June 2, 2025

California Privacy Protection Agency
Attn: Legal Division – Regulations Public Comment
400 R Street, Suite 350 Sacramento, CA 95811

VIA Email: regulations@cppa.ca.gov

Cal Retailers Public Comment on CCPA Updates, Cyber, Risk, ADMT, and Insurance Regulations

Dear Members of the Committee:

The California Retailers Association (Cal Retailers) is submitting the following concerns we have on the modified regulatory text, which is based on the May 1 CPPA hearing, regarding Automated Decision-making Technology ("ADMT"), risk assessments, and cybersecurity. We have also attached our original letter submitted to the CPPA in February for reference.

§ 7001(e) – ADMT Definition:
Request to revise as follows:

- On 1(b), revise: "information that is relevant necessary to make . . ." As currently framed, it is unclear what info would be relevant and may be impossible for a human reviewer to consider all such factors. If a business has a protocol on what info is needed to make a decision or exception, then that should be sufficient.
- On (3), remove "provided that they do not replace human decision making," as it otherwise removes the purpose of the exception. And if a company were to make a decision based solely on a calculator, while perhaps not advisable, it should not be within scope.
- On (3), add "search term software" to exclude when recruiters or employers conduct manual searches using terms to narrow the scope of a recruitment pool.

§ 7001(ddd) – Significant Decision:
Request to revise as follows:

- Under the Significant Decision definition "employment or independent contracting opportunities or compensation" should be removed.
- Under (4), Cal Retailers requests that employment related decisions be limited to hiring or firing—not decisions related to allocation or assignment of work, compensation, bonuses, etc.

§ 7010 (d):
Cal Retailers seeks the removal of this provision. A persistent option for opt-out though a link, versus a just-in-time option, is not as consumer friendly. Consumers are unlikely to have the context needed to know to access the link. The business should determine the appropriate interaction based on its relationship with the consumer and nature of the processing.

The ideal approach would give business flexibility to determine how to offer opt-out; specifically, where PI is being processed for a significant decision, a business should be able to offer the opt-out as part of the user experience that leads to that decision.

§ 7150(b):
Cal Retailers would like the removal of this provision. If this does not prove possible, we have included some feedback along with revisions below.

Feedback:
The rules still regulate the use of ADMT to process publicly available information, as a consumer's presence on a college campus or a grocery store with a pharmacy is not private information. The CPRA otherwise regulates the use of data collected from geo-trackers that identify a consumer's precise geolocation, regardless of the location. As sensitive data, a controller must still conduct a risk assessment (per these regs) and provide an opt out. The overbreadth would capture low risk activities such as providing discounts for (i) prescriptions at specific pharmacies based on a consumer's prior use or (ii) college merchandise based on a student's residence at a specific college.

Potential Revisions:

- (3) – Seek the below revisions to this provision:
    - Using ADMT for a significant decision concerning a consumer that presents a reasonably foreseeable risk of (A) unfair or deceptive treatment of, or unlawful disparate impact on, consumers, (B) financial, physical or reputational injury to consumers, (C) a physical or other intrusion upon the solitude or seclusion, or the private affairs or concerns, of consumers, where such intrusion would be offensive to a reasonable person, or (D) other substantial injury to consumers.
- (4) – Seek removal of this provision.
- (5) – Profiling Sensitive Location – Seek removal of this provision.
- (7) - Seek the creation of a new provision utilizing the below text:
    - A risk assessment completed under another law that is substantially similar to the assessment required under this Article will satisfy the requirements of this Article.

§ 7152(a)(3) – Risk Assessment Requirements:
The purpose of the risk assessment is to make sure the business considers and weighs the privacy harms resulting from certain high-risk processing activities. As businesses continue to innovate, the nature of in-scope processing activities will change. Businesses should retain flexibility in how to approach assessments to make sure that they identify and weigh the right factors. However, the approach in the proposed rules under § 7152 is overly prescriptive and may force businesses to view assessments as a check-the-box exercise rather than focusing on the factors that ultimately matter for the assessment.

The cost of this approach to business and innovation outweighs the privacy benefits to consumers. California companies operate nationally and internationally. Under almost all other privacy laws including GDPR, a business will prepare risk assessments tailored to the processing activity rather than follow the CPPA's formulaic approach. Yet since the proposed rules do not permit a business to rely entirely on an assessment prepared to meet the requirements for another jurisdiction, a business will need to prepare a California-specific supplement for the same processing activity. The CPPA has not explained how this approach will provide incremental benefits to consumer privacy. See § 7156 below for more info.

§ 7156 Interoperability of Risk Assessments:

The modified regulatory text continues to require a California risk assessment to include all the specific requirements under this regulation. Instead, it should follow the approach of all other state privacy laws and permit businesses to rely on assessments prepared for other laws that are reasonably similar in scope and effect. For instance, Colorado requires data protection assessments for (1) processing personal data for targeted advertising (defined as equivalent to "*cross-context behavioral advertising*," not "*behavioral advertising*") and profiling if the profiling presents a reasonably foreseeable risk of (i) unfair or deceptive treatment of, or unlawful disparate impact on, consumers; (ii) financial or physical injury to consumers; (iii) physical or other intrusion on the solitude, seclusion, or private affairs or concerns, of consumers if the intrusion would be offensive to a reasonable person; or (iv) other substantial injury; (2) selling personal data; and (3) processing sensitive data.

Moreover, other state privacy laws that require "*risk assessments*" (i.e. "*data protection assessments*") for high-risk activity limit the scope of activities requiring such assessments in similar circumstances to (1) the processing of "*sensitive data*," which could include location information ***but only when such data is precise geolocation information***, and (2) profiling, ***but only when it presents a reasonably foreseeable risk of the following***: unfair or deceptive treatment of (or unlawful disparate impact on) consumers, financial or physical injury, physical or other intrusion on the solitude or seclusion – or private affairs or concerns – of consumers if it would be offensive to a reasonable person (**hardly the case in a publicly available space, where consumers do not have a reasonable expectation of privacy**), or other substantial injury.

As currently written, the draft rules contemplate interoperability only between similar "*risk assessments*" and do not contemplate "*data protection assessments*." Further, the draft rules are rather stringent in requiring that a company may only forego a risk assessment if (1) the other "*risk assessment*" created for the purpose of complying with another law or regulation "*meets **all** the requirements of this Article*," and (2) if it covers a "*comparable set of processing activities*," defined as processing activities that "*present similar risks to consumers' privacy*." Of course, an entity would not know if an activity presented similar risks until it conducts the risk assessment, thereby the purpose of this provision.

(F): This is an example of how the rule leans more toward the prescriptive rather than functional. Section (a)(1) already requires an assessment to identify the processing purpose, and then (a)(3)(F) requiring mapping those purposes to specific third parties. A business should consider in its assessment both processing purposes and sharing, but this mapping will not always be necessary—especially when a business already discloses categories of data sharing in its privacy notice and the purpose of processing in its agreements with service providers.

