

California Privacy Protection Agency  
Attn: Legal Division – Public Comment Regarding CCPA Updates, Cyber, Risk, and ADMT  
Regulations  
2101 Arena Blvd.  
Sacramento, CA 95834  
regulations@coppa.ca.gov

February 19, 2025

To the Leadership Team and Board of the California Privacy Protection Agency:

We write on behalf of Gibson Dunn & Crutcher LLP's Privacy, Cybersecurity and Data Innovation, Artificial Intelligence, and Tech and Innovation practice groups. Gibson Dunn is subject to the California Consumer Privacy Act and advises clients in many industries on the continuously evolving regulation of data, privacy, cybersecurity, and artificial intelligence. While we offer these comments on our own behalf, and our views may not reflect the views of all our clients, our team includes several former executives at technology companies, and our collective experiences give us unique insight into the practical implications of regulations targeting data practices and technology.

While we appreciate the need for sound regulation, we have significant concerns with the Agency's proposed regulations under the CCPA<sup>1</sup> to govern automated decisionmaking technology ("ADMT"), risk assessments, and cybersecurity audits. As the global epicenter of information technology and artificial intelligence, California has delivered tremendous benefits to society. These benefits are a direct product of Californians' ability to creatively innovate using data and technology. As drafted, however, the proposed rules would impede progress in some of the most promising areas of technological opportunity. They would create headwinds to innovation and stall the engine that has driven so much economic growth in this State.

The net effect of the proposed rules would be to divert resources away from responsible innovation and toward cumbersome and ineffective compliance obligations that do little to protect the privacy and security of Californians. The rules would impose unprecedented burdens on businesses, subjecting them to requirements more onerous than similar regulations in Europe, and putting California out of step with the rest of the country and world. We also fear these regulations would be leveraged to compel a barrage of dense, interruptive disclosures on virtually every commercial

---

<sup>1</sup> As amended by the California Privacy Rights Act ("CPRA").

website and app, disclosures that promise to at best annoy California consumers and more likely confuse, alarm, and mislead them.

The current proposal also exceeds the CCPA's grant of rulemaking authority. Though the CCPA was written to advance focused privacy and data-security objectives, the proposed regulations instead seek to redress complex social issues from civil rights to economic equity that are simply beyond the statutory mandate. Under the guise of regulating automated decisions, the rules propose to cover everyday decisions made by humans simply because those decisions rely in some part on software.

We thus urge the Agency to revisit these regulations to advance instead the privacy and security objectives that animated the CCPA, while allowing businesses to innovate free from exceptional restrictions that would not benefit any California consumer. We write to highlight our most pressing concerns.

## **I. The Proposed Regulations Exceed and Are Inconsistent with the Statutory Authorization**

The proposed regulations must be consistent with the statute that authorized them.<sup>2</sup> And they may not vary from or enlarge the statute's terms.<sup>3</sup> The proposed regulations do not adhere to these principles in certain foundational respects.

The CCPA was originally enacted in 2018 with the stated goal of ensuring the privacy of Californians' personal information. As discussed in more detail below, the 2020 ballot initiative, Prop. 24, amended the CCPA to further strengthen the privacy and security of personal information – including by creating the CPPA to protect, as the Agency's name implies, Californians' privacy.

This 2020 amendment contains two relevant grants of authority. Section 1798.185(a)(14) authorizes the Agency to:

[I]ssu[e] regulations requiring businesses whose processing of consumers' personal information **presents significant risk to consumers' privacy or security** . . . [to] [p]erform a cybersecurity audit on an annual basis . . . [and to] submit to the California Privacy Protection Agency on a regular basis a risk assessment.<sup>4</sup>

---

<sup>2</sup> Gov. Code, § 11342.2 (“No regulation adopted is valid or effective unless consistent and not in conflict with the statute”).

<sup>3</sup> *Credit Ins. Gen. Agents Ass'n v. Payne* (1976) 16 Cal.3d 651, 656.

<sup>4</sup> Civ. Code, § 1798.185, subd. (a)(14)(B) (emphasis added).

And Section 1798.185(a)(15) authorizes the CPPA to:

Issue regulations governing **access and opt-out rights** with respect to a business' use of **automated decisionmaking technology, including profiling** and requiring a business' response to access requests to include meaningful information about the logic involved in those decisionmaking processes, as well as a description of the likely outcome of the process with respect to the consumer.<sup>5</sup>

In several key ways, the proposed regulations stray from these narrow authorizations. They would cover a vast range of technologies, use cases, and perceived harms and would impose unprecedented requirements on virtually every business that uses technology. These requirements do not advance, but instead conflict with, the privacy and security aims of the animating law.

**A. The proposed regulations would improperly regulate *human* decisionmaking under a grant of authority to regulate only *automated* decisionmaking**

Subsection (a)(15) authorizes the Agency to issue targeted regulations governing “automated decisionmaking technology,”<sup>6</sup> a term which is not defined in the statute. The Agency has proposed defining “automated decisionmaking technology” as “any technology that processes personal information and uses computation to” do one of three things: “execute a decision, replace human decisionmaking, or substantially facilitate human decisionmaking.”<sup>7</sup>

This definition conflicts with the statute. The statutory phrase “automated decisionmaking” is a term of art, first introduced in European privacy regulations, which refers to “a decision based solely on automated processing.”<sup>8</sup> The same definition results from giving each word in “automated decisionmaking technology” its plain meaning: “Decisionmaking” is “the process or practice of making choices or judgments, esp. after a period of discussion or thought.”<sup>9</sup> And “automated” means “self-acting or self regulating,” “*without needing human control*.”<sup>10</sup>

---

<sup>5</sup> Civ. Code, § 1798.185, subd. (a)(15) (emphasis added).

<sup>6</sup> Civ. Code, § 1798.185, subd. (a)(15).

<sup>7</sup> Proposed Text of Regulations (Cal. Priv. Prot. Agency, Nov. 2024) (hereafter Draft Regulations), § 7001, subd. (f) (emphasis added).

<sup>8</sup> EU General Data Protection Regulation (GDPR), art. 22.

<sup>9</sup> *Decision-making*, Black's Law Dict. (12th ed. 2024).

<sup>10</sup> *Automated*, Merriam-Webster Dict. (“operated automatically”); *Automatically*, Merriam-Webster Dict. (“done or produced as if by machine . . . having a self-acting or self-regulating mechanism”); *Automated*, Cambridge Dict. (“carried out by machines or computers without needing human control”); *Automated*, Oxford English Dict. (“Converted so as to operate automatically . . . automatic”); *Automatic*, Oxford English Dict. (“self-generated, spontaneous; . . . self-acting; having the power of motion within itself”).

The proposed definition *partially* maps to this plain meaning. One of its three components is “any technology that . . . uses computation to . . . replace human decisionmaking,” which tracks the statutory term. This is an appropriately narrow definition. It may cover, for example, a machine-learning algorithm used by a college to predict the future performance of high school students based on data in their application and then decide, without human input, which students to admit.

But the other two components of the definition do not track the statutory grant of authority. First, the proposed regulations would cover “*executing*” a decision already made by a human. By definition, then, technology in this bucket would not be “making” a decision and so fall outside the authorization. For example, if a law firm decides that associates who work above a certain number of hours will receive a bonus, a program that automatically identifies and notifies associates who are above or below that pre-determined threshold is merely executing the decision already made by the firm. It is not, in any meaningful sense, “making” a decision about who will receive a bonus. But the regulations would apparently cover this use case. The statute does not plausibly regulate this use of technology.

Second, the regulations improperly propose to regulate “human decisionmaking” that is “*substantially facilitat[ed]*” by technology. For instance, the regulations stipulate that “generat[ing] a score about a consumer that [a] *human reviewer* uses as a primary factor to make a significant decision” would be regulated.<sup>11</sup> By its own admission, then, this third proposed definition does not regulate “automated” decisionmaking.<sup>12</sup> Nothing in the CCPA authorizes regulating *human* decisions simply because they are aided or informed by technology.<sup>13</sup> In fact, in recent decades, a significant amount of human decisionmaking has been “substantially facilitated” by “the output of . . . technology.” Take an entity that consults a medical diagnostic to help determine whether someone is eligible for a clinical trial; or a business that consults a review website’s algorithm when choosing what plumber to hire, but ultimately has a human make the final call. Nobody would naturally say that these examples involve “*automated* decisionmaking,” even if an automated process informs a decision that is ultimately made.<sup>14</sup>

---

<sup>11</sup> Draft Regulations, § 7001, subd. (f)(2).

<sup>12</sup> See *Southwest Airlines Co. v. Saxon* (2022) 596 U.S. 450, 457–58 (describing “meaning-variation canon” as “where [a] document has used one term in one place, and a materially different term in another, the presumption is that the different term denotes a different idea”).

<sup>13</sup> “Facilitating” just means “mak[ing] easier” or “help[ing] bring (something) about.” (See *facilitate*, Merriam-Webster Dict.). Like “executing,” “facilitating” does not involve the making of any decisions.

