

THE LEGAL RISKS POSED BY AI CHATBOTS

Unfair and Deceptive Practices
and Other Authorities

Stephanie T. Nguyen is a Senior Fellow at Columbia Law School's Center for Law and the Economy and author of the Division of Field Notes. She is the former Chief Technologist at the FTC.

Adam Teitelbaum is a Senior Fellow at Protect Borrowers and is the former Director of the Office of Consumer Protection at the Office of the Attorney General for the District of Columbia.

I.

Introduction

Growing alarm about existential risks and other harms posed by AI technology have given rise to calls for the government to regulate AI platforms.

This issue brief addresses a related but distinct problem: the use of AI technology by large corporations—including AI platforms themselves—to deploy AI customer-service chatbots¹ in ways that may violate existing consumer protection laws. Law enforcers and the plaintiffs' bar can and should use existing law to stop and deter those violations, and all businesses using AI chatbots should take immediate steps to ensure their AI chatbot practices are compliant with consumer protection laws.

Companies are increasingly deploying AI chatbots as the mandatory gateway between consumers and the help they need. That includes tenants, borrowers, insurance policyholders, and fraud victims seeking assistance with services that can be enormously consequential in their lives. By replacing workers with this technology, companies are handing off critical channels for assistance to untested, unreliable AI systems, effectively turning consumers into guinea pigs as those systems are refined through real-world trial and error. The resulting harms take many forms: inaccurate answers, obstructive doom loops or other friction, or a failure to deliver the assistance the service is supposed to offer.

The issue brief offers enforcers, practitioners, and companies seeking to do right by their customers a framework for assessing AI chatbots for compliance with current law, and recommends further steps that enforcers, law-abiding businesses, and policymakers can take to address these issues. In particular, state and local consumer protection enforcers, as well as bank regulators, should use their existing investigatory and supervisory powers to examine how these systems are designed, deployed, and overseen—and act on any unlawful practices they find.

¹ For simplicity, we use "chatbots" throughout this piece to refer to these interfaces, which can include, for instance, automated voice assistants, AI-guided forms or workflows, and chatbots on websites, apps, or search engines.

II.

Wrong, Blocked & Broken: Potentially Unlawful AI Customer Service Chatbot Practices

Existing consumer protection laws and financial regulations already govern customer service interactions between companies and their customers, and these laws do not exempt AI. Indeed, it is a bedrock principle that companies are generally liable for the conduct of their agents,² and that includes AI chatbots.³ This issue brief details three potentially illegal pathologies associated with businesses' use of AI customer service interfaces, including chatbots: false information, obstructing access, and inadequate service.⁴

A. FALSE INFORMATION

AI chatbots routinely provide customers false or misleading information. Although the federal government has significantly retrenched from enforcing federal bans on unfair and deceptive practices (UDAP)—including Section 5 of the Federal Trade Commission (FTC) Act, 15 U.S.C. § 45; the Consumer Financial Protection Act (CFPA), 12 U.S.C. §§ 5531, 5536, 5552(a); and the Federal Aviation Act,

49 U.S.C. § 41712—all states and the District of Columbia have their own parallel UDAP prohibitions.⁵ These laws squarely reach false or misleading statements about products and services across commerce, including those made by AI chatbots, and they are generally enforceable by both state attorneys general and private plaintiffs. Where the false information concerns the status of a debt, the reasons for a credit denial, a request to modify loan terms, or the outcome of a servicing dispute, it may separately violate consumer protection statutes specific to financial markets—including the Fair Debt Collection Practices Act (FDCPA), 15 U.S.C. § 1692e(2)(A); the Equal Credit Opportunity Act (ECOA), 15 U.S.C. § 1691(d)(2)–(3); and/or the Real Estate Settlement Procedures Act (RESPA), 12 U.S.C. § 2605(e), 12 C.F.R. § 1024.35(e)—all of which are enforceable by state attorneys general,⁶ as well as by private plaintiffs.⁷

