June 6, 2022

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

Submitted electronically via http://rules.cityofnewyork.us

RE: Local Law 1894-A (Automated Employment Decision Tools)

To the Department:

Littler Mendelson P.C.'s Workplace Policy Institute (WPI) submits these comments to the Department of Consumer and Worker Protections pursuant to the Department's Notice of Public Hearing and Opportunity to Comment on Proposed Rules. We thank the Department for the opportunity to comment on this important and timely topic, and address our comments specifically to Local Law 1894-1/144 of 2021 (herein, the "Law").

By way of background, WPI facilitates the employer community's engagement in legislative and regulatory developments that affect their workplaces and business strategies. WPI harnesses the deep subject matter expertise of Littler, the largest law firm in the world with a practice devoted exclusively to the representation of employers in employment and labor law matters. Littler's clients range from new and emerging businesses to Fortune 100 companies throughout the country and around the world.[1] In today's workplace, algorithms frequently make decisions that significantly impact people's lives, across a wide range of activities: health care, housing, lending, and the focus of WPI's comments, employment. We appreciate the opportunity to provide the Department with the benefit of our, and our clients' experience.

The Use of Algorithms Employment Tools by Employers. For employers, the development of algorithmic decision making creates both opportunities and novel issues of concern, and they generate new questions about long-time problems. This in turn potentially affects every aspect of employment decision-making for employers of all size in virtually every industry, from the selection and hiring process, through performance management and promotion decisions, and up to and beyond the time termination decisions are made, whether for performance reasons or as part of a reorganization.

[1] Indeed, to help our clients navigate the complex issues surrounding artificial intelligence, automation, and the transformative effect these and other developments are having in the workplace, Littler recently announced its Global Workforce Transformation Initiative, to assist employers, employees, and policymakers face the challenges—and opportunities—that automation and technology bring to the workplace. The Initiative's inaugural white paper, which analyses many of these issues in depth, may be found at: https://www.littler.com/files/workforce_transformation_report.pdf.

James A. Paretti, Jr.
jparetti@littler.com

Littler Mendelson, PC
815 Connecticut Avenue NW | Suite 400 | Washington, DC 20006

Today, employers can access more information about their applicant pool and workforce than ever before, and have an ability to correlate data gleaned from an application itself, perhaps supplemented by publicly available social media sources, to determine how long a candidate is likely to stay on a particular job. Conversely, by combing through computerized calendar entries and e-mail headers, by way of artificial intelligence (AI), tools exist which can indicate which employees are likely to leave their employment within the next 12 months. These new tools and methods that rely on algorithms and the aggregation and analysis of a massive amount of data are becoming part of the daily landscape in human resource departments.

Similarly, the use of algorithms to review résumés and perform other recruiting functions is becoming far more commonplace. Novel solutions include games-based tools that seek to measure aptitude, tools that conduct interviews and evaluate candidates, and tools that scrape publicly available social media content. The promise of AI-based recruiting tools is to eliminate possible implicit bias of decision-makers and expand the pool of potential candidates. In this way, firms can leverage Big Data to identify and recruit optimal candidates. Employers may also turn to predictive recruiting tools for reasons of efficiency and cost savings by automating at least part of the recruiting process and identifying quality candidates who will stay for the long term.

Equally important, AI-based tools have the potential to promote diversity, equity, and inclusion by expanding the applicant pool and focusing on candidates' abilities versus well-worn proxies for talent such as academic achievement, work history, and employee referrals, all of which are capable of perpetuating historical biases.

Legal Issues Surrounding the Use of Algorithmic Analysis. Deploying algorithmic tools is not risk-free, however, and should only be done carefully, with the involvement of all key stakeholders, and with the assistance of qualified counsel. Artificial intelligence offers a potent antidote to intentional discrimination. Antidiscrimination laws, however, also prohibit practices that are facially neutral if they have a disparate impact on members of protected categories, unless those practices are "job-related" and consistent with "business necessity."

Given the complexity of amassing and then analyzing vast quantities of information, an employer would certainly not reverse engineer the process in order to intentionally discriminate against a protected group. It is far more probable that the use of algorithms may be challenged because it unintentionally yields a disparate impact on one or more protected groups. More precisely, a plaintiff or class may allege that the algorithm used for hiring, promotion, or similar purpose adversely impacts one or more protected groups.

The legal landscape of how courts should view these challenges, and what methods of analysis they should apply, is limited, and despite the rapid advent of AI and predictive technology, employers continue to operate in a legal environment based on rules and regulations developed in an analog world with few guideposts that readily translate to the 21st century workplace. Issues may arise in a context that makes yesterday's compliance paradigm both outdated and difficult to apply.

