September 4, 2026

Colorado Department of Law
Office of the Attorney General
Ralph L. Carr Colorado Judicial Center
1300 Broadway, 10th Floor
Denver, CO 80203

Re: Colorado Automated Decision-Making Technology Act (SB 26-189) Rules

Dear Attorney General Weiser,

On behalf of the American Financial Services Association (AFSA)[1], thank you for the opportunity to comment on the proposed rules implementing the Colorado Automated Decision-Making Technology Act, published August 11, 2026. AFSA submitted informal comments in July, and we thank the Department for the engagement reflected in this draft. The proposed rules align consumer rights response timeframes with the California Consumer Privacy Act's (CCPA) ten-day acknowledgement and forty-five-day resolution structure, narrow the language requirement to the languages a business ordinarily uses rather than a mandated list, and give creditors specific content standards for when a federal notice satisfies Section 6-1-1704's disclosure requirements. We appreciate each of these changes. We encourage the Department to look for similar opportunities to align with the Gramm-Leach-Bliley Act's existing privacy and safeguards framework, which already governs many of the same disclosures at the federal level.

Below are the areas where we believe the proposed rules go beyond what the statute requires, or where the rules could use additional clarity to work as intended. In several places described below, we believe the rules exceed the Department's statutory rulemaking authority outright, not simply the scope we think would be most workable, and we flag those instances specifically. In each one, the gap between what the law requires and what the rule does falls hardest on consumers and on the financial services providers trying to comply with it.

We cannot overstate the concerns of our members on several elements of the proposed rules.

1. Implementation Timeline

Even under the most efficient compliance path, our members need considerably more time to implement these rules than the current schedule allows. The rules are not expected to be final until the fourth quarter of 2026, and the rulemaking hearing itself is not scheduled until October 26 or later, yet both the statute and the rules take effect January 1, 2027. Under that timeline,

[1]Founded in 1916, the American Financial Services Association (AFSA), based in Washington, D.C., is the primary trade association for the consumer credit industry, protecting access to credit and consumer choice. AFSA members provide consumers with many kinds of credit, including direct and indirect vehicle financing, traditional installment loans, mortgages, payment cards, and retail sales finance. AFSA members do not provide payday or vehicle title loans.

1750 H Street NW, Suite 650, Washington, DC, 20006 ▲ 202.296.5544 ▲ www.afsaonline.org ▲ @AFSA_DC

businesses could have as little as two months, and quite possibly less, to build compliance programs once the requirements are actually settled.

That interval is too short to stand up the systems the rules require. Even the least burdensome path we describe below, pairing an existing federal adverse action notice with a separate, compliant Colorado disclosure, still requires drafting that disclosure, integrating it into origination and servicing systems, coordinating with third parties, and training staff. Building a Meaningful Human Review function, with qualified reviewers, training, access to primary evidence, and the tracking and documentation the rules contemplate, is a complex operational capability that cannot responsibly be created in a matter of weeks.

We ask the Department to provide a meaningful implementation period, in the form of a good faith compliance or non-enforcement window of nine to twelve months following the promulgation of the final rules, and to confirm that obligations whose operational detail depends on the content of the final rules are measured from the effective date of those final rules rather than from the statute's original January 1, 2027 date. We also ask the Department to state directly in the rule that, in deciding whether to pursue an investigation of a possible or alleged violation, the Attorney General may consider the amount of time between the effective date of a requirement and the possible violation, along with good faith efforts to comply, consistent with the approach California has taken in its own privacy regulations.

The scale of this operational burden is not hypothetical. One member's internal assessment of a single narrow area, adverse action decisioning in the employment context, estimated four to six months of dedicated compliance work and over $200,000 in cost to implement, for an activity involving an estimated ten to fifteen covered decisions per year. Multiplied across the many separate categories of Consequential Decisions these rules touch, and across our full membership, the cumulative implementation burden is substantial and, in categories like this one, disproportionate to the actual volume of affected decisions.

2. Materially Influence and the De Minimis Factor Standard

The Notice asks which of two proposed standards the Department should adopt for "Materially Influence," and we urge the Department to adopt Standard 2. Standard 1 defines a De Minimis Factor as one with only a trifling, trivial, or incidental impact, and provides that a factor is not de minimis simply because other factors played a more significant role. In a credit context, that reading would treat nearly every model output a decision process touches as material, even where independent factors actually drove the decision. Standard 2 defines a De Minimis Factor as one that is not a substantial factor, and recognizes that an ADMT output may be de minimis where other relevant factors independent of that output played a significantly larger role. That better reflects how consumer credit is actually decided, where a model score is often one input among documented, consistently applied policy rules.

