



April 3, 2024

Assemblymember Rebecca Bauer-Kahan
1021 O Street, Suite 5210
Sacramento, CA 95814

Re: AB 2930 (Bauer-Kahan) - Support, if Amended

Dear Assemblymember Bauer-Kahan,

We, the undersigned worker, consumer, and civil rights organizations, represent thousands of workers and consumers across California. We appreciate your leadership on the critical issue of algorithmic bias and have a “support, if amended” position on AB 2930.

California has the opportunity to lead the nation with innovative policy that places common-sense guardrails on the development and use of AI and automated decision-making tools. Legislating on emerging technology requires thoughtful and deliberative work, and engagement with all the stakeholders who will be impacted. We commend you for taking on this challenge.

We would like to work with you to ensure this bill provides needed protections for the people of California. In its current form, however, we are concerned that it would not provide necessary transparency measures, that some of the definitions would allow companies to side-step the critical

accountability mechanisms this bill creates, and that consumers and job candidates would not receive sufficient notice or explanation of consequential decisions made with the help of automated decision tools, among other concerns. Furthermore, when the bill was reintroduced in 2024 as AB 2930, enforcement provisions were weakened and overbroad protections for “trade secrets” were added.

The most important changes we think should be made include:

1. Revising the bill’s definitions (and its “controlling factor” requirement, in particular) to prevent employers from avoiding compliance with the bill simply by having a human rubber stamp algorithmic recommendations;
2. Excluding worker management and other issues related to ongoing employment because of the diverse sets of laws that apply, the unique power dynamics, and the wide range of harms that go beyond discrimination that can occur when AI is deployed in the workplace;
3. Strengthening the notice and explanation provisions so that workers and consumers have a meaningful understanding of how and why they are assessed by AI systems;
4. Strengthening the impact assessment provisions by requiring unbiased, independent auditors and that all forms discrimination must be assessed and mitigated;
5. Ensuring the bill has adequate enforcement remedies and does not undermine existing laws.

In the attached redline and explanation document, you will find the changes we think are necessary to ensure this bill adequately protects individuals, job candidates and consumers in particular, without impeding innovation. We have also provided an explanation of significant changes, including why each change is important.

We look forward to your response, and to partnering with you on this important issue.

Sincerely,

American Federation of Musicians, Local 7
California Employment Lawyers Association
California Nurses Association
Center for Democracy and Technology
Center on Race and Digital Justice
Consumer Reports
East Bay Community Law Center
Economic Security California Action

Equal Rights Advocates
Greenlining Institute
Legal Aid at Work
Oakland Privacy
RISE Economy
TechEquity Action
SEIU California

AB2930—Summary of Proposed Amendments

Definitions/Scope

1. Eliminate the requirement that a covered ADT be “specifically developed and marketed to, or specifically modified to, make, or be a controlling factor in making,” a consequential decision (Sec. 22756(c)).
2. Revise the definition of covered employment decisions to focus on recruitment, hiring, and other pre-employment decisions (Sec. 22756(d)(1)).
3. Update definition of “artificial intelligence” to match updated OECD definition (22756(b)).
4. Include additional impacts in education, criminal justice, and legal services in the definition of “consequential decision” (22756(d)).

Impact Assessment

1. Require impact assessments to include analyses of all forms of discrimination and provide more detail on mitigating discrimination risks (Sec. 22756.1(a)-(b)).
2. Add a requirement that impact assessments be conducted by an independent auditor--that is, one not connected to the deployer or developer that is free of financial and personal conflicts of interest, and that is allowed to conduct the audit without being subjected to developer- or deployer-imposed restrictions.
3. Require companies to provide auditors with all information relevant to an impact assessment and require auditors to maintain full records of their audits for later review by the enforcement agency.
4. Explicitly require impact assessments to be completed prior to deployment for any tools released after Jan. 1, 2025 (Sec. 22756.1).
5. Require deployer impact assessments to drill down into the validity (in the social scientific sense) and reliability of the ADT by examining whether the attributes an ADT uses are relevant to the decision it is tasked with making/assisting (Sec. 22756.1(a)(2)).
6. Require meaningful stakeholder consultation as part of the impact assessment process (Sec. 22756.1(a)).