(G)(1) [limited to significant decisions]: As described below at § 7220(c)(5), research is still ongoing in how to explain the logic of ADMT models. Moreover, the focus on methodology is not tethered to the risk to the consumer—privacy or otherwise. Instead, the risks are related to an adverse impact of a significant decision and whether a data subject can exercise rights. These are sufficiently addressed by other RA provisions.

§ 7157 Submission of Risk Assessments to the Agency:

The draft rules require companies to submit abridged forms of their risk assessments on an annual basis to the CPPA. Routine submission is not only burdensome, but inconsistent with other state privacy laws and may result in reduced privacy protections as businesses may prepare assessments in a way that is legally protective rather than focused on the right risk-benefit balancing.

However, the CPRA statute mandates that the CPPA issue regulations that require submission at a regular cadence. For instance, the statute does not preclude the CPPA from setting separate standards for what processing activities trigger a risk assessment vs what activities are sufficiently risky to trigger a submission. This would allow the CPPA to focus on those assessments that are highest risk, such as those involving the sale of sensitive data.

Further, the rules require companies to submit risk assessments in the employment context to the regulator, but in most instances, any decisions in the employment context are confidential and not available to competitors. We'd likely need an exception to not require the submission of information that is confidential business/trade secret information.

§ 7157(b):
It is still unclear what the CPPA plans to do with this information. Per our prior points, consider limiting the requirement to certain processing activities only—e.g., sales of sensitive data. Alternatively, seek to limit the substance to metrics, i.e. the number of assessments, and drop (4). Companies are otherwise required to disclose in their privacy notice the types of personal data that they collect, process, and share. Unclear how adding this to the submissions will produce any greater benefit for the Agency.

§ 7150(a)(1) – Significant Decision:
This should be limited in the same manner and for the same reasons above at § 7150(b)(3)(A). § 7200(a)(1) extends application to all uses of ADMTs for decisions regarding provision of, denial of, or access to, employment and employment compensation. Per the regulation, this includes almost all activity within the scope of the employment lifecycle: hiring, promotion/demotion; suspension/termination; and, during employment, allocation/assignment of work; setting of base and incentive compensation; and decisions regarding "other benefits." This also includes "independent contracting opportunities," i.e., the same activities in the IC, gig-economy, and other emerging work contexts.

§ 7150(b)(6) – AI/ADMT Training:
On limiting AI training assessments for models used for significant decisions, we support the revised scope to exclude language covering models that are "capable" of certain purposes and to instead limit to models where the business "intends to use" for those purposes. However, as defined in revised text, the term "intends to use" also covers "permits others to use, plans to permit others to use" and advertising such uses. This language should be removed, as it conflicts with the "intent" language and will bring in scope a wide-range of general use models that are primarily used for other, low-risk purposes. The proposed rules otherwise cover (i) if a deployer intends to use a model to make a significant decision or (ii) a deployer modifies a model with supplemental training that it then intends to use to make a significant decision.

The rules should not extend risk assessments to processing for training a model that is used for emotion recognition, if it does not otherwise involve identifying a specific person (which is already covered). It should also not expressly call out training for models used for biological identification. Risk assessments already extend to processing of sensitive data (which includes biometric data as defined under CPRA). If a deployer is using such a model for biological identification, then RAs already apply. Same if a developer uses biometric data to train a model. But it should not extend to models that are not trained on biometric data—otherwise the rules remove an incentive for developers to minimize the sensitive data that they use in training.

§ 7200(a)(1) – Significant Decision:
See above at 7001(ddd)

The first sentence of 7200(b) should be removed—a business should not have to provide a risk assessment where ADMT was used prior to effective date and is not used on or after the effective date.

§ 7200(a)(2)(A) - Extensive Profiling (Profiling of Employees):
This should be limited in scope.

§ 7200(a)(2)(B) - Extensive Profiling (Publicly Accessible Spaces):
This should be limited in scope.

§ 7200(a)(2)(C) - Extensive Profiling (Behavioral Advertising):
This should be limited in scope.

§ 7150(a)(3) - AI/ADMT Training:
This should be limited in scope.

§ 7220(a) – When Required:
As a threshold matter, the CPRA does not permit regulations on pre-use notice of ADMT—instead, CPRA § 1798.185 calls for regulations "*governing access and opt-out rights*,"
with respect to ADMT. The access right (addressed below) covers the information that a business needs to provide about its ADMT, and so CPPA should not issue separate and overlapping rules on notice.

At minimum, pre-use notice should be limited to where the ADMT processing is otherwise subject to access and opt-out rights. To the extent that one of these customer rights does not apply (e.g., relying on security or fraud prevention exception), then the business should not have an obligation to post this notice. In other words, section 7220(a) should apply subject to the exceptions in § 7221(b) and § 7222(a)(1).

This makes practical sense as forcing a business to make disclosures on how it uses ADMT to perform these functions would undermine the safety and security of consumers and businesses. The notice must provide extensive details about the use of the ADMT, which will be difficult to draft and likely not useful to the consumer, particularly where a business uses ADMT in multiple ways and must provide several notices. Consumers may not be well-equipped to evaluate information about how ADMT works and the logic behind the ADMT.

Suggested Revision:
CPPA did not revise the regulations to limit the pre-use notice requirement to only where ADMT processing is otherwise subject to access and opt-out rights. As a result, businesses will be required to provide such notices even if they use ADMT for exempt purposes for which consumers do not have the right to access or opt out. Amend § 7220(a) as follows:

- A business that uses automated decision-making technology as set forth in section § 7200, subsection (a), and subject to the exceptions in section § 7221(b) and section § 7222(a)(1), must provide consumers with a Pre-Use Notice.

§ 7220(c)(5):
Our preference is to delete pre-use notice entirely. If pre-use notice is required, it should be limited to manageable information.

Suggested Revisions:

- Amend § 7220(c)(5)(A) as follows: The categories of personal information processed by the ADMT.
- Strike § 7220(c)(5)(B)

§ 7220 – Notices:

While the updated rules include the appropriate carveouts (§ 7200(d)), it still requires the notice to include a significant amount of information. The most problematic are (A) and (B), which requires disclosing the type of outputs generated and how the output is used to make a significant decision. Per our original concerns, the CPPA should consider whether this helps the consumer and whether risks are better mitigated through an assessment that requires rigorous testing.

On (A), unclear how it relates to 7222(b)(2) re access right, as one requires disclosing how ADMT processes personal info to decide and other the ADMT logic.

§ 7221(a) – Opt-out of ADMT:

Our preference is to delete the right to opt-out entirely, as this is administratively difficult to implement without significant consumer benefit unless there is an adverse decision. If the right cannot be deleted, it should only apply as a right to appeal in the event of an adverse decision, like the Colorado AI Act and other similar laws, as per the below.

An exception for cybersecurity uses previously included in the regulations should be re-inserted as this type of safe harbor is critical to allowing businesses to safely protect consumer data from unauthorized uses.