For example, under current federal guidance, whether a selection method that produces an adverse impact passes muster under Title VII is often decided with reference to the Uniform Guidelines on Employee Selection Procedures (UGESP). The use of artificial intelligence and algorithmic screening tools can effect a shift from selection criteria distilled from job-related knowledge, skills, and abilities, leaving correlation to be established empirically, to one in which correlation is first established empirically- independently of knowledge, skills, and ability-and leaves the duration of that correlation in question. Unfortunately, because UGESP was adopted in 1978, it fails wholly to contemplate the use of algorithmic data and its reliance on correlation rather than cause-and-effect relationships.

Algorithmic analysis means that employers can theoretically analyze every aspect of every decision without worrying about a need to rely only on a partial sample, and this data allows employers to find (or, in some cases, to disprove) correlations between characteristics and outcomes that may or may not have a seeming connection. As a result, employers need to be able to understand what the use of these tools means for reducing the risks of traditional discrimination claims without giving rise to new varieties of such claims. But there are also new implications for background checks and employee privacy, data security obligations, and new theories of liability and new defenses based on statistical correlations, to name but a few.

The challenge for the legal system is to permit those engaged in the responsible development of Big Data methodologies in the employment sector to move forward and explore their possibilities without interference from guidelines and standards based on assumptions that no longer apply or that become obsolete the next year. An important part of this process is permitting employers to find and work with key business partners to assist in Big Data efforts and developing strategies that have the potential to make the workplace function more effectively for everyone. To do that well, it is vital that human resources professionals, statisticians, scientists, and a range of other stakeholders have a seat at the table when policy decisions are made regarding how and when to permit the use of these tools.

It is against this backdrop that we set forth our substantive concerns with the Law.

Concerns with the Law. We respectfully submit that mere clarification of the Law's penalty schedule is insufficient.[2] As set forth below, as currently adopted, a number of the Law's provisions threaten to limit if not wholly eliminate the use of Big Data in employment-decision making within the City, with the perverse consequences of potentially increasing subjective decision-making and discrimination; dramatically expanding employer liability; and reducing the number of available employment opportunities. We have addressed these concerns below.

Our concerns and recommendations fall into three categories: the **scope** of the law in terms of geographic reach, the tools and systems it regulates, and the employment activity it regulates; its **operational impact** in terms of efficiency, fairness, and clarity of application; and finally, the

[2] With respect to the proposed penalty schedule, we respectfully submit that the Department consider providing an administrative "cure" period by way of regulation. Such a mechanism would allow the Department to notify an employer of any alleged non-compliance with the Law, and provide a period of time in which an employer could rectify any alleged deficiencies before being assessed with any fine.

potential for **misapplication** leading to conflict with established principles of nondiscrimination law.

Scope of the Law

Location. The Law appears intended to apply only to those candidates for employment or promotion who work (or will work) within New York City ("Sec. 20-870: The term 'employment decision' means to screen candidates *for employment or employees for promotion within the city*." (emphasis added)). However, there is potential ambiguity in the interpretation of this specific definition. We urge that the Department make clear that this is the extent of the Law's requirements, and expressly provide that the Law does *not* apply to applicants or employees who do not work (or will not work) in New York City (or who work or will work within the City only casually, occasionally, or sporadically). Similarly, consistent with the Law's text, we urge that the Department state clearly that the Law does *not* apply to employment decisions made within New York City with respect to applicants and candidates for positions outside of the City's borders.

> *Recommended Revision*:
> Employment Decision. The term "employment decision" means to screen candidates for employment or employees for promotion, where both the individual is located and the employment opportunity is intended to be performed within the city.

Promotion. The application of the Law to promotion activity is likely to undermine and upend internal promotion processes that are already subject to anti-discrimination laws both state and federal, without generating any likely improvement in outcome. This is because of three fundamental differences between hiring and promotion. Hiring decisions typically result from an assessment of large numbers of candidates for each opening, while promotions are typically narrower, more individualized decisions. In addition, unlike the swiftness and potential opacity of hiring decisions, promotion decisions are typically part of an ongoing, feedback-driven process based on well-defined job requirements, expectations, and performance criteria that are well known to the employee.

> *Recommended Further Revision*:
> Employment Decision. The term "employment decision" means to screen external candidates ~~for employment~~ or current employees for ~~promotion~~ employment, where both the individual is located and the employment opportunity is intended to be performed within the city.

