TESTIMONY
Presented to the

New York City Council Committee on Technology

on the subject of **Ethical Implications of Using Artificial Intelligence and Automated Decision Systems**

on November 13, 2020

Good afternoon Chair Holden and committee members.

I am pleased to join you today as the Council explores the role of artificial intelligence, or "AI," and automated decision tools that continue to gain prevalence and influence decision making processes and practices that impact New Yorkers. As you realize, employing these systems can offer benefits to New Yorkers, bringing about efficiencies and improving outcomes for our residents. As with other technology tools, without careful application and guidance, these tools may also cause unintended harms. The City shares the Council's interest in preventing any harms that may result from application of these technology tools and sees this as part of a multi-faceted effort to advance the concept of Digital Rights for New Yorkers, building on existing human rights and privacy rights.

Today, I'll share with you the technology context for the AI tools under discussion, update you on the status of the ethics conversations on these technologies, illustrate the City's efforts to prevent harms, including advancing the concept of Digital Rights and creating a framework for managing government use of algorithms, and discuss future protections needed to balance risks.

Evolving Role of Artificial Intelligence and Automated Decision Systems

There are a number of overlapping terms that can create confusion, so permit me to clarify a few points. An "algorithm" is simply a step-by-step recipe for carrying out a task, like rotating a photo 90 degrees or sorting a column in Excel, and the vast majority of algorithms are innocuous.

An "automated decision system," or ADS, currently has no standard definition but can be thought of as a computer program that takes input about a situation and then produces either a result, a recommendation, or prediction to assist a human decision-maker. These can be "fully automated" or "partially automated".

An ADS uses algorithms — both simple and complex — to make or assist decisions about potentially sensitive topics, which is one reason why the City's task force focused on ADS in particular. The term of art that the City uses to refer to an "ADS" is "algorithmic tool," and the City considers the terms "ADS" and "algorithmic tool" to be interchangeable.

Artificial Intelligence and ADS are distinct but related topics. An ADS may or may not use AI-based algorithms, and there are uses of AI that are unrelated to ADS. I will now briefly explain what AI is and some considerations to keep in mind.

Artificial Intelligence is a different way of writing computer programs and is often used in programs

involving prediction. In traditional programs, the author has to provide an explicit recipe for how to carry out the task. AI, on the other hand, is example-driven. Instead of writing an explicit set of rules, data is collected, sorted manually, and then mathematical methods are used to "train" the computer to figure out rules by itself. For example, email spam filters are considered "AI" due to the way they function though they have no general "intelligence".

It can be difficult to understand what an AI system ultimately is doing even if it works well. This is not unlike how cooks may struggle to write down a precise recipe that others can follow because the cooks have learned how to create their dish by trial and error and experience.

AI systems have been in use in virtually every field and aspect of society for a long-while, from consumer financial services to healthcare to housing to transportation and more.

Regardless of type, all systems have some error rate, as they are approximate methods. This is not unlike errors in assumptions or process in human decisions, which can occur absent technology. With technology, however, we need to apply different methods to identify and address errors.

In addressing problematic results for decisions where technology is used, it is extremely important to take into account the specific application of a technology. Technology itself is not inherently biased, however, the ways in which systems are used or how results are interpreted can produce biased outcomes.

Emerging Field of Ethics Related to Technology

AI ethics is an emerging and interdisciplinary field led by academics, practitioners, technologists, and other stakeholders that has become active in the last five years. The term "AI ethics" refers to the study of features in technology systems that affect societal values.

Principles considered in the field of AI ethics are fairness and non-discrimination, accountability, transparency, privacy, and accuracy. In real-world systems, these principles are often in tension with each other and the importance of the human input into technology becomes critical.

For instance, it may be necessary to use sensitive demographic attributes (making the system less privacy-preserving) in order to make a system fairer. Depending upon the human input of the amount and type of information used in the system, results may be more or less fair.

The City is actively engaging with the AI ethics community to learn and gain feedback on how cities can benefit from this important area of work.

Preventing Potential Harms

The City recognizes that as technology tools are more widely used, there is a growing role for local governments in working to ensure that city residents are able to safely access technology and continue to engage in education, employment, community, and other activities utilizing technology systems that produce equitable results. The federal government has begun to grapple with this issue, some states have made progress, but cities are also recognizing the unique role they can play in addressing impacts of emerging technologies.

In 2018, the Mayor, along with the Mayors of Barcelona and Amsterdam, UN Human Rights, UN-Habitat, and others formed the Cities Coalition for Digital Rights. This is a first-ever alliance of local municipalities to advance the concept of rights to protect and empower urban residents in their use of and exposure to digital technologies. The foundation of the approach is the development and concept of "digital rights" principles – which assert protections related to Cybersecurity and Privacy, Equity, Choice, Affordability, Quality, Accountability, and Ethics and Non-Discrimination. The Mayor's Office of the CTO currently uses these principles to guide the City's policy, research, programming, and engagement on both core and emerging technologies. The principles are critical to supporting not only individuals, but also entrepreneurs and small businesses, in navigating our increasingly digital society.

The Cities Coalition for Digital Rights is working with interested local governments, academics, and other experts on an initiative to apply and operationalize Digital Rights related to specific city systems and programs. Thus far, the Coalition is working with a dozen cities in North America and Europe to identify technology-informed practices in relation to observing Digital Rights. This is one of the first multi-city efforts to operationalize Digital Rights at the local level. New York City is serving as an advisor and facilitator for the initiative and will be engaging with leading practitioners, academics, and others on structuring this initiative in the coming months and working to make sure the outcomes of it benefit New York City.

Additionally, in November 2019, Mayor de Blasio signed Executive Order 50 recognizing that government agencies should leverage current technologies that rely on employing algorithms to support agency decision-making, while ensuring fairness and responsible impacts for New Yorkers. This E.O. created a new position of Algorithms Management and Policy Officer, which is a role currently filled by the Director of the Mayor's Office of Operations, Jeff Thamkittikasem. This Officer is responsible for developing citywide policies to guide agencies in the fair, responsible, and transparent use of algorithmic tools, including those using AI. The City has moved forward with this work, publishing introductory policies in September 2020 that are publicly available, and launching the City's first ever agency compliance reporting process. Agencies are currently reviewing their systems to identify those meeting the definition of an algorithmic tool and will report back on their findings. In January 2021, the Officer will publish a public report including information from these agency reports. Through this exercise, the City will have its first-ever look at the scope and scale of algorithmic tools in use by City agencies. This baseline understanding will further aid the Officer in developing additional assessment and complaint-resolution policies in 2021, as required by E.O. 50.

Role of Cities Moving Forward

The role of local governments in balancing the benefits of technology use while protecting residents from unintended harms is only at its beginning stages. While the City has already demonstrated leadership among its U.S. peers, this work will need to evolve along with the development of new technologies and new applications. The City looks forward to continuing collaborations with leading thinkers, practitioners, other stakeholders, and the Council as it puts into practice principles, policies, and protections to enable all New Yorkers to safely and equitably benefit from current and emerging technologies.

Intro 1894

Finally, I will now turn to Intro. 1894. This bill would regulate the use of automated employment decision tools used in the hiring process. The Administration shares in the Council's strong interest of rooting out bias in decision-making systems that use algorithms and artificial intelligence. We have operational, legal, and financial concerns with this bill as written — particularly in light of the various crises the City faces during the COVID-19 response and the current financial situation — and we look forward to working with the Council to address these issues.

Thank you for your time today and for your interest in this important topic.

ACLU of New York

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

Testimony of Daniel Schwarz

On Behalf of the New York Civil Liberties Union

Before the New York City Council Committee on Technology

Regarding Oversight and Regulation of Automated Decision Systems

November 13, 2020

The New York Civil Liberties Union ("NYCLU") respectfully submits the following testimony regarding much-needed oversight and regulation of automated decision systems. 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.

Automated decision systems ("ADS") – software tools or processes that automate, replace, or aid human decision-making – are widely used to administer services, allocate resources, and make inferences about individuals, groups, or places. Whether across government agencies or in private businesses, their ubiquity and opaque deployment risk severely undermining the civil, human, and privacy rights of New Yorkers. The use of ADS is often accompanied by an acute power imbalance between those deploying these systems and those affected by them, particularly given that ADS operate without transparency or even the most basic legal protections. Especially where New Yorker's fundamental rights are at stake – such as in welfare, education, employment, housing, health care, the family regulation (or child welfare) system, or the criminal legal system, 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.

The legislation in front of the Committee, Intro. 1894-2020, laudably attempts to tackle bias in automated employment decision systems. Like many of the aforementioned high-stake areas, ADS in employment and hiring are in urgent need of transparency, oversight, and regulation. Job applicants should not need to worry about being screened by a racist or sexist algorithm when seeking employment. Unfortunately, this bill would not create meaningful protections against such scenarios but instead would give cover to vendors to continue to sell discriminatory systems and offer negligible protection to applicants. The NYCLU therefore opposes Intro. 1894-2020 in its current form and makes recommendations to strengthen the legislation.

ADS result in decisions that impact New Yorker's lives. This Council must act to provide transparency and accountability to automated employment technologies, and ensure that ADS do not operate to digitally circumvent New York City's laws against discrimination.

The Need for Regulation of Automated Decision Systems

While the use of ADS undoubtedly boosts speed and scale, such efficiency is only valuable if the underlying decisions are desirable. Even with the little public information available about ADS, researchers and experts consistently reveal their failures of accuracy and neutrality. Many studies have challenged their opaque or "black box" operation [1] and provided evidence of harmful,[2] discriminatory,[3] sexist,[4] and racist[5] outcomes.

---

[1] *See e.g.:* CATHY O'NEIL, WEAPONS OF MATH DESTRUCTION: HOW BIG DATA INCREASES INEQUALITY AND THREATENS DEMOCRACY (2016); FRANK PASQUALE, THE BLACK BOX SOCIETY (2015).

[2] *See e.g.:* VIRGINIA EUBANKS, AUTOMATING INEQUALITY: HOW HIGH-TECH TOOLS PROFILE, POLICE, AND PUNISH THE POOR (2018); Ed Pilkington, *Digital dystopia: how algorithms punish the poor*, THE GUARDIAN, October 14, 2019, https://www.theguardian.com/technology/2019/oct/14/automating-poverty-algorithms-punish-poor (last visited Nov 10, 2020); Colin Lecher, *A healthcare algorithm started cutting care, and no one knew why*, THE VERGE (2018), https://www.theverge.com/2018/3/21/17144260/healthcare-medicaid-algorithm-arkansas-cerebral-palsy (last visited Nov 10, 2020).

[3] SOLON BAROCAS & ANDREW D. SELBST, *Big Data's Disparate Impact* (2016), https://doi.org/10.2139/ssrn.2477899 (last visited Nov 10, 2020).

[4] *See e.g.:* Jeffrey Dastin, *Amazon scraps secret AI recruiting tool that showed bias against women*, REUTERS, October 10, 2018, https://www.reuters.com/article/us-amazon-com-jobs-automation-insight-idUSKCN1MK08G (last visited Nov 10, 2020); Galen Sherwin, *How Facebook Is Giving Sex Discrimination in Employment Ads a New Life*, AMERICAN CIVIL LIBERTIES UNION , https://www.aclu.org/blog/womens-rights/womens-rights-workplace/how-facebook-giving-sex-discrimination-employment-ads-new (last visited Nov 10, 2020).

[5] *See e.g.:* Kate Crawford, *Opinion | Artificial Intelligence's White Guy Problem*, THE NEW YORK TIMES, June 25, 2016, https://www.nytimes.com/2016/06/26/opinion/sunday/artificial-intelligences-white-guy-problem.html (last visited Nov 10, 2020); Alistair Barr, *Google Mistakenly Tags Black People as 'Gorillas,' Showing Limits of Algorithms*, WSJ (2015), https://blogs.wsj.com/digits/2015/07/01/google-mistakenly-tags-black-people-as-gorillas-showing-limits-of-algorithms/ (last visited Jan 15, 2020).

Software systems are often wrongly perceived as more neutral than humans or as offering a scientific and objective truth.[6] Their proponents are able to make these assertions because the vast majority of ADS are opaque systems, secretly deployed and shielded from independent review due to their proprietary nature. This secrecy obscures the potential errors, outright flaws, biased data, subjective decisions, and personal choices that find their way into these systems. Every ADS is a product of human design, input, and operation.

Obtaining access to ADS's underlying source code and data is difficult and resource intensive, but absolutely critical to understanding the extent to which errors occur and whether they are likely to cause discriminatory harm. For example, it was revealed that a Medicaid ADS in Arkansas had failed to correctly assess care needs of patients with cerebral palsy or diabetes: a fact only discovered through lengthy litigation and subsequent disclosure of the code.[7] And here in New York City, an independent review of the source code of a DNA analysis tool used by the office of the chief medical examiner raised serious questions about its validity, including whether the code may have been intentionally skewed to create more matches.[8]

Many automated systems purport to predict the future by observing the past. Among them are "risk assessment tools," designed to use past policing and court data to "predict" the future behavior of an individual criminal defendant. Specifically, risk assessment tools attempt to determine which attributes are shared by people who previously failed to show up to court. Certain weights are placed on each of the attributes to produce a formula and "score" a person's future risk of flight. Risk assessment tools reflect a troubling philosophy toward criminal justice policy: Using past cases to determine what might happen in future cases disregards time-specific influences that may have affected prior case outcomes and freezes a government judgment in the realities of the past. Critically, it also strips the person who is awaiting trial of independent agency and the ability to make the case that they will appear in court.

But even those who philosophically agree with using past statistics to predict future individual human behavior acknowledge that the value of such a predictive system lies in the value of the data input into it. When an ADS deploys machine learning that relies on large historic datasets to train the underlying models, the quality of that underlying data is of paramount importance. If that data includes false or biased data, every output will repeat this pattern and in turn result in false and biased decision-making. In the context of policing, utilizing data from unconstitutional and racially biased stop-and-frisk practices by the NYPD will create outputs reflecting these practices.[9] This behavior is commonly known by the

---

[6] danah boyd & Kate Crawford, *Critical Questions for Big Data: Provocations for a cultural, technological, and scholarly phenomenon*, 15 INFORMATION, COMMUNICATION & SOCIETY 662–679 (2012).
[7] Litigating Algorithms 2018, AI NOW INSTITUTE, https://ainowinstitute.org/litigatingalgorithms.pdf.
[8] Lauren Kirchner, *Thousands of Criminal Cases in New York Relied on Disputed DNA Testing Techniques*, PROPUBLICA (2017), https://www.propublica.org/article/thousands-of-criminal-cases-in-new-york-relied-on-disputed-dna-testing-techniques (last visited Nov 10, 2020).
[9] Rashida Richardson et al., *Dirty Data, Bad Predictions: How Civil Rights Violations Impact Police Data, Predictive Policing Systems, and Justice,* 94 N.Y.U. L. REV. ONLINE 192 (2019), https://ssrn.com/abstract=3333423.

computer-science idiom "garbage in, garbage out," or in this scenario, as Sandra Mayson coined, "bias in, bias out."[10]

In another recent example, researchers discovered that a widely used health care algorithm used to identify patients' health risks failed to identify many Black patients, making them less likely to be enrolled for medical treatment.[11] And where these systems operate in the dark, people may not even realize that they are suffering at the hands of a flawed machine-learning system: one ADS in Indiana blocked hundreds of thousands of people from receiving vital support services and left them struggling to challenge these decisions.[12]

Given these enormous human impacts that automated systems make on our communities – and the very real possibility of simply automating existing human error and bias – meaningful regulation is the bare minimum our democracy demands. The growing power imbalance between people affected by ADS and those who deploy them is at its height when affected people are not even aware that their lives have been impacted by an ADS. In particular in governmental decision-making, access to information about what systems are in use, whether their accuracy has been studied and their impact assessed, and the mechanisms to obtain redress for harm is essential for the public to be able to engage in a fully-informed discussion regarding what role – if any – these systems should have in our society.

In November 2018, New York City joined the Cities Coalition for Digital Rights and signed its Declaration. It clearly states that people have "sovereignty over their data, including the right to know what happens to their data, who uses it and for what purposes. […] Everyone should have access to understandable and accurate information about the technological, algorithmic and artificial intelligence systems that impact their lives, and the ability to question and change unfair, biased or discriminatory systems."[13] We urge the Council to uphold this promise by enacting legislation that will serve our democratic values and create the regulatory mechanisms necessary to protect against harmful and discriminatory algorithms.

---

[10] Sandra G. Mayson, *Bias In, Bias Out*, 128 YALE LAW JOURNAL (2019), https://www.yalelawjournal.org/article/bias-in-bias-out (last visited Nov 10, 2020). Archived at: http://archive.is/nzP1D.

[11] *See:* Beth Haroules & Simon McCormack, *How an Algorithm Puts Black People's Health in Danger*, NEW YORK CIVIL LIBERTIES UNION (2019), https://www.nyclu.org/en/news/how-algorithm-puts-black-peoples-health-danger (last visited Jan 15, 2020); Ziad Obermeyer et al., *Dissecting racial bias in an algorithm used to manage the health of populations*, 366 SCIENCE 447–453 (2019).

[12] Alyssa Edes & Emma Bowman, *"Automating Inequality": Algorithms In Public Services Often Fail The Most Vulnerable*, NPR.ORG (2018), https://www.npr.org/sections/alltechconsidered/2018/02/19/586387119/automating-inequality-algorithms-in-public-services-often-fail-the-most-vulnerab (last visited Jan 16, 2020); Virginia Eubanks, *We created poverty. Algorithms won't make that go away*, THE GUARDIAN, May 13, 2018, https://www.theguardian.com/commentisfree/2018/may/13/we-created-poverty-algorithms-wont-make-that-go-away (last visited Nov 10, 2020).

[13] Declaration of Cities Coalition for Digital Rights, https://citiesfordigitalrights.org/assets/Declaration_Cities_for_Digital_Rights.pdf.

To close the overwhelming information gap around ADS in New York City, the Council could strengthen and pass Intro. 1806-2019. This legislation would require city agencies to provide basic information about every automated decision system in use. Such disclosures will help the public and policymakers alike understand the current terrain, craft better and more targeted oversight mechanisms, aid people in finding help when they feel they are unfairly impacted by a decision, and drive public education opportunities. Other cities have shown the feasibility of similar efforts: Amsterdam and Helsinki recently launched their respective ADS registries, listing descriptions about their governmental automated decision systems, detailed information regarding the datasets used and how they are processed, assessments for discrimination and harm, and steps for human review.[14]

Yet transparency is only a first step and foundation for more comprehensive and targeted regulation. Effective action will necessarily include mandatory, independent racial and non-discrimination impact assessments, data privacy audits, and holistic consultation with domain experts and people directly affected by the consequences of any ADS – in particular from marginalized groups – prior to any ADS rollout and throughout the entire life cycle. Finally, the Council should recognize that technologies showing significant discriminatory impact against any class protected under the New York City Human Rights Law require outright bans or moratoria – in particular in high-stake areas, including, but not limited to, facial recognition.

Unfortunately, the City's forays into ADS issues have fallen short of these goals. The NYCLU and our partners repeatedly sought to offer input and recommendations through open letters in January 2018,[15] August 2018,[16] March 2019,[17] a comprehensive Shadow Report in December 2019,[18] and have testified before this Committee in January 2020.[19]

---

[14] *See:* City of Amsterdam Algorithm Register, https://algoritmeregister.amsterdam.nl/en/ai-register/ and City of Helsinki AI Register, https://ai.hel.fi/en/ai-register/ (last visited Nov 10, 2020).
[15] Letter to Mayor de Blasio: Regarding NYC Automated Decision Systems Task Force, NEW YORK CIVIL LIBERTIES UNION (2018), https://www.nyclu.org/en/publications/letter-mayor-de-blasio-regarding-nyc-automated-decision-systems-task-force (last visited Nov 10, 2020).
[16] Open Letter to Automated Decision Systems Task Force, NEW YORK CIVIL LIBERTIES UNION (2018), https://www.nyclu.org/en/publications/open-letter-automated-decision-systems-task-force (last visited Nov 10, 2020).
[17] Letter to the Automated Decision Systems Task Force - March 1, 2019, NEW YORK CIVIL LIBERTIES UNION (2019), https://www.nyclu.org/en/publications/letter-automated-decision-systems-task-force-march-1-2019 (last visited Nov 10, 2020).
[18] *See*: Rashida Richardson, ed., *Confronting Black Boxes: A Shadow Report of the New York City Automated Decision System Task Force*, AI NOW INSTITUTE, December 4, 2019, https://ainowinstitute.org/ads-shadowreport-2019.html.
[19] NYC Council Testimony In Relation to Automated Decision Systems Used by Agencies, NEW YORK CIVIL LIBERTIES UNION, Jan 22, 2020, https://www.nyclu.org/sites/default/files/field_documents/20200122-nyclu-testimony-automateddecisionsystems.pdf.