Suggested Revisions:

- Amend § 7221(a) to read as follows: In the event of an adverse significant decision having legal or similarly significant effect, a business must provide consumers with the ability to appeal the decision and in that appeal opt-out of the use of ADMT to make a significant decision concerning the consumer, except as set forth in subsection (b).
- Amend § 7221(b) to add the following sections and text:
    - (1) The business's use of that automated decision-making technology is necessary to achieve, and is used solely for, the security, fraud prevention, or safety purposes listed below ("security, fraud prevention, and safety exception"):
    - (A) To prevent, detect, and investigate security incidents that compromise the availability, authenticity, integrity, or confidentiality of stored or transmitted personal information.
    - (B) To resist malicious, deceptive, fraudulent, or illegal actions directed at the business and to prosecute those responsible for those actions.
    - (C) To ensure the physical safety of natural persons; or
    - (D) To protect property or rights or defend against legal claims.

§ 7221(b)(4) & (5) – Employee Exceptions: Cal Retailers requests that the CPPA remove the limiters "solely" in both exceptions—so long as the ADMT is not used to make another type of significant decision, then the opt out should not apply. As written, it suggests that the exception would not apply to ADMT that is used for both assignment of work and how the business manages its products and services—even though the latter is not a significant decision.

The standard "ensures" sets an unreasonably high bar. Propose revising to say that a business must take reasonable measures to ensure. Also, an ADMT deployer should be able to rely on an assessment or instructions from developer rather than conduct an independent assessment.

§ 7221(f)

Request to Amend entire section to the following text only:

- A business may require a verifiable consumer request for a request to opt-out of ADMT set forth in subsection (a). A business may ask the consumer for information necessary to complete the request, such as information necessary to identify the consumer whose information is subject to the business's use of ADMT.

§ 7222 (a) – When Access Right Applies:

This provision in the modified text still requires that business's responses be specific to the specific consumer making the request (see previous Cal retailers letter attached).

At a minimum, the access right should be limited to an adverse decision. A company should not be required to explain details about when and how it uses technology when no harm is involved, such as where a consumer is pre-approved for credit. This follows the approach of FCRA and the Equal Credit Opportunity Act. We are not aware of any other regulatory regime that requires a company to disclose how it made a non-adverse decision to a consumer. We request the removal of section § 7222 (a)(b)(2).

The same opt-out exceptions for employment uses (under § 7221(b)) should be added here. If the terms ADMT and significant decisions are broadly interpreted to include interim hiring decisions or filtering tools, then it will be impractical to require a business to respond to consumer-specific access requests on how the ADMT was applied to them, regardless of whether it was adverse. This will eliminate the advantages of using automated tools in the first place.

§ 7222(b) – Response to Access Request:

We propose limiting the right to access to situations in which the consumer has actually been subjected to an adverse significant decision. Perhaps including a Pre-Use Notice as a baseline disclosure, and then only upon an adverse decision could a consumer potentially obtain more individualized info.

Suggested Revisions:

- On (1), per above, this should be limited at minimum to where the ADMT use resulted in an adverse decision.

- On (2), this language should be removed, even if limited to were used for an adverse decision. Other regulatory regimes that require a business to explain an (adverse) decision do not require disclosing methodology (eg, under FCRA, the notice when taking an adverse action based on a consumer report must explain that an adverse action occurred, identify the consumer rights, and provide contact info of consumer reporting agency—but not the methodology of the decision making). It does not relate to any privacy risk to the consumer but instead creates a moral hazard in that the primary benefit to consumers of accessing a methodology is to either copy it for their own business needs or use it to game the system, defeating the purpose of the technology and harming all consumers.

- On (3), this is inconsistent with the definition of ADMT, which is limited to technologies that fully replace or substantially replace human decision-making. As framed, it suggests that this requirement applies to interim automated tools. To the extent the rule is seeking to inform consumers about the purpose of the decision, then that is already covered under (1). Again, the outcome at most should be limited to adverse decisions, per 7222(a) above.

§ 7222(k) – Adverse Significant Decisions:
In the employment context, the rules give companies too little time to effectively provide detailed, and in parts, data specific to each individual decision—a bar too burdensome given the broad applicability. In the case of an "*adverse significant decision*" (suspension, demotion, termination, or reduction in compensation), the business must provide notice of the right to access within 15 days of the adverse decision, and with detailed information within 45 days. Upon request, the business must provide "*a plain language explanations*" — requiring interpretation — of the purpose of the ADMT, and more concerningly, (a) the specific outputs the ADMT produced after processing the individual's data; (b) the way in which the business used (and plans to use) the ADMT output and human assessment in making decisions regarding that individual; (c) the "*extensive profiling*", if any, performed by the business using an ADMT; and (d) the precise "*logic*", "*key parameters*", and "*range of possible outputs or aggregate output statistics*" of the ADMT, so the individual can understand the workings of the tool and how the specific decision came to be. Little of this information will be helpful to the individual and will require extensive interpretation on behalf of a company to produce this information in "*plain language*" to an individual—often in each specific "*adverse significant decision*."

Behavioral Advertising – Fallback:
The scope of covered ADMT under § 7200 should not include profiling of a consumer, or at minimum, should exclude behavioral advertising. If not, then the draft should at minimum not apply the access right to this type of ADMT processing.

Suggested Revisions:
See previous comments to strike "*behavioral advertising*."

We also have concerns with the way § 7123(f) is currently drafted. As written, it is effectively useless as it says another audit can only be used if another audit has all the same requirements as the CCPA audit. No other audit regime looks like the CCPA audit, so businesses will always be required to conduct a separate audit for CCPA. Most businesses already conduct annual audits for ISO certification. We suggest the regulations include common security audit frameworks that will be accepted as compliant with these regulations without requiring businesses to make the determination whether they meet all the requirements the agency requires.

Related to this, the specific controls outlined in § 7123(b) risk becoming outdated quickly. Most current cybersecurity audit standards focus on assessing how organizations achieve security outcomes. For example, NIST recommends controls such as: "The confidentiality, integrity, and availability of data-at-rest are protected." In contrast, the proposed regulations mandate specific technical controls to achieve these outcomes—for instance, requiring assessment of "Encryption of personal information, at rest."
As an example, § 7123(b)(2)(A) emphasizes multi-factor authentication (MFA) and password requirements, even though many companies are now transitioning to passkeys and other modern authentication methods. We recommend removing subsections (E) and (P), which we believe are overly prescriptive. Additionally, we request that subsection (O) be revised to exclude third parties, as this could result in businesses being required to audit their peers—raising concerns around feasibility and confidentiality.
We also believe the definition of a "security incident" in § 7123(b)(2)(Q) is overly broad. Specifically, including violations of a business's internal program—rather than focusing on unauthorized access—does not align with industry standards for incidents that may require reporting. This could conflate internal compliance issues with actual security events.
Finally, we recommend the following:

- A limitation on the requirement to submit full audit reports,
- The ability to redact sensitive security and proprietary information, and

- A requirement that the CPPA maintain strict confidentiality and security of submitted reports.

