

Date of Hearing: April 16, 2024

ASSEMBLY COMMITTEE ON PRIVACY AND CONSUMER PROTECTION

Rebecca Bauer-Kahan, Chair

AB 2930 (Bauer-Kahan) – As Introduced February 15, 2024

AS PROPOSED TO BE AMENDED

**SUBJECT:** Automated decision tools

**SYNOPSIS**

*This author-sponsored measure is a significant, meaningful attempt to set state policy with respect to private and public use of automated decision tools (ADTs): systems and services that employ artificial intelligence, and make, or are a controlling factor in making, consequential decisions. Consequential decisions, in turn, are those that materially impact the necessities of a person's life, such as housing, education, employment, and health care (including reproductive health care), as well as the person's involvement with critical public functions such as voting and the criminal justice system.*

*The bill would establish requirements for informed, conscious deployment of ADTs, based on the White House's recently-released Blueprint for an AI Bill of Rights. Broadly speaking, the bill only prescribes one policy outcome: it forbids algorithmic discrimination against individuals on the basis of characteristics protected under state civil rights law, such as race, age, national origin, disability, limited English proficiency, and sex (including pregnancy, childbirth, and related conditions; gender identity; intersex status; and sexual orientation). Much of the bill is instead devoted to ensuring that ADT developers and deployers thoroughly consider the design, testing, and functioning of ADTs before releasing them into the world. This appears to be sound policy, as developers and deployers are the entities that possess the most knowledge about the systems they work with.*

*This bill is supported by the Microsoft Corporation and Workday. It is opposed by a variety of business trade associations, including the California Chamber of Commerce. A broad coalition of organizations, including Tech Equity, Oakland Privacy, and SEIU California, takes a support-if-amended position. Author's amendments respond to some of their requested amendments and make other clarifying changes.*

*If passed by this Committee, this bill will next be heard by the Assembly Judiciary Committee.*

**SUMMARY:** Establishes a statutory framework to further the safe and informed use of ADTs in California. Specifically, **this bill:**

- 1) Requires that developers and deployers of ADTs perform impact assessments on or before January 1, 2026, and annually thereafter. Exempts deployers with fewer than 25 employees unless they deploy ADTs that impact more than 999 people.
- 2) Requires that impact assessments include:
  - a) A statement of the purpose of the ADT and its intended benefits and uses.

- b) A description of the ADT's outputs and how they are used to make consequential decisions, or be a controlling factor in making those decisions.
  - c) A summary of types of data collected from natural persons and processed by the ADT.
  - d) A statement of how a deployer's use of an ADT differs from the ADT purpose as stated by the developer.
  - e) An analysis of potential adverse impacts on the basis of sex, race, color, ethnicity, religion, age, national origin, limited English proficiency, disability, veteran status, or genetic information resulting from a deployer's use of the ADT.
  - f) A description of safeguards implemented by the deployer to address any reasonably foreseeable risks of algorithmic discrimination arising from the use of the ADT.
  - g) A description of how the ADT will be used by a natural person, or monitored when it is used, to make consequential decisions.
- 3) Requires developers and deployers to perform impact assessments if an ADT they develop or deploy receives a significant update.
  - 4) Requires that deployers notify individuals that are subject to a consequential decision made by an ADT, and that they provide subject individuals with relevant information.
  - 5) Requires deployers to accommodate an individual's request to not be subject to an ADT's decision, if technically feasible, and requires the deployer to provide an alternative process. Permits the deployer to capture relevant information from the individual upon accommodating their request.
  - 6) Requires a developer to provide a deployer with a statement regarding the intended uses of an ADT they provide, along with documentation relating to the ADT's known limitations, training data, and how the tool was evaluated for validity and explainability. Does not require disclosure of trade secrets.
  - 7) Requires developers and deployers to establish governance programs, maintained by at least one employee, that assess and address risks of algorithmic discrimination arising from the use of ADTs.
  - 8) Requires developers and deployers to maintain impact assessments performed on ADTs for at least two years.
  - 9) Requires developers and deployers to make publicly available documentation related to the types of ADTs currently used or made available by the entity, as well as how the entity manages algorithmic discrimination risks arising from the use of these ADTs.
  - 10) Prohibits deployers from using discriminatory ADTs.
  - 11) Permits the Civil Rights Department (CRD) to investigate reports of algorithmic discrimination. Requires developers and deployers to provide impact assessments to the CRD upon request.