Intro 1894-2020 - Sale of Automated Employment Decision Tools

The NYCLU commends the sponsor and the Council for raising and attempting to tackle the issue of bias and discrimination in automated employment decision tools. Unfortunately, Intro. 1894-2020 does not sufficiently achieve this goal and we oppose it in its current form, because it would entrench ineffective regulation and could increase the use of harmful ADS technology.

Intro. 1894-2020 erroneously places the sole focus on the *sale* of automated employment decision tools. Any entity that already operates a hiring algorithm would be able to continue using it without any bias audit. Similarly, if a tool is developed in-house – as with Amazon's widely reported sexist hiring ADS[20] – it would not be affected by this legislation.

The legislation would only cover a subset of hiring technologies due to the limited definition of automated employment decision tools. While ADS that rely on statistical theory, such as machine learning, are an important area that requires inclusion, many other assessments and automated tools fall outside this legislation's framing. Furthermore, the bill is drafted to include only the filtering of preferred candidates. Instead, it should include all applicants and employment decisions, including the rejection, removal, and rating of candidates.

At the heart of this legislation lies the mandated bias audit. Unfortunately, the specifications and mandate are too limited, giving much leeway to vendors and therefore risking a biased bias audit. The legislation does not specify the entities who would conduct the bias audits. If vendors are left to their own devices, these audits will hold little value and will suffer from their own biases – vendors have a financial incentive to conduct the audits in-house or to contract with friendly third parties and report no bias. Instead the bill should clearly set out an independent process developed by experts and stakeholders to ensure meaningful testing and assessment as well as mandatory disclosures to the public.

The legislation also includes a notice requirement to candidates, which would require an employer to tell a candidate that an ADS was used to screen a candidate and what "job qualifications or characteristics […] such tool was used to assess in the candidate." This notice is inadequate, because it is unlikely to provide an applicant who was subjected to a biased algorithm with the information they need to understand whether they were discriminated against. Moreover, because some ADS rely on machine learning, meaningful information regarding the qualifications or characteristics they use will be difficult to obtain.

To ensure actual accountability for ADS regulation, a private right of action must be available to affected applicants or workers. Without it, enforcement would be improbable due to the all-too-real capacity limitations of the New York Commission on Human Rights. Furthermore, the Council should mandate the provision of attorneys' fees to New Yorkers who successfully vindicate their right to be free from discriminatory employment decisions. Lastly,

[20] Dastin, *Amazon scraps secret AI recruiting tool that showed bias against women*, s*upra* note 4.

the legislation must include a non-retaliation provision for workers or applicants who exercise their rights protected under this bill.

Every New Yorker has a right to know whether and how automated decision systems are impacting our lives and livelihoods. While we appreciate that Intro. 1894 takes steps to regulate automated employment decision tools, it falls short of ensuring that New Yorkers are free from employment bias masquerading as neutral decision-making. This legislation must be amended and expanded in order to deliver on its promise to mitigate bias and bring justice and equity to the world of hiring ADS. Moreover, as described above, the Council must regulate ADS outside of the employment context as well. Of particular concern are ADS in government that make high-stakes decisions impacting the constitutional or legal rights of New Yorkers.

Conclusion

We thank the Committee for the opportunity to provide testimony and for recognizing the need for oversight and regulation of automated decision systems. The NYCLU urges the Council to pass legislation to create transparency and protections ensuring fair and equitable use of automated decision systems. Unfortunately, Intro. 1894-2020 does not achieve these goals in its current form.

New York City Council
Committee on Technology
11/13/20

Thank you Chair Holden, the Technology Committee, and the New York City Council for permitting me to share our thoughts on such groundbreaking legislation.

I'm Rev. Kirsten John Foy, President & CEO of The Arc of Justice, a national civil rights organization based in Brooklyn. Today I have the privilege of speaking on behalf of a broad coalition of organizations that have come together to support Intro. 1894. We are a diverse group that represents New Yorkers from all walks of life. Signees of this letter include the Bedford Stuyvesant Restoration Corporation, Myrtle Ave BID, Richmond Black Chamber of Commerce, Rockaway Development Corporation, Black Gotham Experience, 3rd Ave BID, NY for Seniors, Warriors in the Garden, Andrew Freedman House, Youth Action Build, Black Bear Brotherhood and the Black Lunch Table.

Each of our groups have seen and experienced the wreckage of the pandemic on New York's public health and economy. And of course we are doing all that we can to get New Yorkers back on their feet, back at jobs, and able to thrive in the city.

But our work alone will hardly be sufficient in ensuring a fair recovery in all five boroughs. That is why we implore the Technology Committee - and the entire NYC Council - to pass Int. 1894.

The shaky recovery is underway and we have no time to waste to make sure it is equitable. Slowly but surely companies in the city are hiring workers once again. And increasingly, companies are using automated technologies to guide their hiring decisions. Large companies in particular can receive hundreds of applications for a single job opening. New technologies can scale the evaluation process without requiring individuals to sift through resume after resume. 36% of companies say they intend on using artificial intelligence for hiring over the next few years.

But just as human biases have led to job discrimination against people of color and women for generations, hiring technologies can also lead to unfair outcomes. However, we are confident that this legislation will reduce discrimination in hiring by requiring that vendors audit their technology annually to show whether their offerings are leading to hiring decisions that do not have an adverse impact on people of color and women. Job applicants seeking employment during one of the worst economic periods ever for our city should not

have the additional worry of being discriminated against. With mandated audits, these technologies will actually be an improvement over the traditional ways most companies now hire workers.

Ultimately, this will be good for the economy. A recent Wall Street Journal study found that the 20 most diverse companies had an average annual stock return of 10% over five years, versus 4.2% for the 20 least-diverse companies.

In addition, job applicants also deserve to know how they are being evaluated for an opening. For so much of our lives, we have assumed that HR reps are reviewing our resumes, assessing our skills during interviews, and calling our references. But with artificial intelligence now playing a role, applicants should know these new systems are conducting the reviews. That's why we are appreciative the City Council is requiring that all job applicants be given notice when employers are subjecting them to automated technologies. This follows similar legislation in Maryland and Illinois.

The legislation comes at a crucial time for our communities. New York's oft-overlooked residents play an integral role in New York's economy, culture, and spirit, but they have never fully enjoyed the fruits of the city's economic largesse. This law will help change that paradigm. We need to encourage companies to choose workers based on their skill, not on how they look.

We are also thrilled that two Black women - Majority Leader Laurie Cumbo and Council member Alicka Ampry-Samuels - are spearheading efforts to pass this bill. Technology should serve everyone equally and we're inspired to see technology policy being led by women of color who are often forced to the sidelines of this industry. We hope their leadership serves as an example for legislators of color throughout the country.

Chair Holden and distinguished Technology Committee Councilmembers, good afternoon and thank you for the opportunity to speak with you today about my support for this bill, which represents an important step forward in the effort to increase transparency and fairness in the hiring process.

I spent more than 40 years working in civil rights for the federal government, first at the Department of Labor's Office of Federal Contract Compliance and then, for the majority of my career, at the Equal Employment Opportunity Commission. I am passionate about equal opportunities for all job seekers and have devoted my career to this mission, trying to make meaningful change for real people.

Over the course of my career I've learned a lot about hiring procedures and employment screens that act as unnecessary barriers to entry for historically disadvantaged groups, most notably non-whites and women. There were many times during the course of my work on a compliance review, an investigation, or a litigation action when it occurred to me that employers tend to be generally unaware of the disparate impact caused by the use of facially neutral employment screens. This is important to understand because it highlights the fact that even if employers want to do the right thing, they may not understand the disproportionate impact that their hiring tools can have on historically disadvantaged groups.

Hiring tools that disproportionately screen out members of particular race or gender groups are not new. In fact, this has been going on for decades. Oftentimes these concerns are discussed only in connection with new technologies like those powered by artificial Intelligence and machine learning. However, any hiring tool or employment screen can cause a disparate impact. It is extremely important that all hiring tools be held to the same, clear standard. Developers and purchasers of these tools must understand the standard and they must be motivated to audit their outcomes for disparate impact.

For example, many large employers use traditional cognitive assessments such as intelligence tests in hiring. Research has demonstrated that traditional cognitive ability tests disproportionately benefit White job applicants. For every 100 White applicants selected via a cognitive test, only 32 African American applicants are also selected (Sackett and Ellingson 1997). Regardless of their adverse impact on women and people of color, employers continue to use these types of hiring tools. This may be simply because they're unaware of these effects or they're unaware of their options for tools with more equitable outcomes.

The federal standard for adverse impact testing has been available to employers since 1978, when EEOC and other federal agencies came together to publish the Uniform Guidelines on Employee Selection Procedures. The Uniform Guidelines was designed to help employers understand how to comply with Title VII of the Civil Rights Act, which outlaws employment discrimination based on race, ethnicity, gender, religion and national origin. The Uniform Guidelines established that some hiring tests, although

having an adverse impact on classes of workers such as African Americans, could be considered acceptable if the test validly captures a bona fide work qualification. The Uniform Guidelines also established that even if a selection device was shown to be valid and job-relevant, the employer had an obligation to use a less discriminatory alternative if one was available.

My experience indicates that employers have generally focused on the job relevance requirements of the Uniform Guidelines rather than on proactive adverse impact testing, and when told by a vendor that a test is valid, the employer does not question whether it is valid for their workforce or it is validated for nondiscriminatory requirements. This has led to outcomes that are damaging to workers and employers alike. Applicants confronting an unfair test have been denied jobs, promotions and other workplace selections and suffered both immediate monetary and future career losses. Employers perhaps suffered more, as their labor pools did not include the most qualified, just the most qualified based on a biased test.

Sometimes this problem for employers can be hard to grasp. But a great American parable is helpful. In 1947, the Brooklyn Dodgers reached out of their normal labor pool and Jackie Robinson became the first African American player in major league baseball. Robinson had a batting average of nearly .300 that year, was the Rookie of the Year, and the Dodgers went on to the World Series. The continued use of a biased labor pool would certainly have made the Dodgers a less talented team.

Just as New York has led the way before, whether it was diversifying baseball or introducing groundbreaking legislation to prohibit questions about salary and criminal history early in the job application process, the city again should be applauded and encouraged to take this action to help ensure that hiring assessments perform fairly for all job applicants.

That's why I am particularly excited about the passage of Int.1894. My hope is that this proactive bias audit requirement will clearly focus employers on fairness outcomes of hiring tests. This clear focus could have a major impact on vendor innovation and inspire an entire industry to refocus on fairness. If I could make one suggestion, it would be to focus the bill squarely on improving fairness across race/ethnic and gender groups in order to ensure that this bill has maximum positive impact and comports with Title VII. This will put the bill on firmer legal ground and ensure immediate compliance. Job candidates have become accustomed to voluntarily identifying themselves based on race/ethnicity and gender, and guidance at the federal level has indicated that employers should monitor applicant flows to ensure fairness on these demographic indicators.

Now more than ever employers must have confidence that their hiring practices are fair, in compliance with federal requirements and are producing the labor pools that will increase productivity. Applicants need to be protected by making certain that selection

devices provide a fair chance to obtain a job, get a promotion and are rewarded in a manner consistent with their real qualifications.

Statement of Professor Mark MacCarthy
On
Int. 1894-2020,
A bill to regulate the sale of automated employment decision tools
Before the Committee on Technology
New York City Council

November 13, 2020

Good afternoon. I would like to begin by thanking Chair Holden and the Technology Committee for holding this hearing, as well as Majority Leader Cumbo and her co-sponsors for bringing forward this precedent-setting proposal.

I greatly appreciate the opportunity to testify on the timely and important issues raised by the bill Int. 1894-2020, introduced on February 27, 2020, and designed to regulate the sale of automated employment decision tools.

I'm Mark MacCarthy. I am a senior fellow with the Institute for Technology Law and Policy at Georgetown Law, and I teach in Georgetown's graduate program in Communications, Culture & Technology. I recently became a Senior Fellow at Brookings Institute's Center for Technology Innovation. I retired from the Software & Information Industry Association (SIIA) in February of 2019, after running their public policy shop for eight years. Prior to that I was a senior public policy official for Visa and ABC and spent eight years as a staff aide for the Energy and Commerce Committee in the U.S. House of Representatives

Relevant to this hearing, I am the author of a law review article on standards of fairness for disparate impact assessment of algorithms and a report for the Brookings Institution on fairness in algorithmic decision making

I support the legislation. It would regulate the sale of "automated employment decision tools" that automatically filter candidates "for hire or for any term, condition or privilege of employment in a way that establishes a preferred candidate or candidates." The bill serves the cause of workplace diversity and the protection of vulnerable groups in the employment process.

Employment decisions need to be improved to ensure that groups who have been subjected to employment discriminated in the past are not subjected to further bias in employment going forward. Automated employment decision tools have the promise of reducing bias introduced by subjective employer decision, but if not properly designed and used they might instead reinforce and even worsen existing patterns of employment discrimination.

The bill requires vendors to conduct bias audits of their employment algorithms to assess their "predicted compliance with…applicable law relating to discrimination in employment." Vendors must conduct these audits in the past year before selling or offering to sell their tools, and must provide to their customers, at no charge, the results of their ongoing annual bias audits.

The intent of the bill is to disclose the extent to which an algorithmic employment decision tool might worsen workplace diversity. Employers cannot manage what they do not measure. They need to know whether prospective employment tools are likely to have discriminatory effects, and the only way they can know that is if vendors conduct disparate impact assessments and convey the results to their potential purchasers.

To ensure that this intent is carried out, I recommend that it be clarified to state clearly that the results of the vendor's bias audit should be disclosed to all potential purchasers as part of the offering of the tool for sale. Vendors should not merely state that they have conducted a required bias audit. They must also disclose the results of that audit to potential buyers. In addition, I recommend that the bill be clarified so that the required bias audit must assess the tool's potential adverse impact on protected classes and that this assessment should be disclosed to potential purchasers.

The standard measurement of adverse effect in employment law is whether a policy, procedure or tool returns positive results for members of a disadvantaged group in the same proportion as for other groups. If the tool preserves statistical parity in this way, it will not worsen outcomes for protected classes. But if it treats members of disadvantaged groups less favorably than members of other groups, it will make things worse for the disadvantaged groups.

This is the statistical notion at the heart of the commonly used 80% rule of thumb in employment law. This rule of thumb says to employers, for instance, that if they hire, say 10% of white applicants, they should be hiring at least 8% of the African American applicants. If they interview 50% of the men applying for a job, they should be interviewing at least 40% of the women. If employers cannot pass this 80% rule of thumb, they should be looking very carefully at the employments methods they are using to ensure that they are not at risk of legal liability for employment discrimination.

Vendors can test whether their automated employment tools are likely to have an adverse impact on protected classes by conducting their initial assessments on different demographic groups and measuring what proportion of people in protected classes receive positive results. These assessments can be kept up to date through the requirement in the legislation for annual bias audits.

Vendors whose employment tools do have an adverse effect on protected class should be allowed to explain in their bias audits that their tools assess job-related qualifications and are consistent with a compelling business necessity. This would allow employers to compare employment decision tools and to shop for an alternative that satisfies their business needs with the smallest possible discriminatory effect.

I applaud the Council on the introduction of this important and practical piece of legislation and on today's hearing to assess it. I encourage the Council to pass the bill with the clarifications suggested in this testimony.

Testimony of Steven Kuyan before New York City Council Committee on Technology regarding Int 1894-2020, “Sale of automated employment decision tools”

November 12, 2020

My name is Steven Kuyan. I am the Director of Entrepreneurship at NYU Tandon School of Engineering where I founded the NYU Tandon Future Labs and Director of the NYU Center for Responsible AI with my colleague Julia Stoyanovich. I am an investor in and advisor to early stage technology companies. I am also a board member of the Business Incubator Association of New York State and a member of the United States Council for Competitiveness. My perspective comes from launching the first New York City sponsored incubator in 2009, initiated as a partnership between New York City Economic Development Corporation and the NYU Tandon School of Engineering. For more than a decade we have helped early stage companies successfully translate their technology advancements — including Artificial Intelligence, Data, and Automated Decision Systems -- into successful commercially viable products for markets hungry for optimization, efficiency, and scale from Deep Technology.

Ensuring that AI can transfer from lab to market is one of the biggest economic opportunities of a generation and one that is instrumental for the health of our startup ecosystem and the competitiveness of our country.[1] Without oversight and regulation, wealth from AI will be concentrated in companies that are able to harness and deploy it. And it’s already happening, just think about the companies we go to for our everyday needs.[2]

It is with this in mind that I support Int. 1894. It's a step in building a roadmap that ensures that all AI be deployed responsibly, ethically, and with transparency. This is necessary for AI to reach its promise of scale in a wide array of critical sectors with public implications such as medicine, mobility, education, banking, law, and as the bill focuses on, employment.

---

[1] https://www.analyticsinsight.net/top-10-countries-leading-the-artificial-intelligence-race/

[2] https://money.usnews.com/investing/investing-101/articles/2018-05-07/robotics-automation-and-ai-are-the-new-fang

Without public oversight and well tailored regulation, current conditions indicate that AI will evolve in the dark, sold on the quality of a marketing campaign. Instead, AI/ADS must be ethically harnessed so it can fulfill its great potential for innovation, and ideally, fairness.

In this statement today, I would like to make two recommendations.

1. **Standardizing Auditing**. Starting with employment tools, there is an incredible opportunity for NYC to help encourage a standardization of auditing of automated decision systems.
2. **Transparent liability**. Current software vendors are not disclosing that those that use their technology and act on the recommendations of their systems are liable for those decisions. This bill must make that a priority.

Recommendation 1: Standardized auditing

Explainable AI, or AI is paramount in industry deployment of AI; however existing methods fail to address this necessity, in part due to a lack of standardisation of explainability methods. Even though one of AI's biggest pitfalls is to be intentionally misused, it doesn't pose such a great danger as it follows a given expected behavior. The real danger lies in the errors that were not accounted for, and that affect millions of people. These errors, or biases, are inherited from the humans building the AI models and introducing their unconscious systemic discrimination belief system. And if we are not cautious, these biases become self-amplifying.

An example of a long streak of biases in technology is the "racist soap dispenser."[3] We cannot really blame AI for being inherently racist, it is simply the way that it was programmed and the data on which it was trained that is flawed. When a group of white researchers developed an automatic faucet, they couldn't predict that it wouldn't recognize the hands of black people as they only tested it on themselves and on people who looked like them.

With every customer data asset of a company is being considered for AI solutions and technology as vendors are developing and deploying such tools without any oversight as to how that data is used to influence future decisions on any dimension other than generating more revenue. As of today, there is no agreed upon, standardized system in place that allows for a third-party, regulatory or industry validation of the technology to assess if it functions as intended, especially without any bias. There is early research and support from the research community to support explainable AI.[4] Without the ability to standardize auditing, companies purchasing AI tools from vendors will have no knowledge of whether the expected results are intended.

[3] https://www.mic.com/articles/124899/the-reason-this-racist-soap-dispenser-doesn-t-work-on-black-skin
[4] https://arxiv.org/pdf/2010.11273.pdf

Recommendation 2: Transparency on liability, responsibility, and accountability

The above recommendation necessitates widespread education on libabily, responsibility, and accountability. Companies across New York City are adopting automated decision systems for hiring without the knowledge that these purchasing decisions are putting them in direct liability they had not accounted for by following the recommendations of systems ill equipped to make *fair* recommendations. Actions taken by companies based on recommendations from automated decisions systems are currently the liability of the companies making those decisions.

To illustrate the point, allow me to share an example of a New York City startup, RentLogic, supporting the efforts of our Center for Responsible AI. RentLogic provides grades to buildings based on available data from New York City agencies and reviews of buildings found through various data sources. Being forward thinking about legal ramifications of giving an algorithmic score to buildings based on potentially biased data, Yale Fox, the founder, understood that building owners will hold them liable for any mistake the system will make. Yale self-imposed a requirement to be transparent and responsible for their results, a rarity for technology companies. It took well over a year to find someone qualified to perform an algorithm and data audit as there are no guidelines that exist for doing so.

The liability is very similar with automated decision systems in hiring. Companies are purchasing these solutions, many of which do not work consistently, without knowing that they are liable for the hiring decisions they make, informed in by what may or may not be Artificial Intelligence. It's imperative that companies are at least aware of their liability before purchasing these tools and it may be the necessary driver of bottom up auditing, where companies themselves demand audits before making purchasing decisions.

