

June 2, 2025

**Tom Kemp**

**Director, California Privacy Protection Agency**  
400 R Street, Suite 350  
Sacramento, CA 95811

**RE: Proposed Regulations on CCPA Updates, Cybersecurity Audits, Risk Assessments, Automated Decisionmaking Technology (ADMT), and Insurance Companies**

Dear Board Members, Executive Director Kemp, and Agency Staff,

The Greenlining Institute (“Greenlining”) appreciates the opportunity to provide comment to the California Privacy Protection Agency (“CPPA” or “Agency”) on proposed regulations for the California Consumer Privacy Act (“CCPA”), Cybersecurity Audits, Risk Assessments, Automated Decisionmaking Technology (“ADMT”), and Insurance Companies. Greenlining’s advocacy is rooted in our commitment to advance a more just economy for communities of color, where we work towards a future where race is never a barrier to opportunity. In order for us to arrive at this future, we recognize that it is imperative for the regulatory agencies governing our emergent technologies to be empowered with the rules and authority to meaningfully protect consumers’ rights.

California voters recognized this as well. In 2020, the majority of voters declared that the information asymmetry between consumers and businesses was inequitable. As a result, they voted in favor of Proposition 24, creating the very CPPA that sits here today. Californians have trusted the Agency with the responsibility of protecting and strengthening their privacy rights against potentially predatory industry practices. The CPPA promised it would work to give consumers *meaningful control* over how their information was used.

Unfortunately, today’s iteration of the proposed regulations does the opposite. Rather than drafting rules that target the harms arising from the largely unregulated use of consumer data—specifically inaccurate or biased automated decision making—the CPPA may have been influenced to soften their previously proposed language to render itself effectively toothless against industry abuses. Wholesale adoption of industry arguments only reinforces the very power asymmetry that Californian voters sought to eliminate in the first place with Proposition 24. We are concerned that the most recently proposed language may indicate the CPPA is prioritizing industry motives over reasonable consumer protections.

In order to reaffirm the CPPA’s commitment to consumer protection, we recommend taking the following actions specific to the definition of ADMT and how risk is being assessed. We also support comments from our partners at ACLU California Action, Privacy Rights Clearinghouse, Consumer Federation, Epic, and the Electronic Frontier Foundation on issues specific to civil rights and other key constituencies.

## Proposed Definition of ADMT Still Leaves Consumers Vulnerable to Algorithmic Bias

The proposed decision to change the definition of ADMT from a technology that will “substantially *facilitate*” human decisionmaking to one that will “substantially *replace*” human decisionmaking is a clear attempt to exempt major ADMT use cases from this regulation. These edits, and the subsequent explanation of “substantially replace,” imply that any automated decision-making process that includes the bare minimum of human involvement is somehow free from risk.

In many major use cases where ADMT is used in tandem with human-in-the-loop decision making, algorithmic discrimination still takes place. For example, in automated home valuation programs, Black homeowners living in formerly redlined neighborhoods are consistently given home appraisals 21-23% lower than similar white homeowners.<sup>1</sup> In hiring processes where AI-enabled recruitment and interview technologies are used, applicants with ‘Black-sounding names’ are placed at a disadvantage.<sup>2</sup> In loan-lending, when algorithms generate interest rates, Black applicants end up with interest rates that are 5.6% higher than their white counterparts.<sup>3</sup>

In each of these cases where algorithms inflict major, life-impacting harm, human involvement still exists: licensed appraisers sign off on automated valuations before presenting them to homeowners; HR executives make final hiring decisions based on AI-filtered candidate pools; and loan officers approve applications with algorithmically-determined interest rates. Humans remain involved in nearly all cases where algorithmic bias causes the most severe harm, yet this ruling fails to address or minimize the documented discrimination we have witnessed over the past decade. The presence of human oversight has not prevented these systematic patterns of bias from occurring.

