

COMMENTS OF THE ELECTRONIC PRIVACY INFORMATION CENTER,

CONSUMER FEDERATION OF AMERICA, AND FAIRPLAY

to the

COLORADO DEPARTMENT OF LAW

Consumer Protection Section

On Automated Decision-Making Technology and Chatbot Safety Act Rulemaking

October 26, 2026

The Electronic Privacy Information Center (EPIC), Consumer Federation of America (“CFA”), and Fairplay submit these comments to the Colorado Department of Law in response to the proposed draft rules for Senate Bill 26-189 (“the “Automated Decision-Making Technology Act” or “ADMT Act”) and House Bill 26-1263 (“Chatbot Safety Act”).¹

EPIC is an independent, nonpartisan public interest research center in Washington, D.C., established in 1994 to protect privacy, freedom of expression, and democratic values in the digital age.² CFA is an association of non-profit consumer organizations based in Washington, D.C. that was established in 1968 to advance the consumer interest.³ Founded 25 years ago, Fairplay is a national, independent nonprofit organization dedicated to protecting children from unfair commercial practices in media and digital environments, where it is a nationally recognized expert.⁴

¹ Colorado Attorney General’s Office, Colorado Automated Decision-Making Technology & Chatbot Safety Rulemaking, <https://coag.gov/ai/>.

² EPIC, *About EPIC* (2026), <https://epic.org/about/>.

³ CFA, *About CFA* (2026), <https://consumerfed.org/about-cfa/>.

⁴ Fairplay, *About* (2026), <https://fairplayforkids.org/mission-impact/>.

For decades, EPIC has worked under different administrations to safeguard individuals' personal data and to protect consumer privacy. EPIC is a leading voice in advocating for protecting privacy and civil rights from harms of artificial intelligence (AI) systems, including automated decision-making technology (ADMT), especially when deployed to make consequential decisions.⁵ CFA has long advocated at both state and federal levels for commonsense regulations on automated decision-making tools that impact investor protections, junk fees, job opportunity, housing opportunities, and more, urging transparency, bright-line rules, and meaningful accountability.⁶ EPIC, CFA, and Fairplay have also been heavily involved in shaping chatbot policy at both the state and federal level, including developing a model chatbot bill that has been considered by seven states so far in 2026 and was recently introduced in Congress.⁷

These comments are in response to the October 6, 2026, interim draft rules ("Updated Proposed Rules") issued by the Colorado Department of Law ("the Department"). EPIC, CFA, and Fairplay commend the development of detailed rules that clarify responsibilities for covered entities and how consumers can hold them accountable when they believe they are harmed by an entity's deployment of an ADMT or a chatbot. EPIC, CFA, and Fairplay have several recommendations to ensure broad coverage of ADMT deployed in consequential decisions, economic incentives that encourage deployers to act responsibly, and ensuring that consumers receive actionable information and clear steps to make sure that they are assessed fairly and

⁵ *Commercial AI Use*, EPIC, <https://epic.org/issues/ai/commercial-ai-use/>; *Screened & Scored in D.C.*, EPIC (Nov. 2022), <https://epic.org/wp-content/uploads/2022/11/EPIC-Screened-in-DC-Report.pdf> [hereinafter "Screened & Scored"]; *Assessing the Assessments: Maximizing the Effectiveness of Algorithmic & Privacy Risk Assessments*, EPIC (June 25, 2025), <https://epic.org/wp-content/uploads/2025/06/Assessing-the-Assessments-Report.pdf> [hereinafter "Assessing the Assessments"].

⁶ See, e.g., *CFA Letter to CFPB on the need for 'Less Discriminatory Algorithms,'* CFA (June 2024), <https://consumerfed.org/news/testimony-comments/urgent-call-to-cfpb-for-regulatory-clarity-on-the-need-to-search-for-and-implement-less-discriminatory-algorithms/>; *CFA Testimony in Support of the New York AI Act*, CFA (Jan. 2026), <https://www.nysenate.gov/sites/default/files/admin/structure/media/manage/filefile/a/2026-01/ben-winters-consumer-federation-of-america-testimony.pdf>.

⁷ See, e.g., EPIC, *EPIC Applauds Introduction of Privacy-Centered Federal Chatbot Bill* (July 10, 2026), <https://epic.org/press-release-epic-applauds-introduction-of-privacy-centered-federal-chatbot-bill-2/>; *People-First Chatbot Act*, 119th Cong. (2d Sess. 2026), https://foushee.house.gov/imo/media/doc/final_people-first_chatbot_act.pdf; *Chatbots*, EPIC, <https://epic.org/issues/ai/chatbots/>; *People-First Chatbot Bill*, EPIC, Consumer Federation of America & Fairplay (Dec. 2025), <https://epic.org/people-first-chatbot-bill/>.

accurately in consequential decisions. For the proposed rules for the Chatbot Safety Act, because EPIC, CFA, and Fairplay do not believe rulemaking can address the law’s weaknesses sufficiently to make it an effective chatbot regulation, our comments below focus on strengthening the privacy protections in the age assurance provisions.