Closing

Please let me close by recognizing that the general public knows and cares about this issue. Not only do we know that because of public education efforts of the Center, but from data. Of consumers surveyed, 62% said they would place higher trust in a company whose AI interactions they perceived as ethical. The inverse was true if the AI was not perceived as ethical.[5] We cannot lose the trust of our public any further.

---

5 https://www.forbes.com/sites/joemckendrick/2019/07/07/ethical-artificial-intelligence-becomes-a-supreme-competitive-advantage/?sh=180b65a31a8f

NYC should enact its bold vision for the deployment of automated decision systems in a way that aligns with societal norms we expect to see in the future we all want to live in. 1894 may not be perfect, but we need to start with a statement that transparency is a necessity for automated decision systems, for better utility, and ultimately, for the continued research and adoption of AI.

The time to do this is now, otherwise we risk getting too far and facing a backlash towards AI that will risk our global position as an innovation hub or worse, not having this opportunity to enforce these system again, similar to the world's inability to have any oversight over AI recommendation tools that power our social media platforms.

NEW YORK UNIVERSITY

Julia Stoyanovich
Assistant Professor of Computer Science and Engineering, Tandon School of Engineering

Assistant Professor of Data Science, Center for Data Science

Department Computer Science and Engineering
Tandon School of Engineering
2 MetroTech Center
Brooklyn, NY 11201
P: 212 998 3251
M: 917 470 8199
stoyanovich@nyu.edu

Testimony of Julia Stoyanovich before New York City Council Committee on Technology regarding Int 1894-2020: "Sale of automated employment decision tools"

November 12, 2020

Dear Chair Holden and members of the Committee:

My name is Julia Stoyanovich. I hold a Ph.D. in Computer Science from Columbia University. I am an Assistant Professor of Computer Science and Engineering at New York University's Tandon School of Engineering, and an Assistant Professor of Data Science at the Center for Data Science. I am the founding Director of the Center for Responsible AI at NYU, together with my colleague Steven Kuyan. 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] Some of the students who took my course are here today. Most importantly, I am a devoted and proud New Yorker.

I would like to applaud the Committee on Technology for their sustained efforts to regulate the use of automated decision systems (ADS) in New York City. The bill we are discussing today represents a potentially transformative opportunity to make the use of ADS in a crucial domain – hiring and employment – responsive to the needs of *all* New Yorkers. I am speaking here in strong support of the bill.

This bill cannot be more timely: The Covid-19 pandemic is hitting members of minority and historically disadvantaged groups particularly hard, with many losing their jobs and being unable to re-enter the workforce. If this bill passes, it will *benefit job seekers*, by ensuring that the unaccountable use of algorithmic decision making in hiring does not further exacerbate these inequities. The bill will also *benefit vendors* of hiring ADS, by helping create an economically and ethically sustainable ecosystem of technological innovation. Finally, the bill will *benefit employers* who use these tools, by helping them evaluate the claims made by vendors during procurement (through auditing), and build trust of job seekers and employees (through public disclosure).

---

[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/

Despite my strong support for the bill, I am of the opinion that much additional work is needed to ensure that, when passed into law, this bill gives rise to auditing and public disclosure procedures that are more than a rubber stamp.

Hiring ADS are a quintessential example of systems that are as impactful – for individuals, population groups, and society at large – as they are contraversial. Vendors of these tools frequently and confidently make claims that, because humans are known to have biases, algorithmic tools are our only viable option. It is dangerous to take such claims on faith: *algorithmic systems themselves, and audits of these systems, are only as good as the standards and objectives to which we hold them*. These objectives are often in competition, and it is not up to data or algorithms to guide us on how to resolve the trade-offs that arise.

Therefore, the meaningfullness of auditing and public disclosure mechanisms for hiring ADS will hang on our ability to support a robust and open dialog about where to use and *not* to use these systems, what kinds of decisions we should be leaving up to them and what kinds are for humans to make, how to state the objectives agains which we check these systems, and how to negotiate the trade-offs that priritize one stakeholder group's intests over another.

I must counter the argument that the City is too resource-constrained, due to the COVID-19 pandemic and otherwise, to engage in the oversight of hiring ADS. We would not be comfortable selling food or medicine if we were unsure about both their actual benefits and their safety. We would not be building a bridge or allowing a car out on the street if we were unsure that these artifacts both work to their specification and are safe. *Algorithmic systems are engineering artifacts.* They have no common sense, no empathy, and no sense of humor. They cannot exercise agency or be held accountable for their actions. To think otherwise is to indulge in magical thinking. Algorithmic systems are what we make of them, and if we decide to use them, then we – all of us collectively – are responsible for how they work. *If we trust algorithmic hiring systems sufficiently to deploy them, then we must ensure both that they deliver on their clearly stated purpose and that they are safe.*

Meaningful oversight of hiring ADS is, of course, is a tall order. But the New York City I know and love does not give up in the face of a challenge. And the City government does not have to do this work alone. The academic community, including the Center for Responsible AI at NYU, are the City's disposal to help make the auditing and public disclosure requirements of the proposed bill actionable.

In my statement today I would like to make three recommendations:

1. **Auditing**: The scope of auditing for bias should be expanded beyond disparate impact to include other dimensions of discrimination, and to also convey information about a tool's effectiveness. Audits should be based on a set of uniform publicly available criteria.
2. **Disclosure**: Information about job qualifications or characteristics for which the tool was used to screen should be disclosed to a job seeker in a manner that is comprehensible and actionable.

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[3] 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.

The entire hiring funnel, as well as each component of the funnel, are examples of automated decision systems (ADS). These systems:

1. Process data about people, some of which may be sensitive or proprietary;
2. Help make decisions that are consequential to people's lives and livelihoods;
3. Involve a combination of human and automated decision making; and
4. Are designed with the stated goals of improving efficiency and promoting, or at least not hindering, equitable access to opportunity.

ADS may or may not use Artificial Intelligence (AI), and they may or may not have autonomy, but they all rely heavily on data.

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

---

[3] 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

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[4] that was broadly circulated by the press[5]. 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 and search 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. Researchers showed that Google displayed 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 in 1964, but it persists in the online ad environment today.

The findings of the 2015 AdFisher study started a line of inquiry into the reasons for gender-based and other types of discrimination in online ad delivery, particularly as they pertain to access to employment and housing opportunities. The current understanding is that there is a multitude of reasons for this, including both biased training data and the advertisement targeting mechanism itself. In their comprehesive 2019 analysis of Facebook's ad delivery mechanisms, Ali et al.[6] explain:

"The enormous financial success of online advertising platforms is partially due to the precise targeting features they offer. Although researchers and journalists have found many ways that advertisers can target—or exclude—particular groups of users seeing their ads, comparatively little attention has been paid to the implications of the platform's ad delivery process, comprised of the platform's choices about which users see which ads.

It has been hypothesized that this process can 'skew' ad delivery in ways that the advertisers do not intend, making some users less likely than others to see particular ads based on their demographic characteristics. In this paper, we demonstrate that such skewed delivery occurs on Facebook, due to market and financial optimization effects as well as the platform's own predictions about the 'relevance' of ads to different groups of users. We find that both the advertiser's budget and the content of the ad each significantly contribute to the skew of Facebook's ad delivery. Critically, we observe significant skew in delivery along gender and racial lines for 'real' ads for employment and housing opportunities despite neutral targeting parameters. Our results demonstrate previously unknown mechanisms that can lead to potentially discriminatory ad delivery, even when advertisers set their targeting parameters to be highly inclusive. This underscores the need for policymakers and platforms to carefully consider the role of the ad delivery optimization run by ad platforms themselves—and not just the targeting choices of advertisers—in preventing discrimination in digital advertising."

---

[4] 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

[5] 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

[6] Ali, Sapiezynski, Bogen, Korolova, Mislove, Rieke, "Discrimination through optimization: How Facebook's ad delivery can lead to skewed outcomes" (2019) https://arxiv.org/pdf/1904.02095.pdf

As a result of this and other lines of research, the US Department of Housing and Urban Development (HUD) is currently investigating Facebook[7], and it is also reported that Google and Twitter are being probed for housing discrimination[8].

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.[9] 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 disparities are most likely due to gender bias in hiring exhibiting itself in the data on which the screening tool was trained: they aligned with, and reinforced, a stark gender imbalance in the workforce at Amazon and other platforms, particularly when it comes to technical roles. Interestingly, despite essentially unlimited data, computational, and human resources, Amazon was unable to fix the problem of bias in hiring by means of a purely technological intervention.

Interviewing: In 2014, it was reported that online personality tests, often used as part of the interviewing process, disproportionately reject candidates suffering from mental illness such as depression and bipolar disorder *even if* they have the right skills for the job.[10] There is much to be said about discrimination based on disability status, which arises due to a multitude of factors, including the under-representation of individuals with disabilities in the training and validation data, and the choice of data representations and effectiveness metrics. Importantly, this type of discrimination is notoriously difficult to detect because individuals with disabilities often do not report their disability status.

In summary, numerous cases of discrimination based on gender, race, and disability status during sourcing, screening, interviewing, and selection[11] stages have been documented in recent reports. These examples show that, if left unchecked, automated hiring tools will replicate, amplify, and normalize results of historical discrimination in hiring and employment.

---

[7] US Department of Housing and Urban Development, Office of Administrative Law Judges (2018) https://www.hud.gov/sites/dfiles/Main/documents/HUD_v_Facebook.pdf

[8] Robertson, "*HUD reportedly also investigating Google and Twitter in housing discrimination probe*", The Verge (2019) https://www.theverge.com/2019/3/28/18285899/housing-urban-development-hud-facebook-lawsuit-google-twitter

[9] 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

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

[11] Emerging Technology from arXiv, *"Racism is Poisoning Online Ad Delivery, Says Harvard Professor*", MIT Technology Review (2013) https://www.technologyreview.com/2013/02/04/253879/racism-is-poisoning-online-ad-delivery-says-harvard-professor/

Recommendation 1: Expanding the scope of auditing

Bias audits should take a broader view, going beyond disparate impact when considering fairness of outcomes. Others will speak 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."[12]

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

A subtle but important point is that even features that can legitimately be used for hiring may capture information differently for different population groups. For example, it has been 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.[13] 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*.[14] 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

---

[12] Yang, "*Ensuring a Future that Advances Equity in Algorithmic Employment Decisions*", Urban Institute (2020) https://www.urban.org/research/publication/ensuring-future-advances-equity-algorithmic-employment-decisions

[13] Reeves and Halikias "*Race gaps in SAT scores highlight inequality and hinder upward mobility*", Brookings (2017) https://www.brookings.edu/research/race-gaps-in-sat-scores-highlight-inequality-and-hinder-upward-mobility

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

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.[15]

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 publicly available.

Recommendation 2: Explaining decisions to the job applicant

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

I recommend to build explanations for job seekers 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.[16]

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.[17]

Recommendation 3: Creating an informed public

To be truly effective, proposed law relies on 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

---

[15] Stoyanovich and Howe, "*Follow the data: Algorithmic transparency starts with data transparency*" (2019) https://ai.shorensteincenter.org/ideas/2018/11/26/follow-the-data-algorithmic-transparency-starts-with-data-transparency

[16] 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

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

standards will strengthen the accountability structures and help build public trust in the use of ADS in hiring and beyond.

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 example. This aligns with a set of recommendations by the Automated Decision Task force[18], on which I served, but we have not yet seen the City act on these recommendations.

We heard from members of the administration that public engagement activities have slowed down due to the COVID-19 pandemic. This does not have to be the case: based on our own experience, described below, there is substantial interest from the public to participate, and an opportunity to effectively use online platforms to educate and engage them. *One of the activities we have been ramping up at the Center for Responsible AI at NYU in recent months has focused specifically on public education and engagement around the use of ADS in hiring.* In collaboration with the Queens Public Library, the Center conducted a series of sessions called "Uncovering Hidden Decisions: AI in Hiring", where we gave an introduction to data, algorithms, and AI, gave actionable advice to job seekers, told participants about Int 1894, and asked for their thoughts on the bias auditing and public disclosure components of the proposed law. The final session will take place on November 17, 2020. I will be happy to share what we learned during these sessions with the Committee on Technology once the series completes.

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".[19] IEEE is the largest professional organization of engineers in the world; I have the pleasure of serving on their AI/AS (Artificial Intelligence / 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

---

[18] See Section 2 of the New York City Automated Decision Systems Task Force Report (2019) https://www1.nyc.gov/assets/adstaskforce/downloads/pdf/ADS-Report-11192019.pdf

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

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.”

Testimony of Manish Raghavan before New York City Council Committee on Technology regarding Int. 1894

Manish Raghavan

November 9, 2020

Dear members of the committee,

My name is Manish Raghavan. I am a researcher at Cornell University studying the societal impacts of algorithmic decision-making, particularly in the context of hiring. I have extensively studied the types of automated employment decision tools being discussed today, and my testimony is largely based on this research.[1] In this testimony, I offer my recommendations regarding Int. 1894, which seeks to regulate algorithmic tools deployed for candidate evaluation.

I appreciate the Council's attention on this important topic. Automated employment decision tools are increasing in prevalence, often with little to no public transparency into their inner workings. In my view, this bill is a step in the right direction. In its current form, it carries some vital provisions to ensure that automated hiring tools are carefully scrutinized for potential discrimination.

At the same time, it's important to recognize the limitations of this bill (and indeed, any attempt to regulate these tools through prospective auditing). In this testimony, I will detail two such limitations:

1. Current interpretations of anti-discrimination law do not preclude all discriminatory behavior that algorithms can exhibit.

2. Audits have limited power to detect discrimination in terms of undisclosed attributes, such as sexual orientation or disability status.

Before diving deeper into these points, it's important to note that hiring tools can perpetuate discrimination even in the absence of explicit bad actors. Due to historical patterns of inequity, algorithms can behave in discriminatory ways simply due to negligence or insufficient attentiveness to these issues. It's crucial that we implement guardrails that protect us from these more subtle, insidious forms of discrimination.

[1] Raghavan, Manish, et al. "Mitigating bias in algorithmic hiring: Evaluating claims and practices." Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency. 2020; Raghavan, M., and S. Barocas. "Challenges for mitigating bias in algorithmic hiring." Brookings. 2019.

Current interpretations of anti-discrimination law do not preclude discriminatory behavior. Vendors of automated employment decision tools, to the extent that they consider issues of bias at all, typically think of anti-discrimination law in terms of the EEOC's 4/5 rule. The 4/5 rule requires that applicants from different protected groups be selected at roughly the same rate—that is, if half of the candidates evaluated are women, then approximately half of the candidates selected by the tool should be women. A violation of the 4/5 rule do not necessarily constitute discrimination, but it can be the basis to open a discrimination suit.

In the absence of specific requirements, it is natural that bias audits will focus on ensuring that tool in question satisfies the 4/5 rule. In my view, this is insufficient, and inconsistent with standards in industrial-organizational psychology.[2]

A particularly important metric to consider is validity, which measures how good a tool is at correctly identifying high- vs. low-performing candidates. How is validity related to bias? One key way in which algorithmic tools can discriminate is via **differential validity**, which occurs when a tool is better at evaluating members of one group than another. For example, if the tool is very good at identifying the top-performing white candidates and not very good at identifying the top-performing African-American candidates, this would be an instance of differential validity.

Even if an assessment satisfies the 4/5 rule, meaning it recommends candidates from all racial groups at roughly equal rates, the top-performing African-American candidates would be more likely to be screened out by the assessment than their white counterparts. Differential validity has been repeatedly found in practical applications of data-driven decision-making,[3] and it's important to ensure that employment decision tools don't perpetuate this form of discrimination.

Assessments that exhibit differential validity are not explicitly illegal, according to current interpretations of the law. However, simply **requiring an auditor to report on measures of differential validity** may induce vendors of automated employment decision tools to ensure that their products work well for everyone, not just those who have been well-represented in historical data. In my view, testing whether a tool performs well across the entire population should be an integral part of any bias audit, and to this end, I believe this bill should explicitly require differential validity testing.

Audits have limited power to detect discrimination in terms of attributes like sexual orientation or disability status. Audits can only be performed with respect to protected attributes on which vendors maintain data. If a vendor doesn't collect data about, say, applicants' sexual orientation, it is impossible for an auditor to know whether a tool produces disparities along these attributes. Nor is it necessarily desirable that vendors maintain this sort of sensitive data; applicants may not feel comfortable divulging this information.

[2]Society for Industrial, Organizational Psychology (US), and American Psychological Association. Division of Industrial-Organizational Psychology. "Principles for the validation and use of personnel selection procedures." 2018.

[3]Buolamwini, Joy, and Timnit Gebru. "Gender shades: Intersectional accuracy disparities in commercial gender classification." Conference on fairness, accountability and transparency. 2018; Koenecke, Allison, et al. "Racial disparities in automated speech recognition." Proceedings of the National Academy of Sciences 117.14 (2020): 7684–7689.

Thus, an audit cannot identify all forms of illegal discrimination, and as such, it's important to be clear on the goals of such an audit. The current language of Int. 1894 refers to compliance with "any . . . applicable law relating to discrimination in employment." In practice, this will not be possible. We should **acknowledge the narrow scope of what is possible through audits**, and what forms of discrimination cannot be detected through these means.

Recommendations. While the above challenges are in a sense inherent to the problem of auditing for bias, there are concrete steps we can take to begin to address them.

1. Set specific standards for what measures should be included in an audit.
2. Require auditors to report on metrics of differential validity.
3. Use caution in interpreting the results of audits. An audit can only test for specific discriminatory behaviors; it cannot certify that a tool is free of bias.

Thank you for your attention.
Manish Raghavan

PARITY

Dr. Rumman Chowdhury
CEO and Founder
www.getparity.ai
rumman@getparity.ai
Testimony Date: November 11, 2020

This testimony is in regards to Int. 1894: A Local Law to amend the administrative code of the city of New York, in relation to the sale of automated employment decision tools

As a leader in the responsible AI community, currently as CEO and founder of Parity, an enterprise ethical AI audit company, and formerly as Accenture's Global Responsible AI lead, I applaud the City's foresight in proactively addressing the harms that can be introduced by automated decision making systems, and in particular, employment algorithms.

In this testimony, I address three critical components of this bill's success - first, constructing an actionable audit, second, instituting methods of citizen redress and highlighting harms, and third, providing a pathway for companies subject to this audit to share their process output without revealing intellectual property or exposing their data. These parts are necessary to create a successful cyclical, self-reinforcing and evolving audit methodology.

First, introducing mandatory audits of HR algorithms is a necessary first step in ensuring responsible use of these systems. However, the current state of practice in our industry is that no clear standards or universally-accepted guidelines exist to perform model audits. To date, audits range from purely qualitative assessments that result in lengthy documents intended for C-suite and management, to purely technical platform-based implementations that assess only the technology and it's outputs. Both are insufficient.

The former ignores the reality of addressing limitations within the technology, often resulting in wishful thinking that is simply not measurable or implementable, and the latter ignores the contextual applications of technology and how it interacts with humans, society, and organizational structure.

My first suggestion to the Council is to create a working group to generate these audit guidelines in collaboration with the multiple bodies that already exist to create responsible AI solutions - organizations like New York University's Alliance for Public Interest Technology, and Algorithmic Advisory Alliance, of which I am a founding member. In my experience, groups like this are the most successful when they incorporate companies, civic organizations, policymakers, and technologists.

In the solutions I have built, I have found that successful audit methodologies include transparent standards and timelines, actionable suggestions for correction, broad stakeholder

input, and the ability to reject an algorithm if standards are not met. The most rigorous audits are often conducted, or at least approved, by third-party organizations that exist independently of implementation.

Audit output needs to be understandable by both technical and non-technical audiences and made available to the public. Implementation requires education. If an audit is enacted, the City has the obligation to ensure city employees are able to utilize audit output appropriately and properly, as end-user errors are often a major source of risk in the implementation of any technical system.

My second suggestion to the Council is to introduce methods for individual citizens to highlight harms. While transparency is laudable, it ignores the power dynamic that exists between employers and potential employees. Notification of algorithmic usage without a clear method to redress perceived harm will not be beneficial, and likely will raise more questions than it will answer.

In developing these solutions, I have found that structured community engagement is key to proactively identifying problems. This can be accomplished via algorithmic harms reporting databases, similar to reporting structures that exist in cybersecurity and other industries. One such initiative is in its nascent stages - the Algorithmic Justice League's CRASH database. Relatedly, structured community interviews and feedback analysis is critical to evolving the audit methodology as new challenges emerge.

Finally, my third suggestion is to create a working environment for companies who would be subject to these audit methodologies to safely and securely share their intellectual property and data with audit developers. Similar 'regulatory sandboxes' have been utilized successfully to create policy for financial regulation and data privacy. Collaborative creation allows for realistic solutions, testing and iteration.