- 12) Permits the CRD to bring an administrative enforcement action not exceeding \$10,000 per violation against any developer or deployer who violates the requirement to disclose impact assessments.
- 13) Permits the Attorney General, district attorneys, a county counsel, city attorneys, and city prosecutors to bring a civil action against a developer or deployer for a violation of the bill. Permits a court to award injunctive relief, declaratory relief, reasonable attorney's fees and litigation costs, and, in cases where a violation involves algorithmic discrimination, a civil penalty of \$25,000 per violation.
- 14) Requires a developer or deployer to be provided with 45 days written notice before an action is commenced against them, during which they can cure the noticed violation.
- 15) Exempts cybersecurity-related technologies from its provisions.
- 16) Defines various terms, including:
  - a) Algorithmic discrimination: the condition in which an ADT contributes to unjustified differential treatment or impacts disfavoring people based on their actual or perceived race, color, ethnicity, sex, religion, age, national origin, limited English proficiency, disability, veteran status, genetic information, reproductive health, or any other classification protected by state law.
  - b) Artificial intelligence: an engineered or machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing a real or virtual environment.
  - c) ADT: a system or service that uses artificial intelligence and has been specifically developed and marketed to, make, or be a controlling factor in making, consequential decisions.
  - d) Consequential decision: a decision or judgment that has a legal, material, or similarly significant effect on an individual's life relating to access to government benefits or services, assignments of penalties by government, or the impact of, or the cost, terms, or availability of several listed categories.
  - e) Deployer: a person, partnership, state or local government agency, or corporation that uses an ADT to make a consequential decision.
  - f) Developer: a person, partnership, state or local government agency, or corporation that designs, codes, or produces an ADT, or substantially modifies an artificial intelligence system or service for the intended purpose of making, or being a controlling factor in making, consequential decisions, whether for its own use or for use by a third party.
  - g) Impact assessment: a documented risk-based evaluation of an ADT.
  - h) Significant update: a new version, new release, or other update to an ADT that materially changes its principal use, principal intended use, or outcome.

**EXISTING LAW:**

- 1) Establishes the Civil Rights Department, and sets forth its statutory functions, duties, and powers. (Gov. Code § 12930.)
- 2) Establishes the Fair Employment and Housing Act. (Gov. Code § 12900 *et seq.*)
- 3) Establishes the Unruh Civil Rights Act. (Civil Code § 51.)
- 4) Defines “trade secret” under the Uniform Trade Secrets Act as information, including a formula, pattern, compilation, program, device, method, technique, or process, that both:
  - a) Derives independent economic value, actual or potential, from not being generally known to the public or to other persons who can obtain economic value from its disclosure or use; and
  - b) Is the subject of efforts that are reasonable under the circumstances to maintain its secrecy. (Civ. Code § 3426.1(d).)
- 5) Establishes the California Privacy Protection Agency (Privacy Agency) and vests it with full administrative power, authority, and jurisdiction to implement and enforce the California Consumer Privacy Act (CCPA) of 2018. (Civ. Code § 1798.199.10.)
- 6) Requires the Attorney General to issue regulations governing access and opt-out rights with respect to businesses’ use of automated decisionmaking technology. Passes along this authority to the California Privacy Protection Agency after a set period of time. (Civ. Code § 1798.185.)

**FISCAL EFFECT:** As currently in print this bill is keyed fiscal.

**COMMENTS:**

1) **Background.** Under the bill, the term “automated decision tool,” or “ADT,” refers to computer systems that are relied upon to make, or to be a controlling factor in making, consequential decisions. ADTs are used to decide whom to hire, what treatments patients should receive, whether to grant parole, and many other life-altering decisions that used to be made entirely by human beings.

Using an ADT is, in many cases, a good thing. Computers can generally draw upon far more digitally-encoded data, and process this data far more quickly, than any human being. A recent *New York Times* article described the development of a breast cancer screening system that uses artificial intelligence:

From the millions of cases the system is fed, the technology creates a mathematical representation of normal mammograms and those with cancers. With the ability to look at each image in a more granular way than the human eye, it then compares that baseline to find abnormalities in each mammogram.

Last year, after a test on more than 275,000 breast cancer cases, Kheiron [Medical Technologies, the system’s developer] reported that its A.I. software matched the performance of human radiologists when acting as the second reader of mammography scans. It also cut down on radiologists’ workloads by at least 30 percent because it reduced the number of X-rays they needed to read. In other results from a Hungarian clinic last year,

the technology increased the cancer detection rate by 13 percent because more malignancies were identified.<sup>1</sup>

However, ADTs are far from flawless. Care should be taken to avoid what New York University Professor Meredith Broussard describes as “the notion that computers are more ‘objective’ or ‘unbiased’ because they distill questions and answers down to mathematical evaluation...an unwavering faith that if the world used more computers, and used them properly, social problems would disappear and we’d create a digitally enabled utopia.”<sup>2</sup> Many aspects of the real world are not reducible to data, and many complex situations present competing demands that human beings remain best equipped to assess.