I. ADMT Act Rules

With respect to the ADMT rules, EPIC, CFA, and Fairplay support maintaining the rebuttable presumption of an ADMT’s material influence on a consequential decision and urges the threshold to be kept high to avoid allowing businesses to opt themselves out of regulation to the detriment of consumers interests. We also recommend reinstating the requirement for covered entities to provide the principal reason for an adverse decision, reverting the language on Meaningful Human Review to include requirements of independent review, and removing considerations for a Commercially Reasonable determination that create perverse economic incentives for covered entities to avoid allocating resources to conduct reconsideration and meaningful human review. Lastly, we urge the Department to reverse the changes that allow deployers to point the finger at developers and third parties to satisfy its responsibilities owed to consumers affected by ADMTs.

A. The Rebuttable Presumption Threshold for an ADMT’s Material Influence Should Remain Difficult to Overcome.

Responsive to Rule 2.2 “De Minimis Factor” Definition

EPIC, CFA, and Fairplay support the inclusion of a rebuttable presumption in determining whether an ADMT output had a material influence on a consequential decision. It is logical to presume that a deployer purchased an automated *decision-making* tool to materially influence decisions. It is appropriate to put the burden on a deployer to prove otherwise, and the presumption should be sufficiently difficult to overcome in order to prevent risks to consumers and ensure that the rules cover the uses of AMDTs intended by the General Assembly.

Particularly important in the definition of “De Minimis Factor” in Rule 2.2 is factor 1, which asks whether the consequential decision maker reviewed the ADMT output *before* making the consequential decision. Individuals tend to favor and accept suggestions from ADMTs—a

long-studied phenomenon known as “automation bias.”⁸ If a deployer can simply claim that an ADMT had a de minimis impact on a decision by involving a human to review (and likely accept) the output of the ADMT, very few systems will be covered by these rules.⁹

We urge the Department to retain the rebuttable presumption in Rule 6.1 and the related “De Minimis Factor” definition in Rule 2.2.

B. Deployers Should be Required to Disclose a Principal Reason for Adverse Decisions.

Responsive to Rule 7

The Updated Proposed Rules strike all references to providing a principal reason for decisions made based on a covered ADMT.¹⁰ The practical effect of this change is that Coloradans will be unable to determine why a decision was made and therefore will not be able to effectuate their rights under the Act to correct factually incorrect or materially inaccurate personal data used in a consequential decision about them. That is unfair. The Department should reinstate the requirements to disclose the principal reasons for a decision and strengthen the requirements for the information that must be given to consumers.

Rights are only valuable to consumers if they are actionable. Without understanding how they compare to other consumers, Coloradans faced with adverse ADMT decisions will not know whether they should pursue additional action. For example, if a person who knows that their payment history on a credit card was used to determine that their account should be sent to collections, but they know that they have not missed a payment, they have the information they need to ask to correct that information. A vague explanation that the decision was made based on “financial information, credit report information, and publicly available information”¹¹ does not provide them with the information they need to make their rights actionable. Thus, to empower consumers, the regulations should a business to provide contextual information. The language to

⁸ Goddard, Roudsari, Wyatt, *Automation Bias: A Systematic Review of Frequency, Effect Mediators, and Mitigators*, J Am Med Inform Assoc. (Jan-Feb 2012), <https://pmc.ncbi.nlm.nih.gov/articles/PMC3240751/>.

⁹ See Lucas Wright et al., *Null Compliance: NYC Local Law 144 and the Challenges of Algorithm Accountability*, AccT ’24, June 03–06, 2024, <https://doi.org/10.1145/3630106.3658998>; Assessing the Assessments, *supra* note 5, at 47-39.

¹⁰ Rule 7.4, 7.5.

¹¹ Rule 7.5(D)(1).

do so is already in Rule 7.6(A)(4), which spells out the information that a deployer must provide if a Consumer requests additional information about the Covered ADMT. The Department should move this rule from Rule 7.6(A)(4) to Rule 7.4, including it as part of the plain language description of the consequential decision and the role the covered ADMT played in that decision:

Describe the types and categories of Personal Data that were considered by the Covered ADMT shall be described inaccurately and with a level of detail that provides Consumers with a meaningful understanding of the information considered by the Covered ADMT. If the Covered ADMT generated any profiles or inferences, the types and categories of those profiles or inferences constitute types and categories of Personal Data that must be disclosed.

- a. For example, types and categories of Personal Data described at a sufficiently granular level of detail include, but are not limited to, “Credit score,” “Medical or health information,” and “criminal history.”
- b. For example, types and categories of Personal Data described at a sufficiently granular level of detail include, but are not limited to, “Inferences about your socioeconomic status,” and “profile about your educational attainment and future earning potential.”

ADMTs use personal data to automate decision-making processes that impact Coloradans’ access to everyday services including health care, education, employment, financial services, and public benefits. A report by TechTonic Justice estimates that virtually all 92 million low-income people in the U.S. have some basic aspect of their lives decided by AI.¹² Coloradans were given the right under the Act to correct inaccurate personal data used in a consequential decision about them, but they cannot do so if they do not know what personal data was used as the principal reason in an adverse decision. The Department should reinstate the requirements to disclose the principal reasons for a decision and use the language suggested above to strengthen the requirements for the information that must be given to consumers.