Additional Issues:
In addition to the concerns outlined above, we respectfully raise the following issues with the proposed regulations, which may create unintended consequences or conflict with existing statutory frameworks:

1. **Over-Inclusive Deletion Threshold:**
   The proposed regulations would require data brokers to honor deletion requests submitted through the DROP if more than 50% of the unique identifiers provided match a single consumer record. This threshold is overly broad and could result in the deletion of personal information for individuals who did not actually submit a request. Additionally, this approach conflicts with existing CPPA regulations, which require verification of deletion requests to a "reasonable" or "reasonably high degree of certainty," depending on the sensitivity of the data. For more sensitive information, verification typically requires at least three matching data points before a business is obligated to act.

2. **Lack of Verification for Authorized Agents:**
   The proposed DROP regulations lack sufficient safeguards to verify that authorized agents are legitimately acting on behalf of consumers. This omission conflicts with existing CCPA regulations, which require agents to provide signed authorization from the consumer and allow businesses to verify the consumer's identity directly or confirm the authorization. Without similar verification requirements in the DROP process, a significant loophole is created that could be exploited by unauthorized agents.

3. **Insufficient Consumer Verification Requirements:**
   The proposed regulations do not mandate adequate verification to confirm that a deletion request is being made by the actual consumer. While there are limited guidelines for verifying residency, there is no requirement to confirm that the individual is a California resident. Moreover, although the regulations allow for verification of specific data elements, they do not require it. This is inconsistent with existing CCPA rules, which obligate businesses to establish reasonable methods for verifying the identity of the requestor and to assess whether the personal information provided is robust enough to prevent fraudulent or spoofed requests.

4. **Mandated Data Standardization Raises Concerns:**
   The proposed rules would require all registered data brokers to reformat their databases to conform to a standardized format prescribed by the CPPA—such as removing capital letters, extraneous characters, and special symbols. This requirement could introduce data security risks by enforcing uniform formatting across systems and may also raise First Amendment concerns by compelling how business's structure and maintain their data.

5. **Improper Expansion of the "Data Broker" Definition:**
   The proposed expansion of the "data broker" definition through the revised interpretation of "direct relationship" exceeds the CPPA's regulatory authority. By including entities that have a first-party relationship with consumers—such as those that sell personal information but also directly interact with consumers—the CPPA is contradicting legislative intent. The California Legislature clearly intended to limit data broker registration and compliance obligations to entities that do not have a direct relationship with consumers.

Again, we appreciate the opportunity to provide comments on the modified regulatory text but continue to urge a thoughtful reconsideration of these regulations to ensure they protect consumers without

unduly burdening businesses or stifling innovation. California's position as a global leader in AI research and development is at stake, and a balanced, well-deliberated approach is crucial for maintaining our competitive edge while safeguarding consumer interests.

If you have any questions or need additional information on our comments included in this letter, please do not hesitate to contact me directly.

Sincerely,

Jacob Brint
Policy Advocate

Original Cal Retailers letter included on following pages.

February 19, 2025

California Privacy Protection Agency
2101 Arena Blvd.
Sacramento, CA 95834

VIA Email: regulations@cppa.ca.gov.

Cal Retailers Comments on CCPA Updates, Cyber, Risk, ADMT, and Insurance Regulations

Dear Members of the Committee:

The California Retailers Association respectfully urges reconsideration of the proposed regulations regarding Automated Decision-making Technology ("ADMT"), risk assessments, and cybersecurity, as they may inadvertently hinder California's economic growth and innovation while potentially falling short of their intended consumer protection goals. We believe a more balanced approach is necessary to safeguard both consumer privacy and the state's economic vitality.

The Standardized Impact Assessment (SRIA) reveals concerning projections: over 52,000 California businesses could face compliance costs, resulting in a $3.5 billion economic impact. This burden may disproportionately affect small businesses, forcing them to divert resources from growth and innovation to legal and compliance needs. The SRIA also forecasts significant job losses, peaking at 126,000 in 2030, and state revenue losses of up to $2.8 billion annually by 2028.

We appreciate the importance of consumer privacy but believe the proposed regulations may exceed the scope of the California Consumer Privacy Act (CCPA). Even Alastair Mactaggart, author of the California Privacy Rights Act, has expressed concerns about the rules' scope. We suggest that AI regulations be developed through a more inclusive process led by the Legislature and the Newsom Administration, ensuring a thorough evaluation of costs, benefits, and budget impacts.

The proposed regulations, particularly those concerning Automated Decision-Making Technology (ADMT), may unintentionally impede online transactions and research. Multiple pop-up notifications could frustrate consumers and hinder their online experiences, potentially harming small and local businesses that rely heavily on e-commerce. We recommend simplifying notice requirements to focus on high-risk activities, benefiting both consumer privacy and business efficiency. Furthermore, the regulations may inadvertently discourage the use of AI technologies that could enhance efficiency, productivity, and growth across various sectors. By treating low-risk AI applications similarly to high-stakes decisions, we risk losing valuable opportunities for innovation and economic advancement.

We respectfully suggest that the California Privacy Protection Agency (CPPA) collaborate closely with Governor Newsom and the Legislature to develop a risk-based approach that addresses genuine consumer risks while fostering innovation. This approach would align with the Governor's Executive Order on AI, which aims to harness AI's benefits for Californians while avoiding a patchwork of conflicting regulations.

We would also like to share very specific examples within the proposed regulations to illustrate why we encourage the board to take time to collaborate with the Governor and the Legislature on this important issue.

1121 L Street, Suite 607, Sacramento CA 95814 • P: 916/443-1975 • www.calretailers.com • cra@calretailers.com

EXAMPLE #1 - **The draft regulations inappropriately attempt to limit first party advertising:** The draft regulations are overly broad and exceed the California Privacy Protection Agency's (CPPA) authority to regulate beyond what was expressly included in the California Consumer Protection Act (CCPA) and the amendments voters approved in the California Privacy Rights Act (CPRA). CCPA clearly exempted information a business acquires through its own interaction with consumers while CPRA amended CCPA to restrict **cross-contextual behavioral advertising** by requiring a business to obtain a consumer's consent before it could share the consumer's personal information with **third parties**. Instead of providing businesses with implementation guidance, the draft regulations inappropriately attempt to broaden the scope of the CPPA's authority by granting consumers a right to opt-out of ADMT and restrict businesses' use of **first party data**. Many consumers expect businesses to provide relevant product recommendations and personalized ads that correspond to items they have previously purchased or considered purchasing to be able to take advantage of special offers and competitive pricing of goods and services. Many businesses also provide customers with opportunities to restrict how a business can use personal information it has collected about them if they choose. The draft regulations inappropriately violate the First Amendment by restricting businesses' free speech right to advertise their products and services without government interference. The draft regulations provide no compelling state interest for restricting speech, and they do not set forth a narrowly tailored solution to achieve their desired outcome.