These authorities clearly apply to false or deceptive customer service information. For example, the Consumer Financial Protection Bureau (CFPB) sued the National Collegiate Student Loan Trusts and their servicer, the Pennsylvania Higher Education Assistance Agency, in 2024, alleging that years of inaccurate customer communications and wrongly denied forbearance requests violated the CFPA's UDAP prohibitions. The companies paid more than \$5 million to settle the matter.⁸ Similarly, a state-law UDAP claim arising from an AI chatbot stating that it was real recently survived a motion to dismiss.⁹

² See, e.g., Rutter Group, Bus. & Professions Code Section 17200 PRACTICE Ch. 6-B (Thomson Reuters ed., 2026) ("The usual rules of common law agency apply [under California's unfair and deceptive practices (UDAP) law]—i.e., a principal is liable for the acts of its agents if such acts are done within the scope of the agent's actual or apparent authority."); *Celtic Life Ins. Co. v. Coats*, 885 S.W.2d 96, 98–99 (Tex. 1994) (holding an insurance company liable for an agent's misrepresentation of terms and benefits of policy pursuant to Texas UDAP law).

³ See, e.g., *Garcia v. Character Techs., Inc.*, 785 F. Supp. 3d 1157, 1173 (M.D. Fla. 2025), *motion to certify appeal denied*, No. 6:24-CV-1903-ACC-DCI, 2025 WL 2581834 (M.D. Fla. July 15, 2025) (declining to dismiss product liability, negligence, and Florida UDAP claims premised on an AI chatbot's actions).

⁴ These categories are not mutually exclusive, and each can lead to catastrophic harms and consequences—such as credit score damage, loss of money, foreclosure, eviction, repossession, spiraling fees with interest, fraud, and denied access to services consumers urgently need. Each of these harms can independently violate consumer protection laws.

⁵ *Consumer Protection Laws: 50-State Survey*, Justia, <https://www.justia.com/consumer/consumer-protection-laws-50-state-survey/> (last visited Sep. 23, 2026).

⁶ 12 U.S.C. § 5552(a)(1).

⁷ 15 U.S.C. § 1692k; 15 U.S.C. § 1691e; 12 U.S.C. § 2605(f).

⁸ Press Release, Consumer Fin. Prot. Bureau, *CFPB Takes Action to Require National Collegiate Student Loan Trusts and Pennsylvania Higher Education Assistance Agency to Pay More than \$5 Million for Student Loan Servicing Failures* (May 6, 2024), <https://wayback.archive-it.org/23481/20251004103921/https://www.consumerfinance.gov/about-us/newsroom/cfpb-takes-action-to-require-national-collegiate-student-loan-trusts-and-pennsylvania-higher-education-assistance-agency-to-pay-more-than-5-million-for-student-loan-servicing-failures/>.

⁹ See *Garcia*, 785 F. Supp. 3d at 1173.

Cases from other jurisdictions are also instructive. When Air Canada’s chatbot falsely told a grieving customer he could claim a bereavement discount after travel, and the airline then denied his refund, a Canadian tribunal rejected the company’s argument that its chatbot was a separate legal entity not responsible for its own representations—finding Air Canada liable for negligent misrepresentation.¹⁰

There is no AI exception to these principles.¹¹ As the CFPB has warned: “Chatbots sometimes get the answer wrong ... [and] being wrong may violate [companies’] legal obligations.”¹²

B. OBSTRUCTING ACCESS

By their very design, customer service chatbots can obstruct consumers from canceling services or reaching customer service—i.e., trapping them in “doom loops”—preventing customers from getting the help they need. This conduct can violate UDAP unfairness prohibitions, as well as deception prohibitions where a company advertises robust service without delivering it. It can also violate consumer financial protections: causing a mortgage servicer to blow through its response deadline in violation of RESPA, 12 U.S.C. § 2605(e); preventing timely investigation of a disputed transaction in violation of the Electronic Fund Transfer Act, 15 U.S.C. § 1693f; or letting a debt collector keep collecting on a debt although the dispute was never processed, in violation of the FDCPA, 15 U.S.C. § 1692g(b).