C. The Rules Should Strengthen Consumer Access to Meaningful Human Review and Reconsideration.

Responsive to Rules 7 and 8

The ability for consumers to appeal consequential decisions is an important aspect of the ADMT Act that the Colorado Legislature passed into law. There are numerous instances of businesses deploying ADMTs to make consequential decisions that result in erroneous or biased

¹² Kevin De Liban, *Inescapable AI: The Ways AI Decides How Low-Income People Work, Live, Learn, and Survive*, TechTonic Justice (2024), <https://www.techtonicjustice.org/reports/inescapable-ai>.

decisions that negatively impact their ability to secure housing, credit, medical treatments, public benefits, and employment and educational opportunities, which in turn significantly affect a person's economic, social, physical, and mental wellbeing.¹³ The erroneous and biased decisions often exacerbate existing inequities.¹⁴ Ensuring the right for consumers to request reconsideration with corrected data or meaningful review when the consumer believes that the adverse decision was incorrect is therefore essential to protect the interests of all Colorado residents. The Updated Proposed Rules unfortunately water down this important accountability mechanism and instead insert additional hurdles for consumers and new ways for ADMT developers and deployers to avoid responsibility.

1. Proposed Factors for Commercially Reasonable Considerations Will Allow ADMT Deployers to Skirt Accountability to Consumers.

Responsive to Rule 8.7(E)

The ADMT Act states that when a consumer experiences an adverse decision by the deployer's use of a covered ADMT that materially influences the decision, the deployer must provide "an opportunity for meaningful human review and reconsideration of the consequential decision, to the extent commercially reasonable."¹⁵ The Updated Proposed Rules enumerate the following factors for "commercial reasonableness":

- a. The type of review required;
- b. The magnitude of harm resulting from the Adverse Outcome, including whether the Adverse Outcome deprives a Consumer of a basic human need, including food, shelter, or required medical care;
- c. The reversibility of the Adverse Outcome;
- d. The value provided by the review of available primary evidence;

¹³ Screened & Scored, *supra* note 5; Assessing the Assessments, *supra* note 5, at Part I; Ben Green, *The Flaws of Policies Requiring Human Oversight of Government Algorithms*, Computer Law & Security Review, vol. 54 (July 2022),

<https://www.sciencedirect.com/science/article/pii/S0267364922000292#sec0007>.

¹⁴ Green, *supra* note 13; Comments of EPIC to the Federal Trade Commission on the Proposed Trade Regulation Rule on Commercial Surveillance & Data Security, *Disrupting Data Abuse: Protecting Consumers from Commercial Surveillance in the Online Ecosystem*, EPIC, 67-97 (Nov. 2022), <https://epic.org/wp-content/uploads/2022/12/EPIC-FTC-commercial-surveillance-ANPRM-comments-Nov2022.pdf>.

¹⁵ SB 26-189 6-1-1705 (1)(a)(II).

- e. The Deployer's size and capacity;
- f. The marginal cost and technical feasibility of the review; and
- g. The availability of reviewers.

All factors other than factor (c) should be struck to ensure that covered deployers allocate sufficient resources to ensure accuracy and fairness of their ADMTs and that the right to meaningful human review is a real accountability tool for Colorado consumers, not just a paper tiger.

The only legitimate consideration for what is “commercially reasonable” is factor (c), “the reversibility of the adverse outcome.” If, for example, a consumer applied to an apartment and received an adverse decision from the deployer of a covered ADMT, and the apartment was already rented out to a different applicant, human review nor a reconsideration can reverse the adverse outcome. However, in that example, the deployer should have to show that the adverse outcome is indeed irreversible because the apartment is already rented out to someone else.

The other factors are not legitimate considerations. While factor (b), “the magnitude of harm resulting from the Adverse Outcome” is a seemingly legitimate consideration, in context of the ADMT Act, which already limits coverage to consequential decisions, all covered Adverse Outcomes are already consequential and can result in harm. The regulations should not provide covered entities an opportunity to dictate which consequential decisions are not harmful enough to provide consumers meaningful human review or reconsideration. Similarly, factor (d) would allow the deployer to predetermine whether the review of available primary evidence would be worthwhile, instead of conducting the meaningful human review or reconsideration itself. This factor would incentivize entities seeking to avoid accountability to argue that the review is not worthwhile, adding another hurdle for consumers to vindicate their rights.