<sup>14</sup> The Agency’s proposed regulations governing the opt-out rights, and specifically the exemptions, underscore this problem. As an initial matter, this “human appeal” exception and the other exemptions in the proposed regulations are unmoored from the statutory purpose of advancing privacy and security, focusing instead on issues like accuracy, fairness, and discrimination. And the human appeal exception in particular demonstrates the overbreadth of the Agency’s definition of ADMT: If a decision is subject to human review, then it is, by definition, not automated; it is ultimately being made by a human. Yet the exception applies only to certain types of decisions, when a human appeal should remove a decision from

Although the draft regulations propose to exempt technologies akin to a “calculator,” this limitation does not do anything. In the same breath, the regulations provide that calculators and the like *are* covered if used to “execute a decision, replace human decisionmaking, or substantially facilitate human decisionmaking.”<sup>15</sup> Since that is just the definition of ADMT reprinted, the “calculator” exception does not change the scope of the regulations’ coverage. And indeed the regulations are replete with supposed examples of “automated decisionmaking technology” that work exactly like calculators. For example, the regulations offer as an example of ADMT “a business’s use of a spreadsheet to run regression analyses” on employees’ performance records.<sup>16</sup> But many calculators have a regression function.<sup>17</sup> It is even possible to calculate a regression on a four-function calculator (or even by hand), using just addition, subtraction, multiplication, and division.<sup>18</sup> If regressions count as ADMT, the purported exclusion of “calculators” cannot mean very much. Likewise, Section 7150(c)(1) contends that the regulations would apply when a rideshare platform assigns rides to drivers, even though rideshare platforms typically allocate work based on human-specified geospatial formulas that calculate which driver is closest to the customer, rather than any sort of automated decision.<sup>19</sup> The lack of real difference between the technologies explicitly included and purportedly excluded under the regulations suggests that in practice, virtually all forms of computation will be covered. Because the CCPA authorizes regulations only of automated decisionmaking, however, these regulations go well past their authorized scope.

Another tell that the regulations exceed the statutory mandate is that their definition of “automated decisionmaking” is out of step with how that term is used internationally. As noted, Europe recognizes that “automated decisionmaking” does not cover decisions that involve humans. Article 22 of the General Data Protection Regulation (“GDPR”), on “Automated Individual Decision-Making, Including Profiling” covers “decisions based *solely* on automated processing.”<sup>20</sup>

---

the scope of the regulations entirely. This further demonstrates that the definition of ADMT is overbroad and strays beyond the statutory mandate.

<sup>15</sup> Draft Regulations, § 7001, subd. (f)(4).

<sup>16</sup> Draft Regulations, § 7001, subd. (f)(4).

<sup>17</sup> *Solution 11918: Calculating and Graphing a Linear Regressions on the TI-83 Plus*, Texas Instruments Knowledge Base (accessed January 31, 2025), <https://education.ti.com/en/customer-support/knowledge-base/ti-83-84-plus-family/product-usage/11918>.

<sup>18</sup> Bobbitt, *How to Perform Linear Regression by Hand*, Statology (May 8, 2020).

<sup>19</sup> Patent No. US12086897, *Dynamic Optimized Reassignment of Providers at Geohash Level*, Applicant: Lyft, Inc., February 3, 2020,

<https://patentimages.storage.googleapis.com/ee/e5/49/b80dd99269e026/US12086897.pdf>; Patent No. US20200072622A1, *Determining Matches Using Dynamic Provider Eligibility Model*, Applicant: Lyft, Inc., February 3, 2020,

<https://patentimages.storage.googleapis.com/4a/3d/da/1a310f2e188a4a/US20200072622A1.pdf>.

<sup>20</sup> GDPR, art. 22 (emphasis added); see also GDPR, recital 71.

Similarly, the U.K. government, in its guidance on the U.K. version of the GDPR, explains that “automated decision-making is the process of making a decision by automated means *without any human involvement*.”<sup>21</sup> Brazil’s equivalent law similarly equates “automated decision[s]” with “decisions made solely based on automated processing.”<sup>22</sup> To interpret California’s law to extend to human decisionmaking using technology would be incongruous and wrong.

The proposal to regulate human decisionmaking – as opposed to an “automated decision” based “solely on automated processing” – thus exceeds the grant of authority that supports the regulations. The references to “executing” and “substantially facilitating” human decisions should be removed from the proposed regulations, and the regulations should be modified to exclude examples, like in Sections 7001(f)(4) and 7150(c)(1)–(2), that do not involve the making of decisions solely by automated technology.

## **B. There is no basis in the statute for keying the regulatory requirements off the overly broad category of “significant decisions”**

The proposed rules extensively regulate businesses that use automation to make any “significant decision,” which the Agency defines to include decisions without any connection to the privacy concerns that establish its authority to regulate here. The category of “significant decisions” is instead defined to cover much of the economy with no privacy tether at all: any decision “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).”<sup>23</sup> When a business uses automation to make a significant decision as the proposed regulations define that term, it must conduct a risk assessment, issue a pre-use notice, and (unless it meets certain exceptions) offer consumers the right to opt out of ADMT and a right of access.

The throughline across these supposedly “significant” decisions is plainly not privacy (and the regulation barely purports to have that theme); it is that these decisions arguably involve a socially important industry. For example, the regulations would govern remote software used to proctor a college-admissions test that processes a consumer’s IP address. Examples like this are covered

---

<sup>21</sup> Information Comm’r’s Off., What is Automated Individual Decision-Making and Profiling? (accessed Jan. 31, 2025),

<https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/individual-rights/automated-decision-making-and-profiling/what-is-automated-individual-decision-making-and-profiling/>.

<sup>22</sup> Lei Geral de Proteção de Dados (LGPD), Art. 20, Official Journal of the Brazilian Government (August 14, 2018).

<sup>23</sup> Draft Regulations, § 7220, subd. (a)(1).



because the Agency considers educational admissions to be important, not because they implicate privacy concerns in any real sense.

But there is no basis in the statute to have these sweeping requirements turn merely on whether a decision is “significant,” without any tether to the statute’s focus on data privacy and security. The CCPA is a *privacy* law, not an all-purpose regulator of automation applications perceived to be socially important. Prop. 24 was titled the “California *Privacy* Rights Act.”<sup>24</sup> And the resulting law is about data privacy from top to bottom. The law mentions “privacy,” “security,” and “personal information” more than 500 times, but “automated decisionmaking” only once, in a single sentence.<sup>25</sup> That sentence is one subsection of one subsection out of Prop. 24’s 31-section, over-20,000-word ballot initiative.<sup>26</sup> It is implausible that in this single sentence, California voters intended to authorize a new legal framework for regulating automated decisionmaking entirely disconnected from privacy concerns.<sup>27</sup> There is nothing in the CCPA to support the idea that the agency is now empowered to enforce it as a general consumer-protection or anti-discrimination statute.<sup>28</sup>

The CPRA’s enactment history further confirms what was (and was not) on California voters’ minds when they approved Prop. 24. As the public debated the law, the *only* concerns presented to them involved privacy and security.<sup>29</sup> The ballot guide explained that Prop. 24 sought to “amend[] consumer privacy laws.”<sup>30</sup> The Attorney General’s official summary promised that the

---

<sup>24</sup> Prop. 24, as approved by voters, Gen. Elec. (Nov. 3, 2020).

<sup>25</sup> Draft Regulations.

<sup>26</sup> Prop. 24, as approved by voters, Gen. Elec. (Nov. 3, 2020).

<sup>27</sup> Indeed, these other concerns are already being addressed by other agencies. The California Civil Rights Department has issued its own proposed regulations concerning the use of “automated-decision systems” in potentially discriminatory ways. (Second Modifications to Initial Text of Proposed Modifications to Employment Regulations Regarding Automated-Decision Systems (Civil Rights Council, Jan. 27, 2025), <https://calcivilrights.ca.gov/wp-content/uploads/sites/32/2025/02/Second-Modifications-to-Proposed-Modifications-to-Employment-Regulations-Regarding-Automated-Decision-Systems.pdf>.)

Such regulations are best left to an agency which has the authority and competence to address discrimination and fairness. The regulations should be narrowed to focus the opt-out right on factors that relate to privacy and security.

<sup>28</sup> It is also no answer that subsection (a)(15) authorizes the Agency to regulate automated decisionmaking. That subsection is prefaced and cabined by section 1798.185, subd. (a), which requires all regulations to “further the purposes of this title.” As we have explained, those purposes all relate to privacy. By contrast, Prop. 24’s “purpose and intent” section does not mention automation or AI even once. Thus, subsection (a)(15) authorizes the agency to regulate automated decisionmaking as necessary to promote data privacy and security. It does not grant a freestanding power to regulate ADMT unmoored from those concerns.

<sup>29</sup> California Secretary of State, Official Voter Information Guide, November 3, 2020, pp. 66–71, <https://vig.cdn.sos.ca.gov/2020/general/pdf/complete-vig.pdf>.

