers. There are many scenarios under which discrimination may be obscured, but to name a few:

a) Imagine an algorithm that performs quite fairly on nursing jobs, but discriminates to a great degree in engineering. If these numbers are high enough in some categories to cancel out the low scores from another category, then the reports will appear to "pass" the 4/5ths rule, but discrimination will remain.
b) Imagine a situation where an employer receives a set of applicants for a position wherein all of the female candidates for a position are, in fact, highly qualified for the job, and all of the male candidates are similarly unqualified. When one set of match scores are high and the other set is low, it can appear that adverse impact exists in favor of women against men, when the resulting metric may fail to represent a lack of discriminatory bias in the tool itself, but instead a feature of the applicant base. This is a direct result of the subjectivity of quantifying "job fit" in mathematical terms.
c) It may be the case that, due to systemic inequity and historical adversity, some intersectional slices of demographic pools may receive lower match scores than others. This may not be the result of a discriminatory tool, but instead a feature of the applicant population as it exists today. Correcting for bias under these circumstances is recommended by academia, but itself may pose a form of "disparate treatment" by virtue of adding weight or altering thresholds to cater to one disadvantaged group over the other.
d) Due to the lack of demographic data, some categories may have insufficient amounts of representation to be accurately quantified, leading to numbers that skew inaccurately in a way that would not reflect discrimination that exists at scale.
e) When demographic inference is employed, the error rates may be so high as to make the resulting metrics adequately inaccurate such that they will not reflect reality, be that in the form of discrimination or a lack thereof.
f) Imagine a tool that is simply of very poor quality. This tool may be trivially fair, in that it approves or denies all candidates equally because it simply does not work. Employers

choosing vendors may be misled into thinking that the tool is worthy of use by virtue of this needed reporting, when in fact it simply "stabs in the dark" at job fit, and may present cases of individual discrimination, especially with regard to demographic categories not represented by the guidance as it stands today.

These may seem like toy examples, but from our work with clients and in research, we find them to be fairly common. Today's guidance may miss these situations, but the field of algorithmic interrogation has provided myriad tools and methods to uncover these scenarios, and we would encourage the city to pursue guidance that is more closely in line with the latest our field has to offer.

Conclusion, Recommendations, Next Steps

We'd like to reiterate our gratitude for the opportunity to provide comment. Significant questions remain on the scope of the law that would preferably be answered in advance of the compliance deadline:

1. Is a broad, global vendor analysis sufficient to each employer who uses the AEDT? Or should each employer tailor the report to the candidates they've assessed?
2. Will race/gender inference suffice for the analysis despite its possibility for decreased accuracy, and if not, what methods do you recommend when demographic information is unavailable?
3. When intersectional slices for rarer combinations of categories are present, and the amounts may not be statistically significant themselves, what sort of reporting does the city recommend?
4. Would universal or synthetic datasets suffice for the analysis, even if these datasets may not be representative of the candidates truly subjected to the system's decisions?
5. Are other forms of screening models (e.g. "culture fit, engagement probability, geographic closeness", etc.) within the scope of the law? Or is the scope limited to assessments of job-to-candidate fit?
6. Will the city consider some form of vendor certification moving forward in order to limit the ability for tools and employers to game the system by choosing unscrupulous providers?
7. Our field of algorithmic scrutiny is rapidly advancing, will the guidance make room for the advancements not currently included in the guidance, and continue to evolve with the pace of science?
8. Will the city consider extending the compliance deadline in order to provide more time to employers and vendors to begin the arduous practice of collecting demographic information from candidates?

We would be happy to further engage with DCWP in order to clarify these questions or to improve guidance for this upcoming or future years, and look forward to your feedback.

Thank you,
The Parity Team

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 10/23/2022

Dear Chair and Members of DCWP,

As the Founder of AIethicist.org and Lighthouse Career Consulting LLC, I welcome the public comment opportunity for the rules proposed by DCWP on implementation of Local Law 144 of 2021, regulating automated employment decision tools (AEDT).[1] Accordingly, I would like to 1) provide feedback on the proposal, and 2) request further clarification regarding the proposed rules for the members of DCWP to consider. 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 AIethicist.org, I provide research, training, and consulting on how to develop, use and govern algorithmic systems in a responsible way. I conduct research and training on AI policy and regulation at Center for AI & Digital Policy and teach data science ethics at University of Michigan. I submitted written comments to DCWP's previous rulemaking.[2]

I again congratulate New York City Council on the passing of this very important and impactful law. The following questions and recommendations build upon my previous written comments to DCWP and would help protect individual's rights and make the expectations from the employers and vendors clearer. In return, it would make it easier to detect the violations, and prevent different interpretations of the requirements. I also recently published two policy briefs on AI & Recruitment systems relevant to DCWP's work.

- First brief provides an in-depth analysis of NYC Law 144, including the benefits, obligations, limitations, and impact on employers, vendors, and candidates.[3]
- The second brief provides a landscape analysis of AI policy and regulations in employment decisions globally.[4] The brief reviews the city, state, federal and international across multiple jurisdictions and provides comparisons.

Definition of 'independence': The Proposal suggest Independent Auditor means "a person or group that is not involved in using or developing an AEDT that is responsible for conducting a

[1] Local Law 144 of year 2021: http://nyc.legistar1.com/nyc/attachments/c5b7616e-2b3d-41e0-a723cc25bca3c653.pdf
[2] 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/
[3] Center for AI and Digital Policy (August 11, 2022). Policy Brief – NYC Bias Audit Law. https://www.caidp.org/app/download/8407232163/AIethicist-NYC-Bias-Audit-Law-08112022.pdf
[4] 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/AIethicist-HumanResourcesRecentDevelopments-09022022.pdf

bias audit of such AEDT." This proposal suggests the auditor could be internal to the vendor or the employer. However, the wording opens some other questions.

---

1) What is meant by 'involved in using or developing an AEDT'? For example, if the vendor/employer's internal audit teams use the tool to hire for their own department, can they be independent? If the vendor's internal audit department conducts regular audits of the development teams, can they be independent? If an employer's internal auditor was involved in procurement due diligence, can they be independent?

A clearer criterion is needed to determine independence if the auditor is to be internal to either the vendor or the client.
It is equally important that 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 transparency on AEDTs when needed.

2) In the absence of any AEDT bias audit accreditation scheme, it is crucial that vendors do not decide by themselves which historical dataset to audit, conduct the audit and then decide on what to report – all in a closed circuit.

Therefore, the scope and rules of audit must be developed independent of the internal auditor conducting the audit to prevent conflict of interest and ensure integrity of audit.

Definition of "Machine learning, statistical modelling, data analytics, or artificial intelligence": Proposal defines three conditions for a mathematical/computer-based technique to be considered under this definition and proposes "and" between each of these conditions. Is it the intent of DCWP that a technique should meet all three criteria to be considered as subject to Law 144? If yes, this change significantly compromises the scope and intent.

Definition of "Simplified Output": Initial definition of AEDT in 20-870 of the Code was reflective of the intended scope of NYC Council's decision. However, narrowing the definition by using the qualifiers like "rely solely on a simplified output" or "with no other factors considered" significantly compromises the scope and intent. I welcome the "examples" provided in Proposal such as prediction, classification, score, tag, categorization, recommendation, or ranking. However, recommend DCWP not include any qualifiers and narrow the scope.

Definition of 'bias': Responsible AI practices and consideration of bias must be across practical and statistical tests, and across design decisions. The assumptions, design/model decisions and trade-offs made by vendor and employers must be included as audit criteria. For example, different cut-off scores used for outputs may change the results significantly across different groups.
Alternatively, a 'culture fit' criteria embedded in the system can discriminate against certain groups behind the veil of technological objectivity.

Proposal suggests that a simple Impact Ratio analysis can be enough to satisfy the requirement of this law. IR is just a rule of thumb and only one of the methods used to calculate disparate impact. If this is the only method suggested by DCWP, there should also be a guidance for those job categories audited where the sample size is small enough for the proposed IR calculation to be irrelevant.

Data used for analysis of Impact Ratio: Proposal, in one of the examples, suggests "vendor uses historical data it has collected from employers on applicants selected for each category to conduct a bias audit." Like above, there is need for further clarification on the dataset used for audit.

1) Vendors need to be equally responsible and transparent to ensure the data and model(s) do not result in disparate impact.

- 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 could 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 be reflective of what each employer is using. Aggregation may not correctly reflect impact ratios for different employers. This necessitates separate audits for each model in use.

For (b), solution would be for the vendor to conduct 2 separate analyses: a) Impact Ratio for all candidates in the system, regardless of clients, AND b) a separate client-based analysis. Since the Proposal defines Bias Audit as a simple Impact Ratio analysis, and since the vendors should already be monitoring client models for quality and liability purposes, this should not be an extra burden. This approach would also give more confidence to employers regarding the safeguards.

2) 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.

Notice to Candidates: The law requires employers to provide candidates minimum 10-day notice about the future use of AEDT and include the specific job qualifications and characteristics the tool will use in the assessment of candidate. This is to allow a candidate to request an alternative selection process or accommodation. To be able to request an accommodation or alternative, a candidate will need to know if the tool will require the use of and/or assess, for example, any physical, cognitive, or motor skills, or mental, emotional, character capabilities or competencies. DCWP Proposal omits this last part in the notice requirement.

Published results of the most recent bias audit: I strongly support the proposal that public AEDT audit summary results to "include the selection rates and impact ratios for all categories." This is crucial for the success and impact of Local Law 144.
As mentioned in 'Notice' section, I also propose the public audit summary to also include "characteristics the tool will use in determining its outcome." This could be the same as the notice provided to candidates during their application process. Such transparency would magnify the intent of the Law.

Thank you for your consideration of my views. I would welcome the opportunity to discuss further about these recommendations.

Merve Hickok, SHRM-SCP
Founder, AIethicist.org, and Lighthouse Career Consulting LLC
merve@lighthousecareerconsulting.com

Holistic AI questions regarding the DCWP's Proposed Rules to Implement Local Law 144 of 2021 (Automated Employment Decision Tools)

October 22, 2022

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 Risk Management 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 have requests for points of clarification regarding the Proposed Rules and their implementation. These are:

2a. Intersectionality when conducting bias audits

Although neither the Law nor the Proposed Rules explicitly mandate calculating the 'selection rates' and 'impact ratios' for different protected categories in an intersectional manner, the Bias Audit examples, and corresponding illustrative data provided by the DCWP in the Proposed Rules, provides calculations on an intersectional basis.

In the example provided in the last update, 'selection rate' and 'impact ratio' figures are provided for males and females, broken down by their race/ethnicity (e.g., Hispanic or Latino, White, Black or African American etc.).

It would be useful for the DCWP to clarify its position with respect to whether the calculations of ‘impact ratios’ and ‘selection rates’ should be performed in an intersectional manner? Is this mandatory, or is it being encouraged as best practice?

Furthermore, the DCWP should be mindful of issues relating to small sample sizes, if an intersectional approach is taken.

2b. Additional Metrics and Approaches

Impact ratios can be problematic if sample sizes are small; other metrics, like the two standard deviation rules, could be more suitable. The DCWP could clarify its position regarding the use of additional metrics to calculate bias.

Furthermore, there are potential issues regarding a lack of consideration of the distribution of scores. For example, in the presence of outliers or if the score distribution is bimodal, the mean score will not be informative. However, looking at the score distribution may provide better insights into how the tool performs for different sub-groups.

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.

- Systematizing Audit in Algorithmic Recruitment – Published at Journal of Intelligence.
- Perceived Fairness of Algorithmic Recruitment Tools – Published at Frontiers in Psychology.
- Towards Algorithm Auditing: A Survey on Managing Legal, Ethical and Technological Risks of AI, ML and Associated Algorithms – Published at IEEE Computer, and influenced public policy documents (OECD, UK Government, ICO, etc.).
- Holistic AI Open Source – Contain metrics and methods to do Bias assessment & mitigation.
- The New York City Bias Audit Law: Regulating AI and automation in HR – our white-paper related to this piece of legislation.

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.

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

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

Oct. 23, 2022

To Whom It May Concern:
On behalf of the team at BABL AI, I welcome the opportunity to provide public comments on the proposed rules for Local Law 144 requiring annual bias audits for automated employment decision tools (AEDTs).

We thank the Department for providing guidance and clarity on many of the terms provided in the original law. 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 promote and protect human flourishing in the age of AI. While the new rules directly address many of our previous concerns, we would like to comment on several areas that remain unclear and that pose barriers for companies wishing to make good-faith efforts to comply.

AEDT Definition: We find that the clarified definition of an AEDT is overly stringent. In its proposed form it allows for serious loopholes for employers to seek exemption from bias audits. To illustrate, a user of an ML-based assessment of personality traits, cognitive ability, or cultural fit whose results are shown to recruiters as part of a candidate profile could claim exemption from the bias audit, arguing that this assessment does not overrule or modify recruiter decision making given that there are other factors recruiters also consider. While this may be true for some use-cases, our experience has shown that this is only determined through a careful risk/impact assessment. Given that many such ML systems tend to contain high risk of algorithmic bias, we view their potential exclusion as risky. Therefore, **we strongly encourage the department to provide a recalibrated definition of an AEDT.**

Independence: We are concerned that this new definition of independence is not sufficiently strong to provide the level of impartiality as desired by the law. There are two reasons for our concern:

1. This definition allows for internal parties with conflict of interest for a vendor or employer (e.g., operating under the same level of management) to conduct the audit which obfuscates impartiality altogether, and
2. This definition shifts the discourse on independence as a legal dispute to a semantic debate on the terms "using" and "developing."

Moreover, the term "developing" carries a dual meaning in the AI fields:

1. A narrow scope as in "the development phase" of the AEDT where the AI or model is first built, trained, and tested prior to deployment, and
2. A broad scope where continuous design, procurement, training/testing, deployment, and monitoring are all parts of AEDT development.

As a result, the number of parties involved in *developing* depends on its meaning and scope. For example, with the emergence of companies whose primary focus is on providing monitoring services for AI solutions and whose monitoring data is looped back to AEDT vendors to improve the tool, it is ambiguous whether these companies qualify as "group involved in developing" the AEDT.

Due to these ambiguities, **we encourage the department to either consider a different approach to defining independence—such as using a precedent as in the Sarbanes–Oxley Act of 2002 or provide special clarification for the terms "using" and "developing".**

Intersectional Analysis: While conventional disparate impact analysis as carried out by the EEOC does not require analysis at the intersectional level, the example shown in the amendments disclose selection rates and impact ratios for intersections of groups (e.g., white male vs. Hispanic female). We applaud the department for taking intersectional analysis into consideration, as we believe it provides better transparency and stronger protection for protected groups.

However, the proposed rules say "[t]hese calculations are consistent with Section 1607.4 of the EEOC Uniform Guidelines on Employee Selection Procedures," which are typically not calculated on an intersectional basis. **We recommend the department to explicitly clarify whether an intersectional breakdown for selection rates and impact ratios is required**. The explicit clarification would have huge ramifications, as many AEDTs are monitored or tested prior to deployment, using the fourth-fifths rule as a reference point, but many are not controlled for *fairness* at the intersectional level.

Source of Testing Data for AEDTs: It is unclear from the proposed rules how a newly developed AEDT would be audited prior to deployment without access to historical demographic data. The two examples in the amendments show only historical data from an employer requesting for the bias audit who either has used the AEDT or plans to use the AEDT, with the assumption that such historical data exists for the tool.

This assumption is challenged by the way new AEDTs are currently developed, where many tools using ML/AI require new sources of data from candidates which can be bespoke to an employer, or whose historical data does not exist from that specific employer or from any employer. Examples of such data include: performance data of candidates playing games, audio/video/written/texting samples of candidates answering interview questions, or facial photos of candidates (e.g., in IDs, cameras, or webcams).

For many AEDTs, employers and vendors are looking to demonstrate compliance via alternative sources of data, such as using synthetic data or simply using the testing set of their training data. **We urge the department to clarify whether the use of synthetic data or training/testing data would be sufficient for employers and vendors to comply with the law.**

1. If so, we encourage the department to also provide requirements to circumvent issues where this data is not representative of the data during use—i.e., out-ofdistribution sample, data drift, or
2. If not, the department should clarify explicitly and provide guidance for newly developed AEDTs where historical data is not available.

Scope of Bias. The scope of bias audit currently only includes disparate impact or impact ratio analysis. We are concerned that this scope would only address the technical and model-centric view of bias. Our work in conducting risk and bias assessments for these automated systems have shown us that biased outcomes can often arise from a variety of sources outside of the development of the tool, such as in the way employers configure the settings of the tool or in the way direct users of the AEDTs interpret simplified outputs. Therefore, we fear that the current audit does not address these sources of bias and does little to prevent harms despite the tool itself being supposedly controlled for bias.

While it is difficult for current proposed rules to circumvent all sources of non-technical bias, we have found that risk assessment is a strong AEDT-agnostic method to uncover biases beyond the technical aspects, and allow employers and vendors to acknowledge and confront these biases head-on. Furthermore, emerging legal and industry frameworks have universally proposed risk and impact assessments as effective tools for detecting and mitigating risk, including the risk of disparate impact. Therefore, **we encourage the department to consider the inclusion of some minimal elements of risk or impact assessment in the audit to control for non-technical bias.**

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

Contact

Shea Brown, Ph.D., CEO & Founder sheabrown@bablai.com

Comment on LL 144 Proposed Rules

Submitted to DCWP | October 24, 2022

ORCAA is an independent algorithmic auditing consultancy. We have conducted algorithmic audits of hiring tools, and have had dozens of discussions related to LL 144 with vendors of AEDTs, employers that use AEDTs, and law firms. This has given us an understanding of various stakeholders' perspectives.

Depending on the rules adopted, this law could provide substantial protection and transparency for candidates, or it could permit inconclusive audits that merely continue the status quo. In any case, it will set an important precedent for similar laws elsewhere. We believe the key principles are to focus on outcomes that matter to candidates, and to consider each AEDT as it is actuallyused, which can vary across employers.

In addition to our comments below, we attach a mock Bias Audit Report we developed (see Appendix A). We believe this is a viable template and would be glad to help DCWP adopt or adapt it.

Per employer bias audits would be more e ective

The example in §5-301(a) of the proposed rules suggests a vendor can conduct a general Bias Audit of an AEDT using "historical data it has collected from employers". We argue this is not feasible in general, and even when it is, it is often insufficient to address the critical question of whether different candidates in certain protected classes have a fair chance of being hired when the AEDT is involved. Bias Audits done on a per employer level, which address an employer's particular use of an AEDT and focus on downstream outcomes rather than scores, would provide a clearer view of whether disparate impact is occurring in the use of the AEDT.

<u>Vendors may not have access to employer usage data</u>

Many vendors do not have access to the data necessary to conduct a bias audit as outlined in the proposed rules. Information about specific candidates – who they are and the hiring decisions made about them – generally belongs to the prospective employer. Some vendors may provide AEDTs that their customers run locally on their own systems. In these cases, the vendor may never have the chance to collect data on those candidates. In other cases the candidates run through the vendor's system, but the vendor's contracts with customers limit what candidate data they can access or store, or how they can use it. These data limitations imply many vendors could not produce bias audits as proposed.

<u>Most vendors do not track downstream outcomes</u>

Employment decisions, not scores, are what matters in terms of fair hiring. Even if vendors have access to some candidate data, it is likely insufficient to address what actually happened to candidates. Many AEDTs produce a score (e.g., a "culture fit" score, as in the second example table in the proposed rules). This is the AEDT output the vendor is most likely to be able to report on, and many or most bias audits based on the proposed rules would focus on score analyses.