ADT usage continues to increase. Global consulting firm McKinsey & Company reports that the adoption of artificial intelligence among the organizations it surveyed more than doubled between 2017 and 2022, with 50 percent now reporting they use the technology.<sup>3</sup>

*ADTs use artificial intelligence.* ADTs use artificial intelligence (AI) to make consequential decisions. AI uses algorithms – sets of rules – to transform inputs into outputs. Inputs and outputs can be anything a computer can process: numbers, text, audio, video, or movement. This is because AI is not fundamentally different from other computer functions. Its novelty lies in its application: unlike standard computer functions, AI is able to accomplish tasks that are normally performed by humans.

*Bias and discrimination can result from poor training data.* There is a famous saying in computer science: “garbage in, garbage out.” The performance of an ADT is directly impacted by the quality, quantity, and relevance of the data used to train it.<sup>4</sup> If the data used to train the ADT is biased, the tool’s outputs will be similarly biased. Over the past thirty years, several industries have been forced to contend with this fact as they have attempted to introduce ADTs into their workflows. Specific examples follow:

*Hiring and recruitment tools.* It is no secret that people of various races, genders, and cultures are not distributed equally throughout the workforce. An ADT that is trained on historical data to make hiring decisions will be predisposed to maintain the ratios it is trained on; as described by Aditya Malik, the Founder and CEO of Valuematrix.ai:

Generative AI, for all its grandeur, has the potential to perpetuate latent biases inherited from human creators. A disconcerting echo of historical prejudices may inadvertently seep into the algorithms. Imagine a scenario where previous senior managers, driven by biases of gender, age, faith or race, rejected candidates for misguided reasons. The AI, if not vigilantly curated,

---

<sup>1</sup> Santariano & Metz, “Using A.I. to Detect Breast Cancer That Doctors Miss,” *New York Times*, Mar. 5, 2023, <https://www.nytimes.com/2023/03/05/technology/artificial-intelligence-breast-cancer-detection.html>.

<sup>2</sup> Broussard, “Artificial Unintelligence,” *MIT Press*, 2019.

<sup>3</sup> McKinsey & Co., “The state of AI in 2022—and a half decade in review,” Dec. 6, 2022, <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-in-2022-and-a-half-decade-in-review>.

<sup>4</sup> Rohit Sehgal, “AI Needs Data More Than Data Needs AI”, *Forbes*, Oct. 5, 2023, <https://www.forbes.com/sites/forbestechcouncil/2023/10/05/ai-needs-data-more-than-data-needs-ai/>.

might misconstrue these patterns as indicators of incompetence, thus exacerbating the exclusion of qualified candidates from underrepresented backgrounds.<sup>5</sup>

This was notoriously experienced by Amazon, who considered automating their hiring practices in the early 2010s. They opted against this approach in 2015 when they realized that their ADT-enabled system was not rating candidates in a gender-neutral way. In fact, their system was excluding women from the pool of acceptable candidates because it had been trained to vet applicants by observing patterns in resumes submitted to the company over a 10-year period. Most came from men, a reflection of inequities across the tech industry.<sup>6</sup>

*Sentencing and bail decisions.* ADTs are frequently used to inform sentencing and bail decisions. These tools are trained using historical data, and the predictions they make can therefore reflect historical bias. A 2016 Propublica study dove into the use of one such tool – COMPAS – in Broward County, Florida.<sup>7</sup> The study determined that Black defendants were far more likely than white defendants to be incorrectly judged to be at a higher risk of recidivism, while white defendants were more likely than Black defendants to be incorrectly flagged as low risk.

The for-profit company that developed this tool, Northpointe, does not publicly disclose the calculations used to arrive at defendants’ risk scores, so it is not possible for either defendants or the public to see what might be driving the disparity. These discrepancies mirror historical injustices perpetuated against Black Americans by California’s criminal justice system.<sup>8</sup>

*Healthcare.* When ADTs are deployed in healthcare, biased historical data can lead to patients being recommended substandard care on the basis of their race or ethnicity. In 2007, an ADT was developed to help doctors estimate whether it was safe for people who had delivered previous children through cesarean section to deliver subsequent children vaginally – a risky procedure. The ADT considered various health relevant factors as it made its decision, such as the woman’s age, her reason for the previous cesarean, and how long ago the cesarean had been performed. However, a 2017 study found that the ADT was biased; it predicted Black and Latino people were less likely to have a successful vaginal birth after a cesarean than non-Hispanic white women. As a result, doctors performed more cesareans on Black and Latino people than on white people.<sup>9</sup> These discrepancies perpetuate historical biases – Black Americans, for example, have historically received a lower standard of healthcare than their white counterparts.<sup>10</sup>

---

<sup>5</sup> Aditya Malik, “AI Bias In Recruitment: Ethical Implications And Transparency,” *Forbes*, Sep. 25, 2023, <https://www.forbes.com/sites/forbestechcouncil/2023/09/25/ai-bias-in-recruitment-ethical-implications-and-transparency/>.