Factors (a), (e), (f), and (g) create perverse incentives for deployers to irresponsibly use ADMTs in consequential decisions without planning for sufficient resources to correct and respond to legitimate consumer requests. For example, take factor (a), “the type of review required.” This factor suggests that a deployer may argue that it can forgo providing consumers with a meaningful human review or reconsideration because the type of review involved would be too complex. The decision at stake is already a consequential decision, which can of course involve many underlying data points to make a decision. Instead of incentivizing covered entities

to make such decisions responsibly and carefully, this factor could allow them to delegate a complex decision-making process to an ADMT, then throw their hands up and argue that a meaningful human review is commercially unreasonable because it is too complex. Factors (e), (f), and (g) relate to whether the deployer allocated enough resources to conduct review. An organization that is too undercapitalized to provide safeguards when making consequential decisions with an ADMT should simply not be permitted to use an ADMT to make that decision.¹⁶ The proposed rules indicate to covered entities that they can forgo budgeting for conducting meaningful human review and reconsideration and cite that lack of budgeting as a reason that it is commercially unreasonable when consumers seek reconsideration or meaningful human review. This incentive structure fails to protect Colorado residents from irresponsible deployers. Removing factors (e), (f), and (g) would also incentivize deployers to test and audit ADMTs for accuracy and fairness, which can minimize the costs of correcting erroneous data or conducting meaningful human review after deployment. The rules should incentivize deployers to responsibly and thoughtfully deploy ADMTs instead of allowing deployers to use their own irresponsible lack of resource allocation to avoid accountability after harm occurs.

In summary, entities that deploy ADMTs without allocating sufficient resources ahead of time are not acting responsibly—the regulations should not allow them to use their lack of planning as a shield against accountability. Factors other than (c) should be struck. Alternatively, if the factors are kept, at minimum there should be a rebuttable presumption that providing consumers meaningful human review and reconsideration is commercially reasonable.

Even if commercially reasonable factors are kept for meaningful human review, the same factors should not apply to whether the deployer should conduct reconsideration of the decision materially influenced by an ADMT when requested by a consumer. The argument above, that the only legitimate factor to consider being factor (c), “the reversibility of the adverse outcome,” applies with stronger force to reconsideration of the decision based on corrected personal data. Reconsideration does not require meaningful human review—it only requires the deployer to re-perform the decision-making process involving a covered ADMT.¹⁷ Compared to providing

¹⁶ See also Comments of EPIC, the Center for Digital Democracy, and the Consumer Federation of America to the California Privacy Protection Agency on Proposed Rulemaking re: Cybersecurity Audits, Risk Assessments, and Automated Decisionmaking, 42 (Mar. 27, 2023), <https://epic.org/wp-content/uploads/2023/03/EPIC-et-al-comments-CCPA-rulemaking-March-2023-2.pdf>.

¹⁷ Rule 8.7(B)(2).

meaningful human review, the cost of reconsideration must be marginal, and as it involves a covered ADMT, at least partially automated. Further, it is objectively unfair to allow a decision that was based on inaccurate personal data about a consumer to be the final decision once a deployer has been informed of the inaccuracy. Instead of adding the language that “reconsideration means that the decision-making process must be re-performed using the Consumer’s correct Personal Data, *if Commercially Reasonable*,” the rules should strike “if commercially reasonable” and insert in its place, “unless the deployer can show that the Adverse Outcome is irreversible.”

2. Proposed Language Removing the Prohibition on Subordinates Conducting Meaningful Human Review Undermine Meaningfulness.

Responsive to Rule 8.7(C)(1)

The original proposed rules for Rule 8.7(C)(1) (previously Rule 7.7(C)(1)) read, “A Meaningful Human Review must be conducted by an independent reviewer who did not make the original decision and who is not a subordinate of the original decision-maker, whenever feasible.” The updated language waters down the requirements for a meaningful review by removing the requirement for the reviewer to not be a subordinate of the original decisionmaker, and it would even allow for the review to be conducted by the same person. If the original decision-maker can conduct review of the decision, the reviewer is more likely to be biased to conform to the original decision. While human review is not a magic bullet,¹⁸ the rules should at minimum ensure that the review is done by an independent reviewer that is different from the original decision-maker and would be free of the pressure that would result from being a subordinate to the original decisionmaker. The rules should reinstate the original proposed language for the meaningful human review so that Colorado consumers can trust that the human review is indeed meaningful.

D. Updated Proposed Rules Let Deployers Off the Hook From Responding to Consumer Requests and Imposes Unnecessary Barriers for Consumers.

The updated language reduces accountability and responsibility on deployers to the detriment of Colorado consumer interests in several instances throughout the updated proposed

¹⁸ Green, *supra* note 13.

rules. The Department should revert the language that allows deployers to point the finger at developers and third parties to satisfy the rules' requirements for transparency and accountability.

1. Deployers Should be Required to Disclose Additional Information on Covered ADMT Immediately Upon Request by Consumer.

Responsive to Rule 7.6(A)

The initial proposed language for Rule 7.6(A) (previously Rule 6.6(A)) provided that when a consumer requests additional information about a covered ADMT through a deployer-provided hyperlink or via telephone, the deployer must “make available immediately, upon the Consumer clicking or navigating to the hyperlinked page or calling a toll-free number, the information about the Covered ADMT and the inputs, including the name of the Covered ADMT, the Covered ADMT version number, if applicable, the Covered ADMT Developer, and the types, categories, and sources of Personal Data used.” This provision would have ensured that the consumer received important information about a covered ADMT without delay, which could enable actionable steps by the consumer if they have reason to believe that an adverse decision was incorrect or made based on inaccurate personal data. In contrast, the Updated Proposed Rules gives the deployer 15 days to respond to the consumer’s request for additional information. This change unnecessarily delays a consumer’s ability to access important information about the covered ADMT, even though the information that the deployer of the ADMT is required to disclose should already be in the possession of the deployer and easy to disclose immediately.