Enforcers have repeatedly applied these authorities to address customer service obstruction—i.e., friction built into the design of customer service channels that makes it hard or impossible for customers to get relief. For example, the FTC successfully secured a settlement against H&R Block after the company made it easy to upgrade to a higher-priced product—just a click—while making it difficult to downgrade to a cheaper one. To switch to a lower-priced product, customers had to contact customer service by phone or chat, and then H&R Block would erase all the tax information they had already entered.¹³ Under the settlement, the company agreed to let consumers downgrade through the same automated, self-service channels used for upgrades, and to preserve their entered data when they do.¹⁴ State attorneys general have brought similar UDAP claims. In October 2025, a coalition of 34 state attorneys general secured a settlement with TFG Holding, Inc., an online subscription-based fashion retailer after finding that the company misrepresented pricing, enrolled customers without adequate disclosures, and maintained cancellation practices that “hindered consumers’ ability to cancel” its VIP membership program, under which customers were automatically billed each month unless they opted out.¹⁵

Companies can also face liability when the design of their systems renders advertised customer protections illusory in practice. For example, the FTC alleged that sneaker marketplace GOAT violated Section 5 by advertising “Buyer Protection Guaranteed” coverage for defective or mismatched items without providing any phone or chat support—instead routing many such claims through automated channels that denied relief or offered only partial credit. GOAT’s settlement requires it to flag these claims so that they are not automatically denied and to resolve them within a commercially reasonable time.¹⁶

¹⁰ Kyle Melnick, *Air Canada Chatbot Promised a Discount. Now the Airline Has to Pay It.*, Wash. Post (Feb. 18, 2024), <https://www.washingtonpost.com/travel/2024/02/18/air-canada-airline-chatbot-ruling/>.

¹¹ See *DoNotPay*, Fed. Trade Comm’n, <https://www.ftc.gov/legal-library/browse/cases-proceedings/donotpay> (last updated Feb. 11, 2025); *Garcia*, 785 F. Supp. 3d at 1173.

¹² Consumer Fin. Prot. Bureau, *Chatbots in Consumer Finance* (June 6, 2023), <https://www.consumerfinance.gov/data-research/research-reports/chatbots-in-consumer-finance/chatbots-in-consumer-finance/>.

¹³ Complaint at 4, *In re H&R Block, Inc.*, No. 9427 (F.T.C. Feb. 23, 2024), https://www.ftc.gov/system/files/ftc_gov/pdf/d09427hrblockadmincomplaintpublic.pdf.

¹⁴ Decision and Order, *In re H&R Block, Inc.*, No. 9427 (F.T.C. Jan. 8, 2025), https://www.ftc.gov/system/files/ftc_gov/pdf/HRBlock-DecisionandOrder.pdf.

¹⁵ Press Release, Md. Office of the Att’y Gen., *Attorney General Brown Secures Settlement with Online Clothing Retailer Regarding Deceptive Advertising and Billing Practices* (Oct. 23, 2025), <https://oag.maryland.gov/News/pages/Attorney-General-Brown-Secures-Settlement-with-Online-Clothing-Retailer-Regarding-Deceptive-Advertising-and-Billing-Practic.aspx>.

¹⁶ Stipulated Order for Permanent Injunction, Monetary Judgment, and Other Relief, *FTC v. 1661, Inc. d/b/a GOAT*, No. 2:24-cv-10329 (C.D. Cal. Dec. 2, 2024), https://www.ftc.gov/system/files/ftc_gov/pdf/GOATProposedOrderreJudgment.pdf.

