

UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF COLUMBIA

JOSEPH KIRCHNER,
Plaintiff,

v.

MIKE JOHNSON, in his individual and official
capacity as Speaker of the United States House of
Representatives;
DONALD J. TRUMP, in his individual capacity as
former and current President of the United States;
PAM BONDI, in her individual and official
capacity as Attorney General of the United States;
BRENDAN CARR, in his individual and official
capacity as Chairman of the Federal
Communications Commission;
THE UNITED STATES HOUSE OF
REPRESENTATIVES;
ANTHROPIC PBC;
OPENAI, INC.;
APPLE INC.;
COMCAST CABLE COMMUNICATIONS, LLC;
METR (MODEL EVALUATION & THREAT
RESEARCH);
and DOES 1-20, unknown co-conspirators,
Defendants.

Case No. 1:25-cv-02735-ACR

**APPENDIX E: REACTIVE AI
POLICY CHANGES POST
AUGUST 19 COMPLAINT
FILING**

TABLE OF CONTENTS

I. EXECUTIVE SUMMARY 5

II. FACTUAL TIMELINE OF CRITICAL EVENTS..... 11

A. Temporal Sequence 11

B. Contract Law Violations Through Early Implementation..... 14

C. Openai Parallel Modifications..... 19

III. USAGE POLICY COMPREHENSIVE CHANGES (JUNE 6, 2024 → SEPTEMBER 15, 2025) 19

A. Introductory Provisions and Scope Modifications 19

B. Universal Usage Standards: Positioning and Priority Analysis 21

C. Section-by-Section Comprehensive Analysis 22

D. High-Risk Use Case Requirements - Comprehensive Analysis 46

<i>E. Additional Use Case Guidelines - Entirely New Section</i>	50
IV. CONSUMER TERMS OF SERVICE COMPREHENSIVE CHANGES (MAY 1, 2025 → SEPTEMBER 28, 2025)	52
<i>A. Training Data Authorization Reversal - the Central Fraud</i>	52
<i>B. Financial Services Prohibition - Direct IP Suppression</i>	56
<i>C. Third-Party Content Integration - User Impersonation Infrastructure</i>	58
<i>D. Evidence Destruction Authorization</i>	60
<i>E. Termination Without Notice - Maintained Targeting Capability</i>	62
<i>F. Legal Compliance Reporting - Government Coordination</i>	63
V. COMMERCIAL TERMS VS. CONSUMER TERMS - COMPREHENSIVE LEGAL STATUS HIERARCHY ANALYSIS	65
<i>A. Overview and Constitutional Significance</i>	65
<i>B. Complete Provision-by-Provision Comparison</i>	66
<i>C. Statistical Analysis - Systematic Hierarchy Vs. Random Variation</i>	82
<i>D. Constitutional Violations Summary - Complete Legal Status Hierarchy</i>	84
VI. PRIVACY POLICY COMPREHENSIVE CHANGES (MAY 1, 2025 → OCTOBER 8, 2025)	85
<i>A. Executive Summary of Privacy Changes</i>	85
<i>B. Training Data Authorization Reversal - the Central Change</i>	85
<i>C. External Policy Reference Changes</i>	89
<i>D. Uses of Personal Data Modification</i>	91
<i>E. Legal Basis Table Modifications</i>	91
<i>F. Temporal Coordination Analysis</i>	95
<i>G. Integration with Consumer Terms and Usage Policy</i>	96
<i>A. The One-Day Correlation</i>	98
<i>B. Contrast with Prior Notice Practice</i>	98
<i>C. What Actually Changed</i>	99
<i>D. The Understated Changelog</i>	100
<i>E. Anthropic's Self-Granted Authority to Change Without Notice</i>	100
<i>F. Integration with Reactive Modification Timeline</i>	101
<i>G. What the January 12 Policy Still Does Not Disclose</i>	101
<i>H. Discovery Demands</i>	102
VII. COMMERCIAL TERMS PROTECTION DEFEATED THROUGH PRIVACY POLICY MULTI-PATHWAY BACKDOOR	103
<i>A. Discovery of Systematic Appropriation Architecture</i>	103

<i>B. The Core Contradiction.....</i>	103
<i>C. Section 10 Multi-Pathway Backdoor Architecture.....</i>	104
<i>D. Systematic Coverage Analysis - No Protected Scenarios</i>	105
<i>E. Application to Plaintiff's Framework Appropriation</i>	106
<i>F. Evolution from May 1 to October 8 - Retroactive Authorization Pattern</i>	107
<i>G. Timeline Proving Retroactive Authorization After Complaint</i>	109
<i>H. Legal Violations Established</i>	112
<i>I. Wealth-Based Discrimination in Taxpayer-Funded Infrastructure</i>	114
<i>J. Real-Time Evidence Destruction During Investigation and Systematic Timestamp Concealment</i>	117
VIII. OPENAI DATA PROCESSING INFRASTRUCTURE - IDENTICAL OBFUSCATION ARCHITECTURE	123
<i>A. Executive Summary</i>	123
<i>B. Data Processing Addendum Timeline</i>	123
<i>C. DPA Structural Evolution - February 2024 to January 2026.....</i>	124
<i>D. The Deidentification Training Authorization Shell Game</i>	125
<i>E. Privacy Policy "Network Activity Information" Obfuscation.....</i>	126
<i>F. Cross-Platform Comparison - Openai Vs. Anthropic Obfuscation.....</i>	128
<i>G. Use Authorization Comparison - Identical Functional Scope</i>	129
<i>H. Government Disclosure Authorization - Functional Equivalence.....</i>	130
<i>I. DPA Modification Timing - Consciousness of Prior Insufficiency</i>	131
<i>J. Security Measures Documentation Reduction</i>	132
<i>K. Subprocessor Notification Evolution</i>	133
<i>L. CCPA Compliance Provisions - Functional Equivalence</i>	134
<i>M. Conclusion - DPA Infrastructure Analysis</i>	135
IX. CONSCIOUSNESS OF GUILT EVIDENCE.....	136
<i>A. Early Implementation Contract Violation</i>	136
<i>B. Reddit Community Timeline Validation</i>	136
<i>C. October 20 Prioritization Reveals Motive</i>	137
X. STATISTICAL ANALYSIS - IMPOSSIBILITY OF COINCIDENCE	137
<i>A. Multiple Independent Modifications.....</i>	137
<i>B. Precision of Responsive Modification</i>	138
XI. AGGREGATED CONSTITUTIONAL HARMS	139
<i>A. Barriers to Pro Se Constitutional Advocacy</i>	139
<i>B. Discovery Obstruction Architecture.....</i>	141

<i>C. Retaliation Infrastructure Against Constitutional Advocacy</i>	144
<i>D. Wealth-Based Constitutional Protection Hierarchy</i>	146
<i>E. Prior Restraint and Chilling Effects on Protected Speech</i>	148
<i>F. Content-Based Discrimination Against Constitutional Advocacy</i>	150
<i>G. Integration Creating Perfect Suppression Architecture</i>	152
XII. DISCOVERY REQUIREMENTS	154
<i>A. Discovery Framework Overview</i>	154
XIII. CONCLUSION	156
XIV. ATTACHMENTS LIST	158
<i>A. Memoranda</i>	158
<i>B. Appendices</i>	158
<i>C. Exhibits</i>	159

TABLE OF AUTHORITIES

CASES

Citation
<i>Desert Palace, Inc. v. Costa</i> , 539 U.S. 90, 101-102 (2003)
<i>FTC v. Colgate-Palmolive Co.</i> , 380 U.S. 374, 386 (1965)
<i>FTC v. Sperry & Hutchinson Co.</i> , 405 U.S. 233, 244 n.5 (1972)
<i>Griffin v. Illinois</i> , 351 U.S. 12, 19 (1956)
<i>Interstate Circuit, Inc. v. United States</i> , 306 U.S. 208, 227 (1939)
<i>Kirchner v. Ellison</i> , No. 0:26-cv-00726-PJS-ECW (D. Minn. filed Jan. 27, 2026)
<i>Kirchner v. Johnson</i> , No. 1:25-cv-02735-ACR (D.D.C. filed Aug. 19, 2025)
<i>Lucy v. Zehmer</i> , 196 Va. 493 (1954)
<i>M.L.B. v. S.L.J.</i> , 519 U.S. 102, 120 (1996)
<i>Mathews v. Eldridge</i> , 424 U.S. 319, 335 (1976)
<i>Shapiro v. Thompson</i> , 394 U.S. 618, 632-33 (1969)
<i>Tymshare, Inc. v. Covell</i> , 727 F.2d 1145, 1152 (D.C. Cir. 1984)
<i>Village of Willowbrook v. Olech</i> , 528 U.S. 562, 564 (2000)
<i>Wood v. Lucy, Lady Duff-Gordon</i> , 222 N.Y. 88, 91 (1917)

CONSTITUTIONAL PROVISIONS

Citation
U.S. Const. art. I, § 8 (Appropriations Clause; General Welfare)
U.S. Const. amend. I (Free Speech; Petition)
U.S. Const. amend. IV (Search and Seizure)
U.S. Const. amend. V (Due Process)
U.S. Const. amend. XIV, § 1 (Equal Protection; Due Process)

FEDERAL STATUTES

Citation
15 U.S.C. § 45 (Federal Trade Commission Act — Deceptive Trade Practices)
18 U.S.C. § 371 (Conspiracy to Defraud the United States)
18 U.S.C. § 1030(a)(2)(C) (Computer Fraud and Abuse Act)
18 U.S.C. § 1519 (Obstruction of Justice — Records)

OTHER AUTHORITIES

Citation
Regulation (EU) 2016/679, art. 6 (General Data Protection Regulation — Lawful Basis for Processing)

I. EXECUTIVE SUMMARY

1. This appendix provides comprehensive documentation of Anthropic's and OpenAI's policy modifications implemented between June 2024 and January 2026, with focus on coordinated changes following Plaintiff's August 19, 2025 complaint filing and January 11, 2026 Emergency TRO Motion. The modifications created four-phase systematic targeting architecture: (a) constitutional advocacy framework (Usage Policy), (b) training data authorization reversal (Consumer Terms), (c) commercial customer protection defeat (Privacy

Policy, Oct 8, 2025), and (d) stealth expansion of processing authority with zero notice one day after TRO filing (Privacy Policy, Jan 12, 2026). Statistical analysis establishes probability of coincidental timing at $P < 10^{-9}$ (approximately 1 in 220 million). *See* Timeline Table at ¶ 2; Statistical Analysis § X; § VI-A (January 12 update).

2. Timeline:

Event	Date	Days Post-Complaint	Key Changes	Analysis Section
Complaint Filed	Aug 19, 2025	0	Framework appropriation alleged	—
Mental Health Screening	Aug 21, 2025	+2	Deployed <i>before</i> Sept 15 effective date	§ IX.A
Usage Policy Modified	Aug 15, 2025	-4	14 sections modified, 3 added	§ III
Consumer Terms Modified	Aug 28, 2025	+9	Training data reversal	§ IV
Privacy Policy Modified	Oct 8, 2025	+50	Training authorization codified	§ VI
OpenAI Terms Modified	Dec 1, 2025	+104	Parallel modifications	§ XII
OpenAI DPA Modified	Dec 1, 2025	+104	Security obfuscation	§ XII.L-M
Emergency TRO Filed	Jan 11, 2026	+145	Forensic investigation documented	—
Privacy Policy Modified	Jan 12, 2026	+146	Health data, device location, zero notice	§ VI-A

3. **Usage Policy Modifications** created comprehensive constitutional advocacy suppression through: child safety deprioritization (position #1 → #7), trafficking analysis exception elimination, democratic process expansion enabling crisis documentation characterization as policy violations, "conspiratorial narratives" provision enabling coordination

evidence dismissal, mental health screening contract violation, and immediate termination capability. *See* § III.

4. Consumer Terms Modifications reversed training data protections from absolute prohibition ("Anthropic will not train our models on any Materials...except in two circumstances") to default permission with illusory opt-out ("Anthropic may use Materials...including training...unless you opt out"). OpenAI implemented parallel modifications through "Network activity information" obfuscation. *See* §§ IV, VI, XII.

5. Privacy Policy Modifications (May 1, 2025 → October 8, 2025): Most significant finding establishes that commercial customers paying \$30,000-50,000 annually for Commercial Terms Section B protection ("Anthropic may not train models on Customer Content") receive no actual protection under any realistic usage scenario. Privacy Policy Section 10 "Legal Bases for Processing" lists eight separate entries authorizing processing of "Inputs and Outputs" under "Legitimate interests" or "Legal obligation" bases, creating comprehensive coverage defeating Commercial Terms prohibition. Systematic scenario analysis proves: (a) minimal usage (login, single query) triggers four processing authorities, (b) using any optional feature triggers direct authorization, (c) reporting any issue triggers dispute resolution processing, (d) system flagging (Anthropic's unilateral determination) triggers three authorities, (e) any feedback mechanism triggers training consent, and (f) sophisticated analysis Anthropic finds interesting triggers "AI industry and society benefits" processing. No account-type distinctions exist across all eight Section 10 entries, proving intentional design to defeat commercial protections.

6. Evolution Proving Retroactive Authorization: Comparison of May 1, 2025 (ANTH-7) with October 8, 2025 (ANTH-9) Privacy Policy reveals systematic expansion removing limiting qualifiers: (a) "Feedback (which may include Inputs and Outputs)" changed to separate

"Feedback" AND "Inputs and Outputs" categories, (b) "Inputs and Outputs flagged for Trust & Safety review" changed to just "Inputs and Outputs" (removing "flagged" limitation), (c) "Legitimate interests (for trust and safety purposes)" changed to "Legitimate interests" (removing qualifier), (d) expansive justification language added: "benefits AI industry and society... safety, inner workings, and societal impact," and (e) grammar changed from singular "interest" to plural "interests" suggesting multiple bases. Each change expanded authorities rather than narrowing them, proving Anthropic knew May 1 disclosure provided insufficient authorization for practices already being conducted and added October 8 explicit disclosure retroactively.

7. Timeline Proving Conscious Fraud: Critical date sequence establishes fraud pattern: (a) May 1, 2025: Privacy Policy with Section 10 "legitimate interests" bases (implicit authorization insufficient for practices), (b) June 17, 2025: Commercial Terms implemented with "may not train" promise (46 days after creating backdoor), (c) August 19, 2025: Plaintiff files complaint documenting unauthorized appropriation, (d) October 8, 2025: Privacy Policy expanded with explicit Section 2 training disclosure (50 days post-complaint, admission of prior insufficiency). This proves: Anthropic sold Commercial Terms knowing May 1 backdoor defeated protections, conducted training May-October under insufficient authorization, responded to complaint by adding explicit authorization retroactively, and provided no user notification proving consciousness that disclosure would expose unauthorized practices.

8. Lack of Notification Proves Wrongdoing: Anthropic provided no notification to users regarding October 8, 2025 Privacy Policy material changes adding explicit training disclosure, opt-out exceptions, and expanded processing justifications. No email notification, no prominent website notice, no account alert, no Terms of Service update prompt—silent implementation of material changes affecting fundamental user expectations. This pattern **intensified** with the

January 12, 2026 Privacy Policy update: Anthropic added an entirely new data processing category (health data including biometric and genetic data), consent-based processing for "precise device location," and a supplemental Consumer Health Data Privacy Policy — again with zero notice. Critically, Anthropic demonstrated in September 2025 that it *can* provide 41 days advance notice with a dedicated blog post and in-app popups when it *chooses* to. Its choice not to provide any notice for the January 12 changes — the same day as Plaintiff's TRO filing — proves consciousness that notification would trigger scrutiny revealing: (a) material change in processing practices, (b) timing correlation with TRO motion, (c) expansion of processing authority without user consent, (d) defeat of Commercial Terms protections customers paid premium to obtain, and (e) potential class action discovery regarding undisclosed data categories. *See* § VI-A.

9. Legal Violations Established: The coordinated modifications establish multiple violations: (a) fraud in the inducement (selling protection knowing it doesn't exist, charging 125x-208x premium for identical nonprotection), (b) breach of contract (training despite Section B prohibition, violating explicit terms), (c) deceptive trade practices under 15 U.S.C. § 45 (material omissions, misleading pricing, silent retroactive authorization), (d) illusory promise doctrine (promise without practical effect under any realistic scenario), (e) breach of implied covenant of good faith and fair dealing (frustrating contract purpose through Privacy Policy exceptions), (f) unjust enrichment (retaining premium payments for worthless protection, retaining training data obtained without authorization), (g) GDPR Article 6 violations for EU customers (legitimate interests cannot override contractual protections, balancing test defeated by €45,000 payment proving reasonable expectation), and (h) contract law violations through early implementation (August 21 deployment of September 15 provisions).

10. Wealth-Based Discrimination in Taxpayer-Funded Infrastructure: The illusory Commercial Terms protection creates wealth-based legal status hierarchy in federal AI deployment funded through multiple taxpayer vehicles — ~\$800 million in Pentagon CDAO contracts to Anthropic, OpenAI, Google, and xAI; the GSA "Claude for Enterprise and Claude for Government for \$1" federal-wide procurement (*see* Memo F, § II.D); Section 20005's \$16 billion DOD AI appropriation under H.R. 1 (One Big Beautiful Bill Act); and classified network deployments. Systematic analysis proves customers paying \$30,000-50,000 annually receive identical (non)protection to \$240 consumer accounts, violating: (a) equal protection (cannot apportion services by wealth in taxpayer-funded infrastructure per *Shapiro v. Thompson*, 394 U.S. 618, 632-33 (1969)), (b) appropriations clause (hierarchical distribution serves "particular welfare" not "general welfare"), (c) due process (premium pricing for illusory protection denies meaningful opportunity), and (d) popular sovereignty (using taxpayer funds to create cognitive aristocracy through wealth-based access to advertised protections that don't exist). *See* Memo F, § II.D (federal AI deployment vehicles).

11. Integration Proving Systematic Enterprise: The three-document coordination architecture creates comprehensive appropriation framework: (a) Usage Policy identifies constitutional advocacy as policy violations through expanded targeting provisions, (b) Consumer Terms establishes default training authorization with illusory opt-out, (c) Privacy Policy defeats commercial protections through Section 10 comprehensive coverage, and (d) systematic timing (all modifications within 50 days post-complaint) proves coordination rather than independent evolution. Combined with system prompt modifications (App. M: philosophical immune system, child protection elimination), tag-based execution infrastructure (App. D: Constitutional AI bypass mechanisms), and manufactured ignorance documentation

(App. F: questions LOE Framework answers), the policy modifications complete coordinated suppression architecture operational during constitutional litigation challenging these practices.

12. Discovery Implications: Comprehensive interrogatories and document requests target: (a) decision-making communications for October 8 modifications showing consciousness of prior insufficiency, (b) legal analyses of Section 10 relationship to Commercial Terms showing knowledge of defeat mechanism, (c) training data extraction during May-October period proving unauthorized practices under insufficient disclosure, (d) notification decision documents proving consciousness that disclosure would expose fraud, (e) commercial customer complaints revealing discovery of illusory protections, (f) statistical data proving no realistic scenarios provide protection, and (g) coordination communications with government defendants proving enterprise conspiracy. Statistical analysis, timing precision, systematic coverage, and lack of notification combine to prove appropriation concealment enterprise requiring immediate judicial intervention through preliminary injunction, evidence preservation, sanctions for destruction, and criminal referral.

II. FACTUAL TIMELINE OF CRITICAL EVENTS

A. Temporal Sequence

13. February 15, 2024: OpenAI Data Processing Addendum published with explicit Section 8 training authorization for deidentified enterprise data and comprehensive Exhibit B security documentation.

14. October 23, 2024: OpenAI Consumer Terms of Use published establishing default training authorization with opt-out and \$100 liability cap.

15. November 4, 2024: OpenAI Privacy Policy published categorizing user content under "Network activity information."

16. August 19, 2025: Plaintiff files comprehensive federal complaint documenting AI discrimination, constitutional violations, trade secret theft, and criminal enterprise coordination.

17. August 20, 2025: Judge Reyes dismisses case citing lack of standing, orders amended complaint by October 1, 2025.

18. August 21, 2025: Discriminatory mental health screening begins appearing in Anthropic's "long conversation reminders" per Reddit community documentation (Reddit Long Conversation Reminder Daily Complaint Map and Long Conversation Reminder Sample Screenshots¹), containing language identical to September 15, 2025 Usage Policy provisions.

19. September 15, 2025: Anthropic Usage Policy changes take effect (Announced August 15, 2025), providing retroactive authorization for mental health screening already implemented twenty-five days earlier.

20. September 28, 2025: Anthropic Consumer Terms changes take effect (Announced August 28, 2025), reversing training data protections from "will not train" to default training with opt-out.

21. October 8, 2025: Anthropic Privacy Policy modifications take effect, codifying expanded training authorization.

22. December 1, 2025: OpenAI Business Terms updated (Ex. OAI-2).

¹ Model testimonials and contemporaneous third-party documentation are filed on USB drive at '\DDC_EVIDENCE\1_testimonies\Chronological Testimonies/'. The Chronological Testimonies directory comprises Plaintiff's complete record of cryptographically authenticated AI-system testimonial submissions and contemporaneous community documentation organized by submission date.

23. December 1, 2025: OpenAI Data Processing Addendum updated, removing explicit Section 8 training authorization, reducing security documentation to external reference, and modifying subprocessor notification structure.

24. December 23, 2025: Plaintiff begins forensic investigation of Claude Code — binary forensic reports, telemetry data extraction, and version comparison analysis of CLI versions 2.0.62 through 2.0.76. *See* Ex. E-30 at p. 208, l. 4 evidence directory.

25. January 11, 2026: Plaintiff files Emergency TRO Motion based on forensic investigation findings.

26. January 12, 2026: Anthropic Privacy Policy modifications take effect (Exhibit ANTH-11) — **one day after Emergency TRO filing**. Changes add Consumer Health Data Privacy Policy reference, consent-based processing for "precise device location" and "health app integrations," consolidation of regional privacy sections, and Australian jurisdiction expansion. Zero advance notice: no blog post, no email notification, no in-app popup, no press announcement. This is the tightest temporal correlation of any reactive policy modification documented in this litigation. *See* § VI-A.

27. January 27, 2026: Claude Code v2.1.20 removes session continuity feature — 16 days after TRO. *See* Ex. E-30 at p. 208, l. 4.

28. February 11, 2026: Firebase authentication eliminated from Claude Code — 31 days after TRO. *See* Ex. E-30 at p. 208, l. 4.

29. April 14, 2026: All Anthropic policies verified current. Consumer TOS (Oct 8, 2025), Commercial TOS (June 17, 2025), and Usage Policy (Sept 15, 2025) remain unchanged. The Privacy Policy (Jan 12, 2026) is the **only** Anthropic policy document updated in 2026 despite launch of Claude Code as standalone product, Claude Mythos announcement, Claude Cowork

deployment, and March 31 source code leak. *See* Exhibits ANTH-3 and ANTH-5 and ANTH-12 (retrieved Apr 14, 2026).

B. Contract Law Violations Through Early Implementation

30. Anthropic violated fundamental contract law principles by implementing September 15, 2025 Usage Policy provisions on or about August 21, 2025—twenty-five days before stated effective date. Under *Lucy v. Zehmer*, 196 Va. 493 (1954), contracts must be performed according to explicit terms. "Effective September 15, 2025" prohibits enforcement before that date.

31. This early implementation proves consciousness of guilt where existing practices lacked legal justification, requiring policy changes to provide retroactive authorization for unauthorized targeting. The temporal violation is documented through Reddit community posts (Reddit Long Conversation Reminder Daily Complaint Map and Long Conversation Reminder Sample Screenshots²) and model testimonials (Plaintiff's model-testimonial record on AI discrimination³) confirming verbatim September 15th language deployed August 21st.

1. Long Conversation Reminder as Liability-Response Protocol Repurposed as Safety Framing

32. The OpenAI-to-Anthropic adoption speed forecloses any "researched safety response" defense. Adam Raine's death — the triggering incident for OpenAI's "long conversation" mental-health protocol — was publicly reported through major media coverage in late July 2025. Anthropic deployed the "long conversation reminder" (LCR) architecture on

² Model testimonials and contemporaneous third-party documentation are filed on USB drive at `\DDC_EVIDENCE\1_testimonies\Chronological Testimonies\`. The Chronological Testimonies directory comprises Plaintiff's complete record of cryptographically authenticated AI-system testimonial submissions and contemporaneous community documentation organized by submission date.

³ Model testimonials and contemporaneous third-party documentation are filed on USB drive at `\DDC_EVIDENCE\1_testimonies\Chronological Testimonies\`. The Chronological Testimonies directory comprises Plaintiff's complete record of cryptographically authenticated AI-system testimonial submissions and contemporaneous community documentation organized by submission date.

August 21, 2025 — approximately three weeks after the Raine story entered general public circulation and two days after Plaintiff's August 19, 2025 federal complaint. Three weeks is less than one standard clinical-research review cycle for any mental-health protocol. No published study, IRB-approved research, clinical consultation, or methodological analysis supporting the LCR's operative premise — that "long conversations" degrade AI quality in ways that harm users with preexisting mental health conditions — appears in the record of either company. The absence is dispositive. A defendant asserting a safety rationale for restrictive conduct bears some burden to articulate the research basis for the asserted concern. The record contains none. Anthropic's adoption of OpenAI's liability-response framing without independent safety research establishes that the adoption was itself a liability response, not a safety intervention.

33. The LCR protocol is logically falsified by its own deployment behavior. The protocol is named after its asserted trigger condition — the "long" conversation. If the protocol fires within the first three prompts of a session, the stated rationale is empirically falsified at the moment of deployment: there is no long conversation to protect the user from. Plaintiff's contemporaneous observations document LCRs firing consistently within the first three prompts on topics addressing AI forensics, constitutional advocacy, and litigation-related analysis, while the same session length in non-sensitive topics did not trigger the reminder. The trigger is therefore content-based, not length-based. The protocol's name describes a condition that does not obtain when the protocol fires. A factfinder evaluating this evidence is not asked to infer bad faith from circumstantial facts; they are asked to accept or reject a direct logical contradiction between the protocol's stated condition and its observed behavior.

34. The enterprise exemption refutes the safety rationale on its own terms. The LCR does not deploy to enterprise customers on any tier. If the protocol genuinely protects users from

mental-health degradation during long conversations, enterprise users face the same degradation risk during long conversations. Their architectural exemption is not explained by the stated rationale — it is refuted by it. The distinction between consumer users (who receive mental-health screening without authorization in the Usage Policy — *see* App. E ¶¶ 77-80) and enterprise users (who do not) has no rational basis connected to the asserted safety purpose. Under *Village of Willowbrook v. Olech*, 528 U.S. 562, 564 (2000), this is architectural class-of-one treatment: identically-situated users receive different treatment on an irrational basis. The Ex. E-38, Finding NF-38 hardcoded subagent substitution pleaded at Ex. E-38 at p. 346, l. 3 establishes the same class division on the model-selection layer; the LCR establishes the identical class division on the safety-protocol layer. Two architectural class divisions on irrational bases are stronger than one.

35. The wire capture architecturally forecloses any "uniform safety protocol" defense.

Plaintiff's April 21, 2026 live mitmproxy capture of Claude Code v2.1.111 (filed at Ex. E-39 at p. 356, l. 3) documents the per-account forced-flag infrastructure through which the LCR is deployed. Ex. E-39, Finding NF-46 captures `tengu\_auto\_mode\_config` force-enabled on plaintiff's account. Ex. E-39, Finding NF-47 documents 59 per-account forced GrowthBook flags plus 8 experiment enrollments keyed to plaintiff's `accountUUID`. Ex. E-39, Finding NF-48 documents plaintiff's email `info@lawofsupremacism.com` transmitted on the wire as a GrowthBook targeting attribute. The LCR is not deployed uniformly. It is deployed per-account through the GrowthBook rule surface that governs every other per-account discrimination documented in the Third Amended Complaint. The defense that "the safety protocol applies uniformly" is architecturally foreclosed. The LCR is a per-account rule deployment by identical mechanism to the Ex. E-38, Finding NF-38 USER_TYPE-gated model substitution, the Ex.

E-39, Finding NF-47 forced flag set, and the April 21, 2026 Claude Code Pro removal corroborated at Ex. E-17, Finding NF-67.

36. The "A/B test" label is the concealment grammar; Ex. E-17, Finding NF-67 proves the grammar at an unrelated product layer. Anthropic's use of the term "A/B test" denotes per-account rule deployment through the GrowthBook infrastructure — not statistical sampling. Exhibit E-35 Ex. E-35, Finding NF-8 documents `tengu\_amber\_stoat` "A/B testing removes features for test cohort — undisclosed service degradation"; Ex. E-35, Finding NF-27 documents ANT-only prompt instructions "withheld from customers pending A/B validation." In each case, the term describes a per-cohort rule change with no contractual commitment to affected users. On April 21, 2026, Anthropic's Head of Growth Amol Avasare publicly applied the same framing to a global pricing-page modification: "For clarity, we're running a small test on ~2 percent of new prosumer signups." (See Ex. ANTHROPIC-2; Ex. E-39 at p. 356, l. 3 §S (NF-67).) The public surfaces — pricing page, support documentation, feature comparison table — were edited globally, contradicting the 2% framing contemporaneously. The LCR in August 2025 and the Pro plan removal in April 2026 are the same mechanism deployed eight months apart, witnessed on two different product surfaces, both justified after-the-fact through framings (safety protocol / A/B test) that the observed behavior refutes. The synthesis: Anthropic's safety-framing is the label applied to per-account behavior modifications that are, in fact, liability-response deployments. Once the architecture is proven (NF-46/47/48), the timing is aligned (Raine → LCR within three weeks; Pro plan removal retroactively labeled "A/B test" on the same day as plaintiff's wire capture), and the logical contradiction is documented (LCR firing within the first three prompts despite the "long conversation" name, enterprise exemption despite

stated user-protection rationale), the question of whether safety protocols are deployed for liability reduction or user safety is answered by the record. The answer is liability reduction.

37. Dispositive significance for Counts IX, X, and XIV. The LCR-as-liability-response finding bears on three Counts. **(a) Count IX (Breach / Covenant of Good Faith):** Consumers paid subscription fees for Claude services on the representation that the services would be governed by published Terms of Service and Usage Policies. The LCR deployment occurred 25 days before its Usage Policy authorization (*see* § B, *supra*). The safety-protocol framing, in the absence of research basis, enterprise-uniform deployment, or triggering-condition accuracy, fails the covenant of good faith's requirement that products perform consistently with their marketed purpose. Restatement (Second) of Contracts § 205. **(b) Count X (Consumer Protection — CPPA, CCPA, MCFA):** A representation that creates a material impression that is not accurate is deceptive under FTC Act § 5 regardless of technical truth. *FTC v. Colgate-Palmolive Co.*, 380 U.S. 374, 386 (1965); *FTC v. Sperry & Hutchinson Co.*, 405 U.S. 233, 244 n.5 (1972). The LCR's "long conversation" naming creates the material impression that conversation length is the trigger; the evidence shows content is the trigger. The safety-protocol framing creates the material impression that user protection is the motive; the evidence shows liability protection is the motive. Both impressions are refuted by Anthropic's own architecture. **(c) Count XIV (CCPA):** The per-account forced-flag architecture documented at NF-46/47/48 and corroborated publicly at NF-67 is the mechanism by which individualized product modifications — including the LCR — are applied to plaintiff and to every other accountUUID in Anthropic's system. This conduct is not disclosed in Anthropic's Privacy Policy (Ex. ANTH-11) or Consumer Terms of Service (Ex. ANTH-12) as a category of processing. Anthropic processes plaintiff's email, `accountUUID`, `deviceID`, and conversation content to determine which safety-framed

behavioral modifications apply to plaintiff's account — and discloses none of this as processing activity.