Notice, Explanation, Transparency, and Opt-Out Rights

1. Expand the notice requirements so that consumers and workers receive meaningful notice of how a tool works as well as an explanation for the decision or recommendation the tool makes (Sec. 22756.2(a)(2)).
2. Include the right to a specific, accurate, and actionable explanation (Sec. - 22756.2).
3. Require notices/explanations to be prepared and presented in a manner that ensures effective disclosure (Sec. 22756.2).

4. Strengthen the right of individuals to opt out of automated decision-making (Sec. 22756.2(b)).
5. Add recordkeeping requirements.

Enforcement and Special Privileges

1. Restore the private right of action from AB 331.
2. Eliminate the right to cure for public attorney enforcement (Sec. 22756.8(b)(1)).
3. Restore AB 331's obligation for developers and deployers to file impact assessments with the Civil Rights Department.
4. Allow additional public attorneys to bring an action for violations of AB 2930 and restore the Civil Rights Department's authority to share impact assessments with other public agencies when appropriate (Sec. 22756.8).
5. Ensure that AB 2930 does not undermine stronger protections that may be available through other laws by clarifying that AB 2930 does not diminish other rights (Sec. 22756.2(b), Sections 22756.1, 22756.2, and 22756.4).
6. Eliminate unique trade secret protections (Sec. 22756.3, Sec. 22756.7(b)).

AB2930—Explanation of Proposed Amendments

Definitions/Scope

- 1. Eliminate the requirement that a covered ADT be "specifically developed and marketed to, or specifically modified to, make, or be a controlling factor in making," a consequential decision (Sec. 22756(c)).**

The “controlling factor” requirement in the current definition of “automated decision tool” would have the practical effect of giving companies the ability to completely avoid complying with the bill. It would be easy for developers to evade compliance by including a disclaimer in their marketing materials stating that a tool “is not designed to be a controlling factor in any decision,” and for deployers to avoid compliance by having a human rubber-stamp algorithmic recommendations. Indeed, a recent study by researchers at Cornell University, Consumer Reports, and Data & Society found that few companies are conducting audits or making disclosures required by New York City's AI hiring law, which similarly is limited to tools that have a dominant effect on the decision-making process.¹ Additionally, the “controlling factor” requirement is inconsistent with existing civil rights protections (including in employment, education, and housing), which apply not only to decisions based solely on unlawful discrimination, but to any decision where impermissible factors influenced the outcome. We recommend that the term “controlling factor” be replaced with the term “contributing factor” in the definition of ADT.

- 2. Revise the definition of covered employment decisions to focus on recruitment, hiring, and other pre-employment decisions (Sec. 22756(d)(1)).**

Covering worker management and other issues relating to ongoing employment is a complex topic involving diverse sets of laws. Automated systems in the workplace result in a wide range of harms that go beyond discrimination. The workplace also involves unique power dynamics that are not present in ordinary consumer settings. Rules for AI used in ongoing employment thus should be developed separately and in coordination with labor unions and other workers' rights experts. At the same time, the current language of the bill is ambiguous regarding whether it covers uses of AI in increasingly prominent pre-employment practices like targeted advertising or identifying LinkedIn profiles of “passive” candidates. Such recruitment and sourcing techniques are already impacting many workers’ access to economic opportunities. The definition of consequential decisions should explicitly cover these practices.

- 3. Update definition of “artificial intelligence” to match current OECD definition (22756(b)).**

¹ See Consumer Reports: New Research: NYC Algorithmic Transparency Law is Falling Short of Its Goals
<https://innovation.consumerreports.org/new-research-nyc-algorithmic-transparency-law-is-falling-short-of-its-goals>

To ensure consistency with other developing standards for AI, the definition should be updated so that it matches the recently revised OECD definition: “‘Artificial intelligence’ means a machine-based system that can, for explicit or implicit objectives, infer, from the input it receives, how to generate outputs such as predictions, recommendations, or decisions that can influence real or virtual environments.”

4. Include additional impacts in the definition of “consequential decision” (22756(d)).

The definition of “consequential decision” should more explicitly cover additional areas in education, criminal justice, and legal services.

Impact Assessment

1. Require impact assessments to include analyses of all forms of discrimination and provide more detail on mitigating discrimination risks (Sec. 22756.1(a)-(b)).