<sup>30</sup> *Id.* at p. 66.

law would let consumers “prevent businesses from sharing personal information,” “correct inaccurate personal information,” and “limit businesses’ use of sensitive personal information.”<sup>31</sup> It also explained that the Agency would “enforce and implement consumer privacy laws.”<sup>32</sup> The Legislative Analyst added that Prop. 24 would “change[] existing consumer data privacy laws” and “provide new consumer privacy rights” concerning the “*sharing* of personal data” and “use of ‘sensitive’ personal data.”<sup>33</sup> He also noted that the CPPA’s authority to “develop[] . . . new regulations” encompassed the power to pass “rules for correcting consumer personal data.”<sup>34</sup> And the arguments for and against Prop. 24 focused exclusively on whether the law would “protect . . . personal information” and how it would impact “privacy rights.”<sup>35</sup>

By contrast, automated decisionmaking and artificial intelligence were not on anyone’s radar. The terms “automated decisionmaking” and “artificial intelligence” do not appear even once in any of the ballot-initiative materials that accompanied Prop. 24.<sup>36</sup> Nor did the Legislative Analyst discuss regulating ADMT, much less for decisions involving non-sensitive information. The complete absence “of such a goal . . . [from the] ballot materials” is a strong tell that the law did not enact it.<sup>37</sup> Indeed, “[i]f this quite significant consequence were consistent with the most reasonable understanding of Proposition [24]’s purpose . . . one would assume there would be some mention of such a goal elsewhere in Proposition [24].”<sup>38</sup> “[E]nactors do not ‘hide elephants in mouseholes.’”<sup>39</sup> And here, that simply cannot be a sound principle of statutory interpretation; Prop. 24’s drafters were forbidden from wedging a comprehensive AI bill into their privacy statute. Under California law, “[a]n initiative measure embracing more than one subject may not be submitted to the electors or have any effect.”<sup>40</sup>

It comes as little surprise then that even the primary advocate for and drafter of Prop. 24, Alastair Mactaggart, has also commented on how the draft regulations have improperly strayed from the privacy mandate.<sup>41</sup> At the November 8, 2024 CPPA board meeting, for instance, Mactaggart

---

<sup>31</sup> *Ibid.*

<sup>32</sup> *Ibid.*

<sup>33</sup> *Id.* at pp. 67–68.

<sup>34</sup> *Id.* at p. 68.

<sup>35</sup> *Id.* at pp. 7071.

<sup>36</sup> Official Voter Information Guide.

<sup>37</sup> *Cal. Cannabis Coalition v. City of Upland* (2017) 3 Cal.5th 924, 940.

<sup>38</sup> *Ibid.*

<sup>39</sup> *Ibid.*

<sup>40</sup> Cal. Const. art. II, § 8(d); see, e.g., *Cal. Trial Lawyers Assn. v. Eu* (1988) 200 Cal.App.3d 351, 359–360 (provision regulating insurers’ campaign contributions was not related to the initiative’s subject of “spiralling insurance costs”).

<sup>41</sup> Nov. 8 CPPA Bd. Hr’g Tr. pp. 99–103.



reminded the Agency that “we should focus on our privacy mandate” after explaining how the draft regulations exceed their authorized scope.<sup>42</sup>

A comparison to Europe’s GDPR also shows why the statute does not authorize the regulation of decisions based solely on their “significance.” As we noted in Part I.A, there is some overlap in the language between Prop. 24 and the GDPR. For example, both laws regulate automated decisionmaking – an indicator that the concept should have similar meaning in both jurisdictions. But the converse is also true: When Prop. 24 conspicuously failed to borrow a certain aspect of the GDPR, that is evidence the voters did *not* intend to import this facet of the European regulations. In this vein, it is telling that, whereas the GDPR regulates the use of automated decisionmaking to make “significant[.]” decisions, Prop. 24 omitted that phrasing from its provision concerning automated decisionmaking, instead keeping the focus on the narrower domain of privacy.<sup>43</sup> Given that the California voters specifically declined to import the “significance” framework, it would be inappropriate for the implementing regulations to reverse course and do just that.

Because the regulations turn on the broad category a business decision falls into, not the degree to which (or even whether) the decision implicates privacy, they are inconsistent with the privacy rationale explicitly stated in Prop. 24 and approved by the voters. And when coupled with the overly broad definition of ADMT, these regulations cover an astoundingly large swath of the economy that Prop. 24 could not have plausibly meant to regulate. The proposed rules plainly exceed their authorization in the CCPA and must instead be revised to cover only decisions with a significant privacy impact.

**C. The provisions limiting how a business can advertise to its own customers based on existing data are not authorized by and are inconsistent with the statute**

The draft regulations impose far-reaching and unauthorized obligations on first-party “behavioral advertising.” The regulations put a raft of requirements – extensive disclosures, burdensome evaluations, and mandatory opt-out rights – on businesses that engage in so-called “extensive profiling,” which, contrary to the plain meaning of those words, is defined to encompass *all* personalized advertising, including advertising based on data a business *already has* through its own transactions with its customers.<sup>44</sup> All these requirements may apply to, for example, a retailer that recommends cleaning supplies to a customer who previously bought them, at a point when

---

<sup>42</sup> *Id.* at p. 106.

<sup>43</sup> See Prop. 24.

<sup>44</sup> Draft Regulations, § 7001, subd. (g) (“‘Behavioral advertising’ means the targeting of advertising to a consumer based on the consumer’s personal information obtained from the consumer’s activity . . . *within the business’s own distinctly-branded websites, applications, or services.*”) (emphasis added).

those supplies may be running low. But the CCPA does not authorize the extensive regulation of this benign conduct; indeed the voters consciously drew a line between such first-party advertising, which they allowed, and cross-context behavioral advertising, which they explicitly gave consumers the right to opt out of.<sup>45</sup>

Indeed, when voters amended the CCPA, they directly addressed the question of how to regulate advertising, leaving no room for the proposed rules. The CCPA, as enacted by the legislature, permitted businesses to use consumers' personal information for advertising and marketing, and gave consumers the right to opt out only from their data being sold to third parties.<sup>46</sup> Prop. 24 expanded that opt-out right to cover both the "selling" and "sharing" of personal information. It specifically identified "cross-context behavioral advertising" – that is, advertising "based on the consumer's personal information obtained from the consumer's activity *across businesses, distinctly-branded websites, applications, or services*" – as a type of "sharing."<sup>47</sup> So while Prop. 24 provided a right to opt out of cross-context behavioral advertising, it did not impose any comparable restrictions on first-party advertising.

Prop. 24's preamble and legislative history further underscore the voters' intent to regulate third-party advertising only. The preamble indicates that voters were focused on the selling or sharing of their personal information with other businesses.<sup>48</sup> The Legislative Analyst confirmed that one of the key rights created by Prop. 24 was to limit the "*sharing* of personal data."<sup>49</sup> Similarly, in describing why Prop. 24 added the concept of "sharing" data and created opt-out rights for "cross-context behavioral advertising," Mactaggart explained that Prop. 24 made it "crystal-clear, when it comes to sharing consumer information for cross context behavioral advertising, that the law gives consumers the right to opt out."<sup>50</sup> On the other hand, he noted that "*first-party data the business has can be used in any way that the business wants with that consumer.*"<sup>51</sup> That was the fundamental balance struck by Prop. 24: consumers were given a right to opt out of *third-party* targeted advertising, but businesses maintained the ability to engage in *first-party* advertising – that is, to advertise to consumers based on information gathered as part of a business's own

---

<sup>45</sup> Draft Regulations, § 7001, subd. (g).

<sup>46</sup> Cal. Assem. Bill No. 375 (2017–2018 Reg. Sess.); Civ. Code, § 1798.140, subd. (d)(4).

<sup>47</sup> Civ. Code, § 1798.140, subds. (k), (ah)(1).

<sup>48</sup> Prop. 24, as approved by voters, Gen. Elec. (Nov. 3, 2020), § 2.I ("Consumers should have the information and tools necessary to limit the use of their information to non-invasive, pro-privacy advertising, where their personal information is not sold to or shared with hundreds of businesses they've never heard of, if they choose to do so.").

<sup>49</sup> California Secretary of State, Official Voter Information Guide, November 3, 2020, pp. 66–71.

<sup>50</sup> Davis + Gilbert LLP, *Alastair Mactaggart's Privacy Perspective: Past, Present and Where We're Headed* (2022), <https://www.mondaq.com/unitedstates/data-protection/1183432/alastair-mactaggarts-privacy-perspective-past-present-and-where-were-headed>.

<sup>51</sup> *Ibid.*

relationship with a consumer. In adding opt outs and burdensome requirements for first-party advertising, the proposed regulations are fundamentally at odds with the voters' intent in approving Prop. 24.

Nor does the mere use of the word “profiling” in the statute justify the scope of the proposed regulations. In explaining its expansive definition of that word, the Agency points to various other state statutes that also regulate “profiling.” But each of these laws – like the CCPA and Prop. 24 – treats profiling and advertising as distinct concepts. Each law creates a right to opt out of profiling in some circumstances.<sup>52</sup> And then each law handles advertising with *separate* statutory language, reflecting the universal understanding that “advertising” and “profiling” are distinct practices.<sup>53</sup> (And in turn, the “advertising” proscriptions in these statutes unflaggingly cover only “targeted advertising” – a term, much like “cross-context behavioral advertising” in Prop. 24, defined to exclude first-party advertising.)<sup>54</sup> It is precisely because Prop. 24 was enacted against a legal background in which “profiling” did not cover “advertising” that Prop. 24 needed to separately address advertising. And when it did, it explicitly carved out first-party advertising from opt-out rights.<sup>55</sup>

The Agency has no authority to include first-party advertising in the draft regulations and should remove all references to first-party behavioral advertising.

---

<sup>52</sup> See, e.g., Va. Code Ann., § 59.1-577, subd. (5)(iii) (providing the ability to opt out of “profiling in furtherance of decisions that produce legal or similarly significant effects concerning the consumer”); Colo. Rev. Stat. Ann., § 6-1-1306, subd. (1)(a)(I)(C) (similar); Conn. Gen. Stat. Ann., § 42-518, subd. (a)(5)(c) (similar); Del. Code Ann. tit. 6, § 12D-104(a)(6)(c) (similar); Fla. Stat., § 501.705, subd. (2)(e)(3) (similar); Ind. Code, § 24-15-3-1, subd. (b)(5)(C) (similar).