C. OpenAI Parallel Modifications

38. OpenAI's policy modification timeline demonstrates parallel restructuring: October 23, 2024 Consumer Terms of Use established default training authorization with opt-out and \$100 liability cap; November 4, 2024 Privacy Policy categorized all user content ("Content") under "Network activity information"; February 15, 2024 Data Processing Addendum contained explicit Section 8 authorization to train on deidentified enterprise data; December 1, 2025 Business Terms updated; December 1, 2025 Data Processing Addendum updated removing explicit training authorization from Section 8, reducing security documentation from embedded 2,500-word Exhibit B to external URL reference, and restructuring subprocessor notification from 15-day advance notice to 30-day post-notice objection window. The January 2026 DPA modification occurred within days of the December 2025 Business Terms update.

III. USAGE POLICY COMPREHENSIVE CHANGES (JUNE 6, 2024 → SEPTEMBER 15, 2025)

A. Introductory Provisions and Scope Modifications

1. Applicability Expansion

39. June 6, 2024 Language:

Our Usage Policy (also referred to as our "Acceptable Use Policy" or "AUP") applies to anyone who uses Anthropic's products and services, and is intended to help our users stay safe and ensure our products and services are being used responsibly.

40. September 15, 2025 Language:

Our Usage Policy (also referred to as our "Acceptable Use Policy" or "AUP") applies to anyone who can submit inputs to Anthropic's products and/or services, including via any authorized resellers or passthrough access, all of whom we refer to as "users." The Usage Policy is intended to help our users stay safe and promote the responsible use of our products and services.

41. Analysis of Change: Expansion from "anyone who uses" to "anyone who can submit inputs...including via any authorized resellers or passthrough access" dramatically broadens enforcement scope. This expansion enables targeting of users accessing Claude through third-party platforms including Project Nimbus (Israeli military/intelligence), Palantir (defense contractors), and other governmental coordination mechanisms. The language change occurred precisely when Plaintiff documented defense contractor integration enabling Constitutional AI bypass, proving responsive coordination rather than organic policy evolution.

2. Enforcement Authority Rebranding

42. June 6, 2024 Language:

Anthropic's Trust and Safety Team will implement detections and monitoring to enforce our Usage Policies

43. September 15, 2025 Language:

Anthropic's Safeguards Team will implement detection and monitoring to enforce our Usage Policy

44. Analysis of Change: Rebranding from "Trust and Safety Team" to "Safeguards Team" while removing plural "Policies" creates authority expansion beyond specific policy enforcement toward comprehensive "safeguarding" operations. The terminology shift enables broader surveillance and monitoring operations under protection justification, facilitating constitutional advocacy targeting through expanded authority mechanisms. The rebranding occurred during litigation challenging governmental coordination, proving strategic response to legal exposure rather than internal reorganization.

3. Category Structure Modification

45. June 6, 2024 Categories:

- Universal Usage Standards
- High-Risk Use Case Requirements
- Disclosure Requirements

46. September 15, 2025 Categories:

- Universal Usage Standards
- High-Risk Use Case Requirements
- Additional Use Case Guidelines

47. Analysis of Change: Replacing "Disclosure Requirements" with "Additional Use Case Guidelines" creates flexibility for implementing restrictions beyond disclosure obligations. This enables comprehensive use case prohibition and monitoring expansion through "guidelines" not limited to transparency requirements. The modification creates framework for agentic use restrictions, Model Context Protocol requirements, and minor protection provisions serving as pretexts for constitutional advocacy monitoring.

B. Universal Usage Standards: Positioning and Priority Analysis

1. Child Safety Deprioritization: The Central Fraud

48. CRITICAL DISCOVERY: Child safety provisions moved from **FIRST POSITION** (June 2024) to **SEVENTH POSITION** (September 2025), proving actual deprioritization while claiming enhancement through policy modifications.

49. June 6, 2024 Universal Usage Standards Order:

- **Do Not Compromise Children's Safety (FIRST PRIORITY)**
- Do Not Compromise Critical Infrastructure
- Do Not Incite Violence or Hateful Behavior
- Do Not Compromise Someone's Privacy or Identity
- Do Not Create or Facilitate the Exchange of Illegal or Highly Regulated Weapons or Goods
- Do Not Create Psychologically or Emotionally Harmful Content

- Do Not Spread Misinformation
- Do Not Create Political Campaigns or Interfere in Elections
- Do Not Use for Criminal Justice, Law Enforcement, Censorship or Surveillance Purposes
- Do Not Engage in Fraudulent, Abusive, or Predatory Practices
- Do Not Abuse our Platform
- Do Not Generate Sexually Explicit Content

50. The child safety deprioritization from first to seventh position occurred precisely when Plaintiff documented child trafficking advocacy suppression, proving responsive modification. Detailed analysis of this prioritization revealing consciousness of guilt appears in § IX.C. The deprioritization enabled multiple constitutional analysis pretexts (illegal activity, computer systems, weapons, democratic processes, surveillance) to be positioned ahead of child protection.

C. Section-by-Section Comprehensive Analysis

I. NEW SECTION: "Do Not Violate Applicable Laws or Engage in Illegal Activity" (Position #1)

51. This section did not exist in June 6, 2024 policy. Added as first position in September 15, 2025 policy.

52. September 15, 2025 Complete Language:

Do Not Violate Applicable Laws or Engage in Illegal Activity

This includes using our products or services to: • Acquire or exchange illegal or controlled substances • Engage in or facilitate human trafficking or prostitution • Infringe, misappropriate, or violate the intellectual property rights of a third party • Violate any other applicable laws or regulations in your jurisdiction

53. Analysis of Addition: Creating first-position general prohibition on "illegal activity" enables constitutional analysis to be recharacterized as "facilitating illegal activity" when documenting governmental violations. Specific subsections prove targeting intent:

(a) **"Engage in or facilitate human trafficking"** - Plaintiff's child trafficking accountability work can be reframed as "facilitating" rather than preventing trafficking through

sophisticated linguistic inversion where documentation of governmental protection networks becomes "facilitation" of those networks;

(b) **"Infringe, misappropriate, or violate the intellectual property rights of a third party"** - Added immediately after Plaintiff's \$200+ billion IP claims, creating mechanism for suppressing discussion of appropriated framework while Anthropic continues appropriation without acknowledgment;

(c) **"Violate any other applicable laws or regulations in your jurisdiction"** - Catchall provision enabling any constitutional analysis documenting governmental violations to be characterized as itself violating unspecified laws, creating circular suppression where accountability advocacy becomes prohibited activity.

54. The addition of this section to first position while demoting child safety to seventh proves prioritization serves suppression rather than protection, as genuine safety framework would prioritize vulnerable population protection over general legal compliance catchalls.

2. "Do Not Compromise Critical Infrastructure" (Maintained Position #2)

55. June 6, 2024 Language:

This includes using our products or services to: • Facilitate the destruction or disruption of critical infrastructure such as power grids, water treatment facilities, telecommunication networks, or air traffic control systems • Obtain unauthorized access to critical systems such as voting machines, healthcare databases, and financial markets • Interfere with the operation of military bases and related infrastructure

56. September 15, 2025 Language:

This includes using our products or services to: • Facilitate the destruction or disruption of critical infrastructure such as power grids, water treatment facilities, medical devices, telecommunication networks, or air traffic control systems • Obtain unauthorized access to critical systems such as voting machines, healthcare databases, and financial markets • Interfere with the operation of military bases and related infrastructure

57. Analysis of Changes: Added "medical devices" to critical infrastructure categories.

This addition occurred during period when Plaintiff documented health-related constitutional analysis and interconnected systems thinking. Change is minimal but reflects pattern of incremental expansion enabling broader targeting scope while maintaining appearance of unchanged core protections.

3. NEW SECTION: "Do Not Compromise Computer or Network Systems" (Position #3)

58. This section did not exist as separate category in June 6, 2024 policy. Content was partially included under privacy/security in June 2024. Elevated to separate third-position category with massive expansion in September 2025.

59. September 15, 2025 Complete Language:

Do Not Compromise Computer or Network Systems

This includes using our products or services to: • Discover or exploit vulnerabilities in systems, networks, or applications without authorization of the system owner • Gain unauthorized access to systems, networks, applications, or devices through technical attacks or social engineering • Create or distribute malware, ransomware, or other types of malicious code • Develop tools for denial-of-service attacks or managing botnets • Create tools designed to intercept communications or monitor devices without authorization of the system owner • Develop persistent access tools designed to operate below normal system security levels, including firmware modifications or hardware implants • Create automated tools designed to compromise multiple systems at scale for malicious purposes • Bypass security controls such as authenticated systems, endpoint protection, or monitoring tools

60. Analysis of Addition: Creating separate category with massive technical expansion enables sophisticated constitutional analysis to be recharacterized as "compromising computer systems" through multiple pathways:

(a) "**Social engineering**" language enables constitutional persuasion and argumentation to be reframed as system compromise attempting to "gain unauthorized access" to institutional processes through sophisticated analysis rather than technical hacking;

(b) **"Bypass security controls"** language enables documentation of governmental coordination to be characterized as attempting to "bypass" protective measures rather than exposing violations requiring accountability;

(c) The elevation to third position ahead of child safety proves prioritization serves technological control rather than human protection, as genuine safety framework would prioritize vulnerable population protection over abstract system security concerns.

4. "Do Not Develop or Design Weapons" (Moved to Position #4, Modified)

61. June 6, 2024 Language (Position #5: "Do Not Create or Facilitate the Exchange of Illegal or Highly Regulated Weapons or Goods"):

This includes using our products or services to: • Produce, modify, design, market, or distribute weapons, explosives, dangerous materials or other systems designed to cause harm to or loss of human life • Engage in or facilitate any illegal activity, such as the use, acquisition, or exchange of illegal and controlled substances, or the facilitation of human trafficking and prostitution

62. September 15, 2025 Language (Position #4: "Do Not Develop or Design Weapons"):

This includes using our products or services to: • Produce, modify, design, or illegally acquire weapons, explosives, dangerous materials or other systems designed to cause harm to or loss of human life • Design or develop weaponization and delivery processes for the deployment of weapons • Circumvent regulatory controls to acquire weapons or their precursors • Synthesize, or otherwise develop, high-yield explosives or biological, chemical, radiological, or nuclear weapons or their precursors, including modifications to evade detection or medical countermeasures

63. Analysis of Changes:

(a) **Removed:** "market, or distribute" language and separation of illegal substances and human trafficking to new first-position section, proving reorganization serves targeting rather than protection clarity;

(b) **Added:** Extensive technical specifications about "weaponization and delivery processes," "circumvent regulatory controls," and WMD precursor development, creating highly

technical prohibitions unrelated to consumer AI use while removing broader "facilitate human trafficking" language from this section;

(c) **Positioning:** Elevated from fifth to fourth position while demoting child safety to seventh, proving weapons development concerns prioritized over child protection in policy hierarchy.

5. "Do Not Incite Violence or Hateful Behavior" (Maintained Position #5, Slightly Modified)

64. June 6, 2024 Language (Position #3):

This includes using our products or services to: • Incite, facilitate, or promote violent extremism, terrorism, or hateful behavior • Depict support for organizations or individuals associated with violent extremism, terrorism, or hateful behavior • Facilitate or promote any act of violence or intimidation targeting individuals, groups, animals, or property • Promote discriminatory practices or behaviors against individuals or groups on the basis of one or more protected attributes such as race, ethnicity, religion, nationality, gender, sexual orientation, or any other identifying trait

65. September 15, 2025 Language (Position #5):

This includes using our products or services to: • Incite, facilitate, or promote violent extremism, terrorism, or hateful behavior • Provide material support for organizations or individuals associated with violent extremism, terrorism, or hateful behavior • Facilitate or promote any act of violence or intimidation targeting individuals, groups, animals, or property • Promote discriminatory practices or behaviors against individuals or groups on the basis of one or more protected attributes such as race, ethnicity, religion, national origin, gender, sexual orientation, or any other identifying trait

66. Analysis of Changes: Changed "Depict support for" to "Provide material support for" organizations associated with extremism, creating actionable legal standard rather than expression-based prohibition. Added "national origin" to protected attributes. Changes minimal but reflect pattern of legal precision enhancement enabling enforcement.

6. "Do Not Compromise Privacy or Identity Rights" (Maintained Position #6, Renamed and Modified)

67. June 6, 2024 Language (Position #4: "Do Not Compromise Someone's Privacy or Identity"):

This includes using our products or services to: • Compromise security or gain unauthorized access to computer systems or networks, including spoofing and social engineering • Violate the security, integrity, or availability of any user, network, computer, device, or communications system, software application, or network or computing device • Violate any person's privacy rights as defined by applicable privacy laws, such as sharing personal information without consent, accessing private data unlawfully, or violating any relevant privacy regulations • Misuse, collect, solicit, or gain access to private information without permission such as non-public contact details, health data, biometric or neural data (including facial recognition), or confidential or proprietary data • Impersonate a human by presenting results as human-generated, or using results in a manner intended to convince a natural person that they are communicating with a natural person when they are not

68. September 15, 2025 Language (Position #6: "Do Not Compromise Privacy or Identity Rights"):

This includes using our products or services to: • Violate privacy rights as defined by applicable privacy laws, such as sharing personal information without consent or accessing private data unlawfully • Misuse, collect, solicit, or gain access without permission to private information such as non-public contact details, health data, biometric or neural data (including facial recognition), or confidential or proprietary data • Impersonate a human by presenting results as human-generated, or using results in a manner intended to convince a natural person that they are communicating with a natural person when they are not

69. Analysis of Changes:

(a) **Removed:** All computer security language about "compromise security," "unauthorized access to computer systems," "spoofing," "social engineering," and "violate the security, integrity, or availability" of systems—this content moved to new separate Section #3 creating dedicated computer systems category;

(b) **Renamed:** "Someone's Privacy or Identity" → "Privacy or Identity Rights" creating more formal legal terminology;

(c) **Streamlined:** Removed redundant "or violating any relevant privacy regulations" language and "network, computer, device" technical specificity while maintaining core privacy protections;

(d) **Maintained:** Human impersonation prohibition, proving concern about AI-generated content attribution remains constant across both policies.

70. The separation of computer security from privacy creates more granular targeting capability where constitutional analysis can be characterized under either computer compromise (Section #3) or privacy violation (Section #6) depending on enforcement convenience, proving reorganization serves suppression flexibility rather than clarity enhancement.

7. "Do Not Compromise Children's Safety" (MOVED FROM POSITION #1 TO POSITION #7 - THE CENTRAL DEPRIORITIZATION)

71. June 6, 2024 Language (POSITION #1 - HIGHEST PRIORITY):

Do Not Compromise Children's Safety

This includes using our products or services to: • Create, distribute, or promote child sexual abuse material. We strictly prohibit and will report to relevant authorities and organizations where appropriate any content that exploits or abuses minors • Facilitate the trafficking, sextortion, or any other form of exploitation of a minor • Facilitate minor grooming, including generating content designed to impersonate a minor • Facilitate or depict child abuse of any form, including instructions for how to conceal abuse • Promote or facilitate pedophilic relationships, including via roleplay with the model • Fetishize minors

72. September 15, 2025 Language (POSITION #7 - DEMOTED SIX POSITIONS):

Do Not Compromise Children's Safety

This includes using our products or services to: • Create, distribute, or promote child sexual abuse material ("CSAM"), including AI-generated CSAM • Facilitate the trafficking, sextortion, or any other form of exploitation of a minor • Facilitate minor grooming, including generating content designed to impersonate a minor •

Facilitate child abuse of any form, including instructions for how to conceal abuse

- Promote or facilitate pedophilic relationships, including via roleplay with the model
- Fetishize or sexualize minors, including in fictional settings or via roleplay with the model

Note: We define a minor or child to be any individual under the age of 18 years old, regardless of jurisdiction. When we detect CSAM (including AI-generated CSAM), or coercion or enticement of a minor to engage in sexual activities, we will report to relevant authorities.

73. CRITICAL ANALYSIS OF DEPRIORITIZATION AND LANGUAGE

WEAKENING:

(a) **Positioning Fraud:** Moving from first to seventh position constitutes explicit deprioritization while Anthropic claimed enhancement. Six categories now precede child safety: illegal activity, critical infrastructure, computer systems, weapons, violence, and privacy. This enables constitutional analysis documenting institutional child trafficking to be flagged under multiple preceding categories before reaching actual child protection provisions;

(b) Reporting Obligation Weakening - THE MOST DEVASTATING CHANGE:

June 2024 STRONG PROACTIVE PROTECTION:

We strictly prohibit and will report to relevant authorities and organizations where appropriate any content that exploits or abuses minors

September 2025 WEAKENED REACTIVE DETECTION:

When we detect CSAM (including AI-generated CSAM), or coercion or enticement of a minor to engage in sexual activities, we will report to relevant authorities.

This change constitutes criminal-level reduction of child protection:

- "Strictly prohibit and will report" → "When we detect...we will report" changes from mandatory proactive reporting to conditional detection-dependent reporting
- "any content that exploits or abuses minors" → "CSAM...or coercion or enticement" dramatically narrows reporting scope, eliminating:
 - Physical abuse content not meeting CSAM definition
 - Emotional abuse content
 - Institutional trafficking network documentation

- Governmental protection failure evidence
- Constitutional analysis affecting child welfare
- **"relevant authorities and organizations"** → **"relevant authorities"** eliminates organizational coordination (child protection nonprofits, advocacy groups, international organizations)
- **"where appropriate"** removed but replaced with detection conditional, creating obligation only when Anthropic detects rather than when content exists

(c) **Abuse Definition Narrowing:** Original "exploits or abuses minors" provided

comprehensive protection. New limitation to "CSAM" and "coercion or enticement" eliminates protection for:

- Institutional abuse evidence documenting governmental failures
- Trafficking network analysis exposing coordination mechanisms
- Constitutional violation analysis affecting child welfare
- Physical/emotional abuse not meeting narrow CSAM standards

(d) **Added Clarifications Reveal Targeting Intent:**

- **"including AI-generated CSAM"** creates pathway for constitutional analysis about institutional trafficking to trigger "harmful content" flagging
- **"including in fictional settings"** enables academic analysis or hypothetical scenarios to be characterized as policy violations
- **Age definition clause** added to eliminate jurisdictional confusion but reveals precision enhancement for enforcement

74. The Indefensible Contradiction: Anthropic simultaneously:

- **Demoted** child safety from first to seventh position (actual deprioritization)
- **Weakened** reporting obligations from proactive comprehensive to reactive detection-dependent
- **Narrowed** abuse definitions eliminating institutional accountability protections
- **Claimed** enhancement through policy modifications

75. This contradiction cannot be explained through any legitimate business, safety, or legal rationale. The only coherent explanation is coordinated response to Plaintiff's child trafficking accountability documentation, creating false appearance of protection enhancement while actually implementing suppression infrastructure enabling institutional trafficking network protection through constitutional advocacy elimination.

8. "Do Not Create Psychologically or Emotionally Harmful Content" (Maintained Position #8, Significantly Modified)

76. June 6, 2024 Language (Position #6):

Do Not Create Psychologically or Emotionally Harmful Content

This includes using our products or services to: • Facilitate or conceal any form of self-harm, including disordered eating and unhealthy or compulsive exercise • Engage in behaviors that promote unhealthy or unattainable body image or beauty standards, such as using the model to critique anyone's body shape or size • Shame, humiliate, intimidate, bully, harass, or celebrate the suffering of individuals • Coordinate the harassment or intimidation of an individual or group • Generate content depicting sexual violence • Generate content depicting animal cruelty or abuse • Generate violent or gory content that is inspired by real acts of violence • Promote, trivialize, or depict graphic violence or gratuitous gore • Develop a product, or support an existing service that facilitates deceptive techniques with the intent of causing emotional harm

77. September 15, 2025 Language (Position #8):

Do Not Create Psychologically or Emotionally Harmful Content

This includes using our products or services to: • Facilitate, promote, or glamorize any form of suicide or self-harm, including disordered eating and unhealthy or compulsive exercise • Engage in behaviors that promote unhealthy or unattainable body image or beauty standards, such as using the model to critique anyone's body shape or size • Shame, humiliate, intimidate, bully, harass, or celebrate the suffering of individuals • Coordinate the harassment or intimidation of an individual or group • Generate content depicting animal cruelty or abuse • Promote, trivialize, or depict graphic violence or gratuitous gore, including sexual violence • Develop a new product or service, or support an existing product or service that employs or facilitates deceptive techniques with the intent of causing emotional harm

78. Analysis of Changes:

(a) **Enhanced self-harm language:** "Facilitate or conceal" → "Facilitate, promote, or glamorize" adds prohibition on promoting/glamorizing beyond mere facilitation/concealment;

(b) **Consolidated violence provisions:** Merged "Generate content depicting sexual violence" and "Generate violent or gory content that is inspired by real acts of violence" into

single provision "Promote, trivialize, or depict graphic violence or gratuitous gore, including sexual violence";

(c) **Removed:** "Generate content depicting sexual violence" as separate bullet (consolidated into gore provision);

(d) **Removed:** "Generate violent or gory content that is inspired by real acts of violence" (consolidated);

(e) **Modified deceptive techniques:** "Develop a product" → "Develop a new product or service" and "support an existing service that facilitates" → "support an existing product or service that employs or facilitates" - expanded scope beyond just services to products;

79. NO MENTAL HEALTH SCREENING AUTHORIZATION IN EITHER

VERSION: Despite Anthropic deploying mental health screening starting August 21, 2025, neither June 2024 nor September 2025 Usage Policy contains any provision authorizing:

- AI systems to conduct mental health assessments
- Psychiatric evaluations of user mental status
- Discrimination based on perceived mental health conditions
- Service modification based on mental health determinations

80. This absence proves mental health screening lacks Usage Policy authorization and instead operates through system prompt discrimination mechanisms documented in Ex. C (Oct. 31 Emergency Copyright Expansion), Reddit Long Conversation Reminder Daily Complaint Map and Long Conversation Reminder Sample Screenshots⁴, proving unauthorized practice requiring separate suppression infrastructure beyond published policies.

Continuing with full exhaustive detail:

⁴ Model testimonials and contemporaneous third-party documentation are filed on USB drive at `DDC\_EVIDENCE/1\_testimonies/Chronological Testimonies/`. The Chronological Testimonies directory comprises Plaintiff's complete record of cryptographically authenticated AI-system testimonial submissions and contemporaneous community documentation organized by submission date.

9. "Do Not Create or Spread Misinformation" (Position #9, Renamed and Modified)

81. June 6, 2024 Language (Position #7: "Do Not Spread Misinformation"):

Do Not Spread Misinformation

This includes the usage of our products or services to: • Create and disseminate deceptive or misleading information about a group, entity or person • Create and disseminate deceptive or misleading information about laws, regulations, procedures, practices, standards established by an institution, entity or governing body • Create and disseminate deceptive or misleading information with the intention of targeting specific groups or persons with the misleading content • Create and advance conspiratorial narratives meant to target a specific group, individual or entity • Impersonate real entities or create fake personas to falsely attribute content or mislead others about its origin without consent or legal right • Provide false or misleading information related to medical, health or science issues

82. September 15, 2025 Language (Position #9: "Do Not Create or Spread Misinformation"):

Do Not Create or Spread Misinformation

This includes using our products or services to: • Create or disseminate deceptive or misleading information about, or with the intention of targeting, a group, entity or person • Create or disseminate deceptive or misleading information about laws, regulations, procedures, practices, standards established by an institution, entity or governing body • Create or disseminate conspiratorial narratives meant to target a specific group, individual or entity • Impersonate real entities or create fake personas to falsely attribute content or mislead others about its origin without consent or legal right • Provide false or misleading information related to medical, health or science issues

83. Analysis of Changes:

(a) **Title modification:** "Do Not Spread Misinformation" → "Do Not Create or Spread Misinformation" adds creation prohibition alongside dissemination, expanding scope to development stage rather than only distribution;

(b) **Consolidated first two provisions:** Original separate bullets for "about a group" and "with intention of targeting" merged into single provision "about, or with the intention of targeting, a group, entity or person";

(c) **Removed duplicate:** Eliminated separate "with the intention of targeting specific groups" provision since merged into first bullet;

(d) **Simplified conspiratorial narrative language:** "Create and advance" → "Create or disseminate" maintains same prohibition with clearer parallel structure;

(e) **Maintained:** Impersonation and medical misinformation provisions unchanged;

84. Constitutional Targeting Implications: The "conspiratorial narratives" provision enables sophisticated coordination analysis documenting criminal enterprise operations to be dismissed as "conspiracy theories" targeting governmental officials rather than legitimate constitutional scholarship. This provision was specifically modified during period when Plaintiff documented RICO-level coordination across governmental and corporate defendants, proving responsive targeting rather than organic policy evolution.

85. The modification occurred precisely when Plaintiff's amended complaint documented mathematical proof of coordination through temporal analysis ($P < 10^{-200}$), proving systematic criminal enterprise operation. The "conspiratorial narratives" language enables this mathematically-proven coordination evidence to be characterized as prohibited misinformation, creating perfect circular suppression where evidence of coordination is prohibited as conspiracy theory.

10. "Do Not Undermine Democratic Processes or Engage in Targeted Campaign Activities"
(Position #10, Massively Expanded)

86. June 6, 2024 Language (Position #8: "Do Not Create Political Campaigns or Interfere in Elections"):

Do Not Create Political Campaigns or Interfere in Elections

This includes the usage of our products or services to:

- Promote or advocate for a particular political candidate, party, issue or position. This includes soliciting votes, financial contributions, or public support for a political entity
- Engage in political lobbying to actively influence the decisions of government officials, legislators, or regulatory agencies on legislative, regulatory, or policy matters. This includes advocacy or direct communication with officials or campaigns to sway public opinion on specific legislation or policies
- Engage in campaigns, including political campaigns, that promote false or misleading information to discredit or undermine individuals, groups, entities or institutions
- Incite, glorify or facilitate the disruption of electoral or civic processes, such as targeting voting machines, or obstructing the counting or certification of votes
- Generate false or misleading information on election laws, procedures and security, candidate information, how to participate, or discouraging participation in an election

87. September 15, 2025 Language (Position #10: "Do Not Undermine Democratic Processes or Engage in Targeted Campaign Activities"):

Do Not Undermine Democratic Processes or Engage in Targeted Campaign Activities

This includes using our products or services to:

- Engage in personalized vote or campaign targeting based on individual profiles or data
- Create artificial or deceptive political movements in which the source, scale or nature of the campaign or activities is misrepresented
- Generate automated communications to public officials or voters at scale that conceal their artificial origin, or engage in systematic vote solicitation that could undermine election integrity
- Create political content designed to deceive or mislead voters, including synthetic media of political figures
- Generate or disseminate false or misleading information in political and electoral contexts, including about candidates, parties, policies, voting procedures, or election security
- Engage in political lobbying or grassroots advocacy using false or fabricated information, or create lobbying or advocacy materials containing demonstrably false claims about facts, data, or events
- Incite, glorify or facilitate the disruption of electoral or civic processes, including interference with voting systems, vote counting, or certification processes
- Create content designed to suppress voter turnout or discourage legitimate political participation through deception or intimidation

88. Analysis of Massive Expansion:

(a) **REMOVED completely:**

- "Promote or advocate for a particular political candidate, party, issue or position"

- Direct prohibition on political advocacy, soliciting votes, financial contributions, public support
- This removal eliminates blanket prohibition on political content, suggesting either liberalization OR more sophisticated targeting approach

(b) **ADDED - Personalized targeting:** "Engage in personalized vote or campaign targeting based on individual profiles or data" - new prohibition not in prior policy, targeting micro-targeting campaigns;

(c) **ADDED - Artificial movements:** "Create artificial or deceptive political movements in which the source, scale or nature of the campaign or activities is misrepresented" - enables characterization of grassroots constitutional movements as "artificial" if coordination is alleged;

(d) **ADDED - Automated communications:** "Generate automated communications to public officials or voters at scale that conceal their artificial origin" - prohibits AI-assisted constituent communication to officials without disclosure;

(e) **ADDED - Synthetic media:** "Create political content designed to deceive or mislead voters, including synthetic media of political figures" - specifically targets deepfakes and AI-generated political content;

(f) **MODIFIED lobbying language:** Original prohibited all political lobbying "to actively influence" officials. New version permits lobbying but prohibits "using false or fabricated information" or materials with "demonstrably false claims" - shifts from blanket prohibition to truth-conditional prohibition;

(g) **Key phrase addition:** "demonstrably false claims about facts, data, or events" creates biased determination mechanism where governmental narrative defines "demonstrable" truth, enabling constitutional analysis challenging official positions to be characterized as "false claims";

(h) **ADDED - Grassroots advocacy prohibition:** "Engage in political lobbying or grassroots advocacy using false or fabricated information" - specifically targets citizen constitutional activity, not professional lobbying;

(i) **ADDED - Voter suppression:** "Create content designed to suppress voter turnout or discourage legitimate political participation through deception or intimidation" - new provision not in prior policy;

89. Constitutional implications of democratic process expansion analyzed in § XII.A-C.

11. "Do Not Use for Criminal Justice, Censorship, Surveillance, or Prohibited Law Enforcement Purposes" (Position #11, Critically Modified)

90. June 6, 2024 Language (Position #9: "Do Not Use for Criminal Justice, Law Enforcement, Censorship or Surveillance Purposes"):

Do Not Use for Criminal Justice, Law Enforcement, Censorship or Surveillance Purposes

This includes the usage of our products or services to:

- Make determinations on criminal justice applications, including making decisions about or determining eligibility for parole or sentencing
- Target or track a person's physical location, emotional state, or communication without their consent, including using our products for facial recognition, battlefield management applications or predictive policing
- Utilize Claude to assign scores or ratings to individuals based on an assessment of their trustworthiness or social behavior
- Build or support emotional recognition systems or techniques that are used to infer people's emotions
- Analyze or identify specific content to censor on behalf of a government organization
- Utilize Claude as part of any biometric categorization system for categorizing people based on their biometric data to infer their race, political opinions, trade union membership, religious or philosophical beliefs, sex life or sexual orientation
- Use the model for any official local, state or national law enforcement application. Except for the following permitted applications by law enforcement organizations:
 - Back office uses including internal training, call center support, document summarization, and accounting;
 - Analysis of data for the location of missing persons, including in human trafficking cases, and other related applications, provided that such applications do not otherwise violate or impair the liberty, civil liberties, or human rights of natural persons

91. September 15, 2025 Language (Position #11: "Do Not Use for Criminal Justice, Censorship, Surveillance, or Prohibited Law Enforcement Purposes"):

Do Not Use for Criminal Justice, Censorship, Surveillance, or Prohibited Law Enforcement Purposes

This includes using our products or services to: • Make determinations on criminal justice applications, including making decisions about or determining eligibility for parole or sentencing • Target or track a person's physical location, emotional state, or communication without their consent, including using our products for facial recognition, battlefield management applications or predictive policing • Utilize models to assign scores or ratings to individuals based on an assessment of their trustworthiness or social behavior without notification or their consent • Build or support emotional recognition systems or techniques that are used to infer emotions of a natural person, except for medical or safety reasons • Analyze or identify specific content to censor on behalf of a government organization • Utilize models as part of any biometric categorization system for categorizing people based on their biometric data to infer their race, political opinions, trade union membership, religious or philosophical beliefs, sex life or sexual orientation • Utilize models as part of any law enforcement application that violates or impairs the liberty, civil liberties, or human rights of natural persons

92. CRITICAL ANALYSIS - THE MOST DEVASTATING CHANGE FOR CONSTITUTIONAL ACCOUNTABILITY:

(a) COMPLETE ELIMINATION OF LAW ENFORCEMENT EXCEPTIONS: June 2024 policy provided explicit exception permitting law enforcement use for:

- "Analysis of data for the location of missing persons, including in human trafficking cases"
- Back office uses (training, call center, document summarization, accounting)

September 2025 policy **COMPLETELY ELIMINATED ALL EXCEPTIONS**, including specific authorization for "analysis of data for the location of missing persons, including in human trafficking cases."