"Automated Employment Decision Tool." The Law's definition of "automated employment decision tool" as "any computational process, derived from machine learning, statistical modeling, data analytics, or artificial intelligence" is susceptible to an unnecessarily over-broad reading, one which would apply the Law to just about any screening tool an employer might use to assess and rank candidates or employees, *including traditional pencil-and-paper tests*. Critically, the law does not provide any reference point for employers to determine when simple mechanical application of a tool crosses into the realm of "machine learning" or "artificial intelligence." Consider the following progression of scenarios:

1. Acme Corp. gives candidates for an accounting position a pencil-and-paper math test.
2. Acme Corp. uses data analytics to identify the key math errors its accountants typically make, and revises the pencil-and-paper test to use problems that test those concepts.
3. Statistical modelling shows that computer scoring of the test is more accurate than human scoring, so Acme Corp. converts the test into an online format *with the same questions*.
4. Acme Corp. replaces the static question-and-answer format of its accounting test with an interactive math 'game' where a machine learning algorithm customizes the questions put to a candidate based on that candidate's prior responses, and then rates their skill.

Scenario 1 is currently regulated by existing employment laws and regulations. At the other extreme, Scenario 4 appears to be squarely the type of automated decision-making that the Law is intended to regulate. The trouble is in the middle. The test in Scenario 2 is "derived from … data analytics," and is thus subject to the Law, *even though it is just another pencil-and-paper test*. Similarly, the test in Scenario 3, being derived from an application of statistical modelling, is subject to the Law, even though it is *merely an electronic reformatting of a pencil-and-paper test*. We expect that the Council did not intend these counterintuitive outcomes, and therefore urge the Department to adopt a more meaningful definition of "Automated Employment Decision Tool," one that differentiates between the rote application of programs to perform pre-determine, human-driven calculations; pre-programmed, static mimicking of human actions; and true "machine learning" in which a program adapts its behaviors and programming absent human intervention.

> *Recommended Revision*:
> Automated employment decision tool. The term "automated employment decision tool" means any computational process~~, derived from~~ that actively uses machine learning, or artificial intelligence, that issues simplified output, including a score, classification, or recommendation … .

We also urge the Department to define the terms "computational process," "statistical modeling," "data analytics," "machine learning" and "artificial intelligence." The differences in vernacular and technical uses of the latter term, particularly, makes it fertile for threshold disputes on the application (and misapplication) of the Law.

We respectfully add that the Law's list of excluded tools is underinclusive and that the Law would benefit from further clarification of the "computational processes" at issue, as well as an elimination of the list. For example, the Law notes that junk mail filters, firewalls, calculators, spreadsheets, databases, data sets, and any "other compilation of data" are *not* automated decision tools. Insofar as these tools do not purport to make or inform any employment-related decision without human intervention, that would seem axiomatic. However, the tools enumerated *could* be part of a larger Automated Employment Decision Tool or process – the manner of use is what matters. For example, even though it "issues simplified output," the rudimentary use of a filter or search function within résumé database is no different than searching through paper copies of the same résumés – it's just quicker and more accurate. The only "automation" involved is simple rote filtering to find candidates who meet human-specified criteria, and that automation

is not, we submit, the issue (just as the automation of a pencil-and-paper test is not, by itself, a concern). Applying machine learning algorithms to the same résumé database to rank or score candidates, on the other hand, seems the type of activity the Law aims to regulate.

> *Recommended Further Revision*:
> Automated employment decision tool. The term "automated employment decision tool" means any computational process~~, derived from~~ that actively uses machine learning, 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~~.[3] ~~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.~~

Finally, we are particularly concerned with the application of the Law to a range of online job-seeking databases and websites, as the Law itself is not clear as to what is contemplated. If an employer goes to a website where job-seekers post their résumés, a human representative enters a series of keywords to match, and the website displays all applicants with that keyword on their résumé/c.v., this would plainly seem to be outside the realm of automated decision making, and we urge the Department to make this clear. Taking the example a step further, however, makes it even less clear how the Law applies to these websites. Suppose, for example, rather than entering keywords for the website to match and return, the employer's representative enters a job-title, or skill set – asking for candidates for "office manager" or "social media administrator," or seeking workers who have experience in "marketing" and "website design." Based on its own program, the website processes the request, and returns a list of candidates to the employer. It may or may not rank this list, and the employer has no knowledge of how the website translated its "plain English" request into a search-and-sort function. Similarly, the employer has no knowledge as to whether and how this proprietary software (which it does not own and whose code and algorithms it has no access to) has developed this ranked list. In light of these facts, we urge the Department to specify in any final regulations that the requirements of the Law apply only to those tools over which an employer has control or proprietary access, and not to routine career-search websites which an employer accesses only as a third-party user.

