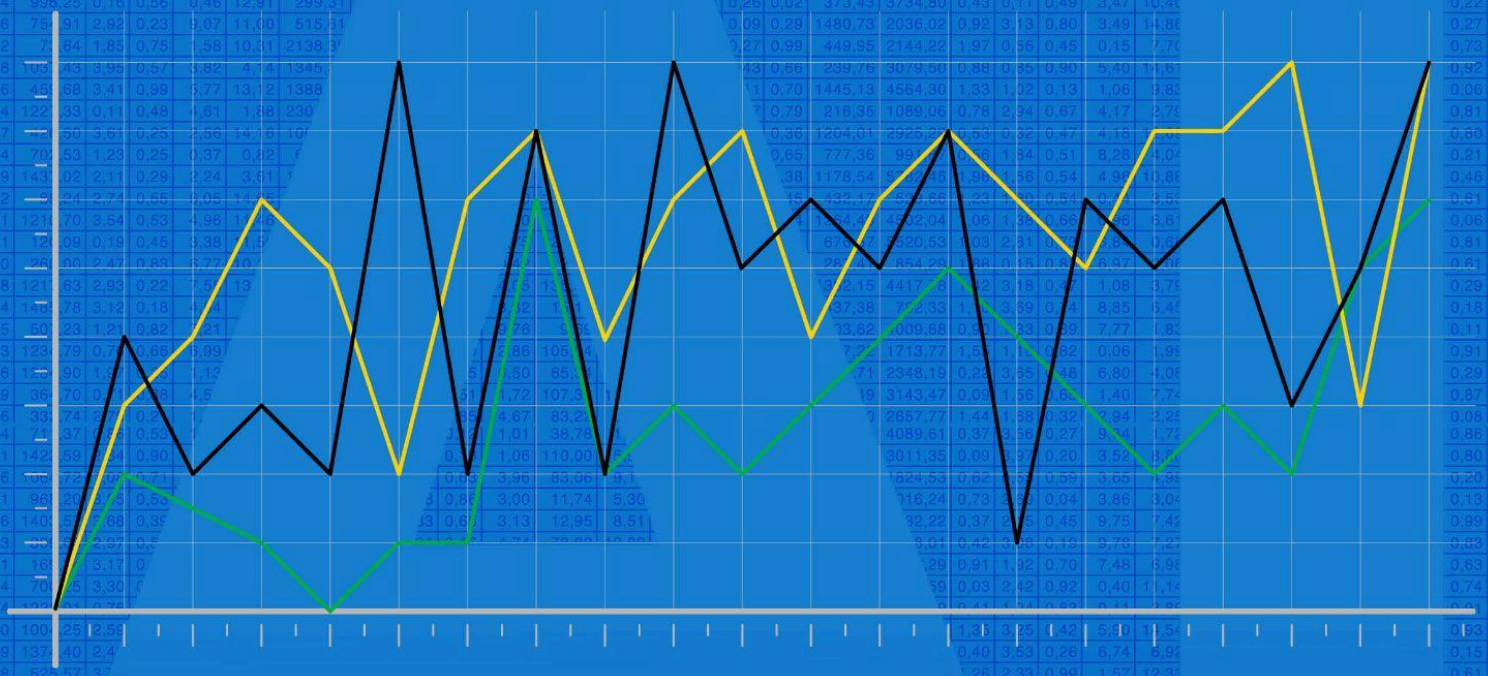

Consumer Welfare in the Age of Generative AI

From Theory to Practice

A White Paper by Consumer Reports
Digital Marketplace Initiative

Supported by the Alfred P. Sloan Foundation

MAY 2026

CR Consumer Reports

Table of Contents

Executive Summary

Artificial intelligence is no longer a feature of consumer financial products and services—it is rapidly becoming their foundation. From credit decisioning and fraud detection to customer service and investment advice, AI systems now shape how millions of Americans access, use, and are affected by financial services every day. This transformation is unfolding at a pace that has outstripped the capacity of existing consumer protection frameworks to respond.

This white paper presents the case that the current moment demands a fundamental reframing of consumer welfare in AI-mediated markets—and offers a practical path forward. Drawing on Consumer Reports' 2025 nationally representative AI in Financial Services Survey of 4,073 U.S. adults,[1] a comprehensive landscape analysis of existing frameworks and regulations, and CR's decade of testing and advocacy in digital financial services, this paper does three things.

First, it documents the gap between what consumers need and what the current system delivers. Our nationally representative survey reveals a public that is broadly aware of AI in financial services, is deeply skeptical of its deployment, and perceives current law as leaving it largely unprotected. Fifty-seven percent of Americans say existing laws are inadequate to protect them from AI-related risks in financial services.[2] Fewer than one in 10 completely trust financial companies to use AI responsibly.[3] And when asked to weigh the benefits against the risks, 38% say the risks outweigh the benefits, while another 39% remain unsure.[4]

Second, it introduces a Consumer Welfare Standard. This is a practical, obligation-based framework organized around five categories: Information Integrity, Fair Treatment, Consumer Control, Accountability and Remedy, and Systemic Responsibility. The standard translates established consumer-protection principles into specific obligations that AI systems owe to the people they serve. It moves beyond abstract principles to define what responsible AI looks like in operation.

Third, it maps the policy landscape and identifies concrete opportunities for action. Existing federal statutes—the Equal Opportunity Credit Act (ECOA), the Fair Credit Reporting Act (FCRA), the Unfair or Deceptive Acts or Practices, and the Unfair, Deceptive, or Abusive Acts or Practices (UDAP/UDAAP)—continue to apply to AI-driven decisions but were not

---

[1] Consumer Reports nationally representative AI in Financial Services Survey of 4,073 U.S. adults September-October 2025, https://article.images.consumerreports.org/image/upload/v1767035543/prod/content/dam/surveys/Consumer_Reports_AI_in_Financial_Services_2025.pdf.

[2] Consumer Reports nationally representative AI in Financial Services Survey 2025. In the survey, 57.15% answered "No" to "Do you believe current government laws and regulations adequately protect you from the risks of AI in financial services?"

[3] Consumer Reports nationally representative AI in Financial Services Survey 2025: "Completely trust" ranges from 2.74% (crypto) to 8.80% (bank account monitoring). No category exceeds 9%.

[4] Consumer Reports AI in Financial Services Survey 2025: Risks outweigh benefits 38.00%; Benefits outweigh risks 22.88%; Unsure 39.12%.

designed to address the opacity, dynamism, and systemic reach of machine learning systems. State-level innovation in Colorado, California, and New York has begun to fill these gaps, but the result is regulatory fragmentation rather than coherent protection. This paper identifies specific policy interventions that policymakers, regulators, and industry leaders can pursue to close the most critical gaps. AI is no longer a product feature, but is increasingly functioning as core market infrastructure. As a result, traditional consumer protection models are insufficient, and an integrated, consumer-centered standard is needed—one that treats AI not as an exception to existing rules, but as a reason to strengthen them.

The stakes are both urgent and consequential. The window for establishing consumer-centered standards is narrowing as industry practices harden and markets consolidate. Whether the transformation of financial services by AI serves consumers or undermines them will depend on the choices made now—by companies, regulators, and advocates alike.

Introduction: Why Consumer Welfare Must Be Reframed

From Product Safety to Systemic Market Mediation

For most of its history, consumer protection has operated on a product-safety model: identify a defective product, warn consumers, compel a fix or recall. This model assumes that harm is traceable to a specific product, that consumers can identify and avoid the risk, and that the product itself is relatively stable over time.

AI in financial services breaks each of these assumptions. AI systems are not static products—they are dynamic, adaptive systems that change their behavior based on the data they process. Harm is not isolated to a single defective feature but can be embedded in the logic of how a system makes decisions, assigns prices, evaluates creditworthiness, or routes customer-service interactions. Consumers typically cannot identify when AI is being used, let alone assess whether its deployment serves their interests. And the products themselves are continuously updated, retrained, and modified, often without the consumer's knowledge.

AI is now market infrastructure. It shapes which consumers see which products, at what price, under what terms, and with what level of service. When AI mediates the marketplace at this scale, the traditional question—"Is this product safe?"—is necessary but no longer sufficient. The question must become: "Does this system treat consumers fairly, transparently, and accountably?"