Companies can further face liability when the design of their customer-facing systems is obstructive or non-responsive. For example, the CFPB sued Navient under the CFPA, 12 U.S.C. §§ 5531, 5536, for compensating call-center representatives based on average call time, which incentivized them to rush borrowers into forbearance rather than income-driven repayment—leading to a \$120 million settlement that banned the incentive structure.¹⁷ The CFPB also sued Comerica Bank for violating Regulation E, 12 C.F.R. Part 1005, by “intentionally terminat[ing] almost 25 million customer service calls,” impeding cardholders from exercising their rights under the law, along with mishandling fraud complaints and charging illegal fees.¹⁸ And the CFPB found that Block, the operator of Cash App, violated Regulation E through deficient investigation and customer service communication processes after Block directed fraud victims to seek reversals of unauthorized transactions through their own banks—only to deny those claims once submitted.¹⁹

Chatbot interface doom loops represent another iteration of this same playbook—they can cause the same harms and violate the same laws.

C. INADEQUATE SERVICE

Even where a chatbot’s design is not itself the problem, chatbots can fail consumers if they are poorly built or inadequately implemented, thus being incapable of doing the job they were in theory designed to do. UDAP laws reach this conduct too: a customer service system incapable of solving the problems it exists to solve is a textbook unfair practice. Where the underlying issue involves a mortgage-servicing error or a disputed electronic transaction, RESPA, 12 C.F.R. § 1024.35(e)(1)(i), and

Regulation E, 12 C.F.R. § 1005.11(c), may apply as well, requiring not just a response but an actual investigation and resolution—obligations a chatbot incapable of resolving anything cannot satisfy.

For example, the CFPB found that the reverse-mortgage servicing operation run by Sutherland Global and its subsidiaries, understaffed to handle as many as 150,000 borrowers, produced “systemic failures to respond to thousands of homeowner requests for assistance,” violating the CFPA’s UDAP prohibitions, along with RESPA. The companies paid \$11.5 million in redress and a \$5 million penalty and were permanently banned from reverse-mortgage servicing.²⁰

Violations may also stem from a company simply failing to invest in adequate customer service staffing or infrastructure. For example, the FTC took action against a telehealth company, NextMed, including for failing to “employ sufficient customer service personnel or maintain adequate technological capacity to address the large volume of customer requests and complaints.”²¹ The Department of Transportation, applying its own UDAP authority, penalized Southwest Airlines a record \$140 million after finding that during the airline’s 2022 holiday meltdown, its call centers became so overwhelmed that stranded customers “were often met with busy signals, hours-long queues to connect with agents, or dropped calls,” compounding the harm from the underlying flight cancellations.²²

These same failures can occur in—and may be exacerbated by—the AI chatbot context, triggering the same liability.

¹⁷ Stipulated Final Judgment and Order, *CFPB v. Navient Corp.*, No. 3:17-cv-00101-RDM (M.D. Pa. Sep. 12, 2024), https://files.consumerfinance.gov/f/documents/cfpb-navient-et-al-stipulated-final-judgment-and-order_2024-09.pdf.

¹⁸ Complaint, *CFPB v. Comerica Bank*, No. 1:24-cv-03401 (N.D. Tex. Dec. 6, 2024), https://wayback.archive-it.org/23481/20250331182046/https://files.consumerfinance.gov/f/documents/cfpb_CFPB-v.-Comerica-Bank_Complaint_2024_12.pdf.

¹⁹ Consent Order, *In re Block, Inc.*, CFPB No. 2025-CFPB-0001 (Jan. 16, 2025), https://files.consumerfinance.gov/f/documents/cfpb_block-inc-consent-order_2025-01.pdf.

²⁰ Consent Order, *In re Sutherland Global Servs., Inc.*, CFPB No. 2024-CFPB-0005 (June 18, 2024), https://files.consumerfinance.gov/f/documents/cfpb_sutherland-global-services-inc-et-al-consent-order_2024-06.pdf.