Beyond these three structural suggestions, a final recommendation to the City is to align with local, national, and international best practices and policies, currently in development. New York City is an influential and notable player in the ongoing global discussion on algorithmic fairness and equity. This law can have positive national and international ramifications if structured appropriately.

Our collective goal is noble - institute forward thinking policies to ensure all constituents reap the benefits of algorithmic decisionmaking systems used in employment decisioning while mitigating and addressing harm. In order to do so, we will need a collaborative effort of citizens, experts, policymakers, and companies developing an audit methodology that draws from existing best practices and evolves to meet the diverse needs of all New York City citizens.

Testimony of Athena Karp, CEO & Founder of HiredScore

Before the New York City Council Committee on Technology

Regarding Int. 1894-2020: Sale of automated employment decision tools

13 November 2020

My name is Athena Karp and I am the CEO and founder of a NYC-based start-up called HiredScore, which could be subject to the proposed legislation. Prior to starting HiredScore, a company that helps employers efficiently and fairly evaluate job candidates, I worked as an investor, with a particular focus on technological innovations in highly-regulated data environments, primarily healthcare. Today, I am grateful for the opportunity to speak to this Committee about my support for this legislation.

As this hearing has made clear, opaque and biased hiring tools can have very real consequences on the lives of New York workers and workforce diversity, as well as job opportunity. Additionally, if improperly designed, automated employment decision tools could create challenges for employers seeking to increase workforce diversity and comply with anti-discrimination laws. As a person who has spent much of my career working in the intersection of next generation technologies and ethical and effective ways to leverage technology, I can attest to the fact that these problems are avoidable. If properly designed, hiring technology can - and should - promote fairness, transparency, and efficiency across the entire process, as well as importantly provide a more respectful and consistent treatment for all job seekers. Unfortunately, it is not always easy for employers to determine which products align with these priorities and which do not. The net effect is that New York workers are unable to benefit from the full potential of advancements in HR technology.

The reason I support this legislation is simple: employers need and deserve better information about the potential implications of the systems they use to evaluate job candidates and the ability to differentiate between tested and explainable solutions from bias-generating or opaque solutions. As purchasers of novel technology, employers should be able to trust the outcome of an audit report so they can make more informed decisions and have a consistent standard applied to employment decision solutions. Without legislation, the HR tech industry will continue to be characterized by ambiguity and employers, especially smaller businesses, will be challenged to define and carry-out audits that demonstrate fair and ethical solutions. We have likely all heard about cases in which algorithms have produced problematic outcomes over the last few years; however, some vendors have invested heavily and spent years proactively designed systems to be explainable, equitable, and auditable. For most employers, it is essential that reporting structures are established to shed light on the fact that some hiring tools take a proactive approach to fairness and others do not. I firmly believe that many New York employers *want* to use technology that can improve the hiring function without introducing bias. I also believe that New

York employers want to use technology that can accelerate achievement of diversity goals. This proposal is an example of how policymakers can help them navigate toward these laudable and important goals.

I would also like to offer a few comments on the utility of the bias audits proposed in this bill. After extensive time working in this space, I fully understand the public concern over new technologies, especially those related to hiring. Given the ever-growing mandate of human resource teams, the availability of data, and new agendas that may not have been carried out before, such as driving diversity of hiring goals, the reality is that technology-based hiring solutions will only grow in importance to employers in efficiently, effectively, and fairly carrying out such work, especially in the unprecedented times in which we are living. Notably, the term "automated hiring system" does not necessarily mean a fully automated hiring solution; it can also refer to a solution that recruiters may use to disqualify or advance candidates. Regardless if an automated hiring system is relying on technology from 1970 or 2020, the companies that build and support these systems need to make countless choices about how to design, build, and invest in innovation to make more explainable, seek and follow external advisors on best practices, and maintain them with vigilance and continued care. With such a wide range of products on the market, employers may not always know which questions to ask to compare their options and the approaches of various vendors. However, one question that is relevant for virtually every type of hiring tool is the extent to which it introduces or increases biases against candidates from different demographic backgrounds. This is the question bias audits can help employers proactively answer in a standardized, consistent, and trusted manner.

Of course, while I firmly believe that well-designed technology can greatly improve the quality of the hiring process for both employers and candidates, this is not to suggest that any HR tool will be ethical or take the steps to test for and ensure no bias from such tools. I started HiredScore in New York because I believe that this city is predominantly made up of employers who are eager to advance while ensuring laws and standards are followed. I believe that New York business-owners want to ensure that job candidates here are treated irrespective of the technology or human processes that are in place. With the leadership exhibited by this bill's sponsors, I am optimistic that the good intentions of these organizations to make hiring more fair will become a reality. Thank you for your attention to this important issue, and for the opportunity to share my perspective.

Testimony of Dr. Frida Polli, CEO & Co-Founder of pymetrics

Before the New York City Council Committee on Technology

Regarding Int. 1894-2020: Sale of automated employment decision tools

13 November 2020

My name is Dr. Frida Polli and I am a former Harvard and MIT neuroscientist turned CEO and co-founder of a NYC-based startup called pymetrics. pymetrics is a vendor of employment selection technology, meaning we are an organization that would be directly affected by this legislation's passing. While tech companies are often known to oppose regulation, this is not the perspective I hope to share with you today. As an organization that is committed to using technology for good, I would like to take this opportunity to explain why pymetrics strongly supports this proposal and applauds its sponsors for taking this important step for fair hiring.

The use of "employment selection tools" is not a frequent topic in the discourse on diversity and inclusion, but the reality is they can have dramatic effects on workforce diversity. Companies like pymetrics essentially build systems that help employers decide which candidates should receive further consideration for jobs. These systems might take the form of static multiple-choice tests, AI algorithms, or resume parsers. In today's globalized and digitized world, large employers may receive hundreds of applications for a single role, so it is inevitable that HR teams implement some means of initially sorting candidates. If the screening procedure is biased, however, the result can be that certain candidates are continually barred from opportunities. Right now, the only way such biases come to light is after an employer has been sued for discrimination. This bill is different because it requires vendors to proactively look for signs of unfairness and essentially label their products accordingly.

As someone who has been building and selling hiring tools for the past several years, I have spent considerable time thinking about the information employers use to decide what kind of HR technology they are going to purchase. In my opinion, there is no reason why clear information about the bias in a hiring tool should not be part of this equation. Over the past several months, countless employers have made commitments to improving workforce diversity. At the same time, many of these same employers are using HR tools that screen out disproportionate numbers of minorities in the hiring process. This bill is a crucial step to overcoming this disconnect: with transparent information about the fairness of hiring tools, well-intended employers will be empowered to implement diversity-friendly systems. Recent advancements in technology make this process of proactively testing tools for bias easier than ever before.

In addition to the transparency this bill will bring to the market of employment selection tools, I strongly support this legislation as someone who believes technology can be a force for good. Today, many people are unsure about the role artificial intelligence (AI) should play in making important decisions. Some even suggest that we abandon automated systems altogether and resort to manual resume reviews. As an academic neuroscientist, I am extremely confident that relying on human decisions to promote fair hiring is not the solution. No matter how well-intended HR personnel may be, humans have unconscious bias that is documented by research studies and is unchanged in 30 years.[1] Additionally, other studies show us that

[1] Quillian, L., Pager, D., Hexel, O., & Midtbøen, A. H. (2017). Meta-analysis of field experiments shows no change in racial discrimination in hiring over time. *Proceedings of the National Academy of Sciences*, 114(41), 10870-10875.

unconscious bias training is ineffective in removing said biases.[2] In contrast, algorithms can be intentionally designed to mitigate bias, and in cases where problems arise, they can be amended going forward. Note that not all, nor even most, algorithms do so at the moment. However, with the bias audits proposed in this bill, employers can essentially differentiate between those hiring tools that were designed with diversity in mind and those that were not.

Of course, in order for this bill to achieve its intended goals, the implementation must be carefully considered. Over the past several years, many prominent voices have produced governance frameworks for ethical AI. While frameworks may appear relevant to this legislation at first glance, I strongly urge the Committee to remember that the scope of this bill is not limited to AI hiring systems. The goal here is to require all vendors to produce comparable fairness metrics about their tools, so it is essential that the structure of bias audits must be relevant for all industry players. Further, if the goal of bias audits is to understand how a hiring tool will affect real New Yorkers, the focus of these audits should be on *outcomes* first and foremost, rather than being concerned with the inner workings of technology. Only a focus on outcomes will allow HR professionals to understand if a tool they are considering, be it AI or not, produces biased results.

Another point I would like to make on the implementation of this bill comes from many conversations I have had with employers who are clients of pymetrics. I have come to believe that many employers in New York City are sincere in wanting to improve the diversity of their workforce and are looking for solutions to help them do so. Nevertheless, they also recognize that employment selection procedures are highly regulated, and the prospect of new technology can introduce legal ambiguity. In order for this legislation to be effective, bias audits must straddle a delicate balance between providing transparent information and not being so arduous as to discourage employers from using hiring technology altogether. As I mentioned, a reversion to human-driven hiring decisions is not a viable path to workforce diversity.

Some of you may be curious about the extent to which fairness-optimized hiring technology can actually make a difference for real New Yorkers. To fully appreciate the effects, it is important to recognize the full extent of bias in the status quo. For example, many employers today still use standardized cognitive ability tests that were first popularized in the nearly 50 years ago. When an HR department uses this type of screening tool though, the degree of bias is staggering: white candidates are selected at *more than three times* the rate of black candidates. In contrast, because pymetrics proactively audits our system for bias, applicants of all races and genders are selected at equitable rates. We are not the only company that takes such steps to promote fairness, but our goal in supporting this legislation is to help employers identify pro-diversity solutions wherever they exist.

On a personal note, I would like to thank the sponsors of this bill for their leadership on this issue. When I founded pymetrics in 2013, I did not necessarily fit the profile of a typical tech entrepreneur. I had spent much of the first 15 years of my career as an academic neuroscientist, but financial necessity ultimately motivated my departure from research. I was a 37-year-old single mom and breadwinner at the time I found myself pitching my start-up idea to investors. In my own way, I have felt the frustration of being stereotyped based on factors that had nothing to do with my abilities or potential. I am optimistic that ethically-designed hiring technology, governed by pragmatic auditing and reporting standards, will make this experience less common for New York workers.

[2] Forscher, P. S., Lai, C. K., Axt, J. R., Ebersole, C. R., Herman, M., Devine, P. G., & Nosek, B. A. (2019). A meta-analysis of procedures to change implicit measures. *Journal of Personality and Social Psychology, 117*(3), 522–559.

ADDENDUM

TO: New York City Council Technology Committee
FROM: Frida Polli, PhD, CEO at pymetrics; Kelly Trindel, PhD, Head of Policy at pymetrics; Sara Kassir, MPP, Senior Policy Analyst at pymetrics
DATE: 13 November 2020
RE: Recommendations for the structure of "bias audits," per Int 1894-2020

The goal of this memo is to propose an appropriate framework to govern the "bias audits" that would be required by NYC INT 1894-2020. Our position is that this bill has genuine potential to improve equity in hiring in New York City, so long as the standards for compliance are clear, demonstrable, and achievable. A few points about the practical implementation of bias audits are especially important. (1) In our view, bias audits are a means of creating credible "labeling" for the HR technology marketplace, so that employers have meaningful information to compare the fairness of different vendors' tools. (2) For such labels to be useful, they must closely align with existing federal laws regarding employment procedures, including Title VII of the Civil Rights Act of 1964 and the Uniform Guidelines on Employee Selection Procedures.[3] (3) We strongly believe that bias audits should be focused on the evaluation of outcomes yielded by a hiring tool, rather than its design or mechanics, to have the best prospects for improving workforce diversity. These tenets and others are explained in the proposed framework detailed below.

BACKGROUND:

A major feature of NYC INT 1894-2020 is the requirement of an annual "bias audit" for employment decision tools. Upon taking effect, this legislation would make New York City one of the first U.S. jurisdictions to require such an assessment as a condition for the legal sale of employment selection technology**.** Three definitions from the proposed statue's text provide the necessary context.

First, the proposed bill defines **"bias audit"** with the following statement:

> "The term 'bias audit' means an impartial evaluation, including but not limited to testing, of an automated employment decision tool to assess its predicted compliance with the provisions of section 8-107 and any other applicable law relating to discrimination in employment."

In other words, the exact nature of what constitutes a bias audit remains unspecified in the text of NYC INT 1984-2020. Unanswered questions include: Who will conduct the audits? What procedures, if any, must they follow? What information must be reported? To whom must the information be reported?

Second, the **scope of technology** subject to this legislation is specified with:

> "The term 'automated employment decision tool' means any system whose function is governed by statistical theory, or systems whose parameters are defined by such systems, including inferential methodologies, linear regression, neural networks, decision trees, random forests, and other learning algorithms, which automatically filters candidates or prospective candidates for hire or for any term, condition or privilege of employment in a way that establishes a preferred candidate or candidates."

[3] 43 FR 38290 et seq. August 25, 1978)

Based on this broad definition of "automated employment decision tool," it is our understanding that this bill is intended to regulate the entire universe of hiring tools, ranging from those built with sophisticated machine learning methods to those have existed for decades and are static in nature. Indeed, we believe that it is very important that bias audits are relevant for the full range of products an employer might purchase to best inform fair hiring strategies.

Finally, "employment decisions" are defined with:

> "The term 'employment decision' means to screen candidates for employment or otherwise to help to decide compensation or any other terms, conditions or privileges of employment in the city."

Notably, while our proposed framework addresses all types of hiring tools, it does not address tools designed to "decide compensation or any other terms, conditions, or privileges of employment." This omission is intentional, because our experience indicates that vendor-supplied tools for automating compensation decisions are exceedingly rare. In fact, we are not aware of any that exist in New York City. Bias audits on compensation decisions would therefore need to be conducted directly on the employer, rather than on external vendors, which is beyond the scope of this bill. Here, we encourage a narrow regulatory scope that best reflects existing products on the market that have real consequences for job seekers.

GUIDING ASSUMPTIONS

a. Statutory Basis:

As written, the proposed amendment states that auditing shall test an employment selection tool's compliance with "the provisions of section 8-107 [of the New York City Administrative Code] and any other applicable law relating to discrimination in employment."

Per Title VII of the U.S. Civil Rights Act of 1964, discrimination in employment can take two forms: *disparate treatment* (e.g., direct, intentional discrimination) or *disparate impact* (e.g., indirect discrimination via a procedure that disadvantages a group or class of people). However, the specific terms of Section 8-107 deal primarily with disparate treatment, which requires inter-personal animus and is less relevant in the context of employment decision tools.

In the absence of city law regarding disparate impact (also known as *adverse impact*), Bias Audits should be informed by relevant federal statutes and guidance published by the U.S. Equal Employment Opportunity Commission. (See III.A. below for further information.)

Determinations of discrimination are notably outside the scope of NYC INT 1894-2020. However, Bias Audits may be thought of as the proactive production of metrics and information that would be considered during a disparate impact lawsuit.

b. Goals of Bias Audits:

Bias Audits are intended to provide clear information regarding employment decision tools, specifically related to their *fairness*, here defined as degree of disparate impact. There are three primary reasons these audits should be undertaken.

First, in light of increased calls to improve workforce diversity, Bias Audits can **provide employers with clear information about the fairness performance of a specific employment selection tool.** Right now, it is exceedingly difficult for employers to proactively consider how different selection procedures might increase or decrease their ability to identify diverse candidates. If implemented properly, Bias Audits serve as a "preview" of how a vendor's tool will fare in practice, better positioning employers to compare across hiring assessments and make informed decisions about which to use.

Second, the requirement of Bias Audits may **encourage vendors who have long neglected fairness considerations to amend their tools to better reflect modern workforce priorities.** Due to existing employment regulations dating back several decades, certain hiring tools that are well-known to disadvantage minority job applicants are technically legal on the basis of "business necessity." While it was significantly more difficult for vendors to design fair hiring tools at the time such regulations were established, this is no longer the case with advancements in technology, and the market should adapt accordingly.

Third, Bias Audits can help **educate the public on the reality that employment selection tools can either help or hinder workforce diversity.** Today, selection tools are largely ignored in public discourse on diversity and inclusion, since many people wrongly assume a "test" is always impartial. Where public narratives do exist, they are largely opposed to technology's role in the hiring process, claiming to prefer "human" evaluations. This widespread perception is extremely problematic because of two facts: 1. human beings are biased[4] and 2. well-designed modern selection tools can provide more equitable evaluations. In order for employers to have latitude to adopt modern hiring tools, public discourse around the issue must be advanced.

c. Balancing Stakeholder Priorities:

Questions of workforce diversity and employment law inevitably interest a broad range of stakeholders, including employers and their legal counsel, employment discrimination plaintiffs, local government agencies, and employment selection tool vendors. Because Bias Audits are effectively reporting requirements placed on vendors, these stakeholders may express different opinions about the appropriate contents of these reports. This conversation is in some ways complicated by opinions from "AI ethics" experts, who are narrowly focused on creating governance structures to root out nefarious technology.

Stated plainly: in order for Bias Audits to be effective, they must provide information in a manner that does not disincentivize employers from purchasing and using employment selection tools that are subject to the statute. It is therefore extremely important that any reporting requirements do not directly introduce legal liability for employers. Further, in order for Bias Audits to be useful, they must facilitate meaningful comparisons across all types of hiring tools. Employers and the public need to be assured that all vendors are reporting analogous information and that those who fail to do so are held accountable.

RECOMMENDATIONS

d. Governance of Bias Audits

Though audits in many contexts are conducted by external professionals, this does not need to be a requirement for Bias Audits of employment selection tools. Third-party audits are generally very useful for

[4] Tversky, A., & Kahneman, D. (1974). Judgment under Uncertainty: Heuristics and Biases. *Science, 185*(4157), 1124-1131.

validating the veracity of a vendor's marketing claims, some of which may or may not be directly related to issues of fairness and discrimination. In contrast, the goal of Bias Audits in this context is more narrow and targeted: to produce accessible and consistent metrics across the market of hiring tools, informed by very specific instructions to vendors on what must be shared with relevant stakeholders.

e. Reporting Requirements

For the purpose of NYC INT 1894-2020, a vendor should be defined as any entity that sells a product indicated for the screening and/or selection of job applicants in New York City (see the Office of Federal Contract Compliance definition of what constitutes a job applicant).[5] Three types of reporting requirements are relevant to vendors.

i. Public Reporting by Assessment Type

In order for vendors to be in compliance with NYC INT 1984-2020, they must submit the following information to the Human Rights Commission of the City of New York by no later than December 31st each year. A standardized Bias Audit form should be used for submissions (see Appendix I for sample). Public reporting need not be provided at the individual employment decision tool level but rather aggregated across individual tools at the assessment type level. Many different assessment types exist, such as cognitive ability tools, personality or job-fit tools, job simulation tools, interview analytic tools and hard skill assessment tools. The vendor should group tools it sells by assessment type and report the results of aggregated fairness outcomes by assessment type publicly.

The Bias Audit should require that vendors of employment selection tools follow three steps to produce estimates of the tool's fairness:

1. Identify a sample of individuals who have been evaluated, scored, or otherwise assessed by tools in the assessment type in the previous 12 months.
    - The sample must:
        - Include individuals with voluntarily supplied, self-identified race/ethnicity and/or gender data
        - Include at least 2% of people in any reported race/ethnic or gender sample[6]
        - Include only anonymized data regarding individual or employer identities (no PII)

2. Differentiate between individuals who were selected ("passed") vs. those who were not selected ("failed") by tools in the assessment type.
    - This process must:
        - Make a binary distinction between passes and fails
        - Be fully objective and replicable using only information collected by the vendor's tool
        - Align with the guidance the vendor generally provides to employers on how to interpret the results of the tool

3. Conduct an adverse impact analysis, as specified by the EEOC's Uniform Guidelines on Selection Procedures (UGESP) on tools in the assessment type.
    - This analysis must:

[5] https://www.dol.gov/agencies/ofccp/faqs/internet-applicants#Q2GI
[6] Ideally sample size for the analysis is 100 or more candidates with at least 30 in each group.

- Include all **racial/ethnic** and **gender** groups representing at least 2% of the demographically-labeled sample
- Calculate selection (pass) rates for each gender and race/ethnic group.
- Produce a single adverse impact ratio for each gender and race/ethnic group by comparing their selection rates as compared to the group with the highest selection rates[7].