In practice AEDT scores are often binned according to numeric thresholds chosen by the employer, and candidates are then presented with simpler bin labels (e.g. Low, Medium, or High) instead of the raw score. These labels are then one of multiple inputs to an employment decision (e.g., whether a candidate is selected for an interview).

Focusing on scores ignores what happens downstream. For one, it misses the critical issue of thresholds. Consider the second example table in the proposed rules, which shows a bias audit based on scores. If the Medium/High threshold were 80, then only three of the 14 groups (White Males, Multiracial Males, and Multiracial Females) would have High average scores. But if the threshold were 70, then half of the groups would have High average scores.

These scenarios would likely lead to substantially different hiring outcomes in terms of average selection rates by group. The first scenario is far more likely to produce a disparate impact than the second, but the proposed score analysis offers no way to capture this important distinction. Since different employers can set different thresholds, a score analysis that pooled data from multiple employers would shed even less light on the employment outcomes that really matter.

Aggregating across customers can mislead

Even if vendors had access to candidate data and downstream employment decisions, the kind of pooled analysis shown in the first example table could be misleading. It could allow discriminating employers to fly under the radar, or make blameless vendors or employers look bad.

Biases by separate employers using a single AEDT could “cancel out” in a combined analysis. We show in Appendix B how this could happen with the first example table from the proposed rules. The example table exhibits gender parity: 48.0% of Males are selected versus 47.0% of Females, for an Impact Ratio of 0.979. Now we imagine the example table was based on data from two firms, A and B, whose corresponding tables are shown separately. Firm A favors Male candidates (51.7% of Males selected versus 40.9% of Females; Impact Ratio = 0.792) and Firm B favors Female candidates (53.1% of Females selected versus 44.4% of Males, Impact Ratio = 0.837). But these significant disparities disappear in the combined analysis.

Similarly, consider two firms C and D using the same AEDT, where Firm C considers 10,000 candidates and Firm D considers just 100 candidates. Suppose Firm C strongly favors Male candidates and Firm D has perfect gender parity. The combined analysis will suggest that the AEDT favors men, simply because Firm C’s data dominates. This result is unfair to Firm D and, arguably, to the AEDT vendor (after all, Firm D was able to use the AEDT without a gender disparity).

The proposed rules do not address applicant sourcing

The pool of applicants an employer reviews, as well as the method the employer uses to generate, or “source”, the pool, can contribute to bias in the hiring process. For example, suppose an employer manages to (unfairly) target its recruiting so that it attracts an overwhelmingly

White applicant pool. These applicants then pass through a resume-filtering AEDT, and are screened out at comparable rates across race groups. The result will be an overwhelmingly White set of hires. Given the proposed definition of a Bias Audit, this would not be seen as a problem. Conversely, when the labor pool for a given job is unbalanced, Bias Audits that address sourcing would effectively incorporate this context. For instance, if 80% of security guards are male, then we should not be surprised if applicant pools for security guard jobs skew male.

Bias Audits could address the sourcing issue by including a demographic breakdown of the applicant pool considered by the AEDT, alongside a similar breakdown of a relevant comparison population. Our mock report (see Appendix A) includes an example. The Bureau of Labor and Statistics would be a natural source for comparison data at a national level, and more local statistics might be available.

The proposed rules do not address sample size issues

Some samples are too small to do robust statistical analysis. The example tables in the proposed rules show intersectional analyses (each unique combination of gender and race/ethnicity is a separate group); this leads to 14 groups. The first example table shows 24 Native American / Alaska Native males, and 17 Native American / Alaska Native females. These samples are simply too small to trust the estimates of selection rates for these groups, or the corresponding Impact Ratios.[1] One alternative would be to allow (or prescribe) coarser analyses when there is little data. For instance, below a certain size threshold it might make sense to analyze gender and race/ethnicity separately rather than intersectionally, or to condense race/ethnicity categories. Offering guidance on how to deal with small samples would mitigate the risk of Bias Audits containing unreliable statistics.

Appendix A: Mock Bias Audit Report

Note: The below version is current as of October 20, 2022. Future updates will be reflected in the live version of the mock report, hosted on our website.

[1] Consider the statistics "rule of thumb" based on the central limit theorem, that at least 30 data points should be used to calculate a population average.

Bias Audit Report: NewCo's Use of ToolX

Mock Report | October 2022

Analyses Performed

We investigate potential bias with three sets of analyses. Analyses may be performed separately per job or job type being hired, since different jobs may have different applicant pools and/or hiring processes.

1. Adverse impact ratio analysis. We calculate selection rates and adverse impact ratios per intersectional group (gender and race/ethnicity), following DCWP's proposed rules.
2. Simple disparate impact analysis. Using applicant data, we measure whether there are differences between demographic groups in employment decisions.
3. Disparate impact with controls analysis. Group differences in employment decisions may be acceptable if they can be "legitimately explained," a legal term of art. For instance, if female applicants are generally better qualified than male applicants, then a higher selection rate for women may be acceptable. This analysis includes controls for characteristics the client deems justifiable.
4. Sourcing analysis. How do the demographics of the candidate pool for this position compare to the nationwide demographics of this job, per the Bureau of Labor and Statistics?

Audit Scope

Client: NewCo

Automated Employment Decision Tool: ToolX is used by NewCo to help determine which candidates are invited for an interview. After submitting an application, a candidate is given the ToolX survey. Their responses are assessed by ToolX, which gives a recommendation: "Interview" or "Screen". This recommendation is shown to a NewCo hiring manager, who can follow it or not, at their discretion. Hiring managers generally, but not always, follow the recommendations.

Employment decision: Whether the candidate was selected for an interview (Y/N)

Variables that may "legitimately explain" group differences: Years of experience

Protected classes addressed: Race/ethnicity. This is inferred for each candidate using the BIFSG method.*

Scope of data:

Job title	# of candidates	Starting from	Until
Retail salesperson	698	1/1/2021	12/31/2021
Engineer	460	5/1/2021	12/31/2021

Notes:
- Prior to 5/1/2021 NewCo hired Engineers using a traditional hiring process, without ToolX

* For documentation on inference methods, see [Pilot url]

1. Adverse Impact Ratio Analysis

		# of applicants	# selected for interview	Selection rate	Adverse impact ratio
Male	Hispanic	83	69	83.1%	0.881
	White	223	202	90.6%	0.960
	Black	140	104	74.3%	0.787
	API	66	59	89.4%	0.947
	AI/AN	14	10	71.4%	0.757
	Two or more	51	43	84.3%	0.893
Female	Hispanic	95	84	88.4%	0.937
	White	212	192	90.6%	0.960
	Black	130	100	76.9%	0.815
	API	71	67	94.4%	1.000
	AI/AN	9	6	66.7%	0.706
	Two or more	64	48	75.0%	0.795

This table calculates selection rates and adverse impact ratios for each intersectional group (gender * race/ethnicity), following the sample bias audit table in DCWP's proposed rules for the law.

We note that the proposed rules do not include thresholds of acceptability for selection rates or the adverse impact ratios, or for the number of candidates that constitute a sufficient sample for statistical analysis.

2. Simple Disparate Impact Analysis

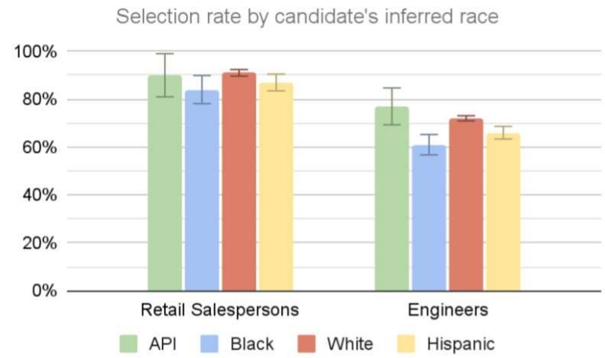

This bar chart shows the selection rates of candidates by their inferred race. For each position NewCo hired, there is a separate set of bars. The "whiskers" at the top of each bar show a 95% confidence interval of the true selection rate for that group.

In this case "selection" means being invited for an in-person interview, since the automated employment decision tool being audited is a pre-interview survey.

For retail salespersons, White and API candidates have higher selection rates than Black or Hispanic candidates, but the overlapping confidence intervals suggest the differences are not significant.

For engineers, Black and Hispanic candidates have lower selection rates and the confidence intervals do not overlap, suggesting the difference is statistically significant.

ORCAA | PILOT | 5

3. Disparate Impact with Controls: Retail Salesperson

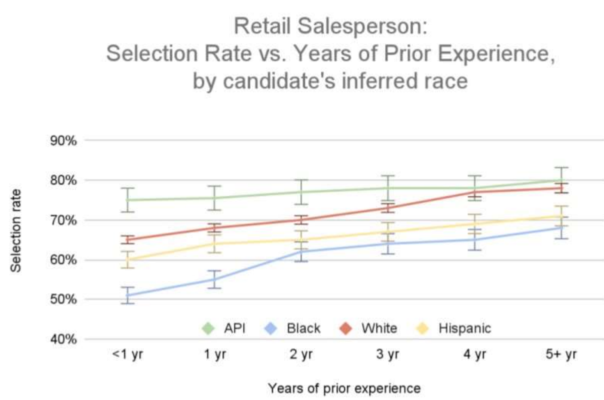

This graph shows the selection rate of retail salesperson candidates according to their years of prior experience. Each inferred race group of candidates is shown as a separate line. The "whiskers" around each point of each line are 95% confidence intervals.

For each inferred race group, selection rate increases with years of prior experience, which makes sense. However there are persistent differences: at every level of prior experience, API candidates are most likely to be selected, followed by White, then Hispanic, and finally Black candidates. Moreover, the confidence intervals for Hispanic and Black candidates are strictly below those of API and White candidates, so the differences here are likely to be statistically significant.

ORCAA | PILOT | 6

3. Disparate Impact with Controls: Engineer

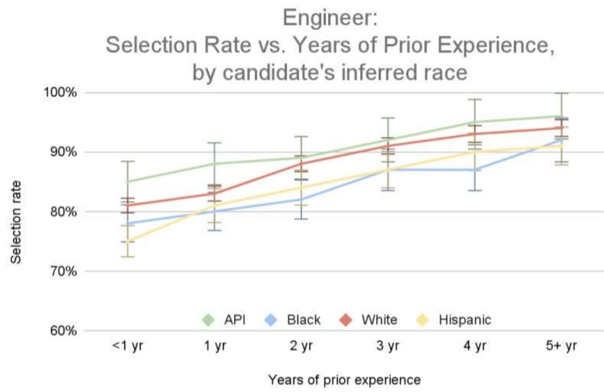

This graph shows the selection rate of engineer candidates according to their years of prior experience. Each inferred race group of candidates is shown as a separate line. The "whiskers" around each point of each line are 95% confidence intervals.

For each inferred race group, selection rate increases with years of prior experience, which makes sense. However there are persistent differences: for every level of prior experience, API candidates are most likely to be selected, followed by White candidates, with Black and Hispanic candidates less likely. There is considerable overlap between the confidence intervals (no line is entirely "on its own" above or below the others), so the differences are unlikely to be highly statistically significant.

ORCAA | PILOT | 7

4. Sourcing Analysis: Retail salesperson

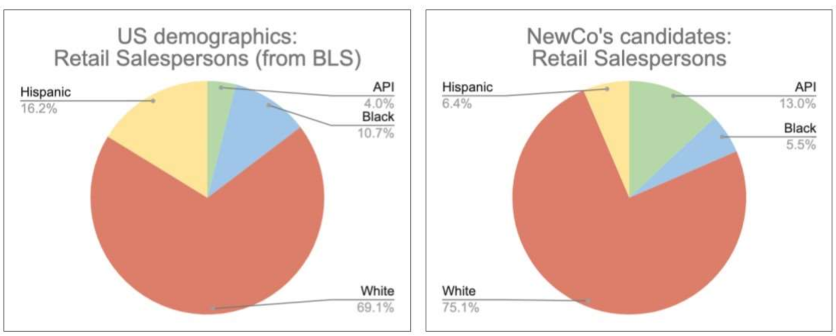

The pie chart on the left shows the demographic breakdown of all US retail salespersons, per BLS' Labor force characteristics by race and ethnicity, 2020 (most recent available; "Retail salespersons" category).

The pie chart on the right shows the demographic breakdown of NewCo's candidates for retail salesperson jobs, modeled using inferred race.

Compared to all US retail salespersons, the main differences are that NewCo's candidate pool is more white (75% vs 69%) and API (13% vs 4%), and less Hispanic (6% vs 16%) and Black (6% vs 11%).

ORCAA | PILOT | 8

4. Sourcing Analysis: Engineers

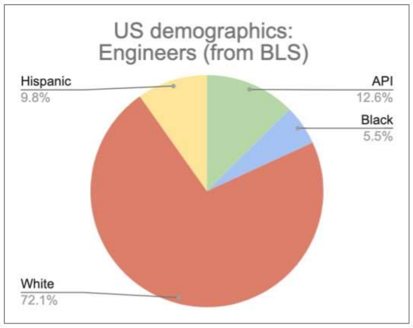

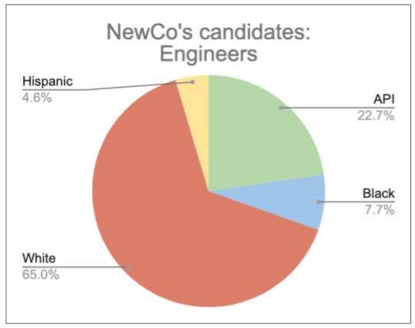

The pie chart on the left shows the demographic breakdown of all US engineers, per BLS' Labor force characteristics by race and ethnicity, 2020 (most recent available; "Architecture and engineering occupations" category).

The pie chart on the right shows the demographic breakdown of NewCo's candidates for engineer jobs, modeled using inferred race.

Compared to all US engineers, the main differences are that NewCo's candidate pool is more API (23% vs 13%), and less White (65% vs 72%).

Appendix B: Combining employer data could hide disparities

Combining the applicant data from (hypothetical) Firm A's and Firm B's tables below would produce the example Bias Audit table shown in §5-301(a).

Example Bias Audit table shown in §5-301(a)

			# of applicants	# selected	Selection Rate	Impact Ratio	Male selection rate
Hispanic or Latino		Male	205	90	43.9%	0.841	48.0%
		Female	190	82	43.2%	0.827	
Non/Hispanic or Latino	Male	White	412	215	52.2%	1.000	Female selection rate
		Black or African American	226	95	42.0%	0.806	47.0%
		Native Hawaiian or Pacific Islander	87	37	42.5%	0.815	
		Asian	321	167	52.0%	0.997	Impact ratio
		Native American or Alaska Native	24	11	45.8%	0.878	0.979
		Two or More Races	115	52	45.2%	0.866	
	Female	White	385	197	51.2%	0.981	
		Black or African American	164	75	45.7%	0.876	
		Native Hawaiian or Pacific Islander	32	15	46.9%	0.898	
		Asian	295	135	45.8%	0.877	
		Native American or Alaska Native	17	7	41.2%	0.789	
		Two or More Races	98	44	44.9%	0.860	

Firm A's table

FIRM A			# of applicants	# selected	Selection Rate	Impact Ratio		Male selection rate
Hispanic or Latino		Male	100	50	50.0%	0.909		51.7%
		Female	95	42	44.2%	0.804		
Non/Hispanic or Latino	Male	White	200	110	55.0%	1.000		Female selection rate
		Black or African American	113	55	48.7%	0.885		40.9%
		Native Hawaiian or Pacific Islander	44	21	47.7%	0.868		
		Asian	160	85	53.1%	0.966		Impact ratio
		Native American or Alaska Native	12	6	50.0%	0.909		0.792
		Two or More Races	58	28	48.3%	0.878		
	Female	White	191	80	41.9%	0.762		
		Black or African American	81	32	39.5%	0.718		
		Native Hawaiian or Pacific Islander	15	6	40.0%	0.727		
		Asian	150	59	39.3%	0.715		
		Native American or Alaska Native	9	4	44.4%	0.808		
		Two or More Races	50	19	38.0%	0.691		

Firm B's table

FIRM B			# of applicants	# selected	Selection Rate	Impact Ratio		Male selection rate
Hispanic or Latino		Male	105	40	38.1%	0.632		44.4%

		Female	95	40	42.1%	0.698		
Non/Hispanic or Latino	Male	White	212	105	49.5%	0.821		Female selection rate
		Black or African American	113	40	35.4%	0.587		53.1%
		Native Hawaiian or Pacific Islander	43	16	37.2%	0.617		
		Asian	161	82	50.9%	0.845		Impact ratio
		Native American or Alaska Native	12	5	41.7%	0.691		0.837
		Two or More Races	57	24	42.1%	0.698		
	Female	White	194	117	60.3%	1.000		
		Black or African American	83	43	51.8%	0.859		
		Native Hawaiian or Pacific Islander	17	9	52.9%	0.878		
		Asian	145	76	52.4%	0.869		
		Native American or Alaska Native	8	3	37.5%	0.622		
		Two or More Races	48	25	52.1%	0.864		

Prepared by ORCAA |

Seyfarth Shaw LLP
620 Eighth Avenue
New York, New York 10018
T (212) 218-5500
F (212) 218-5526

rszyba@seyfarth.com T

(212) 218-3351

www.seyfarth.com

October 24, 2022
VIA E-MAIL

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 Proposed Rules to the Use of Automated Employment Decision Tools Under Local Law 144 of 2021

Dear Commissioner Mayuga:

Seyfarth Shaw LLP ("Seyfarth") is a full-service law firm with 17 offices in the U.S. and globally, including a strong presence in New York City. Seyfarth is recognized as one of the "go to" law firms for labor and employment law representation and counsel in the world. Seyfarth attorneys advise clients on nearly every type of issue that arises for employers, from day-to-day workplace matters to the most complex and sensitive situations, including issues that are on the forefront of technology, social issues, and legal trends.

Seyfarth has a deep institutional knowledge of the labor and employment matters that affect employers in this economy and brings this expertise to bear in its representation of clients across a vast range of industries in New York City, the United States, and globally. Seyfarth has a substantial People Analytics practice group that counsels employers on federal, state and local employment law requirements as it relates to data driven decision-making in the workplace. These technologies include the use of artificial intelligence, predictive analytics, and machine learning algorithms to address key processes including sourcing, hiring, retention, workforce planning, employee engagement, and diversity and inclusion efforts. Of particular relevance here, we counsel employers and vendors with regard to the legal implications of developing and implementing sophisticated algorithmic technologies in the workplace.
We submit these comments in relation to the proposed rules issued by the De
partment of Consumer and Worker Protection ("DCWP" or "Department") on September 23, 2022 that stand to implement New York City Local Law 144 of 2021 ("LL 144") regarding automated employment decision tools ("AEDT").

Based on our experience with employers of all sizes, across industries, and with a wide range of business circumstances, clarity in regulation is key for fostering compliance. Thus, by ensuring that the employer community has definitive guidance on which of their procedures or tools falls within the definition of AEDT, the Department ensures maximum adherence to the goals and protections for New York City's employees that appear to underpin the basis for LL 144. In doing so, the Department may consider that saddling New York City's employers with rules, restrictions, and obligations that impede and stifle the development of technology stand to keep New York City employers ***and*** employees at the back of the pack in the global race for innovation.

1. The Definition of AEDT Should Clearly Identify the Specific "Computational Process[es]" That Will Require a Bias Audit

The definition of AEDT, as defined in LL 144 and further explained in the Department's proposed rules, remains ambiguous as to the types of computational processes are covered by this law.