²¹ Complaint at 17, *In re S. Health Solutions, Inc.*, No. 232-3040 (F.T.C. July 14, 2025), https://www.ftc.gov/system/files/ftc_gov/pdf/ComplaintNextMed.pdf.

²² Press Release, U.S. Dep’t of Transp., *DOT Penalizes Southwest Airlines \$140 Million for 2022 Holiday Meltdown* (Dec. 18, 2023), <https://www.transportation.gov/briefing-room/dot-penalizes-southwest-airlines-140-million-2022-holiday-meltdown>.

III.

What Enforcers, the Plaintiffs' Bar, and Policymakers Can Do To Stop AI Customer Service Chatbot Abuses

Hold Companies Accountable for Violating UDAP and Other Existing Authorities.

State and federal law enforcers, the plaintiffs' bar, and regulators should apply existing law across the AI chat interface ecosystem, scrutinizing both the companies that deploy these systems and the developers that build them. For the plaintiffs' bar, violations are often lying in plain sight, as it is easy to test AI chatbot interfaces for compliance.²³ State attorneys general also have broad subpoena powers under their UDAP statutes, among other authorities. State banking regulators have an even more direct tool when it comes to financial institutions' customer service practices: routine supervisory examination authority that lets them demand information from regulated entities. Enforcers and regulators should seek documents describing the design, deployment, oversight, and impact of these systems and should not hesitate to act where they find the failures described above. Example information requests are located in the Appendix accompanying this issue brief.

Bolster Existing Laws.

New legal requirements should be adopted to complement—not displace—existing legal obligations. In particular, policymakers can build on current law by establishing a “click to get human help” requirement—a version of a human-in-the-loop requirement—but one that guarantees a simple, clearly visible click to reach a human within a reasonable timeframe, not a maze of menus or a doom loop. For example, California's recently enacted Customer Service Chatbots law requires large businesses to provide a clear way to request a human agent and to connect the consumer to that human agent in a timely manner.²⁴ To ensure that “human in the loop” can deliver meaningful outcomes for consumers, policymakers should pair any human-access requirement with substantive accountability measures, such as mandatory public reporting on companies' consumer complaint resolution rates. This can help ensure that human involvement is not simply “procedural theater,”²⁵ but actually helps resolve consumers' problems.

²³ For example, the Center for Law and the Economy recently exposed “Made in USA” fraud through interactions with publicly accessible AI chatbots of Walmart and Amazon. See Erie Meyer & Zachary Harris, *Made in America, Hidden by AI: AI Shopping Chatbots Reveal How Amazon and Walmart Hurt American Manufacturers and Deceive Shoppers* (Ctr. for Law & the Econ., Columbia L. Sch. July 2026), <https://law-and-economy.law.columbia.edu/sites/law-and-economy.law.columbia.edu/files/content/Made%20in%20America.%20Hidden%20by%20AI.pdf>.

²⁴ A.B. 1609, 2025–2026 Leg., Reg. Sess. (Cal. 2026), https://calmatters.digitaldemocracy.org/bills/ca_202520260ab1609. The law provides a start to build upon, though loopholes currently included in it risk undermining its purpose. For example, the law applies only to businesses with at least \$500 million in annual revenue, exempting a large swath of companies deploying customer service chatbots, and it requires only a “good faith effort” to connect consumers to a human within 15 minutes—a standard that leaves companies room to argue substantial compliance even when they fall short. *Id.*

²⁵ Nicolas Spatola, *AI Efficiency Can Undermine Accountability Even With Humans in the Loop*, Tech Pol'y Press (May 5, 2026), <https://www.techpolicy.press/ai-efficiency-can-undermine-accountability-even-with-humans-in-the-loop/>.

Enforcers, regulators, and policymakers can also use guidance, rules, or other public statements to put businesses on notice that current laws generally place responsibility for chatbot harms on both the companies deploying them and the AI developers building them. States may also wish to consider adopting statutory provisions like California's recently enacted law specifying that in any civil action "it shall not be a defense, and the defendant may not assert that the artificial intelligence autonomously caused the harm to the plaintiff."²⁶ As noted above, traditional principles of principal-agent liability already render companies liable for their agent's actions, but provisions like these can save time and resources by foreclosing claims to the contrary by the defense bar.