EXAMPLE #2: The draft regulations inappropriately attempt to regulate Automated Decision-Making Technologies (ADMT): It is inappropriate for CPPA to use its authority to regulate data privacy as justification for regulating ADMTs and apply a data privacy protection framework when this type of technology was not clearly contemplated by CCPA or CPRA. Last year, the California Legislature considered, but did not pass legislation to regulate ADMTs. The Legislature and Governor continue to consider the appropriate restrictions on AI technology. The Governor signed AB 2013 imposing training data transparency requirements, but vetoed other bills attempting to regulate AI. The Governor has noted the importance of striking the appropriate balance between providing industry incentives to innovate without enacting arbitrary restrictions on technology that will stifle competition and has invited continued conversation on this topic with the Legislature. CPPA is encroaching upon the power of the Legislature to legislate with its attempt to usurp authority to regulate ADMTs.

EXAMPLE #3: The draft regulations do not appropriately distinguish between significant decisions and non-significant decisions: Effective AI laws regulate conduct, not the technology itself; otherwise, continued technological advancement renders them obsolete. For example, Colorado's AI law distinguishes between consequential and non-consequential uses of generative AI and grants consumers the right to appeal consequential decisions to ensure that human review is part of any decision pertaining to a consumer's access to education, employment, financial services, housing, health care, or legal services. Individuals are afforded analogous protection under the EU AI Act. The CPPA's draft regulations do not recognize or describe what would be considered a "significant harm" under existing California law, nor do they refer to any examples that would point to these areas of existing law. As a result, the draft regulations do not provide enough clarity on what would be considered a "significant decision" for businesses to meaningfully rely upon to ensure their compliance with California law.

EXAMPLE #4: The draft regulations interfere with regular business operations under existing law: The draft regulations' requirements for cybersecurity audits force a business' board of directors to perform managerial tasks instead of delegating those tasks to business leaders who are better qualified to execute them. For example, the draft regulations require a board member to sign a written statement they have reviewed stating they understand the findings of the cybersecurity audit. This is not an appropriate requirement given the role of a board of directors is to provide strategic planning, leadership,

1121 L Street, Suite 607, Sacramento CA 95814 • P: 916/443-1975 • www.calretailers.com • cra@calretailers.com

and guidance, not to weigh in on day-to-day business decisions for which the board member may or may not have the appropriate level of experience or expertise to meaningfully evaluate.

EXAMPLE #5: The draft regulations impose burdensome compliance requirements without appropriate justification: The purpose of cybersecurity audits is to ensure that businesses who process significant amounts of sensitive personal information have appropriate safeguards in place to protect consumers from the risk of harm of this information becoming public. The draft regulations provide no threshold to evaluate the significance of the risk of harm to consumers before imposing additional costly and burdensome cybersecurity requirements upon the organization. As directed in CCPA and as amended by CPRA, the draft regulations should have provided a methodology to consider the complexity of the business and the type of information it processes before imposing additional cybersecurity requirements.

EXAMPLE #6: The draft rules include several concerning provisions that may mandate businesses to compromise their proprietary info and IP (e.g., how the logic operates and the key parameters that affect the output). The rules should clarify that no provision shall be construed to require the disclosure of trade secrets or confidential or proprietary information about the design or use of an automated system. Also, per § 1798.185(a)(3), the CPPA must issue rules to clarify that companies are not required to disclose trade secretsor proprietary information.

We also have concerns with specific sections of the proposed regulations.

§ 7001(f) – ADMT Definition - The definition of ADMT is overbroad. Including technology that "*substantially facilitates human decision making*" (i.e., "*using the output of the technology as a key factor in a human's decision making*", as when "*a human reviewer uses [an output] as a primary factor to make a significant decision about them*") will require an impossible line-drawing exercise (what is a "*key/primary factor*"? when are other factors considered by the reviewing human "*key/primary*" when they produce the same result recommended by the ADMT?), and will chill use of innovative technologies in California. An ADMT should not "*substantially facilitate human decision-making*" when it (i) performs a narrow procedural task, (ii) improves the result of a previously completed human activity, (iii) detects decision-making patterns or deviations from prior decision-making patterns and is not meant to replace or influence the previously completed human assessment without proper human review, or (iv) performs a preparatory task to an assessment relevant to a significant decision.

§ 7150(b)(3)(A) – Significant Decision - In every other US State law that defines "*profiling,*" such profiling is tied to a legal or similarly significant decision. And in those states, a decision that produces legal or similarly significant effects is a decision that results in the provision or denial of financial or lending services, housing, insurance, education enrollment or opportunity, criminal justice, employment opportunities, health care services, or access to essential goods or services or basic necessities, such as food and water. Providing consumers with the right to opt-out of profiling (and any other associated rights) is not an easy feat. As such, to the extent California will provide consumers with the right to opt-out of profiling (and other similar rights), those rights should be available only when they will significantly impact consumers.

Suggested Revision: Revise as follows – "For purposes of this Article, "*significant decision*" means a decision using information that is not subject to the exceptions set forth in Civil Code sections 1798.145, subdivisions I-(g), or 1798.146, subdivisions (a)(1), (4), and (5), ~~that results in access to, or~~ the provision or denial of, financial or lending services, housing, insurance, education enrollment or opportunity, criminal justice (e.g., posting of bail bonds), employment ~~or independent contracting opportunities or compensation~~, healthcare services, or essential goods or services (e.g., groceries, medicine, hygiene products, or fuel)."

1121 L Street, Suite 607, Sacramento CA 95814 • P: 916/443-1975 • www.calretailers.com • cra@calretailers.com

§ 7150(b)(3)(B)(i) – Extensive Profiling (Profiling of Employees) - It is not clear what the "*extensive profiling*" concept accomplishes that would not be addressed by the privacy risk assessment requirement for "significant decisions" concerning a consumer, which include decisions about provision or denial or educational or employment opportunities. The statute defines **profiling** as automated processing "*to evaluate certain personal aspects concerning that natural person*," and therefore, it seems extensive "*profiling*" would be encompassed by the privacy risk assessment requirement for significant decisions. The CPPA should avoid duplicative requirements that are likely to confuse California businesses and consumers.

Suggested Revision: Strike § 7150(b)(3)(B), or at minimum, subpart (i).

§ 7150(b)(3)(B)(ii) – Extensive Profiling (Publicly Accessible Spaces) - The requirement to undertake risk assessments for extensive profiling is fundamentally in tension with the statutory text, which explicitly exempts "*publicly available information*" (i.e., information made available to the consumer to the public or from widely distributed media or if the consumer has not restricted the information to a specific audience). When a consumer is in public spaces, they have made a deliberate decision not to restrict the information to a specific audience, and moreover, have no reasonable expectation of privacy.

Accordingly, the CCPA makes clear that requirements for businesses, processors, and contractors, including the creation of risk assessments, do not apply to publicly available information, which includes information collected and processed in public spaces.

The definition of publicly available spaces is also unworkably broad and suggests that it encompasses not only parks and sidewalks, but also shopping areas, stadiums, and other places of congregation. The breadth of this definition would be extremely onerous for California businesses, especially small businesses, without a countervailing benefit to consumers.