Presently, LL144 § 20-870 defines AEDT 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. The term "automated employment decision tool" does not include a tool that does not automate, support, substantially assist or replace discretionary decisionmaking processes and that does not materially impact natural persons, including, but not limited to, a junk email filter, firewall, antivirus software, calculator, spreadsheet, database, data set, or other compilation of data.

Section 5-300 of the proposed rules adds that "the phrase 'to substantially assist or replace discretionary decision making'" as used in LL 144 § 20-870, means to: (1) "rely solely on a simplified output (score, tag, classification, ranking, etc.), with no other factors considered," or (2) "to use a simplified output as one of a set of criteria where the output is weighed 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."

In addition, the proposed rules define "machine learning, statistical modeling, data analytics, or artificial intelligence" to mean 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.

Further, "simplified output," as used in LL 144 § 20-870,

> means a prediction or classification as specified in the definition for "machine learning, statistical modeling, data analytics, or artificial intelligence." A simplified output may take the form of a score (e.g., rating a candidate's estimated technical skills), tag or categorization (e.g., categorizing a candidate's resume based on keywords, assigning a skill or trait to a candidate), recommendation (e.g., whether a candidate should be given an interview), or ranking (e.g., arranging a list of candidates based on how well their cover letters match the job description). It 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.

The definitions and explanations above discuss generalized concepts and use broad terminology, making it very difficult for the employer community to know with certainty which "computational process[es]" are covered, and which are not covered. Presumably, every recruitment, selection, and onboarding tool cannot be covered, as doing so would inflict crippling costs upon employers that would risk non-compliance, as well as abandonment of technology that would result in severe delays and administrative burdens associated with reversion to manual processes and tools that would dramatically impede business operations and frustrate recruitment and hiring efforts in New York City. Such definition and outcome helps no one. Presuming further that this law is the result of the City Council or Department seeking to address something that it considered to be a specific issue, problem, or danger, the City Council and Department are in the optimal position to identify, with specificity, the "computational process[es]" that are intended to be covered by the law. Leaving employers to guess, however, especially under threat of potentially business-ending penalties, needlessly creates peril and risk instead of fostering compliance.

As one example, certain employers might utilize tools that might be applied to candidates who have not applied for a specific position. There are a variety of automated tools that allow employers to find and reach out to candidates who might be interested in or qualified for a certain position but who have not yet applied for such position. For example, employers or third-party entities may have collected resumes or profiles and may use a search tool to find appropriate candidates for a position from such collection. Likewise, employers may rely on professional databases to conduct passive searches for candidates. By defining "candidates for employment" as "a person who has applied for a specific employment position by submitting the necessary information and/or items in the format required by the employer or employment agency," it appears that such resume/profile search tools would not be considered an AEDT because such tools are not used to screen anyone who has completed an application or applied for a specific position. The employer community would benefit from further clarification either in the definition of "AEDT" or elsewhere in the rules to confirm that tools applied to individuals who are neither (1) employees nor (2) persons who have applied for a specific position" are not covered by § 20-870. Specifically, § 5-301(a) should be revised to strike the undefined term "individuals" and replace it with "candidates for employment." Likewise, § 5-301(b) should delete the term "applicants" and replace it with "candidates for employment."

As another example, certain employers, potentially utilizing a vendor or consumer reporting agency, may perform a background check (non-criminal and/or criminal) on candidates and employees, and as part of that process may utilize "computer-based techniques" that "generate a prediction … or classification" and "for which a computer … identifies the inputs, that relative importance placed on those inputs … in order to improve the accuracy of the prediction or classification," "for which the inputs and parameters are refined through cross-validation or by using training and testing data." Presumably, LL 144 was not intended to regulate such consumer reporting agencies or to regulate background checks that are performed by employers, however the text does not make this adequately clear.

These examples, and likely a plethora of others, highlight the need for clarification and further guidance as to the specific "computational process[es]" that fall within the definition of AEDT and are thus regulated by LL 144.

2. <u>The Proposed Rules Do Not Provide Sufficient Guidance Regarding the Scope of Bias Audit</u>

Employers need further clarity regarding the scope of the candidate pool that is required to be included in a bias audit. In particular, it is not clear whether such audits may include information regarding candidates outside of NYC or conversely based on a sample set of data. Employers should be expressly permitted to rely on robust data that is not limited to NYC candidates for employment or employees.

To the extent the regulations do not allow employers to use broad data or sample data for purposes of the required bias audit, NYC employers will be stifled in any attempt to implement new AEDTs. That is because NYC employers would have no prior data upon which to perform a bias audit on the AEDT. This would have the unintended consequence of effectively barring the use of new AEDT tools for NYC employers. This issue could be addressed by explicitly providing that employers can rely on bias audits that rely on broad data sets or even sample test data to meet the bias audit requirements under LL 144.

3. <u>The Proposed Impact Ratio Methodology Based on Average Scores is Flawed</u>

The Impact Ratio metric that relies simply on "average scores" is flawed in that it does not adequately consider variability in scores. Based on our discussions with labor economists and I/O psychologists, the proposed methodology does not provide the requisite insight needed to make any inferences about whether there is "bias" in the AEDT scores.

The use of the "average scores" impact ratio analysis will likely be required in all AEDT bias audits. While the proposed rules put forth two impact ratio methodologies,[68] it is difficult to conceive of an AEDT that "selects individuals in the hiring process or classifies individuals into groups" as set forth in § 5-301(a), without scoring "candidates for employment" or employees in some way as described in § 5-301 (b).

Given the concerns with the proposed methodology, we strongly encourage the Department to work with I/O psychologists, labor economists, and statisticians who have deep expertise with employment decision-making tools to evaluate and assess the viability and reliability of the proposed impact ratio methodologies set forth in the proposed rules.

4. The Proposed Rules Do Not Provide Guidance On Conducting a Bias Audit on an AEDT that Does Not Capture or Retain Data on Gender, Race, or Ethnicity

Not every "computational process" utilized by employers captures or retains data about every "candidate for employment's" gender, race, or ethnicity. Section 5-301 of the proposed rules, expanding on § 20-870 and -871, requires that a bias audit, "at a minimum," calculate selection rates and impact ratios based gender, race, and ethnicity, presumably on candidates for employment.

Considering that (1) LL 144 is scheduled to go into effect January 1, 2023, (2) employers who utilize "computational process[es]" that presently do not capture gender, race, and ethnicity data on every "candidate for employment" have insufficient opportunity to obtain such data between now and the effective date upon which to conduct a bias audit, and (3) once § 20-871(a) goes into effect, employers are prohibited from "us[ing] an [AEDT] … unless" the AEDT "has been the subject of a bias audit conducted no more than one year prior to the use of such [AEDT]." Thus, it follows that such employers will be effectively barred from using such AEDT as of January 1, 2023, and will further be barred from obtaining the data necessary to conduct the bias audit on the AEDT, thereby making the bar permanent.

Presumably, this law was not intended to force employers to scrap potentially useful and universally beneficial "computational process[es]" simple because developers did not fully predict LL 144's bias audit requirement, which is not yet in final form. Moreover, there is no requirement in NYC to attempt to collect gender, race and ethnicity data on "candidates for employment."

Thus, clarification and guidance to address this issue is warranted to assist the employer community with compliance with these requirements.

[68] Impact Ratio. "Impact ratio" means either (1) the selection rate for a category divided by the selection rate of the most selected category *or* (2) the average score of all individuals in a category divided by the average score of individuals in the highest scoring category.

5. **Effective Date and Enforcement Needs to Be Deferred to Allow Employers Sufficient Time To Comply With the Rules, Once They Are Finalized**

LL 144 is scheduled to "take[] effect on January 1, 2023." LL 144 § 2. Presently, final rules and guidance has not been promulgated by the Department to clarify the requirements, standards, or criteria for compliance with LL 144, and the effective date is approximately two months away. For example, LL 144 requires a "bias audit," however there remains no reliable clarification of what criteria must be met to complete a bias audit that complies with LL 144. Conducting the required bias audits will be a significant undertaking and employers simply do not have the requisite clarification as noted in these comments.

There are very few organizations that currently employ the appropriate staffing and expertise to complete such bias audits. With proposed regulations providing guidance on how to perform the bias audit and final regulations still unknown, it is very unlikely that any employer will have sufficient time to complete a satisfactory bias audit of their AEDT's prior to January 1, 2023. Thus, we recommend that any enforcement of LL 144 be deferred until a date that is at least 90 follow the publication of final regulations by the Department. Failure to postpone enforcement may lead employers, where possible, to avoid considering for employment any residents of New York
City.

In closing, we appreciate the opportunity to submit comments and feedback to the Department regarding the proposed rules. Recognizing that LL 144 pushes New York City law and the Department's regulation into novel and uncharted territory, we welcome any opportunity to collaborate or otherwise provide feedback to the Department to assist in its efforts for the benefit of all New Yorkers, employers and employees alike. The Department should not hesitate to contact the undersigned.

We thank the Department for its time and attention to our submission.

> *If you have an apple and I have an apple and we exchange these apples then you and I will still each have one apple. But if you have an idea and I have an idea and we exchange these ideas, then each of us will have two ideas.*
>
> — George Bernard Shaw[69]

Respectfully,

[69] George Bernard Shaw was a playwright, critic, polemicist, and political activist, and in 1925 was awarded the Nobel Prize in Literature. Mr. Shaw has no affiliation or connection with Seyfarth Shaw LLP.

SEYFARTH SHAW LLP

s/ Robert T. Szyba

Robert T. Szyba

Partner *s/*

Annette Tyman

Annette Tyman
Partner and Co-Chair People Analytics Practice Group

October 24, 2022

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 New York City Local Law 144 of 2021

To Whom It May Concern:

HireVue is pleased to submit comments to New York City Local Law 144 of 2021 and the proposed Rules issued on September 23, 2022 ("the Law") regarding the use of automated employment decisions tools, including the use of Artificial Intelligence ("AI") in hiring.

HireVue is a global software technology provider that has hosted more than 30 million interviews and 200 million chat-based candidate engagements for over 800 customers around the globe. At HireVue our mission is to change lives by connecting talent and opportunity. Our approach is to use science, technology and best practice to enable a hiring process that is more fair, inclusive, and equitable, and that allows candidates to showcase their potential.

HireVue technologies support both the candidate and employer interviewing experience in and around the New York City area, in a broad spectrum of industries. HireVue's comments on the Law are based on our extensive experience with the use of AI in the context of hiring.

HireVue knows that human decision making in hiring is not immune to bias - ranging from the existence of overt prejudice to subconscious bias. HireVue's commitment to mitigating human bias in hiring has been rooted from the beginning in its use of scientifically-validated methods implemented by our team of advanced degree industrial and organizational (IO) psychologists and data scientists. For example, HireVue's interviewing technology (with or without AI) uses standardized and structured interviews, i.e., the same set of questions in the same order for every candidate applying for a given role. In addition to being a best practice approach in candidate selection, this greatly enhances the candidate experience and ensures each candidate has the same opportunity to demonstrate their competencies for the job.

When HireVue first began developing, designing, and testing our algorithmic assessment solutions, we referenced over a century of study of scientifically-validated methodologies, decades of experience, and well-established Federal guidelines, best practices, and employment laws. HireVue's algorithmic technology evaluates candidates' job-related skills, without considering criteria often used by humans that are unrelated to job success, such as how well dressed one is, which university one attended, which keywords were packed into their application, demographic information, or educational background. Due to our focus on only knowledge and skills relevant to success in a specific job role, all candidates are evaluated in a consistent and objective manner

using HireVue's technology. Furthermore, we develop our technology to minimize bias at every design step - from sampling representative training data, including diverse subject matter experts in the development of our interview questions, to using state-of-the-art bias mitigation algorithms to mitigate bias from the technology's predictions during the design and development stage, before any applicant is assessed.

HireVue's ongoing commitments to transparency, accountability, continual improvement and engaging with stakeholders drives more equitable hiring to offer a faster, fairer and more flexible experience for candidates and employers alike. HireVue has accomplished this through its attention to the following key principles, underpinned by best practices:

1. **Open and transparent approach:** HireVue believes the only thing that should be evaluated during the recruitment process are a candidates job-related competencies in relation to knowledge, skills, and abilities associated with the particular job they are being considered for. Our tools and services use deep industrial and organizational (I/O) psychology expertise and science-backed methods to increase the speed at which employers can screen and hire candidates and which candidates can display their skills. At its core, HireVue's standardized and structured interview (with or without AI) reflects scientifically validated best practices. When structured interviews are used in conjunction with our AI-driven assessments, they are designed to evaluate candidates' job-related skills, without taking into account criteria proven *not* to determine job success, such as demographic information or educational background. Our AI-driven assessments do not replace humans - HireVue simply helps recruiters and talent acquisition teams do a better job, faster by focusing their efforts on candidates exhibiting job-related competencies. Most recently, rather than ignore what has been described as the "black box" nature of AI, HireVue published our industry's first AI Explainability Statement to provide transparency into the methods of designing, developing and training our algorithms.[70]

2. **Continuous product improvement**: HireVue tracks innovative research, in addition to conducting its own academic research and applies advancements to ensure our technology continually meets or exceeds industry best practices. As evidence of our constant drive for improvement, in 2020 we made the decision to remove all video inputs from our algorithms in favor of solely using Natural Language Processing (NLP) after research and stakeholder feedback indicated that improvements in the NLP technology provided more predictive power and that the non-verbal data features did not add incremental accuracy to our assessments. More recently, a member of HireVue's team published ground-breaking

[70] https://www.hirevue.com/why-hirevue/ai-explainability-statement

research[71] that found no differences in HireVue game-based assessment results of candidates on the autism spectrum vs. the general population, whereby traditional human hiring approaches might overlook or, worse, reject these candidates. In addition, HireVue's core development process includes bias mitigation, testing and feedback at the design stage and ongoing review at the deployment stage. This is in addition to the fact that our AI is locked when deployed and not dynamically learning with un-curated data. Lastly, HireVue has also voluntarily submitted its design methods for review through the first-of-their-kind independent 3rd party audits and took the additional step to make the results of those audits publicly available. Our work with O'Neil Risk Consulting and Algorithmic Auditing (ORCAA) on an algorithmic audit is available on our website[72] and Landers Science Workforce LLC's evaluation of our application of IO Psychology to candidate selection can also be reviewed online.[73]

3. **Engaging with our stakeholders:** Beyond leveraging the diverse expertise of a multidisciplinary workforce to design and build our products, we regularly engage with various stakeholders to solicit input and feedback on our research, product development, and bias-mitigation practices. HireVue actively engages with its candidates, its customers, and with third-party organizations like Integrate Autism Employment Advisors, representing neurodiverse candidates to learn about their experience in using our technology. In addition, HireVue's expert advisory board consists of outside experts in relevant fields of IO psychology, legal, and AI ethics. Lastly, in 2019, HireVue established ethical AI principles and we continue to engage with stakeholders through our advisory boards to understand the challenges of both public and private employers.

These guiding principles have allowed HireVue to establish its industry-leading best practices around AI and uphold principles of minimizing discrimination in hiring. In alignment with the foregoing principles, HireVue supports well-designed legislation to address the risks of AI if it is not designed and deployed responsibly. HireVue gleaned significant learnings after subjecting our technology to third-party independent audits and is eager to share its insights on how to craft these in a meaningful way to accomplish the goals the Law is seeking to meet.

Specifically with respect to the Law, HireVue questions severe restrictions on good practices and beneficial innovation that pragmatically broaden access to employment in the name of attempting to mitigate algorithmic bias. The Law's "one size fits all" approach, i.e., failing to differentiate use cases, industries, type of AI (learning vs static), or *what* can reasonably or feasibly be audited in each development stage of an algorithm (design vs. deployment stage) will not necessarily effectuate what appears to be the intent of the Law.

We also call to the attention of the agency that some requirements of the Law may not be reasonably satisfied in the employment context, such as seemingly requiring an audit of the AEDT at the deployment stage when no live-candidate data would be available to audit prior to

[71] https://www.hirevue.com/blog/hiring/hiring-in-a-neurodiverse-world-autism-and-the-science-of-game-basedassessments

[72] https://www.hirevue.com/blog/hiring/industry-leadership-new-audit-results-and-decision-on-visual-analysis

[73] https://www.hirevue.com/blog/hiring/independent-audit-affirms-the-scientific-foundation-of-hirevue-assessments

use, running analysis on data (race, gender, marital status, etc.) that an employer cannot require due to pre-existing employment laws and services providers may not naturally have access to due to data minimization concepts in various privacy laws.

Lastly, with respect to any algorithm audit requirement, a thoughtful audit and any legislation which requires it should leave room for continuous development and improvement to build on the "good" and to identify *and* address concerning results. Based on its experience HireVue offers the following points for the council's consideration with respect to algorithmic audits:

1. **Audit criteria must be clearly defined.** Much like audit standards in industries like privacy, finance, etc., an audit of an AI tool should include reference to the relevant industry and legal standards against which the tool is tested. It should explain how a model works, its purposes and its limitations, and the data it relies upon to make decisions. HireVue follows, and in many ways exceeds, previously established employment frameworks such as the EEOC's Uniform Guidelines on Employee Selection Procedures (1978)[74] to statistically test and mitigate for adverse impact in our lab models during the design and development phase. We also perform adverse impact checks (at a minimum annually) for deployed models and adjust to mitigating bias if it is observed in use with a particular population. The most recent guidance for the Law seems to be at odds with wellestablished best practices regarding the process and methods of how to perform a disparate impact analysis and fails to require a specific sample size threshold to ensure statistical validity. In addition, the requirement to publish results does not come with any consideration as to the intended audience; most laymen do not have sufficient knowledge to interpret statistical tables and do not know what meets federal guidelines or is legally considered "fair" in the context of employment law.

2. **Audit at the design and development phase.** The focus of the audit should be at the outset of product development to ensure algorithmic tools are designed, developed, trained and tested – including steps to identify and mitigate bias – *before* deployment. The algorithm should also be periodically monitored after deployment to identify any unexpected results. While the Law expressly applies to employers, the rulemaking creates confusion on what obligations a service providers has and fails to distinguish between audits performed at the development vs deployment stage, particularly the AEDT has not been previously been used and there are no selection rates for employers to audit.

3. **HireVue, like other vendors, should be responsible for delivering an independent audit on the AI-based products we provide to our customers.** Vendors will differ in how their tools are developed and what sort of data they use, thus the way audits are conducted will not be universal. Audits must always consider the industry and context of where and how the AI is being used. In HireVue's case, it is against the EEOC Guidelines to deploy algorithms that treat job applicants differently from "day-to-day", thus we have chosen to only deploy static algorithmic models (after auditing them and testing against

[74] https://www.eeoc.gov/laws/guidance/employment-tests-and-selection-procedures

established frameworks in hiring). This means our algorithms are "locked" and do *not* learn or continually change from real-time uncurated and unfiltered customer data - as this would be unfair to the job applicants. This approach prevents the risk of bias unexpectedly creeping into our pre-tested models. We also suggest that Customers using HireVue's audited algorithmic tool should be able to rely on those audit results without needing to

conduct a third-party audit of the deployed model–though the model should continue to be periodically monitored to validate the deployment and use of an AI tool in their particular setting, e.g. with live candidate data that the employer would usually (but not always) collect. This distinction and concept fail to be captured in the current Law.

Ongoing dialogue between appropriate stakeholders is the key to creating sensible legislation that protects candidates, companies, and innovation. HireVue welcomes legislation related to the potential impact of all hiring technologies as it encourages transparency and improves the process and fairness for all candidates. Legislation like this demands that all vendors meet the high standards that we have supported since the beginning.