Standard 2's rebuttal factors, however, are framed around a human decision maker's independent review and analysis. In high volume, largely automated credit decisioning, the independent factor that actually governs an outcome, is often a documented, consistently applied policy rule, such as a hard delinquency or bankruptcy trigger or a minimum credit score cutoff, rather than a

1750 H Street NW, Suite 650, Washington, DC, 20006 ▲ 202.296.5544 ▲ www.afsaonline.org ▲ @AFSA_DC

case by case human analysis. Standard 1 already lists previously documented and consistently applied dispositive factors, such as a GPA cutoff, as evidence that rebuts material influence. Standard 2's rebuttal list does not include a comparable example. We ask the Department to adopt Standard 2 and expand its rebuttal factors to expressly recognize documented, consistently applied policy rules, including hard delinquency, bankruptcy, or credit score cutoffs, as the kind of independent, substantial factors that rebut the presumption of Material Influence. That change would appropriately narrow coverage for rule-dominant credit decisions, where an automated process merely replicates a manual rule-based decision to increase efficiency. More broadly, the cleanest way to reach that result is for the Department to pair the De Minimis Factor definition in Standard 2 with the rebuttal provision set out in Standard 1. The two rebuttal provisions differ in three ways that matter. Standard 1 permits rebuttal either by showing that the output was a De Minimis Factor or by showing that it was not consulted, considered in, or incorporated into the decision, while Standard 2 offers only the first route. Standard 1 introduces its evidence list with "for example" and joins the listed items disjunctively, so that any one of them may suffice, while Standard 2 joins its four items conjunctively, which reads as requiring all four. Standard 1 closes with a catch-all for "other persuasive evidence showing judgment independent of the ADMT output," which Standard 2 omits, leaving a closed list. Adopting the definition in Standard 2 together with the rebuttal provision in Standard 1 would address the concern described above and would require no new drafting by the Department.

We also ask the Department to clarify what does not constitute a Consequential Decision in the Financial or Lending Service context. Routine, risk-based pricing, credit underwriting, and compliance screening processes that are already subject to comprehensive federal regulatory oversight under the Fair Credit Reporting Act, the Equal Credit Opportunity Act, and the Gramm-Leach-Bliley Act should not be swept into Covered ADMT simply because they involve a model or scoring tool that has long been part of how those functions are performed.

3. Adverse Outcome Examples Should Track the Rule's Own Definitions

Rule 6.5.D.1 uses, as an illustrative example of a Financial or Lending Services Adverse Outcome, a creditor's decision to refer a delinquent credit card account to collections. We do not believe this example holds up, on two independent grounds.

First, it does not fit the Department's own definition. Rule 2.2 defines "Financial or Lending Service" as five specific things: the extension of credit or a loan, transmitting or exchanging funds, the provision of deposit or checking accounts, check cashing, and installment payment plans. A collections referral is not any of these. It is a decision about how a creditor pursues recovery of a debt the consumer already owes, not a decision about the consumer's access to any of the five enumerated services. Conforming the example to the Department's own text is a narrow fix that should be made regardless of how the Department addresses the scope of the definition itself.

We would also note that the Act's definition of "consumer" turns on an individual acting in an "individual and household context," which differs from the terminology used in other well-established Colorado and federal frameworks. The Colorado Consumer Protection Act, the

1750 H Street NW, Suite 650, Washington, DC, 20006 ▲ 202.296.5544 ▲ www.afsaonline.org ▲ @AFSA_DC

Colorado Consumer Credit Code, and Regulation P each define covered consumer financial products and services as those used primarily for a personal, family, or household purpose. Aligning the Act's "individual and household context" language with that more familiar standard would promote consistency across regulatory regimes and reduce the risk of inconsistent interpretation. We would caution, however, that this standard describes the capacity in which a consumer acts. It does not describe which products or services are financial or lending services, and it should not be imported into Rule 2.2 in a way that would sweep ordinary consumer transactions that are not credit into the enumerated definition.

Separately, we ask the Department to clarify the enumerated term "installment payment plans" in Rule 2.2. Read literally, that term reaches ordinary retail payment arrangements that a seller offers to facilitate the purchase of its own products, including layaway plans, pay-over-time arrangements for merchandise, and equipment payment plans billed alongside a monthly service charge. These are terms of sale for a non-financial product, not financial or lending services offered to the consumer as a product in their own right. We ask the Department to clarify that "Financial or Lending Service" does not include an extension of credit, retail installment sale, lease, or installment payment plan offered by a seller or lessor of goods or services to facilitate a consumer's acquisition of that seller's or lessor's own non-financial goods, equipment, or services, whether or not the arrangement carries a finance charge and whether or not a credit check is used in connection with it. Financing offered by a third-party credit provider, and credit extended as a financial product in its own right, would remain within the definition. Drawing the line according to whether the arrangement finances the provider's own non-financial goods and services, rather than according to the presence of a finance charge, keeps the covered universe of consumer credit products intact while removing transactions the General Assembly is unlikely to have understood as financial or lending services.

Second, and independent of the first point, a collections referral does not fit the statute's own definition of an Adverse Outcome, which requires a decision that denies, terminates, or materially reduces or restricts a consumer's access to a financial or lending service, and could exceed the Department's rulemaking authority for that reason as well. By the time an account is delinquent and referred to collections, the consumer's access to that account has already been affected by the consumer's own nonpayment. The referral itself does not deny or restrict anything the consumer still has, and any referral must, in any event, independently satisfy applicable federal and state law and the terms of the underlying contract, so an automated system is never truly the sole basis for that decision.

A collections referral is also not appropriately treated as an Adverse Outcome because it is a foreseeable consequence of default disclosed in the underlying credit agreement. When a consumer fails to meet their payment obligations, a referral to collections reflects the servicing and recovery of an existing debt rather than a new decision regarding the consumer's access to a covered financial or lending service. As such, it is more properly characterized as routine post-default account administration than as an Adverse Outcome under the Act.