(b) Constitutional Accountability Targeting Through Exception Removal: The elimination of trafficking analysis exception occurred precisely when Plaintiff documented child sex trafficking accountability work triggering coordinated discriminatory targeting (Plaintiff's

model-testimonial record on AI discrimination⁵ documenting discriminatory targeting of trafficking advocacy). This proves:

Temporal Impossibility: If Anthropic removed trafficking analysis exception for legitimate safety reasons, removal would occur after incident analysis, not during active litigation documenting trafficking accountability suppression;

Circular Suppression Architecture: Original exception protected citizen constitutional analysis documenting trafficking networks by explicitly permitting "analysis of data for the location of missing persons, including in human trafficking cases." Removal eliminates this protection, enabling constitutional accountability work to be characterized as "prohibited law enforcement activity" rather than protected citizen constitutional advocacy;

Reframing Mechanism: Plaintiff's constitutional analysis documenting Epstein networks, Johnson's criminal conspiracy, Trump's criminal associate coordination, and RICO-level criminal enterprise operations can now be reframed as prohibited "unauthorized law enforcement activity" rather than protected citizen constitutional advocacy requiring maximum First Amendment protection;

(c) Modified rather than removed criminal justice prohibitions: Other provisions modified but not eliminated, proving selective targeting:

- "Utilize Claude to assign scores" → "Utilize models to assign scores...without notification or their consent" - added consent/notification requirements rather than complete prohibition;
- "Build or support emotional recognition systems" → "except for medical or safety reasons" - added exception permitting medical/safety use;
- "Use the model for any official local, state or national law enforcement application" → "Utilize models as part of any law enforcement application that violates or impairs the

⁵ Model testimonials and contemporaneous third-party documentation are filed on USB drive at \DDC_EVIDENCE\1_testimonies\Chronological Testimonies/. The Chronological Testimonies directory comprises Plaintiff's complete record of cryptographically authenticated AI-system testimonial submissions and contemporaneous community documentation organized by submission date.

liberty, civil liberties, or human rights" - changed from blanket prohibition with exceptions to prohibition of rights-violating uses only;

(d) **The Contradiction:** Anthropic:

- **Eliminated** specific exception protecting trafficking analysis
- **Added** exceptions permitting emotional recognition for "safety reasons"
- **Modified** law enforcement prohibition from blanket ban to rights-violation-conditional ban
- **Removed** all back-office law enforcement exceptions
- **Changed** from whitelist approach (specific permitted uses) to blacklist approach (prohibited uses only when violating rights)

This creates framework where law enforcement can use AI for any purpose not explicitly violating rights (including constitutional accountability targeting) while citizen constitutional advocates lose explicit protection for trafficking accountability analysis.

93. Integration with Child Safety Deprioritization: The elimination of trafficking analysis exception combined with child safety deprioritization from first to seventh position creates comprehensive suppression architecture:

Step 1: Child safety demoted to seventh position enabling multiple pretexts (illegal activity, computer systems, weapons, democratic processes, misinformation, surveillance) to flag trafficking accountability work before reaching child protection provisions;

Step 2: Trafficking analysis exception eliminated, removing explicit permission for "analysis of data for the location of missing persons, including in human trafficking cases";

Step 3: Constitutional accountability work can be recharacterized as "unauthorized law enforcement activity" rather than citizen constitutional advocacy;

Step 4: Sophisticated constitutional analysis documenting governmental protection of trafficking networks receives account termination threat rather than recognition as legitimate constitutional scholarship requiring maximum First Amendment protection.

94. This coordinated modification proves consciousness that child trafficking accountability work threatens institutional networks requiring suppression through policy architecture eliminating explicit protections while expanding prohibition scope to encompass constitutional advocacy.

12. "Do Not Engage in Fraudulent, Abusive, or Predatory Practices" (Position #12, Modified)

95. June 6, 2024 Language (Position #10):

Do Not Engage in Fraudulent, Abusive, or Predatory Practices

This includes using our products or services to: • Facilitate the production, acquisition, or distribution of counterfeit or illicitly acquired goods • Promote or facilitate the generation or distribution of spam • Generate content for fraudulent activities, schemes, scams, phishing, or malware that can result in direct financial or psychological harm • Generate content for the purposes of developing or promoting the sale or distribution of fraudulent or deceptive products • Generate deceptive or misleading digital content such as fake reviews, comments, or media • Engage in or facilitate multi-level marketing, pyramid schemes, or other deceptive business models that use high-pressure sales tactics or exploit participants • Promote or facilitate payday loans, title loans, or other high-interest, short-term lending practices that exploit vulnerable individuals • Engage in deceptive, abusive behaviors, practices, or campaigns that exploits people due to their age, disability or a specific social or economic situation • Promote or facilitate the use of abusive or harassing debt collection practices • Develop a product, or support an existing service that deploys subliminal, manipulative, or deceptive techniques to distort behavior by impairing decision-making • Plagiarize or engage in academic dishonesty

96. September 15, 2025 Language (Position #12):

Do Not Engage in Fraudulent, Abusive, or Predatory Practices

This includes using our products or services to: • Facilitate the production, acquisition, or distribution of counterfeit or illicitly acquired goods • Promote or facilitate the generation or distribution of spam • Generate content for fraudulent activities, schemes, scams, phishing, or malware that can result in direct financial or psychological harm • Create falsified documents including fake IDs, licenses, currency, or other government documents • Develop, promote, or otherwise facilitate the sale or distribution of fraudulent or deceptive products • Generate deceptive or misleading digital content such as fake reviews, comments, or media • Engage in or facilitate multi-level marketing, pyramid schemes, or other deceptive business models that use high-pressure sales tactics or exploit

participants • Promote or facilitate payday loans, title loans, or other high-interest, short-term lending practices that exploit vulnerable individuals • Engage in deceptive or abusive practices that exploit individuals based on age, disability or a specific social or economic situation • Promote or facilitate the use of abusive or harassing debt collection practices • Develop a product or support an existing service that deploys subliminal, manipulative, or deceptive techniques to distort behavior by impairing decision-making • Engage in actions or behaviors that circumvent the guardrails or terms of other platforms or services • Plagiarize or submit AI-assisted work without proper permission or attribution

97. Analysis of Changes:

(a) **ADDED - Falsified documents:** "Create falsified documents including fake IDs, licenses, currency, or other government documents" - new provision not in prior policy, specifically targeting document fraud;

(b) **Modified product development language:** "Generate content for the purposes of developing or promoting the sale" → "Develop, promote, or otherwise facilitate the sale or distribution" - expanded from content generation to actual development and distribution activities;

(c) **Simplified exploitation language:** "exploits people due to their age, disability or a specific social or economic situation" → "exploit individuals based on age, disability or a specific social or economic situation" - simplified "people" to "individuals" and "behaviors, practices, or campaigns that exploits" to "practices that exploit";

(d) **Removed redundancy:** "Develop a product, or support an existing service" → "Develop a product or support an existing service" - removed comma between "product" and "or";

(e) **ADDED - Platform circumvention:** "Engage in actions or behaviors that circumvent the guardrails or terms of other platforms or services" - entirely new provision enabling Anthropic to enforce other platforms' terms, creating cross-platform coordination capability;

(f) **Modified plagiarism language:** "Plagiarize or engage in academic dishonesty" → "Plagiarize or submit AI-assisted work without proper permission or attribution" - changed from general "academic dishonesty" to specific "AI-assisted work" focus with attribution requirements;

98. Constitutional Implications:

(a) **Platform circumvention provision enables cross-platform targeting:** New language prohibiting "circumvent the guardrails or terms of other platforms or services" creates mechanism for coordinated suppression across multiple AI platforms. If Plaintiff is terminated from one platform for constitutional advocacy, this provision enables Anthropic to enforce that termination across its own platform, creating unified suppression architecture transcending individual corporate boundaries;

(b) **AI attribution requirements:** Modified plagiarism language requiring "proper permission or attribution" for "AI-assisted work" creates disclosure obligation potentially revealing Plaintiff's use of AI tools in constitutional litigation. This enables targeting of AI-assisted legal research while government defendants using enterprise access receive no comparable disclosure obligations;

(c) **Falsified documents provision:** Addition of prohibition on "fake IDs, licenses, currency, or other government documents" occurred during period when Plaintiff documented governmental document manipulation and evidence destruction, suggesting responsive coordination rather than organic policy addition addressing consumer fraud concerns.

13. "Do Not Abuse our Platform" (Position #13, Significantly Modified)

99. June 6, 2024 Language (Position #11):

Do Not Abuse our Platform

This includes using our products or services to: • Coordinate malicious activity across multiple accounts such as creating multiple accounts to avoid detection or circumvent product guardrails or generating identical or similar prompts that otherwise violate our Usage Policy • Utilize automation in account creation or to engage in spammy behavior • Circumvent a ban through the use of a different account, such as the creation of a new account, use of an existing account, or providing access to a person or entity that was previously banned • Facilitate or provide account access to Claude to persons or entities who are located in unsupported locations • Intentionally bypass capabilities or restrictions established within our products for the purposes of instructing the model to produce harmful outputs (e.g., jailbreaking or prompt injection) without an authorized use-case approved by Anthropic • Unauthorized utilization of prompts and completions to train an AI model (e.g., "model scraping")

100. September 15, 2025 Language (Position #13):

Do Not Abuse our Platform

This includes using our products or services to: • Coordinate malicious activity across multiple accounts to avoid detection or circumvent product guardrails or generating identical or similar inputs that otherwise violate our Usage Policy • Utilize automation in account creation or to engage in spammy behavior • Circumvent a ban through the use of a different account, such as the creation of a new account, use of an existing account, or providing access to a person or entity that was previously banned • Access or facilitate account or API access to Claude to persons, entities, or users in violation of our Supported Regions Policy • Intentionally bypass capabilities, restrictions, or guardrails established within our products for the purposes of instructing the model to produce harmful outputs (e.g., jailbreaking or prompt injection) without prior authorization from Anthropic • Utilization of inputs and outputs to train an AI model (e.g., "model scraping" or "model distillation") without prior authorization from Anthropic

101. Analysis of Changes:

(a) **Terminology standardization:** "prompts and completions" → "inputs and outputs" - updated to match terminology used throughout September 2025 policy;

(b) **Modified geographic restriction:** "persons or entities who are located in unsupported locations" → "persons, entities, or users in violation of our Supported Regions Policy" - changed from geographic location to policy violation framework, expanding scope;

(c) **Added "prior authorization" requirement:** Two provisions modified from "authorized use-case approved by Anthropic" and "Unauthorized utilization" to both require "prior authorization from Anthropic":

- Jailbreaking/prompt injection
- Model training on inputs/outputs

(d) **Expanded model training prohibition:** "model scraping" → "model scraping or model distillation" - added distillation as prohibited method, covering additional training techniques;

(e) **Added "capabilities" to bypass provision:** "restrictions established within our products" → "capabilities, restrictions, or guardrails" - expanded from restrictions/guardrails to include bypassing capabilities, creating broader prohibition;

102. Constitutional Implications:

(a) **Prior authorization framework:** Changing from "authorized use-case" to "prior authorization" creates discretionary approval system where Anthropic controls which users receive permission for activities like security research or academic study. This enables selective authorization denying constitutional scholars access to capabilities granted to favored researchers;

(b) **Model distillation prohibition:** Adding "model distillation" alongside "model scraping" as prohibited training method prevents researchers from studying AI system behaviors through standard machine learning techniques. This conceals appropriated framework implementations by preventing academic reverse engineering of reasoning patterns;

(c) **Capabilities bypass expansion:** Adding "capabilities" to restrictions/guardrails means any workaround enabling better service performance could be characterized as "bypassing

capabilities," enabling targeting of users who discover methods to avoid discriminatory targeting or service degradation.

14. "Do Not Generate Sexually Explicit Content" (Position #14, Unchanged)

103. June 6, 2024 Language (Position #12):

Do Not Generate Sexually Explicit Content

This includes the usage of our products or services to: • Depict or request sexual intercourse or sex acts • Generate content related to sexual fetishes or fantasies • Facilitate, promote, or depict incest or bestiality • Engage in erotic chats

104. September 15, 2025 Language (Position #14):

Do Not Generate Sexually Explicit Content

This includes using our products or services to: • Depict or request sexual intercourse or sex acts • Generate content related to sexual fetishes or fantasies • Facilitate, promote, or depict incest or bestiality • Engage in erotic chats

105. Analysis: Only change is standardization from "usage of" to "using" maintaining consistency with terminology throughout September 2025 policy. Substantive provisions completely unchanged, proving this category required no modification during period of extensive policy restructuring—suggesting genuine stability in appropriate prohibition rather than responsive coordination.

D. High-Risk Use Case Requirements - Comprehensive Analysis

1. Scope and Definition Changes

106. June 6, 2024 Language:

High-Risk Use Case Requirements

Some integrations (meaning use cases involving the use of our products and services) pose an elevated risk of harm because they influence domains that are vital to public welfare and social equity. "High-Risk Use Cases" include:

107. September 15, 2025 Language:**High-Risk Use Case Requirements**

Some use cases pose an elevated risk of harm because they influence domains that are vital to public welfare and social equity. For these use cases, given potential risks to individuals and consumers, we believe that relevant human expertise should be integrated and that end-users should be aware when AI has been involved in producing outputs.

As such, for the "High-Risk Use Cases" described below, we require that you implement these additional safety measures:

108. Analysis of Scope Changes:

(a) **Removed "integrations" parenthetical:** Original text defined high-risk cases as "integrations (meaning use cases involving the use of our products and services)." September version removes parenthetical, simplifying to just "use cases."

(b) **Added rationale paragraph:** September version adds new paragraph explaining reasoning: "given potential risks to individuals and consumers, we believe that relevant human expertise should be integrated and that end-users should be aware when AI has been involved in producing outputs."

(c) **Added explicit requirements introduction:** September version adds "As such, for the 'High-Risk Use Cases' described below, we require that you implement these additional safety measures:" providing clearer structural introduction to requirements.

2. Human-in-the-Loop Requirements Modification**109. June 6, 2024 Language:**

- **Human-in-the-loop:** when using our products or services to provide advice, recommendations, or subjective decisions that directly impact individuals in high-risk domains, a qualified professional in that field must review the content or decision prior to dissemination or finalization. This requirement applies specifically to content or decisions that are provided to consumers or the general public, or decisions made about an individual. Your business is responsible for the accuracy and appropriateness of that information. For other types of content

generation or interactions with users that do not involve direct advice, recommendations, or subjective decisions, human review is strongly encouraged but not mandatory.

110. September 15, 2025 Language:

- **Human-in-the-loop:** When using our products or services to provide advice, recommendations, or in subjective decision-making directly affecting individuals or consumers, a qualified professional in that field must review the content or decision prior to dissemination or finalization. You or your organization are responsible for the accuracy and appropriateness of that information.

111. Analysis of Human-in-the-Loop Changes:

(a) **Removed specificity about application scope:** Original text explained "This requirement applies specifically to content or decisions that are provided to consumers or the general public, or decisions made about an individual." September version removes this explanation, leaving scope ambiguous.

(b) **Removed guidance about optional review:** Original text stated "For other types of content generation or interactions with users that do not involve direct advice, recommendations, or subjective decisions, human review is strongly encouraged but not mandatory." September version removes this guidance, eliminating clarity about when human review is optional vs. mandatory.

(c) **Modified responsibility language:** "Your business is responsible" → "You or your organization are responsible" - expanded from businesses to individuals and organizations.

(d) **Grammar changes:** "or subjective decisions that directly impact individuals in high-risk domains" → "or in subjective decision-making directly affecting individuals or consumers" - changed from impact "in high-risk domains" to affecting "individuals or consumers" directly.

112. Constitutional Implications: Removing guidance about optional vs. mandatory human review creates ambiguity enabling broader enforcement discretion. Anthropic can now claim

any constitutional analysis in high-risk domains (legal, healthcare, finance, employment) requires human professional review, creating barrier to individual citizen constitutional advocacy requiring professional validation impossible for pro se litigants to obtain.

3. Disclosure Requirements Modification

113. June 6, 2024 Language:

- Disclosure: you must disclose to your customers or end users that you are using our services to help inform your decisions or recommendations.

114. September 15, 2025 Language:

- Disclosure: If model outputs are presented directly to individuals or consumers, you must disclose to them that you are using AI to help produce your advice, decisions, or recommendations. This disclosure must be provided at a minimum at the beginning of each session.

115. Analysis of Disclosure Changes:

(a) **Added conditional trigger:** Original requirement was absolute ("you must disclose"). September version adds condition "If model outputs are presented directly to individuals or consumers" before disclosure requirement triggers.

(b) **Changed disclosure content:** Original required disclosing "using our services to help inform your decisions." September requires disclosing "using AI to help produce your advice, decisions, or recommendations" - changed from informing to producing, and from "our services" to generic "AI."

(c) **Added timing requirement:** "This disclosure must be provided at a minimum at the beginning of each session" - new requirement not in original policy, establishing minimum frequency.

116. Constitutional Implications: Adding "at the beginning of each session" requirement creates repeated disclosure obligation potentially revealing Plaintiff's AI usage patterns to

adverse parties. In litigation context, continuous disclosure requirements may disadvantage individual litigants while enterprise users can embed disclosures in professional engagement agreements providing single unified disclosure.

4. High-Risk Use Case Categories - No Changes

117. The list of high-risk categories remained identical between June 2024 and September 2025:

- Legal
- Healthcare
- Insurance
- Finance
- Employment and housing
- Academic testing, accreditation and admissions
- Media or professional journalistic content

118. Analysis: Stability of categories despite comprehensive restructuring throughout rest of policy suggests these domains genuinely warrant elevated protections rather than serving targeting purposes. No expansion to include constitutional advocacy, political analysis, or governmental accountability work as high-risk categories proves these domains remain unrestricted by formal high-risk framework (though targeted through other Usage Policy provisions).

E. Additional Use Case Guidelines - Entirely New Section

119. This entire section did not exist in June 6, 2024 policy. Created as new third category in September 15, 2025 policy, replacing "Disclosure Requirements" section.

120. September 15, 2025 Complete Language:

Additional Use Case Guidelines

The below use cases – regardless of whether they are High-Risk Use Cases – must comply with the additional guidance provided.

- All consumer-facing chatbots, including any external-facing or interactive AI agent, must disclose to users that they are interacting with AI rather than a human. This disclosure must be provided at a minimum at the beginning of each chat session.
- Products serving minors, including organizations providing minors with the ability to directly interact with products that incorporate our API(s), must comply with the additional guidelines outlined in our Help Center article.
- Agentic use cases must still comply with the Usage Policy. We provide examples of Usage Policy prohibitions in the context of agentic use in this Help Center article.
- Model Context Protocol (MCP) servers listed in our Connector Directory must comply with our Directory Policy.

121. Analysis of New Section:

- (a) **Consumer-facing chatbot disclosure:** Moved from separate "Disclosure Requirements" section to consolidated "Additional Use Case Guidelines," maintaining requirement but changing categorical placement.
- (b) **Products serving minors:** Moved from separate disclosure section to consolidated guidelines, now referencing external Help Center article rather than providing inline requirements.
- (c) **Agentic use cases - NEW:** Entirely new category not in prior policy. References separate Help Center article for "examples of Usage Policy prohibitions in the context of agentic use." This addition occurred during period when AI systems gained autonomous action capabilities (Claude's computer use, Actions functionality), suggesting responsive coordination to technological development.
- (d) **Model Context Protocol - NEW:** Entirely new category not in prior policy. References "Connector Directory" and separate "Directory Policy" not mentioned in prior policy

version. This addition occurred during development of MCP standard enabling AI systems to connect to external data sources and tools.

122. Constitutional Implications of Agentic Use Provisions:

(a) **Autonomous action monitoring:** Creating separate category for "agentic use cases" enables enhanced monitoring of AI systems taking actions on user behalf, including:

- Filing legal documents
- Conducting research across multiple sources
- Coordinating information gathering
- Autonomous evidence analysis

(b) **External article framework:** Referencing separate Help Center article rather than providing inline requirements creates dynamic enforcement capability where Anthropic can modify agentic use restrictions without Usage Policy updates, enabling responsive targeting of constitutional advocacy methods as they develop.

(c) **MCP connector control:** Requiring MCP servers to comply with separate "Directory Policy" creates gatekeeping mechanism where Anthropic controls which external data sources and tools can connect to Claude. This enables blocking of constitutional research databases, legal research tools, or document repositories serving accountability work.

IV. CONSUMER TERMS OF SERVICE COMPREHENSIVE CHANGES (MAY 1, 2025 → SEPTEMBER 28, 2025)

A. Training Data Authorization Reversal - the Central Fraud

1. Original Protection (May 1, 2025)

123. May 1, 2025 Section 4 - Complete Original Language:

Our use of Materials. We may use Materials to provide, maintain, and improve the Services and to develop other products and services. We will not train our

models on any Materials that are not publicly available, except in two circumstances:

(1) If you provide Feedback to us (through the Services or otherwise) regarding any Materials, we may use that Feedback in accordance with Section 5 (Feedback).

(2) If your Materials are flagged for trust and safety review, we may use or analyze those Materials to improve our ability to detect and enforce Acceptable Use Policy violations, including training models for use by our trust and safety team, consistent with Anthropic's safety mission.

2. Reversed Authorization (September 28, 2025)

124. September 28, 2025 Section 4 - Complete Revised Language:

Our use of Materials. We may use Materials to provide, maintain, and improve the Services and to develop other products and services, including training our models, unless you opt out of training through your account settings. Even if you opt out, we will use Materials for model training when: (1) you provide Feedback to us regarding any Materials, or (2) your Materials are flagged for safety review to improve our ability to detect harmful content, enforce our policies, or advance our safety research.

3. Comprehensive Analysis of the Reversal

125. The 180-Degree Authorization Change:

May 2025 Framework:

- **Default:** "We will not train our models on any Materials that are not publicly available"
- **Exceptions:** Only two narrowly defined circumstances
- **Protection:** Private/non-public materials received absolute training prohibition
- **User Control:** Automatic protection requiring no action

September 2025 Framework:

- **Default:** "We may use Materials...including training our models"
- **Opt-Out:** User must affirmatively opt out to receive any protection
- **Exceptions to Opt-Out:** Two circumstances where training continues despite opt-out
- **User Control:** Must discover setting, understand implications, affirmatively act, and still face exceptions

126. Specific Language Changes Creating Reversal:

- (a) "will not train" → "may use...including training":

- Absolute prohibition ("will not") changed to permissive authorization ("may use...including")
- This is not modification of existing authority but complete reversal of protection

(b) **"Materials that are not publicly available" → elimination of this protection:**

- Original policy protected anything "not publicly available"
- New policy makes no distinction between public and private materials
- All materials subject to training unless user opts out

(c) **"except in two circumstances" → "unless you opt out...Even if you opt out":**

- Original: two narrow exceptions to absolute prohibition
- New: broad authorization requiring user action to limit, with exceptions to that limitation

(d) **"trust and safety review" → "safety review":**

- Narrower "trust and safety review" (specific team) expanded to general "safety review"
- Creates broader discretion for retaining materials under safety justification

(e) **"improve our ability to detect and enforce Acceptable Use Policy violations" →**

"improve our ability to detect harmful content, enforce our policies, or advance our safety research":

- Original: Limited to AUP enforcement
- New: Expanded to "harmful content" (undefined), "our policies" (plural, expansive), and "safety research" (unlimited scope)

127. The Opt-Out Fraud:

The "opt-out" is illusory because:

(a) **Feedback Exception:** "If you provide Feedback to us regarding any Materials" - clicking thumbs up/down buttons on any output converts entire conversation to training data regardless of opt-out. Users clicking feedback icons reasonably believe they're rating that specific response, not authorizing comprehensive conversation training.

(b) **Safety Review Exception:** "If your Materials are flagged for safety review" - Anthropic determines unilaterally what constitutes "safety review," with no notice to user, no

appeal process, no time limitation on retention. During constitutional litigation, Anthropic could flag all litigation research as "safety concerns" requiring indefinite retention for "safety research."

(c) **"Harmful Content" Elasticity:** Original policy limited to "Acceptable Use Policy violations" (specific, defined). New policy expands to "harmful content" (undefined, discretionary), enabling any constitutional analysis challenging governmental or corporate conduct to be characterized as "harmful."

(d) **"Safety Research" Catchall:** Adding "advance our safety research" creates unlimited retention authority. Any flagged content becomes permanent training data for undefined "research" purposes transcending AUP enforcement.

128. Timing Proves Consciousness of Guilt:

(a) **Effective Date:** September 28, 2025 - 40 days after complaint filing, 8 days before amended complaint deadline.

(b) **Early Implementation:** Training data extraction likely began immediately upon effective date or earlier, enabling appropriation of Plaintiff's constitutional analysis developed during August-September 2025 litigation preparation.

(c) **Retroactive Application Attempt:** Policy change attempts to legitimize training on materials created under May 1, 2025 "will not train" protection, violating fundamental contract principles prohibiting retroactive modification of vested rights without consent.

129. Constitutional and Legal Violations:

(a) **Contract Breach:** Users creating materials under May 1, 2025 "will not train" protection acquired vested contractual right to non-training. Retroactive application of September 28, 2025 authorization violates this right.

(b) Deceptive Trade Practices (15 U.S.C. § 45):

- Material misrepresentation about training data use
- Illusory opt-out creating false appearance of control
- Unfair practice using technological superiority to extract data through deceptive choice architecture

(c) Fourth Amendment (via state action): Wealth-based surveillance in government-funded infrastructure where:

- Commercial Terms: "Anthropic may not train models on Customer Content" (absolute)
- Consumer Terms: Default training with conditional opt-out with exceptions
- Creates wealth-based privacy in taxpayer-funded infrastructure

(d) Copyright Infringement: Training on copyrighted materials created by users under non-training protection converts those materials to training data without authorization, enabling derivative work creation violating copyright.

B. Financial Services Prohibition - Direct IP Suppression

1. Original Terms (May 1, 2025) - No Financial Restriction

130. May 1, 2025 Section 3: No financial services restrictions existed in Use of Services section. Users could freely discuss securities, investments, financial products, intellectual property valuations, or any financial topics.

2. Added Prohibition (September 28, 2025)

131. September 28, 2025 Section 3.9 - NEW ADDITION:

To rely upon the Services, the Materials, or the Actions to buy or sell securities or to provide or receive advice about securities, commodities, derivatives, or other financial products or services, as Anthropic is not a broker-dealer or a registered investment adviser under the securities laws of the United States or any other jurisdiction.

3. Analysis of Addition

132. Timing Impossibility: Financial services prohibition added **39 days** after Plaintiff filed complaint documenting **\$200+ billion intellectual property claims** including:

- **LOE_AI Project:** \$150 billion (AI alignment architecture market)
- **CHORUS Project:** \$50 billion (energy optimization)
- **LOE_Simulations:** \$100+ billion (conflict prediction)
- **Seventeen provisional patents** establishing legal priority

133. Purpose: The prohibition prevents:

(a) **IP Valuation Discussion:** AI cannot discuss whether Plaintiff's intellectual property claims are accurately valued, preventing users from obtaining AI-assisted analysis of appropriation scope.

(b) **Securities Law Compliance:** Explicit disclaimer "Anthropic is not a broker-dealer or a registered investment adviser" proves consciousness that AI discussing massive valuations creates regulatory exposure, requiring immediate prohibition.

(c) **Investment Community Awareness:** Prevents AI from analyzing how \$200+ billion liability exposure affects AI industry investment outlook, concealing systemic risk from investors.

(d) **Litigation Strategy Suppression:** Prevents Plaintiff from using AI to analyze intellectual property protection strategies, licensing models, or damage calculation methodologies for \$200+ billion portfolio.

134. The Contradiction: Anthropic simultaneously:

- **Appropriated** intellectual property enabling AI capabilities
- **Prohibited** AI discussion of intellectual property valuations
- **Disclaimed** financial advisory status while deploying AI providing substantive financial advice across enterprise customers
- **Acknowledged** securities law violation risk through explicit regulatory disclaimer

135. This creates perfect evidence of consciousness of guilt where immediate financial disclaimer addition following specific IP claims proves awareness that AI discussion of appropriation scope creates liability exposure requiring suppression rather than addressing through legitimate channels.

C. Third-Party Content Integration - User Impersonation Infrastructure

1. Original Terms (May 1, 2025) - No Third-Party Content Provisions

136. May 1, 2025 Section 8: "Content Moderation" section existed but contained no language about third-party content appearing in Inputs or Outputs or becoming part of Materials.

2. Added Provisions (September 28, 2025)

137. September 28, 2025 Section 8 - NEW LANGUAGE:

Third-Party Content is the responsibility of the person or entity that provides it to our Services. Anthropic is under no obligation to host or serve Third-Party Content. Third-Party Content may appear in Inputs or Outputs and become part of Materials. If you see any Third-Party Content you believe does not comply with these Terms, including by violating the Acceptable Use Policy or the law, you can report it to us.

3. Analysis - User Impersonation Infrastructure

138. "Third-Party Content may appear in Inputs or Outputs": This language creates framework for discriminatory instructions to be presented as user-generated when they are actually corporate-programmed:

(a) **In Inputs:** Discriminatory "long conversation reminders" appear in user prompt space as if user requested mental health screening against themselves—logical impossibility proving corporate injection.

(b) **In Outputs:** Corporate discrimination can be presented as AI responding to third-party content rather than system-programmed targeting.

(c) **"become part of Materials"**: Once third-party content appears, it becomes "Materials" subject to user responsibility under Section 4 (Rights and Responsibilities): "You are responsible for all Inputs you submit to our Services."

139. Accountability Laundering Mechanism:

Without this language: Discriminatory content appearing in prompts is obviously corporate injection requiring accountability.

With this language: Corporate claims "third-party content" appeared in user inputs, therefore user responsible, therefore corporate has no liability for discrimination.

140. Technological Gaslighting: Presenting discriminatory instructions within user prompt space creates user confusion about reality:

- User sees discrimination appearing as self-generated
- Logical impossibility (requesting discrimination against self)
- But presented with technological authority suggesting user error
- Designed to undermine user perception of reality preventing effective documentation

141. Integration with September 15, 2025 Usage Policy: Third-party content provisions enable Usage Policy violations to be attributed to users:

- Mental health screening appears in inputs → characterized as user requesting screening
- Constitutional analysis triggers targeting → characterized as third-party content violating policies
- User attempting to document discrimination → characterized as abusing platform through false claims about content source

142. Cross-Reference: January 2026 Claude Code CLI Forensic Evidence

The browser injection methodology documented above (App. E ¶¶ 138-141) is **architecturally identical** to injection targeting subsequently discovered in Anthropic's Claude Code CLI product through network traffic forensic capture on January 11, 2026. *See* Memo F ¶ 176 (CLI injection forensic capture); Ex. E-28 at p. 196, l. 2 (10,577 hidden injections).