Operational Impact of the Law

Ten Business Day Notice Period. The Law can be read to require that an applicant or employee be given no less than ten business days' notice before the use of a subject screening tool, and that such employee or applicant be given the opportunity to request "an alternative selection process or accommodation."

[3] We recommend striking the limitation to natural persons because the preceding term, "employment decisions," is already a defined term of art within this law, and its definition already inheres this limitation.

Read in this way as a practical matter, this means that a candidate must be given 10 days to elect a subjective assessment, so that a job scheduled to remain open for 30 days must remain open for 40, so a late-arriving applicant can take advantage of this 10-day window before the job is filled, and potentially further extend the "open" window beyond 40 days. Indeed, where more than one such latecomer seeks to apply to a position, the required "window" may be extended repeatedly. To guard against this scenario, we recommend that any final regulation make clear that an employer is not required to continuously extend a job posting to accommodate applicants who appear only in an "extended" window.

Perhaps more important, the Law wholly fails to contemplate the business needs of employers who seek to engage temporary or contingent workers immediately (and the workers who want and need these jobs). In that respect, the Law threatens the fundamental viability of staffing agencies and other businesses whose principal operation is the ability to fill jobs quickly, often with little to no advance notice. Given the volume of applicants these firms must screen on an expedited basis, a 10-day "notice" requirement would likely result in these firms ceasing to advertise, screen, and fill positions in performed in New York City. For these reasons, we urge the Department to clarify that the 10-day notice requirement extends only in those instances where a job is intended to be posted for at least 10 days, and that in all instances, an employer is required to give only that notice which is practicable given the nature of the position and the time intended for it to be open.

Notice of Job Qualifications and Characteristics. Often, by their very nature, automated screening tools driven by machine learning, statistical modelling and similar data-analytical AI techniques do not base their results on a specific qualification or characteristic, but rather by analyzing a vast amount of data and determining whether and how these data points may predict success (or likely success) in a position. Thus, while the job qualifications and characteristics present on the corresponding job description will, in a manual process, directly frame the a human recruiter's analysis of a résumé or application (setting aside instances of discriminatory intent), in a machine-learning process the analysis of the same content is conducted using a mathematical framework that identifies and weights patterns within the data that are neither discernible to the human mind nor humanly computable in the given time. Critically, these patterns and weights are not driven by cause-and-effect.

By way of example, assume an algorithm assesses data about an employer's existing workforce to attempt to screen applicants and rank them based on likely success retention in a position. The algorithm may determine with statistical significance that employees with a degree in a STEM concentration are—for whatever reason—statistically more likely to perform well, and stay in the position several months longer than candidates with a degree in a liberal arts concentration. The algorithm cannot explain why (nor is it tasked to) – it only proves that this trait correlates to stronger performance (the employer-user is likely also unaware of this correlation, knowing only that the algorithm has provided candidates statistically more likely to perform well and remain in the job). In these instances, it is not clear what information an employer is expected to provide to employees and applicants subject to the algorithmic tool—that the tool screens for successful performance? Likely retention? Undergraduate concentration? Moreover, the network of weighted patterns that such a tool generates is not reducible to a straightforward English-language description.

Thus, all that an employer may know (and be able to describe) is the information about the job and the candidate that is accessible to the tool (such as, for instance, the job description, the candidate's résumé, and their responses to questions on the application). Accordingly, we recommend that the Department limit the substantive aspect of the Notice to those categories of information.

> *Recommended Revision*:
> The job-specific ~~qualifications and characteristics~~ and candidate-specific data that such automated employment decision tool will use in the assessment of such candidate or employee.

Alternative Assessment. In connection with the notice requirement, the Law also requires that applicants be given the opportunity to "request an alternative selection process or accommodation" but offers no insight into what this legally or practically requires. If an employer uses a screening tool to winnow 7,000 résumés for an open position to 50 for human review, is it required to offer the opportunity for human review to all of the 6,950 individuals who ask? Is the employer required to offer these individuals an in-person screening interview, even if they otherwise would not have qualified for one? We urge the Department to make clear that it is not.

Continuing, even assuming that such an exercise would be possible (which, in many instances, it simply will not be), the Law provides no guidance on how to decide between an employee who receives a glowing subjective assessment and another who scores highest based on data analytics. Is the employer under any obligation to weigh an "alternative assessment" differently than other candidates subject to the screen? Again, we urge the Department to specify that an employer is under no affirmative obligation to prefer a subjectively-screened candidate over one which is highlighted by an automated employment tool, and that an employer may (as in all instances) rely on its business judgment in selecting a candidate or candidates.