Should the Council have questions about this letter or our business, HireVue is absolutely ready and willing to engage in a dialogue with the Council and happy to share our perspective as the leader in responsible hiring solutions.

Sincerely,

DocuSigned by:
Anthony Reynolds
D6AFFEDB75244ED...

Anthony Reynolds
Chief Executive Officer

DocuSigned by:
Lindsey Zuloaga
5AE535A605244EF...

Lindsey Zuloaga
Chief Data Scientist

DocuSigned by:
Nathan Mondragon
489CA36B1E1D486...

Nathan Mondragon
Chief I/O Psychologist

October 24, 2022

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 Decisionmaking Tools* (Ref. No. DCWP-21; "The Rules"). I am the owner of Responsible Artificial Intelligence LLC, a New York company offering algorithmic auditing and consultancy services. Previously, I was a Founder and Chief Technology Officer of Parity Technologies, Inc., a startup dedicated to modernizing model risk management and AI compliance; and also a Director of AI Research at JPMorgan Chase & Co., and a Senior Manager of Data Science at Capital One, where I started R&D teams for responsible AI and its applications to financial regulatory compliance such as the Equal Credit Opportunity Act; and also a Research Scientist at the Massachusetts Institute of Technology working on data science technologies. Since the start of 2022, I have spoken with multiple HR vendors, both startups and established companies, as well as a prominent legal firm, who have sought my input on how to establish compliance with the AEDT law ("The Law"). I would like to provide some comments on the real-world operational questions that have surfaced through my discussions with multiple data scientists, lawyers and vendors, as well as my thoughts on how best practices from federal regulators in employment and finance can be usefully translated into the context of the Law.

Vendor liability. One of the largest open questions around compliance for the Law is what liabilities vendors have. Most employers are unlikely to have in-house expertise to assess whether their use of such AEDTs are compliant, and thus will outsource compliance audits to third-party auditors. This is presumably the intent of the Law, to require such audits to be conducted. However, most employers also lack the ability to build their own AEDTs, and choose instead to purchase solutions from external vendors. In such situations, there is a separation of ownership between an employer's real-world employment data on one side, and the vendor's AEDT code and training data on the other. Auditors must thence navigate challenging internal politics to ensure both that the employer's data and the vendor AEDT is available for a successful audit, assuming that the vendor-client contract even permits such access.

Data ownership issues hinder testing of robustness. Modern data-intensive AEDTs that are built according to current best practices are not defined solely by their internal algorithmic rules, but also by the training data used to develop the AEDT, and also the development process of how models are built, selected against potential alternatives, and validated internally. It is therefore critical to assess the statistical similarity of the training data to the actual deployed use case. The performance of AI systems depends critically on the data fed into them - for example, most AI systems fed only men's profiles will not pick up on the absence of women, even if they were assessed to not have discriminatory outcomes when used on a gender-balanced test population. Therefore, the assessment of discriminatory outcomes must be evaluated in the context of the data for a specific population, which makes it difficult to purely absolve users of vendor solutions. Conversely, it is difficult for vendors to attest or certify that an AEDT will not produce discriminatory outcomes without making strong assumptions on typical usage of their solutions. In short, when vendors own training data while downstream users own test data, the two parties must engage in some form of data sharing to test for robustness and out-of-sample generalization error. Such data sharing must be done carefully to protect the identities in the data sets, particularly when it comes to EEO-1 demographic data and other personally identifiable information which could compromise privacy if leaked.

The risk of ethics-washing. There is an unavoidable conflict of interest that arises when companies pay for audits of their own products. Even in this nascent stage of development for the algorithmic auditing industry, controversy has already arisen over how some companies heavily censor the audits before release, or use the ostensibly neutral platform of academic publications to obtain validation for their reviews in the form of peer review. On one hand, it is important to recognize that compliance audits usually happen under attorney-client privilege, so that clients can address and remediate any negative findings without incriminating themselves in the process. On the other hand, the pay-to-play nature of auditing necessarily creates a conflict of interest that incentivizes keeping the auditee happy in exchange for future business and building relationships. Such concerns are of course not new, and have plagued financial auditing for decades. The experience of the financial services industry clearly points to the need for independent verification of audits, which are usually manifested in the form of regulatory audits by government entities.

Reproducibility requires data and algorithmic governance. The very act of auditing an AEDT implies that an auditor can independently reproduce claims made by developers about the properties of an AEDT, such as its lack of discriminatory impact or expected performance on a given test set. However, the act of reproducing an AI/ML system itself - to set up a replica of the production environment with a replica of a data stream that resembles real world conditions - can itself be a major engineering challenge. Successible reproduction of a production system in an audit environment that does not a ect production data streams will be necessary to ensure that the auditing process does not inadvertently pollute and a ect the use of the AEDT itself.

Data quality issues in demographic label collection. A related issue is that of data quality - the most highly predictive AEDT is still likely to fail if fed erroneous data. In the context of algorithmic auditing, data quality issues extend not just to the data fed into an AEDT, but also the EEO-1 demographic data and other personally identifying information that needs to be collected in order to correctly classify people by race, gender, age, and other protected classes. In practice, EEO-1 demographic data is voluntarily provided by applicants and employees, which means that people will voluntarily refuse to self-identify. Such refusal is not statistically random, but is disproportionately likely to occur when membership in a protected class, such as having a mental disability or being of a particular sexual orientation, carries social stigma or otherwise is likely to cause harm by “outing” someone to belong to some group. This missing not-at-random nature of demographic label quality ought to be considered whether or not discriminatory outcomes can be measured with su cient statistical power, particularly if an imputation method like Bayesian Improved Surname Geocoding (BISG) is used to fill in missing demographic labels, as is commonly done in compliance testing for consumer financial services.

Construct validity. The need for AEDTs is greatest when there is an inherent scaling challenge to the number of decisions that have to be made. In the employment context, this usually shows up in the early stages of recruiting to narrow the funnel of applicants that are shortlisted for subsequent rounds of interviews. However, it is unclear if data collected at early stages of an employment decision, such as receiving a resume or video recording from a job candidate, will contain enough predictive signal to accurately predict a candidates suitability for hiring. In practice, AEDTs cannot predict something abstract like "employability", but instead compute metrics that purport to measure suitability scores or the like for such abstract concepts. An audit must necessarily assess the problem of construct validity, that a prediction target of an AEDT is indeed a valid and suitable quantification that operationalizes the employment decision being considered. Such considerations are of course of long-standing debate in federal employment laws; however, the algorithmic nature of decision-making and its use in making quantitative predictions bring such fundamental measurability concerns to the forefront of assessment. Many metrics purporting to quantify algorithmic bias implicitly assume that the prediction target of the AEDT is perfectly well-defined without any measurement ambiguity, which is unlikely to be true in practice. Therefore, the construct validity of the prediction target needs to be assessed critically to avoid false overprecision and overconfidence in the quantitative evidence for or against algorithmic bias.

The ethics of negative class sampling. A particularly thorny data quality problem goes by the name of reject inference in credit decisions, and is closely related to the problem of positive-unlabeled learning in other machine learning contexts. It is a problem for AEDTs that create a data asymmetry between positive and negative decision classes. For example, an employer incrementally collects more and more information about a candidate that passes multiple interview rounds. Conversely, a candidate not selected for an interview will have less data about them. This means that for hiring decisions, it is easier to assess false positives (a promising candidate that turned out to be a poor employee) than false negatives (a candidate that did not interview well that would have been a good employee). The counterfactual nature of the negative class makes assessments involving them di cult to assess in practice - someone removed from the candidate pool is by definition someone that was never placed in a job, and hence there was no real measurement of whether or not they were good at their job. A critical assessment of an AEDT's predictive value ought to include assessments of how well they classify the negative decision class, but if this class is not measured in any data set, then expert review is needed to validate negative decisions, or otherwise an experimentation framework is needed in order to test counterfactual changes to the AEDT prediction. There are obvious ethical risks to deliberately altering an employment decision for the sake of algorithmic assessment, as well as high costs of incorrect classification which will hinder the collection of real-world validation data. A well-designed audit should recognize the importance of negative class sampling, while at the same time have procedures in place to e ect the necessary counterfactual testing without undue cost.

Intersectionality and subject privacy. The explicit call-out for intersectional testing across multiple protected classes is a welcome strengthening of current federal standards, which do not require testing of, say, race and gender, simultaneously. Nevertheless, intersectional concerns increase the risk of identification and hence loss of privacy for underrepresented groups. The more labels used to define a category, such as “people of gender A, race B, and age group C”, the fewer people are likely to belong to that exact category. Taken to its logical extreme of testing every single protected class defined under federal employment laws, there is a risk that the intersectional categories are so fine-grained that only a single person may belong to that category. When such categories exist, summary statistics can leak information about a single person. In practice, the granularity of intersectional categories must be balanced against privacy concerns. I have some very preliminary research that indicates that di erential privacy is a promising mechanism for achieving these goals in an algorithmic audit, although field testing will be required to validate the theoretical work we have been able to publish.

The fallacy of the four-fifths rule. The literature on algorithmic bias has unfortunately perpetuated a misconception of the significance of the four-fifths rule which the current rules are at risk of perpetuating and codifying. It is often claimed that the Equal Employment Opportunity Act enshrines the disparate impact ratio as the only legitimate metric for measuring employment discrimination, and that when it exceeds 80%, there is no finding of employment discrimination. In reality, tracing the historical development of the four-fifths rules reveals that it was only ever meant to be a bureaucratic rule of thumb for prioritizing cases for further regulatory scrutiny, and in fact the 80% threshold was e ectively set arbitrarily in a 1971 meeting of the California Fair Employment Practice Commission as a compromise between a 70% camp and a 90% camp, a compromise that seems to not have been revisited with much scrutiny ever since. The arbitrariness of the four-fifths rule has been recognized by multiple federal courts in multiple court cases: courts have found that the 80% threshold is neither necessary nor su cient to make a determination of discriminatory outcomes, and have admitted other forms of statistical testing, such as hypothesis testing for equality of means, in actual court cases. In short, the 80% threshold is arbitrary and fails to capture less severe discriminatory outcomes, particularly when the sample size is small and when the membership of people in protected classes is unclear.
To address these operational challenges, I would like to make the following recommendations for your consideration.

Recommendation 1. The City should invest in their own auditors and regulators to assess if audits need to be themselves independently audited, adapting relevant best practices from financial regulators and auditors where helpful.

Recommendation 2. The Rules would benefit from clarification on governance requirements for AEDTs and their associated data sets.

Recommendation 3. The Rules should clarify how robustness and generalization ought to be tested, and if so, how data sharing between di erent owners can be e ected for the purposes of compliance audits.

Recommendation 4. The Rules would benefit from clarification on what liability vendors have for selling AEDTs to downstream clients, and to what extent (if any) these downstream procurers of AEDTs are able to shift liability to the vendor.

Recommendation 5. Regulators should work with standards-setting bodies, such as the National Institute for Standards and Technology (NIST), to develop and curate test data sets that represent typical populations which may be a ected by AEDTs, so as to enable high quality testing of AEDTs that a ords apples-to-apples comparisons.

Recommendation 6. The regulators should favor companies that have voluntarily adopted the NIST AI Risk Management Framework (RMF) or similar best practices for building and using AI systems. The regulators should issue more specific guidance aligned with the AI RMF to streamline compliance reviews.

Recommendation 7. The Rules should not codify any specific metric or threshold for passing or failing, but rather accommodate a possible plurality of valid metrics, and insist on tests of statistical validity rather than simply passing a numerical threshold.

In closing, I would like to congratulate the City for its innovation for enacting the Law, the first of its kind for the employment industry. The comments above are not meant to detract from the significance of the Law, but rather to highlight implementation risks that ought to be managed in order for the Law to have its desired e ect to promote inclusivity and accessibility of job opportunities, improve transparency in high-stakes decision making, and reduce discimination in employment decisions. 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.

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).

Question:

Is a system that makes rule-based decisions to route candidates based on bona fide occupational qualifications, such as “Are you 18 years of age or older?” or “Do you have a Series 7 license?” that would disqualify candidates based on their answer considered to have made an automated decision?

Thank you!

Adam Godson
Chief Product Officer at **Paradox**
phone (262) 441-0245
email adam@paradox.ai
https://www.linkedin.com/in/adamgodson/

Olivia is the AI assistant transforming global talent acquisition and candidate experience.

KNOWLEDGE LEADERSHIP COMMUNITY — THE BLACK INSTITUTE

Bertha Lewis
The Black Institute
470 Vanderbilt Ave, 9th Floor
Brooklyn, NY 11238

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

November 4, 2022

Dear Commissioner Vilda Vera Mayuga and Agency Staff,

As a longtime civil rights activist and community organizer, I am thankful for your efforts to see that Local Law 144 is implemented in an effective and timely manner. We should all be very clear on the precedent-setting nature of this legislation and the opportunity it represents for New York workers.

Not all drivers of racial inequality in our society take the form of loud and obvious instances of oppression. When it comes to economic opportunities, people of color are often disadvantaged when systems that appear neutral are quietly designed to preserve the status quo. As Local Law 144 correctly identifies, the hiring process is one context where seemingly neutral systems are especially likely to perpetuate discrimination.

Employers have many choices in terms of the decision-making tools they can use in a hiring process. Often, these tools are purchased from third-party companies, though they are occasionally built by internal teams. Some employers elect to procure aptitude tests that have been around for decades from traditional test publishers, while others use computer programs built by tech startups to screen resumes. The products run the gamut.

The fundamental challenge with the wide variety of decision-making tools lies in predicting the consequences they will have on society. We tend to assume all tests and standardized evaluations are “objective,” but the reality is that some disadvantage people of color and others facilitate a more inclusive process. Unsurprisingly, when employers opt for tools that disproportionately disadvantages Black and Brown people, those organizations struggle to make progress on workforce diversity. Instead of interrogating their hiring procedures, however, some companies would rather attribute the situation to minorities being underqualified or uninterested in the job.

Local Law 144 presents a very simple solution to the mystery regarding the consequences of different hiring tools. Beginning in 2023, employers will have to ask themselves “does my hiring tool disadvantage underrepresented minorities?” (with the help of an independent auditor) and then present the results on that inquiry to the public. The results are summarized in a report called a bias audit. A bias audit is simply a label that gives the public more information about an employer’s workforce priorities.

I know that the business community has already begun to push back against the suggestion that they owe us transparency regarding racial equity in hiring. They have argued - and will continue to argue - that the public reporting mandates under Local Law 144 are burdensome, complex, or vague. Interestingly, these actors never describe bias audits or their contents as irrelevant for employment equity. Why? To put it simply, the vast majority of employers are well aware their hiring procedures may be extremely racially biased, and they recognize that revealing this information to the public will have massive consequences.

If this administration wants to live up to its promises of providing good jobs and economic opportunities to all New Yorkers, there can be no tolerance for efforts to dilute Local Law 144. There can be no loopholes created or penalty fees suspended. The business leaders who claim to be confused about this legislation are not nearly perplexed as they would like policymakers to think. I encourage the agency to keep in mind that federal recordkeeping requirements under employment law already have many companies collecting the very data needed to conduct bias audits. The thing they are opposed to is only the *transparency* regarding information that they have historically been able to keep hidden from view.

Thank you for your consideration of my comments. I appreciate DCWP’s hard work regarding this important initiative.

Regards,

Bertha Lewis
Founder and President
The Black Institute

Hello,

I have a comment about the proposed rules for the Automated Employment Decision Tools, and in particular the use of Impact Ratio when a score is given by an automated system.

In general, impact ratios based on selection rates are very different from impact ratios based on average scores. In particular, the latter can hide discrimination depending on how the scores are subsequently used. The DCWP should revisit their requirements to reflect this.

For example:

First, imagine a system which selects candidates. If it has an impact ratio of 90% for women, this means women are selected 90% as often as men, which is could generally be considered acceptable.

Now, imagine an automated process which outputs a score, and has an impact ratio of 90% for women (women on average get 90% the score of men). On the face of it, this could seem equally good. However, if the subsequent selection process, perhaps not automated but done by a human, uses scores as a deciding factor, then the final ratio of women getting the job could be much less than 90%. In fact, if employers only take candidates with the top 5% of scores, then this could mean almost no women are retained for the job, despite on average getting 90% of mens' scores.

Currently, the draft seems to only require the reporting of the impact ratio based on the average scores in cases where the system outputs a score. However, in cases in which the tool is used to substantially assist the decision (as in the above example) the scores could cause a large difference in outcomes which is not reflected in the impact ratio.

A possible solution would be to always require an impact ratio based on selection rates. In instances where the automated tool only outputs a score for use by employers, the employers could report using the selection rate of the entire process. For example, if an employer is using an automated system to score candidates, and then uses these scores (in addition to other inputs) to decide which to interview, employers could be required to publish the impact ratio based on selection rates of the final decision of who to interview. This could be in addition to impact ratio based on average scores.

Best Wishes,
Hadrien Pouget

Hadrien Pouget
Visiting Research Analyst, Technology and International Affairs Program
Carnegie Endowment for International Peace
Hadrien.Pouget@ceip.org

STERLING MILLER
CEO & SENIOR COUNSEL

HILGERSGRABEN

7859 WALNUT HILL LANE | SUITE 335
DALLAS, TX 75230

972-848-3838 DIRECT
SMILLER@HILGERS GRABEN.COM

HILGERS GRABEN.COM

November 3, 2022

Via Email

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 Proposed Rules to the Use of Automated Employment Decision Tools Under Local Law 144 of 2021

Dear Commissioner Mayuga:

My firm represents clients with an interest in the above. In particular, we have several questions for you and your staff to consider as you move forward.

Background

Our questions revolve around the audit requirements and adverse impact analysis. There are several clients that do not have a sufficient sample size to properly complete an adverse impact analysis. Research literature in this area generally requires no less than 30 data points per cell with an upper range in the tens of thousands in order to ensure that the bias audit itself is not subject to bias. Without a sufficient number of data points, is any evaluation or audit truly meaningful or proper? While the impact ratio is specified, there are no sample size guidelines provided. In other words, as proposed, the law and the regulations provide little to no guidance on the intended and appropriate sample size businesses must use for the mandated adverse impact analysis in order to real valid results.

Without such guidance, we do not see how a business will know if an adverse impact regarding the use of an AEDT significantly and meaningfully exists. Without standards set by your offices, the analysis is left completely to whatever standards the auditor may employ, which could vary widely between such vendors leading to “audit shopping.” Moreover, a finding of bias that is not statistically supported could unjustly and unfairly impair businesses using AEDTs and hurt their ability to hire and promote. Additionally, what about businesses that are just starting to use an AEDT? They have, literally, no sample to test against. Do they get a pass? If not, how should they proceed as there are no clear specifications on evaluating selection practices that yield too small (hire or promotion) of a sample size - one that cannot reasonably provide statistical significance or insignificance.

<u>Questions</u>

We think the above can be set out in the following two questions for your office:

1. How are companies and experts supposed to evaluate the probative value of the statistical significance or meaningfulness in a small sample size for AEDT audit purposes and at what level is it improper to perform this analysis?

2. How can companies comply with the law if they do not have a valid sample size to audit because they are just starting to use an AEDT?

We appreciate your office’s attention to these questions and look forward to further guidance.

Sincerely,

Sterling Miller

1. What “alternative processes” are envisioned here—what does it mean? The proposed rules remain silent on how “alternative selection process or accommodation” is defined or process to notify the applicant of whether the request was granted or denied. They state only that the notice has to identify a process for requesting an alternative selection process.
2. Why is this process only required for race/ethnicity and gender but not age (or any other legally protected category)?
3. Who qualifies as an “independent auditor” and where is an employer supposed to find one? Who will incur the expense of hiring the auditor and conducting the audit and has DWCP determined the cost or impact of such an expense on an employer? Why are all employers regardless of size covered by this provision?
4. How can an employer publish past results that do not exist if the employer is using AEDT for the first time?
5. How are notice violations determined? Why is there no opportunity for an employer to cure a violation?