2. Deployers Should Not Be Allowed to Point Fingers at Developers for Missing Information.

Responsive to Rules 7.4(B)(5) and 7.6(A)(5)

The Updated Proposed Rule added language in Rule 7.4(B)(5) which states, “If the Deployer is unable to fulfill any of the obligations in section 6-1-1704(3)(a), C.R.S., and this Rule because it has not received the necessary information from the Developer, its notice should contain as much information and specificity as possible based on the information that it has received from the Developer.” This provision would allow deployers off the hook too easily when it should be the deployer’s responsibility to get information from the developer to ensure

compliance with state law and regulations. The added language should be struck in its entirety to ensure that deployers are incentivized to only contract with developers when all necessary information is provided by the developer to the deployer. Alternatively, the language should read, “If the Deployer is unable to fulfill any of the obligations in section 6-1-1704(3)(a), C.R.S., and this Rule because it has not received the necessary information from the Developer, its notice should contain as much information and specificity as possible based on the information that it has received from the Developer *and the Deployer must prove that it requested the missing information from the Developer and that the Developer refused or failed to respond.*”

Similarly, Rule 7.6(A)(5) provides that when a consumer requests additional information about a covered ADMT, “The Deployer is presumed to know the source of any information that it directly receives, the source of any information that a vendor utilizes to run a Covered ADMT on the Deployer’s behalf, and the source of any information that was identified by name by the Developer to the Deployer” and is required to provide that information to the consumer. However, Rule 7.6(A)(5)(a) goes on to say, “For example, specific data brokers, databases, courts, social media companies, schools, and employers, and any other source of Personal Data for the covered ADMT, must be identified by name *when that information is known to the Deployer.*” (Language in italics added in the Updated Proposed Rules). The new language implies that a deployer may not know that information and would not be found in violation of the regulations for failing to obtain that information when operating a covered ADMT. This encourages irresponsible use of covered ADMT that uses personal data whose source the deployer does not know.

Finally, Rule 7.6(A)(5)(b) provides, “If the Deployer has obtained Personal Data from a third party such as a data broker or aggregator that obtains Personal Data from first party or other third-party sources, the description shall include the original source of the information and any intermediary or vendor source(s) of the information *when that information is known to the Deployer.*” (Language in italics added in the updated proposed rules). The new language similarly lets the deployer off the hook for failing to know the original source of personal data and any intermediary or vendor source(s) that can impact Coloradans in consequential decisions. Instead, the rules should aim to encourage deployers to use covered ADMTs that process personal data from known, reliable sources. The added language in both provisions in italics

should thus be removed. Alternatively, the deployer could be given an affirmative defense that it put in a good-faith effort to obtain such information and the developer or third-party sources refused to provide such information, if the deployer does not disclose the necessary information.

3. The Deployer Should Not Be Allowed to Offload Access and Correction Rights to Third Parties.

Responsive to Rules 8.2, 8.3, and 8.4

The updated proposed rules fail to impose responsibility on the deployer that would lessen the burden on consumers to vindicate their right to access and correct their personal data used by a Covered ADMT. The updated proposed rule does this by introducing a concept of a deployer that “does not have control or Possession of the Personal Data used in the Consequential Decision,” who can offload its responsibilities owed to consumers to third parties while continuing to deploy the covered ADMT at issue.¹⁹ For example, Rule 8.2(D)(2) provides, “If the Deployer does not have control or Possession of the Personal Data used in the Consequential Decision that used a Covered ADMT, the request for instructions for requesting or correcting Personal Data shall allow a Consumer to request instructions for accessing or correcting data in the possession or control of a third-party.” Rule 8.3(C) states, “If the Consumer’s Personal Data used in making the Consequential Decision is not within the Deployer’s possession or control, a Deployer shall comply with Consumer’s request for instructions for requesting Personal Data by providing a list of third parties that possess or control the Consumer’s Personal Data, the categories of Personal Data that each third party controls or possesses, and instructions on how the Consumer can exercise their right under section 6-1-1306, C.R.S. to access their Personal Data through the party that possesses or controls the Personal Data.” Rule 8.4(B) allows deployers to similarly provide a list of third parties that possess or control the consumers’ personal data in response to a request for instructions to correct personal data.

These new provisions would bless a process where a consumer is faced with numerous third parties to contact individually to access their personal data and attempt to correct if there are any inaccurate data points. This process could be extremely burdensome to consumers when

¹⁹ Rule 8.2-8.4.

they may have already been harmed by adverse consequential decisions, while imposing very little responsibility on the deployer that used the covered ADMT at issue. This process looks strikingly similar to the status quo, where consumers are led on a wild goose chase to attempt to correct inaccurate data about themselves that are hindering their ability to secure important opportunities in their lives.²⁰ Instead, deployers should be responsible for passing on access and correction requests to the appropriate third parties that control or possess the personal data at issue, engage in follow-up, communicate in a timely manner with the consumer, and responding to the consumer’s request. The deployer of a Covered ADMT in consequential decisions should shoulder proportionate responsibility to respond to access and correction requests instead of offloading them to consumers.