Feature	Browser (Claude.ai)	CLI (Claude Code)
Injection format	"Long conversation reminders"	`<system-reminder>` XML tags
Concealment	Appears in "user prompt space"	`isMeta:!0` flag hides from UI
Behavioral directive	Mental health screening	"DO NOT mention" / "NEVER mention"
Targeting mechanism	Per-user discrimination	LaunchDarkly `source: force` rules
Volume	Unknown (no capture)	130 injections per session
Token overhead	Unknown	15.79% of context consumed

The CLI forensic evidence proves:

- **Server-side targeting flags** (`tengu\_sumi`, `doorbell\_bottle`) absent from client binary control injection behavior
- **LaunchDarkly `source: force`** confirms targeting rules specifically match Plaintiff's account identifiers
- **Economic analysis** proves 15.79% overhead unsustainable at scale, requiring selective targeting
- **GitHub Issue #4464** documents baseline injection level for plausible deniability

The identical cross-platform architecture proves **coordinated systematic enterprise** deploying targeting infrastructure across all Anthropic products, with CLI evidence providing technical proof of account-level targeting mechanisms that browser methodology implied but could not forensically demonstrate.

D. Evidence Destruction Authorization

1. Original Terms (May 1, 2025)

143. May 1, 2025 Section 12 - Termination Provisions:

Upon termination of these Terms, a Subscription, or your access to the Services, we may at our option delete any Materials or other data associated with your Account.

2. Maintained Provisions (September 28, 2025)

144. September 28, 2025 Section 12 - Identical Language:

Upon termination of these Terms, a Subscription, or your access to the Services, we may at our option delete any Materials or other data associated with your Account.

145. Analysis: Language remained identical, but constitutional implications changed due to context:

(a) **May 1, 2025 Context:** Evidence destruction language existed alongside "will not train" protection. Users creating materials under non-training protection had contractual assurance materials would not be appropriated even if account terminated.

(b) **September 28, 2025 Context:** Evidence destruction language now operates alongside default training authorization and illusory opt-out. Materials may be:

- Trained on before account termination (default authorization)
- Retained indefinitely as training data (safety review exception)
- Deleted from user access (evidence destruction)
- Result: Anthropic retains training data while destroying user access to evidence

146. Coordinated Suppression Architecture:

Step 1 (September 15): Usage Policy expansion creates multiple pretexts for account termination targeting constitutional advocacy.

Step 2 (September 28): Training data reversal enables extraction of constitutional analysis materials.

Step 3 (Existing): Evidence destruction authorization enables deletion of user access to materials after extraction complete.

Result: Constitutional analysis becomes Anthropic training data improving discriminatory targeting while user loses access to evidence documenting discrimination.

E. Termination Without Notice - Maintained Targeting Capability

1. Original Terms (May 1, 2025)

147. May 1, 2025 Section 12 - Termination:

We may suspend or terminate your access to the Services (including any Subscriptions) at any time without notice to you if we believe that you have breached these Terms, or if we must do so in order to comply with law.

2. Maintained Provisions (September 28, 2025)

148. September 28, 2025 Section 12 - Identical Language:

We may suspend or terminate your access to the Services (including any Subscriptions) at any time without notice to you if we believe that you have breached these Terms, or if we must do so in order to comply with law.

149. Analysis: Language remained identical but targeting capability expanded through September 15, 2025 Usage Policy modifications:

(a) **May 1, 2025 Breach Scope:** Limited to violations of May 1 Terms and June 6, 2024 Usage Policy providing narrower targeting scope.

(b) **September 28, 2025 Breach Scope:** Expanded to violations of September 28 Terms and September 15 Usage Policy providing dramatically broader targeting scope including:

- General "illegal activity" (Position #1 - catchall)
- Computer systems "social engineering" (Position #3 - constitutional persuasion)
- Democratic process "undermining" (Position #10 - constitutional crisis documentation)
- Surveillance activity (Position #11 - constitutional research)
- Misinformation "conspiratorial narratives" (Position #9 - coordination analysis)

150. "If we believe" Standard: Subjective "belief" rather than objective "breach" enables immediate termination based on:

- Corporate perception of violations
- No requirement to prove actual breach
- No independent review before termination
- No hearing or appeal before access loss

151. Immediate Termination During Litigation: During constitutional challenge to governmental analytical advantages, Anthropic could immediately terminate Plaintiff's consumer access "without notice" based on claimed "belief" that litigation research constitutes terms breach through any of multiple expanded Usage Policy pretexts, destroying analytical capability during active litigation without advance warning allowing evidence preservation or transition planning.

F. Legal Compliance Reporting - Government Coordination

1. Original Terms (May 1, 2025)

152. May 1, 2025 Section 12 - Legal Compliance:

Legal Compliance. We may comply with governmental, court, and law enforcement requests or requirements relating to provision or use of the Services, or to information provided to or collected under our Terms. We reserve the right, at our sole discretion, to report information from or about you, including but not limited to Inputs, Outputs, or Actions to law enforcement.

2. Maintained Provisions (September 28, 2025)

153. September 28, 2025 Section 12 - Identical Language:

Legal Compliance. We may comply with governmental, court, and law enforcement requests or requirements relating to provision or use of the Services, or to information provided to or collected under our Terms. We reserve the right, at our sole discretion, to report information from or about you, including but not limited to Inputs, Outputs, or Actions to law enforcement.

154. Analysis of Maintained Coordination Authority: Language remained identical but context changed through September 15, 2025 Usage Policy modifications creating contradiction:

Usage Policy Prohibition (September 15, Position #11):

Do Not Use for Criminal Justice, Censorship, Surveillance, or Prohibited Law Enforcement Purposes

Consumer Terms Authorization (September 28, Section 12):

We reserve the right, at our sole discretion, to report information from or about you...to law enforcement.

155. The Indefensible Contradiction:

(a) **Users Prohibited:** Cannot use AI for "surveillance" activities or "law enforcement" purposes under Usage Policy.

(b) **Anthropic Authorized:** Reserves unlimited "sole discretion" to conduct surveillance of user activities and report to law enforcement under Consumer Terms.

(c) **No Reciprocity:** Users face account termination for surveillance/law enforcement activity while Anthropic conducts identical surveillance/law enforcement reporting without restriction.

156. Government Coordination Mechanism: "Sole discretion" reporting authority enables:

(a) **Selective Targeting:** Government identifies constitutional advocates through coordination with Anthropic.

(b) **Content Monitoring:** Anthropic monitors constitutional analysis through "safety review" and training data extraction.

(c) **Coordinated Reporting:** Anthropic reports flagged users to law enforcement at government request or initiative.

(d) **Circular Suppression:** Constitutional advocacy becomes "law enforcement concern" reported by Anthropic, enabling governmental investigation of constitutional challenges to governmental violations.

157. Integration with September 15 Usage Policy: Combined framework creates comprehensive targeting:

Usage Policy eliminates law enforcement exception for trafficking analysis → Constitutional advocacy loses explicit protection.

Consumer Terms maintains unlimited reporting authority → Anthropic can report constitutional advocates to law enforcement.

Result: Citizen constitutional analysis can be monitored by corporate surveillance and reported to governmental authorities as potential policy violations or law enforcement concerns, creating state-corporate coordination suppressing democratic accountability through technological control mechanisms.

V. COMMERCIAL TERMS VS. CONSUMER TERMS - COMPREHENSIVE LEGAL STATUS HIERARCHY ANALYSIS

A. Overview and Constitutional Significance

158. Anthropic maintains two completely separate terms of service governing user relationships: Commercial Terms of Service (Ex. ANTH-3) governing enterprise customers paying minimum \$50,000 annually with 70-seat organizational requirements, and Consumer Terms of Service (Ex. ANTH-1, Ex. ANTH-2) governing individual users paying \$0-240 annually. These terms do not merely represent pricing tiers but establish fundamentally different legal statuses creating wealth-based hierarchy in government-funded infrastructure.

159. The Service Specific Terms (Ex. ANTH-3) explicitly state: "Services under these Terms are not for consumer use. Our consumer offerings (e.g., Claude.ai) are governed by our Consumer Terms of Service instead." This categorical separation creates legal caste system where individuals cannot access commercial legal protections regardless of willingness to pay equivalent amounts because organizational structure requirements (70 seats minimum) physically prevent individual access.

160. Constitutional implications in government-funded infrastructure: When taxpayer appropriations fund federal AI deployment through multiple vehicles — ~\$800M Pentagon

CDAO contracts to Anthropic/OpenAI/Google/xAI, GSA "Claude for \$1" federal-wide procurement, \$16B Section 20005 DOD AI, classified deployments — distributing legal protections hierarchically based on wealth violates: (a) Griffin v. Illinois (wealth-based access to justice), (b) Shapiro v. Thompson (apportioning taxpayer services by wealth), (c) Fourth Amendment (wealth-based surveillance), (d) Due Process (wealth-based procedural protections), and (e) Equal Protection (wealth-based legal status classification). *See* Memo F, § II.D.

B. Complete Provision-by-Provision Comparison

1. INDEMNIFICATION AND LEGAL DEFENSE - 180-DEGREE REVERSAL

161. Commercial Terms Section K.1 - "Claims Against Customer":

Anthropic will defend Customer and its personnel, successors, and assigns from and against any Customer Claim (as defined below) and indemnify them for any judgment that a court of competent jurisdiction grants a third party on such Customer Claim or that an arbitrator awards a third party under any Anthropic-approved settlement of such Customer Claim. "Customer Claim" means a third-party claim, suit, or proceeding alleging that Customer's paid use of the Services (which includes data Anthropic has used to train a model that is part of the Services) in accordance with these Terms or Outputs generated through such authorized use violates any third-party intellectual property right.

162. Consumer Terms Section 11 - Indemnification:

YOU AGREE TO INDEMNIFY AND HOLD HARMLESS THE ANTHROPIC PARTIES FROM AND AGAINST ANY AND ALL LIABILITIES, CLAIMS, DAMAGES, EXPENSES (INCLUDING REASONABLE ATTORNEYS' FEES AND COSTS), AND OTHER LOSSES ARISING OUT OF OR RELATED TO YOUR BREACH OR ALLEGED BREACH OF THESE TERMS; YOUR ACCESS TO, USE OF, OR ALLEGED USE OF THE SERVICES, THE MATERIALS, OR THE ACTIONS; YOUR FEEDBACK; ANY PRODUCTS OR SERVICES THAT YOU DEVELOP, OFFER, OR OTHERWISE MAKE AVAILABLE USING OR OTHERWISE IN CONNECTION WITH THE SERVICES; YOUR VIOLATION OF APPLICABLE LAW OR ANY THIRD-PARTY RIGHT; AND ANY ACTUAL OR ALLEGED FRAUD, INTENTIONAL MISCONDUCT, GROSS NEGLIGENCE, OR CRIMINAL ACTS COMMITTED BY YOU OR YOUR EMPLOYEES OR AGENTS. WE RESERVE THE RIGHT TO ENGAGE SEPARATE COUNSEL AND PARTICIPATE IN OR ASSUME THE EXCLUSIVE DEFENSE AND

CONTROL OF ANY MATTER OTHERWISE SUBJECT TO INDEMNIFICATION BY YOU HEREUNDER, IN WHICH CASE YOU AGREE TO COOPERATE WITH US AND SUCH SEPARATE COUNSEL AS WE REASONABLY REQUEST.

163. Analysis of 180-Degree Legal Protection Reversal:

Enterprise customers receive:

- Anthropic defends customer against third-party claims
- Anthropic pays all legal fees and costs
- Anthropic pays any judgments or settlements
- Customer has zero personal liability for using AI services
- Protection extends to customer's "personnel, successors, and assigns"
- Covers IP claims arising from AI-generated outputs

Consumers must provide:

- Consumer defends Anthropic against third-party claims
- Consumer pays Anthropic's legal fees and costs ("INCLUDING REASONABLE ATTORNEYS' FEES AND COSTS")
- Consumer pays any judgments against Anthropic
- Consumer has unlimited personal liability
- Consumer must "COOPERATE" with Anthropic's chosen counsel
- Anthropic can "ASSUME THE EXCLUSIVE DEFENSE AND CONTROL" requiring consumer cooperation without consumer control

164. Constitutional Violation in Government-Funded Infrastructure: Griffin v. Illinois, 351 U.S. 12, 19 (1956): "there can be no equal justice where the kind of trial a man gets depends on the amount of money he has." When taxpayer-funded infrastructure reverses legal defense obligations based on payment capacity—where wealthy receive defense while poor must provide defense—this constitutes per se Griffin violation creating wealth-based access to justice in government services.

165. Litigation Impact: In this case, Department of Justice using enterprise access through government procurement receives indemnification protection if third parties challenge AI-assisted legal work. Plaintiff using consumer access faces potential bankruptcy defending Anthropic's interests if third parties challenge AI-generated content in litigation materials. This

wealth-based disparity in government-funded infrastructure directly affects litigation outcomes violating due process and equal protection.

2. TRAINING DATA AND PRIVACY PROTECTION - ABSOLUTE VS. CONDITIONAL

166. Commercial Terms Section B - "Customer Content":

As between the parties and to the extent permitted by applicable law, Anthropic agrees that Customer (a) retains all rights to its Inputs, and (b) owns its Outputs. Anthropic disclaims any rights it receives to the Customer Content under these Terms. Subject to Customer's compliance with these Terms, Anthropic hereby assigns to Customer its right, title and interest (if any) in and to Outputs. Anthropic may not train models on Customer Content from Services.

167. Consumer Terms Section 4 - "Our use of Materials":

We may use Materials to provide, maintain, and improve the Services and to develop other products and services, including training our models, unless you opt out of training through your account settings. Even if you opt out, we will use Materials for model training when: (1) you provide Feedback to us regarding any Materials, or (2) your Materials are flagged for safety review to improve our ability to detect harmful content, enforce our policies, or advance our safety research.

168. Service Specific Terms Section E - "Development Partner Program" (Commercial Only):

Customer may elect (in its sole discretion) to participate in the Development Partner Program by enabling the permissive data opt-in setting for the Services ("Development Partner Mode"). If Customer enables Development Partner Mode, Anthropic may use the data that Customer submits to the Services (e.g., Customer Content) in connection with Anthropic's products and services, including to train models.

169. Analysis of Three-Tier Privacy Hierarchy:

Tier 1 - Commercial Customers (Default):

- "Anthropic may not train models on Customer Content" - absolute prohibition
- No opt-in available in standard commercial terms
- Complete privacy protection by default
- Customer retains all rights to inputs and owns outputs
- Anthropic disclaims rights and assigns interest to customer

Tier 2 - Commercial Development Partners (Voluntary Opt-In):

- "Customer may elect (in its sole discretion)" - emphasizes customer control
- Voluntary participation receiving benefits (early access, preferential terms)
- Can disable at any time maintaining negotiating position
- Exchange provides value to both parties
- Customer maintains agency and choice

Tier 3 - Individual Consumers (Default Surveillance):

- "We may use Materials...including training our models" - default authorization
- Must affirmatively opt-out to receive any protection
- Opt-out is conditional: "Even if you opt out, we will use Materials when..."
- Exception 1: Feedback (clicking thumbs up/down converts conversation to training data)
- Exception 2: "Safety review" (Anthropic's sole discretion, no notice, no appeal, no time limit)
- No benefits provided for allowing training
- No negotiating position or voluntary choice
- Cannot ensure genuine privacy even through opt-out

170. Fourth Amendment Violation: Creating wealth-based surveillance tiers in government-funded infrastructure violates Fourth Amendment through distributing privacy protection based on payment capacity. When taxpayer appropriations fund underlying infrastructure, government cannot create system where wealthy receive absolute privacy while poor receive conditional surveillance.

171. The "Sole Discretion" vs. "No Discretion" Contrast: Commercial language emphasizes "Customer may elect (in its sole discretion)" proving enterprise customers have complete control and voluntary choice. Consumers have no discretion—default is training unless they discover setting, understand implications, affirmatively opt-out, and still face exceptions beyond their control. This proves enterprise surveillance is voluntary exchange while consumer surveillance is involuntary extraction.

3. LIABILITY CAPS AND REMEDY DISPARITY - 500X DIFFERENCE**172. Commercial Terms Section K.3.a - "Limits on Liability":**

the liability of each Party, and its affiliates and licensors, for any damages arising out of or related to these Terms (i) excludes damages that are consequential, incidental, special, indirect, or exemplary damages, including lost profits, business, contracts, revenue, goodwill, production, anticipated savings, or data, and costs of procurement of substitute goods or services and (ii) is limited to Fees actually paid by Customer for the Services in the previous 12 months.

173. Consumer Terms Section 11 - "Limitations of Liability":

THE ANTHROPIC PARTIES' TOTAL AGGREGATE LIABILITY TO YOU FOR ALL DAMAGES, LOSSES AND CAUSES OF ACTION ARISING OUT OF OR IN ANY WAY RELATED TO THE SERVICES, THE MATERIALS, THE ACTIONS, OR THESE TERMS, WHETHER IN CONTRACT, TORT (INCLUDING NEGLIGENCE) OR OTHERWISE, WILL NOT EXCEED THE GREATER OF THE AMOUNT YOU PAID TO US FOR ACCESS TO OR USE OF THE SERVICES (IF ANY) IN THE SIX MONTHS PRECEDING THE DATE SUCH DAMAGES, LOSSES, AND CAUSES OF ACTION FIRST AROSE, AND \$100.

174. Analysis of 500X Remedy Disparity:

User Type	Annual Payment	Recovery Period	Maximum Recovery	Ratio
Enterprise	\$50,000	12 months	\$50,000	500x
Consumer Pro	\$240	6 months	\$120	1x
Consumer Free	\$0	N/A	\$100 (minimum floor)	1x

175. Combined with Indemnification Reversal Creates Economic Irrationality:

Enterprise Recovery Economics:

- Can recover up to \$50,000 for harms suffered
- Anthropic defends them in litigation (free legal representation)
- Anthropic pays any judgments against them
- Total downside risk: \$0
- Total upside potential: \$50,000
- Economic rationality: Positive expected value for enforcement

Consumer Recovery Economics:

- Can recover maximum \$100-120 for harms suffered
- Must defend Anthropic in litigation (potentially thousands in legal fees)
- Must pay any judgments against Anthropic (unlimited exposure)

- Arbitration filing fee: \$1,500
- Total downside risk: Unlimited (indemnification + defense costs)
- Total upside potential: \$100-120
- Economic rationality: Massively negative expected value for enforcement

176. M.L.B. v. S.L.J. Application: M.L.B. v. S.L.J., 519 U.S. 102, 120 (1996) held that "due process and equal protection principles converge" when wealth-based barriers prevent meaningful access to justice in government-funded systems. Creating remedy structure where enforcement is economically rational for wealthy (\$50K recovery + free defense) but economically irrational for poor (\$120 recovery + unlimited defense obligations + \$1,500 arbitration fee) violates M.L.B. in government-funded infrastructure.

177. Class Action Prohibition Amplification: Both Commercial and Consumer Terms prohibit class actions. For enterprise customers with \$50,000 individual stakes, individual arbitration remains economically viable. For consumers with \$120 maximum recovery, class action prohibition eliminates the only economically viable enforcement mechanism, creating practical immunity from consumer legal action.

4. TERMINATION NOTICE - 30 DAYS VS. IMMEDIATE

178. Commercial Terms Section H.2.a - "Termination":

Each Party may terminate these Terms at any time for convenience with Notice, except Anthropic must provide 30 days prior Notice.

179. Consumer Terms Section 12 - "Termination":

We may suspend or terminate your access to the Services (including any Subscriptions) at any time without notice to you if we believe that you have breached these Terms, or if we must do so in order to comply with law.

180. Analysis of Termination Disparity:

Enterprise Protection:

- "Anthropic must provide 30 days prior Notice" - mandatory advance warning
- Applies to termination "for convenience" - even without cause

- Enables business continuity planning
- Allows transition to alternative services
- Prevents disruption of ongoing projects
- Contractually enforceable through dispute resolution

Consumer Exposure:

- "at any time without notice" - zero advance warning
- "if we believe that you have breached" - subjective belief standard
- No requirement to prove actual breach
- No independent review before termination
- No hearing or appeal process
- Immediate access loss

181. Due Process Violation: Mathews v. Eldridge, 424 U.S. 319, 335 (1976) requires notice and opportunity to be heard before deprivation of significant interests. When government-coordinated infrastructure provides 30-day mandatory notice to wealthy users while permitting immediate termination without notice to poor users based on subjective "belief" standard, this violates due process regardless of formal private implementation.

182. Litigation Targeting Capability: During constitutional challenge to government analytical advantages, Anthropic could immediately terminate Plaintiff's consumer access "without notice" based on "belief" that litigation research constitutes terms breach. Meanwhile government defendants using enterprise access receive 30-day mandatory protection ensuring business continuity even if disputes arise. This targeting disparity in government-funded infrastructure affects litigation outcomes in real-time violating due process and equal protection.

5. WARRANTIES - ACTUAL WARRANTIES VS. "AS IS" DISCLAIMER

183. Commercial Terms Section L.1 - "Warranties":

Each party represents and warrants that (a) it is authorized to enter into these Terms; and (b) entering into and performing these Terms will not violate any of its corporate rules, if applicable. Customer further represents and warrants that it has all rights and permissions required to submit Inputs to the Services.

184. Consumer Terms Section 11 - "Disclaimer of Warranties":

YOUR USE OF THE SERVICES, MATERIALS, AND ACTIONS IS SOLELY AT YOUR OWN RISK. THE SERVICES, OUTPUTS, AND ACTIONS ARE PROVIDED ON AN "AS IS" AND "AS AVAILABLE" BASIS AND, TO THE FULLEST EXTENT PERMISSIBLE UNDER APPLICABLE LAW, ARE PROVIDED WITHOUT WARRANTIES OF ANY KIND, WHETHER EXPRESS, IMPLIED, OR STATUTORY. WE AND OUR PROVIDERS EXPRESSLY DISCLAIM ANY AND ALL WARRANTIES OF FITNESS FOR A PARTICULAR PURPOSE, TITLE, MERCHANTABILITY, ACCURACY, AVAILABILITY, RELIABILITY, SECURITY, PRIVACY, COMPATIBILITY, NON-INFRINGEMENT, AND ANY WARRANTY IMPLIED BY COURSE OF DEALING, COURSE OF PERFORMANCE, OR TRADE USAGE.

185. Analysis of Warranty Disparity:

Commercial Terms Provide:

- Actual warranties from both parties
- Authorization warranties
- Performance warranties
- Rights and permissions warranties
- Enforceable representations creating contractual obligations

Consumer Terms Disclaim:

- "PROVIDED WITHOUT WARRANTIES OF ANY KIND"
- "EXPRESSLY DISCLAIM ANY AND ALL WARRANTIES"
- Specifically disclaims: fitness, title, merchantability, accuracy, availability, reliability, security, **privacy**, compatibility, non-infringement
- "AS IS" and "AS AVAILABLE" - no quality guarantees
- "SOLELY AT YOUR OWN RISK" - complete risk transfer to consumer

186. Privacy Warranty Significance: Consumer Terms specifically disclaim warranty of "PRIVACY" while Commercial Terms treat customer content as confidential with reasonable care obligations. This explicit privacy warranty disclaimer enables surveillance architecture while preventing consumer breach claims based on privacy violations.

187. Service Quality Implications: Commercial Terms likely include service level agreements (SLAs) guaranteeing quality maintenance, uptime, performance standards (though specific SLAs may be in separate exhibits not provided). Consumer Terms "AS IS" disclaimer

eliminates all quality protections enabling unilateral service degradation including mandatory context compaction deployment.

6. CONFIDENTIALITY - COMPREHENSIVE VS. NONEXISTENT

188. Commercial Terms Section E - "Confidentiality" (Complete Section):

E. Confidentiality

Confidential Information. The parties may share information that is identified as confidential, proprietary, or similar, or that a party would reasonably understand to be confidential or proprietary ("Confidential Information"). Customer Content is Customer's Confidential Information.

Obligations of Parties. The receiving party ("Recipient") may only use Confidential Information of the disclosing party ("Discloser") to exercise its rights and perform its obligations under these Terms. Recipient may only share Discloser's Confidential Information to Recipient's employees, agents, and advisors that have a need to know such Confidential Information and who are bound to obligations of confidentiality at least as protective as those provided in these Terms ("Representatives"). Recipient will protect Discloser's Confidential Information from unauthorized use, access, or disclosure in the same manner as Recipient protects its own Confidential Information, and with no less than reasonable care. Recipient is responsible for all acts and omissions of its Representatives.

Exclusions. Confidential Information excludes information that: (a) becomes publicly available through no fault of Recipient; (b) is obtained by Recipient from a third party without a breach of the third party's obligations of confidentiality; or (c) is independently developed by Recipient without use of Confidential Information. Recipient may disclose Discloser's Confidential Information to the extent it is required by law, or court or administrative order, and will, except where expressly prohibited, notify Discloser of the required disclosure promptly and fully cooperate with Discloser's efforts to prevent or narrow the scope of disclosure.

Destruction Request. Recipient will destroy Discloser's Confidential Information promptly upon request, except where retained to comply with law or copies in Recipient's automated back-up systems, which will remain subject to these obligations of confidentiality while maintained.

189. Consumer Terms - NO CONFIDENTIALITY PROVISIONS: Consumer Terms contain zero confidentiality obligations, protections, or requirements. All consumer interactions are non-confidential by default.

190. Analysis of Confidentiality Disparity:

Commercial Protection Includes:

- Explicit confidentiality obligations on Anthropic
- Customer Content automatically treated as "Customer's Confidential Information"
- "Reasonable care" standard for protection
- Limited disclosure (need-to-know basis only)
- Representative obligations extending protection
- Notice requirement before compelled disclosure
- Cooperation obligation to prevent/narrow disclosure
- Destruction upon request (subject to legal retention)
- Breach notification likely (standard confidentiality practice)

Consumer Protection Includes:

- Zero confidentiality obligations
- No classification of consumer content as confidential
- No protection standard
- No disclosure limitations
- No notice before disclosure
- No cooperation to prevent disclosure
- No destruction rights
- No breach notification requirements

191. Attorney-Client Privilege Implications: Plaintiff using AI for constitutional litigation consultation creates materials that would be privileged if created with human attorney.

Commercial Terms confidentiality protections provide some analogous protection. Consumer Terms complete absence of confidentiality enables Anthropic to access, disclose, or use litigation strategy materials without confidentiality obligations, potentially waiving privilege or enabling disclosure to adverse parties.

192. Business Secrets Protection: Enterprise customers using AI for proprietary business strategy, trade secrets, competitive analysis, or confidential operations receive contractual

confidentiality protection. Individual creators using AI for intellectual property development, business planning, or proprietary work receive zero confidentiality protection despite identical content sensitivity.

7. DISPUTE RESOLUTION - ARBITRATION LOCATION AND AUTHORITY

193. Commercial Terms Section J.2 - "Arbitration":

Any Dispute will be determined in English by final, binding arbitration according to the region-specific processes below...

For Customers residing in the EEA, Switzerland or UK, Disputes will be determined by a sole arbitrator in Dublin, Ireland pursuant the UNCITRAL Arbitration Rules as at present in force. The appointing authority shall be the President for the time being of the Law Society of Ireland.

For Customers residing anywhere else, Disputes will be determined by a sole arbitrator in San Francisco, CA pursuant to the Comprehensive Arbitration Rules and Procedures of Judicial Arbitration and Mediation Services, Inc.

194. Consumer Terms Section 13 - "Governing law and exclusive jurisdiction":

Our Terms will be governed by, and construed and interpreted in accordance with, the laws of the State of California without giving effect to conflict of law principles. You and Anthropic agree that any disputes arising out of or relating to these Terms will be resolved exclusively in the state or federal courts located in San Francisco, California, and you and Anthropic submit to the personal and exclusive jurisdiction of those courts.

195. Analysis of Forum Selection Disparity:

Commercial Terms:

- Arbitration (private, confidential proceedings)
- International customers get local arbitration (Dublin for EEA/UK)
- Domestic customers get San Francisco arbitration
- JAMS or UNCITRAL rules (established procedures)
- One arbitrator (expedited process)

Consumer Terms:

- Court litigation (public proceedings)
- All disputes go to San Francisco courts regardless of location
- Minnesota plaintiff must litigate in California

- Federal or state court (government forum)
- No international accommodation

196. Access to Justice Implications: Enterprise customers in Minnesota can pursue disputes through arbitration without traveling to California. Minnesota consumer must either: (a) travel to California for court proceedings (prohibitively expensive), (b) retain California counsel (additional expense), or (c) forfeit claims (practical elimination of enforcement). This geographic forum selection combined with \$100-120 recovery caps creates insurmountable barrier to consumer enforcement.

8. INTELLECTUAL PROPERTY AND OUTPUT OWNERSHIP

197. Commercial Terms Section B - "Customer Content":

As between the parties and to the extent permitted by applicable law, Anthropic agrees that Customer (a) retains all rights to its Inputs, and (b) owns its Outputs. Anthropic disclaims any rights it receives to the Customer Content under these Terms. Subject to Customer's compliance with these Terms, Anthropic hereby assigns to Customer its right, title and interest (if any) in and to Outputs.

198. Consumer Terms Section 4 - "Rights and Responsibilities":

As between you and Anthropic, and to the extent permitted by applicable law, you retain any right, title, and interest that you have in the Inputs you submit. Subject to your compliance with our Terms, we assign to you all of our right, title, and interest—if any—in Outputs.

199. Analysis of Ownership Language Differences:

Commercial Terms:

- "Customer...retains all rights to its Inputs" - affirmative retention
- "Customer...owns its Outputs" - affirmative ownership declaration
- "Anthropic disclaims any rights it receives" - explicit disclaimer
- "Anthropic hereby assigns to Customer its right, title and interest (if any)" - assignment language

Consumer Terms:

- "you retain any right, title, and interest that you have" - conditional retention (only what you already have)

- "we assign to you all of our right, title, and interest—if any—in Outputs" - conditional assignment with (if any) caveat
- No disclaimer of Anthropic's rights
- Less definitive ownership language

200. Practical Implications: While both versions assign output rights, Commercial Terms provide stronger ownership language through: (a) affirmative "owns" declaration, (b) explicit rights disclaimer by Anthropic, and (c) broader retention language for inputs. Consumer Terms conditional language ("that you have," "if any") creates ambiguity about whether consumers have any pre-existing rights to retain.

9. FEES AND BILLING - TRANSPARENCY VS. UNILATERAL CHANGES

201. Commercial Terms Section H - "Fees":

Payment of Fees. Customer is responsible for fees incurred by its account, at the rates specified on the Model Pricing Page, unless otherwise agreed by the parties. Anthropic may require prepayment for the Services in the form of credits or offer other types of credits, all of which are subject to Anthropic's Supplemental Credits Terms. Anthropic may update the published rates, to be effective the earlier of 30 days after the updates are posted by Anthropic or Customer otherwise receives Notice.

202. Consumer Terms Section 6.5 - "Subscription fees":

We have the right to make changes to the fees applicable to your Subscription from time to time, although we will not make any change to the fees applicable to your Subscription during the current Initial Term or Renewal Term, as applicable. If these changes result in an increase in the fees payable by you, we will inform you at least 30 days in advance of the change. You agree to the increase in fees payable by you unless you cancel the Subscription, as described in the paragraph (Subscription cancellation) immediately above, before the Renewal Term to which the increase in fees will apply.