We would also note that automated referral models are typically used to reduce, not expand, the volume of accounts sent to collections. Treating this process as an Adverse Outcome creates an

1750 H Street NW, Suite 650, Washington, DC, 20006 ▲ 202.296.5544 ▲ www.afsaonline.org ▲ @AFSA_DC

incentive to rely less on models of this kind and refer accounts to collections as soon as they become delinquent, which would be worse for consumers, not better.

Treating this kind of automated process as an Adverse Outcome risks a broader, unintended consequence: chilling the use of automated decisioning for actions that benefit consumers. The same dynamic applies to a proactive credit line increase, a reactive credit line increase that lets a consumer make a purchase above an already pre-approved limit, inclusion in a promotional offer such as a short-term interest rate reduction or deferred interest, or the use of alternative data, such as cash flow underwriting, to extend a first credit card or additional credit to a consumer with a thin credit file. Each of these is typically automated, each is designed to benefit the consumer, and each could trigger the same Meaningful Human Review and disclosure obligations that apply to an adverse decision. Treated the wrong way, the rule would give a creditor reason to pull back from exactly the kind of automation that expands access to credit and better terms, particularly for consumers building or rebuilding a credit file.

We ask the Department to remove this example, or replace it with one that reflects an actual eligibility or access decision, and to clarify, consistent with our July comments, that Consequential Decision and Adverse Outcome reach only decisions about whether a consumer obtains a covered product or service, not decisions made in the course of servicing an account the consumer already holds.

Relatedly, we ask the Department to clarify that a process is not Covered ADMT simply because it includes multiple layers of human involvement, or because it automates a filtering step that applies criteria a business already uses in its ordinary course. Two examples in the current draft illustrate the concern. In the employment example at Rule 6.5.F.2, an automated score is one of several inputs alongside two human reviewers, and it is not clear what additional review a third human reviewer could meaningfully add once two humans have already weighed in on the same information. Indeed, a process with that many layers of human involvement is not meaningfully automated in the first place, and should not be treated as ADMT at all. Similarly, screening out an incomplete application before it is ever substantively reviewed is a basic intake function, not a consequential judgment about the applicant, and treating it as Covered ADMT risks pulling routine business processes into a regime meant for decisions with real consequence. We ask the Department to confirm that an application paused or rejected for incompleteness, such as a missing income field or Social Security number, falls within the statutory exceptions to Consequential Decision for routine decisions and for narrow procedural or data-processing functions, since the Covered ADMT's predictive functions play no role in a file being objectively incomplete.

4. Adverse Outcome Disclosure Content Should Track the Statute, Not Exceed It

Several of the itemized content requirements in Rule 6.4.B go beyond what Section 6-1-1704(3)(a) actually requires. That subsection calls for a description of the Consequential Decision and the role the Covered ADMT played in it. It does not contemplate the discrete, granular disclosures the rule adds on top of that standard, such as a requirement to identify and explain specific inferences drawn from Personal Data, or to disclose a consumer's underlying

1750 H Street NW, Suite 650, Washington, DC, 20006 ▲ 202.296.5544 ▲ www.afsaonline.org ▲ @AFSA_DC

profile or risk score. Section 6-1-1704(3)(b) already gives a consumer a separate path to request that kind of detail, including the types, categories, and sources of Personal Data used, and we do not think subsection (3)(a) was meant to require the same level of detail be pushed into the disclosure automatically. Read as a whole, the statute's rulemaking authority for this content is a request for standards, meaning high level parameters for describing an ADMT's role, not a prescriptive list of every category of information a disclosure must itemize.

A creditor relying on an ADMT for a credit decision should be able to use a form notice promulgated under these rules. For example: "An automated system was used to generate electronic scoring metrics and evaluate your credit application." This will allow financial services companies to use defined forms with approved language from the rules and permit consumers to follow the separate path to request more information.

Rule 6.4.B.8's automatic denial factor requirement has a similar problem. Nothing in the statute discusses automatic denials tied to a Deployer's internal standards or practices, and we ask the Department to remove this as a freestanding, itemized requirement. We would also note that Rule 6.4.B.7, requiring the disclosure to accurately describe the decision, the role of the ADMT, and the principal reasons, appears largely redundant with what subsections (A) and (B) already require, and could be removed without any loss of clarity.

Rule 6.4.B.9 also needs to reflect that, for incomplete applications, ECOA and Regulation B give creditors a choice between two federal notices: a notice of adverse action under 12 C.F.R. § 1002.9(a)(1)(ii), or a notice of incompleteness under 12 C.F.R. § 1002.9(c)(2). Section 6-1-1704(6)(a)-(d) already establishes that a creditor providing notice under ECOA and the FCRA is not required to provide a separate or duplicative notice, or one delivered in a manner federal law prohibits. As drafted, Rule 6.4.B.9 could be read to treat the incompleteness notice as the exclusive path, which would restrict a federal option the statute expressly preserves. We ask the Department to clarify that a creditor's choice of either federal notice, consistent with ECOA, satisfies Rule 6.4.B.9.

Finally, Rule 6.4.D.2's requirement that a Deployer provide a single, unified method for a consumer to both request additional information about the ADMT and exercise their ADMT Consumer Rights goes beyond what the statute asks for. Section 6-1-1704(3)(b) asks for instructions and a simple to follow process to request additional information, and section 6-1-1704(3)(c) asks for an explanation of Consumer Rights and how to exercise them. Neither provision requires that these be combined into one unified mechanism, and we ask the Department to remove that requirement.