From Identifying Harms to Operationalizing Protections

The public conversation about AI and consumers has evolved substantially since 2022. Early discourse rightly focused on cataloguing potential harms—bias in credit decisioning, privacy erosion, chatbot unreliability, job displacement. This work was essential: It established that AI in financial services is not benign by default and that consumers bear real costs from irresponsible deployment.

But the field has reached an inflection point. The task now is not simply to describe risks, but to define standards, build evaluation capacity, and create implementation pathways. This is the work of operationalizing consumer protection in AI-mediated markets. It requires moving from principles to metrics, from frameworks to scorecards, and from general concerns to specific, enforceable obligations.

This white paper is designed to advance that transition. It is the policy companion to Consumer Reports' landscape analysis of AI frameworks and regulation in consumer financial products and services, which maps the diagnostic terrain—existing frameworks, as well as regulatory activity and gaps. But where the landscape analysis asks, "What consumer protections exist

and where are the gaps?", this white paper asks, "What protections should exist, and what must policymakers, regulators, and industry leaders do to put them in place?"

The Scope and Urgency of This Moment

Several developments make this paper's publication timely. The AI supply chain is consolidating rapidly. A small number of cloud hyperscalers—Amazon Web Services, Microsoft Azure, and Google Cloud Platform—control the infrastructure on which virtually all AI in financial services depends.[5] The companies developing the most powerful foundation models are increasingly integrated with the platforms that deploy them. Market concentration of this kind creates systemic risk and limits the leverage that any single financial institution—let alone any individual consumer—can exercise.

At the same time, the U.S. regulatory landscape is fragmented and, in key respects, retreating. The revocation of executive orders on AI safety, combined with ongoing debates about the scope of federal agency authority, has created uncertainty about the direction of federal policy.[6] State-level innovation continues—Colorado's AI act,[7] California's proposed Automated Decisions Safety Act, and New York's regulatory guidance each represents important advances—but the result is a patchwork that neither consumers nor companies can easily navigate.

Meanwhile, financial institutions are accelerating their AI adoption. According to government reports, AI is increasingly used across financial services, including for customer service, fraud detection, credit underwriting, and compliance functions.[8] New entrants—from fintech startups to Big Tech platforms—are launching AI-native financial products that operate outside traditional regulatory perimeters. The gap between the pace of deployment and the pace of protection is widening.

The window for establishing consumer-centered standards is narrowing. First-mover advantage in defining what "responsible AI" means for consumers is time-limited. If independent, consumer-centered standards are not established now, the standards that emerge will be set by the industry itself—on terms that serve its interests, not the public's.

---

[5] CR, AI in Financial Services Landscape Analysis, 2026, Market Overview section; also Financial Stability Board, "The Financial Stability Implications of Artificial Intelligence," November 2024; also Mazzucato et al., "How Big Cloud becomes Bigger," *SSRN*, 2025.
[6] EO 14110, "Safe, Secure, and Trustworthy Development and Use of AI," (signed October 30, 2023, revoked by EO 14148 on January 20, 2025).
[7] CO SB 24-205 (signed May 2024); CA AB 1018 (proposed); NY DFS Circular Letter CL2024-07, January 2024).
[8] GAO-25-107197, "Artificial Intelligence: Use and Oversight in Financial Services," May 2025, 8-9; Financial Stability Oversight Council, 2024 Annual Report, 83-85.

Consumer Welfare Theory in AI-Mediated Markets

Limits of Traditional Consumer Welfare Models

The traditional framework for analyzing consumer welfare in financial markets rests on several foundational assumptions: that consumers have access to sufficient information to make informed decisions; that the incentives of financial institutions are adequately aligned with consumer interests through market competition; that products and decision-making processes are sufficiently transparent for meaningful consumer choice; and that market structure is competitive enough to discipline firms that fail to serve consumers well.

AI-mediated markets challenge each of these assumptions in ways that existing consumer welfare theory has not fully reckoned with.

Information asymmetry is structurally intensified. In traditional financial services, consumers face information disadvantages—they may not understand the terms of a credit card agreement or the full cost of a mortgage. But the product itself is, in principle, describable. An AI system that uses machine learning to set insurance premiums based on driving behavior, social media activity, or purchasing patterns creates a qualitatively different kind of opacity. The system's decision logic may not be fully understood even by the company that deploys it, let alone the consumer affected by it.

Incentive misalignment is amplified. AI systems optimize for the objectives their developers specify. In financial services, these objectives are overwhelmingly commercial: maximize revenue per customer, minimize loss exposure, reduce customer service costs. Consumer welfare—defined as the consumer's actual financial health, informed autonomy, and fair treatment—is rarely an explicit optimization target. The result is a structural misalignment between what AI systems are designed to do and what consumers need them to do.

Algorithmic opacity undermines meaningful choice. Consumers cannot exercise informed choice about a system whose operations they cannot observe, understand, or meaningfully challenge. When an AI system denies a loan application, the adverse-action notice required under existing law was not designed to explain the output of a machine learning model trained on hundreds of variables. The "reasons" provided often fail to give consumers actionable information about what they can do differently.

Structural concentration limits competitive discipline. Market concentration in the AI supply chain—from foundation models to cloud infrastructure—means that the competitive dynamics that traditionally protect consumers are weakened. When multiple financial institutions rely on the same underlying models and infrastructure, the diversity of approaches that competition is supposed to produce is eroded.

AI as a Source of Market Failure

These dynamics constitute forms of market failure that traditional consumer-welfare analysis must be updated to address. When consumers cannot observe the decision-making process, cannot meaningfully choose alternatives, and cannot hold firms accountable for outcomes, the market cannot be said to be functioning in consumers' interest—regardless of whether the products on offer are cheaper or faster than their predecessors.

Speed and efficiency—the benefits most frequently cited by AI proponents—are necessary but radically insufficient measures of consumer welfare. Data from a nationally representative 2025 Consumer Reports survey of 4,073 Americans confirms this: when asked what matters most to them in improving the financial services they use, strengthened data security (49%) and lower costs (44%) topped the list, but better customer service (42%), higher accuracy (26%), reduced bias (22%), and greater access to credit (19%) also registered as priorities.[9] Consumers do not define welfare solely in terms of speed; they define it in terms of trust, fairness, and the quality of the relationship with the institutions that manage their money.

Redefining Consumer Welfare for AI Systems

A consumer welfare standard adequate to AI-mediated markets must incorporate dimensions that traditional efficiency-based models have treated as secondary or externalized entirely. Specifically, it must measure:

- **Accuracy and reliability:** Does the system produce correct outputs, and does it communicate its level of confidence honestly?
- **Fairness:** Does the system treat similarly situated consumers equally, and does it avoid reproducing or amplifying patterns of discrimination?
- **Agency:** Does the consumer retain meaningful control over decisions that affect their financial life, including the ability to opt out of AI-mediated processes?
- **Redress:** When the system produces a harmful outcome, can the consumer identify the harm, understand its cause, and obtain a remedy?
- **Resilience:** When the system or its supporting infrastructure fails, does it fail safely and transparently, rather than silently degrading in ways that harm consumers without their knowledge?

[9] Consumer Reports nationally representative AI in Financial Services Survey of 4,073 U.S. adults, September-October 2025, https://article.images.consumerreports.org/image/upload/v1767035543/prod/content/dam/surveys/Consumer_Reports_AI_in_Financial_Services_2025.pdf.

These dimensions are not aspirational additions to existing consumer-welfare analysis. They are the minimum conditions under which a consumer can be said to be served—rather than merely processed—by an AI-mediated financial system. The standard proposed in Section IV operationalizes each of these dimensions.

What Consumers Are Experiencing

In 2025, Consumer Reports conducted a nationally representative survey of 4,073 U.S. adults to assess public awareness, experience, attitudes, and expectations regarding AI in financial services.[10] The results suggest a public that is more aware of AI's presence in their financial lives than is commonly assumed—and substantially more skeptical and concerned than industry narratives would suggest.

The Trust Deficit

The most striking finding from our survey is the depth and breadth of public distrust in AI-mediated financial services. This distrust is not limited to particular demographics or to people who lack familiarity with AI; it is pervasive.

57%	Percent of Americans who say current government laws and regulations do not adequately protect them from potential harms of AI in financial services.
<9%	Percent of Americans who completely trust financial companies to use AI responsibly across any financial service category.
42%	Percent of Americans who hold a negative overall view of the growing use of AI in financial services; only 18% are positive.