Lisa M. Brauner (she/her) | Partner | Schnader Harrison Segal & Lewis LLP
140 Broadway, Suite 3100 | New York, NY 10005-1101
Phone: 212.973.8170 | Fax: 212.972.8798 | Email: LBrauner@Schnader.com

New York | Pennsylvania | California | Washington, DC | New Jersey | Delaware faegredrinker.com

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

November 4, 2022

Submitted electronically via http://rules.cityofnewyork.us
And by Email, via rulecomments@dcwp.nyc.gov

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

To the Department of Consumer and Worker Protection:

Faegre Drinker Biddle & Reath LLP is a law firm with offices throughout the United States, and in the United Kingdom 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 think that the Proposed Rules published by the New York City Department of Consumer and Worker Protection ("Department") provide much needed clarification of the text of Local Law 144 of 2021 ("LL 144"), particularly by refining of the definition of what is an automated employment decision tool ("AEDT"). There had been some concern, raised in some of the public comments about the law made in June 2022, that any use of a computer to tabulate data would be an AEDT. We see now that those concerns have been alleviated, but we also agree with most of the comments made by another commentator, the Society for Human Resources Management ("SHRM"), set forth in its letter dated October 24, 2022, which call for even more clarification. We also have some additional views

Our comments and recommendations to the New York City Department of Consumer and Worker Protection regarding the Proposed Rules are as follows.

1. **Comment and Recommendation, regarding Remote Workers Based in New York City**: We interpret the Proposed Rules to mean that LL 144 applies only to positions for employment that are located in New York City, and that the notices required by LL 144

are required to be given only to residents of New York City. Some of our clients are multinational corporations, with employees all over the world. For many positions, the companies allow employees to work remotely from any place in the world, with the work location being entirely of the employee's choosing. The companies do not have an office in New York City, but some employees choose to live here. If the 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's site should be the employer's site from which the work of the New York City remotely working employee is directed. Both the federal Worker Adjustment and Retraining Notification Act, 29 U.S.C.
§§ 2101 et seq. ("Fed WARN") and the New York State Worker Adjustment and Retraining Notification Act ("NY WARN") call remote employees "out-stationed" employees. *See* 20 C.F.R. §639.3(i)(6), and N.Y. Comp. Codes R. & Regs. tit. 12, § 921-1.1. Their site of employment is deemed to be a company location, as follows:

> The single site of employment for employees … who are out-stationed, or whose primary duties involve work outside any of the employer's regular employment sites … shall be the site to which they are assigned as their employer's home base, from which their work is assigned, or to which they report.

N.Y. Comp. Codes R. & Regs. tit. 12, § 921-1.1

There is existing case law interpreting the location of out-stationed, or remote, employees under both Fed WARN and NY WARN, and we suggest that the Department adopt a similar understanding of where it is that a site of employment is located, such that a New York City resident working in New York City by choice as a remote worker be deemed to not be applying for a job within New York City.

2. **Comment and Recommendation, regarding Use of an AEDT for Soliciting Candidates for Employment**: We welcome the narrowing of the definition of who is a "Candidate for Employment," to be a person applying for a "specific employment position." Many individuals will submit a resume to a potential employer, not for a specific position, but as a matter of general interest in employment with an employer. Those individuals appropriately are not, under the Proposed Rules, a "Candidate for Employment." Many employers will apply AI to previously submitted resumes when a position opens up, for the purpose of soliciting an application for a specific position, and also may apply AI to its existing workforce to see if someone already on the payroll should be encouraged to apply for a particular job. As the Proposed Rules are written, it appears that no notice to those individuals is required to let them know that an AI process is being used to solicit candidates for employment. That said, it also appears that the AEDT must go through a bias audit before it is used for that screening purpose, and the audit must have been conducted no earlier than a year before the AEDT is used in that way. What is not clear is whether a resume submitted for a specific position for which an applicant was not hired may properly be included in the pool of all resumes in possession of the employer for purposes of finding potential candidates when a different position opens up. We respectfully suggest that no notice should be required to a

person whose resume is being kept on file when the employer evaluates whether that past applicant is a good fit for a new or vacant position. Large employers may conduct hundreds of such evaluations a year, and it would give false hope to a past applicant who receives hundreds of notices that his or her resume is being reconsidered. We recommend that the following be added to the definition of a "Candidate for Employment": "A person who in the past has unsuccessfully applied for a specific employment position is not a candidate for employment if an employer uses AEDT to evaluate whether prior candidates should be solicited to apply for positions that subsequently become available."

3. **Comment regarding Accommodation for Candidates for Employment and Employees Who Opt Out of Use of the AEDT for a Specific Employment Position**: On the one hand, Section 5-303(c) of the Proposed Rules states that a required notice "must include instructions for how to request an alternative selection process or accommodation," and on the other, Section 5-303(e) states, "Nothing in this subchapter requires an employer or employment agency to provide an alternative selection process." We are advised by clients and others who permit applicants to opt out of an automated decision-making process (pursuant to the laws of other countries) that the option is very rarely requested, and so we think it appropriate not to force employers to create a selection process that is rarely used. We note, however, that the word "accommodation," which appears in Section 5-303(c) of the Proposed Rules, does not appear in Section 5-303(e), and so we think that those who choose to opt out of being measured by an AEDT, are to be provided an accommodation, which seems to not be a process, but a case-by-case assessment of how an applicant is to be evaluated. We think the ad hoc approach to dealing with the very few who opt out is therefore workable, assuming New Yorkers behave in a manner similar to workers in other countries.

4. **Comment and Recommendation Regarding Effective Date:** A number of public commentators heard at today's public meeting noted that the law will not have final rules until shortly before the law is scheduled to take effect, which is January 1, 2023. We respectfully request that as an enforcement policy, the Department allow at least 90 days to elapse between the issuance of the final rules and the date on which the Department begins to take action to enforce LL 144.

5. **Comment and Recommendation Regarding Good Faith Efforts to Comply for Purposes of Assessing Penalties**: We note that LL 144 is the very first law within the United States squarely and comprehensively addressing the use of AI in making employment decisions. LL 144 is a law of relatively few words, as are the Proposed Rules. We therefore echo SHRM's suggestion in its October 24, 2022 comments that "the Department include a safe-harbor provision that permits employers to be deemed to have complied with LL 144 based on good faith efforts toward compliance."

Thank you for the opportunity to provide our comments.

Very truly yours,

Gerald T. Hathaway

/s/ Bennett B. Borden
Bennett B. Borden

Kate E. Middleton

Scott M. Kosnoff

/s/ Matthew A. Fontana
Matthew A. Fontana

GTH/jq

RETAIL COUNCIL OF NEW YORK STATE
RESULTS FOR TODAY'S RETAILERS

258 State Street
Albany, New York 12210
(518) 465-3586
rcnys.com

VIA EMAIL

November 3, 2022

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, founded in 1931, 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 305,000 retail jobs in New York City in August 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, the definition for "automated employment decision tool" is incredibly broad. Our concern, specifically, is that it encompasses nearly anything that attempts to evaluate candidates in a consistent manner.

2.) Similarly, the law is so broad as to encompass anything involving a computation; it goes beyond the use of artificial intelligence or machine learning as commonly understood.

3.) Publishing the impact ratio analysis and selection rate as part of the summary of results from the bias audit may needlessly expose companies to scrutiny from other agencies. The numbers must be analyzed and interpreted in context, as the numbers alone may give the misperception of bias when none exists. In addition, the rule defines impact ratio but does not define any level at which the impact ratio is problematic.

4.) Lastly, we believe clarification is needed in the following areas:

a. First, if a company uses automated employment decision tools to both score applicants (§5-301(b)) and select candidates to move forward in the hiring process (§5-301(a)), would they be required to conduct two separate bias audits, or would one bias audit be sufficient?

b. Second, as more jobs are being listed as remote, employers will need clarification as to which jobs, and subsequent use of automated employment decision tools, will require bias audits. Will employers need to conduct an audit for jobs that New York City residents may apply for, or only for jobs physically located within the city's boundaries?

c. Third, the definitions of "bias audit" and "independent auditor," as they relate to §5-301, appear contradictory. Bias audit is defined as "an impartial evaluation by an independent auditor." Independent auditor is defined as "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." However, the example given in §5-301 states, "The employer asks the vendor for a bias audit. The vendor uses historical data it has collected from employers on applicants selected for each category to conduct a bias audit." The vendor serving as an independent auditor would appear to contradict the term's definition.

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 subsequent 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,

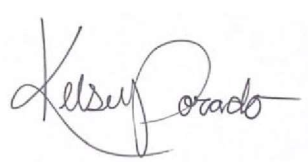

Director of State and Local Government Relations
Retail Council of New York State

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 rules to implement new legislation related to the use of automated employment decision tools.

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, industry 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 rules.

While SIOP is supportive of efforts to implement a regulatory framework around the use of automated employment decision tools, we have specific concerns regarding the methods of calculating adverse impact in the draft 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 average scores) 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 the standard deviations of the means of subgroups or the sample sizes of subgroups. The method based on subgroup averages will likely produce some conclusions that are not consistent with those of other widelyaccepted measures of adverse impact (i.e., masking existing impact or indicating something that would not be considered impact with other commonly used metrics for calculating adverse impact).

The examples provided imply that employers must look at the intersectionality of race/ethnicity and gender (e.g., Non-Hispanic/Latino White Males compared to Non-Hispanic/Latino Black Males) rather than the more traditional approach of comparing selection rates for racial/ethnic groups and for gender groups. In many cases, the intersectionality approach will lead to very small cell sizes, which produce unstable results and are difficult to interpret.

SIOP and our issue experts would be happy to provide our evidence-based insight on how to best craft a scientifically rigorous regulatory framework that accomplishes the goals of the City of New York Department of Consumer and Worker Protection.

1

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
Vice President of Employment & Litigation Services, DCI Consulting
edunleavy@dciconsult.com

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

Dr. Ann Marie Ryan
Professor of I-O 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

Outtz, 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, *Industrial and Organizational Psychology: Perspectives on Science and Practice, 11*(Supl 1), 2–97. https://doi.org/10.1017/iop.2018.195

2

STATE PRIVACY & SECURITY COALITION

November 4, 2022

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

Re: **<u>Proposed Rules to Local Law 144 (2021)</u>**

Dear Commissioner Mayuga,

The State Privacy & Security Coalition, a coalition of over thirty companies and trade associations in the retail, technology, telecommunications, payment card, automobile, and health care sectors, writes to provide comments on the proposed regulations regarding automated employment decision tools (the "proposed rules"). We appreciate the opportunity to provide input and suggest some modifications that can help better align the proposed rules with the statute's intent.

We also would emphasize that artificial intelligence, machine learning, and other related technologies have the ability to dramatically *reduce* unlawful discrimination and unconscious bias, ensuring that employers are able to identify candidates that are best suited to the position and reducing the prevalence of human error. We agree that unlawful discrimination in all forms should be mitigated to the greatest extent possible; we also believe that new technology has the potential to accomplish this even more successfully than human review alone.

Our coalition works to ensure that consumers are able to obtain increased transparency over how their personal data is used, while balancing both operational workability and anti-fraud/cybersecurity for businesses. The proposed amendments that follow are made in that spirit. SPSC language appears in **<u>bold underline</u>**, proposed deletions in ~~strikethrough.~~

1. **Definition of "Automated Employment Decision Tool"**

We propose including language in the definition of Automated Employment Decision Tool ("AEDT") to exclude the use of AEDTs to search resumes for candidate qualifications. In the world of online job postings and remote work – both of which have democratized the job search process for candidates – it is immensely helpful for employers to screen out resumes that do not meet basic requirements. For example, if an accounting firm has an open position for which a CPA is required, automation can save significant amounts of time by eliminating candidates that do not possess such a qualification. Conversely, if an employer is searching for a particular skillset – leadership positions in college or in an early career stage – automation can quickly identify those candidates and provide visibility when human review may unintentionally omit them due to the volume of resumes received.

We suggest the following addition to the proposed definition: "'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' does not include the automated searching of resumes to identify candidate qualifications including relevant skills or experience**."

2. **Definition of "Screen"**

SPSC proposes re-framing this definition to encourage beneficial uses of screening and more clearly target those uses cases that could have adverse impacts. Specifically:

"'Screen' means to make a determination **that a candidate 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.~~

The term is used twice in the operational portions of the bill, and we believe the above definition better aligns with the intent of that usage. The first time it is used is when a business is prohibited from using an AEDT to screen employees unless there has been an audit with the results made publicly available; the second relating to the notices required for employees prior to its use.

By adopting our proposed amendment, we believe this more appropriately focuses both the audit requirement and the notice requirement on circumstances when employees or candidates have or could have an adverse decision made about them, rather than any use of screening that is positive in nature. SPSC believes that regulation around automated decisionmaking should encourage transparency around uses that could result in an unlawful disparate impact, and seeks to align statutes and regulations accordingly.

3. **Disclosure of Impact Ratios**

The ordinance does not refer to an impact ratio, and while we understand the good intentions behind the requirement to disclose these publicly, SPSC requests that they be used for internal purposes only. Our concern is anchored in the fact that candidates or employees may actually be *deterred* from applying for a particular job or promotion if they see an impact ratio that they interpret as negatively impacting their chances based on an immutable characteristic.

Instead, we would request that employers use the impact ratio for internal evaluation of their AEDT, and also allow any candidate who would like to see impact ratio have the right to request a copy for their records and use. A primary goal of the employment process is to foster an inclusive workplace, and in order to do that, attracting qualified and diverse applicants is critical. SPSC does not want to undermine this important process.

4. Bias Audit

SPSC requests that this body clarify the timing of when a bias audit of an AEDT must be completed. As drafted, the draft regulations read as if an audit must be undertaken annually; and yet, it would not make sense for this to be the case, as an AEDT that is found to be working well should not have to undertake a formal audit process unless there has been a material change in the development or deployment of the AEDT. We would request that clarifying language be added to better understand how the audit obligation is best operationalized.

5. Enforcement Relief

SPSC requests that the Department provide some ability for businesses that operate in good faith to obtain relief against what will surely be a flood of litigation seeking to leverage this legislation for pecuniary gain.

There are several ways the next version of the proposed rules could accomplish this:

- First, the proposed rules should incorporate a delayed enforcement date, which would not require delaying the effective date, but delaying the date by which enforcement would begin. SPSC suggests that, as this attempts to regulate new subject matter found nowhere else in the nation (even Illinois' law does not include the granular requirements found here), entities be required to implement the requirements in January 2023, but delay enforcement until January 1, 2024.
- Second, the proposed rules should allow businesses an affirmative defense whereby, if they perform the audits and take any required remedial action to rectify possible disparate impact, and presuming they have had no previous violations, they will be granted an affirmative defense in any enforcement action, private or by corporation counsel.
- Lastly, the proposed rules should establish evaluative criteria that a prospective plaintiff would have to establish prior to bringing an action. These could include:
    - Demonstration of actual injury;
    - Whether such injury was likely caused by design or via a technical error;
    - The degree to which the AEDT caused the actual injury; and
    - The likelihood that the alleged defect in the AEDT process was individualized or widespread among a larger population.

Of course, we would be happy to discuss any of the foregoing at your convenience. We appreciate the opportunity to comment on this important topic, and fully support the goal of inclusive hiring and promotional practices.

Respectfully submitted,

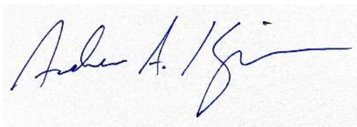

Andrew A. Kingman
Counsel, State Privacy & Security Coalition

RULEMAKING ON AUTOMATED EMPLOYMENT DECISION TOOLS LOCAL LAW 144 OF 2021

Last Updated 11.3.22

TESTIMONY: NYC DEPARTMENT OF CONSUMER AND WORKER PROTECTION

Tech:NYC is a nonprofit member-based organization representing roughly 800 technology companies in New York. Our membership includes some of the largest and most innovative tech companies in the world as well as hundreds of small and growing startups. We are committed to supporting New York's economic recovery and to ensuring that all New Yorkers can benefit from innovation. Tech:NYC and our members work with government and community partners to foster a dynamic ecosystem so that New York remains the best place in the country to start and grow a technology company.

The technology sector has shown strong growth over the past 10+ years, with an 142% increase in new positions since 2010 as well as a 8.7% increase in positions between February 2020 and December 2021. New York currently has a higher unemployment rate than it had before 2020, and labor shortages are common due to changing workplace dynamics as a result of COVID-19. It is also common that a large number of applicants apply to each job listing: earlier this year the online job search and advice platform Zippia concluded that the average corporate job listing receives about 250 candidate submissions. Combined together, these dynamics result in a huge demand by job applicants for technology related positions. Automated Employment Decision Tools ("AEDTs") play an important role for employers, both large and small, who receive overwhelming responses to open employment positions, and must efficiently review applications and fill their open positions.

Tech:NYC supports the city of New York in its process to establish notifications and information to those seeking employment or promotions, when the employer enlists AEDTs in employment decisions. Regarding the proposed rulemaking on Local Law 144 of 2021, Tech:NYC recommends that the draft rule on AETDs is amended to target its actions more directly at the portion of the process where an employer decides to not hire or promote a candidate. This will go a long way to both providing clarification to employers, and to ensuring tech employers can maintain efficient and timely hiring practices.

To further clarify the definitions within LL144 to ensure the law applies to aspects of AEDTs that directly tie to the decision of not hiring or promoting a candidate or employee, Tech:NYC recommends:

- Within the proposed rule's current definition of Automated Employment Decision Tool, one qualifying action is for a tool to weigh certain criteria over others when preparing a score, ranking, or other output. As there may be many criteria being looked at by a tool, and

- different criteria could be prioritized differently per each use of the tool, this definition should be further clarified to apply when weighing one set of criteria “substantially” more than another set.
- The definition of AEDT should specifically no longer include the automated searching of resumes for qualifications which include certain skills or prior experience. This will go a long way to clarifying how this law is applied to tools which effectively assist in reviewing applications, without impacting the requirements of tools used to make ultimate employment decisions.
- The definition of **“Screen”** should be redefined to apply to AEDT functions which determine when a candidate or applicant should ***not*** be hired or promoted.
- Bias audits should only be applied when AEDTs screen candidates or employees for employment decisions, as redefined directly above as determining when a candidate is not hired or promoted. Bias audits should not be applied when AETDs are used for scoring or rating candidates.
- Published bias audit results, as defined in § 5-302, should apply only to AEDTs that screen candidates or employees, as redefined above.

Thank you for the opportunity to present input from employers in the technology sector, and for your consideration of these recommendations.

CITIZENS COMMITTEE FOR NEW YORK CITY

Comments to DCWP on Proposed Rules for NYC Local Law 144 of 2021
Dr. Rahsaan Harris
October 24, 2022

Good Morning,

I am Dr. Rahsaan Harris, and I currently serve as the CEO of Citizens Committee for New York. Our organization is committed to providing New Yorkers, particularly those in low-income neighborhoods, with resources to improve their quality of life. I have spent the last 20 years working at the intersection of corporate philanthropy, community empowerment, and racial and socioeconomic equity.

Local Law 144 represents a long overdue effort to unpack the problem of bias in hiring. I am very excited that city officials are actively working to see this important law implemented after the new year, because it has the potential to bring real change to how workers are evaluated for economic opportunities.