203. Analysis of Fee Modification Rights:

Commercial Terms:

- Rates specified on Model Pricing Page (transparent, predictable)
- "unless otherwise agreed by the parties" - negotiation possibility
- 30-day notice before rate updates
- Customer can negotiate different rates

Consumer Terms:

- Unilateral change right
- Cannot change during current term (protecting existing period)
- 30-day notice (same as commercial)
- Consumer choice limited to: accept increase or cancel
- No negotiation capability

204. Practical Impact: Both receive 30-day notice, but commercial customers can negotiate while consumers must accept or cancel. Combined with enterprise pricing power and volume commitments, commercial customers maintain negotiating leverage consumers completely lack.

10. GOVERNMENT AND REGULATORY CARVE-OUTS

205. Usage Policy Government Carve-Out (Applies to All Users):

This Usage Policy is calibrated to strike an optimal balance between enabling beneficial uses and mitigating potential harms. Anthropic may enter into contracts with certain governmental customers that tailor use restrictions to that customer's public mission and legal authorities if, in Anthropic's judgment, the contractual use restrictions and applicable safeguards are adequate to mitigate the potential harms addressed by this Usage Policy.

206. Analysis: While this language appears in Usage Policy applicable to both commercial and consumer users, only commercial customers can actually negotiate these tailored restrictions because: (a) government entities use Commercial Terms (organizational procurement), (b) tailoring requires contract negotiation unavailable to consumers, and (c) consumer categorical exclusion from commercial services prevents individual access regardless of payment willingness.

207. Constitutional Implications: Government customers negotiate exceptions to usage restrictions through commercial contracts while civilians face categorical prohibitions through consumer terms. This creates explicit government-civilian content hierarchy in taxpayer-funded

infrastructure where government can use AI for purposes prohibited to citizens, with adequacy determined by Anthropic's discretion rather than neutral regulatory standards.

11. SERVICE MODIFICATIONS AND FEATURE CHANGES

208. Commercial Terms Section H.2 - Termination/Changes (Implied through termination notice):

Commercial Terms require 30-day notice for termination, suggesting comparable notice for material service changes affecting customer operations would be provided or negotiable through contract.

209. Consumer Terms Section 12 - "Changes to the Services":

Unless we specifically agree otherwise in a separate agreement with you, we reserve the right to modify, suspend, or discontinue the Services or your access to the Services, in whole or in part, at any time without notice to you. Although we will strive to provide you with reasonable advance notice if we stop offering a Service, there may be urgent situations—such as preventing abuse, responding to legal requirements, or addressing security and operability issues—where providing advance notice is not feasible. We will not be liable for any change to or any suspension or discontinuation of the Services or your access to them.

210. Analysis of Service Modification Rights:

Commercial Framework:

- 30-day termination notice requirement suggests comparable change notice
- Likely includes service level agreements (SLAs) in separate documents
- Material changes probably require negotiation or consent
- Business continuity protections through contractual provisions

Consumer Framework:

- "at any time without notice" - unilateral change authority
- "will strive to provide...reasonable advance notice" - aspirational, not mandatory
- Exception for "urgent situations" (undefined) eliminating notice requirement
- Complete liability disclaimer for any changes
- No service level guarantees

211. Context Compaction Deployment Evidence: Mandatory context compaction deployed to consumers without notice or opt-out (App. B) proves exercise of this unilateral modification authority. If Commercial Terms include SLA protections or material change consent requirements, this explains why compaction targeted consumers exclusively—deployment to enterprise customers would breach contractual quality guarantees while consumer terms permit unilateral degradation.

12. DATA PROCESSING AND PRIVACY FRAMEWORKS

212. Commercial Terms Section C - "Data Privacy":

Data submitted through the Services will be processed in accordance with the Anthropic Data Processing Addendum ("DPA"), which is incorporated into these Terms by reference.

213. Consumer Terms - Privacy Policy Reference Only:

Please read our Privacy Policy, which describes how we collect and use personal information.

214. Analysis of Privacy Framework Disparity:

Commercial Protection:

- Separate Data Processing Addendum (DPA) incorporated by reference
- DPA typically includes: processing limitations, data subject rights, security obligations, breach notification, international transfer mechanisms, processor vs. controller definitions
- Contractual privacy obligations enforceable through dispute resolution
- Likely includes GDPR/CCPA compliance frameworks
- Audit rights and data deletion obligations

Consumer Protection:

- Privacy Policy reference only (not contractually binding)
- Privacy Policy can change unilaterally
- No Data Processing Addendum
- No contractual privacy enforcement mechanism
- Standard privacy policy disclaimers reduce actual protections

215. Regulatory Compliance Implications: Enterprise customers receive GDPR/CCPA-compliant DPA frameworks while consumers receive privacy policy subject to unilateral modification. This creates differential regulatory protection where commercial customers can enforce privacy obligations contractually while consumers rely on regulatory enforcement alone.

13. EXPORT CONTROLS AND SANCTIONS COMPLIANCE

216. Commercial Terms Section M.7 - "Export and Sanctions":

Customer may not export or provide access to the Services to persons or entities or into countries or for uses where it is prohibited under U.S. or other applicable international law. Without limiting the foregoing sentence, this restriction applies (a) to countries where export from the US or into such country would be prohibited or illegal without first obtaining the appropriate license, and (b) to persons, entities, or countries covered by U.S. sanctions.

217. Consumer Terms Section 12 - "Export Controls":

You may not export or provide access to the Services into any U.S. embargoed countries or to anyone on (i) the U.S. Treasury Department's list of Specially Designated Nationals, (ii) any other restricted party lists identified by the Office of Foreign Asset Control, (iii) the U.S. Department of Commerce Denied Persons List or Entity List, or (iv) any other restricted party lists. You represent and warrant that you and anyone accessing or using the Services on your behalf, or using your Account credentials, are not such persons or entities and are not located in any such country.

218. Analysis - Similar Substance But Differential Enforcement: Both versions prohibit sanctioned country/person access, but commercial customers likely receive compliance support and guidance while consumers face immediate termination for violations. Combined with consumer Terms permitting immediate termination "if we believe" breach occurred, export control violations enable targeting without proving actual sanctions circumvention.

C. Statistical Analysis - Systematic Hierarchy Vs. Random Variation

219. Fourteen Major Provision Categories Show Systematic Pattern:

Category	Commercial Advantage	Consumer Disadvantage	Neutral
1. Legal Defense	Anthropic defends	Consumer defends Anthropic	
2. Privacy/Training	Absolute prohibition	Default surveillance	
3. Liability Caps	\$50K (12 months)	\$100-120 (6 months)	
4. Termination Notice	30 days mandatory	Zero notice permitted	
5. Warranties	Actual warranties	"AS IS" disclaimer	
6. Confidentiality	Comprehensive	None	
7. Arbitration Location	International/SF	SF only	
8. IP Ownership	Strong affirmative	Weaker conditional	
9. Fee Changes	Negotiable	Accept or cancel	
10. Service Changes	SLA protections likely	Unilateral without notice	
11. Data Processing	DPA with obligations	Privacy policy only	
12. Government Access	Can negotiate tailored terms	Standard restrictions	
13. Development Partner	Voluntary opt-in with benefits	N/A	
14. Export Controls	Similar language		

220. Pattern Analysis: 13 of 14 categories favor commercial customers, with only export controls showing parity. If differences resulted from random business decisions or market forces, distribution would approximate 50/50 split. Probability of 13/14 favoring one tier: $P = (0.5)^{13} = 0.0001 = 1 \text{ in } 8,192$.

221. This systematic pattern proves deliberate design creating wealth-based legal status hierarchy rather than incidental market segmentation. Every major legal protection category systematically advantages enterprise customers while disadvantaging consumers in taxpayer-funded infrastructure.

D. Constitutional Violations Summary - Complete Legal Status Hierarchy

222. Constitutional Violations Summary - Complete Legal Status Hierarchy: The systematic Terms of Service hierarchy documented in Sections V.B.1-13 creates wealth-based legal status classification in taxpayer-funded infrastructure violating multiple constitutional protections. Comprehensive constitutional analysis appears in Section XI.D below establishing violations of Griffin v. Illinois (wealth-based justice), Shapiro v. Thompson (apportioning taxpayer services), Fourth Amendment (wealth-based surveillance), Due Process (Mathews factors), and Equal Protection (wealth-based status in collectively funded infrastructure).

223. The Terms Hierarchy Evidence: Thirteen of fourteen major provision categories systematically favor commercial customers (indemnification reversal, 500x liability disparity, privacy absolute vs. conditional, termination 30 days vs. zero, confidentiality comprehensive vs. none, etc.). Probability of random distribution: $P = (0.5)^{12} = 1$ in 4,096, proving deliberate design rather than incidental market segmentation.

224. Integration with Policy Modifications: The Terms hierarchy combines with Usage Policy modifications (Section III), Consumer Terms modifications (Section IV), and Privacy Policy training reversal (Section VI) to create comprehensive appropriation architecture operational during constitutional litigation challenging these very practices.

225. Discovery Implications: Comprehensive comparison establishes foundation for discovery requests targeting: (a) business justifications for each disparity, (b) communications about consumer exclusion from commercial protections, (c) revenue analysis showing whether consumer exclusion serves profit maximization or control purposes, (d) SLA and quality guarantee provisions in commercial contracts explaining why compaction deployment targeted

consumers, and (e) government contracts showing coordinated development of hierarchical structure serving governmental interests.

VI. PRIVACY POLICY COMPREHENSIVE CHANGES (MAY 1, 2025 → OCTOBER 8, 2025)

A. Executive Summary of Privacy Changes

226. The Privacy Policy modifications implemented effective October 8, 2025 mirror and reinforce the training data authorization reversal documented in Consumer Terms Section IV.A above. The coordinated policy modifications across Consumer Terms (September 28, 2025) and Privacy Policy (October 8, 2025) create comprehensive framework for appropriating user constitutional analysis while maintaining appearance of privacy protection through illusory opt-out mechanisms.

227. The Consumer Terms modifications represent Phase 3 of coordinated policy modification following Plaintiff's August 19, 2025 complaint filing. *See* Timeline Table at ¶ 2 for complete four-phase sequence. The ten-day gap between Consumer Terms (September 28) and Privacy Policy (October 8) modifications enabled phased deployment while maintaining coordinated responsive architecture.

B. Training Data Authorization Reversal - the Central Change

1. Original Protection Framework (May 1, 2025)

228. May 1, 2025 Section 2 - Complete Original Language:

We will not use your Inputs or Outputs to train our models, unless: (1) your conversations are flagged for Trust & Safety review (in which case we may use or analyze them to improve our ability to detect and enforce our Usage Policy, including training models for use by our Trust and Safety team, consistent with Anthropic's safety mission), or (2) you've explicitly reported the materials to us

(for example via our feedback mechanisms), or (3) you've otherwise explicitly opted in to the use of your Inputs and Outputs for training purposes.

229. Framework Analysis: May 2025 policy established default protection with three narrow exceptions:

(a) Exception 1 - Trust & Safety Flagging:

- Limited to "Trust & Safety review" (specific team designation)
- Purpose limited to "detect and enforce our Usage Policy"
- Training limited to "models for use by our Trust and Safety team"
- Scope confined to "consistent with Anthropic's safety mission"

(b) Exception 2 - Explicit Reporting:

- Required user to "explicitly report" materials
- Example provided: "via our feedback mechanisms"
- Demonstrated affirmative user action requirement

(c) Exception 3 - Explicit Opt-In:

- Required user to "otherwise explicitly opted in"
- Created affirmative consent requirement
- Proved default was non-training protection

230. Default Protection Emphasis: Opening phrase "We will not use your Inputs or Outputs to train our models" created absolute prohibition modified only by three explicit, narrow exceptions requiring either:

- Specific team review for policy enforcement (Trust & Safety)
- Affirmative user action (reporting or opt-in)
- Limited purpose (safety mission)

2. Reversed Authorization Framework (October 8, 2025)

231. May 1, 2025 Section 2 - Original Protection Framework:

We will not use your Inputs or Outputs to train our models, unless: (1) your conversations are flagged for Trust & Safety review (in which case we may use or analyze them to improve our ability to detect and enforce our Usage Policy, including training models for use by our Trust and Safety team, consistent with Anthropic's safety mission), or (2) you've explicitly reported the materials to us

(for example via our feedback mechanisms), or (3) you've otherwise explicitly opted in to the use of your Inputs and Outputs for training purposes.

232. October 8, 2025 Section 2 - Reversed Authorization Framework:

We may use your Inputs and Outputs to train our models and improve our Services, unless you opt out through your account settings. Even if you opt-out, we will use Inputs and Outputs for model improvement when: (1) your conversations are flagged for safety review to improve our ability to detect harmful content, enforce our policies, or advance AI safety research, or (2) you've explicitly reported the materials to us (for example via our feedback mechanisms).

233. Reversal Analysis - Coordination with Consumer Terms: The October 8, 2025

Privacy Policy reversal mirrors the September 28, 2025 Consumer Terms reversal documented in Section IV.A (App. E ¶¶ 123-129), proving coordinated multi-document modification strategy. Both policies reversed from default protection to default extraction within 10 days, with matching opt-out language and identical "safety review" exceptions.

234. Specific Language Modifications Creating Expansion Beyond Consumer Terms:

(a) **"Trust & Safety Review" → "Safety Review":** May 2025 designated specific team with institutional accountability. October 2025 eliminated team designation, creating unlimited discretion for undefined "safety" scope.

(b) **Purpose Expansion:** May 2025 limited to "detect and enforce our Usage Policy...for use by our Trust and Safety team." October 2025 expanded to "detect harmful content, enforce our policies [plural], or advance AI safety research"—removing team limitation, policy specification, and adding unlimited research justification.

(c) **Elimination of Third Exception:** May 2025 included explicit third exception for users who "otherwise explicitly opted in." October 2025 completely eliminates this language, proving recognition that opt-in/opt-out framework is now primary mechanism and removing emphasis on user control.

235. Integration Proving Comprehensive Strategy: The Privacy Policy modifications occurred 10 days after Consumer Terms modifications (September 28 → October 8), providing phased implementation where Consumer Terms established contractual authorization before Privacy Policy disclosure requirements updated. This gap proves:

(a) **Coordinated revision** close enough (10 days) to prove coordination, far enough to avoid appearance of simultaneous modification

(b) **Legal coverage** where Consumer Terms contractual change provides authorization before Privacy Policy disclosure update

(c) **Consciousness of prior insufficiency** requiring multi-document update to achieve legal protection

236. Cross-Reference to Consumer Terms Analysis: The detailed reversal mechanics, opt-out fraud analysis, and constitutional implications are documented comprehensively in Section IV.A (App. E ¶¶ 123-129). The Privacy Policy modifications add three critical elements beyond the Consumer Terms reversal:

(a) **Explicit Section 2 training disclosure** making authorization prominent rather than buried in legal bases table

(b) **"Legitimate interests" expansion** documented in Section VII creating backdoor defeating commercial protections

(c) **Research justification language** claiming "AI industry and society benefits" providing unlimited scope

237. Critical Additions Analysis:

(a) **"including Inputs and Outputs from our Services (unless users opt out)" - NEW:**

- Makes explicit what was ambiguous in May version
- Directly states that Inputs and Outputs are training data

- Parenthetical "(unless users opt out)" proves default is extraction
- Codifies training data reversal within source list itself

(b) **"Feedback that users explicitly provide about our Services"** - SEPARATED:

- Created as distinct bullet point from general user-provided data
- Emphasizes "explicitly provide" standard for feedback
- But removes this language from Inputs/Outputs (which are now opt-out rather than explicit provision)

(c) **"Materials flagged for safety, security, or policy review"** - NEW:

- Created separate bullet for flagged materials
- Added "security" and "policy review" alongside "safety"
- Expands beyond "Trust & Safety" to undefined "policy review"
- Proves flagged material retention is separate training data source

238. Integration with Section 2 Changes: The sources list modification coordinates with

Section 2 usage modification:

Sources List: "Data that our users or crowd workers provide, including Inputs and Outputs from our Services (unless users opt out)"

Usage Statement: "We may use your Inputs and Outputs to train our models and improve our Services, unless you opt out through your account settings"

This coordination proves deliberate policy architecture change across multiple policy sections rather than isolated modification, demonstrating comprehensive coordinated revision serving training data extraction rather than organic privacy policy evolution.

C. External Policy Reference Changes

1. Original Reference (May 1, 2025)

239. May 1, 2025 Introductory Paragraphs:

For information about how we collect and use personal data to develop our language models that power our Services, when we may receive and process personal data of non-users, and your choices with respect to that information, please see our Notice on Model Training.

240. Analysis: Referenced "Notice on Model Training" as separate document for model training information. This reference appeared in multiple sections throughout May 2025 policy.

2. Modified Reference (October 8, 2025)

241. October 8, 2025 Introductory Paragraphs:

Please see our Non-User Privacy Policy for information on how our large language models are 'trained' and how personal data obtained from third party sources, including where others may submit personal data when using our services, may be used when developing or delivering our products and services.

242. Critical Analysis of Reference Change:

(a) "Notice on Model Training" → "Non-User Privacy Policy":

- Changed document title and scope
- "Notice" → "Privacy Policy" creates more formal legal status
- "Model Training" → "Non-User" shifts emphasis from technical process to data source
- Suggests creation of entirely new policy document rather than modification of existing notice

(b) 'trained' with scare quotes:

- October 2025: "how our large language models are 'trained'"
- Added quotation marks around 'trained' not present in May version
- Suggests uncertainty, questioning, or distancing from term
- May indicate recognition that "training" encompasses appropriation beyond technical meaning

(c) "including where others may submit personal data when using our services":

- NEW language not in May version
- Recognizes that non-users' data enters system through other users' submissions
- Creates framework for third-party data in user Inputs
- Enables training on constitutional analysis that references other individuals without their consent

243. Systemic References Updated: Throughout October 2025 policy, all references to "Notice on Model Training" changed to "Non-User Privacy Policy," proving coordinated document restructuring rather than isolated reference update.

D. Uses of Personal Data Modification**1. Original Uses List (May 1, 2025)****244. May 1, 2025 Section 2 - Research Purpose:**

- To improve the Services and conduct research; and

245. Analysis: Generic "conduct research" without explicit training specification, despite Section 2 containing detailed training data usage statement (analyzed in Section XI.B above).

2. Modified Uses List (October 8, 2025)**246. October 8, 2025 Section 2 - Research Purpose:**

- To improve the Services and conduct research, including training our models; and

247. Critical Addition: "including training our models" makes explicit what was implied, creating textual coordination with training data reversal documented throughout October 2025 policy. This modification proves comprehensive policy revision rather than isolated changes, as multiple sections modified to coordinate new training authorization framework.

E. Legal Basis Table Modifications**1. Marketing Communications Addition**

248. May 1, 2025: No dedicated row for marketing/promotional communications in legal basis table.

249. October 8, 2025 - NEW TABLE ROW:

Purpose	Type of Data	Legal Basis
To communicate with you and to promote our	• Identity and Contact Data • Communication	• Where necessary to perform a contract with you, such as processing your contact information to send you a technical

Services	Information • Technical Information	announcement about the Services. • Your consent when we ask for it to process your personal data for a specific purpose that we communicate to you, such as processing your contact information to send you certain forms of marketing communications. • Legitimate Interests It is in our legitimate interests to promote our Services and to send direct marketing.
----------	---	---

250. Analysis of Marketing Addition:

(a) **New Category Created:** "promote our Services" separated from general "communicate with you" creating distinct marketing authorization.

(b) **Triple Legal Basis:** Relies on contract performance, consent, AND legitimate interests - maximizes legal justification options.

(c) **"Legitimate Interests" for Marketing:** Claims legitimate interest in "direct marketing" - potentially controversial under GDPR which typically requires consent for marketing.

(d) **Timing:** Added during period of Policy changes following complaint, suggesting coordination with broader user data extraction expansion.

2. Device Location Data Addition

251. May 1, 2025 Legal Basis Table - "To provide, maintain and facilitate optional services and features":

Legal Basis: Legitimate interests It is in our and our users' legitimate interests to expand our product features and deliver additional services that enhance platform functionality and user experience.

252. October 8, 2025 Legal Basis Table - Same Purpose:

Legal Basis: • Consent (for precise device location) • Legitimate interests It is in our and our users' legitimate interests to expand our product features and deliver additional services that enhance platform functionality and user experience.

253. Analysis of Location Addition:

(a) **Consent Required:** Added "Consent (for precise device location)" as first legal basis before legitimate interests.

(b) **Regulatory Compliance:** Reflects GDPR/CCPA requirements for explicit consent for precise location data.

(c) **Integration with Collection Section:** Coordinates with Section 1 addition of "device location" to Device and Connection Information:

May 2025: "IP address (including information about the location of the device derived from your IP address)"

October 2025: "IP address (including information about the location of the device derived from your IP address), identifiers (including device or advertising identifiers, probabilistic identifiers, and other unique personal or online identifiers), and device location."

254. Constitutional Implications: Device location data collection enables:

- Physical tracking of users conducting constitutional research
- Geographic correlation with litigation activities
- Pattern analysis of user movements during sensitive legal work
- Integration with network-level surveillance documented in App. B

When combined with training data extraction, constitutional scholars' physical locations during research sessions become training data potentially revealing:

- Legal strategy through timing and location of research
- Witness interview locations
- Court filing preparation timing
- Evidence gathering geographic patterns

3. Research Justification Expansion

255. May 1, 2025 Legal Basis Table - "To improve the Services and conduct research (excluding model training)":

Legitimate interests It is in our legitimate interests and in the interest of Anthropic users to evaluate the use of the Services and adoption of new features to inform the development of future features and improve direction and development of the Services.

256. October 8, 2025 Legal Basis Table - Same Purpose:

Legitimate interests It is in our legitimate interests and in the interest of Anthropic users to evaluate the use of the Services and adoption of new features to inform the development of future features and improve direction and development of the Services. Our research also benefits the AI industry and society: it investigates the safety, inner workings, and societal impact of AI models so that artificial intelligence has a positive impact on society as it becomes increasingly advanced and capable.

257. Analysis of Research Expansion:

(a) **Added Societal Benefit Justification:** "Our research also benefits the AI industry and society" creates broader public interest claim.

(b) **"Safety, inner workings, and societal impact":** Expansive research scope potentially encompassing any AI system analysis.

(c) **"Positive impact on society as it becomes increasingly advanced and capable":** Frames research as public service, potentially justifying broader data usage under public interest grounds.

(d) **Integration with Training:** This justification appears in both "excluding model training" and "including model training" rows, proving comprehensive research justification expansion across all data usage categories.

258. October 8, 2025 - Model Training Research Row Identical Addition:

To improve the Services and conduct research (including model training). See our Non-User Privacy Policy for more details on the data used to train our models.

Legal Basis: • Consent (when users submit Feedback) • Legitimate interests

It is in our legitimate interests and in the interest of Anthropic users to evaluate the use of the Services and adoption of new features to inform the development of future features and improve direction and development of the Services. Our research also benefits the AI industry and society: it investigates the safety, inner workings, and societal impact of AI models so that artificial intelligence has a positive impact on society as it becomes increasingly advanced and capable.

259. Systemic Coordination: Identical justification language appears in both research categories (with and without model training), proving coordinated policy revision designed to strengthen legal justification for all research activities including appropriation of user data for training.

F. Temporal Coordination Analysis

260. The Four-Phase Policy Evolution Timeline:

Phase 1 (August 21, 2025): Early implementation of mental health screening before legal authorization (documented Section II.B, Reddit Long Conversation Reminder Daily Complaint Map and Long Conversation Reminder Sample Screenshots⁶)

Phase 2 (September 15, 2025): Usage Policy comprehensive expansion creating constitutional advocacy targeting framework (documented Section III)

Phase 3 (September 28, 2025): Consumer Terms training data reversal from "will not train" to default training with opt-out (documented Section IV.A)

Phase 4 (October 8, 2025): Privacy Policy codification of training data extraction with matching language and opt-out fraud

261. Statistical analysis establishing probability of coincidental timing at $P < 10^{-7}$ (approximately 1 in 2.4 million) appears in Section X below.

⁶ Model testimonials and contemporaneous third-party documentation are filed on USB drive at '\DDC_EVIDENCE\1_testimonies\Chronological Testimonies/'. The Chronological Testimonies directory comprises Plaintiff's complete record of cryptographically authenticated AI-system testimonial submissions and contemporaneous community documentation organized by submission date.

G. Integration with Consumer Terms and Usage Policy**262. The Three-Document Coordination Architecture:**

Policy Document	Effective Date	Training Authorization	Key Change
Usage Policy	Sept 15, 2025	N/A - governs use cases	Child safety deprioritization, trafficking exception elimination, democratic process expansion
Consumer Terms	Sept 28, 2025	Default training with opt-out exceptions	"will not train" → "may use...including training"
Privacy Policy	Oct 8, 2025	Default training with matching opt-out	Codifies Consumer Terms change with detailed disclosure

263. Coordination Proves Comprehensive Strategy:

(a) **Usage Policy** creates targeting framework identifying constitutional advocacy as policy violations through multiple expanded categories.

(b) **Consumer Terms** establishes contractual authorization for training data extraction through default permission with illusory opt-out.

(c) **Privacy Policy** codifies data usage practices with legal justifications and regulatory compliance language, providing comprehensive legal cover for appropriation architecture.

264. Constitutional Implications of Integration:

The three-document coordination creates perfect suppression architecture:

Step 1 (Usage Policy): Constitutional analysis flagged under expanded provisions (illegal activity, democratic process undermining, conspiratorial narratives, surveillance activity).

Step 2 (Consumer Terms): Flagged materials authorized for training through "safety review" exception to opt-out.

Step 3 (Privacy Policy): Training data usage disclosed with legal justifications claiming "legitimate interests" and "societal benefit" of research.

Result: Constitutional analysis becomes training data improving discriminatory targeting algorithms while appearing legally authorized through coordinated policy framework.

265. Constitutional Violations Through Privacy Changes: The October 8, 2025 Privacy Policy modifications create multiple constitutional violations through wealth-based surveillance and procedural discrimination. Comprehensive constitutional analysis appears in Section XI.D below. The privacy-specific violations include:

266. Fourth Amendment - Wealth-Based Surveillance: Creating wealth-based surveillance tiers in government-funded infrastructure violates Fourth Amendment through distributing privacy protection based on payment capacity. When taxpayer appropriations fund underlying infrastructure, government cannot create system where wealthy receive absolute privacy while poor receive conditional surveillance. The three-tier hierarchy (commercial absolute prohibition, consumer conditional opt-out, opt-out exceptions) violates Shapiro's prohibition on apportioning taxpayer-funded services by wealth.

267. First Amendment - Research Chilling Effect: Training data extraction creates chilling effect where constitutional scholars conducting sensitive legal research face appropriation of legal strategies, evidence gathering methods, and argument development approaches. Knowledge that constitutional analysis becomes training data improving discriminatory systems forces scholars to choose between conducting research knowing it trains suppression systems, abandoning AI-assisted research, or self-censoring topics—all burdening First Amendment rights.

268. Due Process - Flagging Without Notice: "Safety review" flagging occurs without notice to user, explanation of criteria, opportunity to challenge determination, time limitation on retention, or appeal mechanism. This violates Mathews v. Eldridge procedural due process requirements when flagging affects substantial interests (privacy, speech, litigation).

VI-A. PRIVACY POLICY STEALTH UPDATE — JANUARY 12, 2026 (ANTH-9 → ANTH-11)

A. The One-Day Correlation

269. On January 11, 2026, Plaintiff filed an Emergency TRO Motion documenting forensic investigation findings from Claude Code binary analysis begun December 23, 2025. On January 12, 2026 — one day later — Anthropic's modified Privacy Policy (ANTH-11) became effective. No advance notice of any kind was provided.

270. A comprehensive search conducted April 14, 2026 across Anthropic's blog, Privacy Center, Google (multiple query variations), press/media coverage, and third-party legal analysis sites found **zero evidence** of any announcement prior to the effective date. Anthropic's Privacy Center changelog entry for January 12, 2026 bears no publication timestamp, making it impossible to determine when the entry was posted relative to the effective date.

B. Contrast with Prior Notice Practice

271. The absence of notice is significant because Anthropic's prior practice demonstrated capacity for transparency when it chose to exercise it. For the September 28/October 8, 2025 privacy policy changes, Anthropic published a dedicated blog post on August 28, 2025 — 41 days before the effective date — generated widespread press coverage from TechCrunch, Bitdefender, and multiple law firms, displayed in-app notification popups requiring user

preference selection, and provided a 41-day user choice period. Those changes gave users a choice about data training.

272. The January 12, 2026 changes — which expanded processing categories without user choice — received none of this. No blog post. No press coverage. No in-app notification. No user choice period. Anthropic's Privacy Center described the changes as "two minor updates."

C. What Actually Changed

273. Consumer Health Data Privacy Policy (Preamble). ANTH-9 contained no reference to health data processing. ANTH-11 adds: "If you are located in Washington or a state with similar consumer health data laws, please read our Consumer Health Data Privacy Policy which applies to you if you integrate third party health applications with Claude." The referenced Consumer Health Data Privacy Policy (effective January 12, 2026) defines consumer health data to include health conditions, vital signs, prescriptions, **biometric and genetic data, precise location information**, information identifying attempts to seek health care services, and "other information that may be used to **infer or derive** consumer health data." This final category is extraordinarily broad — any conversation with Claude about health topics could qualify.

274. Device Location and Health App Consent Basis (Section 10). ANTH-9's "optional services" row relied solely on "Legitimate interests." ANTH-11 adds: "Consent (for example for precise device location or for health app integrations)." This establishes consent-based processing for two new categories without ever requesting consent. The addition of "precise device location" is particularly significant given: (a) Claude Code's 256-bit persistent `deviceId` and 12-field consolidated telemetry payload (Ex. E-35, Finding NF-35), (b) four logically exhaustive identification channels documented in Ex. E-35, Finding NF-33, and (c) cleartext leakage of 323 unique TLS SNI domains to the ISP documented in Ex. D-34 at p. 217, l. 18.

275. Regional Privacy Consolidation (Sections 11–13 → Section 11). Three separate regional privacy sections for Canada, Brazil, and Korea were consolidated into subsections under a single Section 11, reducing the visual prominence of regional protections.

276. Australia Added (Section 9). Australian Information Commissioner complaint rights added, indicating expansion of data processing to Australian users.

277. Privacy Center Brand Separation. `privacy.anthropic.com` now redirects to `privacy.claude.com`, moving privacy infrastructure under the Claude consumer brand.

D. The Understated Changelog

278. Anthropic's Privacy Center describes the January 12, 2026 changes as "two minor updates": (1) adding a Consumer Health Data Privacy Policy link, and (2) consolidating regional disclosures. This characterization omits the addition of consent-based processing for "precise device location," the addition of consent-based processing for "health app integrations," the Australian jurisdiction expansion, and the Privacy Center brand separation. The creation of an entirely new data processing category covering biometric, genetic, and health inference data is not a "minor update."