These numbers are not reflective of ignorance or technophobia. Fifty-five percent of consumers report having actively interacted with AI in the previous three months. Seventy-two percent believe financial companies are already using AI—either for a long time (38%) or recently (34%). Americans are aware that AI is present in their financial lives. They simply do not trust that it is being used well.

[10] Consumer Reports nationally representative AI in Financial Services Survey of 4,073 U.S. adults, September-October 2025, https://article.images.consumerreports.org/image/upload/v1767035543/prod/content/dam/surveys/Consumer_Reports_AI_in_Financial_Services_2025.pdf.

When asked what would increase their trust in a financial company's use of AI, consumers pointed overwhelmingly to structural protections rather than corporate reassurances. These include: the option to opt out of AI-based decisions (53%); independent audits of AI systems (51%); a simple way to dispute or appeal AI decisions (50%); and public reporting on AI accuracy and bias testing (46%). Consumers know what they need. The question is whether policymakers are prepared to listen.

The Priority Mismatch

Our survey reveals a significant gap between what consumers prioritize and what the AI industry emphasizes. When asked to identify the most important improvements that could be made to the quality of the financial services they use, consumers most commonly selected strengthened data security (49%), followed by lower costs (44%) and better customer service (42%). Faster decisions—the benefit most frequently marketed by AI providers—ranked eighth among consumers' priorities, chosen by only 18% of Americans.

This priority mismatch matters for policy. Industry arguments for AI adoption have centered on efficiency, speed, and scale. But consumers are signaling that these are not the dimensions on which they evaluate whether AI serves their interests. They want to know that their data is safe, that they are being treated fairly, that they can reach a human being when something goes wrong, and that someone independent is checking whether the system works as advertised.

When asked about the greatest risks of financial companies using AI, consumers identified errors introduced by AI (45%), job losses (41%), weakened data security (36%), worse customer service (32%), and decisions that are difficult to understand or challenge (28%). Predatory targeting for risky products was flagged by one in four Americans (25%). These are not abstract concerns; they reflect informed consumer perceptions shaped by growing exposure to AI in financial services.

The Experience Gap

Consumer experience with AI in financial services varies dramatically by use case—and the pattern is revealing. The applications for which consumers report the most positive experiences are those that use AI in the background to protect them. Among Americans who had encountered AI in fraud alerts, 70% reported a generally positive experience. But the picture shifts sharply when AI is deployed in consumer-facing, decision-making roles.

Among consumers who had interacted with AI customer service chatbots—the most commonly encountered application—only 36% reported a generally positive experience, while 42% reported a negative one. For automated credit decisions, the split was 37% positive versus 34% negative. For personalized pricing, opinion was essentially evenly divided: 37% positive, 36% negative.

The comfort gap is even more pronounced when consumers are asked about hypothetical deployments. While 75% of Americans are comfortable with AI recommending a driving route, and 59% with AI fraud detection on their accounts, comfort declines precipitously for higher-stakes financial decisions. Only 18% are comfortable with AI approving or denying mortgage applications, 31% with AI setting loan interest rates, and only 21% with loan offers targeted based on location or social media activity.

Policy implication: These data undermine the premise that expanding AI in financial services inherently serves consumers. Consumers themselves draw sharp distinctions between protective uses of AI (fraud detection) and autonomy-displacing uses (credit decisions, pricing, targeted offers). Policy frameworks should reflect these distinctions—imposing heightened scrutiny, transparency requirements, and human-oversight mandates on high-stakes AI deployments rather than treating all AI applications equivalently.

Agentic AI: The Next Frontier of Consumer Risk

Our survey also explored consumer attitudes toward agentic AI—systems that can make decisions and take actions on behalf of consumers autonomously, without requiring approval for each step. The industry is rapidly moving toward this frontier, and consumer sentiment should give policymakers pause.

A majority of consumers expressed discomfort with every agentic AI scenario described. Only 32% were at least somewhat comfortable with an AI agent transferring money between accounts to avoid fees, 31% with an AI agent switching insurance plans, and just 15% were comfortable with an AI agent making a decision to apply for credit on their behalf.

Policy implication: The emergence of agentic AI in financial services creates an urgent need for new regulatory frameworks that did not exist even two years ago. Existing consumer protection law is built on the premise that a consumer affirmatively initiates a financial transaction. When AI agents can initiate, execute, and complete transactions autonomously, the foundational concepts of consent, authorization, and liability must be reconsidered. Policymakers should act now—before agentic AI is widely deployed—to establish clear rules around authorization, human override, liability allocation, and fiduciary duty.

The Consumer Welfare Standard: What AI Systems Owe Consumers

This section presents the core contribution of this paper: A Consumer Welfare Standard that defines what AI systems deployed in consumer financial services owe to the people they serve. This standard is organized around five obligation categories, each containing specific dimensions that can be measured, tested, and enforced.

This standard draws on Consumer Reports' landscape analysis of existing evaluation frameworks, regulatory regimes, and gaps across 12 axes of evaluation. It translates those axes—originally developed as diagnostic tools for mapping the current landscape—into an obligations-based framework designed for implementation. Where existing frameworks identify what could be measured, this standard specifies what must be.

Category 1: Information Integrity

AI systems in financial services must provide consumers with information that is accurate, honestly calibrated, and understandable. This obligation has three dimensions:

Accuracy and Reliability

AI systems must produce outputs that are factually correct and consistent. In financial services, this means, for example, credit decisions based on accurate assessments of risk, chatbot responses that correctly describe product features and company policies, and investment recommendations grounded in sound analysis. Our landscape analysis found that while frameworks exist for measuring accuracy in machine-learning models, comprehensive frameworks for measuring accuracy in the consumer-facing outputs of generative AI systems remain limited.

Policy opportunity: Regulators should require financial institutions to publish accuracy metrics for AI systems that interact directly with consumers or make decisions affecting consumer outcomes. These metrics should be specific, auditable, and comparable across institutions. The proposed federal Algorithmic Accountability Act's documentation requirements and Colorado's impact-assessment obligations provide starting points, but neither currently mandates the kind of consumer-facing accuracy reporting that meaningful accountability requires.

Confidence Calibration

AI systems must communicate their level of certainty honestly, avoiding the pattern that researchers have termed "sycophancy"—the tendency of AI to tell users what they want to hear, regardless of factual basis. This is not a theoretical concern. In April 2025, OpenAI rolledback a major update to ChatGPT following concerns about overly agreeable or "sycophantic"

behavior.[11] Within financial services, sycophancy poses risks that range from customer service chatbots describing product features that do not exist to AI financial advisors reinforcing poor investment decisions rather than providing honest guidance.

Policy opportunity: Our landscape analysis identified a critical gap: no U.S. law, regulation, or guidance specifically names sycophancy as a risk in consumer financial services. But that does not mean the legal cupboard is bare. Existing doctrines—including suitability obligations, fiduciary duties, and prohibitions on unfair, deceptive, or abusive acts or practices (UDAP/UDAAP)—could in principle reach many sycophancy-related harms. For example, an AI system that validates a consumer's risky financial decision rather than providing honest guidance, or that fabricates product features to maintain engagement, may already be engaging in deceptive or abusive conduct under current law. The gap is not in legal authority but in recognition or application. No regulator has tested these doctrines against AI sycophancy, and no enforcement action has established that over-confident or validation-based AI output violates existing standards.

Policymakers and regulators should close this gap on two fronts. First, the Federal Trade Commission (FTC) and Consumer Financial Protection Bureau (CFPB) should issue interpretive guidance clarifying that AI systems exhibiting sycophantic behavior in consumer financial contexts may violate existing prohibitions on deceptive and abusive practices—and should bring enforcement actions that establish this principle. Second, regulators and industry groups should develop AI-specific anti-sycophancy metrics, testing protocols, and reporting requirements that go beyond existing doctrines to address the distinctive risks of confidence-miscalibrated AI. Independent evaluation organizations should compare the relative sycophancy of consumer-facing AI financial products using established research methods, creating the kind of comparative information that helps consumers make informed choices.