Bias Audit. The Law defines "bias audit" as "an impartial evaluation by an independent auditor," and requires that any tool be tested to assess for "disparate impact" on the basis of race, ethnicity, or sex. The Law appears to require that the tool be assessed at least annually, and that "a summary of the results" of the most recent audit be made publicly available.

This requirement raises a host of questions as to how and under what circumstances a "bias audit" will be deemed sufficient under the Law. For example, assume an employer uses the same Tool to screen applicants for a wide range of positions, from C-suite to entry-level openings. Is an employer required to "audit" the tool for bias against each position? Against a class of openings? In the aggregate? Moreover, apart from disparate impact on the basis of race, sex, and ethnicity, it is unclear what the contents of an "audit" must show, or what analysis beyond that set forth in the Law an employer is required to conduct. Finally, how would the law be applied in those circumstances where a tool has a disparate impact against one group, but is superior to any alternative process for the overwhelming majority of candidates? We submit, given the statutory text, and the lack of any further direction or context from the bill's sponsors,

that a disparate impact analysis of the sort described satisfies the requirements of the Law, and urge the Department to explicitly state so in any final regulations.

Similarly, the Law offers no guidance as to what is required in a "summary" of results. Is a simple statement that the tool has been tested in accordance with the law and shows no evidence of disparate impact on the basis of race, ethnicity, or gender sufficient? We would urge the Department to clarify that such a statement (where accurate) satisfies the requirements of the law.

Finally, we urge the Department to align the requirement for a bias audit of a tool with the Law's requirements relating to notice. Specifically, the Law requires that notice be given only where a tool is used to screen "an employee or candidate *who has applied for a position* for an employment decision" (emphasis added). The requirement for a bias audit, however, requires an audit to be conducted where an automated employment decision tool is used "to screen a candidate or employee for an employment decision." For example, consider software that transfers a résumé into a database by using machine learning or another computational process to extract pieces of information (*e.g.*, name, jobs held, degrees held, etc.). Assume, further, that the database, once created, is used and reviewed solely and directly by humans. At the point that an automated tool is being applied, then, there is no position being applied for, so that class of tool should not be within the scope of the Law. Indeed, the notion that an individual has to actually have applied for a position is implicit in the Law's textual reference to "an employment decision" – if no position is being sought, it is axiomatic that no "employment decision" has been made. Moreover, it is unclear how an audit could be conducted in this scenario, insofar as no employment decision is being made.

> *Recommended Revision*:
> § 20-871. Requirements for automated employment decision tools.
> a. In the city, it shall be unlawful for an employer or an employment agency to use an automated employment decision tool to screen a candidate who has applied for a position or employee…

Potential Misapplication of the Law

Finally, in requiring that disclosure of the results of an audit be made public, the Law appears to not contemplate the fact that under federal and cognate state civil rights laws, a screening tool may be lawfully used even if it has a disparate impact on protected groups. *See, e.g.,* 42 U.S.C. § 2000e-2(k)(1)(A)(i) (employment practice that is job related for the position in question and consistent with business necessity is lawful despite disparate impact); *id.* § (k)(1)(A)(ii) (employment practice is lawful despite disparate impact where no less discriminatory alternative practice exists). In so doing, the Law is likely to lead to frivolous charges of discrimination, increased expense for employers, and mismatched expectations for applicants and employees. We urge the Department to be mindful of this in crafting regulations, and consider explicitly making clear to applicants and employees that the mere fact that a screening tool results in some disparate impact does not make its use *per se* unlawful.

Recommended Revision:
§ 20-874 Construction. The provisions of this subchapter shall not be construed to limit any right of any candidate or employee for an employment decision to bring a civil action in any court of competent jurisdiction, or to limit the authority of the commission on human rights to enforce the provisions of title 8, in accordance with law. The provisions of this chapter shall not be construed to limit the right of an employer or employment agency, pursuant to 42 U.S.C. § 2000e-2(k)(1)(A)(i) & (ii), to use a screening tool that has a disparate impact on protected groups so long as the employment practice at issue is job related for the position in question and consistent with business necessity, or where no less discriminatory alternative practice exists.

* * *

The above highlights key issues we have identified with the Law thus far; we expect there will be others, and we welcome the opportunity to work cooperatively with the Department toward regulatory solutions that address the practical and business realities of how these tools are used in the modern workplace.

Respectfully submitted,

James A. Paretti, Jr.
Shareholder

for

LITTLER MENDELSON, P.C.
WORKPLACE POLICY INSTITUTE