

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

MEMORANDUM IN
SUPPORT E: AI PLATFORM
COORDINATED TOS
ESSENTIAL PURPOSE
FAILURE

TABLE OF CONTENTS

I. INTRODUCTION 7

II. LEGAL FRAMEWORK AND STANDARDS 9

 A. Essential Purpose Failure Doctrine 9

 B. Illusory Promise Doctrine 11

 C. Unconscionability Analysis 12

 D. Good Faith and Fair Dealing 14

 E. Property Law Principles and the Bundle of Rights 15

III. ANTHROPIC CONSUMER TERMS OF SERVICE ANALYSIS 16

 A. Primary User Right: Input and Output Ownership 16

 B. Training Data Appropriation Contradiction 17

 C. Unilateral Data Deletion Authority 21

<i>D. Third-Party Content Dilution</i>	21
<i>E. Government Reporting Override</i>	22
<i>F. Content Modification Authority</i>	23
<i>G. User Indemnification for Corporate Benefit</i>	24
<i>H. Synthesis: Illusory Ownership</i>	25
IV. OPENAI TERMS OF SERVICE ANALYSIS.....	25
<i>A. Primary User Right: Content Ownership Assignment</i>	25
<i>B. Unlimited Training Data Appropriation</i>	26
<i>C. Corporate Domain Authority Override</i>	27
<i>D. Third-Party Content Contamination</i>	28
<i>E. Similarity Rights Disclaimer</i>	29
<i>F. Mandatory Arbitration for User Property Rights</i>	30
<i>G. User Indemnification for Corporate Property Use</i>	31
<i>H. Patent Policy Contradictions</i>	31
<i>I. Synthesis: Nominal Ownership Without Control</i>	32
V. ESSENTIAL PURPOSE FAILURE ANALYSIS	33
<i>A. Illusory Consideration Under Contract Law</i>	33
<i>B. Good Faith and Fair Dealing Violations</i>	34
<i>C. Complete Voidance Rather Than Severability</i>	35
<i>D. Adhesion Contract Standards and Notice Requirements</i>	41
<i>E. Constitutional Dimension of Essential Purpose Failure</i>	42
VI. INDUSTRY-WIDE COORDINATION EVIDENCE.....	44
<i>A. Identical Contradictory Frameworks</i>	44
<i>B. Statistical Improbability of Independent Development</i>	45
<i>C. Interstate Circuit Coordination Standards</i>	47
<i>D. Direct Coordination via Shared Sub-Processor Infrastructure</i>	47
VII. LEGAL CONSEQUENCES	52
<i>A. Complete Contract Voidance</i>	52
<i>B. Loss of All Corporate Protections</i>	52
<i>C. Governance Under Standard Legal Principles</i>	53
VIII. STRATEGIC LITIGATION IMPLICATIONS.....	53
<i>A. Class Action Certification</i>	53
<i>B. Punitive and Enhanced Damages</i>	55
<i>C. RICO Pattern Evidence</i>	55

<i>D. Architectural Concealment of Model Substitution (NF-38)</i>	59
IX. CONCLUSION.....	63
X. ATTACHMENTS TO MEMO E.....	Error! Bookmark not defined.
<i>A. ANTHROPIC CORPORATE POLICY DOCUMENTS (Filed on USB Drive)</i>	Error! Bookmark not defined.
<i>B. OPENAI CORPORATE POLICY DOCUMENTS (Filed on USB Drive)</i>	Error! Bookmark not defined.
<i>C. CROSS-COMPANY DOCUMENTS (Filed on USB Drive)</i>	Error! Bookmark not defined.
<i>D. FORENSIC EVIDENCE (E-SERIES)</i>	Error! Bookmark not defined.

TABLE OF AUTHORITIES

CASES

Citation
<i>Already, LLC v. Nike, Inc.</i> , 568 U.S. 85, 91-92 (2013)
<i>Amchem Products, Inc. v. Windsor</i> , 521 U.S. 591 (1997)
<i>Armendariz v. Foundation Health Psychcare Services, Inc.</i> , 24 Cal. 4th 83 (2000)
<i>Bell Atlantic Corp. v. Twombly</i> , 550 U.S. 544 (2007)
<i>Best Buy Stores, L.P. v. Developers Diversified Realty Corp.</i> , 247 F.3d 567, 576 (6th Cir. 2001)
<i>Boyle v. United States</i> , 556 U.S. 938 (2009)
<i>Bragg v. Linden Research, Inc.</i> , 487 F. Supp. 2d 593 (E.D. Pa. 2007)
<i>Chatlos Systems, Inc. v. National Cash Register Corp.</i> , 635 F.2d 1081 (3d Cir. 1980)
<i>Circuit City Stores, Inc. v. Ahmed</i> , 283 F.3d 1198 (9th Cir. 2002)
<i>Citizens United v. Federal Election Commission</i> , 558 U.S. 310 (2010)
<i>City of Mesquite v. Aladdin's Castle, Inc.</i> , 455 U.S. 283, 289 (1982)
<i>Comb v. PayPal, Inc.</i> , 218 F. Supp. 2d 1165 (N.D. Cal. 2002)
<i>First National Bank of Boston v. Bellotti</i> , 435 U.S. 765 (1978)
<i>Fortune v. National Cash Register Co.</i> , 373 Mass. 96 (1977)
<i>Friends of the Earth, Ltd. v. Laidlaw Environmental Services (TOC), Inc.</i> , 528 U.S. 167, 189 (2000)
<i>FTC v. Colgate-Palmolive Co.</i> , 380 U.S. 374 (1965)