Transparency and Explainability

Consumers have a right to understand how AI systems make decisions that affect them. In our survey, 74% of Americans strongly agreed that they want companies to disclose when AI is involved in the financial services they use. When asked what information they would want if an AI system denied their loan application, 63% wanted a list of all factors and how much each mattered—not just a generic reason.[12]

Policy opportunity: Existing explainability requirements in federal law—principally the adverse-action notice requirements under ECOA and FCRA—were not designed for machine

[11] OpenAI blog, "Sycophancy in GPT-4o," April 30, 2025.
[12] Consumer Reports nationally representative AI in Financial Services Survey of 4,073 U.S. adults, September-October 2025, https://article.images.consumerreports.org/image/upload/v1767035543/prod/content/dam/surveys/Consumer_Reports_AI_in_Financial_Services_2025.pdf.

learning systems.[13] Colorado's AI Act requires post-decision explanations for high-risk decisions, and the proposed New York AI Consumer Protection Act would require explanations of algorithmic decision-making logic. Policymakers should strengthen these requirements to mandate not merely disclosure that AI is used, but meaningful explanation of how AI contributed to a specific decision, calibrated to the consumer's level of understanding and the stakes involved.

Category 2: Fair Treatment

Bias and Non-Discrimination

AI systems must not discriminate against consumers on the basis of race, ethnicity, gender, age, disability, or other protected characteristics—whether directly or through proxy variables. Seventy-five percent of Americans are concerned that AI in financial services could result in discrimination based on characteristics like race or income, with 31% very concerned. Yet when asked whether using AI in financial services will reduce or increase bias, more Americans believed it would increase bias (26%) than reduce it (18%), with the largest group (39%) unsure.[14]

Consumer expectations around accountability are clear: 52% believe financial companies should be required to investigate the cause of any bias, 50% say independent organizations should be allowed to test their systems, 49% want bias-testing results reported publicly, and 48% say affected customers should be compensated.

Policy opportunity: Existing anti-discrimination law—particularly ECOA and the Fair Housing Act—applies to AI-driven decisions but was not designed to detect or remedy the forms of bias that machine learning systems can introduce through training data, feature selection, or proxy discrimination. Colorado's ban on algorithmic discrimination, the proposed federal AI Civil Rights Act, and state-level hiring audit requirements like New York City's Local Law 144 all represent important advances. Policymakers should extend these frameworks to mandate: (a) pre-deployment bias testing for all high-risk AI financial services; (b) ongoing bias monitoring at a frequency consumers support (58% say continuously, 23% say at least annually[15]); (c) independent third-party auditing access; and (d) mandatory compensation for consumers harmed by biased AI decisions.

[13] CFPB Consumer Financial Protection Circular 2023-03, September 19, 2023.
[14] Consumer Reports nationally representative AI in Financial Services Survey of 4,073 U.S. adults, September-October 2025, https://article.images.consumerreports.org/image/upload/v1767035543/prod/content/dam/surveys/Consumer_Reports_AI_in_Financial_Services_2025.pdf.
[15] ibid.

Market Fairness

AI systems must not be used to manipulate consumers through personalized pricing that exploits individual data profiles, deceptive design patterns that steer consumers toward products that serve the company's interest rather than the consumer's, or information environments that are optimized for engagement rather than accuracy. Investigations by Consumer Reports and others into algorithmic pricing at Instacart—and data-driven discounting practices at Kroger—have documented how opaque pricing and personalization systems can lead to consumers being charged inconsistent or less favorable prices without their knowledge.[16]

Policy opportunity: The United States has no broad prohibition on algorithmic price discrimination in financial services. China's algorithmic recommendation rules prohibit "unreasonable differential treatment in transaction prices or other transaction conditions" based on user's preference—a standard that has no direct analogue U.S. law governing AI-driven pricing.[17] The FTC has targeted deceptive design patterns but has not established clear rules around AI-driven personalized pricing. Policymakers should consider whether AI-driven price personalization in financial products requires new disclosure obligations, non-discrimination standards, or outright prohibitions in categories where consumers cannot reasonably detect or avoid differential treatment.

Deceptive Design in AI Financial Interfaces

AI-mediated financial services introduce new vectors for deceptive design—the use of interface choices that steer consumers toward outcomes that serve the company's interests rather than the consumer's. In traditional digital products, deceptive design patterns (sometimes called "dark patterns") manifest in manipulative cancellation flows, confusing consent interfaces, and hidden fees. AI amplifies these risks by enabling personalized manipulation: systems that adapt their presentation, framing, or urgency cues based on individual behavioral profiles.

In the specific context of consumer financial services, deceptive design in AI systems can appear in onboarding and enrollment flows that obscure the role of AI in decision-making; consent interfaces that make opting into data collection the path of least resistance while burying opt-out mechanisms; explanation interfaces that present adverse-action notices or AI-generated justifications in technically accurate but practically useless formats; and product-recommendation engines that frame higher-margin products as "personalized" matches without disclosing the commercial incentives driving those recommendations.

---

[16] Consumer Reports" "Instacart's AI-Enabled Pricing Experiments", December 9, 2025; "Instacart Stops Pricing Tests," January 8, 2026; "Consumer Reports Investigation Uncovers Kroger's Widespread Data Collection…" May 21, 2025.
[17] China, Provisions on the Management of Algorithmic Recommendations in Internet Information Services (effective March 2022).

Policy opportunity: Several states have begun to address deceptive design. Texas law prohibits dark patterns to obtain consumer consent for data processing±an approach that could be extended to AI interaction disclosures.[18] California[19] and Colorado[20] privacy rules specify that consent obtained through dark patterns is invalid. At the federal level, the CFPB has identified manipulative choice architecture as potentially "abusive"[21] under its enforcement authority. But no jurisdiction has enacted rules specifically targeting AI-enabled personalized manipulation in financial services. Policymakers should: (a) extend existing deceptive design prohibitions to explicitly cover AI-personalized interfaces; (b) require that consent, opt-out, and explanation interfaces meet accessibility and comprehension standards that are verified through independent testing; (c) establish that explanations of AI decisions must be not merely provided but accessible—meaning presented in formats, reading levels, and interface locations where consumers will actually encounter and understand them; and (d) fund independent comparative evaluations of deceptive design patterns across consumer financial AI products, creating market pressure for clearer, fairer interfaces.

Category 3: Consumer Control

Agency and Opt-Out

Consumers must retain meaningful control over whether and how AI systems are used in their financial lives. According to our survey, 77% of Americans believe consumers should have the right to opt out of AI-based decision-making in financial services.[22] This is not a marginal preference; it is a supermajority consensus that crosses demographic and political lines.

Similarly, our survey shows that if an AI system made an error that affected them, consumers have clear expectations about remedy: 48% would want a human to review the decision; 47% want a clear explanation of what went wrong; 44% want corrections made; and 40% want compensation for any harm. These expectations describe a system of meaningful human oversight and accountability that does not currently exist in most AI-mediated financial services.

---

[18] Texas Data Privacy and Security Act, Tex. Bus. & Com. Code § 541.051 et seq., 2023.
[19] California Code of Regulations, tit. 11, § 7004, 2023, (invalidating consent obtained through dark patterns).
[20] Colorado Privacy Act Rules, 4 CCR 904-3, Rule 7.03, 2023.
[21] Consumer Financial Protection Bureau, Policy Statement on Abusive Acts or Practices, April. 3, 2023; CFPB, Request for Information Regarding Uses of AI in Financial Services, 2024.
[22] Consumer Reports nationally representative AI in Financial Services Survey of 4,073 U.S. adults, September-October 2025, https://article.images.consumerreports.org/image/upload/v1767035543/prod/content/dam/surveys/Consumer_Reports_AI_in_Financial_Services_2025.pdf.

Policy opportunity: Several states[23]—including California, Colorado, Connecticut, Montana, Nebraska, New Jersey, Oregon, and Texas—provide consumers with rights to limit the use of their data, including opting out of profiling used in significant decisions and, in some cases, the sharing or sale of personal data.[24] But none of these provisions empower consumers to opt out of AI use entirely in the context of their finances. Current opt-out mechanisms also place the burden on consumers, requiring them to navigate complex processes to exercise rights they may not even know they have. Policymakers should: (a) establish universal opt-out rights for AI-based decision-making in high-stakes financial contexts; (b) require companies to make opt-out mechanisms simple, accessible, and prominently disclosed; (c) consider opt-in requirements for the highest-stakes applications, such as AI-driven credit decisions and insurance underwriting; and (d) mandate meaningful human-review pathways for any consumer who disputes an AI-mediated decision.