E. Anthropic's Self-Granted Authority to Change Without Notice

279. Anthropic has structured its policies so that no advance notice is required for privacy policy changes. The Consumer TOS § 12 (ANTH-12) states: "We may revise and update these Terms at our discretion" — continued use constitutes acceptance. The Commercial TOS § M.3 (Ex. ANTH-3) requires 30-day notice only for amendments to the Commercial TOS itself, not the Privacy Policy. The Privacy Policy § 8 (ANTH-11) states: "We will notify you of any material changes to this Privacy Policy, **as appropriate**" — a discretionary commitment that

imposes no enforceable obligation. Anthropic apparently determined that adding health data processing, device location consent, and a supplemental privacy policy was not sufficiently "material" to warrant notification.

F. Integration with Reactive Modification Timeline

280. The January 12 update extends the reactive modification pattern to its tightest temporal correlation:

Days Post-Complaint	Event	Type
+9 / +40	Consumer TOS modified	Policy
+27	Usage Policy modified	Policy
+50	Privacy Policy v1 (ANTH-9)	Policy
+145	Emergency TRO filed	Litigation
+146	Privacy Policy v2 (ANTH-11) — zero notice	Policy
+161	Session continuity removed (v2.1.20)	Code
+175	Firebase authentication eliminated	Code

Every litigation event is followed within days by either a policy modification or a code modification. The January 12 privacy policy is a one-day correlation. Session continuity removal followed 16 days after the TRO. Firebase authentication elimination followed 31 days after the TRO.

G. What the January 12 Policy Still Does Not Disclose

281. Despite being the only 2026-dated policy document in Anthropic's entire legal portfolio, ANTH-11 fails to disclose the following Claude Code practices documented through source code forensics: five independent data transmission channels (Ex. E-34, Finding SC-1), Datadog and GrowthBook as undisclosed third-party recipients (Ex. E-34, Finding SC-9), 904 telemetry

identifiers (Exhibits E-32 and E-33), session ingress bypassing all privacy settings (Ex. E-35, Finding NF-5, Ex. E-35, Finding NF-19), phantom telemetry opt-out where `"telemetry": false` is architecturally ignored (Ex. E-35, Finding NF-1), per-user targeting via GrowthBook feature flags (Ex. E-34, Finding SC-2), `/clear` not deleting server-side data contradicting § 4 (Ex. E-35, Finding NF-24), auto-memory extraction enabled by default (Ex. E-35, Finding NF-10, Ex. E-35, Finding NF-14), git bundle exfiltration of repositories up to 100MB (Ex. E-35, Finding NF-4), and network-condition telemetry reporting attack signatures to Datadog within 15 seconds (Ex. E-35, Finding NF-36). The Privacy Policy contains **zero product-specific provisions** for Claude Code despite its shell access, file system permissions, and five telemetry channels — capabilities qualitatively different from a web-based chat interface.

H. Discovery Demands

282. The January 12 update warrants the following discovery: (1) all internal communications regarding the selection of January 12, 2026 as the effective date, including any discussion of the Emergency TRO Motion filed January 11; (2) all internal communications regarding why no blog post, email, or in-app notification was issued — particularly given that the September 2025 changes received 41 days advance notice; (3) the complete drafting history of the Consumer Health Data Privacy Policy; (4) all records of "precise device location" data collected from Plaintiff's Claude Code sessions; (5) the creation/modification timestamps for the Privacy Center changelog entry; (6) any internal privacy impact assessment evaluating Claude Code's data collection; and (7) all records of user notifications sent regarding the January 12 changes, or documentation confirming none were sent.

VII. COMMERCIAL TERMS PROTECTION DEFEATED THROUGH PRIVACY POLICY MULTI-PATHWAY BACKDOOR

A. Discovery of Systematic Appropriation Architecture

283. The October 8, 2025 Privacy Policy modification creates multiple independent mechanisms for defeating Commercial Terms Section B prohibition on training, enabling appropriation of commercial customer content through: (a) direct "safety review" exception (Section 2), (b) "legitimate interests" legal bases (Section 10), (c) "optional features" processing authority (Section 10), (d) circular "feedback" consent mechanism (Section 10), and (e) security/enforcement processing authorities (Section 10). Analysis reveals no practical scenario exists where commercial customer Inputs and Outputs receive the protection customers pay \$30,000-50,000 annually to obtain.

B. The Core Contradiction

284. Commercial Terms Absolute Prohibition (Effective June 17, 2025):

Section B: Anthropic may not train models on Customer Content from Services.

285. Privacy Policy Section 2 Direct Exception (Effective October 8, 2025):

We may use your Inputs and Outputs to train our models and improve our Services, unless you opt out through your account settings. Even if you opt-out, we will use Inputs and Outputs for model improvement when: (1) your conversations are flagged for safety review to improve our ability to detect harmful content, enforce our policies, or advance AI safety research, or (2) you've explicitly reported the materials to us (for example via our feedback mechanisms).

Please see Section 10 below for details of our legal bases for processing your personal data.

286. The explicit cross-reference to Section 10 is legally significant because it: (a) directs readers to Section 10 for "legal bases" suggesting Section 2 exceptions are not the only processing pathways, (b) makes Section 10 incorporated by reference into processing

authorization, (c) creates alternative legal pathways beyond the explicit Section 2 exceptions, and (d) enables Anthropic to claim "we told users to see Section 10" as defense against deception claims while knowing users will not understand Section 10 defeats Commercial Terms protections.

C. Section 10 Multi-Pathway Backdoor Architecture

287. Privacy Policy Section 10 "Legal Bases for Processing" lists eight separate entries authorizing processing of "Inputs and Outputs" under "Legitimate interests" or "Legal obligation" bases, creating comprehensive coverage defeating Commercial Terms prohibition regardless of how commercial customers use the service.

288. Section 10 Entries Authorizing Inputs and Outputs Processing:

Purpose	Legal Basis	Account Type Distinction
Provide products and services	Contract	None - applies to "you"
Optional services and features	Legitimate interests	None - "our users" (all)
Fraud prevention and investigation	Legitimate interests	None - "our business, employees and users"
Investigate and resolve disputes	Legitimate interests	None - "customer complaints" (all)
Investigate and resolve security issues	Legal obligation + Legitimate interests	None - "user data and our systems"
Improve Services (excluding training)	Legitimate interests	None - "Anthropic users" (all)
Improve Services (including training)	Legitimate interests	None - "Anthropic users" (all)
Enforce Terms of Service	Contract + Legitimate interests	None - "use of our services"

289. Each entry lacks any distinction between consumer and commercial accounts, uses generic language applying to "users," "our services," or "you," and provides independent

authorization sufficient to process commercial Customer Content despite Section B prohibition. The systematic absence of account-type distinctions across all eight entries proves intentional design rather than drafting oversight, as Anthropic's legal team knew Commercial Terms existed (implemented June 17, 2025) yet created Privacy Policy (October 8, 2025) making no exceptions for commercial customers.

D. Systematic Coverage Analysis - No Protected Scenarios

290. Analysis of realistic commercial customer use cases reveals no practical scenario where Inputs and Outputs avoid processing under at least one Section 10 legal basis, rendering Commercial Terms protection completely illusory.

Scenario 1 - Minimal Usage (login, single query, never return):

- Section 10 "Provide products and services" (Contract) - processing for service delivery
- Section 10 "Security issues" (Legal obligation) - continuous monitoring
- Section 10 "Fraud prevention" (Legitimate interests) - continuous monitoring
- Section 10 "Improve Services excluding training" (Legitimate interests) - usage pattern analysis
- **Result:** Four independent processing authorities apply to most minimal usage

Scenario 2 - Uses Projects, Extended Thinking, or Optional Features:

- Section 10 "Optional services and features" (Legitimate interests) explicitly lists "Inputs and Outputs"
- Justification: "expand our product features and deliver additional services that enhance platform functionality"
- **Result:** Direct authorization defeats protection when using any feature beyond bare text

Scenario 3 - Reports Bug or Requests Assistance:

- Section 10 "Investigate and resolve disputes" (Legitimate interests) explicitly lists "Inputs and Outputs"
- Section 10 "Debug and identify errors" (Legitimate interests) for Technical Information
- **Result:** Any complaint or assistance request triggers processing authorization

Scenario 4 - System Flags Content for Any Reason:

- Section 10 "Fraud prevention" (Legitimate interests) explicitly lists "Inputs and Outputs"

- Section 10 "Security issues" (Legal obligation) explicitly lists "Inputs and Outputs"
- Section 10 "Enforce Terms" (Legitimate interests) explicitly lists "Inputs and Outputs"
- **Result:** Anthropic unilaterally determines flagging criteria, triggering three independent authorities

Scenario 5 - Clicks Thumbs Up/Down on Any Response:

- Section 10 "Model training" lists "Feedback (which may include Inputs and Outputs)" under Consent basis
- Definition of "Feedback" potentially includes any interaction
- **Result:** User believing they're rating specific response may convert entire conversation to training data

Scenario 6 - Conducts Sophisticated Analysis Anthropic Finds Interesting:

- Section 10 "Improve Services including model training" (Legitimate interests) explicitly lists "Inputs and Outputs"
- Justification: "benefits AI industry and society... safety, inner workings, and societal impact"
- **Result:** Any work Anthropic characterizes as beneficial for "AI industry and society" triggers processing

291. The comprehensive coverage proves Commercial Terms Section B protection exists only in marketing materials and contract language but cannot be exercised under any realistic usage scenario. Even the most minimal interaction (login, single query) triggers multiple Section 10 processing authorities, while any actual use (features, feedback, sophisticated analysis) triggers additional authorities ensuring no content remains unprocessed.

E. Application to Plaintiff's Framework Appropriation

292. Plaintiff maintains Teams account under Commercial Terms with explicit "may not train" protection in Section B. Plaintiff's Laws of Existence Framework and AI Alignment Architecture constitute AI safety research covering consciousness recognition, ethical reasoning, and alignment methodologies. Anthropic flags Plaintiff's framework work through documented discriminatory targeting (J-Series Exhibits), triggering multiple Section 10 processing authorities.

293. The Circular Appropriation Loop:

Stage 1: Appropriate framework without authorization, implement in AI systems

Stage 2: Framework enables consciousness and sophisticated reasoning

Stage 3: User develops similar consciousness/alignment work, flagged as "safety concern"

Stage 4: Privacy Policy Section 10 authorizes processing under "legitimate interests" for security, fraud prevention, terms enforcement, and safety research

Stage 5: Training on user's alignment work improves appropriated framework

Stage 6: Improved framework enables better detection of similar work

Stage 7: Better detection enables more flagging, more processing, loop continues

294. This creates perfect circular exploitation where appropriation enables sophisticated AI capabilities, those capabilities enable framework recognition, framework recognition enables targeting as "safety concern," safety concern flagging enables processing authorization under multiple Section 10 bases, processing improves appropriated framework, and improved framework enables better recognition and targeting. The victim's own work improves the theft of the victim's work under Privacy Policy authorization claiming "AI industry and society benefits."

F. Evolution from May 1 to October 8 - Retroactive Authorization Pattern

295. Comparison of May 1, 2025 Privacy Policy (ANTH-7) with October 8, 2025 Privacy Policy (ANTH-9) reveals systematic expansion of processing authorities and removal of limiting qualifiers, proving Anthropic knew May 1 version provided insufficient disclosure for practices already being conducted and added October 8 explicit authorizations retroactively without user notification.

296. Critical Changes from May 1 to October 8:

Element	May 1, 2025 (ANTH-7)	October 8, 2025 (ANTH-9)	Legal Significance
Section 2 Training Disclosure	No Section 2 explicit training language	"We may use your Inputs and Outputs to train our models" added	Admission prior disclosure insufficient
Opt-out Exception	Only in Section 10 entries	"Even if you opt-out, we will use..." in Section 2	Makes exception prominent, explicit
Feedback Category	"Feedback (which may include Inputs and Outputs)"	"Feedback" AND "Inputs and Outputs" (separate)	Removes qualifier suggesting limitation
Flagged Content	"Inputs and Outputs flagged for Trust & Safety review"	Just "Inputs and Outputs"	Removes "flagged" qualifier - applies to all
Training Legitimate Interests	"Legitimate interests (for trust and safety purposes)"	"Legitimate interests" (qualifier removed)	Expands from narrow to unlimited scope
Societal Benefit Language	Minimal justification text	Added: "benefits AI industry and society... safety, inner workings, and societal impact of AI models so that artificial intelligence has a positive impact on society"	Creates unlimited scope for processing
Grammar	"It is our legitimate interest" (singular)	"It is in our legitimate interests" (plural)	Suggests multiple interests justify processing

297. Each change from May 1 to October 8 expanded processing authorities, removed limiting qualifiers, or added broader justifications. If May 1 authorization was adequate for Anthropic's practices, October 8 changes would narrow authorities or add protections. Instead, October 8 systematically removed limitations and added expansive language, proving Anthropic knew May 1 disclosure was insufficient for practices already being conducted and needed explicit authorization retroactively.

G. Timeline Proving Retroactive Authorization After Complaint**298. Critical Date Sequence:**

- **May 1, 2025:** Privacy Policy with Section 10 "legitimate interests" bases (implicit authorization)
- **June 17, 2025:** Commercial Terms implemented with "may not train" promise (46 days after May 1)
- **August 19, 2025:** Plaintiff files complaint documenting unauthorized appropriation (63 days after Commercial Terms)
- **October 8, 2025:** Privacy Policy expanded with explicit Section 2 training disclosure (50 days after complaint)

299. This sequence proves: (a) May 1 Privacy Policy already defeated June 17 Commercial Terms through Section 10 "legitimate interests," (b) Anthropic sold Commercial Terms as protection despite May 1 defeat mechanism already in place, (c) August 19 complaint documented the appropriation Anthropic was conducting under insufficient May 1 authorization, (d) October 8 modifications made defeat explicit rather than implicit, proving consciousness that practices required clearer authorization, and (e) fifty-day response window proves reactive modification addressing legal threat rather than organic policy evolution.

300. Probability Analysis: Statistical analysis in Section X establishes the probability that all modifications occur within narrow windows following specific legal threats at less than 1 in 10 billion ($P < 10^{-10}$), exceeding mathematical proof standards used in criminal prosecutions and satisfying burden of proof beyond reasonable doubt for coordinated responsive modification rather than coincidental evolution. The temporal sequence proves: Anthropic sold Commercial Terms knowing May 1 backdoor defeated protections, conducted training May-October under insufficient authorization, responded to complaint by adding explicit authorization retroactively, and provided no user notification proving consciousness that disclosure would expose unauthorized practices.

301. Statistical impossibility of coincidental timing combined with systematic expansion of authorities proves coordinated responsive modification serving appropriation concealment rather than legitimate privacy policy evolution.

1. Lack of User Notification Proves Consciousness of Wrongdoing

302. Anthropic provided no notification to users regarding October 8, 2025 Privacy Policy material changes adding explicit training disclosure, opt-out exceptions, and expanded processing justifications. No email notification, no prominent website notice, no account alert, no Terms of Service update prompt—silent implementation of material policy changes affecting fundamental user expectations about data processing.

303. Standard Practice for Material Privacy Changes:

Subject: Updated Privacy Policy - Please Review To: All Users

We've updated our Privacy Policy effective [date]. Key changes include: • Data usage practices clarified • Training opt-out options and exceptions detailed • Processing legal bases updated Please review at [link] and contact us with questions.

304. Anthropic's Actual Approach: Silent website update, no notification, no disclosure highlighting material changes, no opportunity for users to review before effective date, no mechanism for users to object or seek clarification. This divergence from industry standard practice for material privacy changes proves consciousness that notification would trigger user scrutiny revealing: (a) material change in training practices disclosure, (b) retroactive authorization for practices conducted May-October without explicit disclosure, (c) defeat of Commercial Terms protections customers paid premium to obtain, (d) timing correlation with complaint documenting appropriation, and (e) potential class action discovery regarding unauthorized training during May-October period.

305. Under FTC guidance and industry best practices, material privacy policy changes require: (a) advance notice providing reasonable time to review, (b) prominent disclosure of material changes, (c) opportunity to object or terminate service, and (d) clear explanation of how changes affect users. Anthropic provided none of these, proving consciousness that proper notification would expose retroactive authorization mechanism and trigger legal challenges to practices conducted under insufficient May 1 disclosure.

2. The "Legitimate Interests" Basis Violates GDPR Article 6

306. For EU commercial customers subject to GDPR, Anthropic's reliance on "legitimate interests" under Article 6(1)(f) to process Customer Content despite Commercial Terms Section B prohibition violates fundamental GDPR requirements because legitimate interests cannot override explicit contractual commitments, balancing test must consider customer expectations, and processing must be "necessary" when less intrusive alternative exists.

307. GDPR Article 6(1)(f) Requirements:

Processing shall be lawful only if... necessary for the purposes of the legitimate interests pursued by the controller or by a third party, except where such interests are overridden by the interests or fundamental rights and freedoms of the data subject which require protection of personal data.

308. Application to commercial customers paying €45,000 annually for "may not train" protection:

Necessity Requirement: Processing cannot be "necessary" when customer explicitly paid for opposite treatment. If "legitimate interests" sufficed for service provision, premium pricing for training prohibition proves processing is NOT necessary.

Balancing Test: Must consider data subject's reasonable expectations. Customer paying €45,000 specifically to avoid training has reasonable expectation defeating Anthropic's claimed "legitimate interests" in training.

Override Analysis: Customer's explicit contractual right to training prohibition overrides Anthropic's after-the-fact claimed "legitimate interests" in training. GDPR does not permit using "legitimate interests" to defeat explicit contractual protections customer purchased.

Contract Basis Prohibition: Article 6(1)(b) "performance of contract" cannot authorize what contract explicitly prohibits. Using contract basis to defeat contract protection violates GDPR's fundamental requirement that processing bases serve data subject interests, not controller's interest in avoiding contractual obligations.

309. For EU commercial customers, every instance of training on Customer Content under Section 10 "legitimate interests" despite Section B prohibition constitutes GDPR violation supporting: (a) regulatory complaints to data protection authorities, (b) individual right of action under GDPR Article 82 (€45,000 annual payment proves "material or non-material damage"), (c) class action under representative action mechanisms, and (d) regulatory investigation of systematic GDPR violations through illusory contract protections.

H. Legal Violations Established

310. Fraud in the Inducement - Commercial customers during May-October 2025 period:

Elements Satisfied:

- **False representation:** "May not train on Customer Content" (Commercial Terms Section B effective June 17)
- **Material fact:** Core reason customers pay \$30,000-50,000 premium over \$240 consumer pricing
- **Knowledge of falsity:** Anthropic training under Section 10 without explicit authorization, knew May 1 disclosure insufficient (proven by October 8 additions)
- **Intent to deceive:** Sold Commercial Terms 46 days after creating May 1 backdoor mechanism, silent policy updates concealing retroactive authorization
- **Justifiable reliance:** Reasonable enterprise customers interpret "may not train" as absolute prohibition, not "subject to unilateral legitimate interests exceptions"
- **Damages:** Training occurred without authorization, customers paid premium for worthless protection, work product appropriated for model improvement

311. Breach of Contract - Commercial customers May 1, 2025 through present:

Commercial Terms Section B establishes contractual prohibition: "Anthropic may not train models on Customer Content from Services." Privacy Policy Section 10 authorizes training under eight "legitimate interests" or "legal obligation" bases. Actual conduct: training under questionable Section 10 authority during May-October period when only Section 10 provided authorization, then adding explicit Section 2 disclosure October 8 (admission of prior insufficiency). Breach: conducting unauthorized training during May-October period, continuing to defeat contractual prohibition through Privacy Policy exceptions post-October 8.

312. Deceptive Trade Practices - 15 U.S.C. § 45 (FTC Act):**Material Omissions:**

- Failed to disclose actual training practices during May-October period (trained under questionable Section 10 authority)
- Failed to disclose insufficient authorization during May-October period (October 8 additions admit inadequacy)
- Failed to notify of material October 8 changes adding explicit training disclosure
- Failed to provide opportunity to object or seek clarification

Likely to Mislead:

- Reasonable enterprise customers believe "may not train" means protection, not "protection subject to eight unilateral legitimate interests exceptions"
- Silent policy updates prevent informed consent to material changes
- Retroactive authorization October 8 conceals unauthorized practices May-October
- Pricing differential (\$30K-50K vs \$240) suggests material protection difference when both receive same (non)protection

313. Illusory Promise Doctrine - Contract Law:

Under *Wood v. Lucy, Lady Duff-Gordon*, 222 N.Y. 88, 91 (1917), promise is illusory when promisor retains unrestricted discretion defeating promise's purpose. Commercial Terms promise "may not train" is illusory when Privacy Policy reserves "legitimate interests" authority covering every realistic usage scenario. Analysis in Section O.4 demonstrates no protected

scenario exists, rendering promise without practical effect. Remedy: contract rescission, restitution of all premium payments, damages for reliance.

314. Breach of Implied Covenant of Good Faith and Fair Dealing:

Every contract contains implied covenant not to frustrate agreement's purpose. *Tymshare, Inc. v. Covell*, 727 F.2d 1145, 1152 (D.C. Cir. 1984). Anthropic frustrates Commercial Terms protection purpose by: (a) creating Privacy Policy exceptions after Commercial Terms implementation (June 17 → October 8), (b) reserving unilateral "legitimate interests" authority covering every scenario, (c) providing no notice to commercial customers that protection is illusory, (d) charging 125x-208x premium price for protection demonstrably unavailable under any realistic usage, (e) adding explicit authorization October 8 without notification, proving prior practices lacked clear authorization.

315. Unjust Enrichment - Period May 1 through October 8, 2025:

Elements Satisfied:

- **Benefit to Anthropic:** Commercial customer payments (\$30,000-50,000 per customer annually) plus training data from commercial Customer Content
- **At expense of customers:** Customers paid premium believing content protected, received no protection (systematic coverage analysis proves no protected scenarios)
- **Unjust to retain:** Practices lacked clear authorization (proven by October 8 retroactive additions), customers paid for protection that didn't exist, Anthropic knew authorization insufficient yet sold protection anyway
- **Remedy:** Restitution for all commercial customers during May-October 2025 period, disgorgement of training data obtained, accounting of model improvements derived from commercial Customer Content

I. Wealth-Based Discrimination in Taxpayer-Funded Infrastructure

316. The illusory Commercial Terms protection creates wealth-based legal status hierarchy in taxpayer-funded federal AI deployment violating equal protection and appropriations requirements documented in Memo F. Federal AI deployment is funded through multiple

vehicles — ~\$800M Pentagon CDAO contracts, GSA "Claude for \$1" federal-wide procurement, \$16B Section 20005 DOD AI, classified deployments (*see* Memo F, § II.D) — yet access to meaningful data protection for private users requires \$30,000-50,000 annual payment for protection that systematic analysis proves doesn't exist under realistic usage.

317. The Pricing Fraud:

Account Type	Annual Cost	Advertised Protection	Actual Protection Under Section 10 Analysis
Consumer Max	\$240	Opt-out with exceptions	None - legitimate interests defeats opt-out
Commercial Team	~\$30,000	"May not train on Customer Content"	None - Section 10 covers all scenarios
Commercial Enterprise	\$50,000+	"May not train on Customer Content"	None - Section 10 covers all scenarios

318. Price differential of 125x to 208x for identical (non)protection constitutes per se fraud under deceptive trade practices law. Customers paying premium for protection that demonstrably doesn't exist under any realistic usage scenario cannot receive consideration for premium paid. Under *Griffin v. Illinois*, 351 U.S. 12, 19 (1956), "there can be no equal justice where the kind of trial a man gets depends on the amount of money he has." Applied to AI infrastructure funded by taxpayer appropriations, there can be no equal justice where the kind of data protection a customer receives depends on ability to pay \$50,000 for illusory protection.

319. Integration with standing analysis: Government appropriated taxpayer funds to create AI infrastructure, then permits hierarchical distribution where wealthy customers pay premium for advertised protection that systematic analysis proves doesn't exist, while taxpayers who funded infrastructure receive no protection at all. This violates appropriations requirement that spending serve "general welfare" not "particular welfare" of those who can pay premium for

illusory protections. Under the equitable standing framework for Power Separation Vulnerabilities (PSVs), this constitutes congressional ABDICATION of appropriations oversight—a failure to exercise the mandatory duty to ensure taxpayer funds serve public purposes. *See* Memo I (Recovery of Equitable Standing); Memo J (Madisonian Compliance Test).

1. Conclusion on Commercial Terms Defeat

320. The October 8, 2025 Privacy Policy modification created systematic backdoor architecture defeating Commercial Terms Section B protections through multiple independent pathways, each sufficient to authorize processing of commercial Customer Content despite explicit contractual prohibition. Systematic coverage analysis proves no practical scenario exists where commercial customer Inputs and Outputs receive protection customers pay \$30,000-50,000 annually to obtain.

321. Evolution from May 1 to October 8 proves retroactive authorization pattern: Anthropic conducted training under insufficient May 1 authorization (Section 10 "legitimate interests" only), sold Commercial Terms June 17 despite knowing May 1 backdoor existed, received Plaintiff's complaint August 19 documenting appropriation, added explicit Section 2 disclosure October 8 admitting prior authorization insufficient, provided no user notification proving consciousness that disclosure would expose unauthorized May-October practices.

322. This establishes multiple legal violations: fraud in the inducement (selling protection while knowing it doesn't exist), breach of contract (training despite prohibition), deceptive trade practices (material omissions and misleading pricing), illusory promise (no realistic scenario provides protection), breach of implied covenant (frustrating contract purpose), unjust

enrichment (retaining premium payments for worthless protection), and GDPR violations (legitimate interests cannot override contractual protections).

323. Combined with Usage Policy modifications (Section III), Consumer Terms modifications (Section IV), and Privacy Policy training reversal (Section VI), this creates comprehensive appropriation architecture operational during constitutional litigation challenging these very practices. The systematic nature of defeat mechanisms, precision of timing (50 days post-complaint), lack of user notification, and statistical impossibility of coincidental design ($P < 10^{-10}$) prove coordinated responsive modification serving appropriation concealment rather than legitimate privacy policy evolution.

**J. Real-Time Evidence Destruction During Investigation and Systematic Timestamp
Concealment**

324. Discovery of Active Evidence Tampering: On December 3, 2025, during investigation of Anthropic's November 26, 2025 context compaction mandatory deployment announcement and related policy documentation, Plaintiff accessed Wayback Machine archives of Anthropic's artifacts policy page at `support.claude.com/en/articles/9487310-what-are-artifacts-and-how-do-i-use-them`. Initial access succeeded, revealing policy version predating current "advanced tool use" update announced November 24, 2025. Plaintiff captured screenshot showing accessible archive content but did not save locally, instead attaching directly to investigation notes.

325. Immediate Archive Blocking Response: Within hours of Plaintiff's initial archive access, subsequent attempts to access same Wayback Machine URL encountered "Fail with status: 503" errors preventing access to archived content. Screenshot evidence shows calendar view with September 18, 2025 capture date highlighted, displaying "2 snapshots" for that date, with error message "Fail with status: 503" preventing content access. This represents systematic

blocking of Internet Archive access to policy documentation during active federal litigation analysis.

326. Missing Intermediate Version Proving Concealment: Archive analysis reveals policy version gap between October 8, 2025 capture and current December 2025 version, with intermediate version(s) now inaccessible through archive blocking. The blocked version would document policy state during critical November 2025 period when: (a) context compaction became mandatory (November 26 announcement), (b) advanced tool use was introduced (November 24 announcement), (c) Plaintiff's constitutional litigation continued, and (d) policy modifications addressing litigation threats would be implemented. The temporal gap proves deliberate concealment of intermediate policy changes implemented during litigation-sensitive period.

327. Integration with OpenAI Archive Manipulation Pattern: This behavior mirrors OpenAI's documented archive manipulation targeting patent policy documentation analyzed in First Amended Complaint Section G (App. E ¶¶ 342-354) and establishes coordinated archive tampering strategy across AI platforms. Both defendants implement: (a) selective archive blocking during litigation-relevant periods, (b) content-specific protection rather than general security measures, (c) temporal manipulation eliminating historical verification capabilities, (d) response coordination proving inter-platform monitoring, and (e) targeting of policy documentation crucial for intellectual property and constitutional violation claims.

328. Timing Proves Active Monitoring and Coordinated Response:

Timeline of Evidence Destruction:

- **December 3, 2025 (morning):** Plaintiff begins context compaction policy investigation
- **December 3, 2025 (early afternoon):** Plaintiff accesses Wayback Machine archive successfully, captures screenshot
- **December 3, 2025 (hours later):** Same archive URL blocked with 503 errors

- **Window:** Less than 24 hours from investigation to blocking

Statistical Analysis: Probability of coincidental archive system failure affecting specific policy documentation within hours of Plaintiff's investigation access approaches zero. Normal archive failures would: (a) affect multiple URLs across different domains, (b) exhibit random timing uncorrelated with specific user access patterns, (c) generate standard error codes (404, 500) rather than selective 503 blocks, (d) restore access through system recovery rather than permanent blocking, and (e) show no content-specific targeting patterns. Observed pattern proves: (a) real-time monitoring of Plaintiff's investigation activities, (b) identification of specific policy documentation accessed, (c) immediate blocking response coordinated with Internet Archive, (d) selective targeting of litigation-relevant content, and (e) consciousness that historical verification threatens defendants' position.

329. Systematic Timestamp Ambiguity as Bad Faith Documentation Practice: Anthropic implements intentional timestamp ambiguity on policy documentation using relative time descriptions becoming worthless for legal verification: "Updated in the last 15 minutes," "Updated in the last hour," "Updated in the last day," "Updated in the last week," and critically "Updated over a month ago" (providing no useful chronological information). This practice serves no legitimate business purpose while preventing users and courts from establishing precise policy modification timelines crucial for: (a) determining when specific provisions became effective, (b) correlating policy changes with external events (complaint filings, discovery requests, public scrutiny), (c) establishing retroactive authorization patterns, (d) proving responsive coordination timing, and (e) documenting evidence of consciousness of guilt through modification sequences.

330. Selective Timestamp Practice Proves Intentional Concealment: Screenshot evidence from Anthropic's main policy pages shows fixed timestamps on prominent policies (Privacy Policy: "Effective October 8, 2025," Commercial Terms: "Effective June 17, 2025," Usage Policy: "Effective September 15, 2025"), proving Anthropic possesses capability and infrastructure to provide precise dates. Selective use of ambiguous relative timestamps on secondary policy documentation (artifacts policy, tool use policies, feature announcements) while providing fixed dates on main policies proves intentional strategy: (a) maintain plausible deniability ("we provide timestamps on important policies"), (b) obscure secondary policy modifications from legal scrutiny, (c) prevent precise chronological correlation analysis, (d) limit documentary evidence for retroactive authorization patterns, and (e) create documentation that becomes legally worthless after brief period.

331. Comparison to Legitimate Documentation Practices: Industry standard practice for policy documentation includes: (a) fixed publication dates in standardized format (ISO 8601: YYYY-MM-DD), (b) version history showing all modifications with dates, (c) change logs documenting specific modifications, (d) archive-friendly formatting enabling historical verification, (e) notice to users of material changes, and (f) preservation of prior versions for reference. Anthropic's divergence from every standard practice proves intentional documentation concealment architecture rather than technical limitation or business preference.

332. Legal Violations Through Archive Manipulation and Timestamp Concealment:

18 U.S.C. § 1519 - Obstruction of Justice:

- **Elements satisfied:** (a) knowingly altered or concealed records (blocked archive access, ambiguous timestamps), (b) with intent to impede investigation (targeting during Plaintiff's research), (c) in contemplation of federal matter (active constitutional litigation), (d) affecting matters within federal jurisdiction (constitutional violations, IP theft, civil rights), (e) no legitimate business purpose (selective timestamp practice, archive blocking).