Suggested Revision: Strike § 7150(b)(3)(B). If any part of this subpart (ii) remains, it would be helpful to clarify that "*publicly accessible place*" excludes the "*internet*" (similar to the EU AI Act), by clarifying that it refers to a ***physical*** place that is open to or serves the public.

§ 7150(b)(3)(B)(iii) – Extensive Profiling (Behavioral Advertising) - While other state privacy laws require a risk assessment for "*targeted ads*," the draft rules significantly expand and alter what would be required—(i) it would extend to all behavioral ads rather than only CCBA (the focus of the statute) and (ii) impose detailed requirements that are not calibrated to the potential risk, since no consumer data is being shared with third parties or combined with other third-party data.

Suggested Revision: Strike §7150(b)(3)(B), or at minimum, subpart (iii).

§ 7150(b)(4) – AI/ADMT Training - ADMT/model training should not be a category subject to heightened obligations (risk assessments, notice, opt out).

(1) Training a model is not "*automated decision-making*" in its core—because the "*training*" does not involve a decision that has an impact on a specific consumer—and so should be out of scope for these rules. The rules aim to cover certain high-risk AI/ADMT applications, such as when used to make a significant decision. But here, the rules would also cover developing tools that could provide lots of low-risk processing, but would still be in scope because they could one day be used for a higher risk application.

The actual use of ADMT/AI systems for these higher-risk applications would still be covered under these rules, and so extending obligations to the training of such tools is both misplaced and unnecessary. In other words, this training category expands the type of technologies that are subject to these obligations because many if not all models "*could*" be used to make a significant decision.

1121 L Street, Suite 607, Sacramento CA 95814 • P: 916/443-1975 • www.calretailers.com • cra@calretailers.com

This "*theoretical*" approach is inconsistent with other risk-based frameworks focused on automated decision-making used to make a significant decision. This is a different issue because training a model on personal data is different from making a decision about that person (or otherwise creating any risk for them).

(2) This also exceeds the subject matter of what the CCPA contemplates, i.e., the privacy risk that may result from the processing of personal data. The statute and rules already provide ways for consumers to control how their data is used for training—they can opt out of ADMT that results in legal or similarly significant effects, access the data that a business processes about them, correct their data, and delete their data. The CPPA should not use this rulemaking to impose risk assessments that regulate AI training more broadly, when untethered to the privacy risk.

(3) The CPPA significantly departs from the approach taken in other state privacy frameworks, which neither mention nor provide heightened requirements for the use of personal information for model training. It also differs from the one other AI law (CO), where model training is not considered a high-risk decision. Also, California passed AB 2013 this year, which already imposes disclosure requirements on training data.

Suggested Revision: Strike §7150(b)(4)

§ 7154(a) – Prohibition of Certain Activities: The draft rules will prohibit processing of personal info for any covered activity if the risks outweigh the benefits. This is an extreme prohibition, and goes far beyond other AI regulations (e.g., the EU AI Act bans very limited categories of uses). It will also discourage innovation since the balance between benefits and risks is highly subjective and may be close depending on what perspective is applied (e.g., what qualifies as a risk or benefit varies wildly among experts). As an alternative, consider formulation used for unfairness under other legal regimes is the processing likely to cause substantial harm to consumers that is not reasonably avoidable (e.g., an opt out after reasonable notice and option for human review), and the injury is not outweighed by the benefit to consumers. This formulation limits the restriction to only processing that causes "substantial harm" rather than where there is only a non-material impact, and acknowledges that through the opt out, consumers can make their own choice about whether to permit the activity (rather than regulators making the choice for them).

Suggested Edit: Strike § 7154(a).

§ 7156 Interoperability of Risk Assessments: The rules governing "*risk assessments*" should align and be interoperable with the requirements for data protection assessments in other states. For instance, Colorado requires data protection assessments for (1) processing personal data for **targeted advertising** (defined as equivalent to "*cross-context behavioral advertising*," not "*behavioral advertising*") and **profiling** if the profiling presents a reasonably foreseeable risk of (i) unfair or deceptive treatment of, or unlawful disparate impact on, consumers; (ii) financial or physical injury to consumers; (iii) physical or other intrusion on the solitude, seclusion, or private affairs or concerns, of consumers if the intrusion would be offensive to a reasonable person; or (iv) other substantial injury; (2) **selling** personal data; and (3) processing **sensitive data**.

Moreover, other state privacy laws that require "*risk assessments*" (i.e. "*data protection assessments*") for high-risk activity limit the scope of activities requiring such assessments in similar circumstances to (1) the processing of "*sensitive data*," which could include location information ***but only when such data is precise geolocation information***, and (2) profiling, ***but only when it presents a reasonably foreseeable risk of the following***: unfair or deceptive treatment of (or unlawful disparate impact on) consumers, financial or physical injury, physical or other intrusion on the solitude or seclusion – or private affairs or concerns – of consumers if it would be offensive to a reasonable person (**hardly the case in a publicly available space, where consumers do not have a reasonable expectation of**

1121 L Street, Suite 607, Sacramento CA 95814 • P: 916/443-1975 • www.calretailers.com • cra@calretailers.com

privacy), or other substantial injury.

As currently written, the draft rules contemplate interoperability only between similar "*risk assessments*" and do not contemplate "*data protection assessments*." Further, the draft rules are rather stringent in requiring that a company may only forego a risk assessment if (1) the other "*risk assessment*" created for the purpose of complying with another law or regulation "*meets **all** the requirements of this Article*," and (2) if it covers a "*comparable set of processing activities*," defined as processing activities that "*present similar risks to consumers' privacy*." Of course, an entity would not know if an activity presented similar risks until it conducts the risk assessment, thereby the purpose of this provision.

§ 7157 Submission of Risk Assessments to the Agency: The draft rules require companies to submit abridged forms of their risk assessments on an annual basis to the CPPA. Routine submission is not only burdensome, but inconsistent with other state privacy laws and may result in reduced privacy protections as businesses may prepare assessments in a way that is legally protective rather than focused on the right risk-benefit balancing.

However, the CPRA statute mandates that the CPPA issue regulations that require submission at a regular cadence. For instance, the statute does not preclude the CPPA from setting separate standards for what processing activities trigger a risk assessment vs what activities are sufficiently risky to trigger a submission. This would allow the CPPA to focus on those assessments that are highest risk, such as those involving the sale of sensitive data.

Further, the rules require companies to submit risk assessments in the employment context to the regulator, but in most instances, any decisions in the employment context are confidential and not available to competitors. We'd likely need an exception to not require the submission of information that is confidential business/trade secret information.