5. The Federal Notice Accommodation Needs a Clearer, Broader Scope

We appreciate that Rule 6.7.A now defines what a federal ECOA, Regulation B, or FCRA notice must include to also satisfy Section 6-1-1704. Having a defined standard is a meaningful improvement. Our concern now is narrower but still significant, in three respects.

First, satisfying Rule 6.7.A as drafted will, in practice, require creditors to add Colorado specific content directly onto the Regulation B model form. Creditors have relied on that model form's

1750 H Street NW, Suite 650, Washington, DC, 20006 ▲ 202.296.5544 ▲ www.afsaonline.org ▲ @AFSA_DC

compliance safe harbor for decades, and many of our members are reluctant to alter the form itself for fear of forfeiting that safe harbor under federal law.

We ask the Department to clarify that a creditor may satisfy Rule 6.7.A by pairing the unmodified federal model form with a separate, compliant Colorado disclosure, rather than by integrating Colorado specific content into the model form itself. Regulation B's "clear and conspicuous" standard limits the disclosures a creditor may include on the model form to those Regulation B itself sets out. See 12 C.F.R. § 1002.4(d)(1); 12 C.F.R. Part 1002, Supp. I, cmt. 4(d)-1. Regulation B expressly authorizes the addition of only one other disclosure, the Fair Credit Reporting Act's adverse action disclosure. See 12 C.F.R. Part 1002, App. C, cmt. 2; 15 U.S.C. § 1681m. Absent a similar accommodation for Colorado's disclosure, a creditor that adds it directly to the model form risks violating Regulation B and forfeiting the safe harbor that comes with using a Regulation B model notice. See 12 C.F.R. Part 1002, App. C. That approach gives creditors the certainty the accommodation is meant to provide, and the flexibility to implement it, without asking them to choose between two forms of regulatory risk. The Department could remove additional legal uncertainty here by publishing its own safe harbor Colorado disclosure form, consistent with the approach other states and the federal government have taken, that creditors could pair with the federal form with confidence.

Second, we ask the Department to confirm that the accommodation covers account management Adverse Outcomes and not only application decisions, wherever a federal adverse action or risk-based pricing notice is provided for the same decision.

Third, we ask the Department to confirm that the accommodation works together with the FCRA risk-based pricing and credit score disclosure exception, so that a creditor who satisfies its federal risk-based pricing obligation through the model credit score disclosure notice is not also required to generate a separate Colorado disclosure for the same decision.

More broadly, even where the specific disclosures required by these rules and by ECOA or the FCRA are substantively identical, requiring a separate Colorado notice on top of a compliant federal notice creates exactly the kind of duplicative disclosure that Section 6-1-1704 was not intended to require. We ask the Department to confirm that this accommodation should be read broadly, consistent with that intent, so that compliance with existing federal adverse action and credit reporting notice requirements generally satisfies Section 6-1-1704 without requiring a second, largely duplicative notice. We also ask the Department to acknowledge the residual risk of interference or conflict between these rules and existing federal disclosure requirements more generally, so that Deployers are not left to reconcile inconsistent obligations without guidance.

6. Delivery of Adverse Outcome Disclosures Should Not Require Duplicate Notices

Rule 6.2.A.1 requires an Adverse Outcome disclosure to be delivered through at least two methods whenever a Deployer has two ways of reaching a consumer, which in practice means almost every creditor with both a mailing address and an email address on file. We think this requirement exceeds what the statute calls for and, in practice, provides little benefit to consumers while imposing real costs.

Many of our members already run a delivery process built around the exact concern this rule is meant to address. A disclosure is sent electronically, and an established process identifies consumers who do not receive or do not open that notice, so a follow up can be sent by mail to reach them. A blanket two method requirement forces a second, redundant mailing in the large majority of cases where the first method already reached the consumer, generating paper, postage, and processing cost without any gain in consumer awareness. Indeed, most of our members' customers consent to receiving their notices electronically, and they have the opportunity to print or save the electronic notice. Electronic delivery should remain an option a Deployer can choose to use, not a mandatory build for institutions that already rely on other established, compliant channels. We ask the Department to revise Rule 6.2 to allow a Deployer to satisfy the delivery standard through any single permitted method.

This same issue arises for creditors relying on the federal notice accommodation described above. Rule 6.7.A conditions substitute compliance on what the notice says and when it goes out, but is silent on how it must be delivered, while Rule 6.2.A.1 separately imposes its own two method requirement. The Department has already relieved education deployers of this exact tension: Rule 6.7.B permits a FERPA subject deployer to use the channels FERPA itself allows, and Rule 6.5.B.1 repeats that accommodation. No comparable language currently appears in Rule 6.7.A for creditors. We ask the Department to confirm that a notice delivered consistent with ECOA or FCRA delivery requirements also satisfies Rule 6.2, on the same logic already extended to education deployers under FERPA.

Rule 6.3's use of "delivers" alongside "mailing" in the same provision also creates ambiguity about whether a mailed disclosure is deemed delivered when it is placed in the mail or only upon actual receipt. Because these disclosures will often be combined with a federal adverse action notice, that ambiguity matters: if actual delivery is required, creditors may need to use certified mail to confirm receipt, adding meaningful cost to a process that federal law does not require. We ask the Department to clarify that a disclosure is deemed delivered when it is mailed, consistent with how federal adverse action notices are treated.