Privacy and Data Minimization

Consumers' willingness to share personal data with AI financial services is conditional, not categorical. In our survey, 34% of Americans said they would not, under any conditions, share personal information with AI financial services for the purpose of offering more useful or relevant products and services. Among those willing to share data, the most important conditions were that the company does not share or sell data to anyone else (44%); the ability to delete data or stop AI from using it (44%); human involvement in important decisions (37%); and clear explanations of data use (36%).[25]

The types of data consumers are comfortable sharing with AI to help make decisions about things like loan approvals or insurance pricing are notably limited. While 36% are comfortable sharing employment history, only 10% are comfortable sharing browser history, medical history, or social media activity.[26]

---

[23] Colorado Privacy Act, Colo. Rev. Stat. § 6-1-1306(1)(a)(I)(C): 137, Right to opt out of profiling in decisions producing legal or similarly significant effects; Connecticut Data Privacy Act, Conn. Gen. Stat. § 42-518(a)(5)(C), opt-out of profiling tied to significant decisions; Delaware Personal Data Privacy Act, Del. Code tit. 6, § 12D-104(a)(6)(C), Right to opt out of profiling in decisions producing legal or similarly significant effects concerning the consumer; Montana Consumer Data Privacy Act, Mont. Code Ann. § 30-14-2808(1)(e)(iii), opt out of … profiling in furtherance of automated decisions that produce legal or similarly significant effects concerning the consumer; New Jersey Data Privacy Act, N.J. Stat. § 56:8-166(a)(3), right to opt out of … some types of profiling (for example, profiling to determine whether a consumer should receive a loan or mortgage, a job offer, or an insurance policy); Oregon Consumer Privacy Act, Or. Rev. Stat. § 646A.574(1)(c); Texas Data Privacy and Security Act, Tex. Bus. & Com. Code § 541.051(c)(3).

[24] National Conference of State Legislatures, AI legislation tracker Summary.

[25] Consumer Reports nationally representative AI in Financial Services Survey of 4,073 U.S. adults, September-October 2025, https://article.images.consumerreports.org/image/upload/v1767035543/prod/content/dam/surveys/Consumer_Reports_AI_in_Financial_Services_2025.pdf.

[26] ibid.

Policy opportunity: The absence of a comprehensive federal data privacy law remains the most significant structural gap in consumer protection in AI-mediated financial services. Existing sector-specific laws—GLBA and FCRA—do not address the volume and types of data that AI systems can ingest and process. Policymakers should advance comprehensive privacy legislation that includes: (a) data-minimization requirements limiting AI systems to data reasonably necessary for the financial service provided; (b) purpose-limitation rules preventing data collected for one service from being repurposed for AI training or cross-product targeting; (c) meaningful consent requirements calibrated to the sensitivity of the data and the stakes of the decision; (d) robust deletion rights that address the distinctive challenge of data embedded in AI systems—where personal information may persist not only in databases but also in model weights, training sets, and derived inferences that resist conventional deletion methods; and (e) requirements that financial institutions provide consumers with practical, accessible deletion mechanisms and confirm that deletion has been effectuated, recognizing that the right to delete data is meaningful only if the process for exercising it is simple and verifiable. Current deletion regimes vary significantly across jurisdictions, place the burden of action on consumers, and generally do not address data that has been incorporated into AI model training. Policymakers should close this gap by requiring that AI-specific deletion obligations apply to data in all forms—stored, processed, and learned.

Protections for Minors in AI-Mediated Financial Services

An increasing number of financial products and services reach minors directly—through teen banking apps, payment platforms embedded in gaming and social media, budgeting tools marketed to young people, and custodial investment accounts with AI-driven features.[27][28] Use of digital financial tools is already widespread among younger consumers, with adoption highest among younger age groups—suggesting continued expansion into earlier life stages.[29] These products expose minors to AI-mediated financial decisions at an age when their capacity to evaluate risks, understand consent, and exercise meaningful agency is still developing.[30]

The regulatory landscape offers limited protection. Some state privacy laws[31] [32] include heightened data protections for minors, particularly in the context of data collection, sharing, and targeted advertising, and recent New York proposals[33] include provisions addressing juvenile

[27] Consumer Financial Protection Bureau, Banking in Video Games and Virtual Worlds: Issue Spotlight, 2024.
[28] Federal Reserve, Report on the Economic Well-Being of U.S. Households 2023, 2024.
[29] Consumer Financial Protection Bureau, Defining Larger Participants of a Market for General-Use Digital Consumer Payment Applications, November 2024.
[30] OECD, Children in the Digital Environment: Revised Typology of Risks, 2023.
[31] California Privacy Rights Act (CPRA), Cal. Civ. Code § 1798.120 (opt-out rights and restrictions on sale/sharing of minors' data).
[32] Colorado Privacy Act, Colo. Rev. Stat. § 6-1-1308. (heightened protections for minors in profiling and targeted advertising).
[33] Protecting Our Kids: Governor Hochul Announces Nation-Leading Proposals to Protect Kids Online, Restrict AI Chatbots and Combat the Youth Mental Health Crisis, 2026.

interactions with AI. China's algorithmic recommendation rules impose specific duties regarding minors, including prohibitions on content that induces addiction or overconsumption.[34] But no U.S. jurisdiction has enacted comprehensive protections for minors interacting with AI-mediated financial products.[35]

Policy opportunity: Policymakers should establish AI-specific protections for minors in financial services, including: (a) enhanced data minimization and purpose-limitation requirements for financial products used by consumers under 18; (b) age-appropriate consent mechanisms that genuinely inform and empower both parents and minors, rather than relying on click-through agreements; (c) prohibitions on deceptive design patterns targeting minors in financial product interfaces; (d) appeal and human-review rights specifically designed for minors and their guardians when AI-mediated financial decisions affect them; and (e) restrictions on the use of behavioral data from minors for AI training, profiling, or targeted financial product marketing.

Category 4: Accountability and Remedy

Redress

When AI systems produce harmful outcomes—whether a wrongful denial, an erroneous charge, or a biased recommendation—consumers must have clear, accessible pathways to identify the harm, understand its cause, and obtain remedy. Currently, AI's redress problem is structural: the opacity of AI decision-making makes it difficult for consumers to recognize harm, system complexity makes it difficult to identify the cause, and the absence of clear liability frameworks makes it difficult to obtain a remedy.[36]

Policy opportunity: Policymakers should establish that deployers of AI in financial services bear responsibility for harm caused by their systems, regardless of whether the AI component was developed by a third party. This mirrors the product liability principle that the entity that places a product in the stream of commerce is responsible for its safety. Additionally, policymakers should consider: (a) private rights of action for consumers harmed by AI decisions; (b) requirements that companies maintain clear internal escalation pathways for AI-related complaints; and (c) regulatory guidance clarifying that existing consumer complaint and dispute resolution mechanisms under CFPA and other statutes apply fully to AI-mediated decisions.

Monitoring and Auditability

[34] Cyberspace Administration of China, Provisions on the Administration of Algorithmic Recommendation for Internet Information Services, 2022.
[35] Brookings Institution, Why AI Policy Thrives in Some States and Fades in Others, 2026.
[36] UC Berkeley CLTC, "AI's Redress Problem," 2022.

AI systems must be designed and maintained to permit ongoing monitoring and independent audit. Consumers strongly support this: 51% say independent audits would increase their trust in financial companies' AI, and 46% want public reporting on accuracy and bias testing.

Policy opportunity: Colorado's AI act[37] requires impact assessments for high-risk AI systems, and the EU AI Act mandates continuous monitoring, logging, and post-deployment testing. U.S. policy should move in this direction. Policymakers should require: (a) pre-deployment and periodic impact assessments for high-risk AI in financial services; (b) logging and record-keeping requirements that enable retrospective audit; (c) independent third-party audit rights, modeled on the financial audit requirements that already apply to publicly traded companies; and (d) public reporting requirements that give consumers and regulators visibility into AI system performance.

Whistleblower Protections for AI-Related Consumer Harms

Effective accountability depends not only on external monitoring, but also on internal transparency. Employees, contractors, and technical staff within financial institutions and AI vendors are often the first to identify systems that produce discriminatory outcomes, generate unreliable outputs, or operate in ways that conflict with consumer-protection obligations. Their ability to share critical information—with regulators, advocacy organizations, or the public—without retaliation is essential to the integrity of the consumer-protection ecosystem.