§ 7150(a)(1) – Significant Decision: This should be limited in the same manner and for the same reasons above at § 7150(b)(3)(A). § 7200(a)(1) extends application to all uses of ADMTs for decisions regarding provision of, denial of, or access to, employment and employment compensation. Per the regulation, this includes almost all activity within the scope of the employment lifecycle: hiring, promotion/demotion; suspension/termination; and, during employment, allocation/assignment of work; setting of base and incentive compensation; and decisions regarding "other benefits." This also includes "independent contracting opportunities," i.e., the same activities in the IC, gig-economy, and other emerging work contexts.

§ 7200(a)(2)(A) - Extensive Profiling (Profiling of Employees): This should be limited in scope.

§ 7200(a)(2)(B) - Extensive Profiling (Publicly Accessible Spaces): This should be limited in scope.

§ 7200(a)(2)(C) - Extensive Profiling (Behavioral Advertising): This should be limited in scope.

§ 7150(a)(3) - AI/ADMT Training: This should be limited in scope.

§ 7220(a) – When Required: As a threshold matter, the CPRA does not permit regulations on pre-use notice of ADMT—instead, CPRA § 1798.185 calls for regulations "*governing access and opt-out rights*," with respect to ADMT. The access right (addressed below) covers the information that a business needs to provide about its ADMT, and so CPPA should not issue separate and overlapping rules on notice.

At minimum, pre-use notice should be limited to where the ADMT processing is otherwise subject to access and opt-out rights. To the extent that one of these customer rights does not apply (e.g., relying on

1121 L Street, Suite 607, Sacramento CA 95814 • P: 916/443-1975 • www.calretailers.com • cra@calretailers.com

security or fraud prevention exception), then the business should not have an obligation to post this notice. In other words, section 7220(a) should apply subject to the exceptions in § 7221(b) and § 7222(a)(1).

This makes practical sense as forcing a business to make disclosures on how it uses ADMT to perform these functions would undermine the safety and security of consumers and businesses.

Suggested Revision: Amend § 7220(a) as follows: *A business that uses automated decision-making technology as set forth in section § 7200, subsection (a), and subject to the exceptions in section § 7221(b) and section § 7222(a)(1), must provide consumers with a Pre-Use Notice.*

§ 7220(c)(5) – Explainability: The draft requires businesses to explain, in plain language, the logic used in the automated decision-making technology, including the key parameters that affect the output of the automated decision-making technology. We recommend deletion of this provision as it is effectively an explainability requirement. Research in the field of explainable ADMT is progressing rapidly, but many complex AI models (which tend to be the most useful ones) are not yet fully explainable. Indeed, requiring an explanation now could result in consumer confusion. CPPA should therefore consider whether it would benefit California to impose this requirement, or whether there are other better methods to mitigate risk such as human review and rigorous testing.

The draft rules are also in tension with the statute's explicit recognition that the CCPA's requirements do not require the business to disclose trade secrets (Cal. Civ. Code 1798.100(f)). The exception under § 7220(c)(5)(C) is too narrow. This is particularly important in the HR context, because HR deals with not only employee confidential data, but also beta and pilots for products that should be excluded as confidential business/trade secret information.

Suggested Revision: Strike § 7220(c)(5).

Behavioral Advertising – Fallback: For reasons stated above, the scope of covered ADMT under § 7200 should not include profiling of a consumer, or at minimum, should exclude behavioral advertising. If that is not excluded, then any notice requirements should be tailored to the reduced risk and different circumstances of advertising. For instance, as drafted, the rules would impose more detailed disclosures than required for higher-risk cross-context behavioral ads under § 7013.

AI/ADMT Training – Fallback: For reasons stated above, the scope of covered ADMT under § 7200 should not include the use of ADMT for training. If that is not excluded, then at minimum such processing should not trigger pre-use notice. It will contribute to notice fatigue without reducing risk as consumers will struggle to understand this notice. Instead, the draft rules already require AI/ADMT deployers (i.e., "users") to conduct risk assessments and companies may invest in accuracy/testing safeguards that better demonstrate trustworthiness.

§ 7221(b) – Exceptions General: Expand the set of exceptions under § 7221(b) to include conducting internal research, fixing technical errors, effectuating product recalls, and performing internal operations consistent with the consumer's expectations (like other privacy laws).

§ 7221(b)(1) - Security and Fraud Prevention: The fraud and security exception should be unencumbered by whether ADMT is "necessary" for the purposes of security, fraud prevention, or safety. Businesses should be free to choose the most effective and reasonable method of security, fraud prevention, or safety without regard for whether a particular method is "necessary."

Suggested Revision: Amend § 7221(b)(1) as follows: (b)(1) The business's use~~s of~~ that automated decision-making technology **~~is necessary~~** **solely** to achieve, **~~and is used solely for~~**, the security, fraud prevention, or safety purposes listed below…

§ 7221(i) – Single ADMT Opt Out: The draft rules would require a business to offer a single option to opt-out of all covered ADMT, although businesses may present consumers with a choice to allow specific uses. Consumers will struggle to comprehend the impact of a general opt out in the abstract or to analyze a vast range of potential use cases (e.g., behavioral advertising vs screening for health risks). This will likely result in consumers making opt-out elections to avoid certain high-risk use cases, and then losing out on significant beneficial opportunities that would likely approve when presented with the specific use case. Instead, business should be required to surface an opt out that is targeted to the specific use case so the consumer can decide in real time and in context rather than in the abstract. Also, mandating a single opt out presumes that the use of ADMT is generally harmful to consumers or lacks benefits, and is antithetical to California's support for innovation, efficiency, and tools that reduce human error and bias.

Suggested Revision: Amend § 7221(i) as follows: In responding to a request to opt-out of ADMT, a business may present the consumer with the choice to allow specific uses of automated decision-making technology ~~as long as the business also offers a single option to opt-out of all of the business's use of automated decision-making technology set forth in subsection (a).~~

§ 7221(b)(4) & (5) – Employee Exceptions: The draft rules provide some exceptions to the ADMT Opt-Out for "*allocation/assignment of work and compensation decisions*" and *"work or educational profiling."* However, these exceptions require companies to conduct *"an evaluation"* and implement expensive and burdensome *"accuracy and nondiscrimination safeguards."* Employers generally are allowed flexibility to ensure employees are working and productive. This would stifle employers' ability to ensure employees are properly working or the company is properly staffed. It can also affect customer service where companies look at these tools to identify when to route calls to employees to ensure 1) they are working and 2) ensure they are not overwhelmed with the volume of calls. Moreover, it is not clear what "*work or educational profiling*" means. The proposed rules refer to "*extensive profiling*," not "*work or educational profiling*."

Behavioral Advertising – Fallback: The scope of covered ADMT, under § 7200, should not include profiling of a consumer, or at minimum, should exclude behavioral advertising. If not, then the draft should at minimum not apply an opt-out right to all behavioral ads as it conflicts with the statutory opt-out framework that is limited to CCBA. As a final fallback, the rules should be clear that any opt out that applies to behavioral ads is limited to targeting ads based on inference preferences from consumers based on their personal data. The draft rules potentially sweep in contextual ads since the definition of behavioral ads is not limited to activity "*over time*."