<sup>53</sup> See, e.g., Va. Code Ann., § 59.1-577, subd. (5)(i) (providing the ability to opt out of “targeted advertising”); Colo. Rev. Stat. Ann., § 6-1-1306, subd. (1)(a)(I)(A) (similar); Conn. Gen. Stat. Ann., § 42-518, subd. (a)(5)(A) (similar); Del. Code Ann. tit. 6, § 12D-104, subd. (a)(6)(a) (similar); Fla. Stat., § 501.705, subd. (2)(e)(1) (similar); Ind. Code § 24-15-3-1(b)(5)(A) (similar).

<sup>54</sup> See, e.g., Va. Code Ann., § 59.1-575 (“[t]argeted advertising’ does not include . . . [a]dvertisements based on activities within a controller’s own websites or online applications”); Colo. Rev. Stat. Ann., § 6-1-1303, subd. (25) (similar); Conn. Gen. Stat. Ann., § 42-515, subd. (39) (similar); Del. Code Ann. tit. 6, § 12D-102, subd. (33) (similar); Fla. Stat., § 501.702, subd. (33) (similar); Ind. Code, § 24-15-2-30 (similar).

<sup>55</sup> The Federal Trade Commission distinguishes between first-party data use and third-party data sharing as well, singling out the latter for enforcement. See, e.g., *In re Gateway Learning Corp.* (July 7, 2004), FTC No. 042-3047,

<https://www.ftc.gov/sites/default/files/documents/cases/2004/07/040707agree0423047.pdf>;

*In re Chitika, Inc.* (Mar. 14, 2011), FTC No. 1023087,

<https://www.ftc.gov/sites/default/files/documents/cases/2011/03/110314chitikaagree.pdf>.

**D. The regulations related to “physical or biological identification or profiling” are unauthorized**

The draft regulations seek to impose multiple unwarranted requirements on “physical or biological identification or profiling.” The regulations define “physical or biological identification or profiling” to mean “identifying or profiling a consumer using information that depicts or describes their physical or biological characteristics, or measurements of or relating to their body.”<sup>56</sup> A business who uses “physical or biological identification or profiling” for a “significant decision” or “extensive profiling” must “conduct an evaluation” of its “identifying or profiling to ensure that it works as intended” and “does not discriminate”; and “must implement policies, procedures, and training to ensure” that the “identifying or profiling works as intended.”<sup>57</sup> The regulations would grant consumers a complete right to opt out of the use of their personal information for any training of ADMT that is capable of being used “for physical or biological identification or profiling.”<sup>58</sup> These regulations are incompatible with the statute.

To start, although the Agency has apparently proposed these regulations under its Subsection (a)(15) power to regulate “access and opt-out rights with respect to a business’ use of automated decisionmaking technology, including profiling,” the regulations fly past this grant of authority in two ways. For one thing, they regulate far more than access and opt-out rights. They set substantive criteria that “identification or profiling” must satisfy and compel testing and quality-assurance procedures. There is no basis for this substantive aspect of the regulations. The regulations also exceed the statutory requirement that they concern “automated decisionmaking technology, including profiling.” The regulations cover, in addition to profiling, the mere “identifying” of a consumer using biometrics.<sup>59</sup> “Identifying” is not “profiling.”<sup>60</sup> The draft

---

<sup>56</sup> Draft Regulations, § 7001, subd. (gg).

<sup>57</sup> Draft Regulations, § 7201.

<sup>58</sup> Draft Regulations, §§ 7200, subd. (a), 7221, subd. (a)–(b).

<sup>59</sup> No other comprehensive state law includes “identification” in the definition of “profiling.” See, e.g., Va. Code Ann., § 59.1-575 (“‘Profiling’ means any form of automated processing performed on personal data to evaluate, analyze, or predict personal aspects related to an identified or identifiable natural person’s economic situation, health, personal preferences, interests, reliability, behavior, location, or movements.” Identification does not fall under this definition because identification does not require businesses to “evaluate, analyze, or predict . . . personal aspects” like “health” or “personal preferences,” but rather to verify or confirm one’s identity.); Ind. Code, tit. 24, § 24-15-2-23 (defining profiling as “solely” automated processing but similarly excluding “identification” because it is not an “evaluat[ion], analy[sis], or predict[ion] relating to “personal aspects” like “health records,” “interests,” or “movements”).

<sup>60</sup> It does not appear that the Agency has tried to justify this regulation under the authority to regulate “automated decisionmaking.” And for good reason: identification does not entail making a decision. When an online grocery store uses a scanner to check the ID of someone buying medicine, or a college’s anti-cheating software automatically verifies the student ID of a remote exam taker, to say that anyone

regulations define “profiling” as processing personal information to “analyze or predict aspects concerning that natural person’s intelligence, ability, aptitude, performance at work, economic situation; health, including mental health; personal preferences, interests, reliability, predispositions, behavior, location, or movements”<sup>61</sup> – in short, predicting someone’s behavior or personal characteristics. Someone’s identity, however, is not a behavior or characteristic. Other parts of the CCPA bolster this distinction between “identifying” and “profiling.” For example, the CCPA grants consumers the right to opt out of certain uses of their biometric information, but not if a business has collected this information “without the purpose of inferring characteristics about a consumer.”<sup>62</sup> And even the portion of the regulations ostensibly directed at “profiling” exceeds the statutory limit. The statute authorizes at most a right for consumers to opt out of having their data used to profile them – not the right created by the regulations, a right to opt out of having their data used merely to train a technology that theoretically could be used to profile other people.<sup>63</sup>

The regulations also conflict with the statute by erecting a confusing scheme for regulating biometric information that competes with a different one already created by the statute. The statute already defines a category called “sensitive personal information,” which includes “the processing of biometric information for the purpose of uniquely identifying a consumer.”<sup>64</sup> The statute then guarantees consumers the right to limit the use of their sensitive personal information.<sup>65</sup> But this right is highly qualified. Consumers cannot opt out of businesses’ using their data to “improve, upgrade, or enhance the service[s]” they offer.<sup>66</sup> The statute also authorizes additional rules qualifying this right of consumers in order to protect the “legitimate operational interests of businesses.”<sup>67</sup>

The draft regulations conflict with this carefully balanced scheme. For example, under the draft regulations, a user may opt out of the use of her biometric data to “improve [a business’s] algorithm.”<sup>68</sup> This is irreconcilable with the statute’s express safe harbor allowing businesses to use sensitive personal information to improve the services they offer. And putting this specific glaring conflict aside, given that the statute already lays out an approach to biometric regulation and does so using a specific statutory term, the statute cannot be plausibly read to authorize the

---

has made a “decision” would be strained. There has been no judgment or weighing of options; the identifications are no more a “decision” than when a calculator determines whether two values are equal.

<sup>61</sup> Draft Regulations, §7001, subd. (kk).

<sup>62</sup> Civ. Code, § 1798.121, subd. (d).

<sup>63</sup> Civ. Code, § 1798.185, subd. (a)(15).

<sup>64</sup> Civ. Code, § 1798.140, subd. (ae)(2).

<sup>65</sup> Civ. Code, § 1798.121.

<sup>66</sup> Civ. Code, §§ 1798.121, subd. (a), 1798.140, subd. (e)(8).

<sup>67</sup> Civ. Code, § 1798.185, subd. (a)(18)(C).

<sup>68</sup> Draft Regulations, §§ 7200, subd. (a), 7221, subd. (a)–(b).

Agency to define a new similar, overlapping term and design a separate scheme of rights associated with that term.<sup>69</sup>

The proposed regulations of “physical or biological identification or profiling” should therefore be removed. At the very minimum, “identification” and “identifying” should be deleted from the definition.

### **E. The “Pre-Use Notice” requirements are not authorized by the statute**

Even though the enabling provision authorizes “regulations governing *access* and *opt-out* rights” for automated decisionmaking, the proposed regulations invent an entirely new category of requirements.<sup>70</sup> Specifically, businesses engaged in ADMT must provide a “prominent and conspicuous” pre-use notice with extensive information, including: a “plain language explanation of the specific purpose for which the business proposes to use the automated decisionmaking technology”; an explanation of any exceptions to the right to opt out that the business relied on; “information about how the automated decisionmaking technology works,” such as the “logic,” “key parameters,” and “intended output” of the ADMT; and information about the role of humans in the decision.<sup>71</sup>

These mandated disclosures conflict with the CCPA. Not only does the statute nowhere mention them, it explicitly handles consumer notice differently. When discussing consumers’ right to “information about [an algorithm’s] logic,” the law specifically couches that right in terms of an “access” request rather than any sort of pre-use notification. Meanwhile, other parts of the law require businesses to give notice, in some form, of what personal information they collect and how it is used “at or before the point of collection”<sup>72</sup> – but as other parts of the regulations make clear, this flexible requirement can be satisfied by providing consumers with a link to a section of its

---

<sup>69</sup> Further illustrating that implausibility is that in addition to conflicting with the statute, the draft regulations conflict sharply with the existing regulations fleshing out limitations on the use of “sensitive information.” Under the existing regulations, businesses have the right to use sensitive information like biometrics to “verify or maintain” the quality of the business’s products and “improve, upgrade, or enhance” their service or device (§ 7027, subd. (m)). By contrast, under the draft regulations, a business may not use biometrics to “improve [its] algorithm” if a user opts out (Draft Regulations, §§ 7200, subd. (a), 7221 subd. (a)–(b)). It is inevitable that having two separate regulations of essentially the same activity will lead to conflicts like this – not to mention unsettle the expectations of businesses that have already invested money complying with the first set of regulations – which is further evidence the statute did not authorize that.