Strengthen Enforcement and Supervisory Capacity.

Policymakers should adequately resource state attorneys general so that existing laws are actually enforced, especially as AI automation becomes widespread and the scope and scale of violations increases. For decades, agencies have been stretched thin or slashed, even as economic conditions have heightened the need for their action. To help meet the scale of the problem, policymakers should increase state attorneys general's budgets, and enforcers across the country should hire more corporate accountability cops, increase in-house technical expertise, and improve the efficacy of their consumer complaint infrastructure.

Open Regulatory Information and Processes to the Public.

Agencies that accept consumer complaints should publish them, so that regulators, the press, industry, and advocates can spot emerging trouble. Indeed, the CFPB's own complaint system—which has helped more than 12 million people and returned \$93.5 million to consumers in 2024 alone—shows this model works at scale,²⁷ even as the agency moved in August 2026 to stop publishing the very narratives that made it useful in the first place.²⁸ Components—including the design and logic of the complaint system as well as the tools used to redact personal details from the complaints—are open source and available for use by other jurisdictions.²⁹

Regulators and enforcers should also crowd-source information about current market practices, emerging trends, and key open questions, such as by publishing research questions for outside experts and by holding official agency meetings to receive public testimony and evidence. Real-world experiences—including with AI chatbots—should shape and inform the work.

²⁶ Act of Oct. 13, 2025, ch. 672, 2025 Cal. Stat. 672 (codified at Cal. Civ. Code § 1714.46).

²⁷ Erie Meyer, *Complaints.gov: Building a Modern Infrastructure for Accountability*, Vanderbilt Policy Accelerator (Dec. 2025), https://cdn.vanderbilt.edu/vu-URL/wp-content/uploads/sites/412/2025/12/15201553/Complaints.gov_.pdf.

²⁸ Consumer Fin. Prot. Bureau, *The CFPB to Cease Discretionary Publication of Complaint Narratives and Visualizations* (Aug. 14, 2026), <https://www.consumerfinance.gov/about-us/newsroom/the-cfpb-to-cease-discretionary-publication-of-complaint-narratives-and-visualizations/>.

²⁹ CFPB Open Tech, <https://cfpb.github.io/> (last visited Sep. 23, 2026).

The team is grateful to the expert colleagues and staff who reviewed and contributed to this piece, as well as to those whose prior work helped inform it. Thank you to the following individuals, listed in alphabetical order by last name: Winston Berkman-Breen, Atur Desai, Douglas Farrar, Seth Frotman, Hannah Garden-Monheit, Eric Halperin, Lina M. Khan, Min Su Kim, Brad Lipton, Lev Menand, Erie Meyer, Sarah Miller, Maxwell Niles, Mike Pierce, Lorelei Salas, and Abby Etta Smith.

Appendix: Building the Record — Questions Regulators Should Ask

This document outlines a set of requests regulators can adapt for a civil investigative demand (“CID”), subpoena, or supervisory information request to any company relying on AI-related chatbots for customer-facing functions. The goal is to build a robust and actionable record: what the company deployed the chatbot for, how (or whether) it was tested, what happened when it failed, and what internal accountability mechanisms existed to catch failures when they occurred. For simplicity, this appendix uses “AI system” to refer broadly to a system that processes data automatically and produces output used in consequential decisions.³⁰

1. Deployment and Scope

- a. Documents sufficient to identify all AI-related customer service tools deployed, by product line and channel (phone, chat, email), and the dates of deployment;
- b. Documents sufficient to identify which types of customer interactions are handled (i) solely by AI without human review, (ii) solely by a human, and (iii) by a combination of AI and human handling, broken down by product line or channel, for the relevant period; and
- c. Organizational charts or documents identifying who is responsible for the design, deployment, and oversight of these systems.