## Automation Bias Still Leads to Algorithmic Bias

Industry lobbyists may also make the argument that incorporating human oversight into automated decision-making processes renders additional regulatory intervention unnecessary, arguing that human reviewers can effectively identify and correct biased algorithmic outputs. This position assumes that human actors possess both the technical expertise and institutional awareness needed to recognize problematic decisions and intervene appropriately. Unless the entirety of California’s workforce is suddenly imbued with a deep understanding of algorithmic infrastructure and decades of robust experience in their respective industries overnight, this assumption does not hold water.

In practice, many automated systems are deployed in emerging use cases where established expertise may not exist, or are overseen by personnel who lack the specialized knowledge required to effectively audit algorithmic outputs. The average human reviewer—whether a customer service representative, loan

---

<sup>1</sup> <https://www.naacpldf.org/appraisal-algorithmic-bias-racial-discrimination/>

<sup>2</sup> <https://ojs.aaai.org/index.php/AIES/article/view/31748/33915>

<sup>3</sup> <https://www.sciencedirect.com/science/article/abs/pii/S0304405X21002403>

officer, home appraiser, or HR staff member—typically receives minimal training on the technical aspects of the automated systems they're supposed to oversee. It is unrealistic to expect average users to possess the domain expertise necessary to identify subtle forms of bias or systematic errors that require intervention. The burden of this expectation should fall onto developers.

Furthermore, as artificial intelligence becomes the new workplace standard, the risk for automation bias grows. When confronted with ADMT-generated outputs, human users—both novice and expert—will often over rely on the algorithms' provided decision, rather than trusting their own judgement.<sup>4</sup> Research shows that this automation bias is more likely to take place when the generated outputs adhere to preexisting stereotypes about groups and social identities, like race and gender; irresponsible systems that leave the door open for complacency and automation bias are more likely to harm communities of color.<sup>5</sup>

Combatting automation bias requires deliberate product design, decision-making transparency, and reasonable guidelines about the volume of tasks that the users are expected to complete. The proposed regulations offer developers no mandate to implement these components into their ADMT. Without an explicit mandate, companies use ADMT to maximize the output of their employees without regard for the harms of automation bias. For example, when the health insurance company Cigna implemented ADMT in their claims-approval process, doctors spent an average of 1.2 seconds scanning patient cases before "reviewing" each request.<sup>6</sup> Former Cigna doctors testified that the system was used to quickly deny claims without any substantive medical review.

Developers and industry lobbyists may try to argue that these ADMT are less discriminatory than humans and, therefore, humans *should* have less oversight in major decision-making processes. Whether or not these claims are true, the fact of the matter is that a biased or inaccurate ADMT can create significantly more harm at mass scale than one bad human decision-maker. In the case of Cigna, after policyholders appealed their denied requests, it was revealed that the ADMT had an 80% error rate—an oversight that was rubber stamped by the 1.2 seconds of menial "human involvement."

With only the vague definition of "human involvement" proposed in sections § 7001.e.1.A-C, companies are given immense discretion as to what bare minimum human involvement can look like. The proposed regulations also fail to provide any specification as to what it would mean for a human actor to "know how to interpret and use the technology" or "review and analyze the output of the decision." There are no benchmarks listed about the human actor being able to articulate the ADMT's decisionmaking logic, identify biases and vulnerabilities in the data, recognize flawed target variables and mitigation strategies. In its current state, the proposed regulations assume these humans are qualified simply by virtue of being there. This is self-regulation in its most irresponsible and vague form.

---

<sup>4</sup> <https://journals.sagepub.com/doi/10.1177/0018720810376055>.

<sup>5</sup> <https://academic.oup.com/jpart/article/33/1/153/6524536?login=false>

<sup>6</sup> <https://apnews.com/article/cigna-california-health-coverage-lawsuit-4543b47cd6057519a7e8dc6d90a61866>

In the event that the CPPA retains this proposed definition of ADMT, we urge the agency to develop a more rigorous definition of "human involvement" that accounts for and mitigates the well-documented risks of automation bias.