Additionally, the Bias Audit should require that vendors of employment selection tools answer the following questions for each assessment type:

1. Are all job applicants for a particular position assessed using the same individual assessment type?
   If no, please explain why
2. Do any employment decision tools consider race, ethnicity, or gender? Or are proxy variables for these categories included in the device?
   If yes, which variables?
3. Has the employment decision tool undergone a third-party audit?
   If yes, are the results of the third-party audit available to the purchaser? Are they public?
4. How often is the employment decision tool updated or refreshed? An update or a refresh means any change to the assessment items or the decision scoring mechanism.
   Are updates recurring on a regular basis? Describe how updates are done.
5. What job qualifications or general characteristics does the employment decision tool assess?
   a. *How are candidates informed of these job qualifications or characteristics?*

See Appendix I for a sample form to be completed by the vendor for public reporting by assessment type.

ii. Information Shared with the Purchaser Prior to Use

Vendors should supply to purchasers at the time of sale a copy of their most recently-submitted public forms. In addition to this annual public reporting from vendors of employment decision tools aggregated by assessment type, NYC INT 1984-2020 should require that vendors provide employers with additional information to ensure they are well-informed about the expected fairness performance of the particular employment decision tool that they intend to utilize. While public reporting should be aggregated by assessment type, direct-to-purchaser reporting should include greater specificity by individual tool or customized algorithm. This information should be provided to the purchaser prior to use of a particular tool on applicant data.

For prospective clients, vendors have a responsibility to provide useful information to aid the employer in understanding the fairness implications of the tool in their use case. Vendors are not permitted to charge for these services. If the vendor does not have a sufficient sample to provide this fairness performance information prior to use by the purchaser, they must clearly communicate this to the purchaser, indicating the reason for this and plans to develop such a sample.

See Appendix II for a sample form to be completed by the vendor for direct-to-purchaser reporting prior to use.

[7] An example, illustrating how to calculate adverse impact ratios, is included in Appendix I. See: https://www.shrm.org/resourcesandtools/tools-and-samples/toolkits/pages/avoidingadverseimpact.aspx

iii. Information Shared with the Purchaser Ongoing and Annually

Once an employer has used a vendor's tool to evaluate actual job applicants, vendors have a responsibility to compile adverse impact metrics for the employer and provide this post-deployment report in a timely manner. In other words, it should not be permissible for a vendor to simply send hiring tests to an employer without monitoring the subsequent fairness performance.

In order for a post-deployment Bias Audit report to be completed, the vendor must complete the following steps:

1. Develop a system for collecting the voluntary self-identification of demographic information from applicants who complete the test or assessment. If this cannot be done, then develop a system for collecting anonymized voluntary self-identification information from the purchaser specific to applicants who complete the test or assessment. The transfer of Personally Identifiable Information (PII) is not required in this process.
2. Report to the purchaser on binary (pass/fail) outcomes of the employment decision tool by race, ethnicity, and gender. The binary outcome should mirror the client's actual use of the tool as understood by the vendor.
3. Complete this process at least annually as long as the client utilizes the tool for employment decisions and the sample size reaches at least x. Include in the analysis results for each demographic group that reach 2% or more of the total sample size.

The Human Rights Commission will not proactively collect post-deployment reports from vendors, but random audits may include a review of evidence that such reports were completed and shared with employers on an annual basis.

See Appendix III for a sample form to be completed by the vendor for direct-to-purchaser reporting ongoing and annually.

CONCLUSION

We strongly support Int 1894-2020 as a mechanism for providing employers and the public with information about the fairness of employment selection systems. We believe that audits that focus more on outcomes than on the inner-workings of a given system will be most useful for making meaningful comparisons across the HR tech marketplace. Additionally, bias audit reports should have multiple audiences to balance the need for transparency with employer concerns about legal liability, as we have demonstrated in the appendices of this document. If this legislation is pragmatically implemented in a manner that aligns with existing employment law, we are very optimistic about the prospects for New York City to set an important precedent for the future of fair hiring technology.

APPENDIX I
SAMPLE BIAS AUDIT FORM
FOR PUBLIC REPORTING BY ASSESSMENT TYPE
(To be completed by vendor at least annually)

DATE OF SUBMISSION
COMPANY NAME

1. Identify the assessment 'type' for this analysis. Indicate the approximate number of individual employment decision tools aggregated within this 'type.'

2. Describe how the sample sourcing the information below was collected, identified, or otherwise determined. If the sample is not comprised of previous job applicants, explain why not.

3. If the sample is not comprised of previous job applicants, describe the general similarities and differences between the sample sourcing the information below and the typical applicant pool for relevant jobs.

4. Describe how individuals in this sample were designated as "selected" (e.g., passed) or "not selected" (e.g., failed) by your system.

5. Using the process you described above, complete the below table for adverse impact testing. (See figure below for example)

		Gender		*Race/Ethnicity*			
	All	Male	Female	White	Black	Hispanic	Asian
n Total							
n Selected (passed)							
n Not Selected (failed)							
Selection Rate (%)							
Impact Ratio							

6. Are there multiple tools within this assessment type designed to assess or consider candidates for the same position or job?
   If yes, please explain why

7. Do any employment decision tools consider race, ethnicity, or gender? Or are proxy variables for these categories included in the device?
   If yes, which variables?

8. Have any employment decision tools in this assessment type undergone a third-party audit?
   If yes, are the results of the third-party audit available to the purchaser? Are they public?

9. How often are the employment decision tools in this assessment type updated or refreshed? An update or a refresh means any change to the assessment items or the decision scoring mechanism.
   Are updates recurring on a regular basis? Describe how updates are done.

10. What general characteristics do the employment decision tools in this assessment type designed to assess?
    How are candidates informed of these job qualifications or characteristics?

Step 1: Impact ratio for hiring—calculate the rate of selection

EEO Group	Applicants	Hires	Percent Hired
African-American	108	25	23
Latino	78	24	31
Caucasian	325	114	35

The group with the highest selection rate is Caucasian, with 35 percent. Next, calculate the impact ratio.

Step 2: Impact ratio for hiring—calculate the impact ratio

EEO Group	Percent Hired	Divide	Impact Ratio	Adverse Impact?
Caucasian	35			
African-American	23	23/35	66%	Yes; it is less than 80%
Latino	31	31/35	89%	No; it is more than 80%

https://www.shrm.org/resourcesandtools/tools-and-samples/toolkits/pages/avoidingadverseimpact.aspx

APPENDIX II
SAMPLE BIAS AUDIT FORM
FOR DIRECT-TO-PURCHASER REPORTING BY INDIVIDUAL TOOL
(To be completed by vendor prior to launch on applicant flow data)

DATE OF SUBMISSION
COMPANY NAME

1. Identify the assessment 'type' for this analysis. Indicate the approximate number of individual employment decision tools aggregated within this 'type.'

2. Identify the specific decision tool for this analysis and the recommended use case for the specific decision tool.

3. Describe how the sample sourcing the information below was collected, identified, or otherwise determined.

4. Describe how individuals in this sample were designated as "selected" (e.g., passed) or "not selected" (e.g., failed) by your system.

5. Using the process you described above, complete the below table for adverse impact testing.

		Gender		*Race/Ethnicity*			
	All	Male	Female	White	Black	Hispanic	Asian
n Total							
n Selected (passed)							
n Not Selected (failed)							
Selection Rate (%)							
Impact Ratio							

6. Is there more than one version of this tool in production to assess or consider candidates for the same position or job?

If yes, please explain why

7. Does this employment decision tool consider race, ethnicity, or gender? Or are proxy variables for these categories included in the device?
   If yes, which variables?

8. How often will this employment decision tool be updated or refreshed? An update or a refresh means any change to the assessment items or the decision scoring mechanism.
   Are updates recurring on a regular basis? Describe how updates are done.

9. What general characteristics does this employment decision tool assess?
   How are candidates informed of these job qualifications or characteristics?

APPENDIX III
SAMPLE BIAS AUDIT FORM
FOR DIRECT-TO-PURCHASER REPORTING ON APPLICANT FLOW ANNUALLY
(To be completed by vendor annually)

Please note that this form is meant to facilitate the vendor providing custom results to the employer on the outcome of adverse impact testing on their actual applicant flow. Best practice is for the vendor to report on both statistical and practical definitions of adverse impact. These reports should be completed annually. Ideally sample size for the analysis is 100 or more candidates with at least 30 in each group.

DATE OF SUBMISSION
COMPANY NAME

1. Identify the assessment 'type' for this analysis. Indicate the approximate number of individual employment decision tools aggregated within this 'type.'

2. Identify the specific decision tool for this analysis and the vendor's understanding of the use case for the specific decision tool.

3. Describe how the sample sourcing the information below was collected, identified, or otherwise determined.

4. Describe how individuals in this sample were designated as "selected" (e.g., passed) or "not selected" (e.g., failed) by your system.

5. Using the process you described above, complete the below table for adverse impact testing.

		Gender		*Race/Ethnicity*			
	All	Male	Female	White	Black	Hispanic	Asian
n Total							
n Selected (passed)							
n Not Selected (failed)							
Selection Rate (%)							
Impact Ratio							

Statistical Significance Test Result (p-value)							

6. Is there more than one version of this tool in production to assess or consider candidates for the same position or job?

If yes, please explain why

7. Does this employment decision tool consider race, ethnicity, or gender? Or are proxy variables for these categories included in the device?

If yes, which variables?

8. How often will this employment decision tool be updated or refreshed? An update or a refresh means any change to the assessment items or the decision scoring mechanism.

Are updates recurring on a regular basis? Describe how updates are done.

9. What general characteristics does this employment decision tool assess?

How are candidates informed of these job qualifications or characteristics?

10. What statistical significance test was used to determine adverse impact outcomes in the above? Why was this test utilized?

syndio
fairness at work.

November 12, 2020

My name is Zev J. Eigen. I am the founder and Chief Science Officer of Syndio Solutions (www.synd.io), the world's leading Equity Tech company that empowers organizations to make fair and equitable decisions about how to value and reward employees. To date, Syndio's pay equity software has helped level the playing field for 1.3M employees, totaling more than $85M in remediation paid by employers. Prior to founding Syndio, I led the analytics and data science function of the world's largest labor and employment management-side law firm, I served as a full-time faculty member at Northwestern University School of Law, visited as a faculty member at NYU Law and Yale Law School, worked as in-house counsel for Twentieth Century Fox Film Corporation, and have served as an expert on matters relating to technology and workplace compliance for many organizations. I hold a PhD from MIT, a JD from Cornell Law School and a BS from Cornell University's School of Industrial and Labor Relations. I also grew up in New York City from an early age until college.

The intent of this legislation appears to be to incent employers to apply technology responsibly by ensuring that "automated decision tools" are used in a fair and consistent way without inuring undue benefit to one group of a protected category over another. In my experience evaluating many systems on behalf of many employers over the years, "automated employment decision tools" applying "algorithmic methodologies" to filter candidates for hire or to make decisions regarding terms and conditions of employment vary in their degree of validity and transparency and in the diligence of the companies that make such tools in adhering to best practices in testing and proactively mitigating bias. So, in the broadest sense, I agree and support New York City's efforts to ensure that vendors are held to a high standard (by proxy) by requiring their customers (employers) to pay more attention to their output and in clarifying the government's expectations and standards for auditing of such solutions.

It is worth noting that there are software solutions designed to identify discriminatory practices, such as Syndio's PayEQ application that employs algorithms to pinpoint areas in which employers may be compensating or otherwise rewarding employees differently due to their protected category status. The software helps employers fashion fair and equitable remedies for these issues at the individual, group, and policy levels. I mention this because the law appears to assume a separation between "automated decision tools" employers may apply, and an audit of those tools to ensure that they are fair and equitable. That may be the case for most systems and vendors. However, as evidenced by the example of Syndio's PayEQ software and others, some tools actually provide a way for employers to audit themselves and hold themselves to a standard of reducing and eliminating bias in employment decision making. I would encourage lawmakers to ensure that employers are not penalized or burdened unnecessarily if and when they adopt such technologies. Additionally, the law makes no reference to auditing analog, or manual tools that employers commonly employ to guide employment decisions that often carry bias and contribute to discriminatory practices. Holding additional hearings to understand how Equity Tech vendors as well as other vendors in this space operate would benefit New York City employers and employees alike.

Testimony of

Christopher Boyle

Director of Data Research and Policy

New York County Defender Services

Before the

Committee on Technology

Oversight Hearing – Ethical Implications of Using Artificial Intelligence and Automated Decision Systems

&

Intro. 1894-2020

November 13, 2020

My name is Christopher Boyle and I am the Director of Data Research and Policy at New York County Defender Services (NYCDS). We are a public defense office that represents New Yorkers in thousands of cases in Manhattan's Criminal and Supreme Courts every year. I have been a New York City public defender for more than twenty years. Thank you to Chair Holden for holding this hearing on the use of artificial intelligence and automated decisions systems.

Automated Decision Systems are routinely used to inform actions at every step of the legal system. From the locations to which police are deployed to who gets released pretrial; from access to treatment and programs to the length of one's sentence or their eligibility for parole; algorithms are significantly influencing important criminal justice decisions. While a primary objective of such programs is to eliminate the effects of race or class biases, numerous studies have shown that without proper oversight, "risk assessments unintentionally amplify [these]... under the guise of science."[1]

[1] Picard, Sarah, Matt Watkins, Michael Rempel, and Ashmini G. Kerodal. "Beyond the Algorithm: Pretrial Reform, Risk Assessment, and Racial Fairness." Center for Court Innovation. July 2019. https://www.courtinnovation.org/sites/default/files/media/documents/2019-06/beyond_the_algorithm.pdf.

New York City has spent the past three years reviewing and discussing how city agencies use automated decision systems and artificial intelligence. But it feels like, despite proposed legislation, public hearings, and task force reports, we still have a long way to go towards true transparency. This summer, the Council passed the Public Oversight of Surveillance Technology (POST) Act (Int. 487A), a bill that requires the NYPD to disclose their use of surveillance technologies. The first disclosure by the NYPD will be due in early 2021. This bill is critical to help us understand what technology the NYPD relies on to surveil our clients and communities.

The POST Act is a long overdue reform, but the City still has a long way to go. At present, we do not have access to information regarding how many ADS are used in New York City in the criminal legal system, nor do we know for what purposes they are being implemented. This must change. Earlier this year I testified before this committee about two bills, Intros 1447-2019 (in relation to an annual inventory of agency data) and 1806-2019 (in relation to reporting on automated decision systems used by city agencies). To date, these bills have not passed, and I urge this Committee to also bring these bills to a vote.[2]

The bill on today's agenda is another step forward, but we urge the committee today to expand the bill in scope and substance.

Int. 1894-2020 - A Local Law to amend the New York city charter, in relation to the sale of automated employment decision tools

NYCDS strongly supports Intro 1894-2020, a bill on today's agenda that would regulate the use of automated employment decision tools (AEDTs) in the hiring process. The bill would

1. Require both pre-sale *and* free yearly post-sale "audit(s) for bias";
2. Require employment candidates to be notified within 30 days if an AEDT was used to assess their candidacy, and for what specific purpose; and
3. Impose a penalty for non-compliance

We believe that Int. 1894 will help to protect people from bias by AEDTs. Yet we urge the Council to consider two things.

First, we notice that this bill, which is limited to the employment/hiring context, goes a lot further than the POST Act or the transparency bills considered by this committee in January. We strongly believe that people should be protected from bias when seeking employment, but we also believe that this same level of protection should be extended to people facing the loss of their liberty in the criminal legal system.

The technology covered by the POST Act doesn't cover algorithmic tools created and used by non-police actors in the system, such as risk assessment tools whose outputs are used by judges to make bail determinations, or the DNA software, STRmix, licensed by our crime lab to attempt to interpret complex DNA mixtures in criminal cases. We urge you to consider introducing and passing a bill similar to this one to apply to the criminal legal system.

[2] You can read my complete testimony from January 22, 2020 at https://nycds.org/wp-content/uploads/2020/02/2020.01.22-Tech-Committee-Algorithms-Testimony-FINAL.pdf.

Second, you should consider amending this bill to not only include the language about regular audits for bias, but also to make clear that these AEDTs, even where proprietary, must be subject to under-the-hood examination by independent experts without non-disclosure agreements or other such impediments to a full and fair evaluation.

Flawed algorithmic decision systems can have real-life consequences. For example, STRmix is software that uses algorithmic systems to interpret complex DNA mixture analyses. In 2015, an error in the underlying STRmix code led to problems in 60 criminal cases in Australia.[3] The problem was only discovered in the midst of a criminal trial where prosecutors sought to include its faulty results as evidence.[4] As defense attorneys, we require access to the source code to ensure that STRmix analyses should be relied upon in court. Yet the New Zealand company that created STRmix insists on keeping the source code under wraps due to trade secrets-type interests, so the independent scientific community has not been able to evaluate that software properly.

An audit for bias, while important, is not sufficient to protect against harm. We would ask for this same protection to be included in any legislation extending the types of protections in this bill to algorithms used at various stages in the criminal legal system process.

If you have any questions about my testimony, please contact me at cboyle@nycds.org.

[3] David Murray, "Queensland authorities confirm 'miscode' affects DNA evidence in criminal cases," *The Courier-Mail*, March 20, 2015, available at https://www.couriermail.com.au/news/queensland/queensland-authorities-confirm-miscode-affects-dna-evidence-in-criminal-cases/news-story/833c580d3f1c59039efd1a2ef55af92b.
[4] Lauren Kirchner, "Where Traditional DNA Testing Fails, Algorithms Take Over," *ProPublica*, Nov. 4, 2016, available at https://www.propublica.org/article/where-traditional-dna-testing-fails-algorithms-take-over.

SURVEILLANCE TECHNOLOGY OVERSIGHT PROJECT

@ the Urban Justice Center:
40 Rector Street, 9th Floor
New York, New York 10006
www.S.T.O.P.Spying.org | (646) 602-5600

STATEMENT OF
ALBERT FOX CAHN, ESQ.
EXECUTIVE DIRECTOR
SURVEILLANCE TECHNOLOGY OVERSIGHT PROJECT ("S.T.O.P.")

BEFORE
THE COMMITTEE ON TECHNOLOGY,
NEW YORK CITY COUNCIL

FOR A HEARING ON
THE ETHICAL IMPLICATIONS OF USING ARTIFICIAL INTELLIGENCE AND AUTOMATED DECISION SYSTEMS

PRESENTED
NOVEMBER 13, 2020

1. Introduction

Good morning, my name is Albert Fox Cahn, and I serve as the Executive Director of the Surveillance Technology Oversight Project ("S.T.O.P."). S.T.O.P. advocates and litigates for New Yorkers' privacy, fighting discriminatory surveillance. My thanks to Chair Holden and the committee staff for the opportunity to discuss the dangers that artificial intelligence poses to New York's workers.

2. The Danger of Automated Decision Systems in the Workplace

While few New Yorkers fully understand how Automated Decision Systems ("ADS") impact hiring, promotion, and other employment decisions, thousands of employees are already having their livelihoods decided (at least in part) by these systems. ADS, including Artificial Intelligence and machine learning, is increasingly used by employers big and small. But even as these systems pose a growing threat to workers' rights, there are relatively few protections against errors, bias, and discrimination.

New York City has some of the strongest legal protections against employment discrimination anywhere in the country. But while these laws have been robustly enforced against human discrimination for years, they have yet to be meaningfully applied to many Automated Employment Systems.

Algorithmic discrimination occurs in numerous ways. One example is Amazon's attempt to build ADS recruitment tool in 2014.[1] After training the ADS with the CVs from years of successful applicants, the system learned to simply emulate the hiring discrimination of human employers, ranking female-presenting applicants lower for attending all-women's colleges or listing groups with "women's" in the title.[2] In another example, a widely-used healthcare algorithm drawing from healthcare cost data – in which less money is spent on Black patients' healthcare than white patients' – under-identified Black patients for complex care by more than half.[3]

ADS are sold to the public as "objective" and "scientific", but they are frequently just as biased as human decision makers, if not more so. Only ADS often discriminate opaquely, leaving victims without any legal redress. Even worse, one biased ADS can impact thousands, even millions of employees and job candidates, having a far larger discriminatory impact than any one human employer could.

These concerns are why we agree with many council members that it is urgent for New York to legislate against ADS discrimination. Unfortunately, we believe the current language of Introduction

---

[1] James Vincent, *Amazon reportedly scraps internal AI recruiting tool that was biased against women*, THE VERGE (Oct. 10, 2018), https://www.theverge.com/2018/10/10/17958784/ai-recruiting-tool-bias-amazon-report.
[2] *Id.*
[3] Ziad Obermeyer, Brian Powers, Christine Vogeli, Sendhil Mullainathan, *Dissecting racial bias in an algorithm used to manage the health of populations*, SCIENCE (Oct. 25, 2019), https://science.sciencemag.org/content/366/6464/447

1894 falls far short of this laudable goal, and we fear that if the bill is passed in its current form, it may have dire unintended consequences.