<sup>70</sup> Civ. Code, § 1798.185, subd. (a)(15).

<sup>71</sup> Draft Regulations, § 7220. While Civ. Code, § 1798.185, subd. (a)(15) authorizes the Agency to issue regulations requiring “meaningful information about the logic involved in those decisionmaking processes,” that is only in connection with “response[s] to access requests,” not a “pre-use notice.”

<sup>72</sup> Civ. Code, § 1798.100, subd. (a).



privacy policy.<sup>73</sup> Elsewhere, the CCPA does expressly require businesses to issue certain “prominent” disclosures, but notably not here.<sup>74</sup> The legislature and voters thus know how to create a “pre-collection” notice regime, and even created an intricate one. They chose not to authorize the Agency to create yet another.<sup>75</sup>

And for good reason. Especially given the scope of the regulations, users would be bombarded with the proposed pre-use notifications constantly. As detailed in Part II below, copious social-science research confirms that consumers are likely to suffer from this information overload. The California law, correctly interpreted, does not allow this anti-consumer result. The Agency has no authority to include a pre-use notice requirement in the draft regulations and should remove the requirement.

## II. The Regulations Are Not Supported by Substantial Evidence

Regulations must be reasonably necessary to implement the statute authorizing them,<sup>76</sup> and the proposed regulations are not. Although the draft regulations would impose unprecedented burdens on California businesses and consumers, there is not substantial evidence that they are necessary to effectuate the goals of the CCPA. Those goals, as we have noted, were explicit. Prop. 24 states that “the rights of consumers and the responsibilities of businesses should be implemented with the goal of strengthening consumer privacy, while giving attention to the impact on business and innovation.”<sup>77</sup> The proposed regulations advance many concerns unrelated to privacy and security while impeding innovative product development.

This is why Mactaggart, now a member of the CPPA’s board, has expressed concern about the “overreach” of the draft regulations,” that they “undermine[] privacy rather than protecting it,” and that they mandate obligations inconsistent with the “privacy and security” focus of the statute.<sup>78</sup> As explained more below, the overly burdensome demands of the regulations are likely to lead

---

<sup>73</sup> Cal. Code Regs., tit. 11, § 7012, subd. (f).

<sup>74</sup> Specifically, “prominent and robust” notice is required when a business transfers personal information to a third party as part of a “merger, acquisition, bankruptcy, or other transaction” and the third party “materially alters how it uses or shares the personal information.” (Civ. Code, § 1798.140, subds. (ad)(2)(C), (ah)(2)(C).) Third parties are permitted, but not required, to satisfy their notice obligations by “prominently and conspicuously” “providing the required information . . . on the homepage of its internet website.” (Civ. Code, § 1798.100, subd. (b)); Civ. Code, § 1798.130, subd. (a)(5)(C)).

<sup>75</sup> *Hamdan v. Rumsfeld*, (2006) 548 U.S. 557, 578 (“A familiar principle of statutory construction . . . is that a negative inference may be drawn from the exclusion of language from one statutory provision that is included in other provisions of the same statute.”).

<sup>76</sup> Gov. Code, § 11342.2 (“No regulation adopted is valid or effective unless . . . reasonably necessary to effectuate the purpose of the statute.”).

<sup>77</sup> Prop. 24, § 3, subd. (C)(1).

<sup>78</sup> Nov. 8 CPPA Bd. Hr’g Tr., pp. 99–103.

businesses to divert limited resources from effective privacy protections, resulting in a net reduction in actual privacy and security protections for consumers. As Mactaggart put it, “this just creates a regulatory burden that I think has a negative impact on privacy.”<sup>79</sup>

**A. There is no basis for regulating human decisionmaking merely because it is assisted by technology**

The Agency has not put forward substantial evidence to support its definition of ADMT, which imposes onerous requirements on uses of technologies that only “execute” or “substantially facilitate” decisions made by humans. California businesses have used algorithms, artificial intelligence, “regression analyses,” “computation,” and other technology to assist with human decisions for decades. As Mactaggart noted, the proposed “definition of ADM[T] includes the use of almost any computerized technology in a way that describes how humans have used computers for 30 or 40 years.”<sup>80</sup> Businesses have deployed these techniques to execute or inform countless “significant decisions” and instances of “extensive profiling” (as the regulations define those terms), and the use of this technology is essential to California’s economy.<sup>81</sup> Yet the Statement of Reasons does not cite any evidence that decisions executed by technology or substantially facilitated by technology put consumers at a heightened privacy or security risk and must be regulated.

Instead, the Statement merely notes that its definition of ADMT “is informed by other frameworks addressing the use of ADMTs,” including the Biden Administration’s now-rescinded Blueprint for an AI Bill of Rights, an EEOC guidance document, and an academic article that discusses government uses of ADMT.<sup>82</sup> These policy documents do not support the proposed definition, however, since none defines ADMT to include the mere “execution” or “substantial facilitation” of a human decision or contends that those activities present privacy concerns.<sup>83</sup> To the contrary, such a broad scope would put California out of step with other states, including Connecticut,<sup>84</sup>

---

<sup>79</sup> Nov. 8 CPPA Bd. Hr’g Tr., p. 106.

<sup>80</sup> Nov. 8 CPPA Bd. Hr’g Tr., p. 100.

<sup>81</sup> Additional longstanding practices now covered by these regulations include: use of software or programs derived from statistics or other data-processing techniques (§ 7001, subd. (f)(1)); a business’s use of a regression analysis to evaluate employee performances (§ 7001, subd. (f)(4)); a dating app’s provision of geolocation, ethnicity, and medical information from a consumer’s profile to its analytics service provider (§ 7150, subd. (c)(3)); a grocery store’s use of wifi tracking within its stores to observe consumer shopping behavior (§ 7150, subd. (c)(5)); an educational provider’s use of software that automatically screens a student’s work for plagiarism (§ 7220, subd. (d)(3)).

<sup>82</sup> California Privacy Protection Agency, Initial Statement of Reasons (hereafter ISOR), (July 2024) p. 14.

<sup>83</sup> ISOR at p. 14 n.64.

<sup>84</sup> Conn. Gen. Stat. Ann., § 42-518.

Delaware,<sup>85</sup> Indiana,<sup>86</sup> Montana,<sup>87</sup> Rhode Island,<sup>88</sup> Maryland,<sup>89</sup> Texas,<sup>90</sup> Florida,<sup>91</sup> Nebraska,<sup>92</sup> Tennessee,<sup>93</sup> and New Hampshire,<sup>94</sup> which all provide a right to opt out of profiling in furtherance of “solely” automated decisions. By producing no evidence of privacy harms stemming from the broader range of activities it seeks to cover, the Agency fails to justify the scope of its regulation.<sup>95</sup>

**B. There is no basis to define “significant decisions” and “extensive profiling” to cover everyday uses of technology that pose no privacy concerns**

The Statement of Reasons does not contain substantial evidence to support the regulations’ broad definitions of “significant decisions” or “extensive profiling.” In fact, the Statement contains no evidence that the far-reaching scenarios covered by these definitions present any risk to the privacy or security of personal information – much less “substantial evidence” that regulating ADMT in these contexts is necessary.

The Statement offers only high-level explanations for its sweep, without linking the categories the regulations would cover to real privacy concerns. For example, while the Statement cites a generalized concern about the “lack of consumer control over their personal information,”<sup>96</sup> it does not link this concern to examples of a “significant decision” or “extensive profiling,” and especially not to examples of first-party behavioral advertising. Nor does the Statement attempt to tie this putative privacy harm to any specific ADMT use (let alone the uses that the Agency characterizes as “significant”) or explain why the alleged harms are not adequately addressed by the CCPA and numerous sector-specific laws.<sup>97</sup>

---

<sup>85</sup> Del. Code Ann., tit. 6, §12D-104, subd. (a)(6)(c).

<sup>86</sup> Ind. Code, § 24-15-23.

<sup>87</sup> Mont. Code Ann., § 30-14-2808.

<sup>88</sup> 6 R.I. Gen. Laws, § 48.1-5, subd. (e)(4).

<sup>89</sup> Md. Code Ann. Com. Law, § 14-4605, subd. (b)(7)(iii).

<sup>90</sup> Tex. Bus. & Com. Code, § 541.001, subd. (24).

<sup>91</sup> Fla. Stat., § 501.702, subd. (25).

<sup>92</sup> Neb. Rev. Stat., § 87-1102, subd. (25).

<sup>93</sup> Tenn. Code Ann., § 47-18-3201, subd. (21).

<sup>94</sup> N.H. Rev. Stat. Ann., § 507-H:4, subd. (I)(e).

<sup>95</sup> During the November 8, 2024 CPPA board meeting, Mactaggart stated, “If a human is materially involved in a decision, no opt-out should be required. And . . . again, I think we should focus on our privacy mandate.” (Nov. 8 CPPA Bd. Hr’g Tr., p. 106–107.)