<sup>6</sup> Jeffrey Dastin, “Amazon scraps secret AI recruiting tool that showed bias against women,” *Reuters*, Oct. 9, 2018, <https://www.reuters.com/article/amazoncom-jobs-automation/insight-amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSL2N1VB1FQ/>.

<sup>7</sup> Julia Angwin, Jeff Larson, Surya Mattu and Lauren Kirchner, “Machine Bias”, *ProPublica*, May 23, 2016, <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>.

<sup>8</sup> California Task Force to Study and Develop Reparation Proposals for African American, “Final Report,” California.

<sup>9</sup> Caleb J Colón-Rodríguez, “Shedding Light on Healthcare Algorithmic and Artificial Intelligence Bias,” *U.S. Department of Health & Human Services Office of Minority Health*, Jul. 12, 2023, [minorityhealth.hhs.gov/news/shedding-light-healthcare-algorithmic-and-artificial-intelligence-bias](https://minorityhealth.hhs.gov/news/shedding-light-healthcare-algorithmic-and-artificial-intelligence-bias).

<sup>10</sup> California Task Force to Study and Develop Reparation Proposals for African American, “Final Report,” California.

*Credit and loan approval.* Financial tools that utilize ADTs are similarly susceptible to bias and discrimination. An investigation by The Markup (and co-published by the Associated Press) revealed that in 2019, lenders were more likely to deny home loans to people of color; in particular, lenders were 40 percent more likely to turn down Latino applicants for loans, 50 percent more likely to deny Asian/Pacific Islander applicants, and 70 percent more likely to deny Native American applicants than similar white applicants. Lenders were 80 percent more likely to reject Black applicants than similar white applicants. In every case, the prospective borrowers of color looked almost exactly the same on paper as the white applicants, except for their race.<sup>11</sup>

The premises of this bill are, first, that Californians ought to be much more informed about the ADTs being used to make consequential decisions in their lives, and second, that developers and deployers of ADTs ought to be much more conscious and deliberate about what the impacts of their tools might be. In short, California ought to get a handle on ADTs while it still can before they become yet another phenomenon— like social media, greenhouse gases, guns, and urban sprawl— that could have been dealt with deliberately and intelligently in the past, but was instead allowed to expand without any constraints past the point of being manageable.

2) **Author’s statement.** According to the author:

AB 2930 protects individuals from algorithmic discrimination by requiring developers and users to assess automated decision tools (ADTs) that make consequential decisions and mitigate any discovered biases. The use of ADTs have become very prominent within different sectors such as housing, employment, and even in criminal justice sentencing and probation decisions. The algorithms used within ADTs can be prone to unrepresentative datasets, faulty classifications, and flawed design, which can lead to biased, discriminatory, or unfair outcomes. These tools can exacerbate the harms they are intended to address and ultimately hurt the people they are supposed to help. As the use of decision making via algorithm becomes more prevalent in our daily lives, it is crucial that we take the necessary steps to ensure that they are used ethically and responsibly.

3) **What this bill would do.** The provisions of this bill are derived from the principles outlined in the *Blueprint for an AI Bill of Rights*, issued by the White House Office of Science and Technology Policy on October 4, 2022.<sup>12</sup> The *Blueprint* identifies five principles that should “guide the design, use, and deployment of automated systems to protect the American public in the age of artificial intelligence.” The applicable principles, which this measure generally seeks to implement, are as follows:

*Safe and Effective Systems.* The *Blueprint* summarizes this principle as meaning that “you should be protected from unsafe or ineffective systems.” The bill would implement this principle by requiring developers and deployers to complete impact assessments for any ADT that they develop or use. These impact assessments require thoughtful consideration of what ADTs are intended to do, what their effects might be, and how they are validated. Once completed, these impact assessments can be requested by the Civil Rights Department. To shield small businesses from the costs associated with this bill, deployers with fewer than 25 employees are exempted

---

<sup>11</sup> Emmanuel Martinez and Lauren Kirchner, “The Secret Bias Hidden in Mortgage-Approval Algorithms,” *Markup*, Aug. 25, 2021, [themarkup.org/denied/2021/08/25/the-secret-bias-hidden-in-mortgage-approval-algorithms](https://themarkup.org/denied/2021/08/25/the-secret-bias-hidden-in-mortgage-approval-algorithms).