7. The Meaningful Human Review Standard Needs to Reflect How Credit Is Actually Extended

Rule 7.7's standard for Meaningful Human Review is difficult to reconcile with how consumer credit is actually extended at scale. The example at Rule 7.7.E.7, a consumer applying for store credit at checkout, illustrates a much broader category: point of sale financing and similar products are built around a decision the consumer receives in seconds, at the counter, not after a queued human review. These same timing pressures apply with added force where a single application is routed to multiple finance sources at once, each rendering its own independent decision under its own timeline. A business extending point-of-sale credit is not in a position to hold the underlying transaction open for up to forty-five days while a Meaningful Human Review plays out, and it is not clear what is supposed to happen if a review concludes, after the fact, that a consumer should have received a better offer once the transaction has already closed on different terms. We ask the Department to ensure that its Meaningful Human Review and

1750 H Street NW, Suite 650, Washington, DC, 20006 ▲ 202.296.5544 ▲ www.afsaonline.org ▲ @AFSA_DC

Adverse Outcome timing rules account for these multi-source, point-of-sale credit models, rather than assuming a single-lender, single-timeline decisioning process.

We have significant concerns about Rule 7.7.E.3's presumption that review is commercially reasonable whenever the harm is a severe and irreversible denial of a basic human need. We have not located a definition of "basic human need" anywhere in Colorado law. Common usage of the term generally involves food, shelter, clothing, and medical care and should not include lending decisions. However, the housing example at Rule 7.7.E.5 risks suggesting the presumption could also reach ordinary mortgage lending. If this term is going to serve as a threshold for a commercial reasonableness presumption, it needs a clear and objective definition that excludes lending decisions, and we ask the Department to provide one, or, in the alternative, to remove this presumption altogether, since it goes beyond what the statute itself requires.

When it comes to fair lending, a validated, monitored, automated underwriting model can be tested for disparate impact and applied consistently to similarly situated applicants. That consistency is itself a fair lending safeguard. The statutory definition of Meaningful Human Review, by contrast, calls for a reviewer who does not default to the system's output, which is another way of describing unstructured human discretion, a well documented source of disparate treatment risk under ECOA and the Fair Housing Act. A rule that pushes deployers toward routine, live human override of tested models risks making outcomes less consistent and less auditable, not more fair, and we encourage the Department to use its rulemaking authority to minimize friction between this standard and the fair lending guidance our members already have to follow.

We ask the Department to clarify that commercially reasonable does not require a live human reviewer to be available at the moment of an instantaneous credit decision, consistent with the example at Rule 7.7.E.7, and that a deployer using a validated, monitored automated model, with a defined process for a consumer to request reconsideration afterward, satisfies the standard without staffing real time override capability at the point of sale. We use the word reconsideration deliberately: many of our members already operate a reconsideration process for credit decisions, and would prefer to rely on that existing infrastructure rather than build a new one to match unfamiliar terminology in the rule.

We also ask the Department to confirm, in the alternative, that existing consumer review, dispute, complaint, servicing, and regulatory processes a Deployer already has in place can satisfy the Meaningful Human Review requirement, without requiring Deployers to build an entirely new, standalone review mechanism.

We also ask the Department to confirm two related points. First, that the commercial reasonableness factors in Rule 7.7.E account for the realities of automated, high volume credit decisioning, so the analysis reflects the marginal cost and feasibility of live review across large volumes of requests. Second, that the limitation on ADMT assisting a Meaningful Human Review does not bar decision support tooling used only to assemble, retrieve, and organize the primary evidence a human reviewer considers, as distinct from tooling that generates, scores, or re-derives the decision itself. Reviewers handling requests at scale need reliable tools to gather

1750 H Street NW, Suite 650, Washington, DC, 20006 ▲ 202.296.5544 ▲ www.afsaonline.org ▲ @AFSA_DC

the record in front of them, and barring those tools would make review slower and less accurate without making it any more independent.

Relatedly, Rule 7.7.C provides that reviewers should be independent of the original decision maker. Avoiding conflicts of interest is an important objective, but the rule should not require formal organizational independence. In institutions where thousands of consumer decisions are made each month, review is often conducted within established underwriting, servicing, quality assurance, or complaint handling functions. Requiring reviewers to sit in a separate reporting line from the original decision maker could require new governance structures that add little benefit to consumers. The rule should instead permit review by qualified personnel with the authority and expertise to conduct an objective review, regardless of where they sit within the organization.

We ask the Department to reconsider Rule 7.7.G.1's requirement that an Adverse Outcome be stayed, where possible, pending Meaningful Human Review. This mechanic does not translate well to decisions that are not readily reversible in the interim, such as a decision not to hire an applicant or not to extend an offer of credit for a specific purchase a consumer needs to make promptly. We also ask the Department to reconsider Rule 7.7.D.3's suggestion that a full reversal of the original decision indicates that human review was meaningful. A reversal could result from many things, including an ADMT that did not function correctly the first time, a reviewer who did not follow the rule's own standards, or a later automated process overriding the first. None of those circumstances says anything about whether the review that produced the reversal was itself meaningful, and we do not think the rule should suggest otherwise.