Additionally, the proviso in the "*behavioral advertising*" definition makes it seem like measurement (e.g., attribution) could be covered as well. It unhelpfully copies some CPRA language, but without the additional context from the CPRA that clarifies that measurement is not in scope.

AI/ADMT Training – Fallback: § 7221(b) provides a limited set of exceptions to the opt-out right, but does not extend them to AI/ADMT training. For reasons stated above, the scope of covered ADMT under § 7200 should exclude ADMT training uses. If that is not excluded, then the exceptions should apply to this use, in particular, the fraud and security exception under (b)(1) and the evaluation exception under (b)(3), (4), and (5). Businesses should be encouraged to evaluate whether the ADMT is discriminatory or working as intended and the best way to do that is to train on representative samples of data.

If AI/ADMT training is not excluded entirely from Article 11, then another fallback should exclude or limit the right to opt out. Generally, model training (especially for the largest models trained on internet-scale data) personal data included in the training corpus is incidental. Requiring implementation of an opt out process runs counter to data minimization best practices, as it may require individual identification.

1121 L Street, Suite 607, Sacramento CA 95814 • P: 916/443-1975 • www.calretailers.com • cra@calretailers.com

Suggested Revisions: The preferred option would be to strike § 7150(b)(4) – AI/ADMT Training altogether. If not, amend § 7221(b) by striking (b)(6).

§ 7222 Requests to Access ADMT: General concerns - CPRA § 1798.185 instructs the CPPA to issue regulations on access rights with respect to a business's use of ADMT, and requiring responses to include "*meaningful information about the logic involved in those decision-making processes, as well as a description of the likely outcome of the process with respect to the consumer*." It does not call for separate notice. Based on these statutory requirements, the CPPA should consider a single set of rules about how businesses need to provide meaningful information about their use of ADMT.

Businesses should have the option to provide this information in a notice (rather than a response to a specific request). The rules should not require businesses to provide consumer-specific responses, as this (i) is not required under the statute, (ii) is not practical or feasible in many cases, and (iii) would be difficult to comply without disclosing confidential information or allowing consumers to game the process, which would have dangerous implications where the ADMT is used to make a significant decision. Consumers already have separate access rights under the CPRA, and so can obtain any personal information processed by the company, including ADMT inputs and outputs containing their personal information.

§ 7222 (a) – When Access Right Applies: The draft rules create a right to access ADMT when a business uses ADMT for significant decisions (§ 7200(a)(1)) or extensive profiling (§ 7200(a)(2)). It does not apply to AI/ADMT training. This right should be limited to where ADMT is used to make a significant decision. The access right serves to allow consumers to understand whether they want to exercise their opt-out right, and to allow them to correct any inaccurate input concerning their personal information. This may assist consumers when presented with an ADMT offering that will assist in making a significant decision, but does not apply to "extensive profiling" such as profiling for behavioral advertising. For those uses, the consumer can decide whether to opt out regardless of how technology works. Businesses should not be required to publicly disclose confidential information about their technological processes absent any direct consumer benefit.

Suggested Revisions: Amend § 7222(a) as follows: *Consumers have a right to access ADMT when a business uses automated decision-making technology as set forth in § 7200, subsections (a)(1~~)~~ ~~(2)~~ .... Strike § 7222(b)(3)(B).*

§ 7222(b) – Response to Access Request: Per above, the responses should not be tailored to specific consumers. At minimum:

(b)(2) should be modified so that the business must provide only the range of potential outputs and not the specific output as it relates to the consumer. If the output itself contains personal information related to the consumer, then it would be subject to the separate, broader access right under the CPRA.

(b)(4) should be clarified to not require the business to explain how the ADMT operated with respect to a specific consumer.

§ 7222(k) – Adverse Significant Decisions: In the employment context, the rules give companies too little time to effectively provide detailed, and in parts, data specific to each individual decision—a bar too burdensome given the broad applicability. In the case of an "*adverse significant decision*" (suspension, demotion, termination, or reduction in compensation), the business must provide notice of the right to access within 15 days of the adverse decision, and with detailed information within 45 days. Upon request, the business must provide "*a plain language explanations*" — requiring interpretation — of the purpose of the ADMT, and more concerningly, (a) the specific outputs the ADMT produced after processing the individual's data; (b) the way in which the business used (and plans to use) the ADMT output and human assessment in making decisions regarding that individual; (c) the "*extensive profiling*",

if any, performed by the business using an ADMT; and (d) the precise "*logic*", "*key parameters*", and "*range of possible outputs or aggregate output statistics*" of the ADMT, so the individual can understand the workings of the tool and how the specific decision came to be. Little of this information will be helpful to the individual and will require extensive interpretation on behalf of a company to produce this information in "*plain language*" to an individual—often in each specific "*adverse significant decision*."

Behavioral Advertising – Fallback: The scope of covered ADMT under § 7200 should not include profiling of a consumer, or at minimum, should exclude behavioral advertising. If not, then the draft should at minimum not apply the access right to this type of ADMT processing.

Suggested Revisions: See previous comments to strike "*behavioral advertising*."

We also have concerns with the way § 7123(f) is currently drafted. As written, it is effectively useless as it says another audit can only be used if another audit has all the same requirements as the CCPA audit. No other audit regime looks like the CCPA audit, so businesses will always be required to conduct a separate audit for CCPA. Most businesses already conduct annual audits for ISO certification. We suggest the regulations include common security audit frameworks that will be accepted as compliant with these regulations without requiring businesses to make the determination whether they meet all the requirements the agency requires.

Related to this, the specific controls in § 7123(b) run the risk of quickly becoming outdated. Most existing cybersecurity audit standards call for the assessment of how organizations achieve outcomes (e.g., NIST recommends as a security control, "*The confidentiality, integrity, and availability of data-at-rest are protected.*"). The proposed regulations instead require specific security controls to achieve certain outcomes (e.g., requiring assessment of "*Encryption of personal information, at rest*"). For example, (b)(2)(A) focuses on MFA and passwords when most companies are increasingly moving to passkeys.

Finally, we suggest a limitation on the submission of full audits, allowance for redaction of sensitive security information and other information, and a requirement that the CPPA keep the reports secure and confidential.

General Comments and Concerns on AI/Privacy Regulations Impact on Emergencies
Weave in how these regulations will negatively impact the supply chain or small business recovery for those who are trying to rebuild after emergencies like the recent wildfires in Los Angeles County.

Again, we appreciate the opportunity to provide comments on the proposed regulations, but urge a thoughtful reconsideration of these regulations to ensure they protect consumers without unduly burdening businesses or stifling innovation. California's position as a global leader in AI research and development is at stake, and a balanced, well-deliberated approach is crucial for maintaining our competitive edge while safeguarding consumer interests.

If you have any questions or need additional information on our comments included in this letter, please do not hesitate to contact me directly.

Sincerely,

Sarah Pollo Moo
Policy Advocate

1121 L Street, Suite 607, Sacramento CA 95814 • P: 916/443-1975 • www.calretailers.com • cra@calretailers.com