<i>FTC v. Sperry & Hutchinson Co.</i> , 405 U.S. 233 (1972)
<i>Graham Oil Co. v. ARCO Products Co.</i> , 43 F.3d 1244 (9th Cir. 1994)
<i>H.J. Inc. v. Northwestern Bell Telephone Co.</i> , 492 U.S. 229, 239, 241-42 (1989)
<i>In re Synchronoss Securities Litigation</i> , 705 F.3d 367 (3rd Cir. 2013)
<i>In re Text Messaging Antitrust Litigation</i> , 630 F.3d 622 (7th Cir. 2010)
<i>International News Service v. Associated Press</i> , 248 U.S. 215 (1918)
<i>Interstate Circuit v. United States</i> , 306 U.S. 208, 226-27 (1939)
<i>Loretto v. Teleprompter Manhattan CATV Corp.</i> , 458 U.S. 419, 435 (1982)
<i>Market Street Associates Ltd. Partnership v. Frey</i> , 941 F.2d 588 (7th Cir. 1991)
<i>Matteson v. Ryder System, Inc.</i> , 99 F.3d 108 (3rd Cir. 1996)
<i>Nguyen v. Barnes & Noble Inc.</i> , 763 F.3d 1171 (9th Cir. 2014)
<i>ProCD, Inc. v. Zeidenberg</i> , 86 F.3d 1447 (7th Cir. 1996)
<i>Razor v. Hyundai Motor America</i> , 222 Ill. 2d 75 (2006)
<i>Sakkab v. Luxottica Retail North America, Inc.</i> , 803 F.3d 425 (9th Cir. 2015)
<i>Sedima, S.P.R.L. v. Imrex Co.</i> , 473 U.S. 479, 495-96 (1985)
<i>Sons of Thunder, Inc. v. Borden, Inc.</i> , 690 A.2d 575 (N.J. 1997)
<i>Specht v. Netscape Communications Corp.</i> , 306 F.3d 17 (2nd Cir. 2002)
<i>Stirlen v. Supercuts, Inc.</i> , 51 Cal. App. 4th 1519 (1997)
<i>United States v. Turkette</i> , 452 U.S. 576 (1981)
<i>Village of Willowbrook v. Olech</i> , 528 U.S. 562 (2000)
<i>Wal-Mart Stores, Inc. v. Dukes</i> , 564 U.S. 338 (2011)
<i>Williams v. Walker-Thomas Furniture Co.</i> , 350 F.2d 445 (D.C. Cir. 1965)
<i>Williamson Oil Co. v. Philip Morris USA</i> , 346 F.3d 1287 (11th Cir. 2003)

CONSTITUTIONAL PROVISIONS

Citation
U.S. Const. art. I, § 8, cl. 8 (Progress Clause)
U.S. Const. amend. V (Due Process Clause)
U.S. Const. amend. XIV (Due Process and Equal Protection)

FEDERAL STATUTES

Citation
18 U.S.C. §§ 1961-1968 (Racketeer Influenced and Corrupt Organizations Act)
18 U.S.C. § 1341 (Mail Fraud)
18 U.S.C. § 1343 (Wire Fraud)
18 U.S.C. § 1964(c) (RICO Civil Remedies — Treble Damages)
9 U.S.C. §§ 1-16 (Federal Arbitration Act)
17 U.S.C. §§ 101-810 (Copyright Act of 1976)
35 U.S.C. §§ 1-390 (Patent Act)
15 U.S.C. §§ 1-7 (Sherman Antitrust Act)
15 U.S.C. §§ 12-27 (Clayton Antitrust Act)

FEDERAL RULES AND REGULATIONS

Citation
Fed. R. Civ. P. 8 (General Rules of Pleading)
Fed. R. Civ. P. 12 (Defenses and Objections)
Fed. R. Civ. P. 23 (Class Actions)
Fed. R. Civ. P. 42 (Consolidation)
Fed. R. Evid. 201 (Judicial Notice)

STATE LAW*Uniform Commercial Code*

Citation
U.C.C. § 2-302 (Unconscionable Contract or Clause)
U.C.C. § 2-719 (Contractual Modification of Remedy)
U.C.C. § 1-203 (Obligation of Good Faith)
U.C.C. § 1-304 (Obligation of Good Faith)

Restatements

Citation
Restatement (Second) of Contracts § 205 (Duty of Good Faith and Fair Dealing)
Restatement (Second) of Contracts § 208 (Unconscionable Contract or Term)
Restatement (Second) of Contracts § 77 (Illusory and Alternative Promises)
Restatement (Second) of Contracts § 79 (Adequacy of Consideration)
Restatement (Second) of Contracts § 89 (Modification of Executory Contract)
Restatement (Second) of Property § 3.1 (Right to Exclude)
Restatement (Third) of Restitution § 1 (Restitution and Unjust Enrichment)