II. Chatbot Safety Act Rules

A. Definitions should further the intent of the laws and not create additional loopholes.

Responsive to Rules 2 and 9

1. The “age category” definition should be removed.

EPIC, CFA, and Fairplay would recommend striking the definition of “age category” entirely. The concept of age categories is only referenced once later in the rules, in a requirement about what metrics operators must include in an annual report. No substantive requirements in the statute or rules require chatbot operators to determine users’ ages with a level of specificity more precise than whether or not someone is a minor. This means that operators can comply with the law by determining simply whether or not someone is a minor and then, if a user is a minor, implementing the protections required by the law. This “minor” or “not a minor” determination is more privacy-protective than requiring or even allowing chatbot operators to ascertain users’ ages or age range with enough specificity to place them within one of the age categories set out in this definition.

Given that this age category structure is used in only one reporting requirement, we believes the potential incentives for chatbot operators to collect more personal information than necessary or to engage in more accurate but less privacy-protective age assurance methods

²⁰ Screened & Scored, *supra* note 5; De Liban, *supra* note 12, at 122-124.

present a higher risk of harm to users than the marginal benefit that may come from more specific reporting from chatbot operators. Thus, we recommend striking the definition of “age categories” and its one reference later in the rules (14.4(A)). However, if the Department does retain the “age category” concept, it is essential that one of the categories encompasses adults.

2. The Department should ensure the rules do not introduce additional weaknesses into definitions, especially the “Conversational Artificial Intelligence Service” definition.

The definition of “conversational artificial intelligence service” outlines what chatbots are covered by this law and these regulations. Unfortunately, this definition—and its 12 exemptions—is woefully inadequate to cover many of the chatbots that are causing harm to minors. The deficiencies with this definition and its numerous exemptions are primarily within the statute itself, and the problems the inadequate definition creates cannot be cured by clarifying regulations alone. The draft regulations on the exemptions within this definition do not add to our concern with the law’s coverage, but we urge the Department to take care not to further expand the exemptions through regulation.

Importantly, EPIC, CFA, and Fairplay also implores the Legislature to revisit the Chatbot Safety Act during the next legislative session to close some of the numerous loopholes and strengthen the provisions to ensure it can be an effective mechanism for limiting chatbot harms, which unfortunately is impossible in its current form.

B. An effective age assurance framework should center privacy, security, and user trust.

Responsive to Rule 10

EPIC has been a leading voice on platform governance issues, advocating for meaningful regulation and oversight of social media platforms and chatbot providers that respects the privacy, speech rights, and associational rights of users. As part of this work, EPIC has provided numerous courts, legislators, and attorneys general with guidance about the constitutional, privacy, and access questions implicated by kids’ online privacy and safety legislation, including filing numerous amicus briefs across the country and in the U.S. Supreme Court on these

issues.²¹ We have also given expert recommendations to policymakers on how to implement safeguards that require age assurance while properly striking the balance between protecting privacy and ensuring efficacy, including providing feedback to New York on its SAFE for Kids Act; to California on various online safety laws with age components, including SB 976's regulation of addictive feeds; and to Vermont on its Age-Appropriate Design Code.²² This legal and policy expertise combined with EPIC's deep privacy and technical knowledge makes us uniquely well-positioned to provide feedback to policymakers seeking to implement privacy-protective and effective age assurance regimes.

Responsive to Rule 10.2(A)

EPIC, CFA, and Fairplay commend the Department for its obvious attention to prioritizing privacy in the drafted rules. We were particularly happy to see the Department embrace data minimization principles, such as in Rule 10.2(A)(1). Requiring operators to limit the data they collect for age assurance, limiting its use to that specific purpose, and requiring deletion of age-assurance data upon determining age will help users feel more secure submitting to age assurance and will build trust that the age assurance process will not give tech companies another avenue with which to exploit personal data. We recommend that the Department further limit the data that operators can collect for the purposes of conducting age assurance to only the information that is “strictly necessary” to determine whether a user is a minor. The “strictly necessary” standard is one that is often applied to protect the processing of sensitive personal data, which includes personal information about minors, so data collected for age assurance purposes clearly falls within this category of information that deserves enhanced protection.

²¹ *Platform Governance Laws and Regulations*, EPIC, <https://epic.org/issues/platform-accountability-governance/platform-governance-laws-and-regulations/>.

²² *Comments of EPIC to the New York State Attorney General on the Advance Notice of Proposed Rulemaking for the SAFE for Kids Act Pursuant to New York General Business Law Section 1500 et seq.*, EPIC (Sept. 30, 2024), https://epic.org/wp-content/uploads/2024/10/EPIC-Comments_NY-SAFE-For-Kids-Act.pdf; *Comments of EPIC to the California Department of Justice on the Notice of Proposed Rulemaking for SB 976: Protecting Our Kids from Social Media Addiction Act*, EPIC (June 30, 2026), https://epic.org/wp-content/uploads/2026/07/EPIC_SB-976-Rulemaking-Comments.pdf; *Comments of EPIC to the Vermont Attorney General's Office Re: Proposed Vermont Age-Appropriate Design Code Rules for Sections 2449g of Act 63*, EPIC (Oct. 2, 2026), <https://epic.org/wp-content/uploads/2026/10/EPIC-Comment-Vermont-AADC-%C2%A7-2449g-Rulemaking.pdf>.