<sup>12</sup> The White House, “Blueprint for an AI Bill of Rights,” Oct. 2022, <https://www.whitehouse.gov/ostp/ai-bill-of-rights/>.

from this requirement, unless an ADT used by the deployer has impacted more than 999 people per year. Developers and deployers would have to complete new impact assessments any time a significant update, as defined, is made to an ADT.

*Algorithmic Discrimination Protections.* The *Blueprint* summarizes this principle as meaning that “you should not face discrimination by algorithms[,] and systems should be used and designed in an equitable way.” The bill would implement this principle in at least three ways:

- First, developers and deployers would be required to prepare governance programs containing reasonable administrative and technical safeguards to map, measure, manage, and govern the reasonably foreseeable risks of algorithmic discrimination associated with the use of an ADT.
- Second, at least one employee would have to be designated to oversee and maintain the governance program and compliance with this bill’s requirements. That employee would have the authority to assert to their employer a good faith belief that the design, production, or use of an ADT fails to comply with one or more requirements of the bill. The employer would then be required to conduct a prompt and complete assessment of the compliance issues raised.
- Third, the bill would flatly prohibit a deployer from using an ADT that contributes to algorithmic discrimination on the basis of any category protected under California’s civil rights laws.

*Notice and Explanation.* The *Blueprint* summarizes this principle as meaning that “you should know that an automated system is being used and understand how and why it contributes to outcomes that impact you.” The bill would implement this principle by requiring a deployer to notify any natural person that is the subject of a consequential decision that an ADT will be used to make, or will be a controlling factor in making, that decision. The notification must be made at or before the time the ADT is used.

*Human Alternatives, Consideration, and Fallback.* The *Blueprint* summarizes this principle as meaning that “[y]ou should be able to opt out, where appropriate, and have access to a person who can quickly consider and remedy problems you encounter.” The bill would implement this principle by requiring a deployer to accommodate a natural person’s request to not be subject to an ADT and to instead be subject to an alternative selection process or accommodation. The bill limits this requirement in two significant ways, however. First, the alternative is only available in cases in which a consequential decision is to be made solely based on the output of an ADT; i.e., one over which there is no human discretion. Second, the alternative must be technically feasible.

Ultimately, AI-powered ADTs are increasingly influencing the lives of Californians. These tools are being used to assess eligibility for loans and jobs, and influencing outcomes in healthcare and criminal justice. These tools can, when poorly designed and improperly used, undermine privacy and equity. Without a robust legal framework for their use, ADTs risk perpetuating systemic biases through opaque decision-making and a general lack of accountability. This bill seeks to tackle these risks by enforcing transparency and ensuring that decision-making tools are scrutinized and pressure-tested for discriminatory outcomes. Such an approach safeguards fundamental rights and fosters trust in the technologies that are reshaping society.



4) **Comparisons to AB 331.** This bill derives from AB 331, introduced in 2023 by the same author. AB 331 was passed out of two Assembly policy committees – Privacy and Consumer Protection, and Judiciary – but was ultimately held in the Assembly Appropriations Committee. Key differences between the bills include:

- First, AB 331 granted individual Californians the right to bring a civil action against a deployer who violated the bill’s provisions. The current bill does not provide for a private right of action, and gives public prosecutors civil enforcement authority.
- Second, AB 331 required developers and deployers to provide impact assessments to the Civil Rights Department within 60 days of their completion. By comparison, the current bill only requires developers and deployers to provide impact assessments to the Civil Rights Department upon request, within 7 days of the request being made.
- Third, AB 331 granted the Civil Rights Department the ability to share impact assessments with other state entities as appropriate. The current bill does not allow for sharing of impact assessments in this way. As a result, it is unclear how public prosecutors seeking civil liability are meant to support their claims of algorithmic discrimination.

#### 5) **Stakeholder positions.**

In support, Microsoft Corporation writes:

AB 2930 aligns with the Microsoft Responsible AI framework and recognizes the promise and potential pitfalls that the next generation of AI may bring through the use of Automated Decision Tools (ADTs). The bill requires developers and users of ADTs to conduct and record impact assessments to evaluate the intended uses of ADTs including an analysis of potential adverse impact on the basis of race, color, ethnicity, sex, religion, age, national origin, or any other classification protected by state law. We believe AB 2930 establishes appropriate legal guardrails on the development and deployment of ADT systems through the use of impact assessments and design evaluations.