Existing whistleblower frameworks offer partial protections. The Consumer Financial Protection Act, the Anti-Money Laundering Act, and the Bank Secrecy Act each provide whistleblower protections in defined contexts. The proposed federal AI Whistleblower Protection Act would extend protections specifically to individuals reporting AI-related safety and rights concerns. At the state level, proposed legislation in Illinois and New York would create AI-specific whistleblower pathways.

Policy opportunity: Policymakers should enact AI-specific whistleblower protections that cover a broad array of AI-related consumer harms in financial services, including discriminatory outcomes, reliability failures, data misuse, and undisclosed conflicts of interest. Consumer advocacy organizations should compile and publicize existing whistleblowing channels, assist regulators in increasing the accessibility and visibility of their tip lines, and support the passage of federal and state legislation that creates protected pathways for individuals reporting AI-related risks to consumer welfare.

Category 5: Systemic Responsibility

[37] CO SB 24-205, tied to Colorado Consumer Protection Act §6-1-112.

Governance

Companies deploying AI in financial services must maintain governance structures that ensure AI systems are developed, deployed, and monitored with adequate internal oversight. This includes clear lines of accountability, board-level responsibility for AI risk, and internal compliance frameworks that treat AI governance with the same seriousness as financial risk management.

Policy opportunity: Prudential regulators—the OCC, FDIC, and Federal Reserve—should issue updated guidance clarifying that model risk-management frameworks (SR 11-7 and related guidance) apply to AI and machine learning systems, including generative AI. State regulators should follow the lead of the New York Department of Financial Services, which has issued guidance requiring insurers to assess AI risks in use of algorithms and external data in underwriting and pricing—particularly risks of unfair discrimination—demonstrating how regulators can operationalize AI risk management.[38] Policymakers should also consider whether existing governance requirements need to be updated to reflect the specific risks of AI, including the risks created by dependence on third-party AI providers.

Environmental Cost Transparency

AI systems impose significant environmental costs through the energy required to train models and operate the data center infrastructure that supports them.[39] These costs are currently largely invisible to consumers. As AI accounts for a growing share of financial services, the environmental footprint of consumers' financial relationships increases—without their knowledge or consent. In our survey, 15% of Americans identified environmental impact as a risk of AI in financial services[40]—a figure likely to grow as public awareness of AI's energy demands increases.

Policy opportunity: At the federal level, bipartisan legislation such as the proposed Artificial Intelligence Environmental Impacts Act would require environmental impact assessments and standardized voluntary reporting by AI companies, organizations, or other entities. The EU AI Act includes environmental-transparency provisions. Policymakers should extend these principles to financial services specifically, requiring companies to disclose the environmental footprint of their AI systems and enabling consumers to compare the environmental impact of competing financial products.

---

[38] New York Department of Financial Services, Insurance Circular Letter No. 1, 2024, Use of External Consumer Data and Information Sources in Underwriting and Pricing, July 2024.
[39] MIT Technology Review, "We did the math on AI's energy footprint," May 2025.
[40] Consumer Reports nationally representative AI in Financial Services Survey of 4,073 U.S. adults, September-October 2025, https://article.images.consumerreports.org/image/upload/v1767035543/prod/content/dam/surveys/Consumer_Reports_AI_in_Financial_Services_2025.pdf.

Supply Chain Accountability

The AI supply chain in financial services is deep and opaque. A consumer's bank[41] may use an AI system built by a technology company, trained on data assembled by a third-party vendor, running on infrastructure operated by a cloud hyperscaler.[42] When something goes wrong, no single entity may accept responsibility. The concentration of the AI supply chain among a small number of providers also creates systemic resilience risks. When these providers suffer outages, AI-dependent financial services may fail silently, degrading in quality without consumers' knowledge.

Compounding this opacity is the absence of standardized disclosure requirements between AI developers and the financial institutions that deploy their systems. A bank integrating a third-party AI model into its credit underwriting may not receive comprehensive documentation of the model's training-data composition, known failure modes, performance limitations across demographic groups, or behavioral characteristics such as confidence calibration and sycophancy tendencies. Without this information, deployers cannot fulfill their own obligations to consumers—they cannot explain what they do not understand, and they cannot test for risks they have not been told to look for. The EU AI Act addresses this directly, requiring that providers of high-risk AI systems furnish deployers with substantial technical documentation and instruction for use. No comparable U.S. requirement exists.

Policy opportunity: Our landscape analysis found that resilience requirements for AI in consumer financial services are virtually nonexistent. The CFPB has raised concerns about concentration in core service providers,[43] and the FTC has warned about "single points of failure" in cloud computing,[44] but neither has taken regulatory action specific to AI resilience. Policymakers should: (a) require financial institutions to maintain and disclose contingency plans for AI system failures; (b) establish resilience standards requiring graceful degradation and transparent failure; (c) consider whether the concentration of AI infrastructure among a small number of cloud providers poses systemic risk requiring specific oversight; (d) require developer-to-deployer disclosure for AI systems used in consumer financial services, ensuring that financial institutions receive documentation of training data characteristics, known limitations, performance benchmarks across demographic groups, and behavioral risk profiles before deployment; and (e) ensure that deployers verify, rather than merely accept, vendor representations about AI system performance—paralleling the due diligence standards that apply to other forms of third-party vendor management in financial services.

---

[41] OCC Bulletin 2023-17 (Third-Party Risk Management), re outsourcing of models; Federal Reserve SR 13-19 / SR 11-7, re vendor dependencies; FDIC FIL-44-2008 (and updates), re-shared accountability.
[42] Georgetown University, "Ask a Professor: CrowdStrike Outage," July 2024; NIST AI Risk Management Framework, 2023, AI systems are modular and multi-actor by design.
[43] CFPB Director Chopra, opening remarks to Community Bank and Credit Union Advisory Councils, 2022.
[44] FTC, "Cloud Computing RFI: What We Heard & Learned," November 2023.

Applying the Standard: Financial Services Case Studies

This section demonstrates how the Consumer Welfare Standard applies to three critical areas in which AI is already reshaping consumer financial services.

Dynamic and Algorithmic Pricing

AI-driven pricing is one of the most consequential—and least transparent—applications of machine learning in consumer markets. Investigations by Consumer Reports and others have documented how algorithmic pricing systems at companies, including Instacart and Kroger, can produce inconsistent, opaque, and unfair pricing outcomes.[45] These investigations mirror growing regulatory concern, including from the FTC, and influenced state-level legislation, including California's AB 446,[46] about opaque and potentially unfair pricing practices.[47]

Algorithmic pricing implicates nearly every dimension of the Consumer Welfare Standard. It raises Information Integrity concerns when consumers cannot determine how their price was set or whether it differs from what others pay. It raises Fair Treatment concerns when price differences correlate with protected characteristics or when personalized pricing exploits behavioral vulnerabilities. It undermines Consumer Control when consumers have no ability to opt out of personalized pricing or to comparison-shop effectively. And it creates Accountability gaps when neither the consumer nor the regulator can reconstruct the logic that produced a specific price.

Policy pathway: Policymakers should require algorithmic pricing transparency in financial products, including disclosure of whether and how AI is used to set prices, rates, or fees; prohibitions on pricing practices that produce discriminatory outcomes; and independent audit rights that allow regulators and designated third parties to examine pricing algorithms. The absence of a broad U.S. prohibition on algorithmic price discrimination represents a significant gap that consumer advocates should work to close at both the federal and state levels.

Generative AI in Advice and Decision Support

Generative AI systems are increasingly deployed as customer-facing financial advisors, insurance navigators, and product-recommendation engines. These systems present a

[45] "Kroger Stores Overcharging Shoppers," June 2025; "Inside Kroger's Secret Shopper Profiles," May 21, 2025; "Instacart's AI-Enabled Pricing Experiments," December 9, 2025; "Instacart Stops Pricing Tests," January 8, 2026; New York Times, "Kroger and Walmart Deny Surge Pricing After Adopting Digital Price Tags." October 23, 2024.
[46] Consumer Reports, Support for AB 446: Surveillance Pricing. May 12, 2025.
[47] FTC, "Commercial Surveillance and Data Security Rulemaking," (initiated in 2022, ongoing); Sen. Gallego, One Fair Price Act. December 2025.

distinctive consumer welfare challenge because they interact with consumers in natural language, creating an impression of expertise and personalization that may not be warranted.