3. Improving Introduction 1894 to Better Protect New York's Workers

We urge council members to address several limitations in the current version of Introduction 1894 that dramatically undercut its stated purpose. In these remarks, I outline numerous high-level responses to the existing text, but I would welcome the opportunity to work with council members to draft revised statutory language.

As part of revisions, we urge the Council to dramatically expand the definition of "automated employment decision tool." The current statutory language is artificially narrow, included only a small subset of the ADS already being marketed to employers. In regulating novel uses of machine learning and other forms of artificial intelligence, we mustn't ignore the harm inflicted by less cutting-edge, and more commonly-used, forms of ADS.[4]

We believe that ADS should include "any software, system, or process that aims to automate, aid, or replace human decision-making relevant to employment. Automated Employment Decision Tools can include both tools that analyze datasets to generate scores, predictions, classifications, or some recommended action(s) that are used by employers to make decisions regarding employees, contractors, and jobs candidates."[5] We have already seen a broad consensus from civil society groups here in New York to adopt such a broad definition in response to last year's ADS Taskforce report.

Similarly, the definition of "employment decision" should also be broadened to include every type of employment decision made by automated employment decision tools. This should include not just hiring decisions, but promotions, scheduling, raises, and more.

The audit process at the heart of this legislation must be better defined to create an enforceable legal standard. Currently, there is no meaningful guidance on how to conduct such an audit with many forms of ADS. Additionally, any audit must be completed by an independent auditing firm, providing ways for workers and other stakeholders to understand how employer ADS operate.

But no matter how audits are conducted, they aren't enough on their own. Notably, Introduction 1894's current language fails to do the most crucial thing needed to prevent use of biased ADS: the bill fails to outlaw such systems. Rather, the bill merely requires such systems to be "audited" for bias. But that audit provides no meaningful protections on its own. A company whose audit reveals biased outcomes could freely sell its product by carrying out this *pro forma* step.

---

[4] *See* Rashida Richardson, ed., "Confronting Black Boxes: A Shadow Report of the New York City Automated Decision System Task Force," P. 20, AI NOW INSTITUTE, December 4, 2019, https://ainowinstitute.org/ads-shadowreport-2019.html.
[5] *Id.*

Similarly, software firms and employers must not be allowed to hide evidence that their software is biased. Rather than just conducting an audit, firms must be required to report their results (good or bad) to the City's Commission on Human Rights or another designated agency. Introduction 1894 should be amended to ban any software tool that has reported evidence of bias (as defined in an updated statute) in the prior year.

This legislation should not only require an audit, but it should attach meaningful penalties to any vendor that sells biased ADS and for any employer that uses such a system. Additionally, liability should jointly and severably apply to both vendors and employer. This will help ensure that victims of automated discrimination are able to recover compensation even when their employer or a software vendor is otherwise judgement proof.[6] Additionally, liability under this ordinance must extend to the City itself. New York City's 325,000 municipal workers must be empowered to bring the same claims as their private sector counterparts.[7]

For those who suffer ADS discrimination, Introduction 1894 puts their rights completely at the whim of enforcement agencies. This simply is not enough to deter misconduct, especially as we see a potential surge in ADS in New Yorkers' workplace hiring[?]. We urge the Council to supplement agency enforcement under this section with a private right to sue employers and vendors who violate this statute. This "force multiplier" will supplement agency actions,[8] but only if this legislation also provides attorneys' fees for a prevailing party.[9] Without attorneys' fees, those most at risk of algorithmic discrimination will be least likely to have their day in court..[10]

Again, we commend the Council for the spirit of Introduction 1894, but we urge you to work with us and other stakeholders to amend the draft. If we fail to pass a revised and strengthened bill, New Yorkers will face increasingly powerful and prevalent ADS without any meaningful legal protections. I look forward to working with the members ensure that Introduction 1894 lives up to the lofty goals that motivated this legislative effort.

---

[6] § 1:25.Joint and several liability, 1 Comparative Negligence Manual § 1:25 (3d ed) ("The joint and several liability doctrine, which applies when more than one defendant tortiously contributed to the plaintiff's injury, allows a nonnegligent plaintiff to recover the full amount of the damages arising from the tortiously caused injury from any one or any combination of the defendants who tortiously contributed to the injury. It has been said that joint and several liability shifts the chore of seeking contribution to the person who perpetrated the harm rather than its innocent recipient.").

[7] Maria Doullis, *The Growth of NYC Employee Headcount*, CBC (May 18, 2020), https://cbcny.org/research/growth-nyc-employee-headcount.

[8] Cameron F. Kerry, John B. Morris, Jr., *In privacy legislation, a private right of action is not an all-or-nothing proposition*, BROOKINGS (Jul. 7, 2020), https://www.brookings.edu/blog/techtank/2020/07/07/in-privacy-legislation-a-private-right-of-action-is-not-an-all-or-nothing-proposition/.

[9] See: 42 U.S.C. § 1988 (2000).

[10] See: § 973 GUARDING AGAINST EXCESSIVE FEES INCURRED AS PART OF FLSA LAWSUITS, 2014 WL 12883902 ("In this sense, the FLSA provision for attorney's fees serves an important public policy goal -- empowering those without the means to finance litigation.").

Attachment A

November 11, 2020

Hon. Laurie A. Cumbo
New York City Council Majority Leader
250 Broadway, Suite 1833
New York, NY 10007
via email

RE: Int. 1894-2020 - Sale of Automated Employment Decision Tools.

Dear Council Member Cumbo:

We, the undersigned civil rights, labor, and civil society organizations commend you for your leadership in tackling the discriminatory threat of automated employment decision tools. We urge the Council to require employers and hiring technology vendors to proactively measure and remediate disparate impacts, and consider less discriminatory alternatives. While we are glad to see this issue getting much needed attention, we are quite concerned that the current language of Int. 1894 could prove counterproductive in the fight against algorithmic discrimination.

We have flagged a number of concerns with the existing language below, and we'd welcome the opportunity to meet with you and your staff to discuss potential changes:

- **Definition of "automated employment decision tool":** Currently, this definition is underinclusive, capturing only a small portion of the technologies and processes that are currently or potentially used in employment settings. We would recommend a more expansive definition that would capture the full range of hiring technologies deployed in New York City, including applicant tracking systems, digital versions of psychological and personality assessments, and other complex procedures that do not fit cleanly within Int. 1894's current scope.
    - One possible formation is: "Automated Employment Decision Tools are any software, system, or process that aims to automate, aid, or replace human decision-making relevant to employment. Automated Employment Decision Tools can include both tools that analyze datasets to generate scores, predictions, classifications, or some recommended action(s) that are used by employers to make decisions regarding employees, contractors, and jobs candidates."[1]
- **Definition of bias audit:** Today, relatively little is publicly known about hiring technology vendors' auditing processes.[2] Existing law and federal agency guidance also do not provide clear and robust standards for reviewing the discriminatory impacts of hiring tools and

[1] See Rashida Richardson, ed., "Confronting Black Boxes: A Shadow Report of the New York City Automated Decision System Task Force," P. 20, AI Now Institute, December 4, 2019, https://ainowinstitute.org/ads-shadowreport-2019.htm

[2] See Manish Raghavan, et al., "Mitigating Bias in Algorithmic Hiring: Evaluating Claims and Practices," https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3408010.

processes.[3] We are concerned that the current language would allow employers and vendors to comply with the law by conducting a pro forma, internal audit, without any meaningful opportunity for third party review. In addition to mandating that annual bias audits be conducted by independent third parties, we recommend that workers be given the opportunity to audit any hiring process for bias. The Committee will need to work, together with a range of stakeholders, to define auditing procedures that include statistical testing, accessibility testing, and proactive consideration of less discriminatory alternatives.

- **Thorough disparate impact audits must involve both vendors and employers:** Compliance with §8-107 cannot be established through a pre-sale audit alone. That law dictates that disparate impacts be measured with respect to the relevant applicant pool or available workforce for a particular job. Such measurement requires data from employers. Similarly, the "business objective" defense turns on its relationship to the particular job and employer.

- **Liability for biased tools:** Currently, no provision of this bill would penalize the sale or use of an Automated Employment Decision Tool that is found to be biased. While such a system may create liability for the vendor and employer under existing New York Human Rights Laws, we urge you to also establish liability here.

- **Definition of "employment decision":** Currently, this definition is underinclusive, capturing only a small subset of the employment decisions that are made by automated employment decision tools.

- **Private right of action:** We fear that even the best possible automated employment decision tool law will be little more than a dead letter in the absence of a private right of action. In addition to the existing civil penalties, we would urge you to include a private right of action for any employee, contractor, or applicant who is subjected to a biased automated employment decision tool.

- **Attorneys' fees:** To ensure that all New Yorkers are able to avail themselves of a private right of action under this law, we would also urge you to provide attorneys' fees for prevailing plaintiffs. This will ensure that low-income employees, contractors, and job applicants will be able to have their day in court.

- **Non-exclusivity:** We urge you to clarify that compliance with Int. 1894 does not preclude a private right of action or agency enforcement action under any other provision of New York City law. In short, compliance with Int. 1894 should be a floor, not a ceiling, for compliance with non-discrimination protections.

- **Reporting:** We urge you to require mandatory reporting to the New York Commission on Human Rights, disclosing the results of any Automated Employment Decision Tool audits.

[3] For example, the Uniform Guidelines on Employee Selection Procedures (UGESP) suggest using a four-fifths impact ratio as a general rule for measuring disparate impact, both the EEOC and OFCCP use additional measures, such as statistical significance tests, when investigating disparate impacts, and courts have refused to adopt a single arithmetic measure of discrimination, acknowledging that the right measurement depends on the context.

The Commission should provide test results to the public to the full extent possible, as well as maintaining a "banned list" of any Automated Employment Decision Tool found to be biased in the prior year.

- **Government hiring:** We urge you to ensure that this legislation applies with full force to any Automated Employment Decision Tool used by New York City agencies. Government hiring must not be held to a lower standard for fairness than what we require for the private sector.

To reiterate, we are grateful for your leadership on this matter, and we hope that we can work with your office to draft language that ensures the spirit of this legislation is fully realized in the years ahead. Unfortunately, these concerns will also make it impossible for us to support passage of Int. 1894 as currently drafted

Sincerely,

AI Now Institute at NYU
BetaNYC
Cryptoharlem
Data for Black Lives
The Legal Aid Society of NYC
NAACP Legal Defense and Educational Fund
National Employment Law Project
New York Civil Liberties Union
New York Communities For Change
OceanHill Brownsville Alliance
S.T.O.P. - The Surveillance Technology Oversight Project
Upturn

CC: Intro. 1894 Co-Sponsors
New York City Council Technology Committee Members

New York City Council
Committee on Technology

Ethical Implications of Using Artificial Intelligence and Automated Decision Systems
November 13, 2020

Written Testimony of
Dr. Sarah Myers West, Postdoctoral Researcher, AI Now Institute.

Good afternoon Chairman Holden and members of the Committee on Technology. My name is Dr. Sarah Myers West and I am a Postdoctoral Researcher at the AI Now Institute, an interdisciplinary research institute at New York University that focuses on the social implications of artificial intelligence. I respectfully submit the following testimony on Int. 1894-2020, the "Fair Shot Act", a law that would regulate the sale of automated employment decision tools.

The City Council's scrutiny of this space is particularly important in a moment where the pandemic has introduced unprecedented challenges for workers in the city. As a growing number of New Yorkers are grappling with the prospect of unemployment, it is critical that employers and technology companies be held accountable for ensuring that workers' rights are protected, both in their hiring practices and in the development, acquisition and use of technologies that mediate their relationships with current and future workers.

The automated employment decision tools that the Fair Shot Act aims to regulate are already in wide use across a range of industries and job categories, from tools used to source potential candidates, to the software used to screen them, such as systems that automatically scan the resumes of applicants, to gamified assessments designed to evaluate personality traits for managerial roles. Even in the absence of clear standards of oversight and evaluation, these systems are already being used to make important decisions throughout the hiring process, from who gets targeted with a job ad, to who may be called in to an interview, to what salary might be offered to a candidate.[1] Job candidates are often unaware when these systems are in use, what qualifications were taken into account in making decisions about hiring, what bias a given system might encode, or why they didn't get the job, and are thus unable to identify or marshal evidence to identify when discrimination takes place - let alone aggregate the data across multiple candidates to necessary to challenge it.

Even more, we lack basic evidence to understand whether and how these tools work, to what extent they are effective in their evaluation, and the risk of errors influencing their outcomes. What we do know, according to research by a team of scholars at Cornell, is that while these companies express a sensitivity to issues of bias and discrimination there is a worrying lack of well-defined best practices as to appropriate methods for debiasing or audits, and with no transparency about the validity of their assessments, assuming a system is assessed at all.[2] Indeed, there is currently no standard for algorithmic

[1] Bogen, M. and Rieke, A. (2018). Help Wanted: An Examination of Hiring Algorithms, Equity, and Bias. *Upturn.* Retrieved from https://www.upturn.org/reports/2018/hiring-algorithms/.
[2] Raghavan, M., Barocas, S., Kleinberg, J. and Levy, K. (2019). Mitigating Bias in Algorithmic Hiring: Evaluating Claims and Practices. ACM Conference on Fairness, Accountability, and Transparency (FAT*), 2020. Retrieved from https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3408010.

bias audits more generally, and the field of research examining these issues is still very young. Other studies indicate significant doubt as to whether these systems work as advertised, and even more concerningly, that they may in fact introduce other, new forms of employment discrimination.[3]

In July, the AI Now Institute joined 23 other civil rights, employment and privacy organizations by signing a set of Civil Rights Principles for Hiring Assessment Technologies. These principles state that hiring assessments:

- Should not discriminate
- Should measure traits and skills that are important to job performance
- Should be transparent to job applicants
- Should be thoroughly and regularly audited
- Should be subject to meaningful oversight by state and federal regulators

For these reasons we are particularly supportive of the City Council's attention to this important issue.

Building on these principles, this testimony makes two primary points: one, that this is a space in urgent need of increased accountability and oversight. Two, bias as it surfaces in these tools cannot be separated out from historic and present-day patterns of employment discrimination - and in fact, research suggests that these tools could introduce new forms of bias. Thus, while in its intent the Fair Shot Act addresses an area sorely in need of close scrutiny, we are concerned that at present this bill could provide a rubber stamp for the perpetuation of discriminatory hiring practices, in a way that makes hiring discrimination much harder to see and could, in the end, could even compound its effects.[4]

Automated Employment Decision Tools Deserve Greater Scrutiny.

First, automated employment decision tools deserve greater scrutiny than they are currently subject to. We know next to nothing about how the models used in automated employment decision tools work and their outcomes, beyond the marketing claims provided by the companies. We have reasons to doubt these claims: one such company, Hirevue, is currently facing charges in front of the Federal Trade Commission of engaging in unfair and deceptive business practices.[5] Opening up these tools to expert review and scrutiny is fundamental to any meaningful program of accountability and oversight.

Given their significant impact, candidates and employers alike deserve independent, third party evaluations of these hiring systems. But as written, the bill doesn't provide workers the information that they need in order to challenge these tools' use. Many job applicants are evaluated by automated employment decision tools without ever knowing that they were used, or if a human even looked at their

[3] Kim, P. (2017). Data-Driven Discrimination At Work. *William & Mary Law Review, 48,* 857-936. Retrieved from https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2801251
[4] Ibid.
[5] Electronic Privacy Information Center. (2019). In Re HireVue. Retrieved from https://epic.org/privacy/ftc/hirevue/EPIC_FTC_HireVue_Complaint.pdf

application. This diminishes workers' rights, leaving them unable to request reasonable accommodations or to make corrections when errors are introduced in their data.

As we highlighted in the 2019 AI Now report, it is important to note that it is employers, not workers, who are the "customers" that these companies seek to court with promises of efficiency and fewer worries about accountability and liability. In fact, several prominent vendors of these tools actively offer to cover any of their customers' legal fees or liabilities that might arise from the use of their products or services.[6] If a goal of this legislation is to combat longstanding patterns of employment discrimination, it is critical that it be designed to serve the needs of those who are discriminated against, rather than the firms responsible for perpetuating discriminatory hiring practices.

Bias Mitigation Efforts Cannot Be Separated from Other Forms of Employment Discrimination

Transparency and disclosure is an important step toward ensuring accountability in the use of automated employment decision tools, but it is only the first of many.[7]

Some of the companies selling these systems say that they conduct audits that demonstrate their validity. But in many cases we lack sufficient information about how these audits work in practice to make an independent assessment of whether and to what extent they are effective in mitigating discriminatory outcomes.[8] The gold standard for auditing would be conducted by an independent third party who can provide an impartial and accountable view into the system.[9] But even still, it's important to note that the field of algorithmic auditing is very young, and there is currently no accepted standard to which such audits adhere, with different approaches often using very different mathematical definitions of "fairness" and "bias," some that may capture the kinds of employment discrimination these tools may perpetuate, and some that may not.

Moreover, in order to be effective, audits need to take into account the actual data used by each employer, and each model, in order to surface potential forms of discrimination. A black box audit of a system without the actual data used both to train a given model, and as input on the employer side, will be able to tell us very little about its real-world activity - this is roughly akin to fumbling in the dark without a flashlight.[10] By design, algorithmic systems mirror patterns identified in large sets of data. In practice, this means that the introduction of automated employment decision tools into a company's hiring practices

---

[6] Crawford, K. et al. (2019). AI Now 2019 Report. Retrieved from https://ainowinstitute.org/AI_Now_2019_Report.pdf.

[7] Civil Rights Principles for Hiring Assessment Technologies. (2020). Retrieved from https://civilrights.org/resource/civil-rights-principles-for-hiring-assessment-technologies/.

[8] Sanchez-Monedero, J., Dencik, L. and Edwards, L. (2020). What does it mean to 'solve' the problem of discrimination in hiring? Social, technical and legal perspectives from the UK on automated hiring systems. ACM Conference on Fairness, Accountability, and Transparency (FAT*), 2020. Retrieved from https://arxiv.org/pdf/1910.06144.pdf; Raghavan, M., Barocas, S., Kleinberg, J. and Levy, K. (2019). Mitigating Bias in Algorithmic Hiring: Evaluating Claims and Practices. ACM Conference on Fairness, Accountability, and Transparency (FAT*), 2020. Retrieved from https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3408010.

[9] Civil Rights Principles for Hiring Assessment Technologies. (2020). Retrieved from https://civilrights.org/resource/civil-rights-principles-for-hiring-assessment-technologies/.

[10] Raji, I.D. and Buolamwini, J. (2019). Actionable Auditing: Investigating the Impact of Publicly Naming Biased Performance Results of Commercial AI Products. Association for the Advancement of Artificial Intelligence. Retrieved from https://dam-prod.media.mit.edu/x/2019/01/24/AIES-19_paper_223.pdf

will be likely to exacerbate, rather than diminish, historical patterns of employment discrimination reflected in the data a given model is trained on, reflecting and amplifying existing conditions of workplace inequality.[11] Pre-sale audits don't enable this kind of assessment, let alone account for the kinds of ongoing tweaks and adjustments that are involved in calibrating many of automated employment decision tools once they are in use in a given firm or institution.

On their own, bias mitigation techniques such as auditing won't be enough to prevent employment discrimination.[12] Bias in these tools cannot be separated from historical and present day employment discrimination. The criteria used to tailor the design of these systems to individual positions often still involves assessments of candidate 'fit', criteria that have long reinforced discriminatory practices.[13] Audits should thus be holistic in nature, considering the use of automated employment decision tools not in isolation, but in relation to employers' overall hiring practices and histories.

Even if we were to take at face value the claims that 'debiasing' processes could mitigate discrimination against protected groups, we know that these biases may be reintroduced through proxy variables.[14] For example, zip code can serve as a proxy for race due to historical patterns of housing discrimination. As the Civil Rights Principles highlight, "Machine learning algorithms can discover subtle correlations and proxies for protected characteristics, even when they are purposefully omitted from the model-building process...Merely removing demographic data from the model-building process will not accomplish this goal."[15]

This is especially important because the research suggests that in fact, the use of automated employment decision tools could even introduce new forms of bias, including on the basis of protected categories like race, gender, age, and ability.[16] For example, some assessments introduce a number of means to discriminate on the basis of accessibility, through interfaces that discriminate against people with

---

[11] Bogen, M. and Rieke, A. (2018). Help Wanted: An Examination of Hiring Algorithms, Equity, and Bias. *Upturn.* Retrieved from https://www.upturn.org/reports/2018/hiring-algorithms/.