History teaches us that transformative technologies like AI require new rules of the road. Proactive, self-regulatory efforts by responsible companies will help pave the way for these new laws, but we know that not all organizations will adopt responsible practices voluntarily. In our view, effective AI regulations should center on the highest risk applications and be outcomes-focused and durable in the face of rapidly advancing technologies and changing societal expectations. To spread the benefits of AI as broadly as possible, regulatory approaches around the globe will need to be interoperable and adaptive, just like AI itself. We believe AB 2930 moves California in the right direction to ensure that AI is built and used responsibly and ethically.

A coalition of opponents from business trade associations sets forth what it sees as the key problems with the bill:

- **Scoping problems.** The bill applies to every industry imaginable and is not sufficiently limited to high-risk uses of ADT within those industries. Both the scope and various key

terms (in particular consequential decisions, automated decision tools, algorithmic discrimination, and controlling factor), all require additional clarity and/or narrowing.

- Unworkable obligations in relation to the mandated impact assessments and notice requirements. In particular, the “opt-out” related requirements, including the requirement that businesses allow consumers to opt out of an ADT when “technically feasible”.
- Overly narrow confidentiality protections. Impact assessments must be exempt from the California Public Records Act to encourage candor and avoid concerns around assessments becoming fodder for litigation.
- Allowing multiple enforcement entities will invariably create confusion for compliance. Providing for a single enforcer (the Attorney General) will promote consistent interpretation and application across the state.
- Insufficient preemption protections and related concerns around regulatory activity. Preemption is needed both to prevent conflict with localities and, similarly, to prevent concerns around departments/agencies over-regulating this technology and getting ahead of the Legislature and Governor. These issues are too important to Californians across the state and our struggling economy to significantly delegate and defer to unelected officials.

Finally, a coalition of organizations, including Tech Equity, Oakland Privacy, and SEIU California, has adopted a support-if-amended position, seeking the following changes:

1. Revising the bill’s definitions (and its “controlling factor” requirement, in particular) to prevent deployers 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 of discrimination must be assessed and mitigated;
5. Ensuring the bill has adequate enforcement remedies and does not undermine existing laws.

As described below, author’s amendments respond to 1), 3), and 5).

6) **Author’s amendments.** In response to feedback from various stakeholders, the author has offered several amendments to clarify and strengthen the provisions of this bill. In broad strokes, the amendments would do the following:

- Clarify the categories of personal information in impact assessments in order to align with California Consumer Privacy Act:

- (3) A summary of the ~~type of data~~ *categories of information* collected from natural persons and processed by the automated decision tool when it is used to make, or be a ~~controlling~~ *substantial* factor in making, a consequential ~~decision~~ *decision, including, but not limited to, all of the following:*
- (A) *Each category of personal information identified by reference to the applicable subparagraph enumerated under paragraph (1) of subdivision (v) of Section 1798.140 of the Civil Code.*
  - (B) *Each category of sensitive personal information identified by reference to the applicable paragraph and subparagraph enumerated under subdivision (ae) of Section 1798.140 of the Civil Code.*
  - (C) *Each category of information related to a natural person's receipt of sensitive services, as defined in Section 56.05 of the Civil Code, identified by reference to the specific category of sensitive service enumerated in the definition.*
- Require a deployer to provide natural persons with sufficient information to better enable them to take advantage of the alternative selection process or accommodation, in lieu of being subject to a decision made by an ADT:
    - (D) *Information sufficient to enable the natural person to request to be subject to an alternative selection process or accommodation, as applicable, in lieu of the automated decision tool, as provided in subdivision (b).*
  - Require developers to provide deployers with a description of the deployers' responsibilities under the bill, whenever an ADT is transferred from a developer to a deployer. This will foster compliance with the bill's requirements:
    - (4) *A description of the deployer's responsibilities under this chapter.*
  - Require impact assessments to be maintained for five years, rather than two years. This will create a more robust evidentiary record for any possible enforcement actions under the bill:
    - (6) *Maintain for ~~two~~ five years after completion the results of an impact assessment.*
  - Permit the Civil Rights Department to provide an impact assessment it receives to a public prosecutor for the purposes of assisting that prosecutor in initiating or litigating a civil action:
    - (c) *The Civil Rights Department may provide an impact assessment it receives to a public prosecutor listed in subdivision (a) of Section 22756.8 to assist that public prosecutor in initiating or litigating a civil action under that section*
  - Grant the Civil Rights Department the authority to bring a civil action against a deployer or developer. This change fosters a more complete enforcement scheme:
    - 22756.8. (a) (1) Any of the following public ~~attorneys~~ *entities* may bring a civil action against a deployer or developer for a violation of this chapter:

(A) The Attorney General in the name of the people of the State of California.