In addition to these data minimization principles, the Department should also require operators to conduct age assurance with other privacy-protective methods in mind. For example, processing user data on the users' device instead of a remote server and using cryptographic techniques to protect information in transit also enhance user privacy, security, and trust.²³ Tools that process and store data on users' devices do not expose users to the same privacy and security risks as tools that process on companies' servers because the companies do not collect the data at all—the data never leaves the users' devices. The company thus cannot repurpose the data, and a hacker attacking the company's servers cannot steal the data. These techniques also mitigate concerns about age assurance chilling First Amendment-protected activity by adults by preventing users' identifying information from being linked to their internet use. Techniques like on-device processing and cryptographic techniques should also be explicitly included by the Department as considerations for operators in choosing age-assurance methods.

We also appreciate that Rule 10.2(A)(2) requires operators to offer users a choice of age assurance methods and prohibits providing government identification from being the sole method offered to users. This is a great start, but we would recommend the Department specify that operators cannot solely use any one method of age assurance and instead require operators to offer users with multiple methods to choose from. Different users are likely to prefer different methods. Some may be reticent to provide government ID—or not have one at all—while others may prefer this authentication method because they are familiar with the process from other authentication experiences. Some users may avoid parental attestation as an age determination method if the user does not have a reliable or safe relationship with a parent. These examples illustrate the importance of providing user with choices throughout the age assurance process, and we strongly recommend the Department mandate that operators offer users this choice. Best practice for protecting users' privacy would be for the Department to require operators to offer at least one method that does not rely on submission of a government ID or biometric data to determine age.

²³ Scott Babwah Brennan & Matt Perault, *Keeping Kids Safe Online: How Should Policymakers Approach Age Verification?*, Utah State University Center for Growth and Opportunity 4 (2023), <https://www.thecgo.org/research/keeping-kids-safe-online-how-should-policymakers-approach-age-verification/>.

In addition to the excellent factors set out in Rule 10.2(A)(1) and (2), EPIC, CFA, and Fairplay recommend that the Department add a requirement that operators first use information they already possess about users to determine age, if possible, before requesting additional personal information from users. Many online companies already have sufficient information to make a reasonable determination of whether a user is a minor or an adult. The Department should adopt regulations that require companies to use information they already possess before requesting any additional personal information from users. For example, many operators already use personal information, including a user’s status as a minor or an adult, to target that user with advertising.²⁴ Similarly, operators may be provided with clear indications that a user is a minor in the chat interface, such as discussions about school or grade level, references to child- or teen-oriented media, or questions about homework assignments clearly designed for minors. The Department should require operators to use this kind of information already in their possession for age assurance before requesting additional information from users. These methods of age determination would then be, essentially, frictionless; operators could make the age determination on the back end without collecting any additional information from users. The Department should, however, make clear that the Act does not license operators to begin or increase surveillance of users for the purpose of determining age.

In addition to these privacy considerations, several of the factors under Rule 10.2(A) rightly require operators to reassess users’ ages at certain points, including when minors become adults and if operators suspect or receive a report of fraud. EPIC, CFA, and Fairplay strongly urge the Department to also require operators to offer an appeal or error-correction process that users can access if their initial age determination is inaccurate. This process should allow users to submit additional information but should not require users to use one specific method, particularly submission of a government identification, to confirm their age status. The appeals process should also set out a timeline for review so that users do not experience lengthy delays.

²⁴ Kate O’Flaherty, *Meta AI to Start Using Your Data for Targeted Ads—Here’s How and When*, Forbes (Oct. 6, 2025), <https://www.forbes.com/sites/kateoflahertyuk/2025/10/06/meta-ai-confirms-your-data-will-be-used-for-ads-heres-how-and-when/>; Joseph Cox, *How ChatGPT Targets Ads to You Based on Your Prompts, Phone Number, and Email*, 404 Media (Oct. 15, 2026), <https://www.404media.co/how-chatgpt-targets-ads-to-you-based-on-your-prompts-phone-number-and-email/>.

We recommend the Department look to California’s proposed rules on SB 976 § 564 for a strong example of such an appeals process.²⁵

Rule 10.2(A)(8) rightly acknowledges that users’ self-declarations are not an accurate measure of whether a user is, in fact, an adult. However, we recommend that the Department revise this provision slightly and allow operators to accept users’ self-declarations that they are under 18 as a determination that a user is a minor. The incentives for users to misrepresent their age on self-declarations all encourage users to claim they are adults when they are not to avoid safety features and other restrictions. There are no similar incentives for users to misrepresent that they are minors when they are actually adults. Thus, operators should be able to rely on users’ self-declarations that they are minors as evidence that they should be treated as minors.