[12] Raghavan, M., Barocas, S., Kleinberg, J. and Levy, K. (2019). Mitigating Bias in Algorithmic Hiring: Evaluating Claims and Practices. ACM Conference on Fairness, Accountability, and Transparency (FAT*), 2020. Retrieved from https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3408010.; Corbett-Davies, S. and Goel, S. (2018) The Measure and Mismeasure of Fairness: A Critical Review of Fair Machine Learning. arXiv preprint. Retrieved from https://arxiv.org/abs/1808.00023.

[13] Sanchez-Monedero, J., Dencik, L. and Edwards, L. (2020). What does it mean to 'solve' the problem of discrimination in hiring? Social, technical and legal perspectives from the UK on automated hiring systems. ACM Conference on Fairness, Accountability, and Transparency (FAT*), 2020. Retrieved from https://arxiv.org/pdf/1910.06144.pdf; Raghavan, M., Barocas, S., Kleinberg, J. and Levy, K. (2019). Mitigating Bias in Algorithmic Hiring: Evaluating Claims and Practices. ACM Conference on Fairness, Accountability, and Transparency (FAT*), 2020. Retrieved from https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3408010.

[14] Raghavan, M., Barocas, S., Kleinberg, J. and Levy, K. (2019). Mitigating Bias in Algorithmic Hiring: Evaluating Claims and Practices. ACM Conference on Fairness, Accountability, and Transparency (FAT*), 2020. Retrieved from https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3408010.

[15] Civil Rights Principles for Hiring Assessment Technologies. (2020). Retrieved from https://civilrights.org/resource/civil-rights-principles-for-hiring-assessment-technologies/. See also: Bogen, M. and Rieke, A. (2018). Help Wanted: An Examination of Hiring Algorithms, Equity, and Bias. *Upturn.* Retrieved from https://www.upturn.org/reports/2018/hiring-algorithms/.

[16] Kim, P. (2017). Data-Driven Discrimination At Work. *William & Mary Law Review, 48,* 857-936. Retrieved from https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2801251; Ajunwa, I. (2018). Age Discrimination by Platforms. 40 Berkeley J. Emp. & Lab. L. 1 (2019), Retrieved from: https://ssrn.com/abstract=3142979 https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3142979; Whittaker, M. et al. (2019). Disability, Bias and AI. *AI Now Institute,* Retrieved from https://ainowinstitute.org/disabilitybiasai-2019.pdf.

disabilities such as blindness, deafness or speech disorders, or through the use of algorithmically-mediated personality tests that screen out people with disabilities related to mental health and autism.[17] Facial and voice analysis technologies work less well for people of color, English speakers with non-native accents, and trans people.[18] Automated employment decision tools may also exacerbate the digital divide, by necessitating that job candidates have the ability to connect to high-speed internet in order to complete an assessment, or have access to a smartphone or another expensive device.

In conclusion, greater transparency around the development, data, use, and logics of these automated hiring tools is urgently needed, not only by employers, but by workers and by regulators. Without such transparency, it will be impossible to ensure accountability and effective oversight of a space of growing importance to the hiring landscape. And while audits can be a useful tool, it is important that, at a minimum, they be conducted by independent third parties, incorporate the evaluation of the actual data used by developers and employers, and that they go beyond statistical evaluations to account for how automated employment decision tools are incorporated into hiring practices. Even still, auditing and mandating disclosure of the use of automated employment decision tools will not be enough to combat discrimination in the hiring process, and in some instances, the use of these tools risks making discrimination worse.

Given their prevalence, it is deeply concerning how little we know about whether automated employment decision systems work, let alone what kinds of harms they introduce. In a moment where workers are facing increasing precarity in the wake of the pandemic, the City Council's intervention on their behalf is badly needed. It's likewise critical that regulation of this space be designed to provide the support that workers will need in protecting themselves against employment discrimination.

[17] Fruchterman, J. and Mellea, J. (2018). Expanding Employment Success for People with Disabilities. *Benetech.* Retrieved from https://benetech.org/about/resources/expanding-employment-success-for-people-with-disabilities/

[18] Civil Rights Principles for Hiring Assessment Technologies. (2020). Retrieved from https://civilrights.org/resource/civil-rights-principles-for-hiring-assessment-technologies/.

NYC Council Testimony – Technology Committee
11/13/20

Good afternoon. My name is Andrew Hamilton, President of the Metro New York National Black MBA Association.

I am here to share my support 1894-2020, Sale of automated employment decision tools.

The Black MBA Association is all about uplifting Black professionals and connecting members with jobs they deserve. Our initiatives include career fairs, education opportunities, and mentorship.

Since Covid, we have expanded our efforts by tripling the frequency of our virtual job fairs in New York. Attendance has also tripled with over 500 attending one of our fairs in September. My role is a volunteer position, but I have worked tirelessly with over 200 companies to try and help companies diversify their c-suite workforce during this racial justice reckoning.

Int. 1894 perfectly aligns with our mission by ensuring that all job applicants are treated fairly and not judged by the color of their skin, but by their qualifications.

This bill is particularly essential now, because as has been extensively covered, companies are increasingly evaluating workers using new automated tools. Our members are seeing this all the time as they apply for jobs. While some of these tools are built with fairness in mind, some are not, including those that use facial recognition to screen candidates.[1]

However, facial recognition concerns in hiring date back to well before this outbreak. Specifically, in 2018, then-Senator Kamala Harris, along with Senators Elizabeth Warren and Patty Murray, wrote a letter to the federal Equal Employment Opportunity Commission to express their concerns over the "mounting evidence" of facial analysis leading to biases against vulnerable populations.[2]

The letter stated that facial recognition "amplifies" discrimination and that "such disparities can encode and magnify gender, racial, and other biases that exist in our society…"

To combat this potential problem, the senators requested that the EEOC "develop guidelines for employers on the fair use of facial recognition technologies and how this technology may violate anti-discrimination laws."

---

[1] In the Covid-19 jobs market, biased AI is in charge of all the hiring, Wired, 10/6/20

[2] Letter to EEOC, 9/17/18.

Unsurprisingly, since this prescient letter was sent, Washington has failed to make any real effort or progress in regulating this new hiring space. Perhaps that will change with Ms. Harris now Vice President-elect, but we cannot wait around to find out.

The issue isn't just facial recognition. Look at what happened when Amazon sought to automate resume review, which was already a problematic way of assesses job applicants. When someone reviews a resume, they tend to base someone's qualifications on racial and gender cues such as their name, where they live and where they went to school. Studies have shown that white workers get far more callbacks for interviews than Black workers with identical resumes. Amazon ultimately shut down the program, after realizing it would inevitably exclude women from many opportunities.

The fact that some hiring technologies are harmful, and others are not is why this New York City legislation is so important. Without a bill such as Int. 1894, there is real risk that discrimination against people of color and women will increasingly impact New York workers for the foreseeable future. While many employers have recently made commitments to increasing diversity in their workforces, they do not currently have the information they need to judge whether a hiring tool will help or hinder progress. We need NYC Council members to step up and make the distinction between fair and unfair technology clearer.

What Int. 1894 specifically does is to force vendors of hiring technology to proactively evaluate their products for bias and report on the results. The legislation will ensure that employers who use these products receive annual reports about any potential discrimination occurring in their candidate pipelines. It also mandates that job applicants be given notice about the nature of the technology used in the hiring process. New York City workers deserve such transparency and protections. A worker shouldn't just feel grateful that they have been given an interview. They should trust the process.

The bill comes at an important time for the city, which is already struggling with sky-high unemployment. As jobs slowly but surely return, we need all workers at all types of jobs to be sure that they are being given a fair shot to simply get a foot in the door. An economy built on equity will help the city return even stronger.

21 of your colleagues have joined the bill and we are deeply appreciative of their willingness to tackle this essential, but complex issue. We hope that you join your colleagues and help New York City make history.

On behalf of so many different groups in New York City, I thank you for taking on this complex, but essential issue.

New York City Council Testimony
Arva Rice, New York Urban League

Thank you to Chair Holden and the Technology Committee for holding a hearing on such a forward-thinking bill that will no doubt better protect workers who have been historically discriminated against.

My name is Arva Rice and I am Executive Director of the New York Urban League. Our organization is currently celebrating its centennial, an achievement we are proud of and one we are using to redouble our efforts to bring meaningful improvements to the lives of Black New Yorkers. To mark the occasion, we are releasing a comprehensive report on November 22nd, "The State of Black New York." The report reveals deep racial disparities in the city and also makes specific policy recommendations to guarantee a fairer New York City emerges from the ashes of this pandemic.

Among our core policy proposals is getting behind the passage of Intro. 1894. That's because I see discrimination in hiring all the time. Specifically, the New York Urban League works with too many talented and bright young professionals who never get callbacks. I recently referred one young person, who was bilingual, to a fintech company, but they couldn't even get a returned call. They ultimately landed a job elsewhere, where they were promoted twice in less than a year.

Indeed, far too many companies fail to take Black and brown job applicants seriously. Earlier this fall, Wells Fargo's CEO blamed the lack of diversity in corporate America on the lack of "qualified" Black workers in the talent pipeline. The assumption here is that there is nothing inherently wrong with the processes large companies use to screen job candidates. According to this perspective, employers can set whatever standards they want for prospective applicants, even if those standards perpetuate systemic inequality.

In addition, I recently explained the importance of this legislation in a Blavity opinion piece I co-authored with your colleague, Council member Alicka Ampry-Samuel. As Council member Ampry-Samuel has shared, she would occasionally "whiten" her resume from Alicka to Alicia prior to her career as an elected official to increase her chances of getting a callback for an interview. Her experience is backed by numerous studies. When a person with a white-sounding name submits their resume, they receive 50% more callbacks over someone with a Black-sounding name. Many people of color have internalized this bias against us and it damages our self-esteem and self-worth as a result.

So whether or not all employers are willing to admit it, there is a real problem with how Black workers are evaluated. Beyond unconscious bias creeping into how resumes are traditionally reviewed, referral programs also keep certain workers down. Over one-third of U.S. workers get their current job from a referral, but a Black woman is 35% less likely to get this kind of boost than a white man in a similar position.

So after decades of not seeing real improvement in workplace diversity or pay equity, we need to accelerate our efforts and push our culture to re-evaluate hiring so that all New Yorkers, not just the privileged few, have a fair shot in hearing the magic words - "you're hired."

Hopefully this testimony has made plain that employers need to stop blaming "unqualified" candidates and start questioning their hiring processes. And that's why this bill is so unique - it will guide companies towards a better and more equitable way of choosing job candidates. Even when an employer wants to use a system that prioritizes diversity, there is no useful way for them to compare the options. Marketing claims are never checked and neither employers nor candidates know who to trust. The legislation simply asks vendors of employment evaluation tools to audit their products for adverse impact before selling them in New York City. That way our city can help guarantee that the job candidate who is most qualified, and not just the person who looks the part, will receive the job offer.

To be sure, we need New York City to lead here. At the federal level, the Equal Employment Opportunity Commission has been slowly bled of resources over the past several years and has not kept pace with the myriad ways employers evaluate talent. Instead, we need to continue to bolster the anti-discrimination mission of the city's Human Rights Commission. The agency has done amazing work at leveling the playing field and keeping companies accountable. But we cannot just tackle instances of bias and discrimination after the incident happens. We need actionable transparency and accountability earlier in the process. This bill does exactly that.

The city is right to embrace a growing tech sector. It will be a lynchpin of our recovery. But our public policy should also keep pace with consequential social and economic changes that workers are experiencing. We must seek to regulate technology to ensure it provides real social benefit across the five boroughs.

Thank you for the opportunity to share my support for Int. 1894-2020.

Testimony of Lauren D'Arinzo before New York City Council
Committee on Technology, regarding The Use of Automated Decision Systems in Hiring and Int 1894
November 13th, 2020

Good Afternoon and thank you Chair Holden and the Committee,

My name is Lauren D'Arinzo and I am a Master's Student at New York University where I study Data Science and AI. I am also part of a team at New York University conducting research on using data science responsibly as well as the bias and stability of hiring algorithms. In this testimony, I would like to express my support for Bill 1894 and propose some suggestions to add to its concreteness and improve its intended goals.

First, I would like to highlight the need for the regulation of automated decision systems used in hiring spaces. As a current job seeker and applicant, it is unsettling that a future employer might disregard my application based on the output of an algorithm that has not been rigorously tested for unfair impact or unstable results by an independent third party, and that I might not even be informed of its use.
As a student I have received my fair share of job and internship rejections in my undergraduate and graduate careers -- how many of them were because my output from an automated decision system did not meet the threshold of the ideal output? How are employers even defining these thresholds for a position when many of these tools output is a prediction of personality traits? Why are vendors advertising predictive personality assessments as proxies for qualities of a good employee? How, specifically, are these tools measuring the accuracy of their predictions? These are all questions that policy makers should be asking both the vendors who make these tools and the employers that use them before they are allowed to impact someone's ability to get a job.

What worries me most is that had I not been recruited into a project team explicitly doing research in this space, I would likely not have even known that these types of tools are regularly used by Fortune 500 companies. How many job applicants have important life outcomes, employment decisions, influenced by the output of these tools and aren't even aware of it?

To supplement my support of this bill, I would also like to suggest concrete mechanisms of how to audit these tools, not just with regard to bias as mentioned in the bill but also to stability. Many of these tools claim to predict personality, which behavioral psychology literature will support is something that remains relatively stable over time. As such, if a candidate is assessed by the tool at two different time points, the output should be similar. It is also important for a tool predicting personality traits to be platform agnostic -- that is, if a candidate's Resume was used in the system, it should produce similar results as the output of their LinkedIn profile being used in the system. Otherwise, how can there be confidence that either output is an accurate description of a candidate's personality and their then perceived fitness for a position?

Without regulation, automated decision systems that affect real people's livelihood can have adverse consequences. In summary, I recommend that the City Council adopt a form of Bill 1894, but with stronger language surrounding what it means to audit a tool, with specific detail about both fairness and stability.

Thank you.
Lauren D'Arinzo

Testimony of Kelsey Markey before the New York City Council Committee on Technology regarding Sale of Automated Employment Decision Tools
Int. 1894-2020

November 13, 2020

Good afternoon, my name is Kelsey Markey. I am a Master's student at New York University where I study data science and AI. I am also a part of a research team there investigating bias and stability in hiring tools.

I am very excited to hear that New York City is considering this important bill. As a soon-to-be-graduate, I am frequently subjected to various types of these hiring systems. These tools are used on me and other applicants to screen our resumes, to run background checks, or to analyze our social media profiles. They're also sometimes used as personality tests to see if a candidate will be a "good fit" at a company, and they are even used for video interviews.

One particular example that comes to mind for me is when I applied to an internship this spring. After passing the initial data-science skills test, the company asked me to do an additional assessment to examine my communication skills. This assessment asked me to record myself on video responding to questions and gave me very specific guidance. This guidance included things like to speak naturally and dress as I would for an in-person interview, standard things that would convey my level of professionalism and communication skills. However it also had some more unusual points like how my video should be well-lit with a neutral background, how I should not wear any prints or have any clutter around, and that I should maintain eye contact and smile throughout the video. As a data scientist, these things popped out as me as suggestions that might help a computer to better detect and analyze my actions.

After further investigation I learned that video tools such as these were increasingly common in analyzing potential job candidates. That likely, my facial movements, word choice, and speaking voice was being used to compare me to other applicants and to give me a score based on how "employable" I was. As someone who has struggled with anxiety and nervousness much of my life, I have to admit I was worried. Would a nervous shake in my voice or tremble in my hands mark me as non-employable? What about nonnative speakers with accents or different vocabularies? And was I told not to wear distracting prints because the system was also parsing out my skin color or body shape?

Time and time again, I have seen in my courses, research and work how easy it is to make a biased decision system. Bias can be introduced in many parts of the data pipeline, and often is done without intention or malice. To use my video interview tool as an example, bias could be introduced in the data that was used to build the system, as we might expect to see in the dataset for a workforce that has historically lower numbers of women and people of color. Bias can also be introduced in the features that the system uses for analysis. This video interview tool could have used any of the features that I mentioned earlier like body movements or voice tone, or it could have learned any others from the data that it found to be indicative of employability. Because of this it is important to consider if the features might be introducing bias themselves, for example if they were suggestive of protected classes like gender, race, or disability.

Bias can also be introduced in the validation step through the data or method that is used to determine if the algorithm is "working" as expected. Finally, it can also be introduced in the technical implementation of a system, such as if the tool is applied to all candidates or just some of them.

As part of my data science education I've learned how to assess systems for potential bias throughout the data pipeline. However, when tools such as these are completely unavailable to the public, none of these questions can be answered. Transparency of these hiring systems is essential because it ensures accountability to the public and facilitates audits by experienced computer and data scientists. This country has long said that discrimination is not welcome in hiring, so why are we not also holding these algorithms to the same standards?

My recommendations for this bill are as follows:

1) As suggested, these tools should be subject to an audit for bias, at a minimum of every year, by an impartial outside auditor trained in issues of data ethics and responsibility. I also recommend that these audits have predetermined metrics that they must meet for quantifying what is an acceptable level of bias in the system.
2) The bill suggests that companies should make known "the job qualifications or characteristics for which the tool was used to screen". However, I recommend making clear not just what the tool was looking for, but also which features are being given to the system to determine these qualifications.
3) I recommend that the use of such tools should be disclosed to candidates, and ideally *before* the tool is used. This is important not only for a candidate's own knowledge and preparation, but also to ensure transparency in the often opaque hiring process.
4) And finally, I recommend the creation of a thoughtful mechanism by which members of the public can report suspected uses of these hiring tools.

Thank you very much for your time!

Kelsey Markey
kelseymarkey@gmail.com

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

P 202 371 0910

Writers email: eellman@cdiaonline.org
Writer's direct dial: +1 (202) 408-7407
CDIAONLINE.ORG

November 12, 2020

The Honorable Robert F. Holden
Chair, Committee on Technology
New York City Council
250 Broadway
New York, NY 10007

Dear Chair Holden,

I write on behalf of the Consumer Data Industry Association (CDIA) to oppose Intro 1894.[1] The proposal attempts to regulate technology-driven algorithms in hiring even though they often help reduce bias in hiring.[2] The bill does nothing to address the subjective, unconscious bias that can sometimes exist when mental algorithms in human hiring are used. 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 in employment. Technology in hiring meets the urgent need to hire quickly, hire smartly, and hire remotely. Intro. 1894 stands in the way of tools to reduce hiring bias, and it stands in the way of getting people working again.

1. Federal, state, and local statutes, rules, and guidance protect job applicants

The Fair Credit Reporting Act. When employers use a third-party background check company, called a consumer reporting agencies (CRAs), federal Fair Credit Reporting Act (FCRA) and the New York fair credit reporting law will protect applicants and employees. Among other things, these laws have strict accuracy constraints, require adverse action notices, allow consumer access to the consumer report, and demand rigorous regimes to process disputes. The law imposes obligations on both employment screening companies and employers. When an employer does not use a background check company, these federal and state protections do not apply.

---

[1] 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

[2] The bill would (1) Regulate the use of automated employment decision tools, which, for the purposes of this bill, encompass certain systems that use algorithmic methodologies to filter candidates for hire or to make decisions regarding any other term, condition or privilege of employment; (2) Prohibit the sale of such tools if they were not the subject of an audit for bias in the past year prior to sale, were not sold with a yearly bias audit service at no additional cost, and were not accompanied by a notice that the tool is subject to the provisions of this bill; and (3) Require any person who uses automated employment assessment tools for hiring and other employment purposes to disclose to candidates, within 30 days, when such tools were used to assess their candidacy for employment, and the job qualifications or characteristics for which the tool was used to screen.

Equal Employment Opportunity Commission (EEOC) Guidance. Under federal Guidance from the EEOC, automatic denials for applicants with criminal histories are generally prohibited.

New York City ban-the-box law. New York City has one of the strictest ban-the-box laws in the country, a law that strictly regulates screening for and denial of applicants.

Federal Trade Commission AI Guidance. The Federal Trade Commission (FTC) issued Guidance for users of artificial intelligence and algorithms. The guidance tells users (1) Not to deceive consumers about how they use automated tools; (2) To be transparent when collecting sensitive data; (3) If they make automated decisions based on information from third-party vendors, they may be required to provide the consumer with an "adverse action" notice; (4) If they deny consumers based on algorithmic decision-making they should explain why; (5) If they use algorithms to assign a score to a consumer, they should disclose the key factors that affected the score, in rank order; and (6) If they change the terms of a deal based on an automated tool, they should tell consumers.

