

January 14, 2025

BY EMAIL

California Privacy Protection Agency  
Legal Division – Regulations Public Comment  
2101 Arena Blvd. Sacramento, CA 95834

RE: Public Comment on California Consumer Privacy Act (CCPA) Updates,  
Cyber, Risk, ADMT, and Insurance Regulations

Resolution Economics, LLC, a consulting firm with offices in Los Angeles, Washington, D.C., New York, Chicago, Charlotte, and Austin, makes this submission in response to the Notice of Public Hearing and Opportunity to Comment that was issued by the California Privacy Protection Agency regarding the Proposed Regulations on the California Consumer Privacy Act (“CCPA”).

We have specific experience in the area of risk assessment and evaluation of AI-enabled tools. We have been and are currently advising clients on how to evaluate the impact of and navigate compliance obligations around AI-enabled and other automated decision making and selection tools. Our experts provide independent audits that assess whether the use of AI-enabled, algorithmic, and other automated tools results in disparate outcomes with respect to race, gender, ethnicity, age, and/or other demographic categories and intersectional identities.

Resolution Economics partner Victoria A. Lipnic, head of our Human Capital Strategy Group and former Acting Chair of the U.S. Equal Employment Opportunity Commission (EEOC), led the Artificial Intelligence Technical Advisory Committee (“AI TAC”) convened by the Institute for Workplace Equality. Several Resolution Economics Directors were members of the AI TAC. This multi-disciplinary group of 40 experts included labor economists, data scientists, industrial-organizational psychologists, attorneys, civil society advocates, AI vendors, employers, and former officials from the EEOC and the Department of Labor’s Office of Federal Contract Compliance Programs (“OFCCP”). In December 2022, the AI TAC released the report *EEO and DEI&A Considerations in the Use of Artificial Intelligence in Employment Decision Making*. This pioneering document, one of the first to address the key issues around automatic decisionmaking technology, analyzes how professional standards, legal precedents, and principles of transparency and fairness apply to AI tools in employment decisions. It offers recommendations on data collection, employee selection procedures, statistical analysis, and addressing adverse impacts.<sup>1</sup>

It is our view that, when properly formulated and implemented, AI audit and assessment requirements can play a critical role in addressing the multifaceted challenges posed by artificial intelligence systems. By subjecting AI-enabled systems to rigorous evaluations and audits,

---

<sup>1</sup> <https://irp.cdn-website.com/b44ff977/files/uploaded/AI-TAC%20Report%20-%20Final%20December%2021%2C%202022.pdf>

potential biases and discriminatory practices can be identified. Where biases or considerable differences across demographic groups are detected in the outcome of AI-enabled systems' use, audits and assessments provide a foundation and a framework for corrective actions. The developers and/or deployers of those systems can implement appropriate remedial measures to address identified problems, enhance the fairness and inclusivity of their AI-enabled systems, and prevent future occurrences of similar issues.

Due to Resolution Economics' expertise in employment-related AI use, we have reviewed the proposed amendments to the CCPA regulations with a particular focus on the employment-related automated decisionmaking technology ("ADMT") provisions. We submit the following comments and questions regarding the Proposed Rules:

1. The current proposed regulations are unclear as to what is required for compliance
  - A. What is required when evaluating ADMT systems used in employment decisions for non-discrimination safeguards?

The proposed regulations require businesses to implement several safeguards when using ADMT for employment decisions. Businesses must conduct an evaluation of the ADMT to ensure it works as intended for the business' proposed use and does not discriminate based on protected classes (§7201). This evaluation is required when ADMT is used for significant decisions concerning employment or independent contracting opportunities or compensation or extensive profiling in employment decisions (including hiring, allocation or assignment of work and compensation, promotion, demotion, suspension, and termination). After evaluation, businesses must implement policies and procedures to ensure the ADMT works as intended and does not discriminate. Businesses must continue to provide training on these policies and procedures.

Notably, the proposed regulations:

- Don't specify what methods must be used for evaluation
- Don't define standards for determining discrimination
- Don't detail what constitutes adequate safeguards
- Don't specify what must be included in policies, procedures, or training

Thus, while the regulations mandate evaluation and safeguards they do not provide sufficiently clear or specific standards for compliance.<sup>2</sup>

---

<sup>2</sup> As a simple example regarding specific questions that may arise when evaluating an ADMT used in employment decisions, consider the issue of missing demographic information. Not all ADMTs seek data regarding race, ethnicity or gender. Even where an ADMT does ask for such information, an increasing number of individuals choose not to disclose their race, ethnicity and/or gender. The proposed regulations do not provide any guidance as to how to take into account such situations when evaluating ADMTs to ensure they work as intended for the business' proposed use and do not discriminate based upon protected classes. For instance, is imputation allowed? Should individuals who choose not to identify race, ethnicity or gender be excluded from the respective race or gender analyses?

B. Which responsibilities lie with the vendor (the provider or developer) and which lie with the user of the ADMT (the employing entity)?

The proposed regulations impose specific obligations on providers and developers of ADMTs, who must provide “all facts necessary” to businesses purchasing their systems for risk assessment purposes (§7153(a)) and must supply “plain language explanation” of system requirements and limitations (§7153(b)). Businesses using third-party ADMT must also review and validate vendor evaluations (§7152(a)(6)(B)), maintain independent safeguards regardless of vendor assurances (§7152(a)(6)(B)), and supplement inadequate vendor information (§7156(b)).