The current text would not require employers to check for forms of discrimination other than disparate impact. Disparate impact analysis does not capture many forms of unlawful discrimination, such as disparate treatment, the use of proxy traits that correlate with protected characteristics, or barriers to accessibility that may violate the rights of disabled consumers and workers. Additionally, the current provisions regarding mitigating identified discrimination risks (22756.1(a)(6) and (b)(5)) are vague, and do not explicitly require companies to mitigate identified risks. This is especially important for individuals with disabilities who may need accommodations.

2. Add a requirement that impact assessments be conducted by an independent auditor--that is, one not connected to the deployer or developer that is free of financial and personal conflicts of interest, and that is allowed to conduct the audit without being subjected to developer- or deployer-imposed restrictions.

Existing research and academic reviews of auditing companies have underscored the importance of third-party auditors (and how they are distinct from first or second party auditors) due to the potential pitfalls and lack of accountability inherent in first and second party auditing systems.²³ Third-party auditors who have full access to AI systems and are free of conflicts of interest are more likely to analyze and publish truthful assessments. We recommend amendments that will help prevent the evolution of a “revolving door” where auditors have recent relationships with the audited entity that may affect the impartiality of the impact assessment.

²See AI NOW: Algorithmic Accountability: Moving Beyond Audits

<https://ainowinstitute.org/publication/algorithmic-accountability#audit-economy>

³Costanza-Chock, S., Raji, I. D., & Buolamwini, J. (2022). Who Audits the Auditors? Recommendations from a field scan of the algorithmic auditing ecosystem. 2022 ACM Conference on Fairness, Accountability, and Transparency.

<https://doi.org/10.1145/3531146.353321>

3. Require companies to provide auditors with all information relevant to an impact assessment and require auditors to maintain full records of their audits for later review by the enforcement agency.

Similar to independent auditor requirements—transparency requirements will help ensure that all relevant information is available and considered during the course of impact assessments.

4. Explicitly require impact assessments to be completed prior to deployment for any tools released after Jan. 1, 2025 (Sec. 22756.1).

Under the current language, an employer could conceivably start using a tool on (e.g.) January 3, 2025 and not complete an impact assessment on it until January 2, 2026.

5. Require deployer impact assessments to drill down into the validity (in the social scientific sense) and reliability of the ADT by examining whether the attributes an ADT uses are relevant to the decision it is tasked with making/assisting (Sec. 22756.1(a)(2)).

A company should have to ensure that an ADT validly and reliably measures the characteristics or attributes of interest before automating a decision-making process that impacts consumers' or workers' lives. Planning for validation and reliability testing should be part of both the developer and deployer impact assessments, as is already required under current anti-discrimination laws.⁴

6. Require meaningful stakeholder consultation as part of the impact assessment process (Sec. 22756.1(a)).

Frameworks for ethical AI frequently recommend stakeholder engagement as part of the development, assessment, and deployment process for AI.⁵ This is particularly crucial when AI is used to make important decisions such as those that AB 2930 would cover.

Notice, Explanation, Transparency, and Opt-Out Rights

1. Expand the notice requirements so that consumers and workers receive meaningful notice of how a tool works as well as an explanation for the decision or recommendation the tool makes (Sec. 22756.2(a)(2)).

⁴ See UNIFORM GUIDELINES ON EMPLOYEE SELECTION PROCEDURES, <https://www.govinfo.gov/content/pkg/CFR-2018-title29-vol4/xml/CFR-2018-title29-vol4-part1607.xml>

⁵ See, e.g., Alun Preece, et al., Stakeholders in Explainable AI (2018), <https://arxiv.org/pdf/1810.00184.pdf>. See also, e.g., NIST AI Risk Management Framework 1.0 at 9 (Jan. 2023) (“Identifying and managing AI risks and potential impacts – both positive and negative – requires a broad set of perspectives and actors across the AI lifecycle”); IEEE’s Ethically Aligned Design: A Vision for Prioritizing Human Well-being with Autonomous and Intelligent Systems, at 65 (2019), https://standards.ieee.org/wp-content/uploads/import/documents/other/ead_v2.pdf.

The current text of AB 2930 vaguely requires deployers to provide impacted individuals with a pre-evaluation notice that includes "a plain language description of" the ADT and includes no post-assessment notice or explanation requirements. This would not provide consumers or workers with the information needed to determine whether a tool has violated or may violate their legal rights.