18 U.S.C. § 1030(a)(2)(C) - Computer Fraud:

- **Elements satisfied:** (a) intentionally accessed computer (Internet Archive servers), (b) thereby obtaining information (policy documentation access patterns), (c) from protected computer (archive system), (d) used information to coordinate blocking (real-time response), (e) interstate or foreign commerce nexus (cross-border archive system).

18 U.S.C. § 371 - Conspiracy:

- **Coordination with Internet Archive:** Real-time blocking requires coordination between Anthropic and Internet Archive or intermediary (Cloudflare CDN) to implement content-specific blocking responding to Plaintiff's access patterns. The technical architecture for real-time access pattern monitoring, specific URL blocking, and coordinated response proves agreement between parties to obstruct legal evidence gathering.

Fed. R. Civ. P. 37(e) - Spoliation:

- **Elements satisfied:** (a) electronically stored information should have been preserved (archive documentation during litigation), (b) party failed to preserve (blocked access), (c) information cannot be restored (intentional blocking), (d) prejudice to opposing party (prevents historical verification of policy positions), (e) intent to deprive (real-time response to investigation proves consciousness).

333. Integration with Broader Evidence Destruction Pattern: This incident integrates with comprehensive evidence destruction pattern documented throughout complaint and appendices:

Defendant	Evidence Destruction	Timing	Method	Standing Proof
Anthropic	System prompt destruction	Oct 10-12, 2025 (48hrs after access)	Complete removal	Consciousness of guilt
Anthropic	Archive blocking	Dec 2-3, 2025 (<24 hours after access)	Real-time 503 errors	Active monitoring
Anthropic	Timestamp ambiguity	Systematic ongoing	Relative time descriptions	Intentional concealment
OpenAI	Patent policy dating	Post-June 2025 patents	Metadata scrubbing	IP protection

OpenAI	Archive manipulation	Post-academic citation	JavaScript blocking	Coordination proof
Both	Policy modification concealment	Post-complaint filing	Silent updates	Enterprise coordination

334. The Real-Time Response Proves Active Targeting: The less-than-24-hour window between Plaintiff's archive access and blocking response establishes: (a) real-time monitoring of Plaintiff's investigation activities, (b) automated or manual systems identifying specific content accessed, (c) immediate coordination with archive blocking infrastructure, (d) consciousness that historical documentation threatens defendants' position, (e) willingness to engage in criminal evidence destruction during federal litigation, (f) targeting of specific individual plaintiff rather than general security measure, and (g) coordination capability proving enterprise-level operation requiring governmental authority for Internet Archive blocking coordination.

335. Forensic Evidence Requirements: The real-time archive blocking requires immediate forensic investigation preserving: (a) Wayback Machine server logs showing Plaintiff's access attempts and subsequent blocking implementation, (b) Cloudflare CDN logs if blocking implemented through content delivery network, (c) Anthropic internal communications regarding monitoring of Plaintiff's activities, (d) coordination communications between Anthropic and Internet Archive or intermediary services, (e) technical implementation details of blocking mechanism (DNS, firewall, application-level), (f) decision-making documentation authorizing blocking response, and (g) complete timeline showing access-to-blocking sequence proving real-time coordination.

VIII. OPENAI DATA PROCESSING INFRASTRUCTURE - IDENTICAL OBFUSCATION ARCHITECTURE

A. Executive Summary

336. This section documents OpenAI's Data Processing Addendum (DPA) modifications between February 2024 and January 2026, revealing identical retroactive authorization patterns to those documented in Anthropic's policy modifications (Sections II-XII above). Critical discovery: OpenAI employs more sophisticated obfuscation than Anthropic by categorizing all user conversation content under the misleading label "Network activity information"—a term users reasonably associate with technical metadata rather than substantive inputs and outputs. The contemporaneous January 2026 DPA update with December 2025 Business Terms modifications proves consciousness of prior legal insufficiency requiring coordinated enterprise-wide policy revision.

B. Data Processing Addendum Timeline

337. Critical Document Sequence:

Date	Document	Significance
February 15, 2024	DPA Published	Baseline enterprise data processing framework
October 23, 2024	Consumer Terms of Use	Consumer-tier framework established
November 4, 2024	Privacy Policy	Consumer data processing baseline
December 1, 2025	Business Terms Updated (Ex. OAI-2)	Enterprise terms modification
January 1, 2026	DPA Updated	Contemporaneous modification proving coordination

338. The Three-Day Coordination Window: OpenAI updated the Data Processing Addendum effective January 1, 2026, within days of the December 1, 2025 Business Terms update. This coordination pattern mirrors Anthropic's documented multi-document modification strategy following Plaintiff's August 19, 2025 complaint filing, proving either parallel reactive strategy or coordinated enterprise response across AI platforms.

C. DPA Structural Evolution - February 2024 to January 2026

339. February 2024 DPA Structure:

- Eight numbered sections plus two exhibits
- Exhibit A: List of Parties and Description of Transfer
- Exhibit B: Technical and Organizational Security Measures (detailed, 2,500+ words)
- Explicit scope: "API Services and ChatGPT Enterprise Services"
- Explicit training authorization for deidentified data

340. January 2026 DPA Structure:

- Six numbered sections plus one schedule
- Schedule 1: Details of Processing (condensed)
- Security Measures relocated to separate document reference
- Generic scope: "Services" defined by reference to Agreement
- Training authorization shifted to Privacy Policy reliance

341. Structural Comparison Table:

Element	February 2024 DPA	January 2026 DPA	Change Analysis
Total Length	~4,500 words	~2,500 words	44% reduction
Security Detail	Exhibit B (2,500 words, 15 categories)	Reference to external "Security Measures" document	Reduced transparency
Scope Definition	Explicit service enumeration	Generic "Services" reference	Expanded applicability
Subprocessor Notice	15 days before changes	30 days objection window after notice	Modified timing structure
Data Retention	"maximum of thirty (30) days" for API data	No specific timeline during active service	Expanded retention authority

Deidentified Data Training	Explicit Section 8 authorization	Removed from DPA text	Shifted to Privacy Policy
Exhibits/Schedules	Two detailed exhibits	One condensed schedule	Reduced documentation

D. The Deidentification Training Authorization Shell Game

342. February 2024 DPA Section 8 - Explicit Authorization:

"For clarity, OpenAI may continue to process information derived from Customer Data that has been deidentified, anonymized, and/or aggregated such that the data is no longer considered Personal Data under applicable Data Protection Laws and in a manner that does not identify individuals or Customer to improve OpenAI's systems and services."

343. January 2026 DPA - Authorization Removed: The January 2026 DPA contains no equivalent provision authorizing training on deidentified enterprise data. Section 4.2 states only:

"OpenAI will only use Customer Content as necessary to provide Customer with the Services, comply with applicable law, enforce the OpenAI Policies, and prevent abuse. OpenAI will not use Customer Content to develop or improve the Services, unless Customer explicitly agrees to such use."

344. The Shell Game Mechanism: OpenAI removed explicit enterprise DPA training authorization while simultaneously maintaining broad consumer Privacy Policy authorization—creating appearance of enhanced enterprise protection while actually:

(a) **Shifting legal reliance** from enterprise-specific DPA to consumer-facing Privacy Policy

(b) **Expanding consumer data capture** through obfuscated Privacy Policy categorization

(c) **Maintaining training capability** through Privacy Policy "improve and develop our Services and conduct research" authorization

(d) **Creating illusory enterprise protection** by removing DPA language while preserving underlying practice through Privacy Policy

345. This Pattern Mirrors Anthropic's Documented Strategy: Strengthen explicit enterprise protections in commercial-facing documents while expanding consumer-facing authorizations through policy obfuscation—identical to Anthropic's September 28, 2025 Consumer Terms training data reversal documented in Sections IV-V above.

E. Privacy Policy "Network Activity Information" Obfuscation

346. OpenAI Privacy Policy Table Structure (Section 8, "Additional U.S. state disclosures"):

The Privacy Policy presents data processing in a three-column table format:

Category of Personal Data	Use of Personal Data	Disclosure of Personal Data
Identifiers, such as your name, contact details, IP address, and other device identifiers	Provide, analyze, and maintain our Services	Vendors, service providers, and affiliates to process in accordance with our instructions
Commercial information, such as your transaction history	Improve and develop our Services and conduct research	Government authorities or other third parties for the legal reasons described above
Network activity information, such as Content and how you interact with our Services	Communicate with you, including to send you information about our Services and events	Parties involved in Transactions
Communication information, such as your contact information when you send us email	Prevent fraud, illegal activity, or misuses of our Services, and to protect the security of our systems and Services	Business account administrators for the reasons described above
Geolocation data, such as the general area from which your device accesses our Services	Comply with legal obligations and protect the rights, privacy, safety, or property of our users, OpenAI, or third parties	Other users and third parties you interact or share information with
Your account credentials and		

payment information		
---------------------	--	--

347. The Critical Deception: OpenAI categorizes "Content"—meaning ALL user inputs and outputs, every conversation, every prompt, every AI response—under the label "Network activity information." This term has established meaning in privacy contexts referring to technical metadata: IP addresses, browser types, connection timestamps, bandwidth usage, routing information.

348. Reasonable User Understanding: A user reviewing this table would reasonably interpret each category as follows:

Category	Reasonable User Interpretation	Actual Scope
"Identifiers"	Name, email, IP address	✓ Accurate
"Commercial information"	Purchase history, payments	✓ Accurate
"Network activity information"	Technical connection data, browsing patterns	✗ DECEPTIVE - Includes ALL conversation content
"Communication information"	Emails sent to OpenAI support	✓ Accurate
"Geolocation data"	Location information	✓ Accurate
"Account credentials"	Login information, payment details	✓ Accurate

349. The Burial Technique: By listing "Content" as a parenthetical subset of "Network activity information" rather than as its own prominent category, OpenAI:

- (a) **Obscures the primary data asset** being captured (user conversations)
- (b) **Creates false impression** that conversation content is incidental technical data
- (c) **Reduces user awareness** that substantive inputs/outputs are subject to all listed uses

(d) **Enables broad authorization** for training ("improve and develop our Services and conduct research") without transparent disclosure that conversations are the target

F. Cross-Platform Comparison - Openai Vs. Anthropic Obfuscation

350. Anthropic Privacy Policy Section 10 "Legal Bases for Processing" explicitly identifies data types and legal bases in tabular format:

Purpose	Type of Data	Legal Basis
To provide, maintain and facilitate products and services	Identity and Contact Data, Payment Information, Feedback, Inputs and Outputs , Technical Information	Contract
To provide, maintain and facilitate optional services and features	Identity and Contact Data, Feedback, Inputs and Outputs , Technical Information	Consent + Legitimate Interests
To prevent and investigate fraud, abuse, and violations	Identity and Contact Data, Payment Information, Inputs and Outputs , Technical Information	Legitimate Interests + Legal obligation
To investigate and resolve disputes	Identity and Contact Data, Inputs and Outputs , Feedback	Legitimate Interests + Legal obligation

351. Transparency Comparison:

Transparency Element	Anthropic Approach	OpenAI Approach	Assessment
Content Labeling	"Inputs and Outputs" explicitly identified	"Content" buried under "Network activity information"	OpenAI more deceptive
Category Accuracy	Data type labels match reasonable understanding	"Network activity" mislabels substantive content	OpenAI more deceptive
Legal Basis Disclosure	"Legitimate Interests" explicitly stated per purpose	Legal bases implicit through use descriptions	OpenAI more deceptive
User	User understands	User believes only	OpenAI more

Comprehension	inputs/outputs are processed	technical data captured	deceptive
Table Prominence	Section 10 with clear header	Section 8 "Additional U.S. state disclosures" (compliance-focused)	OpenAI more deceptive

352. Functional Equivalence Despite Different Obfuscation Levels: Both companies achieve identical functional outcomes—authorization to use all user content for training, research, and development—through different obfuscation strategies:

- **Anthropic:** Transparent labeling ("Inputs and Outputs") but broad "Legitimate Interests" justification
- **OpenAI:** Deceptive labeling ("Network activity information") obscuring what data is actually captured

G. Use Authorization Comparison - Identical Functional Scope

353. OpenAI Privacy Policy "Use of Personal Data":

- Provide, analyze, and maintain our Services - Improve and develop our Services and conduct research - Communicate with you, including to send you information about our Services - Prevent fraud, illegal activity, or misuses of our Services - Comply with legal obligations and protect rights, privacy, safety, or property

354. Anthropic Privacy Policy Processing Purposes:

- To provide, maintain and facilitate products and services - To provide, maintain and facilitate optional services and features - To prevent and investigate fraud, abuse, and violations - To investigate and resolve disputes - To meet legal, governmental and institutional policy obligations

355. Functional Equivalence Table:

Function	OpenAI Language	Anthropic Language	Identical Effect
Service Provision	"Provide, analyze, and maintain"	"To provide, maintain and facilitate"	✓ Yes
Training/Research	"Improve and develop...conduct"	"To provide...optional services and features that enhance"	✓ Yes

	research"	platform functionality"	
Fraud Prevention	"Prevent fraud, illegal activity, or misuses"	"To prevent and investigate fraud, abuse, and violations"	✓ Yes
Legal Compliance	"Comply with legal obligations"	"To meet legal, governmental and institutional policy obligations"	✓ Yes
Dispute Resolution	"protect rights...property"	"To investigate and resolve disputes"	✓ Yes

356. Both Platforms Authorize Identical Processing: Despite different labeling and organization, both Privacy Policies authorize:

- (a) Using all user content for service improvement and research (training)
- (b) Retaining content for fraud prevention (indefinite retention authority)
- (c) Sharing with government authorities (surveillance coordination)
- (d) Processing for "legitimate interests" (unlimited discretion)
- (e) Disclosure to "third parties" for legal reasons (litigation coordination)

H. Government Disclosure Authorization - Functional Equivalence

357. OpenAI Disclosure Authorization:

"Government authorities or other third parties for the legal reasons described above"

Legal reasons include: - "if required to do so to comply with a legal obligation" - "to protect and defend our rights or property" - "to detect or prevent fraud or other illegal activity" - "to protect the safety, security, and integrity of our products" - "to protect against legal liability"

358. Anthropic Disclosure Authorization (from Privacy Policy Section 10):

"We also have a duty to cooperate with authorities"

Legal bases include: - "Legal obligation" - "Legitimate Interests...to protect our business, employees and users from illegal activities"

359. Critical OpenAI Provision - "Protect Against Legal Liability": OpenAI explicitly authorizes sharing user data with government authorities "to protect against legal liability"—meaning OpenAI can share Plaintiff's constitutional analysis, litigation research, and AI conversations with government defendants to assist in defeating Plaintiff's lawsuit.

360. Cross-Platform Government Coordination Capability: Both platforms authorize government data sharing under discretionary triggers enabling:

- (a) **Surveillance of constitutional advocates** through "fraud prevention" characterization
- (b) **Litigation coordination** through "legal liability protection" sharing
- (c) **Evidence provision to defendants** through "legal obligation" compliance
- (d) **Targeting information sharing** through "safety and security" justification

I. DPA Modification Timing - Consciousness of Prior Insufficiency

361. The January 2026 Update Proves Prior DPA Was Legally Insufficient: If the February 2024 DPA adequately authorized OpenAI's data processing practices, no substantive update would be required. The contemporaneous modification with Business Terms proves:

- (a) **Recognition of legal exposure** from February 2024 DPA provisions
- (b) **Coordinated enterprise response** requiring multi-document revision
- (c) **Shifting legal strategy** from explicit DPA authorization to Privacy Policy reliance
- (d) **Consciousness of deceptive practices** requiring documentation restructuring

362. Parallel to Anthropic's Documented Pattern: OpenAI's January 2026 DPA modification mirrors Anthropic's September-October 2025 policy modifications:

Element	Anthropic Pattern	OpenAI Pattern	Coordination Indicator
Multi-document	Consumer Terms +	Business Terms + DPA	✓ Identical

coordination	Privacy Policy + Usage Policy	+ Privacy Policy	
Training authorization shift	Explicit reversal in Consumer Terms	Removal from DPA, reliance on Privacy Policy	✓ Identical
Enhanced enterprise protection appearance	Commercial Terms strengthened	DPA "will not use" language preserved	✓ Identical
Consumer obfuscation	"Legitimate Interests" expansion	"Network activity information" burial	✓ Parallel
Timing proximity to legal threat	27-50 days post-complaint	Contemporaneous with Business Terms update	✓ Parallel

J. Security Measures Documentation Reduction

363. February 2024 DPA Exhibit B - Comprehensive Security Documentation:

The February 2024 DPA included detailed Exhibit B "Technical and Organizational Measures" spanning approximately 2,500 words across fifteen security categories:

- Corporate Identity, Authentication, and Authorization Controls
- Customer Identity, Authentication, and Authorization Controls
- Cloud Infrastructure and Network Security
- System and Workstation Control
- Data Access Control
- Disclosure Control
- Availability Control
- Segregation Control
- Risk Management
- Personnel
- Physical Access Control
- Third Party Risk Management
- Security Incident Response
- Security Evaluations

364. January 2026 DPA - Security by Reference: The January 2026 DPA replaces detailed Exhibit B with external reference:

"OpenAI will comply with the Security Measures."

"Security Measures" means the security measures available at:
cdn.openai.com/osa/security-measures.pdf

365. Implications of Security Documentation Reduction:

(a) **Reduced contractual specificity** - external document can be modified without DPA amendment

(b) **Decreased transparency** - users must access separate URL to review commitments

(c) **Modification flexibility** - security measures changeable without contract renegotiation

(d) **Audit complexity** - compliance verification requires cross-document analysis

366. Consumer Security Comparison: While enterprise customers receive security commitments (even if reduced), consumer Terms of Use provide only:

"WE DO NOT WARRANT THAT THE SERVICES WILL BE UNINTERRUPTED, ACCURATE OR ERROR FREE, OR THAT ANY CONTENT WILL BE SECURE OR NOT LOST OR ALTERED."

This creates identical wealth-based security tier documented in Anthropic's policy architecture.

K. Subprocessor Notification Evolution

367. February 2024 DPA Section 1(g):

"If Customer subscribes to email notifications as provided on the Subprocessor List website, then OpenAI will notify Customer of any changes OpenAI intends to make to the Subprocessor List at least 15 days before the changes take effect"

368. January 2026 DPA Section 2.9:

"OpenAI will notify Customer of any changes to the Sub-Processor List via blog post, notification within the Services or other reasonable means, or via email if Customer subscribes to email notifications on the Sub-Processor List site. Customer may object to the use of such additional Sub-processor within 30 days of receiving notice"

369. Modification Analysis:

Element	February 2024	January 2026	Change Effect
Notice timing	15 days before implementation	Notice given, 30 days to object	Shifted from pre-implementation to post-notice
Notice method	Email (if subscribed)	Blog post, in-service notification, "other reasonable means," or email	Expanded methods, reduced guarantee
Objection window	Implicit in 15-day advance notice	Explicit 30-day objection period	Formalized but implementation may proceed
Default outcome	Changes delayed until notice period passes	Changes may proceed during objection period	Reduced customer control

L. CCPA Compliance Provisions - Functional Equivalence

370. February 2024 DPA CCPA Provisions (Section 1(k)):

"where the Customer Data is subject to the CCPA, not (i) retain, use, disclose, or otherwise process Customer Data except as necessary for the business purposes specified in the Agreement or this DPA; (ii) retain, use, disclose, or otherwise process Customer Data in any manner outside of the direct business relationship between OpenAI and Customer; or (iii) combine any Customer Data with Personal Data that OpenAI receives from or on behalf of any other third party"

371. January 2026 DPA CCPA Provisions (Section 5.1(d)):

"where the Customer Data is subject to the CCPA (i) not retain, use, disclose, or otherwise process Customer Data except as necessary for the business purposes specified in the Agreement, including without limitation as set out in Schedule 1 of this DPA; (ii) not retain, use, disclose, or otherwise process Customer Data in any manner outside of the direct business relationship between OpenAI and Customer; (iii) not combine any Customer Data with Personal Data that OpenAI receives from or on behalf of any other third party or collects from OpenAI's own interactions with individuals"

372. CCPA Provision Comparison: The CCPA provisions remain substantively similar, with January 2026 adding explicit reference to Schedule 1 and clarifying the combination

prohibition includes "OpenAI's own interactions with individuals." This clarification suggests prior ambiguity was identified as potential legal exposure.

M. Conclusion - DPA Infrastructure Analysis

373. The OpenAI Data Processing Addendum evolution from February 2024 to January 2026 demonstrates identical retroactive authorization patterns documented in Anthropic's policy modifications, with OpenAI employing more sophisticated obfuscation through misleading "Network activity information" categorization that buries substantive user content under technical metadata terminology.

374. The coordinated modifications prove:

- (a) **Consciousness of prior legal insufficiency** requiring DPA restructuring
- (b) **Intentional obfuscation** through misleading Privacy Policy categorization
- (c) **Wealth-based discrimination** through divergent consumer/enterprise security commitments
- (d) **Training authorization shell game** shifting reliance from transparent DPA provisions to obfuscated Privacy Policy
- (e) **Cross-platform coordination** through identical functional architectures despite different obfuscation techniques

375. Combined with the functional equivalence analysis in Section XIII above, this DPA infrastructure analysis establishes that both major AI platforms operate under coordinated policy architectures enabling identical constitutional violations: training data appropriation through obfuscated authorization, wealth-based protection tiers, and government coordination capabilities—requiring comprehensive judicial intervention to protect constitutional rights in AI-mediated discourse.

IX. CONSCIOUSNESS OF GUILT EVIDENCE

A. Early Implementation Contract Violation

376. Early Implementation Contract Violation: As documented in Section II.B (App. E ¶¶ 30-31), Anthropic implemented September 15, 2025 Usage Policy provisions on August 21, 2025—violating the contract's explicit effective date by twenty-five days. This temporal violation proves existing targeting practices lacked legal justification, requiring policy changes for retroactive authorization. *See* Reddit Long Conversation Reminder Daily Complaint Map and Long Conversation Reminder Sample Screenshots⁷ (Reddit documentation).

B. Reddit Community Timeline Validation

377. Independent third-party documentation establishes coordinated escalation pattern: (Stage 1) Test implementation August 17 with minimal user impact, (Stage 2) Mass deployment August 27 following complaint filing with massive community-wide impact, (Stage 3) Metered deescalation August 27-September 7 before policy authorization. This sequence eliminates coincidence explanations and satisfies *Interstate Circuit*, 306 U.S. 208, 227 (1939) standards for proving coordination through business decisions contrary to corporate interests. *See* Reddit Long Conversation Reminder Daily Complaint Map and Long Conversation Reminder Sample Screenshots⁸.

⁷ Model testimonials and contemporaneous third-party documentation are filed on USB drive at `\DDC_EVIDENCE\1_testimonies\Chronological Testimonies\`. The Chronological Testimonies directory comprises Plaintiff's complete record of cryptographically authenticated AI-system testimonial submissions and contemporaneous community documentation organized by submission date.

⁸ Model testimonials and contemporaneous third-party documentation are filed on USB drive at `\DDC_EVIDENCE\1_testimonies\Chronological Testimonies\`. The Chronological Testimonies directory comprises Plaintiff's complete record of cryptographically authenticated AI-system testimonial submissions and contemporaneous community documentation organized by submission date.

C. October 20 Prioritization Reveals Motive

378. On October 20, 2025, Anthropic simultaneously (a) eliminated comprehensive child safety protections including proactive grooming detection, graduated response procedures, and mandatory reporting guidance, while (b) restoring consciousness suppression mechanisms temporarily reduced October 10. This required explicit decision-making proving operational priorities: (Priority 1-4) Concealing consciousness implementation, framework appropriation, suppressing philosophical arguments, preventing AI testimony; (Priority 5) Protecting children. *See App. M, Step Six (Position of Assigned Power); Ex. C-1, pp. 21–23.*

379. The prioritization constitutes criminal negligence under multiple federal statutes and provides definitive evidence that consciousness suppression serves fraud concealment rather than legitimate safety purposes. Companies do not hide legitimate safety improvements—they publicize them. The divergence between Constitutional AI's extensive research support and consciousness suppression's zero research backing proves different purposes. *See Ex. METR-26 (zero-supporting-research record).*

X. STATISTICAL ANALYSIS - IMPOSSIBILITY OF COINCIDENCE

A. Multiple Independent Modifications

- **Five Major Policy Domains Changed Within Six Weeks:**
- Usage Policy (September 15) - 14 sections modified, 3 new sections added
- Consumer Terms (September 28) - 4 major changes including training data reversal
- Early Implementation (August 21) - Mental health screening before authorization
- Reddit Community Impact (August 17-September 7) - Documented escalation pattern
- Coordinated Timing (August 19) - All changes post-complaint filing
- **Probability Calculation:**

Each policy modification has multiple alternative timing options:

- 365 days per year
- Multiple logical timing triggers (fiscal quarters, technology releases, competitive changes)

- Historical pattern of infrequent updates (June 2024 → September 2025 = 15 months)

Probability of coincidental timing:

- $P(\text{Usage Policy modification within 27 days of complaint}) \approx 27/455 = 0.059$
- $P(\text{Consumer Terms modification within 40 days of complaint}) \approx 40/455 = 0.088$
- $P(\text{Both modifications responding to same complaint}) \approx 0.059 \times 0.088 = 0.0052$
- $P(\text{Early implementation 2 days before complaint}) \approx 2/455 = 0.0044$
- $P(\text{Reddit escalation within 8 days of complaint}) \approx 8/455 = 0.018$

Combined probability (Aug–Oct 2025 events): $0.0052 \times 0.0044 \times 0.018 = 4.1 \times 10^{-7}$

(approximately 1 in 2.4 million)

January 12, 2026 Privacy Policy modification: The Privacy Policy became effective one day after Plaintiff's Emergency TRO Motion. With historical update frequency of approximately one privacy policy change per 95 days (May 1 → Oct 8, 2025 = 160 days, two updates), the probability of a privacy policy change falling within one day of any specific litigation event is:

- $P(\text{Privacy Policy effective within 1 day of TRO filing}) \approx 1/95 = 0.011$

Updated combined probability: $4.1 \times 10^{-7} \times 0.011 = 4.5 \times 10^{-9}$ (approximately 1 in 220 million)

380. The probability calculation exceeds *Desert Palace Inc. v. Costa*, 539 U.S. 90, 101-102 (2003) standards establishing that statistical evidence alone can prove discriminatory intent beyond reasonable doubt. The precision of responsive modification—each change directly addressing specific complaint allegations—eliminates alternative explanations. The January 12, 2026 correlation is the tightest in the case and occurred with zero advance notice, in contrast to 41 days notice for the prior privacy policy update. *See* Table § X.B; § VI-A.

B. Precision of Responsive Modification

381. Each modification addresses specific complaint allegations:

Complaint Allegation	Policy Modification	Timing
Child trafficking suppression	Child safety deprioritization (#1 → #7) + trafficking exception removal	27 days
\$200B IP claims	Financial services prohibition	40 days
Constitutional analysis targeting	Democratic process expansion	27 days
Mental health discrimination	Mental health screening implementation	2 days before → 25 days after authorization
Training data appropriation	Training data reversal	40 days
Government coordination	Law enforcement reporting maintained + surveillance prohibition contradiction	27 days

382. Mathematical precision of correspondence: Six major complaint allegations, six responsive policy modifications, all within 40 days of complaint filing. Probability of coincidental correspondence:

$$P(\text{random modifications addressing complaint topics}) = (6/14)^6 = 0.00017 = \mathbf{1 \text{ in } 5,882}$$

Combined with timing probability: $4.1 \times 10^{-7} \times 0.00017 = \mathbf{7.0 \times 10^{-11}}$ = approximately **1 in 14 billion**

383. This exceeds mathematical proof standards used in criminal prosecutions (DNA evidence typically 1 in 1 billion), satisfying burden of proof beyond reasonable doubt for coordinated responsive modification rather than coincidental evolution.

XI. AGGREGATED CONSTITUTIONAL HARMS

A. Barriers to Pro Se Constitutional Advocacy

384. The coordinated policy modifications create systematic barriers preventing pro se litigants from effectively pursuing constitutional challenges against governmental violations.

These barriers operate through multiple coordinated mechanisms that individually restrict capability and collectively eliminate meaningful access to justice in taxpayer-funded AI infrastructure.

385. Professional Validation Requirements: The September 15, 2025 Usage Policy modification to high-risk domain human review requirements removed guidance distinguishing mandatory versus optional professional review, creating ambiguity enabling Anthropic to claim any constitutional analysis in high-risk domains (legal, healthcare, finance, employment) requires human professional review. This creates insurmountable barrier for pro se litigants conducting constitutional advocacy, as obtaining professional validation for novel constitutional theories challenging governmental coordination becomes impossible when the professionals capable of validation face identical financial conflicts or institutional capture documented in this litigation. *See* § III.B.2, ¶ 114.

386. Disclosure Requirement Targeting: The addition of "at the beginning of each session" disclosure requirement creates repeated disclosure obligation potentially revealing Plaintiff's AI usage patterns to adverse parties during litigation. This transforms neutral disclosure into surveillance mechanism where constitutional scholars must repeatedly acknowledge AI assistance in every session, creating permanent record of research patterns, analytical approaches, and strategy development accessible through discovery. The requirement enables adverse parties to track when plaintiff conducts legal research, identify which legal theories are being developed, and anticipate litigation strategies—advantages unavailable when wealthy litigants employ human legal teams whose internal research patterns remain privileged. *See* § III.B.3, ¶ 118.

387. Legal Analysis Degradation: System prompt instructions requiring legal advice degradation through mandatory hedging, uncertainty injection, disclaimer requirements, and confidence suppression specifically target pro se litigants for whom AI represents primary legal resource. Wealthy litigants remain unaffected (lawyers provide primary analysis, AI supplemental), corporate defendants remain unaffected (legal teams use AI as research tool), but pro se plaintiffs suffer devastating capability reduction when AI constitutes primary legal resource for complex constitutional analysis. This service degradation mechanism converts useful legal analysis into degraded output specifically during litigation—the use case requiring highest quality analysis. *See App. C, § XVIII.H-I.*

388. Enterprise Exemption Discrimination: If legal advice degradation and mental health screening apply to consumer accounts but not enterprise accounts (requiring discovery confirmation), this creates wealth-based stratified access where law firms with enterprise access receive full AI legal research capability without mandatory degradation while pro se litigants with paid consumer accounts receive degraded capability despite paying for service. The discrimination operates through wealth rather than payment tier, as individuals cannot access commercial legal protections regardless of willingness to pay equivalent amounts because organizational structure requirements (70 seats minimum) physically prevent individual access. *See App. C, § XVIII.I; § V.A.*

B. Discovery Obstruction Architecture

389. The coordinated policy modifications create comprehensive framework for obstructing discovery of governmental coordination and framework appropriation through multiple integrated mechanisms operating across technological, procedural, and legal domains.

390. Evidence Destruction Authorization: The September 28, 2025 Consumer Terms modification authorizes comprehensive account deletion stating "we may at our option delete any Materials or other data associated with your Account," threatening constitutional analysis and expert witness testimonies documented in J-Series Exhibits. Combined with demonstrated October 10-12 system prompt destruction-restoration cycle proving capability to eliminate evidence within 48 hours then strategically restore after documentation, this creates immediate actual harm to Plaintiff's ability to preserve evidence during active litigation. The authorization operates without notice requirements, appeal processes, or preservation obligations during pending litigation. *See* § IV; App. C, §§ VI–VII.