Finally, Rule 7.7.H requires detailed documentation of each Meaningful Human Review, including reviewer identity, authority, training, and timing. Financial institutions already maintain extensive documentation of decisions like these through existing systems, including complaint management, underwriting review, servicing, audit, quality assurance, and case management platforms. We ask the Department to clarify that existing business records may satisfy Rule 7.7.H, provided those records reasonably demonstrate the review performed and the outcome reached, rather than requiring institutions to build new and duplicative documentation processes solely for this purpose.

8. Protecting Proprietary Models, Fraud Controls, and Developer Trade Secrets

Rule 5.3 gives Developers a defined path to withhold trade secret information. Rules 7.3, 6.6, and 7.6 give Deployers no comparable protection when a consumer requests the Personal Data used in a Consequential Decision. Those rules call for the specific pieces of data considered, including score, rank, and all inputs, made available immediately through a hyperlink or within ten business days by other means. For an institution using a proprietary underwriting model, that means disclosing the model's inputs and outputs on request, with no mechanism to protect what is, for many of our members, a significant competitive asset built over years of development, iterative refinement, and investment.

The same level of detail that would expose a proprietary model also gives bad actors a working blueprint for reverse engineering it. A requirement designed to help a consumer understand a

1750 H Street NW, Suite 650, Washington, DC, 20006 ▲ 202.296.5544 ▲ www.afsaonline.org ▲ @AFSA_DC

decision made about them can, at that level of specificity, just as easily help someone trying to defeat a fraud model. These risks come without a corresponding benefit to consumers: a consumer is unlikely to be able to reverse engineer a creditor's scoring model on their own, and to the extent a model relied on inaccurate information, the consumer's existing right under the Fair Credit Reporting Act to obtain a copy of their credit report and dispute inaccurate information already provides a remedy.

We share this same concern in the context of fraud and consumer authentication protections, where the receipt of highly specific information could help fraudsters learn to circumvent fraud, security, and other models used by financial institutions to protect consumers. We ask the Department to more clearly document, consistent with Section 6-1-1704(7), that a Deployer is not required to make a disclosure, provide an explanation, or furnish information that would compromise fraud prevention, cybersecurity, or authentication protections.

We ask the Department to extend a trade secret and security sensitive information limitation to Rules 7.3 and 7.6 comparable to what Rule 5.3 already provides Developers, so Deployers can withhold or present in aggregate the specific inputs that would reveal proprietary scoring logic, while still giving consumers a meaningful, plain language explanation of the decision.

A related concern, described further in Section 4 above, is that requiring disclosure of a consumer's actual score and any automatic denial threshold under Rule 6.4.B.6 can itself create fair lending and safety and soundness tension. Publishing exact score cutoffs invites gaming and aids reverse engineering, and telling a consumer the precise score and numeric threshold that governed a decision adds little the consumer can act on beyond the specific principal reasons for the outcome. Federal law already sets a workable standard here: under ECOA and Regulation B, as the CFPB has confirmed applies even to creditors using complex algorithms, a creditor must provide accurate, specific principal reasons for an adverse action. We ask the Department to align the Adverse Outcome content standard with that established federal framework, rather than imposing a more granular, standalone Colorado standard that mandates disclosure of an exact score and an automatic denial cutoff.

This same trade secret concern extends upstream to Developer disclosures. The Notice asks whether the rules should require Developers to disclose any testing performed and its results, affirmatively state when no testing was performed, and specify a required level of granularity for describing training data categories. Many of our members sit on both sides of this question: they receive documentation from model vendors in their capacity as Deployers, and they are themselves Developers of models that must produce this kind of documentation for others. An overly granular or open ended disclosure mandate would compel disclosure of competitively sensitive and trade secret information without a corresponding benefit to the Deployers receiving it, and a rigid requirement to disclose all testing and complete results could expose proprietary validation methodologies. We ask the Department to adopt proportionate, trade secret protective disclosure standards: training data categories and known limitations should be described at a level of generality sufficient for a Deployer to use the model responsibly, a Developer should be able to describe testing at a summary level without disclosing proprietary methodologies or

complete results, and the trade secret protection in Rule 5.3 should apply across all Developer disclosures.

9. Personal Data Should Be Disclosed by Category, Not by Name

Rule 6.6.A.3 requires a Deployer to identify each source of Personal Data by name, including original sources and any intermediary or vendor sources in the chain. For financial institutions, this is not practical. Similarly, Rule 7.3 requires a Deployer to provide a Consumer with specific Personal Data (final rank, score, classification, etc.) and all ADMT inputs. Data may come from credit bureaus, fraud vendors, identity verification providers, affiliates, public records sources, and internal systems, and tracing each piece of data back to a named source, including any fourth party sources further up the chain, is often not something any single business in the chain is able to do. It would also create its own security, fraud, proprietary, and legal exposure: many of these vendor relationships are governed by confidentiality obligations, and identifying vendors by name or providing this data to a consumer in a consumer facing disclosure increases the attack surface for anyone trying to map out a business's data supply chain for fraudulent purposes.

The statute itself asks for a description of the types, categories, and sources of personal data, not the names or identities of those sources, and modern privacy frameworks at the state level and abroad generally recognize that long lists of vendor names do not meaningfully benefit consumers. We ask the Department to permit disclosure of sources and types of Personal Data by category rather than by name and specific pieces of Personal Data. This approach gives consumers meaningful information about where their data originates while avoiding data obligations that are not workable in practice.