2. Design, Testing, and Validation

- a. Documents sufficient to describe how the AI system was tested for accuracy and reliability prior to deployment, including any accuracy benchmarks, pass/fail criteria, or pilot results;
- b. All internal analyses, audits, or assessments evaluating the AI system’s accuracy, consistency, or failure rate, before and after deployment; and
- c. Documents describing the underlying model(s) used, including whether built in-house, licensed from a third party, or a general-purpose LLM adapted for customer service.

3. Accuracy, Errors, and Consumer Harm

- a. Documents sufficient to identify the percentage or volume of customer interactions handled solely by AI without human review, by month, broken down by product line or channel, for the relevant period;
- b. All documents relating to instances in which the AI system provided inaccurate, inconsistent, incomplete, or misleading information to consumers, including documents reflecting how those issues were investigated, escalated, addressed, or resolved;
- c. All documents relating to consumer complaints—whether received directly by the company or via regulators (CFPB, state AG, BBB, etc.)—referencing the AI system, including documents reflecting how each was investigated, escalated, addressed, or resolved; and
- d. Documents sufficient to show how the company tracks or measures whether the AI system actually resolves consumer inquiries (e.g. resolution rate, repeat-contact rate, escalation rate), and for any such metric, the underlying data for every month during the relevant period, broken out separately for each of the three interaction categories identified in 1(b).

4. Human Access and Escalation

³⁰ See, e.g., *FTC Announces Crackdown on Deceptive AI Claims and Schemes*, Fed. Trade Comm’n (Sep. 25, 2024), <https://www.ftc.gov/news-events/news/press-releases/2024/09/ftc-announces-crackdown-deceptive-ai-claims-schemes>.

4. Human Access and Escalation

- a. Documents sufficient to describe all methods by which a consumer can reach a human representative, and any limitations on that access (hours, wait times, availability by channel);
- b. Data sufficient to show average, median, and maximum wait or response times for consumers seeking to escalate from AI system to a human, by month, for the relevant period; and
- c. Documents sufficient to identify all categories of inquiries for which the AI system is prohibited from serving as the sole point of contact, if any, and how that prohibition is enforced.

5. Staffing and Cost Reduction

- a. Documents sufficient to show changes in human customer service staffing levels, by function, before and after AI deployment;
- b. All communications concerning the business case, cost savings, or projected human staffing reductions associated with AI deployment; and
- c. Documents reflecting any decision to reverse, pause, or scale back AI-driven customer service due to performance or accuracy concerns.

6. Vendor and Contractual Relationships

- a. All contracts, statements of work, or other agreements with third parties relating to AI-related customer service tools; and
- b. All communications with any such third parties concerning error rates, limitations, or performance issues.

7. Consumer Communications and Marketing

- a. All communications to consumers describing or promoting [Company's] customer service functions, including any representations regarding the quality, availability, responsiveness, or reliability of customer service;
- b. All communications to consumers relating to [Company's] use of AI-related customer service tools, including any claims regarding accuracy, efficiency, or capability; and
- c. Documents sufficient to identify what disclosures, if any, are made to consumers when they are interacting with an AI-related customer service tool.

8. Surge Capacity and Crisis Readiness

- a. Documents sufficient to describe how each AI-related customer service tool and any supporting human staff are expected to perform under a significant increase in customer inquiries, including any stress testing, load testing, or capacity planning conducted for the system;
- b. Documents sufficient to show current staffing levels for human customer service representatives, and whether those levels would be sufficient to handle a significant increase in inquiry volumes if any AI-related customer service system underperforms or fails; and
- c. Documents relating to [Company's] ability to rapidly re-hire, redeploy, or contract additional human customer service capacity in the event of a significant increase in customer inquiries, or in the event any AI-related customer service systems underperforms or fails.