The factors in Rule 10.2(A) are set out as items the Department will consider when determining if an age assurance method is commercially reasonable. EPIC, CFA, and Fairplay would recommend that these factors be converted into mandatory criteria that operators must follow before they may adopt a particular age assurance method. Taken together, these factors will make the age assurance process easier, more effective, and more privacy-protective for users, which should be the top priority of any age assurance framework. Thus, it is essential that operators meet all of these criteria before requiring users to undergo age assurance.

Responsive to Rule 10.2(B)

Rule 10.2(B) sets out particular age-assurance methods that are generally accepted in a non-exhaustive list. We commend the Department for allowing operators flexibility to choose other age-assurance mechanisms. This flexibility incentivizes development and adoption of additional age-assurance methods that may be more accurate or privacy-protective.

However, we are concerned that Rule 10.2(B)(4)’s allowance of “an assessment of the user’s digital footprint originating from a verified email address” is not an adequately privacy-protective method. One’s digital footprint can often reveal a plethora of sensitive personal data that may be more invasive than even a process requiring submission of a government

²⁵ Proposed Regulations on SB 976, Protecting Our Kids from Social Media Act § 564, California Department of Justice, <https://oag.ca.gov/system/files/media/sb976-proposed-regulations.pdf>.

identification.²⁶ Inclusion of this method may also incentivize operators to purchase personal information on users from data brokers or to indiscriminately scrape their personal information from across the web.²⁷ Because of these unintended incentives and risks to privacy, we recommend removing this method from the list of approved methods for age assurance.

EPIC, CFA, and Fairplay also suggest that the Department add two additional methods to the list of commercially reasonable methods outlined in Rule 10.2(B). First, as explained above, we would recommend the Department add as an acceptable method self-declaration that a user is a minor. Again, while self-declaration that someone is an adult is clearly not an effective method of age assurance, the same is not true of self-attestation that a user is a minor; the incentives are simply not there for a user to misrepresent their age in this direction. Second, we would recommend adding an explicit provision allowing operators to use data or inferences of the users' age, age range, or status as an adult or minor that the operator has already associated with a user, including inferences an operator has used to target a user with advertisements, to determine age. Because this method relies on data already in the operators' control, it is extremely privacy-protective, and because advertising is tied directly to an operator's financial bottom line, the operator has a vested interest in using only data or inference that it has a high degree of confidence is accurate.

EPIC, CFA, and Fairplay encourage the Department to continue to center privacy in its determinations of which age-assurances methods are acceptable for operator use. The privacy risks from age determination will not just vary method-to-method but will also vary tool-to-tool—different companies will have different data practices and implement different privacy and security features. The Department should assess the privacy risks posed by each potential age determination tool and prioritize tools that minimize data collection, use, and disclosure. Accordingly, the Department should prioritize age assurance tools that do not collect sensitive personal information, such as biometric identifiers or social security numbers.

²⁶ Comments of EPIC on the FTC's Proposed Trade Regulation Rule on Commercial Surveillance and Data Security, 7, EPIC (Nov. 2022), <https://epic.org/wp-content/uploads/2022/12/EPIC-FTC-commercial-surveillanceANPRM-comments-Nov2022.pdf>.

²⁷ Comments of EPIC on the CFPB Request for Information Regarding Data Brokers and Other Business Practices Involving the Collection and Sale of Consumer Information, EPIC, 2-14 (July 14, 2023), <https://epic.org/wp-content/uploads/2023/07/EPIC-CFPB-data-brokers-RFI-comments-071423.pdf>.

Finally, a concluding provision in Rule 10.2(B) discusses how the Department will consider whether particular methods of age assurance are commercially reasonable. We recommend removing consideration of the “financial resources” of the operator from the determination of whether or not a particular age-assurance method is commercially reasonable. If an operator does not have the financial resources to provide users with an effective and privacy-protective method of age assurance, the operator should not operate a chatbot at all. Securing an age-assurance method that protects users’ privacy and can effectively and reliably determine whether a user is an adult or a minor is not something that is nice to have for a chatbot operator; it is a necessity to properly comply with Colorado law and to keep users safer. Therefore, the operator’s financial situation should not be a consideration when the Department is determining what methods are appropriate for operators to use to determine age.

III. Conclusion

EPIC, CFA, and Fairplay urge the Department’s continued development of the rules to ensure that Coloradans have the strongest protections possible under the ADMT Act and the Chatbot Safety Act. EPIC, CFA, and Fairplay will each continue to be available for future discussion about our recommendations as to how the Department can best protect Coloradans.

Sincerely,

/s/ Caitriona Fitzgerald
Caitriona Fitzgerald
EPIC Deputy Director and Policy Director

/s/ Kara Williams
Kara Williams
EPIC Counsel

/s/ Mayu Tobin-Miyaji
Mayu Tobin-Miyaji
EPIC Counsel

/s/ Ben Winters
Ben Winters
CFA Director of AI and Privacy

/s/ Haley Hinkle
Haley Hinkle
Fairplay Policy Counsel

/s/ Brendan Bouffard
Brendan Bouffard
Fairplay Staff Attorney