10. A Consumer's Bare Assertion Should Not Override Third-Party Furnished Data

Rule 7.4.C.4 provides that where a Deployer did not receive Personal Data directly from the consumer and has no documentation supporting its accuracy, the consumer's assertion that the data is inaccurate is sufficient, by itself, to establish that it is inaccurate. Applied to financial services, this reaches the most common scenario in the industry: a creditor relies on a credit report or other data furnished by a consumer reporting agency and generally does not independently hold the underlying documentation for a given tradeline or public record. As drafted, the rule would require a creditor to treat a delinquency, a bankruptcy, or any other reported item as inaccurate based solely on a consumer's say so, and to reevaluate the ADMT decision on that basis.

We would also note some tension with the rule's own example at Rule 7.7.E.7, where a reviewer closes out a Meaningful Human Review because the ADMT functioned properly, without independently verifying the credit score behind the decision. It is not clear how that same rule can, elsewhere, require a creditor to treat an unverified consumer assertion as sufficient to override the same kind of third party furnished data.

This would require creditors to make credit decisions on a record they know may not be accurate, which cuts against safe and sound underwriting and could ultimately work against the consumer, by either forcing an extension of credit the consumer cannot repay or forcing a

reevaluation on a record no more reliable than the one it replaced. It also sits awkwardly next to a process that already exists for exactly this situation: FCRA gives consumers a direct right to dispute inaccurate information with the consumer reporting agency that furnished it, and that dispute process involves an actual investigation, not an unverified assertion standing in for a correction. Because the accuracy of bureau sourced credit data is already governed by the FCRA dispute process, a consumer's bare assertion should not, by itself, establish inaccuracy or compel re-decisioning of a credit outcome based on that data.

We ask the Department to revise Rule 7.4.C so that, for third party furnished data, a Deployer may direct the consumer to the applicable FCRA dispute process, or be given a reasonable period to investigate, rather than treating the consumer's unsupported assertion as conclusive.

We also ask the Department to clarify Rule 7.4's correction process to confirm that a Deployer's instructions may direct a consumer to the applicable consumer reporting agency's dispute process for third party furnished data, since the statute requires only that Deployers provide instructions for correction, not that a Deployer itself be the sole party that can act on a correction request.

11. Multiparty Arrangements, Program Managers, and Technology Vendors Need Workable Rules

The Notice asks whether the rules should define Deploy and ADMT Vendor, whether a Deployer that runs an ADMT through a vendor must fulfill every Deployer obligation itself, whether vendors must assist with consumer data access and correction requests, and whether vendor duties should track the processor obligations already established under the Colorado Privacy Act. In our members' businesses, consequential credit decisions are frequently made in arrangements where more than one party touches the ADMT. A rule that treated every participant as a full Deployer would create duplicative and potentially conflicting obligations, and could leave consumers receiving multiple, inconsistent notices about the same decision.

We ask the Department to adopt a rule that permits a single Deployer to be designated per Consequential Decision in a multiparty arrangement, so that one accountable party owns the consumer facing obligations for that decision, confirms that a Deployer may reasonably rely on an ADMT Vendor's documentation in meeting its own obligations, and adopts a bounded, processor-style-vendor-assist obligation, modeled on the Colorado Privacy Act's processor duties, under which an ADMT Vendor assists the Deployer in responding to consumer data access and correction requests rather than being treated as a full co-Deployer subject to the entire obligation set.

Relatedly, we ask the Department to clarify Rule 4.2's treatment of Midstream Developers. As drafted, it is not clear whether a party that simply inputs its own data into a third party's ADMT product has thereby engaged in the kind of configuration that triggers Midstream Developer obligations. We ask the Department to confirm that merely inputting data does not constitute configuration for these purposes, and that the obligation applies only where a party takes an affirmative step to adapt or configure the underlying tool itself.

1750 H Street NW, Suite 650, Washington, DC, 20006 ▲ 202.296.5544 ▲ www.afsaonline.org ▲ @AFSA_DC

12. Identity Verification and Response Timing Need to Reflect Operational Reality

Two features of the consumer rights rules will require our members to build new processes and pull fraud and identity protection teams into routine request handling. The rules require a Deployer to authenticate a consumer, return the specific pieces of Personal Data and inputs used in a decision, and, in the correction context under Rule 7.4.C, assess and potentially verify information. Verifying a requester's identity and then returning the exact data and inputs that drove a credit or fraud decision is not a clerical task. It necessarily involves the same fraud and identity teams that protect against account takeover and synthetic identity abuse, because a request channel that authenticates a person and then returns underlying decision data is itself an attractive target. This compounds the proprietary model and fraud control concerns described above, and means every request carries an operational and fraud review cost the rules do not appear to account for.

Building the technical infrastructure the rules contemplate is also a significant undertaking. Rule 6.4.C requires an Adverse Outcome disclosure to provide a request method through an online hyperlink or URL, and Rule 6.6 requires that, once a consumer uses that link or a toll free number, the required information about the ADMT and its inputs be made available immediately. Building secure, consumer specific request mechanisms, populating them with accurate, decision specific data on demand, and integrating them with authentication and case tracking systems is a substantial development effort that cannot be completed on the timeline currently contemplated. We ask the Department to confirm that a Deployer may apply authentication and fraud control measures proportionate to the risk of a given request before returning decision data, and to coordinate the consumer rights process with existing fraud and identity controls, rather than treating it as a separate channel with its own standalone requirements.