The Board should align with the existing California definition of automated decisionmaking systems, which more appropriately captures systems that can cause harm to consumers and workers. The "Alternative 1" definition from the April 4, 2025 board meeting,<sup>7</sup> aligns with existing state definitions of automated systems, and would capture harmful systems that the current narrow definition excludes. For example, an AI system that generates hiring and interview recommendations based on applicant profiles—even if a human reviewer technically makes the "final" decision—would be covered under this definition but could easily escape regulation under the draft definition if a company claims minimal human review exempts them from the definition of ADMT.

If the Board proceeds with the current definition, it must ensure that "human involvement" is truly meaningful rather than perfunctory. The current definition's requirements are insufficient to prevent companies from implementing token human oversight. The proposed regulation should require that human reviewers have sufficient resources, and time, in addition to authority, to meaningfully review automated decisions. The Board should incorporate this language requiring that human involvement be substantive, not merely procedural. The text should read:

- (1) For purposes of this definition, to "substantially replace facilitate human decisionmaking" means a business uses the technology's output to make a decision without human involvement.
- (2) Human involvement requires the human reviewer to:
  - (A) Know how to interpret and use the technology's output to make the decision;
  - (B) Review and analyze the output of the technology, and any other information that is relevant to make or change the decision, including a thorough description of the technologies' decisionmaking logic provided by the developer; and
  - (C) Have the sufficient authority, resources, and time to make or change the decision based on their analysis in subsection (B).

### **Consumers Deserve Notice of Adverse Significant Decisions**

Consumers ought to receive additional notice and access to an ADMT when they are subject to an adverse significant decision. This is the bare minimum for establishing consumer protection against biased ADMT and meaningful control over sensitive data. We urge the CPPA to restore § 7222.k.1-3 and codify this right.

---

<sup>7</sup> A computational process derived from machine learning, statistical modeling, data analytics, or artificial intelligence that issues simplified output, including a score, classification, or recommendation, that processes personal information and is used to assist or replace discretionary human decisionmaking and materially impacts consumers.

Adverse significant decision notifications are essential to minimizing the effects of algorithmic discrimination. When individuals are denied a job, loan, housing, or public benefit due to an automated system, clear notice ensures they understand that an algorithm played a role and enables them to assess whether the decision was fair or lawful. Without notice, affected individuals have no way to identify potential bias, request an explanation, or contest the outcome—effectively stripping them of due process. These are protections that similar data disclosure rules, such as the Fair Credit Reporting Act, offer consumers.<sup>8</sup> Adverse significant decision notifications should be the baseline of consumer protection.

Moreover, notice requirements promote accountability among system developers and deployers by creating a feedback loop incentivizing them to monitor for discriminatory outcomes and improve system fairness. In policy terms, notice is not only a matter of transparency but a necessary condition for oversight, equitable treatment, and the enforcement of civil rights in the digital age.

We urge the CPPA to restore § 7222.k.1-3 and empower consumers' fundamental right to notice, as laid out in the expectations of Proposition 24.

## Conclusion

California has the opportunity to lead the nation in establishing meaningful protections against algorithmic bias. We respectfully urge the CPPA to strengthen, rather than weaken, these critical consumer safeguards by maintaining robust notification requirements, rejecting overly narrow definitions that exclude consequential automated systems, and recognizing that true consumer protection requires transparency and accountability—not merely the presence of humans who may defer to biased algorithmic outputs.

The stakes are too high, and the evidence of harm too clear, to retreat from the comprehensive approach that these regulations originally promised. We ask the CPPA to prioritize consumer protection over industry convenience and ensure that California's regulations fulfill their intended purpose of preventing algorithmic discrimination.

With Regards,

Angel Lin  
**Tech Equity Policy Fellow**

**Mobile:** (425) 623-4720

**Email:** [angel.lin@greenlining.org](mailto:angel.lin@greenlining.org)

**Pronouns:** she/her

---

<sup>8</sup> <https://www.consumercomplianceoutlook.org/2013/second-quarter/adverse-action-notice-requirements-under-ecoa-fcra/>