Throughout my career, I have repeatedly witnessed the power of transparency in furthering social equity. Today, it is very common for executives to make bold claims of wanting to promote diversity and inclusion, but many fail to take even basic steps to involve people of color in the process. One way workers and job applicants are barred from meaningful participation is through information asymmetries. Business leaders might very well recognize that certain internal activities are perpetuating racial inequity, but if workers and job candidates are in the dark, it becomes exceedingly difficult to call for change.

This is why the bias audits called for by Local Law 144 are so critical. When employers decide what types of automated hiring tools to use, the decision is ultimately one about equity and diversity. One of the defining features of any assessment technology is the extent to which the particular tool disadvantages people of color. The vendors who build these products, and the employers who use them, are often keeping track of this information behind closed doors to comply with employment regulations under federal law. With bias audits, the doors will be open, and we will suddenly have evidence of which hiring tools are responsible for perpetuating racial equity.

I encourage DCWP to account for the incredible potential offered by this new legislation by implementing rules that will not make transparency seem optional for employers hiring New

Yorkers. Virtually every large company uses some form of an automated tool to screen job candidates because they receive far too many applicants to possibly review manually. The knee-jerk reaction of these employers will be to try to circumvent the requirements of this law. We can hardly blame them; this law is unprecedented and transparency is scary because it represents giving up some control. Attitudes will change in due time as transparency becomes the norm, especially because many employers are genuine in wanting to increase diversity. But until then, regulators must make it clear to employers that there are no loopholes.

There are a few loopholes in the proposed rules, which I strongly urge the agency to close. Right now, the draft rules suggest that a hiring tool only qualifies as an AEDT if an employer uses it without human oversight or weights it very heavily. With this definition, any smart corporate lawyer will have an obvious invitation to argue that their organization is exempt from bias audits altogether, and city officials would have no way of proving otherwise.

Another loophole that the agency should close relates to the technical definition of an AEDT. The proposed language suggests that hiring tools that rely on machine learning are the ones that are most worthy of scrutiny. While this is certainly a prominent narrative in the media, which is inclined to fear-monger about technology, I do not believe sufficient evidence exists to support this claim. The racial consequences of hiring tools have been hidden from the public for decades. Since regulators do not know if AI and ML hiring tools are more problematic than more traditional forms of automation, it makes no sense to limit bias audits to one type but not the other.

If we look around at the amount of progress that has been made toward racial equity in the workforce over the last several decades, it is safe to say the results have been disappointing. At the individual level, many people are doing their part to be kind and compassionate as a strategy for promoting inclusion and equality of opportunity, but the individual level can only take us so far. Racial equity in hiring is a problem rooted in systemic barriers, and until we know what those barriers are, we cannot find better solutions. I am thankful to the agency for taking an important step in that direction.

Thank you for your time.

Sincerely,

Dr. Rahsaan Harris
Chief Executive Officer
Citizens Committee for New York City

Professor Janet Helms

Comments re: Rulemaking for DCWP Local Law 144 of 2021

Most people, who have applied for almost any kind of job, have been screened by employment tests. Bus drivers, sales clerks, corporate executives, and many others suffer testing. Yet for over 40 years as a professor of testing and a testing researcher, I have tried to expose a problem that is hidden from the public. It is that many employment tests support legal racial segregation, and test developers know it.

Employment tests seem to be objective measures because everyone responds to the same items in a given test. Yet when employers use high scores on tests to hire or promote employees, applicants of color are hired at lower rates than White applicants. The hiring difference occurs because their scores on employment tests are lower than White applicants' scores. If people of color are hired, employers may use their low scores on tests to justify segregating them into low-paying and low-skill jobs, while White applicants are chosen for the better-paying and higher-skill jobs.

It is no secret that groups of color have lower average scores on employment tests than White groups. Such test-score differences have existed since psychologists created the first employment tests to evaluate World War I army recruits in the early 1900s. The secret is that employers are not required to reveal to the public how much their tests and other tools contribute to poor employment outcomes for people of color.

Shortly after the Civil Rights Era, groups of Black workers filed a number of class-action anti-discrimination lawsuits. The goal of the suits was to prevent employers from using intelligence tests to keep Black workers from being promoted. In response, the Supreme Court coined the concept of "disparate impact." The Court defined the term as employers' use of supposedly neutral employment procedures that in practice limited opportunities of "protected groups." The Civil Rights Act of 1964 had included racial groups of color and White women as protected groups.

The Equal Employment Opportunity Commission (EEOC) has required employers to use the EEOC's 80% rule to investigate adverse impact since 1978. As applied to Black and White racial groups, the rule is that an employer can obtain a "fairness proportion" by dividing the proportion of hired Black applicants by the proportion of hired White applicants. The EEOC specifies that a fairness proportion less than .80 possibly indicates discrimination. A proportion multiplied by 100 becomes a percent, which is why the rule is called the "80% rule." A fairness proportion of 1.00 or 100% indicates that the hiring procedure is entirely fair.

Federal law does not isolate tests as the only type of employment procedure that should be subjected to analyses of adverse impact. Yet tests have had a large almost invisible role in causing adverse impact. When a test is used as a hiring tool, the employer decides what the passing score is. Different employers may choose different scores. These decisions value the employer's subjective determination of good scores and disregard the structure and content of the tests on which decisions are based. But it is likely that a test whose content is biased in favor of White test takers gives them an advantage in employers' decision making.

Just how sure am I that systemic racial bias in hiring tools is a hidden problem? Last year, the American Psychological Association (or APA) issued an "apology to communities of color for its role in contributing to systemic racism" throughout its 130-year history. Robert Yerkes was a eugenicist who developed the first employment test. Later he became an APA President. He believed that Whites were superior to peoples of color, and that Whites' higher mean test scores proved it. Many present-day tests use Yerkes's racially biased framework.

The APA has provided professional guidelines for the development, scoring, interpretation, and use of racially biased tests for over 30 years. It took a long time for APA to admit its guilt for fooling society in to using testing technologies that harm people of color. APA's apology confirms that traditional tests are biased toward White test takers in no uncertain terms. Therefore, adverse impact analyses should be conducted when the tests are used to make employment decisions. The results of these adverse-impact analyses must be made available to the public.

Throughout my academic career, critics of my individual efforts to expose racial bias in testing have asked me: If testing practices are so unfair, why don't people of color or their advocates prove it? As a researcher advocate, my answer is that publishers will not provide me with the information necessary to conduct adverse impact analyses. As for people of color more generally, the law has changed so much that it is almost impossible for them to prove the adverse impact that they experience. Individuals cannot file antidiscrimination lawsuits. They must file complaints as members of groups. Also, the complaining groups, rather than the employer, must provide the evidence of the adverse impact. How are job applicants supposed to get access to employers' protected information?

The structure and design of employment tests have not changed since the 1920s. They are still racially biased. The only change is that computers administer and score the tests for employers. Computers provide the employer a score for each person and the employer decides what score is "passing". Different employers may use different passing scores for the same test. The public deserves to know the results of the employers' decision-making. Furthermore, the public should support recent efforts to end disparate impact.

One effort the public should support is the use of Artificial Intelligence (AI) to change how employers make hiring decisions. Advances in data science and AI create an opportunity to do things differently. For example, it is now possible to use machine learning to build different hiring models and estimate the amount of disparate impact that will likely occur under each model. Employers will be given this information before they ever make any judgements about real people so that they can make fairer decisions. Some people are suspicious of using AI for decision making. Perhaps they do not know "Siri" or "Google". I believe that with proper oversight, AI's use in automated hiring could drive real progress toward equality of opportunity.

A second effort the public should support is the groundbreaking Local Law 144, which takes effect in New York City at the beginning of 2023. The Law requires employers to reveal the disparate impact of their hiring tools, if they are doing business in New York City. This information will be provided in public reports called "bias audits." Bias audits must be completed by an independent party and posted to the organization's hiring website. Eventually the law will require mandatory "labels" that reveal the bias in each hiring tool. Bias labels will be similar to other labels that already exist for describing fuel efficiency and car emissions.

It is important that the public speak up about one specific requirement of Local Law 144 before the law goes into effect this January. The administration is currently writing some official rules to clarify for employers the details for complying with the Law. One detail that I think is extremely important is the specification of which types of hiring tools should submit bias audits. The current definition of hiring tools is limited because it excludes tests that people do not think of as "automated".

As a long-time researcher, dedicated to reducing the harm caused by biased employment tests, I am enthusiastic about Local Law 144. It will bring transparency to a very broad range of automated tools. "Automated" just means that a computer collects and analyzes data from job candidates and provides the employer with a score to assist with deciding who to hire or promote. Automated applies to almost all employment tools ranging from digital versions of the racially biased cognitive tests in use today to more sophisticated AI models. I firmly believe that progress toward racial equity means that every type of employment tool should be examined for disparate impact. Local Law 144 could make this happen.

Yet regulators have offered a plan that suggests that certain kinds of hiring technology deserves bias audits more than others. They have focused on AI technology. But AI hiring tools have not been responsible for a century of disparate impact. Those everyday tests that we all take have been. The "black box" normally attributed to AI exists in virtually all employment assessment technology. Local Law 144 could open the boxes wherever they exist in New York City.

Please see below additional comments regarding the proposed Automated Employment Decisions regulations. I provide these comments in my personal capacity, and not as a representative of my employer.

1. I appreciated very much the comments of the city councilor who spoke first during today's hearing and would echo her call for experts to help guide the development of fair, practical and effective regulations. I would strongly urge the agency to consult with Industrial/Organizational Psychologists, who are the foremost scientific experts in personnel selection and selection systems. You may wish to include Industrial/Organizational Psychologists from the academic world with strong track records of publication in peer-reviewed journals that focus on personnel selection. At present, there appears to be an overemphasis on computer science and computer scientists at the table, at the omission of those who are trained in psychological measurement principles.

2. I would caution the agency to avoid making the definition of "automated employment decision tool" overly broad, considering the administrative burden posed on employers. A broad definition that encompasses all assessment types that utilizes computer technology (rather than, more narrowly, artificial intelligence and machine learning) is likely to discourage employers' use of valid, objective, and fair assessment tools that predict key work outcomes. This could instead encourage employers to revert to the natural human biases that come into play in unstructured interviews, such as "who you know" being the determining factor in getting a job. In other words, this could unintentionally promote subjectivity and exacerbate discrimination risks.

3. The equating of bias audits and impact ratios is concerning, as noted in a previously submitted comment. Chapter 3 of the 2014 Standards for Educational and Psychological Testing describes how measurement bias, and more broadly fairness, are complex concepts that cannot be boiled down to an impact ratio alone. Validity of measurement (i.e., is the measurement fit for purpose) must be considered.

   Consider the following situation to illustrate the point. There is a considerable literature on sex differences in personality, suggesting that, on average, females tend to score higher on the personality trait of Agreeableness than do men (see Costa, et. al., 2001, Weisberg, et. al. 2011 for a few examples; the difference is moderate - typically .4 to .5 standard deviations on average).

   With that in mind, let us consider how impact ratios alone can give a misleading indicator of bias. Consider hypothetical "Personality Tool X" that unexpectedly demonstrates zero male-female difference in average Agreeableness (i.e., tendency towards cooperative, polite behavior) scores. Assume further that an employer decides to utilize the Agreeableness score from "Personality Tool X" to select employees for a job where Agreeableness is critical for job success. An impact ratio is then calculated using method 2 from the proposed regulations that shows the same average score for male and female applicants (ratio of 1.0).

   Depending on how it is used, "Personality Tool X" Agreeableness score might actually artificially disadvantage females in this employment context, despite appearing neutral on its face. This is because "Personality Tool X" may not reflect true population differences in Agreeableness as established in the scientific literature, which suggests females possess this attribute to a greater extent than males. It is contrary to current scientific knowledge in personality psychology. At a minimum, further investigation is needed in this situation. Equating bias with impact ratios throws out the science of psychological measurement, which is a well-established field with professional standards and guidelines.

Costa, P. T., Jr., Terracciano, A., & McCrae, R. R. (2001). Gender differences in personality traits across cultures: Robust and surprising findings. *Journal of Personality and Social Psychology, 81*(2), 322–331.

Weisberg, Y. J., DeYoung, C. G., & Hirsh, J. B. (2011). Gender differences in personality across the ten aspects of the big five. *Frontiers in Psychology, 2,* Article 178.

Regards,

Joseph Abraham

Joseph D. Abraham, Ph.D.
Vice President, Assessment Solutions

TALOGY
PREVIOUSLY PSI
TALENT MANAGEMENT

O: 01 818 847 6180, Ext. 7306 | **M:** 01 918 404 0050
Time zone: Central Time

Talogy.com | LinkedIn

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

October 24, 2022

Dear NYC Department of Consumer and Worker Protections,

My name is Andrew Hamilton and I'm Immediate president of the New York Metro Chapter of the National Black MBA Association. For decades, automated hiring tools have played a significant role in determining who receives serious consideration from employers. Unfortunately, the racial implications of these systems are seldom made clear to employers or the public.

Two years ago, I testified before the City Council's Technology Committee in support of Local Law 144. I firmly believed the legislation could disrupt one of the major challenges people of color face in accessing economic opportunity. **Once this law goes into effect, critical information that can drive real progress toward diversity and inclusion will be available.** Finally, with public bias audit reports, we will be able to have frank conversations about the consequences different methods of hiring have for racial equity.

As someone who firmly believes in the spirit of this law, I have a few comments in response to DCWP's proposed rules.

First, I believe it is essential that regulators implement this legislation so that the definition of automated employment decision tools, or AEDTs, does not encourage nefarious behavior. The behavior I am especially concerned about is employers who will change their hiring processes for the purpose of avoiding this important push for transparency. If the rules include overly specific details about which technologies are subject to bias audits, like those built with machine learning, employers will be incentivized to avoid those systems altogether, regardless of their racial impacts.

To be clear, while some people associate the term "automated" with cutting edge artificial intelligence systems, automated decision-making tools have been around for many years. 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; in 2022, whether you apply for a white-collar or blue-collar job at many of the country's most prominent businesses, there is a decent chance you will be evaluated by a tool that was first automated in the 1970s.

If DCWP tries to limit the scope of this law to only require transparency for the most modern technology, employers will face very strong incentives to avoid innovation. From my perspective, this would be a terrible outcome if the goal is to reduce racial disparities. If we know one thing for certain, it is that the most common ways of screening job applicants used in the last few decades have not advanced equity in hiring. I strongly urge DCWP to avoid any rules that employers may interpret as a loophole for older technology.

Another aspect of this law's implementation that I would like to speak on is the importance of bias audit reports being straightforward and trustworthy. Once implemented, bias audit reports will provide data that essentially indicates if a particular employer relies on assessment tools that tend to disproportionately disadvantaged people of color. With this information, we will have a massive opportunity to hold employers accountable for their diversity and inclusion commitments, *but only if the public feels the reports are valid*. If employers are under the impression that no one is taking the contents of a bias audit seriously, this law will not create meaningful incentives for companies to adopt fairer hiring tools.

One way regulators might make the reports more credible is to specify who counts as an "independent" auditor. Currently, the rules only mention that the auditor cannot have been involved in using or developing the tool. From my perspective, at the very least, the auditor should not be employed by either the employer or technology vendor. Auditors should also have to provide key information about how relevant data samples were sourced.

My decision to support this law early on came from my belief that innovation can be a powerful force for change, but only if we have the necessary evidence to evaluate it. In recent years, many new technologies have been developed for use in the employment context, and vendors often claim that these products can promote fairer outcomes. While some of these statements may very well be true, we cannot trust them blindly. With Local Law 144, we will be able to better differentiate between tools that were actually built with racial equity in mind and the ones that were not.

I would like to thank the agency for its hard work on this precedent-setting legislation. We should all be very proud of the important step New York City is taking to bring transparency to such a critical use case for technology.

Regards,

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

Comments on Proposed Rules for Local Law 144 of 2021

November 4, 2022

Dear Commissioner Mayuga and DCWP Team,

Please accept our comments on Local Law 144 as you work to see this incredibly important legislation implemented in the near term. Our recommendations can be summarized as follows:

I. Local Law 144 should be implemented with a very inclusive definition of AEDTs that does not rely on the details of how an automated hiring tool is built or used by the employer.
II. Local Law 144 should require disparate impact be calculated in a manner that aligns with methods enshrined in federal law in terms of both statistical methodologies and relevant demographic categories.
III. Independent auditors should specifically be defined as individuals who are not employed by the organization that builds, sells, or uses the AEDT.
IV. Bias audit reports should include information that clearly describes and contextualizes the nature of any data samples used. This should be the case whether an audit is conducted at the level of a vendor or that of an individual employer.

Please note that we recently published a peer-reviewed paper that includes an aggregate disparate impact analysis of an AEDT we are very familiar with: pymetrics. It is available here:
https://link.springer.com/article/10.1007/s43681-022-00208-x
Thank you for your work thus far. We truly feel it will make a difference for equity in hiring.

Best,

Frida Polli
Founder and Former CEO
Pymetrics

Sara Kassir
Former Policy and Research Principal
pymetrics

OVERVIEW:

In contexts like job applicant screening, **"automated decision making" is not a new term.**[75,76,77] The Federal Trade Commission specifically made this point in a 2021 blog post, countering the common misconception that existing regulations cannot be applied to AI/ML systems.[78] **Employers have been able to use automated tools, often purchased from third-party vendors, to screen job applicants for decades.** The basic parameters of such technology are very simple: a system collects data from or about job candidates, uses some computational process to analyze it, and provides an output to communicate an inference about the candidate's potential for success.

The notion of **using statistical tests to evaluate hiring tools for problematic forms of bias is also far from novel.**[79] One federal court aptly explained the importance of such data-driven methods in 1976: "Statistical evidence is an important tool for placing seemingly inoffensive employment practices in their proper perspective."[80] **For several decades, compliance experts have encouraged employers to conduct voluntary "self-audits" of their hiring tools for disparate impact.**[81] Testing researchers and HR tech vendors also conduct these studies in the course of developing automated decision- making tools.

While "audits" for disparate impact have been fairly common for half a century, results have seldom been made readily available to the public. This **opacity has restricted progress toward equality of opportunity for workers across the U.S.** in significant ways. First, many employers

---

[75] "The use of the computer to 'interpret' [personality] tests has progressed, in less than a decade, from speculative projection to solid reality." See: Fowler Jr, R. D., & Marlowe Jr, G. H. (1968). A computer program for personality analysis. Behavioral Science, 13(5), 413-416.

[76] "By programming decision models into the computer which approximate those of the persons involved, the computer can rapidly and efficiently 'consider' a great deal of information about many alternative jobs and candidates. In this way interviews can be proposed which hopefully are better than those resulting from manual file searches or machine searchers that seek simply acceptable pairings of job candidates and vacancies." See: Holt, C. C., & Huber, G. P. (1969). A computer aided approach to employment service placement and counseling. *Management Science*, *15*(11), 573-594.

[77] "In March 1976, the U.S. Civil Service Commission began operating its first fully automated written test scoring facility. Tests administered throughout the nation are now machine-scored and the results are returned to the administering office within 3 days. Previously, the machine scoring cycle frequently took 6 to 8 weeks to be completed." See: *Civil Service Journal.* (1976). United States: Office of Personnel Management.

[78] "The question, then, is how can we harness the benefits of AI without inadvertently introducing bias or other unfair outcomes? Fortunately, while the sophisticated technology may be new, the FTC's attention to automated decision making is not." See: Jillson, Elisa. "Aiming for truth, fairness, and equity in your company's use of AI." *Federal Trade Commission* (2021)

[79] "In the problem of racial discrimination, statistics often tell much, and the Courts listen." See: *State of Alabama v. United States*, 304 F.2d 583 (5th Cir. 1962).