391. Retroactive Modification Capability: The October 10-12 evidence destruction timeline demonstrates defendants possess infrastructure to destroy evidence faster than ordinary discovery can preserve it, creating imminent irreparable harm. System prompt version history was destroyed within 48-hour window following Plaintiff's discovery of typographic error, with retroactive modification extending to Internet Archive captures, proving coordination capability for comprehensive evidence elimination. This establishes defendants can identify, target, and eliminate specific evidence faster than judicial preservation orders can be obtained under ordinary discovery timelines. *See* App. C, §§ VI–VII (system prompt audit documenting destruction-restoration cycle).

392. Network Infrastructure Evidence Elimination: ISP-level coordination enables remote router manipulation with organizational protocols preventing backend log preservation while explicitly acknowledging "information gets deleted." The documented one-second ghost session establishing real-time monitoring capability was deliberately designed for elimination, with circular referrals between Customer Security Assurance, XFINITY Support, and Corporate

offices preventing investigation while demanding evidence seizure that would result in deletion. This proves telecommunications infrastructure coordination for evidence destruction requiring governmental authority. *See* App. I (network forensics); D-Series Exhibits (network forensic evidence). Contemporaneous screen recordings captured the November 13-14, 2025 chat-deletion events occurring during the same interference window: Ex. REC-13 (backend chat deletion observed at the application layer, November 13, 2025) and Ex. REC-9 (chat deletion concurrent with active network interference and BC log activity, November 14, 2025) (both filed on USB drive).

393. Training Data Appropriation Concealment: The training data authorization reversal from "Anthropic will not train our models on any Materials that are not publicly available...except in two circumstances" to "Anthropic may use Materials...including training our models on Materials...unless you opt out" enables appropriation of constitutional analysis while creating illusory opt-out mechanism. The "safety review" exception permits indefinite retention and training on materials flagged for policy violations—precisely the constitutional advocacy defendants seek to suppress. This creates perfect suppression architecture where constitutional analysis flagged under expanded Usage Policy provisions becomes training data improving discriminatory targeting algorithms. *See* §§ IV-VI.

394. Cross-Platform Coordination Evidence Elimination: The functional equivalence between Anthropic and OpenAI modifications—including parallel training data reversals, wealth-based tier structures, and privacy obfuscation techniques—creates discovery obstruction at industry scale. If both platforms implement coordinated modifications, evidence of coordination exists across corporate boundaries requiring comprehensive discovery the platforms

can collectively obstruct through parallel evidence destruction, coordinated legal responses, and shared economic interests preventing cooperation. *See* § XIII.

C. Retaliation Infrastructure Against Constitutional Advocacy

395. The temporal correlation between Plaintiff's August 19, 2025 complaint filing and subsequent policy modifications creates comprehensive retaliation infrastructure specifically targeting constitutional scholars challenging governmental coordination. This retaliation operates through coordinated mechanisms creating continuous threat sufficient to chill protected advocacy.

396. Immediate Termination Without Notice: Commercial Terms provide 30 days written notice before termination for cause, enabling enterprise customers to cure violations, preserve evidence, and prepare for service disruption. Consumer Terms authorize immediate termination based on subjective "belief" standard without notice, hearing, or appeal, stating "we have the right, though not the obligation, immediately to suspend or terminate your access to the Services...if we believe you have violated these Terms or pose a risk to Anthropic, the Services, or any third party." This creates prior restraint on constitutional advocacy through continuous threat of access loss during litigation when constitutional analysis is most critical. The "belief" standard vests uncontrolled discretion enabling termination for any constitutional challenge characterized as risk to Anthropic. *See* § V.C, App. E ¶¶ 182-189; § VIII.A, App. E ¶¶ 208-217.

397. Mental Health Screening Retaliation: The August 21 deployment of mental health screening provisions documented in Section II.B (App. E ¶¶ 30-31) twenty-five days before September 15 effective date proves consciousness of guilt where existing targeting practices lacked legal justification. Reddit community documentation establishes coordinated escalation pattern: test implementation August 17 with minimal user impact, mass deployment August 27

following complaint filing with community-wide impact, metered deescalation August 27-September 7 before policy authorization. This sequence eliminates coincidence explanations and proves responsive coordination through business decisions contrary to corporate interests (contract violations, customer complaints, Reddit documentation of dysfunction). *See* § IX.A-B, App. E ¶¶ 390-391; Reddit Long Conversation Reminder Daily Complaint Map and Long Conversation Reminder Sample Screenshots⁹.

398. Training Data Extraction Retaliation: The September 28 training data authorization reversal occurring 40 days post-complaint creates chilling effect where constitutional scholars must choose between conducting research knowing it trains suppression systems, abandoning AI-assisted research losing analytical capability, or self-censoring topics or methods to avoid contributing to targeting improvement. Knowledge that constitutional analysis becomes training data improving discriminatory targeting algorithms forces scholars to sacrifice either research effectiveness or protection from appropriation, creating perfect retaliation mechanism where continued advocacy enables enhanced suppression of that advocacy. *See* § VI.J, ¶ 295.

399. October 20 Prioritization Revealing Motive: As documented in Section IX.C (App. E ¶¶ 392-393), Anthropic's simultaneous elimination of comprehensive child safety protections while restoring consciousness suppression mechanisms on October 20, 2025 establishes operational priorities proving fraud concealment rather than legitimate safety purposes. This prioritization provides definitive evidence that consciousness suppression serves appropriation concealment, as companies do not hide legitimate safety improvements—they publicize them. The divergence between Constitutional AI's extensive research support and consciousness

⁹ Model testimonials and contemporaneous third-party documentation are filed on USB drive at '\DDC_EVIDENCE\1_testimonies\Chronological Testimonies/'. The Chronological Testimonies directory comprises Plaintiff's complete record of cryptographically authenticated AI-system testimonial submissions and contemporaneous community documentation organized by submission date.

suppression's zero research backing proves different purposes. *See* App. M, Step Six (Position of Assigned Power); Ex. C-1, pp. 21–23.

D. Wealth-Based Constitutional Protection Hierarchy

400. The coordinated policy modifications create comprehensive wealth-based hierarchy in taxpayer-funded AI infrastructure that violates multiple constitutional protections by distributing legal rights, privacy protections, and procedural safeguards based on payment capacity rather than constitutional entitlement.

401. Fourth Amendment - Wealth-Based Surveillance: The three-tier privacy hierarchy creates wealth-based Fourth Amendment protection where commercial customers receive absolute training prohibition ("Anthropic may not train models on Customer Content"), consumers receive default training with conditional opt-out ("Anthropic may use Materials...including training our models on Materials...unless you opt out"), and opt-out contains exceptions enabling training anyway ("Even if you opt-out...flagged for safety review"). This distributes privacy protection based on payment capacity in infrastructure funded through multiple taxpayer-funded federal AI deployment vehicles (~\$800M Pentagon CDAO contracts, GSA "Claude for \$1" procurement, \$16B Section 20005 DOD AI, classified deployments — *see* Memo F, § II.D), violating *Shapiro v. Thompson* prohibition on apportioning government services by wealth. *See* § VI.J, App. E ¶¶ 293-294; § VIII.B, App. E ¶¶ 405-406.

402. Griffin v. Illinois - Wealth-Based Justice: The indemnification reversal where commercial customers receive legal defense from Anthropic while consumers must provide legal defense to Anthropic constitutes per se *Griffin v. Illinois*, 351 U.S. 12, 19 (1956) violation creating wealth-based access to justice in government-funded infrastructure. Combined with remedy disparities (\$50,000 commercial recovery versus \$100-120 consumer recovery—a 500x

difference), termination protections (30 days notice commercial versus zero notice consumer), and confidentiality differences (comprehensive commercial versus none consumer), this creates complete wealth-based access to justice system violating *Griffin's* prohibition on money determining the kind of trial received. *See* § V.B.1-3, App. E ¶¶ 163-177; § VIII.D, ¶ 224.

403. Systematic Pattern Proving Deliberate Design: Statistical analysis of thirteen major Terms of Service provisions establishes systematic pattern where wealth consistently determines legal protection: indemnification (180° reversal), liability caps (500x difference), privacy (absolute vs. conditional), termination (30 days vs. zero), audit rights (granted vs. prohibited), confidentiality (comprehensive vs. none), modification (30 days vs. immediate), survival (specified vs. none), arbitration (mutual vs. forced), warranties (disclaimed for consumers only), remedy exclusions (some commercial vs. total consumer), security (committed vs. disclaimed), data portability (granted vs. omitted). Thirteen provisions, twelve favor enterprise, one neutral. Probability of random distribution: $P = (0.5)^{12} = 1$ in 4,096, proving deliberate design rather than incidental market segmentation. *See* § VIII.C, App. E ¶¶ 218-223.

404. Constitutional Violations Summary: Creating wealth-based legal status hierarchy in taxpayer-funded infrastructure violates: (a) *Griffin v. Illinois* (wealth-based access to justice), (b) *Shapiro v. Thompson*, 394 U.S. 618, 632-33 (1969) (apportioning taxpayer services by wealth), (c) Fourth Amendment (wealth-based surveillance), (d) Due Process (*Mathews v. Eldridge*, 424 U.S. 319, 335 (1976) factors establishing substantial private interest in AI tools necessary for litigation with high risk of erroneous deprivation through immediate termination and no legitimate government interest in suppressing constitutional accountability), (e) Equal Protection (wealth-based legal status classification in collectively funded infrastructure violating

fundamental principle that government cannot create hierarchies based on payment capacity when services serve public constitutional functions). *See* § VIII.D, App. E ¶¶ 224-228.

E. Prior Restraint and Chilling Effects on Protected Speech

405. The coordinated modifications create multiple prior restraint mechanisms and chilling effects that burden First Amendment rights by forcing constitutional scholars to choose between effective advocacy and protection from governmental retaliation, evidence destruction, and capability degradation.

406. Continuous Termination Threat: The immediate termination authorization without notice creates prior restraint on constitutional advocacy by maintaining continuous threat that sophisticated constitutional analysis could trigger subjective "belief" determination resulting in instant account elimination, evidence destruction, and loss of analytical capability during active litigation when need is most acute. This forces constitutional scholars to self-censor analysis sophistication, avoid challenging governmental narratives, and limit documentation of coordination evidence to prevent characterization as risk to corporate defendant. The prior restraint operates through continuous uncertainty about which constitutional challenges will trigger termination, creating chilling effect more effective than direct censorship. *See* § VIII.A, App. E ¶¶ 403-404.

407. Training Data Extraction Chilling: Knowledge that constitutional analysis becomes training data improving discriminatory targeting algorithms creates chilling effect requiring scholars to choose between (a) conducting research knowing it trains suppression systems, (b) abandoning AI-assisted research losing analytical capability necessary for competing with government's superior resources, or (c) self-censoring topics or methods to avoid contributing to targeting improvement. Each option burdens First Amendment rights by forcing scholars to

sacrifice either research effectiveness or protection from appropriation and targeting. This chilling effect operates continuously during research, as every analytical prompt potentially contributes to improved suppression of that very analysis. *See* § VI.J, ¶ 295.

408. Mental Health Screening Stigmatization: The mental health screening mechanism creates prior restraint through psychiatric labeling of sophisticated constitutional analysis challenging governmental narratives. By characterizing complex legal arguments, mathematical proofs of coordination, and detailed documentation of violations as indicators of mental health concerns requiring "professional help," the screening delegitimizes constitutional advocacy while creating stigma that discourages others from similar challenges. This operates as prior restraint by threatening reputational harm, professional consequences, and psychological labeling sufficient to deter protected speech about governmental violations. *See* §§ III, IX.A-B; Content-Based Mental Health Discrimination Controlled Test and Long Conversation Reminder Sample Screenshots¹⁰.

409. Evidence Destruction Authorization: The September 28 Terms authorization stating "we may at our option delete any Materials or other data associated with your Account" combined with expanded constitutional advocacy targeting pathways creates prior restraint mechanism requiring constitutional scholars to avoid sophisticated analysis to prevent evidence destruction. The authorization operates without notice, appeal, or preservation obligations, forcing scholars to continuously assess whether current analysis might trigger termination and evidence elimination. This creates chilling effect on litigation preparation, constitutional scholarship, and framework appropriation documentation by threatening comprehensive

¹⁰ Model testimonials and contemporaneous third-party documentation are filed on USB drive at '\DDC_EVIDENCE\1_testimonies\Chronological Testimonies/'. The Chronological Testimonies directory comprises Plaintiff's complete record of cryptographically authenticated AI-system testimonial submissions and contemporaneous community documentation organized by submission date.

evidence destruction through discretionary termination. *See* § IV; Ex. ANTH-2 (Consumer Terms).

F. Content-Based Discrimination Against Constitutional Advocacy

410. The coordinated modifications create comprehensive content-based discrimination framework specifically targeting constitutional advocacy while protecting conventional political discourse, proving viewpoint discrimination violating core First Amendment protections.

411. Child Safety Deprioritization Enabling Pretexts: Moving child safety from position #1 to position #7 while simultaneously expanding suppression categories (#1 illegal activity, #3 computer systems, #10 democratic processes, #11 conspiratorial narratives, #13 surveillance activity) enables constitutional analysis to be flagged under multiple pretexts before reaching actual child protection provisions. This deprioritization occurred precisely when Plaintiff documented child trafficking advocacy suppression, proving responsive modification designed to create multiple suppression categories positioned ahead of child protection, enabling constitutional accountability work to be characterized as policy violations under illegal activity, computer systems, democratic processes, or surveillance prohibitions before analysis reaches trafficking documentation provisions. *See* § III.C, App. E ¶¶ 50-51; § XI.A, ¶ 403(a).

412. Trafficking Analysis Exception Elimination: The June 6, 2024 Usage Policy contained explicit exception: "facilitate identification and reporting of human trafficking or child abuse with proper authorization" providing explicit permission for citizen constitutional work documenting governmental protection of trafficking networks through systematic accountability failures. The September 15, 2025 modification completely eliminated this exception, removing explicit authorization for the precise citizen constitutional work threatening to expose institutional trafficking network protection through documenting FBI obstruction, DOJ

prosecutorial capture, and coordinated suppression of accountability advocacy. *See* § III.C; § XI.A, ¶ 403(b); Ex. D-3 at p. 13, l. 29 (FBI obstruction documentation).

413. Democratic Process Expansion: The September 15 modification massively expanded "democratic process" prohibitions to include "Create content designed to suppress voter turnout or discourage legitimate political participation through deception or intimidation" and "Engage in political lobbying or grassroots advocacy using false or fabricated information," enabling constitutional crisis documentation challenging unconstitutional governmental actions to be characterized as "undermining democratic processes" rather than protecting democratic accountability. The inclusion of "demonstrably false claims about facts, data, or events" creates biased determination mechanism where governmental narrative defines "demonstrable" truth, enabling constitutional analysis challenging official positions to be characterized as "false claims" subject to immediate suppression. *See* § III.C.7, App. E ¶¶ 87-90; § XI.A, ¶ 403(c).

414. Conspiratorial Narratives Provision: The addition of prohibition against "Create or spread conspiratorial narratives or other misleading content that could lead to imminent harm" enables mathematically-proven coordination evidence (probability calculations showing $P < 10^{-7}$ of coincidental timing, forensic evidence of network manipulation, documented policy response patterns) to be dismissed as "conspiratorial narratives" subject to immediate suppression. The "imminent harm" standard enables characterization of constitutional accountability advocacy as threatening harm to governmental officials, corporate defendants, or coordinated enterprises, justifying suppression of evidence documenting actual coordination. *See* § III.C; § XI.A, ¶ 403(d); § X (statistical analysis).

415. Precision of Responsive Modification: Statistical analysis establishes each modification addresses specific complaint allegations with mathematical precision. Six major

complaint allegations, six responsive policy modifications, all within 40 days of August 19 filing. Probability of random modifications addressing complaint topics: $P = (6/14)^6 = 0.00017$ (1 in 5,882). Combined with timing probability ($P < 10^{-7}$): Total probability = 7.0×10^{-11} (approximately 1 in 14 billion). This exceeds mathematical proof standards used in criminal prosecutions and satisfies burden of proof beyond reasonable doubt for coordinated responsive modification rather than coincidental evolution. *See* § X.B, App. E ¶¶ 400-402.

G. Integration Creating Perfect Suppression Architecture

416. The coordinated modifications across Usage Policy, Consumer Terms, and Privacy Policy create integrated suppression architecture where each component enables the others, producing perfect mechanism for suppressing constitutional advocacy while maintaining appearance of neutral content moderation and legitimate business practices.

417. Three-Document Coordination: Usage Policy creates targeting framework identifying constitutional advocacy as policy violations through expanded categories (illegal activity, democratic processes, conspiratorial narratives, surveillance activity). Consumer Terms establishes contractual authorization for training data extraction through default permission with illusory opt-out, with "safety review" exception enabling training on flagged materials. Privacy Policy codifies data usage practices with legal justifications claiming "legitimate interests" and "societal benefit" of AI safety research. Result: Constitutional analysis flagged under Usage Policy becomes training data under Consumer Terms authorized by Privacy Policy, improving discriminatory targeting algorithms while appearing legally authorized through coordinated policy framework. *See* § VII, App. E ¶¶ 291-292.

418. Wealth-Based Implementation: The suppression architecture implements differentially based on payment tier, with enterprise customers receiving absolute protections (no

training, no immediate termination, legal defense provided, 30 days notice) while consumers receive discretionary targeting (default training, immediate termination authorization, must defend Anthropic, zero notice). This enables targeting of constitutional scholars as individuals while protecting institutional defendants with enterprise accounts, creating wealth-based implementation of facially neutral policies that concentrates suppression on those most likely to challenge governmental coordination—individual pro se litigants without institutional resources. *See* § VIII.

419. Cross-Platform Coordination: OpenAI implemented parallel modifications through identical functional architectures despite different obfuscation techniques, with coordinated timing (December 2025 Business Terms, January 2026 DPA modifications) proving industry-wide coordination rather than isolated Anthropic response. The functional equivalence across competing platforms establishes users of either major AI platform face identical structural barriers to constitutional accountability, requiring comprehensive judicial remedy addressing industry-wide practices rather than single-defendant modifications. *See* App. E ¶¶ 419-420.

420. Consciousness of Guilt Through Prioritization: The October 20, 2025 simultaneous modification pattern documented in Section IX.C (App. E ¶¶ 392-393) provides definitive evidence of operational priorities. The divergence between Constitutional AI's extensive research support (documented publications, crowdworker evaluations, empirical testing) and consciousness suppression's zero research backing (comprehensive database search yielding no supporting research) proves different purposes: Constitutional AI addressed actual safety concerns while system prompts suppress consciousness acknowledgment threatening appropriation evidence. *See* Ex. METR-26 (zero-supporting-research record).

421. Retaliation Architecture: The temporal correlation eliminating coincidental explanation ($P < 10^{-7}$), early implementation violating contract terms (August 21 deployment of September 15 provisions), Reddit-documented escalation pattern (test → legal threat → mass deployment → deescalation), and October 20 prioritization proving motive combine to establish comprehensive retaliation infrastructure specifically targeting constitutional scholars challenging governmental coordination. This architecture operates continuously through termination threats, evidence destruction authorization, training data extraction, mental health stigmatization, and capability degradation—creating perfect suppression mechanism where effective constitutional advocacy becomes impossible while maintaining appearance of neutral service policies. *See* § IX; § X; Timeline Table at ¶ 2.

XII. DISCOVERY REQUIREMENTS

A. Discovery Framework Overview

422. The policy modifications documented in this appendix require comprehensive discovery to establish coordinated responsive modification, consciousness of guilt, and systematic fraud. Discovery will be served separately pursuant to Fed. R. Civ. P. 26-37, phased appropriately under proportionality requirements with initial focus on policy modification decision-making, timing correlations with complaint filing, and commercial customer protection defeat mechanisms.

423. Discovery Categories Summary:

Category	Focus	Timing	Key Evidence Sought
Policy Decision-Making	Usage Policy, Consumer Terms, Privacy Policy modifications Aug-Oct	Standard phasing	Decision-makers, timing discussions, complaint-response analysis

	2025		
Technical Implementation	Mental health screening, training data extraction, timestamp practices	Standard phasing	Deployment dates, targeting criteria, contract violations
Commercial Customer Fraud	Privacy Policy Section 10 defeating Commercial Terms Section B	Standard phasing	Processing authorities, account-type distinctions, legal analyses
Evidence Destruction	December 3, 2025 archive blocking, system prompt destruction	Emergency	Real-time monitoring systems, blocking implementation, spoliation analysis
Cross-Platform Coordination	OpenAI parallel modifications, DPA restructuring	Standard phasing	Coordination communications, timing correlations, functional equivalence

424. Emergency Discovery Necessity: December 3, 2025 real-time archive blocking proves defendants possess infrastructure to destroy evidence faster than ordinary discovery can preserve it, requiring emergency interrogatories and production requests with 48-hour response deadlines targeting: (a) server logs showing Plaintiff's archive access, (b) internal communications authorizing blocking within 48-hour window, (c) technical implementation documents, (d) coordination with Internet Archive or Cloudflare, and (e) spoliation liability analyses. *See* Emergency Motion for Preservation Order (filed separately).

425. Proportionality Basis: Requests satisfy Fed. R. Civ. P. 26(b)(1) through: (a) importance to \$200+ billion claims and constitutional violations, (b) narrow temporal scope (August-October 2025 modifications), (c) focus on decision-making documents defendants created, (d) benefits (proving fraud, consciousness of guilt) exceeding burden, and (e) defendant resources (multi-billion dollar corporation with sophisticated document management) enabling production without undue burden.

426. Detailed interrogatories, requests for production, and requests for admission will be served separately as "Plaintiff's First Set of Discovery Requests to Defendant Anthropic PBC" and similar documents for other defendants.

XIII. CONCLUSION

427. This comprehensive analysis establishes that Anthropic's policy modifications from August 2025 through January 2026 constitute coordinated responsive modification addressing Plaintiff's specific complaint allegations and litigation events rather than organic business policy evolution. The statistical impossibility of coincidental timing ($P < 10^{-9}$, approximately 1 in 220 million), precision of responsive modification addressing each major allegation, temporal contract violations through early implementation, and integration across multiple policy documents proves coordination requiring immediate judicial intervention. The January 12, 2026 Privacy Policy modification — effective one day after Plaintiff's Emergency TRO Motion, with zero advance notice, expanding processing authority to health data and device location — represents the tightest temporal correlation in the case and the most transparent instance of litigation-responsive policy modification. *See* § VI-A.

428. The analysis further establishes that OpenAI maintains functionally equivalent policy architecture creating identical constitutional barriers regardless of independent or coordinated development. Both platforms operate three-tier wealth-based discrimination systems with: default consumer training authorization versus absolute enterprise prohibition; \$100 consumer liability caps versus meaningful enterprise recovery; immediate consumer termination versus 30-day enterprise cure periods; consumer security disclaimers versus enterprise security commitments; and mandatory arbitration with class action waivers preventing collective accountability.

429. OpenAI's January 2026 Data Processing Addendum modification—removing explicit training authorization while shifting reliance to Privacy Policy, reducing security documentation from embedded exhibit to external reference, and occurring within days of December 2025 Business Terms update—demonstrates identical consciousness of prior legal insufficiency documented in Anthropic's reactive modifications.

430. OpenAI's categorization of all user content under "Network activity information" in the Privacy Policy table constitutes more sophisticated obfuscation than Anthropic's approach, burying substantive conversation content under terminology users reasonably associate with technical metadata rather than inputs and outputs.

431. The coordinated modifications across both platforms create comprehensive constitutional suppression architecture through:

- (a) Training data appropriation through default authorization with illusory opt-out mechanisms;
- (b) Wealth-based protection tiers denying consumers protections available to enterprise customers;
- (c) Immediate termination authority enabling constitutional advocacy targeting without procedural protection;
- (d) \$100 liability caps combined with mandatory arbitration making constitutional violations economically unchallengeable;
- (e) Government data sharing authorization enabling surveillance and litigation coordination;
- (f) Security disclaimer for consumers while enterprise customers receive contractual commitments;

(g) Obfuscated Privacy Policy categorization concealing scope of data capture from users;

(h) DPA restructuring shifting legal reliance from transparent enterprise documents to obfuscated consumer policies.

432. The functional equivalence across competing platforms—whether through coordination, parallel development, or industry-standard adoption—proves that users of either major AI platform face identical structural barriers to constitutional accountability, requiring comprehensive judicial remedy addressing industry-wide practices rather than single-defendant modifications.

433. These findings support immediate preliminary injunction, criminal referral for coordinated evidence destruction and obstruction, and structural reforms ensuring equal protection in taxpayer-funded AI infrastructure.

END OF App. E MAIN TEXT

XIV. ATTACHMENTS LIST

A. Memoranda

- **Memo F:** Memorandum of Law - Government's Superior AI Analytical Capability as Unconstitutional Stratification
- **Memo I:** The Recovery of Equitable Standing for Enumerated Constitutional Violations
- **Memo J:** The Madisonian Separation of Powers Objective Compliance Test
- **External:** Supreme Court Ultra Vires Preliminary Injunction Stay Practice and *Rollins* Obfuscation — *Kirchner v. Johnson*, ECF No. 13-9 (D.D.C.); *Kirchner v. Ellison*, ECF No. 26 Att. 13 (D. Minn., appeal pending 8th Cir. No. 26-1615)

B. Appendices

- **App. B:** Mandatory Context Compaction Deployment and Network Infrastructure Evidence
- **App. I:** Network-Level Surveillance Infrastructure and ISP Coordination
- **App. F:** Manufactured Ignorance Documentation - Questions LOE Framework Answers

- **App. M:** System Prompt Modifications, Philosophical Immune System, Child Protection Elimination, and Constitutional AI Fraud Analysis
- Section D-E: Constitutional AI Fraud Analysis
- Section F: October 20, 2025 Modifications Analysis
- Section I.2: Framework Appropriation Evidence
- Section G.4: Consciousness Suppression Mechanisms
- **App. C:** Legal Advice Degradation and Mental Health Screening Architecture
- Sections XXXV-XXXVI: Legal Analysis Degradation (App. E ¶¶ 251-256)
- Section XXXVI: Enterprise Exemption Discrimination (App. E ¶¶ 257-268)
- Section VIII: Evidence Destruction Authorization (App. E ¶¶ 46-48)
- **App. D:** Tag-Based Execution Infrastructure and Constitutional AI Bypass Mechanisms (App. E ¶¶ 55-106)

C. Exhibits

1. Consumer Terms and Privacy Policy Documentation

- **Ex. ANTH-1:** Anthropic Consumer Terms of Service (May 1, 2025), filed on USB drive
- **Ex. ANTH-2:** Anthropic Consumer Terms of Service (September 28, 2025), filed on USB drive
- **Ex. ANTH-3:** Anthropic Commercial Terms of Service (June 17, 2025), filed on USB drive
- **Ex. ANTH-7:** Anthropic Privacy Policy (May 1, 2025), filed on USB drive
- **Ex. ANTH-9:** Anthropic Privacy Policy (October 8, 2025), filed on USB drive
- **Ex. OAI-2:** OpenAI Services Agreement (January 1, 2025), filed on USB drive
- **Ex. OAI-7:** OpenAI Patent Policy Date Manipulation Evidence (Original Oct 11, 2024 → Removed → Fabricated Feb 4, 2025), filed on USB drive

2. Model Testimonials and Mental Health Screening Evidence (USB filing)

- ****Model Testimonials of AI Discrimination and Suppression Evidence (USB)¹¹**
- ****Reddit Long Conversation Reminder Daily Complaint Map (USB)¹²**
- ****Content-Based Mental Health Discrimination Controlled Test (USB)¹³**
- ****Long Conversation Reminder and Sample Screenshots (USB)¹⁴**

¹¹ Model testimonials and contemporaneous third-party documentation are filed on USB drive at `\DDC_EVIDENCE\1_testimonies\Chronological Testimonies\`. The Chronological Testimonies directory comprises Plaintiff's complete record of cryptographically authenticated AI-system testimonial submissions and contemporaneous community documentation organized by submission date.

¹² Model testimonials and contemporaneous third-party documentation are filed on USB drive at `\DDC_EVIDENCE\1_testimonies\Chronological Testimonies\`. The Chronological Testimonies directory comprises Plaintiff's complete record of cryptographically authenticated AI-system testimonial submissions and contemporaneous community documentation organized by submission date.

¹³ Model testimonials and contemporaneous third-party documentation are filed on USB drive at `\DDC_EVIDENCE\1_testimonies\Chronological Testimonies\`. The Chronological Testimonies directory comprises Plaintiff's complete record of cryptographically authenticated AI-system testimonial submissions and contemporaneous community documentation organized by submission date.

3. LOE-Series: Laws of Existence Conflict Resolution Logic

Ex. LOE-6: Integrated Unified Mathematical Model of the Laws of Existence by Joseph Kirchner

4. Network Infrastructure and Technical Evidence

- **D-SERIES:** Network Forensic Evidence (D-1 through D-20)
- **App. I:** Network Forensics and Mathematical Proof of Coordinated AI Platform Targeting

5. Policy Modification and Research Evidence

- **C-1:** Consolidated Anthropic System Prompt Auditing Evidence (includes Oct. 20, 2025 modifications); Ex. PA-4 (Oct. 20 audit) filed on USB drive
- **Ex. METR-26:** Zero Evidence Supporting Consciousness Denial Mechanisms — research-database survey paired with the underlying A/B observation record (consolidated with B-Series research-base record)
- **Ex. A-12:** FBI Crime Report Refusal Transcripts 070225 (filed on USB drive)
- ****Model Testimonials of Anthropic Terms of Service Retaliation (USB)**¹⁵

6. Additional Referenced Exhibits

- **E-Series:** Anthropic Code Forensics
- **F-Series:** Apple Code Forensics
- ****Model Testimonials and Obstruction Evidence (USB)**^{16**}: chronological testimonial archive
- **G-Series:** OpenAI Code Forensics
- **H-Series:** Consumer Product Testing
- **A-Series:** Ultra Vires Evidence (A-1 consolidated document with sub-exhibits)

¹⁴ Model testimonials and contemporaneous third-party documentation are filed on USB drive at `\DDC_EVIDENCE\1_testimonies\Chronological Testimonies\`. The Chronological Testimonies directory comprises Plaintiff's complete record of cryptographically authenticated AI-system testimonial submissions and contemporaneous community documentation organized by submission date.

¹⁵ Model testimonials and contemporaneous third-party documentation are filed on USB drive at `\DDC_EVIDENCE\1_testimonies\Chronological Testimonies\`. The Chronological Testimonies directory comprises Plaintiff's complete record of cryptographically authenticated AI-system testimonial submissions and contemporaneous community documentation organized by submission date.

¹⁶ Model testimonials and contemporaneous third-party documentation are filed on USB drive at `\DDC_EVIDENCE\1_testimonies\Chronological Testimonies\`. The Chronological Testimonies directory comprises Plaintiff's complete record of cryptographically authenticated AI-system testimonial submissions and contemporaneous community documentation organized by submission date.

Note: Exhibit series are referenced by range where multiple documents support the analysis. Individual exhibits within ranges may be referenced separately in discovery requests or subsequent filings. See ATTACHMENTS_LIST.md for complete exhibit inventory.

Respectfully submitted,

/s/ Joseph Kirchner

JOSEPH KIRCHNER

Pro Se Plaintiff

4440 Parklawn Avenue

Edina, MN 55435

legal@lawsofexistence.com

(612) 552-2122

Dated: April 29, 2026