We separately ask the Department to revisit the specific response deadlines themselves. Rule 7.2 requires a Deployer to provide instructions within twenty four hours when a consumer uses a toll free number or address, and Rule 6.6 requires information about the ADMT and its inputs to be made available immediately upon use of a hyperlink or toll free number, or within ten business days by other means. These windows are not achievable for a regulated financial institution that must first verify identity and assemble decision specific data, particularly by mail, where a twenty four hour turnaround is not physically possible. They are also markedly shorter than the ten day acknowledgement and forty five day resolution timing the Department itself adopted for the broader consumer rights process under Rules 7.6 and 7.7, producing an internally inconsistent set of deadlines, with the most burdensome ones attached to exactly the requests that most need identity verification and fraud review.

We ask the Department to align all response windows to the same ten day acknowledgement and forty five day resolution structure, with the ability to extend for complex requests, to remove the twenty four hour and immediate availability requirements in Rules 7.2 and 6.6, and to confirm that response times are measured from the point at which a consumer's identity has been successfully authenticated. We would also ask the Department to confirm that Rule 7.5.E's fee prohibition does not require a Deployer to compensate a consumer for notarization costs in every

1750 H Street NW, Suite 650, Washington, DC, 20006 ▲ 202.296.5544 ▲ www.afsaonline.org ▲ @AFSA_DC

case where notarization is used for authentication purposes, since requiring a notarized affidavit is not itself charging a consumer a fee to exercise a right. Finally, we ask the Department to clarify Rule 7.6.D's reference to an appeal process under section 6-1-1306(3), which establishes an appeal right under the Colorado Privacy Act rather than the Colorado ADMT Act, so Deployers have a clear understanding of what, if anything, is required here.

13. The Consumer Communication Standards in Rule 3 Should Stay Workable

Rule 3.2 sets general standards for disclosures and communications to consumers, and several aspects of it, applied literally, would be difficult for any business to operationalize at scale. Requiring that a communication consider the vulnerabilities or unique characteristics of the specific consumer receiving it asks Deployers to do something they are generally not in a position to do for each and every consumer or user, particularly given the balance of the requirements already in place to ensure meaningful, accessible disclosures. Similarly, requiring that communications be provided in the language a business has previously used with a specific consumer would require tracking and notating the language of every past interaction with every consumer, which is not practical at scale, and could discourage businesses from offering service in additional languages as a courtesy in the first place, since doing so would create a new standing obligation going forward.

Finally, the requirement that a disclosure be readable on all devices is broader and vaguer than it needs to be. We ask the Department to adopt a more straightforward standard, such as requiring that a disclosure use a format that makes it readable, consistent with the approach other states have taken in analogous privacy rules.

14. Federal Preemption and Interaction with Federal Banking and Consumer Credit Law

More broadly, we ask the Department to confirm that compliance with the extensive federal framework already governing consumer credit decisioning, including ECOA and Regulation B's adverse action requirements and the FCRA's right to dispute and correct credit report information, should generally be treated as sufficient to satisfy the ADMT Act's parallel requirements, not simply as a starting point for a separate, standalone Colorado compliance regime. Further, any implementation of the ADMT Act and these rules may only be applied consistent with applicable federal consumer credit and banking law, including for national banks, and may not impose state-law requirements that prevent or significantly interfere with federally authorized banking powers (including credit operations, servicing and collections functions and related customer decisioning).

The Department should also clarify that the rules are not intended to require any notice, explanation, or disclosure where doing so would be prohibited by federal law, including the ECOA, Reg B, and the FCRA, would require disclosure of nonpublic personal information in a manner that would violate the Gramm-Leach-Bliley Act, or would compromise the confidentiality or integrity of federally required or regulated programs such as cybersecurity, fraud prevention, AML/CFT, or sanctions compliance.

1750 H Street NW, Suite 650, Washington, DC, 20006 ▲ 202.296.5544 ▲ www.afsaonline.org ▲ @AFSA_DC

Conclusion

AFSA is grateful for the Department's willingness to engage with the industry throughout this process, and for the changes already reflected in this draft. The recommendations above are aimed at the same goal the Department has already shown in this rulemaking: rules that protect consumers without imposing burdens the statute never called for. Adopting Standard 2 for Materially Influence, conforming the Adverse Outcome examples and disclosure content to the statute's own terms, preserving and clarifying the scope of the federal model form safe harbor, allowing sensible delivery and timing standards, calibrating Meaningful Human Review to how credit is actually extended, protecting proprietary models and trade secrets from disclosure, permitting source disclosures by category, grounding the correction standard in the dispute processes that already exist, providing workable rules for multiparty arrangements and technology vendors, clarifying the interaction between these rules and existing federal law, and allowing an adequate implementation period would go a long way toward that goal.

We welcome the chance to discuss any of the above further before the rules are finalized. If you have any questions or would like to discuss this further, please do not hesitate to contact me at erayhan@afsamail.org or (805) 501-8873 at your convenience.

Sincerely,

Elora Rayhan
State Government Affairs
American Financial Services Association
1750 H Street, NW, Suite 650
Washington, DC 20006-5517
erayhan@afsamail.org