(B) A district attorney, county counsel, or city attorney for the jurisdiction in which the violation occurred.

(C) A city prosecutor in any city having a full-time city prosecutor, with the consent of the district attorney.

(D) *The Civil Rights Department.*

- Clarify that the protections offered under the bill are not intended to undermine or overwrite existing rights and remedies. Although nothing in the bill in print suggests otherwise, this has been a point of confusion for many stakeholders:

*22756.11. The rights, remedies, and penalties established by this chapter are cumulative and shall not be construed to supersede the rights, remedies, or penalties established under other laws including, but not limited to, Chapter 6 (commencing with Section 12940) of Part 2.8 of Division 3 of Title 2 of the Government Code and Section 51 of the Civil Code.*

- Replace the phrase “controlling factor” with “substantial factor” throughout the bill for purposes of determining when an artificial intelligence system rises to the level of an ADT subject to the bill. A number of stakeholders have expressed concerns that “controlling” sets too high of a bar, excluding circumstances in which an ADT may have influenced but not decided the outcome. On the other hand, the author does not want to capture circumstances in which an algorithm plays a minor role that is incapable of altering the outcome of the decision:

(c) “Automated decision tool” means a system or service that uses artificial intelligence and has been specifically developed ~~and marketed~~ to, or specifically modified to, make, or be a ~~controlling~~ *substantial* factor in making, consequential decisions.

- Clarify that developers are prohibited from making available discriminatory ADTs. Previously, while developers were required to run impact assessments, there was no clear prohibition on the release of discriminatory ADTs:

*(b) A developer shall not make available to potential deployers an automated decision tool that results in algorithmic discrimination.*

- Make various other small changes to improve clarity.

7) **Additional considerations.** Going forward, there are two key issues that the author may wish to address.

*Application of administrative penalties.* While this bill empowers the Civil Rights Department to bring administrative enforcement actions against developers or deployers that violate certain provisions, committee staff have been informed that the Civil Rights Department does not currently possess the necessary staff or resources to conduct administrative enforcement

hearings. At a minimum, this would likely require new positions for administrative law judges to preside over the hearings as well as an initial investment to adopt and implement procedures for the presentation of evidence, arguments, and testimony related to an alleged violation. Going forward, the author may wish to examine whether the bill should expressly address this issue.

*Relationship to the CCPA and CPPA.* Proposition 24 (2020) created the California Privacy Protection Agency (Privacy Agency) in California, vested with full administrative power, authority, and jurisdiction to implement and enforce the CCPA and the CPRA. The Agency’s responsibilities include updating existing regulations, and adopting new regulations.

In March 2020, the Privacy Agency issued draft regulations governing the use of automated decision-making technology (ADMT) by businesses that are subject to the CCPA.<sup>13</sup> These regulations would require businesses that use ADMTs for “significant decisions” to: conduct and submit risk assessments to the Privacy Agency; provide consumers pre-use notice, the right to opt out, and certain information about the ADMT; and prohibit discrimination based on protected classes for businesses that use biometric data for significant decisions. Thus, it appears that there are significant overlaps between the proposed regulations and some provisions of this bill with respect to risk assessments, pre-use notice, and the right to opt out of an automated decision-making process. Going forward, the author may wish to examine ways of aligning these two parallel efforts at regulating automated decisionmaking processes.

8) **Related legislation.** AB 2885 (Bauer-Kahan, 2024) would establish a uniform definition for “artificial intelligence” in California’s codes.

AB 331 (Bauer-Kahan, 2023) was substantially similar to this bill. The bill was held in Assembly Appropriations.

AB 302 (Ward, Chap. 800, Stats. 2023) required the California Department of Technology (CDT) to conduct an inventory of all high-risk automated decision systems being used in state agencies. The bill was chaptered on October 8<sup>th</sup>, 2023.

SB 313 (Dodd, 2023) would have established an Office of Artificial Intelligence within CDT, with “the powers and authorities necessary to guide the design, use, or deployment of automated systems by a state agency to ensure that all AI systems are designed and deployed in a manner that is consistent with state and federal laws and regulations regarding privacy and civil liberties and that minimizes bias and promotes equitable outcomes for all Californians.” The bill was held in Senate Appropriations.