<sup>96</sup> ISOR at p. 60.

<sup>97</sup> See Civ. Code, §§ 1798.110, 1798.120. Consumer-privacy concerns are already addressed by existing sector-specific laws. See Health Insurance Portability and Accountability Act, 45 C.F.R., § 164.502; see also Fair Credit Reporting Act, 15 U.S.C., § 1681, subd. (b); see also Equal Employment Opportunity Commission, 29 C.F.R., § 1635.9.

Although a broader policy debate has recently emerged around the potential benefits and harms of fully automated decisionmaking and AI, this debate has not been principally focused on privacy concerns.<sup>98</sup> Rather, these technologies implicate fairness considerations and broader philosophical questions around the appropriate role of technology in everyday life. This discussion has tended toward the theoretical, emphasizing the potential harms to society if technology is left to its own devices – but with very few examples of real harms related to the Agency’s privacy-and-security mandate.<sup>99</sup>

A comparison to Europe’s GDPR helps underscore why the regulations here are inappropriately broad. The GDPR covers a broader range of applications (though even then, only with respect to *solely* automated decisions), but it does so in order to implement sweeping human-rights objectives. The GDPR frames its purposes in all-encompassing terms: to “serve mankind,” and protect all manner of “freedoms” and “fundamental rights,” ranging from “freedom of expression and information” to “diversity.”<sup>100</sup> And the GDPR is itself grounded in the European Union’s Charter of Fundamental Rights, which enshrines principles such as human dignity, nondiscrimination, and due process.<sup>101</sup> It is no surprise, then, that the GDPR covers all manner of decisions with a legal or similarly significant effect.<sup>102</sup> The CCPA, by contrast, was never meant to promote such a diverse array of human-rights or policy priorities, beyond privacy. It does not establish a comprehensive rights-based framework. As detailed above, it was enacted to enhance transparency, provide consumers with greater control over their personal information, and regulate how businesses collect, share, and sell that information.<sup>103</sup> And thus it cannot carry the weight that the draft regulations seek to put on it.

The references to “behavioral advertising” should be deleted, and as discussed in Part I.B, the regulations should be revised to cover only decisions with a significant privacy impact.

---

<sup>98</sup> Krupa and Brandstätter, *UK data reform nurtures innovation but ensures safeguards to ensure EU adequacy, officials say* (November 21, 2024), Mlex, <https://www.mlex.com/mlex/articles/2264157/uk-data-reform-nurtures-innovation-but-ensures-safeguards-to-ensure-eu-adequacy-officials-say> (on UK proposed reform); Kern, *Humans versus machines: Who is perceived to decide fairer? Experimental evidence on attitudes toward automated decision-making* (October 14, 2022), *Patterns*, Vol. 3, Iss. 10, <https://www.sciencedirect.com/science/article/pii/S2666389922002094>.

<sup>99</sup> Chakravorti, *AI’s Trust Problem* (May 3, 2024) *Harv. Bus. Rev.*, <https://hbr.org/2024/05/ais-trust-problem>.

<sup>100</sup> GDPR, recital 4.

<sup>101</sup> *Charter of Fundamental Rights of the European Union* (Dec. 7, 2000) O.J. (C 364).

<sup>102</sup> *Ibid.*

<sup>103</sup> See Prop. 24.

### C. There is no basis to support the burdensome pre-use notice and request-to-access requirements

Similarly, the detailed and burdensome disclosure obligations contained in the proposed regulations are not necessary to protect consumers' privacy or security.<sup>104</sup> To the contrary, substantial evidence demonstrates that mandating extensive "conspicuous" notices in the course of routine consumer interactions would *undermine* privacy and security by overwhelming consumers and leading them to tune out important disclosures. At the same time, the enormous compliance burden on businesses will be a headwind on innovation.

The Agency has not put forward any evidence that the pre-use notices or access rights will help consumers. The Statement's discussion of pre-use notices is bereft of any evidence justifying the invention of this requirement.<sup>105</sup> And its justification of the "request to access" regulations is nearly as sparse. On that score, the Statement points only to consumers' right to access how credit scores are calculated.<sup>106</sup> But discrete information about credit score calculations is a far cry from the detailed disclosures required here.

Worse still, the regulations are likely to backfire for consumers, because the pre-use notice requirements will result in a highly disruptive online experience. Given the staggering proposed coverage of the "automated decisionmaking" regulations, consumers would be bombarded with pre-use notifications constantly. And given the dense list of required information, the notices will be long. Businesses will need to pepper users with numerous detailed categories of information, ranging from the fine details of how the automation works (its "logic" and "parameters") to a non-generic (that is, long) explanation of the purpose behind the automation, to a list of rights.<sup>107</sup> What is worse, users *must* be presented with most of these details before they even interact with the business or product; this is not like a warning label on a microwave that they may exercise autonomy over whether to read. So it is inevitable that many users will be force-fed excessive information they do not want.

Abundant social science confirms the intuition that overloading consumers with this information will be bad for them. Studies show that forcing consumers to view "excessive information" will overwhelm them and "degrade the quality" of their choices.<sup>108</sup> One reason is that "mandated

---

<sup>104</sup> ISOR at pp. 85, 91–92.

<sup>105</sup> ISOR at pp. 83–86.

<sup>106</sup> ISOR at pp. 91–97 & nn. 141–143.

<sup>107</sup> Draft Regulations, § 7220, subd. (c).

<sup>108</sup> See Latin, *Good Warnings, Bad Products, and Cognitive Limitations* (1994) 41 UCLA L.Rev. 1193, 1214–15, <https://heinonline.org/HOL/LandingPage?handle=hein.journals/uclalr41&div=41&id=&page=>; see also Zheng et al., *How Causal Information Affects Decisions* (2020) 13 Cogn. Res. Princ. Implic.,

disclosure can crowd out useful information” and focus users on irrelevant considerations.<sup>109</sup> For example, an FTC study showed that a “proposed disclosure of brokerage fees” caused consumers to focus overly on those fees, and thus “overestimate the total cost of loans.”<sup>110</sup> Mandatory disclosures are also often too complicated for consumers to understand.<sup>111</sup> And the situation becomes even worse when disclosures accumulate across products: each decreases the effectiveness of every other one, as they “compete[] for . . . time and attention with [each other].”<sup>112</sup> “Even if [consumers] wanted to read all the disclosures relevant to their decisions, they could not do so proficiently,” and they will “soon learn their lesson and give up any inclination they may have had to devote their lives to disclosures.”<sup>113</sup> The upshot is that both the “use of encyclopedic warnings” and the “overuse of warnings” “may, in fact, decrease the effectiveness of all warnings.”<sup>114</sup> Excessive disclosures may also lead consumers to simply shut down and avoid interacting with covered businesses at all.<sup>115</sup>

Here, consumers will at best tune out the annoying barrage of similarly sounding pre-use notices they see every day, and at worst be distracted from the details they actually need to know, like the features and price of a product, the admissions criteria of a university, or an employer’s personnel policies. In no way will they benefit. Consider perhaps the closest analogy to the proposed disclosures, the now-ubiquitous cookie banner that websites display to comply with European regulations. The cookie banner has been a consensus failure for consumer privacy and empowerment, because Internet users have been so inundated with the disclosures that they simply disregard them.<sup>116</sup>

---

<https://pubmed.ncbi.nlm.nih.gov/32056060/> (documenting a psychological experiment showing that giving consumers certain “information can actually lead to worse decisions”); Dalley, *The Use and Misuse of Disclosure as a Regulatory System* (2007) 34 Fla. St. U. L.Rev. 1090, 1115,

<https://ir.law.fsu.edu/lr/vol34/iss4/2/> (describing “information overload” and how an excess of information can lead decisionmakers to make ill-informed decisions).

<sup>109</sup> Ben-Shahar and Schneider, *The Failure of Mandated Disclosure* (2011) 159 U.Penn.L.Rev. 647, 737, <https://www.jstor.org/stable/41149884>.

<sup>110</sup> Craswell, *Taking Information Seriously: Misrepresentation and Nondisclosure in Contract Law and Elsewhere* (2006) 92 Va. L.Rev. 565, 584, <https://virginialawreview.org/articles/taking-information-seriously-misrepresentation-and-nondisclosure-contract-law-and/>.

<sup>111</sup> Ben-Shahar and Schneider at pp. 665–672.

<sup>112</sup> *Id.* at p. 689.

<sup>113</sup> *Id.* at p. 690.

<sup>114</sup> Schwartz and Driver, *Warnings in the Workplace: The Need for a Synthesis of Law and Communication Theory* (1983), 52 U.Cin.L.Rev. 38, 43.

<sup>115</sup> See Craswell at p. 584; Accenture, *The Empowered Consumer* (2024), <https://www.accenture.com/us-en/insights/consulting/empowered-consumer> (finding that in a three-month period, three quarters of consumers “walked away from purchases simply because they felt overwhelmed” by information).