2. There are significant operational, technical, and legal problems with Intro. 1894.

A. Augmented technology helps limit bias and protect applicants and employees, but the bill would upend tools to help job applicants

Many people have become more informed about the unconscious bias that can exist in hiring decisions. Technology in hiring reduces unintentional discrimination in employment. Algorithms can remove the unconscious bias that we know exists in some human decisions. Augmented tools can complement human interaction. Unconscious bias does not mean bad or knowing intentions, nor are most hiring managers intentionally discriminating, but the term does mean what the term says -- unconscious bias. To reduce discrimination against a suspect class, employers have embraced objective rules that specify when criminal history will disqualify from proceeding and when they will move an applicant forward. Objective criteria in technology-assisted decisions can prevent excessive weight from being applied to a candidate's criminal history. Any government regulation should seek to preserve augmented intelligence, rather than erect more barriers to hiring.

When an employer uses a background check company, the background checks are regulated by the federal and state fair credit reporting laws. Additional regulation, like that in Intro 1894, will make hiring difficult in a climate where hiring is already exceedingly challenging.

Federal and state background check laws, like the FCRA and the ban-the-box law, and EEO laws impose heavy obligations on background check companies and employers and afford job applicants with substantial rights.

B. Augmented technology helps applicants, employees, and employers move in jobs more quickly, but the bill would upend tools to help job applicants

We are all critically aware that jobs are scarce, and hiring is slow. New York needs to get moving again. During the pandemic, remote hiring is vital to filling jobs, and that may still be true when the pandemic passes. Temporary and gig hiring is on the rise. Jobs are being filled across city, county, and state lines for remote work in New York and elsewhere. To fill jobs quickly, employers need modern tools for a modern economy to validate identity and to verify whether objective criteria are met to get someone working right away. Technology, like algorithms, can quickly fill essential positions and remove barriers to hiring.

Technology can emphasize higher degree offenses and recent offenses for employers to aid their decision making. These tools allow employers to focus on the records that matter the most. Technology-assisted decisions can reduce unconscious bias for information that is not relevant to the job position and enable a quicker, more consistent, and fair hiring process.

Technology adds value to help employers and consumers because augmented intelligence can improve matching logic to properly align the right consumer with the correct record. Technology, which is often grounded in human-supervised machine learning, can help ensure that background checks meet the legal obligations for maximum possible accuracy. Automation can help standardize how criminal charges are reported in a more consistent format despite the variations in court records.

C. The bill contains substantial drafting problems in the bill that are hard to cure

"Augmented employment decision tool." The definition of an "automated employment decision tool" is flawed and needs a closer review. The term means

> any system whose function is governed by statistical theory, or systems whose parameters are defined by such systems, including inferential methodologies, linear regression, neural networks, decision trees, random forests, and other learning algorithms, which automatically filters candidates or prospective candidates for hire or for any term, condition or privilege of employment in a way that establishes a preferred candidate or candidates.

Read logically, human decisions with latent, unconscious bias are not subject to regulation under the introduction, but automated decision tools that are transparent and neutral would be regulated. This deeply flawed application of legislation would do a disservice to employers and job applicants across the city.

"Decision." The definition of a "decision" is flawed. This term "means to screen candidates for employment or otherwise to help to decide compensation or any other terms, conditions, or privileges of employment in the city." Under the FCRA, for 50 years, there has been a commonly accepted definition in federal and state hiring law called an "adverse action." Adverse actions include denial of work or promotions when the decisions are made in whole or in part on information from background check services. "Decisions" is a term already covered in the definitions of "adverse actions" and the legislation overlooks long-standing federal and state definitions.

"Help to decide." The phrase "help to decide" in the definition of an "employment decision" is so broad that it means nothing and everything all at the same time. Augmented decision tools can help remote, and gig workers get back to work faster, especially in a word where personal contact is limited, if not impossible. Technology tools can be used to verify identity, authenticate criteria such as identity documents, current licensing, and to identify if a worker is subject to sanctions or restrictions. These are critical tools to keep essential workers in the field and to replace human interaction where it is not available.

D. The proposal requires serious, unnecessary, and unwanted invasions of consumer privacy

The measure could lead to substantial privacy invasions of job applicants. The introduction requires users of automated decisions to test an automated decision engine against each of the protected classes enumerated in § 8-107 of the Human Rights Law. Such testing would require companies to obtain from applicants (via hiring managers) information about actual or perceived age, race, creed, color, national origin, gender, disability, marital status, partnership status, caregiver status, sexual and reproductive health decisions, sexual orientation, uniformed service or alienage or citizenship status. It is hard to imagine that the government, job applicants, or companies want that information to be shared and stored as contemplated by the proposal.

E. The proposal is against the EEOC Guidance; Compliance is improved with AI

Not only is Intro. 1894 contrary to EEOC Guidance, compliance with EEOC Guidance is improved with the use of augmented intelligence. The EEOC recommends employers against a "one-size-fits all" screening mechanism for all jobs in all situations. To avoid that blanket approach, employers must first conduct a targeted criminal background screen and then engage in an individualized assessment of persons with criminal background records. This initial step - a targeted screen of individuals with criminal records involves – is aided by the use of AI. In conducting this initial step, employers are instructed to look at three factors, referred to as the "*Green* factors" (*Green v. Missouri Pac. Railroad*). The *Green* factors require employers to consider (1) the nature and gravity of the offense or conduct; (2) the time that has passed since the offense or conduct, or the completion of a criminal sentence; and 3) the nature of the job currently held or sought.

3. Conclusion

When used appropriately by employers, a combination of human and technology tools can help root out 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 in employment. Technology in hiring meets the urgent need to hire quickly, hire smartly, and hire remotely. Intro. 1894 stands in the way of tools to reduce hiring bias, and it stands in the way of getting people working again.

I hope that this information is helpful to you. Please let me know if you have any questions or need additional information.

Sincerely,

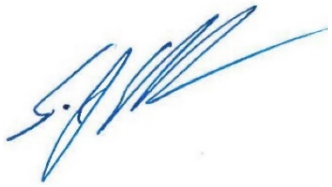

Eric J. Ellman
Senior Vice President, Public Policy & Legal Affairs

cc: Committee on Technology

=

University of Massachusetts Amherst

Center for Employment Equity

Donald Tomaskovic-Devey

Professor of Sociology

tomaskovic-devey@soc.umass.edu

TO: New York City Council

FROM: Donald Tomaskovic-Devey, Director Center for Employment Equity

DATE: 11/5/2020

RE: NYC Bias Audit bill testimony

My testimony today is broadly supportive of the Bias Audit Bill being considered by the NYC Council. I am the Director of the Center for Employment Equity (CEE) at the University of Massachusetts, Amherst. CEE provides scientifically careful analyses and curated data to the community of citizens, employers, and policy makers concerned with promoting equitable workplaces. Personally, I have studied the causes and consequences of gender and racial disparities in employment diversity for four decades. I've submitted an abbreviated resume with my written testimony.

The best organizational research shows that the most effective approach to promoting equal opportunity in employment decisions is to develop appropriate goals and metrics, share them with stakeholders, and embrace accountability for outcomes.[1] When thinking about hiring technologies this implies both demonstrating their connection to the actual work being performed and ensuring that the results of their recommendations are not biased for or against particular demographic groups. It also requires transparency to users – both job seekers and employers – as to the results of these bias audits so that they understand the potential for bias and can choose technologies that minimize, or preferably, eradicate bias in recommendations.

Prior studies of diversity policy efficacy have found that accountability structures lead to clear improvements in the representation of white women, Black women, and Black men in management.[2] If no one is accountable, change is unlikely. Transparency can serve as a powerful foundation for accountability and empower decision makers and employees alike.[3]

When organizations make employment *processes* transparent both managers and job applicants better understand how decisions happen. It is essential that purchasers of hiring

[1] Elizabeth Hirsh and Donald Tomaskovic-Devey. Metrics, Accountability and Transparency: A Simple Recipe to Increase Diversity and Reduce Bias. Kellogg Foundation, What Works Project. Center for Employment Equity, University of Massachusetts. 2020.

[2] Alexandra Kalev, Frank Dobbin, and Erin Kelly, "Best Practices or Best Guesses? Assessing the Efficacy of Corporate Affirmative Action and Diversity Policies," *American Sociological Review* 71, no. 4 (2006): 589–617.

[3] Emilio J. Castilla, "Accounting for the Gap: A Firm Study Manipulating Organizational Accountability and Transparency in Pay Decisions," *Organization Science* 26, no. 2 (2015): 311–333.

tools and technologies understand how these selection devices work, that they have an opportunity to look "under the hood" as needed to understand potential sources of bias with regard to race or gender and most importantly, that they are armed with clear, solid metrics associated with both expected and actual performance on job candidates.

Race and gender bias can also arise based on its association with proxy variables such as social class. For example, if the employee screening technology is based on a firm's current workforce and that firm has tended to hire from Ivy league schools, rather than CUNY, an unsupervised algorithm will tend neglect the good CUNY candidates. This result is not inevitable, as some hiring technologies are designed to discover this type of bias and root it out. However, without proper transparency and auditing focused on outcomes this type of result is likely to occur all too often.

At this point it is clear to many in the AI/machine learning community that algorithms can be biased at the level of design, data input, and algorithmic decision making.[4] But it is also clear that these systems can be audited in terms of bias and then be reengineered to remove the bias from decision making, what computer scientists refer to as equivalent positivity rates.[5] This is not a difficult technical problem, one simply looks at the result of an initial AI implementation and examines its correlation with protected categories such as gender or race and then reengineers the algorithm to remove the correlation with the protected category. There are already firms that do this, but they are in a distinct minority in this industry.[6]

An additional danger with automated employment selection tools is removing managers from the initial decision making process may generate the misleading impression that the opportunity for bias in human decision making has been eliminated. This idea of 'fairness by blindness' sets a potentially false sense of security that can actually lead managers to assume, rather than promote, bias free workplaces.[7] Organizations are better served by making the *outcome* of hiring technologies transparent to stakeholders so that both individuals within the organization (i.e., job applicants, employers) and external audiences (e.g., regulators, the public) can assess the demographic consequences and disparities of automated hiring tools.

I support this bill becoming law because it introduces clear metrics and transparency and empowers decision makers to reduce bias in employment decisions. I also think it is clear that this industry has already demonstrated both that it is unlikely to police itself and audit practices are already available and implemented by some firms. Importantly, New York employers will not bear any cost in implementing these auditing requirements, but rather those costs will be born by vendors in this already competitive industry.

---

[4] Barocas, Solon, and Andrew D. Selbst. "Big data's disparate impact." *California Law Review.* 104 (2016): 671; Kleinberg, Jon, Sendhil Mullainathan, and Manish Raghavan. "Inherent trade-offs in the fair determination of risk scores." *arXiv preprint arXiv:1609.05807* (2016).

[5] Hardt, Moritz, Eric Price, and Nati Srebro. "Equality of opportunity in supervised learning." In *Advances in neural information processing systems*, pp. 3315-3323. 2016.

[6] For example Pymetrics, https://www.pymetrics.ai/

[7] Castilla, Emilio J., and Stephen Benard. "The paradox of meritocracy in organizations." *Administrative Science Quarterly* 55, no. 4 (2010): 543-676.

Donald Tomaskovic-Devey

a. Professional Preparation

B.A. 1979 Sociology, Fordham University, Bronx, New York, Summa Cum Laude
Ph.D. 1984 Sociology, Boston University, Boston, MA.

b. Appointments

2017- present Director, Center for Employment Equity, University of Massachusetts, Amherst
2005-present Professor of Sociology, University of Massachusetts, Amherst
1999-2016 Intergovernmental Personnel Act research appointment to the U.S. Equal Employment Opportunity Commission (various years)
1984-2005 Assistant to Full Professor, Sociology, North Carolina State University

c. Synergistic Activities

Founder of EEODataNet a network of scholars using EEOC data, including annual meeting at EEOC headquarters in collaboration with EEOC researchers and litigators, 2013-2016.
Co-Founder (2017) Center for Employment Equity, University of Massachusetts, which focuses on the development of scientific research, policy application, and data visualizations utilizing EEOC employment and charge data.
Founder and convener (2016) of Comparative Organizational Inequality Network, which includes 30+ scientists from fifteen countries and five disciplines analyzing linked employer-employee administrative data.

d. Recent Funded Research

Donald Tomaskovic-Devey (PI), "Workplace Earnings Polarization: A Four-Country Panel Analysis" National Science Foundation, July 2019-June 2021 ($277,272).
Jasmine Kerrisey (PI) and Don Tomaskovic-Devey (co-PI) and Steve Boutcher (co-PI), Public Sector Precarity: The Joint Role of Institutional and Organizational Processes. National Science Foundation, July 2019-June 2021 ($245,269).
Fidan Kurtulus (PI) and Don Tomaskovic-Devey (co-PI), Estimating the Effects of Randomized OFCCP Audits on Workplace Female and Minority Composition: A Random Control Trial. Arnold Foundation January 2019- December 2020 ($149,903).
Donald Tomaskovic-Devey (PI), "Moving the Needle on Racial Employment Equity. Annie E. Casey Foundation, October 2019-October 2020 ($71,545).
Don Tomaskovic-Devey, Lee Badgett, Fidan Kutulus. EEO Research and Data Center. W. K. Kellogg Foundation. January 2017-December 2019 ($656,588)
Don Tomaskovic-Devey, Andrew Penner, Dustin Avent-Holt. The Organizational Production of Earnings Inequalities, National Science Foundation, September 2015-August 2018 ($97,000 UMass portion)
Don Tomaskovic-Devey (PI) and Lee Badgett and Fidan Kurtulis (Co-PIs). Building an Interdisciplinary Equal Employment Opportunity Research Network and Data Capacity, September 2013-September 2016. National Science Foundation ($255,820).

e. Monographs

Donald Tomaskovic-Devey and Dustin Avent-Holt. 2019. *Relational Inequalities: An Organizational Approach* NY: Oxford University Press.
Kevin Stainback and Donald Tomaskovic-Devey. 2012. *Documenting Desegregation: Racial and Gender Segregation in Private Sector Employment since the Civil Rights Act.* NY: Russell Sage Foundation. *[Analyzes EEO-1 data from 1966-2005]*

Donald Tomaskovic-Devey, 1993. *Gender and Racial Inequality at Work: The Sources and Consequences of Job Segregation*. Ithaca, NY: ILR Press.

f. Peer Reviewed Papers Using EEOC Data

Safi Shams and Donald Tomaskovic-Devey, 2019. "Racial and Gender Trends and Trajectories in Access to Managerial Jobs." *Social Science Research*. doi:10.1016/j.ssresearch.2018.12.020

Fidan Kurtulus and Donald Tomaskovic-Devey, 2012**. "**Do Female Top Managers Help Women to Advance? A Panel Study Using EEO-1 Records" *Annals of the American Academy of Political and Social Science*. 639: 173:197.

Kevin Stainback and Donald Tomaskovic-Devey, 2009. "Intersections of Power and Privilege: Long-Term Trends in Managerial Representation." *American Sociological Review*. 74:800-820.

McTague, Tricia, Kevin Stainback, and Donald Tomaskovic-Devey, 2009. "Organizational Response to Institutional Pressures for Equal Employment Opportunity since the Civil Rights Act of 1964." *Social Forces*. 87:1499-1527.

Yucel, Recai.M., Hong Ding, Ali K. Uludag, and Donald Tomaskovic-Devey, 2008. Multiple imputation in multiple classification and multiple-membership structures. In *Proceedings of the Section on Bayesian Statistical Science of the American Statistical Association*.

Donald Tomaskovic-Devey and Kevin Stainback. 2007. "Discrimination and Desegregation: Equal Opportunity Progress in U.S. Private Sector Workplaces Since the Civil Rights Act. *The Annals of the American Academy of Political and Social Science*. 609:49-84.

Donald Tomaskovic-Devey, Catherine Zimmer, Kevin Stainback, Corre Robinson, Tiffany Taylor, and Tricia McTague, 2006. "Documenting Desegregation: Segregation in American Workplaces by Race, Ethnicity, and Sex 1966-2000." *American Sociological Review*. 71:565-588.

Kevin Stainback, Corre Robinson, Donald Tomaskovic-Devey, 2005. "Race and Workplace Integration: A Politically Mediated Process?" *American Behavioral Scientist*. 48:1200-1229.

Corre L. Robinson, Tiffany Taylor, Donald Tomaskovic-Devey, Cathy Zimmer, and Matthew Irvin, 2005 "Studying Race and Sex Segregation at the Establishment-Level: Methodological Concerns and Substantive Opportunities in the Use of EEO-1 Data." *Work & Occupations*. 32:5-38.

g. Other Peer Reviewed Employment Diversity and Discrimination Papers

Andreja Poje, Aleksandra Kanujou-Mrčela, and Donald Tomaskovic-Devey. 2019. Equal Pay for Equal Work or Work of Equal Value in Practice: The Case of the Professions of Medical Nurse/Medical Technician, Police Officer, and University Professor." *Teorija in Praska*. 56,1:134-161.

Laurel Smith-Doerr, Sharla Alegria, Kaye Husbands Fealing, Debra Fitzpatrick, and Donald Tomaskovic-Devey. 2019. "The Value of Women's Work in Science Policy: Occupational Sex Segregation in US Science Agencies and Effects on Pay, 1994-2008." *American Journal of Sociology* 125: 534-576. Sociology 125, no. 2 (September 2019): 534-576. https://doi.org/10.1086/705514

Melzer, Silvia Maja, Donald Tomaskovic-Devey, Reinhard Schunck, Peter Jacobebbinghaus. 2018. "A Relational Inequality Approach to First- and Second-Generation Immigrant Earnings in German Workplaces." *Social Forces*. 97: 91–128.

Anja-Kristin Abendroth, Silvia Maja Melzer, Alexandra Kalev, Donald Tomaskovic-Devey. 2017. “Women at Work: Women’s Access to Power and the Gender Earning Gap.” *Industrial and Labor Relations Review*. *70*: 190-222.

Donald Tomaskovic-Devey, Martin Hällsten, and Dustin Avent-Holt. 2015. “Where do Immigrants Fare Worse? Modeling Workplace Wage Gap Variation with Longitudinal Employer-Employee Data.” *American Journal of Sociology*. 120:1095-1143.

Dustin Avent-Holt and Donald Tomaskovic-Devey. 2012. **“**Relational Inequality: Gender Earnings Inequality in US and Japanese Manufacturing Plants in the Early 1980s.” *Social Forces*. 91: 157-180.

Dustin Avent-Holt and Donald Tomaskovic-Devey. 2010. “The Relational Basis of Inequality: Generic and Contingent Wage Distribution Processes.” *Work and Occupations*. 37:162-19.

Donald Tomaskovic-Devey, Dustin Avent-Holt, Catherine Zimmer and Sandra Harding, 2009. “The Categorical Generation of Organizational Inequality: A Comparative Test of Tilly’s Durable Inequality.” *Research in Social Stratification and Mobility*. 27:128-142.

Donald Tomaskovic-Devey and Sheryl Skaggs, "Workplace Gender and Racial Composition and Productivity: An Establishment Level Test of the Statistical Discrimination Hypothesis." *Work & Occupations*. 26:422-445. 1999.

h. Examples of Honors, Awards, Appointments

German Institute for Economic Research, SEOP Scientific Advisory Boars

Samuel F. Conti Faculty Fellowship, University of Massachusetts, September 2016-August 2017.

Anneliese Maier Research *Award*, Alexander von Humboldt Foundation, January 2014-December 2018.

Visiting Scientist, Copenhagen Business School, University of Lujbljana, Sciences Po, Stockholm University, Queensland University of Technology, Utrecht University, Bielefeld University.

BETA NYC

We want to thank Director Thamkittikasem for reaching out early on to talk about AMPO. We were excited about the Executive Order, but it is disappointing that the city says that it is supportive of digital rights, transparency, and accountability and yet hasn't posted anything publicly on its website.

The Mayor's Office of Operations hasn't updated the AMPO page nor the News page since September 2019.

Transparency builds trust. This was one of our biggest concerns with the previous Taskforce and it is disappointing that we're one year into AMPO and we still don't have the type of transparency we asked for last year.

If the city is sharing its insights with international partners, could this Administration share these insights with Council and the public?

We have signed on to the collective letter with AI Now, Data for Black Lives, The Legal Aid Society of NYC, NAACP Legal Defense and Educational Fund National Employment Law Project, New York Civil Liberties Union, New York Communities For Change, & S.T.O.P. – The Surveillance Technology Oversight Project Upturn.

We think this bill is notable but as several shortcomings
* First, there needs to be private right of action.
* This will ensure that low-income employees, contractors, and job applicants will be able to have their day in court.
* Government hiring must not be held to a lower standard for fairness than what we require for the private sector.

Majority Leader Cumbo, when you're looking to review technology, there is a great book named "data feminism" By Catherine D'Ignazio and Lauren F. Klein. A new way of thinking about data science and data ethics that is informed by the ideas of intersectional feminism.
https://mitpress.mit.edu/books/data-feminism