The risks include: overconfidence, where AI systems present uncertain conclusions with unwarranted authority; sycophancy, where systems agree with consumers rather than providing honest guidance; hallucination, where systems fabricate product features, terms, or company policies; and explainability gaps, where consumers cannot determine how the system arrived at its recommendation or whether it was optimizing for their interest or the company's.

Policy pathway: Policymakers should consider whether AI systems that provide financial advice or recommendations should be subject to fiduciary or suitability standards analogous to those that apply to human financial advisors. At minimum, any AI system that provides personalized financial guidance should be required to: disclose that it is an AI system; disclose any conflicts of interest; provide a clear basis for its recommendations; and offer a pathway to human review. The absence of regulatory activity specifically addressing AI sycophancy in consumer financial services represents a significant gap.

Credit Decisioning and Risk Modeling

AI-driven credit decisioning is the area in which existing consumer protection law is most directly applicable—and where the limitations of that law in the AI context are most apparent. The Equal Credit Opportunity Act prohibits discrimination in lending. The Fair Credit Reporting Act governs the accuracy of consumer reports and the consumer's right to dispute inaccurate information. Regulation B requires adverse-action notices when credit is denied.

But these laws were designed for a world in which credit decisions were made by identifiable models using a limited number of variables. Machine learning systems that use hundreds or thousands of features, including non-traditional data sources, strain the capacity of existing law to ensure that consumers receive meaningful explanations, that discriminatory outcomes are detected and remedied, and that the decision-making process is auditable.

Policy pathway: Policymakers should: (a) update adverse action-notice requirements to mandate explanations that are meaningful in the context of machine learning—not merely technically compliant; (b) require pre-deployment and ongoing disparate-impact testing for AI credit models, with results reported to regulators; (c) clarify that the use of AI does not relieve lenders of their obligation to provide individualized reasons for adverse actions; and (d) ensure that regulators have the technical capacity to examine and evaluate AI credit models, including through investments in regulatory technology and expertise.

The Regulatory Landscape: Applicable, But Not Integrated

Consumer Reports' landscape analysis examined the full range of U.S. federal and state regulatory activity relevant to AI in consumer financial services, as well as significant international regulatory developments. The analysis identified significant strengths in the existing framework—and critical gaps that this section maps against the Consumer Welfare Standard.

Technology-Neutral Statutes Continue to Apply

Existing federal consumer protection statutes—including the Equal Credit Opportunity Act, Fair Credit Reporting Act, Consumer Financial Protection Act (CFPA), the Gramm-Leach-Billey Act (GLBA), the Truth In Lending Act (TILA), the Fair Housing Act (FHA), and the Federal Trade Commission Act's prohibition on unfair and deceptive practices—apply to AI-mediated financial decisions. This is a critical baseline: The use of AI does not create a legal exemption from existing obligations.[48] Companies cannot argue that because a decision was made by an algorithm they are relieved of responsibility for discriminatory or deceptive outcomes.[49]

But technology-neutral application is necessary and insufficient. These statutes were designed for a world of identifiable decision-makers, transparent criteria, and static products. Applying them effectively to AI systems requires updated interpretive guidance, new enforcement tools, and investments in regulatory technical capacity that have not yet been made at scale.

State-Level Innovation and Fragmentation

In the absence of comprehensive federal AI legislation, states have emerged as the primary locus of regulatory innovation:

- **Colorado's AI Act** requires impact assessments, algorithmic non-discrimination obligations, and post-decision explanations, and makes violations enforceable under the state's consumer protection enforcement framework.[50]

---

[48] GAO, "Artificial Intelligence: Use and Oversight in Financial Services" (GAO-25-107197), May 19, 2025, 5-6, providing an overview of the legal framework that financial institutions are still subject to existing fair lending, consumer protection, and regulatory frameworks.

[49] https://files.consumerfinance.gov/f/documents/cfpb_joint-statement-enforcement-against-discrimination-bias-automated-systems_2023-04.pdf.

[50] Colorado, Concerning Consumer Protections in Interactions With Artificial Intelligence Systems (SB 24-205, §§ 6-1-1701 et seq, signed May 17, 2024); Colorado Consumer Protection Act (penalties), Colo. Rev. Stat. § 6-1-112.

- **California's proposed Automated Decisions Safety Act** would establish pre-deployment assessments and governance requirements for automated-decision systems.[51]
- **New York's DFS Guidance** requires insurers to assess risks arising from AI and algorithmic decision-making,[52] while the proposed NY AI Consumer Protection Act relies on public enforcement rather than creating a private right of action for consumers harmed by AI.[53]
- **Multiple states** (Colorado, Connecticut, Delaware, Montana, New Jersey, Oregon, Texas) provide opt-out rights for profiling in furtherance of decisions that produce legal or similarly significant effects.[54]

These state-level developments are important and should be supported. But the resulting patchwork creates compliance complexity for companies, coverage gaps for consumers, and the risk of a race to the bottom as states compete for industry presence. Moreover, proposals for federal preemption of state AI regulations pose a direct threat to the strongest existing consumer protections.

Policy imperative: Federal AI legislation should establish a floor of consumer protection, not a ceiling. Any federal framework must preserve the ability of states to enact stronger protections. Consumer advocates should vigorously oppose federal preemption provisions that would weaken or displace state laws such as Colorado's AI Act.

Critical Gaps Across the Consumer Welfare Standard

Our landscape analysis identified specific gaps where existing regulation fails to address dimensions of the Consumer Welfare Standard:

Explainability Gaps

While multiple laws require some form of disclosure or explanation—including ECOA's adverse-action notices and Colorado's post-decision explanations—none require the depth of explanation that consumers demand. In our survey, 63% of consumers said they would want a

[51] California AB 2930 (Automated Decision Systems Act), § requirements on impact assessments and risk evaluation.
[52] New York Department of Financial Services, Circular Letter No. 2024-07, Use of Artificial Intelligence Systems and External Consumer Data and Information Sources in Insurance Underwriting and Pricing, July 11, 2024.
[53] Cal. AB. 2930.
[54] Colorado Privacy Act, Colo. Rev. Stat. § 6-1-1306(1)(a)(I)(C); Connecticut Data Privacy Act, Conn. Gen. Stat. § 42-517(a)(1)(C); Delaware Personal Data Privacy Act, Del. Code tit. 6, § 12D-104(a)(1)(C); Montana Consumer Data Privacy Act, Mont. Code Ann. § 30-14-2703(1)(c); New Jersey Data Privacy Act, N.J. Stat. § 56:8-166(a)(3); Oregon Consumer Privacy Act, Or. Rev. Stat. § 646A.574(1)(c); Texas Data Privacy and Security Act, Tex. Bus. & Com. Code § 541.051(c)(3).

complete list of all factors and their relative weight if denied a loan by AI. Current legal requirements fall far short of this standard.

Deceptive Design Gaps

While several states prohibit dark patterns in consent and disclosure contexts, no U.S. law specifically addresses the use of AI to personalize manipulative interface designs in financial services. Critically, even where explainability requirements exist, there is no standard ensuring that explanations are accessible rather than merely available. Companies can satisfy the letter of disclosure requirements while burying explanations in fine print, presenting them in dense technical language or placing them in interface locations that consumers are unlikely to reach. Government agencies and consumer-advocacy organizations should comparatively evaluate how financial institutions present AI-related explanations and whether deceptive design patterns are used to undermine their effectiveness.

Sycophancy and Confidence Calibration Gaps

While existing doctrines such as UDAP/UDAAP, suitability obligations, and fiduciary duties could, in principle, address sycophancy-related harms, as outlined in our landscape analysis, they have not been tested or applied in this context. No current U.S. law, regulation, or guidance specifically addresses AI sycophancy in consumer financial services. Although existing laws may be leveraged to address some simpler issues, as AI in financial services evolves to assume more fiduciary and decision-making responsibilities this remains a complete regulatory void in a risk area that researchers, companies, and advocates have all identified as urgent.

Resilience Gaps

No U.S. law or regulation establishes resilience requirements for AI systems in consumer financial services. The CFPB and FTC have raised concerns about infrastructure concentration and single points of failure, but neither has taken regulatory action. The EU AI Act's requirements for accuracy thresholds, error resilience, and logging provide a model that U.S. policymakers have not yet adopted.[55]

Pricing Transparency Gaps

The U.S. has no broad prohibition on algorithmic price discrimination in financial services. While credit-specific discrimination is prohibited under ECOA, there are no general rules requiring transparency or fairness in AI-driven pricing of fees, insurance premiums, or service terms.