<sup>116</sup> See, e.g., Utz et al., *(Un)informed Consent: Studying GDPR Consent Notices in the Field* (2019), <https://arxiv.org/abs/1909.02638> (studying user behavior in reaction to cookie banners and noting the “[r]ecurring theme[]” “that the notices were ‘annoying . . . , so [users] just ignore them out of



The regulations will also be a costly drag on business. Generating the required disclosures for the pre-use notifications and access rights will be an exceedingly complex task. The proposed regulations require an explanation of “the output of the automated decisionmaking technology with respect to the consumer,” “the role the output played in the business’s decision and the role of any human involvement,” and “how the automated decisionmaking technology worked with respect to the consumer.”<sup>117</sup> These disclosures will apparently have to be individualized to each consumer. This poses an immense data-governance and retention challenge. Businesses will have to store detailed information regarding every single “significant decision” made using ADMT, and will have to build systems that can, upon request, parse that data to construct a usable individualized response. This is orders of magnitude more challenging than responding to a request to know or a request to correct, under California law, given the inherent complexity of automated processing. Despite that, the regulations do *not* provide any exceptions when compliance would involve “disproportionate effort” – even though similar exceptions exist for requests to correct, delete, or know.<sup>118</sup> Maintaining and processing this data for the entire range of “significant decisions” would necessarily stifle the innovative engines that drive California’s economy. But neither the Agency’s statement of reasons nor its economic analysis addresses these concerns.

And there are yet more reasons why the disclosures will hurt the public that the Statement does not grapple with. To start, the regulations would compel businesses to make statements that are confusing and even misleading. Disclosing the “logic” and “key parameters” of an ADMT in “plain language” may often be an impossible task. The most advanced AI models today have *billions* or even *trillions* of parameters. Their internal logic is just “a long list of numbers.”<sup>119</sup> Translating these numbers into human-understandable explanations is far from trivial.<sup>120</sup> The field

---

frustration”); O. Kulyk et al., *Has The GDPR Hype Affected Users’ Reaction to Cookie Disclaimers* (2020) 6 J. Cybersecurity, <https://academic.oup.com/cybersecurity/article/6/1/tyaa022/6046452> (studying web users’ behavior and concluding that “participants considered the cookie disclaimer as a nuisance” and so “tend[ed] to accept cookie disclaimers blindly to get rid of it”); M. Nouwens et al., *Dark Patterns after the GDPR: Scraping Consent pop-ups and Demonstrating their Influence* (2020), <https://dl.acm.org/doi/10.1145/3313831.3376321> (“[T]he frequency of the pop-ups caused frustration and consent fatigue.”).

<sup>117</sup>Draft Regulations, § 7222, subd. (b).

<sup>118</sup> Draft Regulations, §§ 7022, subd. (b)–(c), 7023, subd. (f), 7024, subd. (h).

<sup>119</sup> Anthropic, *Mapping the Mind of a Large Language Model* (May 21, 2024), <https://www.anthropic.com/research/mapping-mind-language-model>.

<sup>120</sup> See J. Woods, *Machine Learning Interpretability: New Challenges and Approaches* (Mar. 14, 2022)

Vector Institute, <https://vectorinstitute.ai/machine-learning-interpretability-new-challenges-and-approaches/>;

See generally R. Dwivedi, *Explainable Ai (XAI): Core Ideas, Techniques, and Solutions* (2023), 55 ACM Computing Surveys, <https://dl.acm.org/doi/10.1145/3561048>.

of research devoted to this task has made promising advances.<sup>121</sup> But even when sophisticated researchers get a handle on how an advanced AI model works, their explanations have been long and jargon-filled.<sup>122</sup> And researchers have struggled to convert these explanations into a form understandable by non-expert humans.<sup>123</sup> So in many circumstances, any “plain language” explanation of the model’s logic will be overly simplistic and misleading. It is never proper for the government to direct a business to mislead its customers.<sup>124</sup>

There is also ample reason to be concerned that such a disclosure regime could be misused to gain access to confidential business or consumer information. For example, it would be plainly inappropriate to compel the admissions office of a private college to disclose the “logic” and underlying “assumptions” of its admissions policy. A university may reasonably want to keep this information private, to prevent prospective students from gaming the system. But if a school implements or informs its admissions decisions in part using an automated system (as colleges fielding hundreds of thousands of applications necessarily will), it now may have to reveal exactly that confidential information.

Worse still, the disclosure requirements can be misused by malicious actors to gain unauthorized access to personal information. An unfortunately common scenario is that malicious actors use social engineering to obtain consumers’ login credentials for a service.<sup>125</sup> Under a compelled-

---

<sup>121</sup> See Anthropic, *supra*; K. Wang et al., *Interpretability in the Wild: A Circuit for Indirect Object Identification in GPT-2 Small* (Nov. 1, 2022), <https://arxiv.org/abs/2211.00593>.

<sup>122</sup> See, e.g., Wang, *supra* (twelve technical pages to explain how a large language model predicted a single word in a sentence).

<sup>123</sup> See H. Siu et al., *STL: Surprisingly Tricky Logic (for System Validation)*, (May 26, 2023), <https://arxiv.org/abs/2305.17258>.

<sup>124</sup> Cf. *Barton v. Neeley* (6th Cir. 2024) 114 F. 4th 581, 592 (explaining that the First Amendment protects the “right to decide what to say and what not to say, and accordingly, the right to reject governmental efforts to require [someone] to make statements he believes are false”), and *Massachusetts Ass’n of Priv. Career Sch. v. Healey*, 159 F. Supp. 3d 173, 199–200 (D. Mass. 2016) (holding that a regulation requiring a business to make misleading statements was subject to heightened scrutiny under the First Amendment).

<sup>125</sup> See, e.g., Pavur & Knerr, *GDPArrrrr: Using Privacy Laws to Steal Identities*, Blackhat USA (2019), <https://i.blackhat.com/USA-19/Thursday/us-19-Pavur-GDPArrrrr-Using-Privacy-Laws-To-Steal-Identities-wp.pdf> (noting that “social engineers can abuse right of access requests as a scalable attack vector for acquiring deeply sensitive information about individuals”); IBM, *IBM Security X-Force Threat Intelligence Index 2024* at p. 9, <https://www.ibm.com/reports/threat-intelligence> (noting that “the focus has shifted towards logging in rather than hacking in, highlighting the relative ease of acquiring credentials compared to exploiting vulnerabilities or executing phishing campaigns”); *Verizon 2023 Data Breach Investigations Report* (2023) at p. 8, <https://www.verizon.com/about/news/media-resources/attachment?fid=65e1e3213d633293cd82b8cb> (noting that “74% of all breaches include the human element, with people being involved either via Error, Privilege Misuse, Use of stolen credentials or Social Engineering”); Stahie, *Billions of Leaked Credentials Available on the Dark Web*, Bitdefender (2020) (noting 15 billion credentials available on the dark web), <https://www.bitdefender.com/en-us/blog/hotforsecurity/billions-of-leaked-credentials-available-on-the-dark-web>.

disclosure regime, an attacker with these stolen credentials may now be able to learn even more information about his victim by obtaining the inferences a business has made about her and use that ill-gotten information in furtherance of identity theft or targeted phishing attacks. In this way, the regulations may be more harmful to privacy than enhancing of it.

#### **D. There is no basis to require the onerous risk assessments**

The Statement does not contain substantial evidence demonstrating that the extremely detailed and burdensome risk assessments are necessary to further consumers' privacy. Per the statute, the purpose of the risk assessment is to evaluate which instances of data processing have elevated "risks to privacy."<sup>126</sup> But many of the activities that must be addressed by the risk assessment have no impact on privacy at all. For example, the draft regulations would require each business to discuss the "completeness, representativeness, timeliness, validity, accuracy, consistency, and reliability" of its information sources and the "logic" of certain algorithms. None of these requirements bears any relationship to privacy or security concerns. The Statement does not explain otherwise.

Not only is there no evidence that risk assessments are necessary to advancing privacy and security, but the overbroad compliance regime proposed here would undermine privacy and security.<sup>127</sup> The risk assessments must address dozens of discrete issues. Undertaking such an extensive assessment anytime ADMT is used for a broad category of "significant decisions" would be enormously resource-intensive. Companies throughout the economy would need to divert resources, including engineering talent, away from substantive risk mitigation and toward producing burdensome risk assessments with little relation to privacy or security. The Statement denies any tradeoff with the blanket statement that "risk assessments are cost effective."<sup>128</sup> But its only source discusses not the regulations here, but the burdens of complying with Europe's GDPR, an entirely different set of requirements. And even with respect to those requirements, the source does not support the point: it acknowledged that the cost of the GDPR's data-protection assessments may already be "prohibitive," particularly for smaller companies that otherwise could

---

<sup>126</sup> Cal. Civ. Code, § 1987.1785(a)(14)(b).

<sup>127</sup> Nov. 8 CPPA Bd. Hr'g Tr. 99 ("With respect to the risk assessments, I think these proposed regulations will make the inclusion criteria for risk assessments so broad that we will end up hurting the cause of privacy, not helping it. The scope of these regulations effectively mandates risk assessments for almost any business using software. This spread will hurt businesses and overwhelm our agency with, I think, largely form paperwork, diminishing our focus – our ability to focus on enforcement. There's no chance we'll be able to review tens and tens of thousands of multi-page risk assessments at this stage with our current resources.").

<sup>128</sup> ISOR, p. 71–72.

substantially benefit from automation.<sup>129</sup> The Agency must promulgate regulations that balance the enhancement of privacy with the promotion of innovation, and since the risk-assessment requirements would do little to improve privacy and stifle innovation, the significant cost imposed by risk assessments is unsupported and unnecessary.<sup>130</sup>

### **E. There is no basis for the rigid cybersecurity audit requirements**

The cybersecurity audit requirements are overly simplistic, in both when they apply and what they entail. The Statement of Reasons fails to show that the draft regulations' blunt requirements are necessary or appropriate.