AB 2269 (Chau, 2020) would have established the Automated Decision Systems (ADS) Accountability Act of 2020. The Act would have required businesses in California that provide a person, as defined, with a program or device that uses an automated decision system (ADS) to take affirmative steps to ensure that there are processes in place to continually test for biases during the development and usage of the ADS, conduct an ADS impact assessment on its program or device to determine whether the ADS has a disproportionate adverse impact on a protected class, as specified, examine if the ADS in question serves reasonable objectives and furthers a legitimate interest, and compare the ADS to alternatives or reasonable modifications

---

<sup>13</sup> California Privacy Protection Agency, “Draft Risk Assessment and Automated Decisionmaking Technology Regulations,” Mar. 2024, [https://cppa.ca.gov/meetings/materials/20240308\\_item4\\_draft\\_risk.pdf](https://cppa.ca.gov/meetings/materials/20240308_item4_draft_risk.pdf).

that may be taken to limit adverse consequences on protected classes. The bill was held in Assembly Privacy and Consumer Protection.

***ARGUMENTS IN SUPPORT:***

Microsoft Corporation states:

AB 2930 aligns with the Microsoft Responsible AI framework and recognizes the promise and potential pitfalls that the next generation of AI may bring through the use of Automated Decision Tools (ADTs). The bill requires developers and users of ADTs to conduct and record impact assessments to evaluate the intended uses of ADTs including an analysis of potential adverse impact on the basis of race, color, ethnicity, sex, religion, age, national origin, or any other classification protected by state law. We believe AB 2930 establishes appropriate legal guardrails on the development and deployment of ADT systems through the use of impact assessments and design evaluations.

Secure Justice states:

We believe AB 2930 is a critical need in this fast-emerging world of AI decision making. Sufficient research exists to demonstrate that many individuals are being actively harmed today as AI proceeds without sufficient oversight, transparency into the decision making, and lack of accountability when discriminatory decisions occur. AB 2930 addresses each of these concerns.

Workday states:

We believe that AB 2930 would establish a strong, pragmatic, and thoughtful regulatory framework for high-risk AI in California. Given the pace of innovation and AI adoption, we believe a “wait and see” approach when it comes to AI and public policy is unworkable and sensible guardrails are needed now to address the AI trust gap. Similarly, we believe that attempting to address every policy issue raised by AI’s development and deployment in a single legislative measure too often leads to proposals that collapse under their own weight and result in inaction. AB 2930 is a sensible way forward that addresses algorithmic discrimination, establishes rigorous AI governance requirements, and builds public trust in AI.

***ARGUMENTS IN OPPOSITION:***

California Chamber of Commerce writes:

We agree that it is critical that companies take care to reduce bias and discrimination in consequential decisions impacting people. At the same time, we believe these issues exist whether these decisions are human made from start to finish, or a byproduct of using or incorporating new technologies in the decision-making process. As such, it is critical that any regulatory efforts proceed with caution and precision, particularly as technology is still developing and has the potential to reduce, if not one day eliminate, such undesirable outcomes.

Electronic Frontier Foundation writes:

A.B. 2930, in covering private and government entities and several areas of consequential decisions, seeks to offer a one-size-fits all solution to an issue that is too complex for a single set of standards. While there may be general common principles for evaluating automated decisionmaking tools, implementing such standards depends heavily on context.

Technet writes:

While existing supervisory risk management and operational resiliency expectations may not expressly address AI, existing risk management principles provide a framework for financial institutions implementing AI to operate in a safe, sound and fair manner – as they would for the use of any other technology. Because AB 2930 is duplicative of the robust rules that our licensed members already follow, and due to the potential negative impacts to California consumers, we must respectfully oppose AB 2930.

## **REGISTERED SUPPORT / OPPOSITION:**

### **Support**

Microsoft Corporation  
Secure Justice  
Workday

### **Support If Amended**

American Federation of Musicians, Local 7  
Bsa the Software Alliance  
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  
Seiu California  
Techequty Action

### **Opposition**

American Council of Life Insurers  
American Property Casualty Insurance Association  
American Staffing Association  
Association of California Life and Health Insurance Companies  
Association of National Advertisers  
California Association of Realtors  
California Bankers Association

California Chamber of Commerce  
California Community Banking Network  
California Credit Union League  
California Financial Services Association  
California Hotel & Lodging Association  
California Manufacturers & Technology Association  
California Mortgage Bankers Association  
California Retailers Association  
California Staffing and Recruiting Association  
Chamber of Progress  
Civil Justice Association of California  
Computer and Communications Industry Association  
CTIA  
Electronic Frontier Foundation  
Insights Association  
Los Angeles Area Chamber of Commerce  
Mortgage Bankers Association  
National Association of Mutual Insurance Companies  
Personal Insurance Federation of California  
State Privacy and Security Coalition, INC.  
Technet

**Oppose Unless Amended**

Internet.works

**Analysis Prepared by:** Slater Sharp / P. & C.P. / (916) 319-2200