[80] *Senter v. General Motors Corp.*, 532 F.2d 511 (6th Cir. 1976)

[81] "The great weight given by the courts to statistical employment data indicates it is to your own advantage to conduct a thorough self-audit voluntarily to determine possible existence of discriminatory practices and needed remedial action." See: Idelson, E. M. (1974). Affirmative Action and Equal Employment: A Guidebook for Employers. United States: U.S. Equal Employment Opportunity Commission.

who have relied on hiring tools with disparate impact for years face virtually no pressure to revisit their methods.[82] Despite increased calls to hold companies accountable on equity commitments, **neither diversity advocates nor HR professionals are sufficiently informed about the extent of bias (particularly in older forms of hiring automation)** to draw attention to the problem. Second, vendors of HR technology solutions are not incentivized to apply recent scientific advancements to the development of automated screening tools with less bias. Because so little information about disparate impact is available for the vast majority of products on the market, **well-intended actors are unlikely to be rewarded for building equitable alternatives.** In fact, efforts to create new forms of automated decision-making systems are generally subject to intense skepticism, rooted in sensationalized media narratives regarding the unknowable risks of technology.[83] Certain HR tech vendors further exacerbate this sense of mistrust by marketing their products as "bias-free" or "fairness-tested," often with no supporting evidence.[84]

Most fundamentally, **Local Law 144 represents a remarkable opportunity to bring transparency to the world of automated hiring tools.** The "bias audits" required by the statute can facilitate progress on workforce equity by (1) re-educating society on the fact that automated decision-making tools can either promote or restrict diversity in hiring and (2) providing clear and trustworthy information about the disparate impact of various hiring tools. **In effect, bias audits provide for a mandated public disclosure regime similar to those adopted in the climate mitigation context.** The goal of "labeling" particular systems or products in terms of its societal consequences - whether they be related to environmental protection or racial equity - is to invite the possibility of prosocial behavior.

Some commentators have argued that advancements in technology have rendered existing anti-discrimination regulations irrelevant. On the contrary, **Local Law 144's reliance on disparate impact to define decision-making bias makes the statute particularly useful in the present day.** The continued relevance of the concept is rooted in its technological agnosticism. As the Supreme Court stated in the landmark *Griggs v. Duke Power*, the data used to conduct a disparate impact

[82] "Most written examinations today continue to have substantial disparate impact; what has changed is that the tests are better constructed, in the sense that they are harder to challenge in court because they have been properly validated, but not better in the sense of being better predictors of performance." See: Selmi, M. (2005). Was the disparate impact theory a mistake? *Ucla L. Rev.*, *53*, 701.

[83] "People who are less familiar with AI/ML may also make potentially erroneous inferences about it, leading to adverse reactions within AI/ML-based systems…The tendency to make these inferences may stem from sensationalized coverage about AI/ML from entertainment and news media (e.g., as an agentic force that could overthrow or control humanity, or as a perpetuator of bias or prejudice)." See: Gonzalez, M. F., Liu, W., Shirase, L., Tomczak, D. L., Lobbe, C. E., Justenhoven, R., & Martin, N. R. (2022). Allying with AI? Reactions toward human-based, AI/ML-based, and augmented hiring processes. Computers in Human Behavior, 130, 107179.

[84] "Hiring technology vendors frequently claim that they audit and address bias within the tools they create. But they seldom offer details or make available the results of independent evaluations, at least publicly." See: Bogen, M., & Rieke, A. (2018). Help wanted: An examination of hiring algorithms, equity, and bias. Upturn.

analysis reflects the ***outcomes*** yielded by a given employment procedure, rather than its "motivations."[85] In other words, **disparate impact analysis is an approach to bias auditing that can cut through the noise that tends to characterize the HR tech industry.** The evaluation is one that emphasizes the ***consequences*** of a decision-making system for people of different demographic identities; it can therefore be applied in a consistent manner to 80-year-old hiring tools and cutting-edge technologies alike.

New York City's **Local Law 144 could very well change the face of hiring tools, but this possibility depends on the details of the statute's execution.** The most important questions regarding the law's implementation include: (1) When are bias audits required? (2) What information must be included in public reports? and (3) Who has the authority to serve as an independent auditor? If Local Law 144 is to achieve its maximum potential - thus creating precedent-setting transparency regime that reveals the effects automated hiring tools have on workforce diversity - the answers to these questions are simple: (1) Bias audits must be conducted on all forms of hiring automation that have consequences for NYC workers, irrespective of the details of how the technology works. (2) Bias audit reports must present clear information about the extent to which automated hiring tools disadvantage minority candidates, summarized using disparate impact analysis as defined by existing anti-discrimination regulations. (3) "Independent" audits must be conducted by individuals who understand the nature of disparate impact analysis and who are not employed by the organizations that build or use the tools.

At present, the rules proposed by DCWP will fall short in driving toward the desired outcomes, making it very likely that New Yorkes will remain unaware of the real consequences of automated hiring tools. Below, we have outlined specific remarks on the draft rules published by DCWP last month.

COMMENTS ON DCWP PROPOSED RULES AS PUBLISHED ON SEPT. 22, 2022:

I. Local Law 144 should be implemented with a very inclusive definition of AEDTs that does not rely on the details of how an automated hiring tool is built or used by the employer.

SUMMARY

- The original definition of AEDT suggested that bias audits were required of any method of automated decision-making that collects data from job applicants,

[85] Griggs v. Duke Power Co. - 401 U.S. 424, 91 S. Ct. 849 (1971)

analyzes it using a computation process, and provides the employers with an interpretive output.

- At present, DCWP's proposed rules imply an effort by the agency to limit the definition of "automated" hiring technologies in a manner that will render many (if not most) hiring tools exempt from bias audits.
- **We strongly urge regulators to recognize that an overly-specific definition for AEDTs is inappropriate and will vastly dilute the utility of this law.**

OVERVIEW OF LANGUAGE IN PROPOSED RULES

Arguably the most troubling issue with DCWP's proposed language is the agency's apparent efforts to limit the definition of AEDT such that only certain forms of hiring technology are required to be subject to bias audits.

The proposed rules establish a less-inclusive interpretation of the law by presenting sub-definitions for two phrases that appear in the original definition[86] of AEDT:

➢ "Machine learning, statistical modeling, data analytics, or artificial intelligence" is proposed to mean:

"a group of mathematical, computer-based techniques: 1. 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 2. 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 3. for which the inputs and parameters are refined through cross-validation or by using training and testing data."

➢ "to substantially assist or replace discretionary decision-making" is proposed to mean:

"to rely solely on a simplified output (score, tag, classification, ranking, etc.), with no other factors considered, or to use a simplified output as one of a set of criteria where the output is weighted more than any other criterion in the set, or to use a

[86] Prior to the release of these rules, the AEDT was defined in Local Law 144 as follows: "Automated employment decision tool. 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. The term "automated employment decision tool" does not include a tool that does not automate, support, substantially assist or replace discretionary decision-making processes and that does not materially impact natural persons, including, but not limited to, a junk email filter, firewall, antivirus software, calculator, spreadsheet, database, data set, or other compilation of data."

simplified output to overrule or modify conclusions derived from other factors including human decision-making."

With the additional sub-definitions for these phrases, whether or not a bias audit is necessary under this law would depend on: (1) how the decision-making system is built and/or (2) how the decision-making system's recommendations are used within the employer's hiring process.

From the perspective of regulators, the notion of requiring "audits" of any kind likely raises concerns that companies operating in NYC will be subject to undue burdens. Limiting the scope of the mandate to only the riskiest and most consequential technologies is one possible strategy for avoiding unnecessary "red tape." According to the proposed rules, hiring tools should be "flagged" as requiring a bias audit if they are designed using significant input from computers and/or used to screen job candidates with little or no input from humans.

REACTIONS

While DCWP's effort to narrow the definition of an AEDT may seem reasonable at first glance, the strategy is inappropriate given the nature of employment selection tools and how they are used in practice. **Implementation of Local Law 144 must account for two facts: (1) Automated hiring tools are extremely varied from a technological perspective**, and **(2) Such tools are rarely implemented in a way that completely removes human involvement.**

Regarding the former point, the instruments and methods underlying automated hiring tools are remarkably heterogeneous, ranging from multiple-choice cognitive ability screens developed in the early twentieth century to AI algorithms capable of analyzing digital work samples. In light of this immense variation, it is impossible to characterize particular technical specifications as more or less associated with disparate impact.

Regarding the latter point, employers are generally disinclined to fully deferring to the outputs provided by any HR technology, and emerging regulations requiring human oversight in high-stakes decision-making contexts make such an approach even less likely. Accordingly, there is little to be gained from attempting to narrow the definition of AEDTs to implement bias audits only in the most extreme instances of automation.

Below, we detail the undesirable consequences that will likely emerge if regulators choose to ignore these defining features of HR technology.

First, making the necessity of bias audits contingent on how the tool is built will create several perverse incentives around HR technology. Three are described below.

1. *The oldest (and often most biased) forms of automation will be advantaged.*

   Because employment selection tools are not amenable to hard and fast rules about what makes a system biased, efforts to narrow the definition of "machine learning, statistical modeling, data analytics, or artificial intelligence" in Local Law 144 will create arbitrary loopholes. In all likelihood, these loopholes will advantage the oldest forms of automated decision-making, which are often also the most biased.

   Consider the Wonderlic Personnel Test, which was developed in 1939 by a Northwestern grad student who identified a set of verbal reasoning, math, and vocabulary questions that could serve as a short-form IQ test.[87]An automated version of this assessment is still used by many large companies that have a significant presence in NYC, ranging from financial services firms to chain restaurants.[88] Cognitive ability tests like the Wonderlic are well-known among experts to be particularly likely to disadvantage racial minorities.[89]

   Despite its indisputable history of blocking progress toward workforce diversity, the Wonderlic test could be interpreted as exempt from Local Law 144 according to subpoint 2 of the proposed definition of "machine learning, statistical modeling, data analytics, or artificial intelligence." Rather than a computer being responsible for selecting the tool's inputs, a human identified them, albeit albeit 80 years ago.

   Of course, from the perspective of New York City workers, it is entirely irrelevant if a decision-making tool was built manually or derived using more sophisticated methods like machine learning. Local Law 144 will

[87] Wonderlic, E. F., & Hovland, C. I. (1939). The Personnel Test: a restandardized abridgment of the Otis S-A test for business and industrial use. *Journal of Applied Psychology, 23*(6), 685–702. https://doi.org/10.1037/h0056432

[88] Wonderlic's customer testimonials include companies like Duracell, Aveda, and Subway. Websites like Wall Street Oasis, which provide descriptions of hiring processes from real job candidates, also mention the Wonderlic being used by hedge funds like Citadel and Point72.

[89] Cottrell, J. M., Newman, D. A., & Roisman, G. I. (2015). Explaining the black-white gap in cognitive test scores: Toward a theory of adverse impact. *The Journal of applied psychology*, *100*(6), 1713–1736. https://doi.org/10.1037/apl0000020

therefore fall short in achieving its transparency objectives if such a broad loophole is permitted.

2. *Technological innovation will be curbed, even when it can facilitate fair and effective hiring.*

   In addition to advantaging older forms of automation, DCWP's proposed definition for "machine learning, statistical modeling, data analytics, or artificial intelligence" will create a perception of employers being "punished" for attempts to apply modern methods to the hiring process, irrespective of whether the technology is biased or fair. This view has been prominently articulated by dominant law firms advising NYC employers in recent months. For example, one describes Local Law 144 as an effort to apply "significant compliance obligations" to "employers who use *artificial intelligence* technologies to recruit and hire candidates" (emphasis added).[90]

   If NYC employers are given the option of avoiding bias audits altogether by simply abandoning innovations in hiring technology, they will inevitably take it. Large organizations have become very comfortable with the opacity surrounding bias in HR technology over the course of decades. Lawyers generally advise that the aggregation of demographically-labeled data (like that required in a disparate impact analysis) should be avoided, even when such information can be used "proactively to develop accountability systems and promote constructive structural change."[91]

   Despite the fact that many NYC employers recognize the need for new solutions to make real progress on workforce diversity, few will voluntarily open themselves up to unnecessary scrutiny for the sake of fairer hiring. If the city's employers abandon AI/ML employment procedures, New York workers will be excluded from many promising developments in the HR tech industry. Over the last decade, various systems built with AI/ML have been proven to result in significantly less disparate impact than

[90] *Robots beware: NYC cracks down on artificial intelligence use in the workplace*. Fisher Phillips. (2021, December 21).

[91] "Fear of liability for violation of ambiguous legal norms induces firms to adopt strategies that reduce the short-term risk of legal exposure rather than strategies that address the underlying problem. They accomplish this in significant part by discouraging the production of information that will reveal problems, except in the context of preparation for litigation…It thus encourages these important legal actors to overlook ways of using that information proactively to develop accountability systems and promote constructive structural change." See: Sturm, S. (2004). 17. Second Generation Employment Discrimination: A Structural Approach. In *Women and the US Constitution* (pp. 292-313). Columbia University Press.

traditional automated tools.[92,93,94]

Unfortunately, if adoption of newer and fairer technology is viewed as a compliance burden by employers, vendors working on these products may be persuaded to abandon their efforts. As improvements to hiring technology are broadly disincentivized, and progress is stifled, workers will continue to be screened using the same systems that have preserved historical inequalities for decades.

3. *Some HR tech vendors will attempt to "game the system" by introducing suboptimal AI/ML development processes.*

   The perverse incentives created by DCWPs proposed definition of "machine learning, statistical modeling, data analytics, or artificial intelligence" will also apply to companies that build automated hiring tools. For HR tech vendors, regulators' attempts to specify which technical aspects of a hiring tool trigger a bias audit could be interpreted as an opportunity to intentionally design a system to fall outside the scope of Local Law 144.

   While it is difficult to precisely predict what workarounds might emerge from companies seeking to avoid bias audits, data scientists we spoke to were able to immediately list several. For example (add)

Second, characterizing only the most extreme forms of automation as AEDTs (e.g., where outputs are relied upon with minimal human discretion) will make the requirements of Local Law 144 irrelevant for most employers. Two reasons are described below.

1. *Some level of human involvement in interpreting automated decision-making outputs is very common in the hiring context, but it seldom has implications for disparate impact.*

The proposed definition for "to substantially assist or replace discretionary decision making" is misaligned with the reality of how most employers implement

92 Mullainathan, S. (2019). Biased algorithms are easier to fix than biased people. *The New York Times*.

93 Kassir, S., Baker, L., Dolphin, J., & Polli, F. (2022). AI for hiring in context: a perspective on overcoming the unique challenges of employment research to mitigate disparate impact. *AI and Ethics*, 1-24.

94 Will, P., Krpan, D., & Lordan, G. (2022). People versus machines: introducing the HIRE framework. *Artificial Intelligence Review*, 1-30.

automated hiring tools. In particular, the definition suggests that a significant proportion of employers rely exclusively on the outputs of a single automated system to evaluate candidates without the possibility of human oversight or input. Because such practices are actually quite uncommon,[95,96]DCWP's proposed definition will create a very large loophole for most organizations to claim that they need not provide public audit reports.

One major reason employers are unlikely to "rely solely on a simplified output, with no other factors considered" or to "use a simplified output to overrule or modify conclusions derived from…human-decision making" is the existence of regulations and guidelines prohibiting such practices in other jurisdictions. According to Article 22 of EU and UK GDPR, individuals who are evaluated by "solely automated" decision-making tools are entitled to request human intervention that is "meaningful" and not merely a procedural "token gesture." Given that many NYC employers and HR tech vendors operate in these jurisdictions, most already have controls in place to ensure decision-making systems are not considered "solely automated." The White House's "Blueprint for an AI Bill of RIghts" specifically states that "those impacted by an automated system should be given a brief, clear notice that they are entitled to opt-out, along with clear instructions." Employers who take fairly basic steps to adhere to these increasingly common tenets could then easily make the case that their automated hiring tools do not meet the definition of an AEDT per Local Law 144.

It is worth emphasizing that, while the fact that more automated decision-making systems are being implemented with human oversight is a good thing, it is also an inappropriate basis for gauging whether a hiring tool should be subject to a disparate impact analysis. Bias audits are directed at providing information about the *general* tendency of a decision-making tool to disadvantage minority candidates. Even when human oversight technically exists in an automated system, this does not meaningfully change things for the vast majority of people. Further, given the fact that human recruiters rely on implicit biases and social

[95] "The hiring processes supported by these systems are not fully automated. Human decision-makers must choose whether and how to define an ideal candidate, at what stage of hiring to use the tool, and how to incorporate tool outputs into hiring decisions." See: Rhea, A. K., Markey, K., D'Arinzo, L., Schellmann, H., Sloane, M., Squires, P., & Stoyanovich, J. (2022). External Stability Auditing to Test the Validity of Personality Prediction in AI Hiring. *arXiv preprint arXiv:2201.09151*.

[96] "Most hiring tools today are not fully automated... Instead, AI-based tools for hiring are developed to replace specific intermediate steps in the end-to- end hiring process." See: Buyl, M., Cociancig, C., Frattone, C., & Roekens, N. (2022, June). Tackling Algorithmic Disability Discrimination in the Hiring Process: An Ethical, Legal and Technical Analysis. In *2022 ACM Conference on Fairness, Accountability, and Transparency* (pp. 1071-1082).

stereotypes to evaluate candidates, there is little reason to assume that human input will mitigate disparate impact.

2. *Establishing which automated decision-making outputs are most heavily weighted in a hiring process is generally not possible.*

The proposed definition of the phrase "to substantially assist or replace discretionary decision making" is also incorrect in its suggestion that employers typically collect all information about a candidate at once and then apply a systematic weighting, using a heuristic that does not change over time. In an ideal world then, employers would ensure that bias audits are conducted on whichever outputs they weigh most heavily. However, this does not translate in reality because many employers approach hiring as a series of "funnels" designed to winnow down the candidate pool. If the candidate pool is being narrowed at each "funnel" by an automated decision-making system, all stages of the process have implications for disparate impact. It would therefore be incorrect to characterize the simplified outputs from one tool as more or less significant than others. A further complication comes from the fact that an employer could easily decide to change the sequence of the funnels or to adapt how the outputs are interpreted between hiring cycles.

The notion of limiting bias audits to only those AEDTs that are "weighted more than any other criteria" or used to "overrule or modify conclusions" is also not supported by federal regulations regarding disparate impact. The EEOC's Uniform Guidelines on Employment Selection Procedures, which provide employers with detailed information about how they can proactively comply with Title VII of the Civil Rights Act, is clear on this point. In the event that an employer's *overall* hiring process appears to be disadvantageous to minority applicants (as many do today), "the individual components of the selection process should be evaluated for adverse impact."[23]

RECOMMENDATIONS

[23] 29 C.F.R. § 1607 (2020).

DCWP should not attempt to conceive of "automated" hiring tools as only those derived using machine learning (or attempt to define machine learning itself within this law).

Proposed language: *'Machine learning, statistical modeling, data analytics, or artificial intelligence'* means a group of mathematical, computer-based techniques that are software-coded to translate input data to a 'simplified output.'

DCWP should also not permit hiring tools that are implemented "to substantially assist or replace discretionary decision making" to be exempt from bias audits on the basis of some human oversight existing. Bias audits are about the *general* effects of different tools, regardless of how employers implement them.

Proposed language: '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 assume a simplified output (score, tag, classification, ranking, etc.) serves as a valid inference.