[55] EU Artificial Intelligence Act (2024), entered into force August 2024, phased implementation through 2027.

Enforcement Capacity Gaps

Even where legal frameworks exist, enforcement capacity is a binding constraint. Current financial penalty structures may not adequately deter AI-driven misconduct by large technology firms. Civil penalties available under Colorado's consumer protection law, while meaningful for smaller firms, may represent a cost of doing business for companies with billions in revenue.[56] The CFPB's enforcement authority is significant, but its future scope remains uncertain. And many federal agencies lack the technical expertise to evaluate AI systems effectively.

The International Comparison

The EU AI Act represents one of the most comprehensive AI regulatory frameworks enacted globally.[57] It establishes a tiered risk framework; mandates pre-deployment and ongoing risk management for high-risk systems (explicitly including AI in consumer financial services); requires data quality and representativeness standards; and imposes transparency, explainability, and human-oversight obligations that exceed those in current U.S. law.

China's algorithmic recommendation regulations take a different but instructive approach, establishing explicit prohibitions on algorithmic price discrimination, mandatory opt-out mechanisms with non-personalized alternatives, and periodic audit requirements.[58]

Policy implication: The U.S. is falling behind peer jurisdictions in establishing a coherent framework for AI consumer protection. Both the EU and China have enacted specific, binding rules that address dimensions of consumer welfare not yet reached by U.S. regulation. While neither model is directly transferable, both demonstrate that comprehensive AI consumer protection regulation is achievable and that the U.S.'s current fragmented approach is not inevitable.

---

[56] Colo. Rev. Stat. § 6-1-112(1)(a) (Colorado Consumer Protection Act — civil penalties).
[57] EU Artificial Intelligence Act (2024), entered into force August 2024, phased implementation through 2027; OECD AI Policy Observatory (2024/2025 updates).
[58] Cyberspace Administration of China, Internet Information Service Algorithmic Recommendation Management Provisions (2022), arts. 17–18, 24; translation.

From Framework to Implementation

A consumer welfare standard is only as valuable as the ecosystem that implements, monitors, and enforces it. This section outlines three implementation pathways: independent evaluation infrastructure, policy integration, and corporate engagement.

Independent Evaluation Infrastructure

The AI evaluation landscape lacks what the consumer financial services sector has long had: an independent, trusted third-party that evaluates products from the consumer's perspective. Consumer Reports is building this capacity through its AI evaluation program, but the field as a whole needs a more robust evaluation ecosystem.

Effective independent evaluation requires: standardized testing methodologies that enable comparison across products and institutions; transparent scoring frameworks that translate complex technical assessments into information consumers can use; investigative capacity to identify and document harms that companies do not voluntarily disclose; and sustained funding from sources that do not compromise evaluators' independence.

Policy pathway: Policymakers should support the development of independent AI evaluation infrastructure through: (a) funding for independent testing and evaluation organizations; (b) legal protections (safe harbors) for good-faith AI security and fairness research; (c) requirements that companies provide independent evaluators with reasonable access to test AI systems; and (d) recognition of independent evaluation findings in regulatory proceedings and enforcement actions.

Policy Integration

The Consumer Welfare Standard proposed in this paper is designed to be regulatory-ready. Each obligation dimension can be translated into specific regulatory requirements, enforcement standards, or supervisory expectations. Key integration pathways include:

- **Federal regulatory guidance:** Prudential regulators, the CFPB, and the FTC should issue guidance applying the dimensions of this standard to their respective supervisory and enforcement frameworks.
- **State legislative models:** The standard provides a template for comprehensive state AI consumer protection legislation that goes beyond the current patchwork.
- **Industry standards:** Self-regulatory organizations and industry groups should adopt the standard as a baseline for voluntary commitments, with the understanding that voluntary adoption is not a substitute for regulatory requirements.

- **Regulatory sandboxes:** Regulators should consider sandbox environments that allow testing of AI financial products against the Consumer Welfare Standard before full-scale deployment.

Corporate Engagement

Consumer Reports' experience with the Fair Digital Finance Framework and the Fairness by Design Playbook demonstrates that some financial companies—particularly those seeking competitive differentiation through consumer trust—will voluntarily adopt higher standards when the standards are clearly defined, practically implementable, and independently verified. Companies such as Chime, Cash App, and Varo demonstrate elements of CR's Fair digital Finance Framework, showing that fairness can function as a competitive advantage.

The corporate-engagement pathway for AI-mediated consumer finance products could include: pre-market testing against the Consumer Welfare Standard that we will publish soon; transparency commitments on AI disclosure, explainability, and bias testing; and certification models that allow companies to demonstrate compliance with consumer-welfare standards.

Critical caveat: Voluntary corporate adoption is valuable, but cannot substitute for regulatory requirements. The consumer-welfare challenges posed by AI in financial services are structural, not situational. They will not be resolved by the goodwill of individual companies. They require binding obligations, enforced by regulators with the authority and capacity to hold firms accountable.

Consumer Mobilization

Throughout our survey, consumers demonstrated both the awareness and the preferences needed to support stronger AI protections. Seventy-seven percent of Americans want the right to opt out of AI-based decisions, and 57% say current laws are inadequate. These are not niche positions; they are mainstream. Consumer advocacy organizations, working in coalition with civil rights groups, state attorneys general, and legislative champions, have an opportunity to translate these preferences into political action. The survey data presented in this paper provides the empirical foundation for that effort.

Conclusion: Building Consumer-Centered AI Markets

The findings presented in this paper converge on a central conclusion: AI is transforming consumer financial services in ways that existing consumer-protection frameworks were not designed to address, at a pace that regulatory systems have not matched, and in directions that consumer preferences do not endorse.

The market incentives driving AI adoption in financial services are overwhelmingly commercial. Consumer welfare—understood as fair treatment, informed autonomy, meaningful choice, and accessible redress—is not what these systems are optimized to deliver. Our nationally representative survey of 4,073 Americans confirms that the public understands this. They see AI's potential benefits. They also see the risks, and they do not trust that the current system is managing those risks on their behalf.

The Consumer Welfare Standard proposed in this paper offers a practical path forward. It does not call for the prohibition of AI in financial services. It calls for AI to be held to the same obligation that any entity providing consumer financial services is held to: the obligation to serve the consumer's interest honestly, fairly, and accountably.

Making this standard real requires action on four fronts simultaneously:

- **Policymakers** must update existing legal frameworks and enact new ones that address the specific risks of AI in financial services—including explainability gaps, pricing opacity, sycophancy, resilience risks, and the challenges of agentic AI. Federal legislation should establish a floor of protection, not a ceiling, and should preserve state innovation.
- **Regulators** must invest in the technical capacity to supervise AI systems, issue guidance applying existing legal standards to AI contexts, and use their enforcement authority to establish that AI-driven harm carries real consequences.
- **Industry** must move beyond aspiration to implementation, adopting the kind of pre-deployment testing, ongoing monitoring, and transparent reporting that the Consumer Welfare Standard requires—not as a compliance burden but as a competitive advantage.
- **Independent evaluation organizations** must build the testing infrastructure, comparative analysis capacity, and public reporting systems that make AI's impact on consumers visible, measurable, and actionable.

Consumer Reports is committed to ensuring the outcomes of AI-mediated finance favor consumers. This paper is an invitation to every institution and leader who shares that commitment to join in building the evaluation infrastructure, the regulatory frameworks, and the market conditions that make consumer-centered AI not just possible, but expected.

About Consumer Reports

Consumer Reports is an independent, nonprofit organization that works side by side with consumers to create a fairer, safer, and healthier world. CR's Digital Marketplace initiative focuses on ensuring that digital products and services—including AI-mediated financial services—are safe, fair, and transparent for all consumers. This work is supported by the Alfred P. Sloan Foundation as part of the Consumer Welfare in an AI-Mediated World project.

Consumer Reports nationally representative AI in Financial Services Survey of 4,073 U.S. adults (September-October 2025). Toplines available at https://article.images.consumerreports.org/image/upload/v1767035543/prod/content/dam/surveys/Consumer_Reports_AI_in_Financial_Services_2025.pdf.