2. Include the right to a specific, accurate, and actionable explanation (Sec. - 22756.2).

An algorithmic tool that is used to make consequential decisions about consumers should be able to produce specific and accurate explanations of why it generates the outputs it generates, such that consumers understand the output, and understand if it is based on inaccurate information or inferences about them. The right to an explanation is included in numerous standards and recommendations regarding governance of automated decisions, from government, industry standards groups, and civil society published in the last five years.⁶

3. Require notices/explanations to be prepared and presented in a manner that ensures effective disclosure (Sec. 22756.2).

The current text does not provide adequate detail on what constitutes effective notice under sec. 22756.2. Notice and explanations should be in languages and formats that clearly and effectively convey the required information.

4. Strengthen the right of individuals to opt out of automated decision-making (Sec. 22756.2(b)).

The current text of 22756.2(b) provides an opt-out right for decisions made "solely" through an ADT. It would be easy for a deployer to avoid this provision completely simply by designating a human to rubber-stamp automated decisions.

5. Add recordkeeping requirements.

Absent such recordkeeping obligations, regulators and individuals who are subjected to discrimination may find themselves unable to obtain adequate discovery or meet their burden of proof even in cases where clear discrimination occurred.

Enforcement and Special Privileges

1. Restore the private right of action from AB 331.

⁶ See the OSTP's Blueprint for an AI Bill of Rights, NIST's Four Principles of Explainable Artificial Intelligence, IEEE's Ethically Aligned Design: A Vision for Prioritizing Human Well-being with Autonomous and Intelligent Systems, Civil Rights Principles for Hiring Assessment Technologies, and Civil Rights Standards for 21st Century Employment Selection Procedures.

The rights provided in this bill will only be meaningful if the individuals who are intended to be protected by this bill have a way to enforce them. The public attorneys given enforcement authority under this bill simply do not have the resources or capacity to take on all of the investigations and enforcement actions that would be necessary to properly enforce this bill. The private right of action should be reinserted into the bill and should be consistent with existing frameworks for plaintiffs bringing discrimination claims.

2. Eliminate the right to cure for public attorney enforcement (Sec. 22756.8(b)(1)).

Recent experience in California illustrates this: the right to cure in the California Consumer Privacy Act (CCPA) significantly delayed and hampered enforcement — for the first four years the law was in effect, there was only one enforcement action. Fortunately, the CCPA's right to cure was removed earlier this year, in recognition of the perverse incentives it created.

3. Restore AB 331's obligation for developers and deployers to file impact assessments with the Civil Rights Department.

This is an important accountability step that would ensure the state agency with authority to enforce most of the state's antidiscrimination laws gets information regarding ADTs that may impact millions of California residents in a timely manner.

4. Allow additional public attorneys to bring an action for violations of AB 2930 and restore the Civil Rights Department's authority to share impact assessments with other public agencies when appropriate (Sec. 22756.8).

Given the incredibly broad scope of AB 2930, which covers an enormous range of decisions on many aspects of Californians' lives, additional public attorneys with appropriate subject matter or regional expertise (e.g., the Civil Rights Department, Department of Education, etc.) may be better-positioned to prosecute certain actions or classes of actions. The CA AG should be able to designate those public attorneys to help with enforcement.

5. Ensure that AB 2930 does not undermine stronger protections that may be available through other laws by clarifying that AB 2930 does not diminish other rights (Sec. 22756.2(b), Sections 22756.1, 22756.2, and 22756.4).

A savings clause should be added to the bill to ensure that this bill does not diminish existing antidiscrimination laws, e.g., the Fair Employment and Housing Act, the Unruh Act, etc. Additionally, it should be clear that compliance with the impact assessment, notice, and governance requirements does not create a defense to other unlawful conduct by deployers and developers.

6. Eliminate unique trade secret protections (Sec. 22756.3, Sec. 22756.7(b)).

Given the likelihood that developers will attempt to designate AI and other automated systems as proprietary trade secrets, this creates an unacceptable risk that developers will withhold information

that deployers will need to meet their impact assessments and disclosure obligations. Developers can protect their legitimate intellectual property interest through the contractual and technical measures that technology companies have used for years when selling and distributing novel software, systems, or other technologies. Additionally, the special carve-outs under the California Public Records Act are unnecessary. The Public Records Act already contains provisions protecting sensitive and privileged information from disclosure. See, e.g., Cal. Gov't Code §§ 7927.700-705. Developers and deployers of AI systems can rely on those protections and do not need special protections that other businesses do not enjoy.