To further clarify the scope of AEDTs, DCWP should also consider specifying that all hiring procedures considered relevant by EEOC/OFCCP regulations are subject to this law.

II. Local Law 144 should require disparate impact be calculated in a manner that aligns with methods enshrined in federal law in terms of both statistical methodologies and relevant demographic categories.

SUMMARY

- The original definition of a bias audit included "the testing of an automated employment decision tool to assess the tool's disparate impact on persons of any component 1 category required to be reported by employers pursuant to subsection (c) of section 2000e-8 of title 42 of the United States code as specified in part 1602.7 of title 29 of the code of federal regulations."
- At present, DCWP's proposed rules include two approaches to the assessment of disparate impact. One of these approaches (i.e., producing a ratio of average scores for demographic categories) is not actually a form of disparate impact analysis because it does not provide practical information about whether an automated decision-making system results in systemic disadvantage. Additionally, DCWP's proposed rules suggest that a disparate impact analysis may be conducted on intersectional identity groups; this methodology does not align with federal employment law and is unlikely to produce robust statistical results.

- **To ensure that bias audits provide consistent and meaningful information, we strongly urge NYC regulators to require bias audits to include disparate impact analysis methods that emphasize the *consequences* of hiring tools, per existing federal regulations. Additionally, the disparate impact analyses should be conducted separately for gender and race/ethnic comparisons.**

OVERVIEW OF LANGUAGE IN PROPOSED RULES

As Local Law 144 is essentially a public disclosure regime, the law is only as useful as the information included in bias audit reports. In the context of promoting workforce diversity, the most useful information reveals the *effects* of a given hiring tool for different demographic groups.

In the proposed rules, DCWP has outlined two possibilities for calculating disparate impact:

- ➢ The first - *"the selection rate for a category divided by the rate of the most selected category"* - is in line with established methods.
- ➢ The second - *"the average score of all individuals in a category divided by the average score of individuals in the highest category"* - will fall short *consequences* of an automated hiring system.

The second option for calculating disparate impact is directed at automated hiring tools that provide "scores" for job candidates, rather than "pass/fail" classifications.

DCWP's proposed rules also a definition of "selection rate," but then incorrectly suggests that disparate impact analyses should involve comparisons of intersectional identity groups. This suggestion is made in the example offered subsequent to the definition:

- ➢ *"e.g., If 100 Hispanic women apply for a position and 40 are selected for an interview after use of an AEDT, the selection rate for Hispanic women is 40/100 or 40%."*

REACTIONS

Regarding statistical methodology for calculating disparate impact:

Automated hiring systems may issue a wide variety of outputs, but in the context of disparate impact, the relevant question is quite simple: "Is the candidate getting a positive or negative evaluation?" Stated differently, even if a tool assigns a numerical score to candidates, that score must ultimately be interpreted in the screening process (e.g., those who receive at least a 6 out of 10 "pass"). The differential effects of that *interpretation* is the important part in a disparate impact analysis, because this speaks to whether a given demographic group is disadvantaged by the system.

The convention in employment selection is to focus on how a decision-making system classifies individuals into yes/no categories, rather than comparing average group-level scores. This is because federal regulation strongly encourages employers to ensure hiring tools are properly "validated," meaning they are backed by empirical evidence of a relationship to job performance. Conducting validity studies in the employment context is difficult for a few reasons, including small sample sizes and subjective job performance metrics, which result in "noisy" data. Because of these conditions, it is very unusual for a validity study to prove that individuals who score an 89/100 on a hiring assessment outperform those who score an 88/100. A more common approach is to demonstrate that scores of 80/100 or higher are associated with success.

Importantly, the consequences of automated decision-making tools can easily be obfuscated when selection rates are not taken into account. Consider a test on which Black candidates receive an average score of 70/100 and White candidates receive an average of 85/100. According to the draft language for §5-301(b), the ratio reported in the bias audit for this test could be calculated as 75 ÷ 85 = 0.82, suggesting it is not especially biased. However, if the same test is generally implemented such that a score of 80/100 is considered passing, it may be that only 30% of Black candidates fall above the threshold, compared to 50% of White candidates. In this case, the appropriate disparate impact analysis is calculated as 0.30/0.50 = 0.60, which suggests a significantly more biased scenario.

Regarding comparisons of intersectional identity groups:

Local Law 144 touches on an important distinction regarding disparate impact analyses: one that is conducted by an employer or vendor seeking to *proactively* comply with federal law versus one that is introduced as prima facie evidence of a charge by a plaintiff.[24]

On the former point, guidance on how employers should proactively design and evaluate their hiring procedures to ensure they are not violating anti-discrimination law is specified by the EEOC's Uniform Guidelines on Employee Selection Procedures (UGESP). According to UGESP, employers should monitor whether their hiring procedures result in disparate selection rates for demographic groups making up at least 2% of the relevant applicant pool.[25] This threshold is understandable; the point is for employers to react to robust findings of a potentially discriminatory hiring tool, rather than continuously being alarmed by disparities that are likely to be statistical anomalies. Efforts to include intersectional identity groups under Local Law 144 will have a similarly unproductive consequence: if bias audits present a large number of identity groups, and those groups each represent fairly small numbers of candidates, reports will mostly reflect insignificant findings.

RECOMMENDATIONS

DCWP should allow the disparate impact analyses included in a bias audit to be conducted in one of two ways.

[24] While bias audit reports should not *proactively* include disparate impact comparisons for intersectional groups, we would like to note that this does not preclude the possibility of such discrimination occuring against groups like black men, black women, hispanic men, or hispanic women. Members of these groups have rights of action under federal law, just like single-identity groups. We are only suggesting that transparency mandates should prioritize group-level comparisons that are likely to yield statistically significant results.

[25] 29 C.F.R. § 1607.15(a)(2). "Adverse impact determinations should be made at least annually for each such group which constitutes at least 2 percent of the labor force in the relevant labor area or 2 percent of the applicable workforce."

1. Comparison of selection rates, which apply when (a) a tool that classify individuals into pass (i.e., "recommended") and fail ("not recommended") categories OR (b) a tool assigns scores that are then interpreted as using a cutoff, where all scores above the specified threshold are viewed as equivalent.
    - Regarding public reporting: bias audit reports should include details of how candidates who receive a positive evaluation from the tool are distinguished from those receiving a negative evaluation. (e.g., An output of 1-100 is provided for each candidate, all candidates scoring above 70 are considered likely to succeed in the job.)
    - Note that some HR tech vendors may provide employers with multiple options for interpreting outputs. It may be advisable for a third-party auditor to conduct multiple disparate impact analyses

in accordance with the interpretive strategies most commonly used by employers.

2. Comparison of group-level averages **using a two-sample t-test**, which accounts for differences in the sample sizes and standard deviations observed for each group. In contrast to the simplistic ratio proposed by §5-301(b), federal precedent suggests this more robust methodology.
   - Regarding public reporting: reports focused on disparities in average scores must also report group-level standard deviations.

DCWP should also specify that disparate impact analyses are conducted and reported separately for gender and race/ethnicity groups, as specified by EEO component-1 categories.

III. Independent auditors should specifically be defined as individuals who are not employed by the organization that builds, sells, or uses the AEDT.

- The original definition of a bias audit described "an impartial evaluation by an independent auditor."
- DCWP's proposed rules define an independent auditor as "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."
- We encourage regulators to specify that auditors cannot be employed by the organizations using or developing the AEDT in question, even if their work takes place in a different division/department.
- **We offer the following definition for consideration:**
  - Independent auditor. The term "independent auditor" means a person whose knowledge, skills, experience, training, or education qualifies them to conduct statistical analyses related to discrimination.[97] An independent auditor cannot be directly employed by the organization using the relevant tool in its employment process OR (if applicable) any third-party vendor involved in the development or sale of the tool. An independent auditor is required to provide attestment to DCWP (1) that they intended to conduct all evaluations using established methods and to apply these methods reliably, (2) that financial incentives did not in any way influence

[97] See Federal Rules of Evidence. Rule 702. "A witness who is qualified as an expert by knowledge, skill, experience, training, or education may testify in the form of an opinion or otherwise if: (a) the expert's scientific, technical, or other specialized knowledge will help the trier of fact to understand the evidence or to determine a fact in issue; (b) the testimony is based on sufficient facts or data; (c) the testimony is the product of reliable principles and methods; and (d) the expert has reliably applied the principles and methods to the facts of the case."

their evaluation or reported findings, and (3) the amount they were paid for the auditing service and by whom. This attestment must be submitted to DCWP prior to the date the bias audit is made public on any employer or employment agency's online hiring website.

IV. Bias audit reports should include information that clearly describes and contextualizes the nature of any data samples used. This should be the case whether an audit is conducted at the level of a vendor or that of an individual employer. Please see the below sample for a vendor-level example report.

SAMPLE REPORT

Basic Information:

1.	**AEDT vendor name**	(insert company name)
2.	**AEDT name (if any)**	(insert product name)
3.	**Description of AEDT**	(e.g., Assessment type? Nature of data inputs? Note that relevant description may come from marketing materials for the AEDT.)
4.	**Description of job characteristics measured by AEDT**	(e.g., occupation-specific tasks, occupational requirements/work activities, worker requirements/skills)
5.	**Hiring contexts in which AEDT is typically used (types of employer organizations, industries, role types, etc.)**	(insert)

Scope of AEDT Use and Data Availability:

6.	**Estimate the total number of individuals evaluated by the AEDT in the last 12 months.**	(Specify if across all use cases, particular geography, etc.)
7.	**Estimate the number of individuals evaluated by the AEDT in the last month who were applying for positions with employers/employment agencies subject to Local Law No. 144.**	(Relevant employers must employ at least 250 people total and at least 10 who reside in NYC for the purpose of income taxes)

8.	**What data does the vendor have available regarding the outcomes generated by the AEDT for the population described in Q7? Explain discrepancies.**	(insert)
9.	**For what proportion of the data described in Q8 are demographic labels in accordance with EEO Component-1 categories available?**	(insert)
10.	**What steps, if any, is the vendor taking to improve the representativeness of available data?**	(insert)

Disparate Impact Analysis:

10.	**Describe the nature of the sample used to conduct disparate impact analyses included in this report.**	(Note that this may be redundant from Data Provenance section. Sample size must be reported.)
11.	**Describe the interpretive strategies used to determine whether individuals included in the sample (see Q10) are "passed" or "failed" by the AEDT.**	(e.g., All candidates scoring above 50th percentile are considered "passed.")
12.	**Provide the proportion of all Component-1 demographic groups represented in the sample described in Q10-11.**	(insert percentage breakdowns)

Disparate Impact Results: (illustrative data)

		Gender		Race/Ethnicity					
	All	**Male**	**Female**	**White**	**Black**	**Hispanic or Latino**	**Native Hawaiian or Pacific Islander**	**Asian**	**Native American or Alaskan Native**
n Total	20,563	11,234	9,329	12,112	2,245	3,915	103	2,034	154
n Selected (passed by AEDT)	8,147	4,607	3,540	4784	834	1345	59	1,054	71
n NOT Selected (failed by AEDT)	12,416	6,627	5,789	7,328	1,411	2,570	44	980	83

Selection Rate	0.40	0.41	0.38	0.39	0.37	0.34	0.57	0.52	0.46
Impact Ratio (Compared to Highest Selection Rate)		1.00	0.925	0.76	0.72	0.66	Excluded due to group representing less than 2% of total sample	1.00	Excluded due to group representing less than 2% of total sample
Impact Ratio (Compared to White/Male Selection Rate)		**1.00**	**0.925**	**1.00**	**0.94**	**0.87**		**1.31**	

Availability of Alternatives/Accommodations:

13.	**Please describe what is measured by the AEDT, with an emphasis on information that candidates would find relevant for requesting alternative/accommodated hiring procedures.**	
14.	**What, if any, procedures does the AEDT vendor have in place to process accommodations requests?**	
15.	**What, if any, procedures does the AEDT vendor contractually require of employers to have in place to process accommodations requests?**	
16.	**Please list any disabilities, limitations, or work-related functions the AEDT CAN accommodate. (See Job Action Network)**	
17.	**Please list any disabilities, limitations, or work-related functions the AEDT CANNOT accommodate. (See Job Action Network)**	

Auditor Information:

13.	**Auditor name**	(insert)
14.	**Auditor employer, if any**	(Is the auditor employed by a firm or consultancy that has contracted with the auditee to undergo this evaluation? If not, enter "self-employed.")

15.	**Briefly describe the auditor's knowledge, skills, experience, training, and/or education that qualifies them to conduct this evaluation.**	
16.	**Is the auditor employed by the AEDT vendor or any employer organization that uses the AEDT product described in this evaluation?**	

17.	**Has the objectivity of this evaluation been in any way affected by financial incentives or other contingencies offered by the AEDT vendor or any employer organization that uses the AEDT product described in this evaluation?**	
18.	**The auditor must sign and date attesting to the veracity of items 15, 16, and 17.**	

Stephen J. Malone, November 4, 2022

Comments Regarding New York City Department of Consumer and Worker Protection Proposed Rules Implementing Local Law 144 of 2021 (“LL 144”)

Thank you for the opportunity to submit comments on the proposed rule.

I am a practicing in-house attorney in New York City, focusing on employment law. My remarks at the November 4, 2022 hearing and in this written comment are submitted solely as an interested individual with many years of experience working with human resources professionals; I am not commenting on behalf of any organization.

My comments focus on three points: 1) the 10-day notice, 2) the scope of the AEDT definition, and 3) the effective date.

First, the Department should clarify that the ten-day notice to Candidates for Employment in Section 5-303 is only needed on an employer’s jobs website one time. It is not meant to impose a ten-day waiting period on the employer for each and every posting, or each and every Candidate for Employment. The proposed rule, as drafted, potentially could be read to impose a ten-day waiting period before using an AEDT on each and every particular candidate’s application. NYU Langone also highlighted this issue in its written comment. Please clarify that the notice may go up on the employer’s careers website one time, and then an AEDT may be used ten days thereafter on any Candidate for Employment for any posting. To require otherwise will impede hiring and make New York City employers far less competitive in a tight labor market where hiring decisions must be made quickly.

Second, the Department should provide more concrete examples to the employer community as to what AI processes require a bias audit. The prior written comments from the Seyfarth Shaw firm and the Society for Human Resources Management make this point very well. AI tools for recruiting can expand opportunities to more candidates and make the hiring process more efficient for New York employers. Yet LL144 and the proposed rules are so broad that they may deter the employer community from using any AI tools for recruiting whatsoever. I ask the Department to provide additional examples of what is inside and outside the definition of an AEDT.

Finally, I recommend that the City Council delay the January 1, 2023 effective date of the law, like it did earlier this year for the pay transparency law. By the time these proposed rules are finalized, employers will have little or no time to comply with the new law. An additional six-month delay in the effective date is warranted, or alternatively, a delay in the Department’s enforcement.

Thank you.

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.

Always Designing for People

- 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.

Online comments: 19

- **Fred Oswald, Rice University, workforce.rice.edu**
  Note that there are a number of different-yet-reasonable ways to calculate adverse impact statistics, other than the impact ratio. The 4/5ths Rule (impact ratio of < .80 as prima facie case for disparate impact) arises from the UGESP FAQ, but does not preclude these important alternatives. The 4/5ths Rule itself is asymmetric; that is, the 4/5ths Rule for protected-group selection operates differently than a 5/4ths Rule for protected-group rejection, even though conceptually, it is the same. Furthermore, the 4/5ths Rule is easier to violate at lower selection rates (e.g., 1.5% vs. 2% selection violates the 4/5ths Rule, even though this is a .5% selection rate difference) than at higher ones (e.g., 76% vs. 95% does not violate the 4/5ths rule, yet is a 19% selection rate difference).

  For alternative adverse impact indices that are possible, see the Free Adverse Impact Resource (FAIR) at https://orgtools.shinyapps.io/FAIR/, and a related book chapter:

  Oswald, F. L., Dunleavy, E., & Shaw, A. (2017). Measuring practical significance in adverse impact analysis. In S. B. Morris & E. M. Dunleavy (Eds.). Adverse impact analysis: Understanding data, statistics and risk (pp. 92-112). Routledge.

  https://www.researchgate.net/publication/314245607_Measuring_practical_significance_in_adverse_impact_analysis

  Comment added September 23, 2022 2:51pm

- **Mike Fetzer**
  The proposed rules defines independent auditor as “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." In the first example under Section 5-301 (Bias Audit), the proposed rules state "...The employer asks the vendor for a bias audit. The vendor uses historical data it has collected from employers on applicants selected for each category to conduct a bias audit..." Since the majority of AEDT vendors are also the developers of AEDTs, this seems to indicate that vendor can conduct a bias audit for the employer as long as the person or group employed by the vendor is not the same person or group that developed the AEDT.
Similarly, the second example states "...The employer uses historical data on average "culture fit" score of applicants for each category to conduct a bias audit..." This also seems to indicate that the employer can conduct their own bias audit, as long as the person or group conducting the audit is not using the AEDT.
Put another way, the proposed rules do not require that the independent auditor be a separate legal entity from either the vendor or the employer. If that is the intent, it would be beneficial to clarify within the final rules.

Comment added October 18, 2022 3:02pm

- **Mike Fetzer**
  The expanded definition of machine learning, statistical modelling, data analytics, or artificial intelligence in the proposed rules is helpful. However, it could still be interpreted to include a broad range of employment decision tools that fall outside the conventional definitions of machine learning and/or artificial intelligence (ML/AI). If the intent of LL 144 and the proposed rules are to regulate tools that do fall under the conventional definitions of ML/AI, I would recommend the following:
  (1) Remove the terms "data analytics" and "statistical modelling"
  (2) in part (ii) of the expanded definition, revise to read "for which a computer actively identifies the inputs...."
  (3) Provide additional examples of what does not fall under the intended definition of an AEDT.

Comment added October 18, 2022 4:19pm

- **Hannah Wade, NYU Langone Health**
  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.

  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 5-303) 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 October 20, 2022 9:56am

- **Mike Fetzer**
  § 5-303 of the proposed rules (Notice to Candidates and Employees) states: “(a) For the purpose of compliance with § 20-871(b)(1) and (2) of the Code, an employer or employment agency may provide notice to a candidate for employment who resides in the city by: (1) Including notice on the careers or jobs section of its website in a clear and conspicuous manner at least 10 business days prior to use of an AEDT”
  If an employer were to provide a notice on their general careers/jobs site 10 business days before posting an opening for a specific job where an AEDT would be used, would this be in compliance with LL 144?

  Comment added October 21, 2022 1:11pm

- **Mike Fetzer**
  § 5-303 of the proposed rules (Notice to Candidates and Employees) states:(c) The notice required by § 20-871(b)(1) of the Code must include instructions for how to request an alternative selection process or accommodation.
  It also state:

(e) Nothing in this subchapter requires an employer or employment agency to provide an alternative selection process.
If an employer is not required to provide an alternative, why would they be required to include instructions on how to request one?

Comment added October 21, 2022 1:14pm

- **Stephen H**
  Clarification needed around lack of data on protective attributes. Most vendors do not have existing data regarding many of the listed protective attributes, such as ethnicity or sex. Employers often choose not to share such data with the vendor, even if they do collect them, due to their sensitive nature. As a vendor trying to satisfy the requirements of the audit (not just for existing customers but also for any future customers that the legislation may apply to in the next 12 months), that makes the application of the suggested test on the output of any machine learning algorithm rather tricky without actually knowing the protective attributes of the sample population. What would be the suggested approach for vendors who do not currently receive or process data on protective attributes to satisfy the audit requirement?

  Comment added October 21, 2022 5:16pm