These requirements give rise to key questions, including:

- Whose responsibility is it to resolve potential disparities before and after implementation (which can potentially become much more complicated in multi-party ADMT implementations)?
- How do businesses establish a standard approach for what constitutes “adequate” vendor evaluation and independent safeguards to comply with the proposed regulations?
- What documentation and evidentiary standards will be deemed compliant to demonstrate that a business has sufficiently “reviewed and validated” a third-party ADMT’s evaluation?

The proposed regulations lack specificity to provide clear guidelines differentiating vendor and user evaluation responsibilities, establish detailed protocols for managing shared responsibility scenarios, and define precise parameters for delegating monitoring duties. These gaps create significant compliance challenges for organizations seeking to implement ADMT systems while maintaining regulatory compliance.

2. The proposed regulations lack specific guidance on acceptable methodological standards for assessing ADMTs’ potential discriminatory outcomes

The proposed regulations require businesses to evaluate “the automated decisionmaking technology to ensure it works as intended for the business’s proposed use and does not discriminate based upon protected classes” and “where a business obtains the automated decisionmaking technology from another person, the business must identify the following:

1. Whether it reviewed that person’s evaluation of the automated decisionmaking technology, and whether that person’s evaluation included any requirements or limitations relevant to the business’s proposed use of the automated decisionmaking technology.
2. Any accuracy and nondiscrimination safeguards that it implemented or plans to implement” [§7152(a)(6)(B)].

The proposed regulations’ ADMT evaluation requirements lack the specificity needed to answer key considerations such as:

- What specific methodological standards should be applied when identifying potential discriminatory outcomes in automated decisionmaking technology, particularly for complex machine learning systems where discrimination may not be immediately apparent?
  - What specific methodological standards should be applied when businesses assess and ensure the “quality of personal information” used in ADMT systems (as required by §7152(a)(6)(B)), particularly considering that the regulations’ definition of quality includes completeness, representativeness, timeliness, validity, accuracy, consistency, and reliability of sources?
  - If insufficient data is available to conduct an evaluation, may synthetic/test data be used instead? If so, what features should these synthetic/test data possess?
3. The proposed regulations lack clarity about who should conduct ADMT evaluations

When it comes to who should review ADMT systems – and with what level of independence – the proposed regulations provide different levels of specificity for different types of assessments. They provide the most detailed standards in regard to cybersecurity audits [§7122]. For such audits, the regulations make clear that auditors may be internal or external but must exercise objective and impartial judgment, completely free from business influence. The requirements explicitly require that any cybersecurity auditor, external or internal, report directly to the board of directors, ensuring a level of organizational detachment that prevents potential biases.

The proposed regulations appear to take a different approach regarding risk assessments [§7151], primarily focusing on internal evaluation. The regulations require businesses to engage “relevant individuals” directly involved in the processing activity, typically from product, fraud-prevention, or compliance teams. No specific guidance is provided regarding how to ensure objectivity or independence in such evaluations.

The proposed regulations offer the least amount of guidance when it comes to who should conduct technology evaluations [§7201] to ensure technological performance and prevent discrimination across protected classes. This raises several key questions:

- a. For internally developed ADMT, businesses must conduct their own comprehensive evaluation. The proposed regulations, however, do not address who is to perform those evaluations. Can a business use an internal or external auditor? And to whom should such auditors report?
- b. When using vendor-provided technologies, the proposed regulations appear to give businesses the option to either conduct an independent assessment or review and validate the vendor’s existing evaluation. However, the proposed regulations are mute about the specific methodological standards that should be applied to assess a vendor’s evaluation for compliance. For example, is a vendor evaluation

study acceptable if it is based on a use-case and data from another business? What about if it is based on a use-case and data from a business in a different industry?

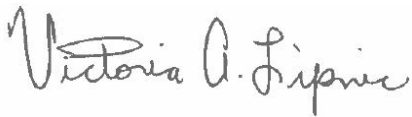
4. The proposed regulations present significant challenges in addressing the nuanced differences between AI model types

Finally, the proposed regulations present significant challenges in addressing the nuanced differences between AI model types. While ADMT systems based on *predictive* AI models typically use established statistical inference methods, *generative* AI models create new content that requires more complex non-discrimination assessments. Hybrid AI models combining predictive and generative approaches pose the most significant evaluation challenges.

The regulations' requirement for a "plain language explanation" of ADMT logic [§7220(c)(5)] oversimplifies the complexity of modern AI architectures. From simple rule-based systems to neural networks with millions of parameters, AI models operate through intricate, often non-linear processes that resist straightforward explanation. For instance, large language models generate decisions through sophisticated interactions across interconnected nodes, where causality is probabilistic rather than deterministic. The proposed regulatory framework creates a fundamental tension between technical complexity and transparency requirements. While mandating explanation of key parameters and logic [§7220(c)(5)], the regulations do not provide concrete guidance on translating complex, high-dimensional computational processes into comprehensible terms. This approach risks forcing companies to produce explanations that are either misleadingly reductive or incomprehensibly technical.

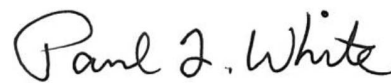
We appreciate the opportunity to provide these comments as part of the California Privacy Protection Agency's proposed rulemaking process.

For Resolution Economics, LLC:



---

Victoria A. Lipnic, Esq.  
Partner  
Resolution Economics, LLC  
Washington, DC



---

Paul White, Ph.D  
Partner  
Resolution Economics, LLC  
Washington, DC

cc: Ali Saad, Ph.D., Resolution Economics, Los Angeles, CA  
Margo Pave, Esq., Resolution Economics, Washington, DC  
Gurkan Ay, Ph.D., Resolution Economics, Washington, DC