The thresholds for when an audit is required are unjustified. The thresholds are based on blunt indicators, a business's revenue and number of consumers whose data is processed.<sup>131</sup> These simplistic conditions fail to account for how cybersecurity practices and the need for an audit vary across different industries. For example, strict compliance checklists may be appropriate for a mature institution with a predictable workflow, but counterproductive for a software company with a rapidly evolving product and headcount.<sup>132</sup> The draft regulations could lead to disproportionate compliance costs for businesses without lowering true risks to consumer security. The Statement does not address this concern.

---

<sup>129</sup> Iwaya et al., *Privacy Impact Assessments in the Wild: A Scoping Review* (2024), <https://www.sciencedirect.com/science/article/pii/S2590005624000225>.

<sup>130</sup> The risk assessments, as envisioned by the proposed regulations, also run afoul of the First Amendment. Courts have repeatedly rejected recent attempts to require disclosures about a company's use of technology and its opinions on whether and how this use maps to ambiguous and often pejorative characterizations. The Ninth Circuit made this point twice in just the last year while striking down remarkably similar California laws. In one case, the law demanded, akin to the present regulations, that certain website operators report on whether "the design of the[ir] online product . . . could harm children" in various specific ways. (*NetChoice v. Bonta* (9th Cir. 2024) 113 F. 4th 1101, 1109.) The requirement was invalid because it compelled "covered businesses to opine on potential harm" of their product outside the context of any specific transaction. In the other case, the State compelled businesses to "implicitly opin[e] on whether and how certain controversial categories of content should be moderated." (*X Corp. v. Bonta* (9th Cir. 2024) 116 F. 4th 888, 901.) Yet this request too was invalid, because the government had no authority to make a company offer "opinions about and reasons for" its policies. The only difference here is that there is nothing "implicit" about what the new regulation asks for. It flat-out tells companies to express an opinion on whether or not their technology fits within the vague and value-laden categories in the regulations and, if so, the merits and drawbacks of their own policies. But this is well past the range of speech that a government can legitimately compel.

<sup>131</sup> Draft Regulations, § 7120.

<sup>132</sup> Wallace, *The Importance of Cybersecurity by Industry*, <https://www.uscybersecurity.net/the-importance-of-cybersecurity-by-industry>; Cristiano and Prenio, *Regulatory approaches to enhance banks' cyber-security frameworks* (2017), <https://www.bis.org/fsi/publ/insights2.pdf>.

And when audits are required, the mandated components are problematically rigid. The particular approaches that work in one industry or for one particular size of business may backfire elsewhere.<sup>133</sup> Moreover, the detailed cybersecurity audit requirements set forth in the regulations – including dozens of discrete requirements – would, at best, introduce a box-checking exercise and, at worst, distract businesses from focusing on actually optimizing security and keeping sensitive information safe.<sup>134</sup>

### III. The Proposed Regulations Lack Clarity

Regulations must be easy to understand and follow,<sup>135</sup> and “due process also requires that regulations be written with sufficient clarity so that those subject to the law can understand what is required or prohibited.”<sup>136</sup> But complying with the proposed regulations will require herculean guesswork. The regulations leave California businesses to puzzle over whether and when the regulations apply and, if they do, how to comply.

First, the **definition of ADMT** is troublingly vague. The flexible terms “execute,” “substantially facilitate,” and “key factor” provide little guidance to businesses about what qualifies as ADMT. It may be difficult to assess whether a particular output of a technology plays a “substantial” or “key” role in a decision, particularly when the technology merely informs human decisionmaking; there may be no agreed-upon way to quantify the weight that a factor plays in a human decision. The examples only compound this indeterminacy. Section 7001(f)(2) states that ADMT “substantially facilit[es] human decisionmaking” when it is used “to generate a score about a consumer that a human reviewer uses as a primary factor to make a significant decision.” But in Section 7001(f)(4), the regulations indicate that using technology to “calculate” a “score that [a] manager will use to determine which [employee] will be promoted” is not even a use of ADMT. The regulation appears to discern between “generating a score” for the purpose of guiding a human

---

<sup>133</sup> Etoom, *Strategising cybersecurity: Why a risk-based approach is key* (2023), <https://www.weforum.org/stories/2023/04/strategizing-cybersecurity-why-a-risk-based-approach-is-key/>; Boehm et al., *The risk-based approach to cybersecurity* (2019), <https://www.mckinsey.com/capabilities/risk-and-resilience/our-insights/the-risk-based-approach-to-cybersecurity>.

<sup>134</sup> Marotta and Madnick, *Convergence and divergence of regulatory compliance and cybersecurity* (2021), [https://doi.org/10.48009/1\\_iis\\_2021\\_10-50](https://doi.org/10.48009/1_iis_2021_10-50) (“regulatory compliance can negatively affect cybersecurity”); Sjouwerman, *5 Reasons Why Compliance Alone Is Not Efficient at Reducing Cyber Risks* (2022), <https://www.corporatecomplianceinsights.com/compliance-not-enough-cybersecurity-risk/>; Internet Security Alliance, *Cyber Regulations Are Counter-Productive to True Security* (2021), <https://isalliance.org/cyber-regulations-are-counter-productive-to-true-security/>.

<sup>135</sup> Gov. Code, § 11349, subd. (c); *FCC v. Fox Television Stations, Inc.* (2012) 567 U.S. 239, 253 (same under Due Process clause).

<sup>136</sup> See *FCC v. Fox Television Stations, Inc.* (2012) 567 U.S. 239, 253.

decision and “calculating a score” for that same purpose, but without any meaningful explanation of how the two are different.

Section 7001(f)(4) likewise creates confusion as to what “technology” is in scope. It alternately says that “calculators,” “spreadsheets,” and “similar technologies” are *not* ADMT, then asserts that the “use of a spreadsheet to run regression analyses” *is* ADMT if used by humans evaluating job performance, but then says that it is *not* ADMT if it “merely . . . organize[s] human . . . evaluations.” As we noted above, the distinction between “regressions” and “calculators” is wholly unclear, and a business has little hope at guessing which side of the line its software falls on. The Agency’s attempt to explain the regulation only adds confusion because “the language of the regulation conflicts with the agency’s description of the effect of the regulation.”<sup>137</sup> These artificial distinctions underscore the unworkability and ambiguity of the proposed definition of ADMT.

Second, the term “**significant decision**” also lacks clarity. The specific categories that count as “significant” are problematically vague. For example, what does it even mean for a decision to “result[] in access to, or the provision or denial of . . . criminal justice”? The regulations do not say, beyond offering the single example of the “posting of bail bonds.” Suppose a security firm guarding a semiconductor factory uses an AI tool to decide which visitors must go through extra screening. Since the security screening could theoretically discover evidence of a crime and lead to a prosecution, does the company’s use of AI fit the definition? It is likewise unclear what decisions count as affecting “housing.” If a college assigns roommates using software that considers students’ personal preferences, does it have to conduct a risk assessment and offer an opt-out? Or does housing extend only to the purchase or lease of real property? And what counts as an “essential good or service”? The regulations provide a handful of examples (groceries, medicine, hygiene products, or fuel) but what else should be considered “essential” and how is that decided? Is Internet access essential? Cultural opportunities? Firearms? And even if a good is unequivocally “essential,” which decisions affect “access” to it? Do the regulations cover every single transaction related to that good (for example, a grocery store’s denying a consumer access to one particular foodstuff on one occasion)? Or does a decision count only when it wholesale excludes a consumer from the good (like if the only utility company that services a consumer’s home disconnects the power)? The proposed definition of “significant decision” creates more questions than it answers.

---

<sup>137</sup> Office of Administrative Law, OAL Review for Compliance with the Six Substantive Standards of the Administrative Procedure Act, § 3.03 (Apr. 2023), <https://oal.ca.gov/wp-content/uploads/sites/166/2023/04/OAL-Review-for-6-APA-Standards.pdf>.



Third, the proposed **pre-use notice** and **right to access** regulations likewise fail to explain how those disclosures would function. The pre-use notice and any response to a request to access may not “describe the purpose in generic terms” and must include information about the logic, key parameters, and output of the ADMT, which must be in “plain language.” But, as discussed above, automated decisionmaking technology, including artificial intelligence, often involves dynamic and constantly evolving, highly technical systems that can consider hundreds of inputs of variable weights that lead to a range of different outputs. And businesses may be constantly tweaking and testing their technology to optimize for different circumstances or to account for changes in the marketplace. And as discussed above, translating any given iteration of an ADMT system into plain English may be an impossible task. The regulations provide no guidance on how to provide accurate and digestible information given this highly complex backdrop.

\* \* \*

We appreciate this opportunity to share some of our concerns with the Agency and hope that the Agency will revise the proposed regulations to focus on the privacy and security concerns expressed by the People of California in approving Prop. 24.

Respectfully submitted,




---

Ashlie Beringer  
Partner  
Co-Chair of the Privacy, Cybersecurity and Data Innovation Practice Group



---

Jane Horvath  
Partner  
Co-Chair of the Privacy, Cybersecurity and Data Innovation Practice Group



---

Cassandra Gaedt-Sheckter  
Partner  
Co-Chair of the Artificial Intelligence Practice Group