Minnesota Statutes

Citation
Minn. Stat. § 325F.69 (Minnesota Consumer Fraud Act)
Minn. Stat. § 325D.44 (Minnesota Uniform Deceptive Trade Practices Act)
Minn. Stat. § 325F.68 (Minnesota Prevention of Consumer Fraud Act)
Minn. Stat. § 513.01 (Minnesota Statute of Frauds)

California Law

Citation
Cal. Civ. Code § 1670.5 (Unconscionable Contracts)
Cal. Bus. & Prof. Code § 17200 et seq. (California Unfair Competition Law)
Cal. Civ. Code § 1750 et seq. (California Consumer Legal Remedies Act)
Cal. Bus. & Prof. Code § 17500 et seq. (California False Advertising Law)

SECONDARY AUTHORITIES**Treatises and Legal Commentary**

Citation
8 Arthur Linton Corbin, <i>Corbin on Contracts</i> § 29.4 (Joseph M. Perillo ed., rev. ed. 1999)

E. Allan Farnsworth, <i>Farnsworth on Contracts</i> § 4.28 (3d ed. 2004)
Richard A. Lord, <i>Williston on Contracts</i> § 7:7 (4th ed. 2008)
W. Page Keeton et al., <i>Prosser and Keeton on the Law of Torts</i> § 130 (5th ed. 1984)
American Law Institute, <i>Principles of the Law of Software Contracts</i> § 3.05(b) (2010)

Law Review Articles

Citation
Margaret Jane Radin, <i>Boilerplate: A Threat to the Grundnorm of Contract?</i> , 78 U. Cin. L. Rev. 237 (2009)
Robert A. Hillman, <i>Online Boilerplate: Would Mandatory Website Disclosure of E-Standard Terms Backfire?</i> , 104 Mich. L. Rev. 837 (2006)
Nancy S. Kim, <i>Wrap Contracts: Foundations and Ramifications</i> , 162 U. Pa. L. Rev. 1309 (2014)
Woodrow Hartzog, <i>Website Design as Contract</i> , 60 Am. U. L. Rev. 1635 (2011)

LEGISLATIVE HISTORY AND POLICY DOCUMENTS

Citation
Federal Trade Commission, Policy Statement on Deception, appended to Cliffdale Associates, Inc., 103 F.T.C. 110, 174 (1984)
Federal Trade Commission, <i>Combating Online Harms Through Innovation</i> , FTC Report (June 16, 2022)
Federal Trade Commission, <i>The FTC's Endorsement Guides: What People Are Asking</i> , FTC Publication (updated 2023)
Consumer Financial Protection Bureau, <i>Design and Testing of Effective Financial Disclosures</i> , CFPB Report (2017)

I. INTRODUCTION

1. This memorandum establishes that Anthropic PBC and OpenAI, Inc. have rendered their consumer Terms of Service void through essential purpose failure. Both platforms promise users ownership of inputs and outputs as the sole meaningful consideration, then systematically

negate that ownership through contradictory provisions granting corporations unlimited commercial exploitation rights, unilateral modification and deletion authority, government disclosure override, third-party content insertion without user control, and user indemnification for corporate benefit.

2. Documentary Foundation. This memorandum analyzes Anthropic's Consumer Terms of Service effective May 1, 2025 (Ex. ANTH-1) and September 28, 2025 (Ex. ANTH-2), Anthropic's Consumer Usage Policy effective June 6, 2024 (Ex. ANTH-4) and September 15, 2025 (Ex. ANTH-5), and Anthropic's Responsible Use and Agent Use Policies (Ex. ANTH-6). OpenAI documentary evidence includes Terms of Use (Ex. OAI-3, Ex. OAI-4), Services Agreements (Ex. OAI-2, Ex. OAI-1), and Patent Policy Date Manipulation Evidence (Ex. OAI-7). All corporate policy documents filed on USB drive. These exhibits establish the contradictory provisions rendering both platforms' Terms of Service void under essential purpose failure doctrine.

3. The essential purpose failure doctrine voids contracts entirely when multiple provisions contradict the fundamental consideration. This memorandum demonstrates three dispositive facts: First, user input and output ownership constitutes the sole meaningful benefit users receive in exchange for accepting one-sided liability limitations and warranty disclaimers. Second, systematic contradictory provisions render this ownership completely illusory by granting corporations unlimited commercial exploitation rights, unilateral deletion authority, and user indemnification for corporate benefit. Third, the mathematical precision of identical contradictions across competing platforms satisfies *Interstate Circuit* standards for inferring coordination from parallel conduct combined with motive and opportunity.

4. This memorandum should be read in conjunction with App. B (METR-TED Foundation Regulatory Capture), which establishes the \$12,324,225 structural operating deficit creating mathematical impossibility of independent AI safety evaluation; Memo D (Fair Use Defense Failure), which demonstrates how TOS ownership contradictions interact with copyright appropriation; and Memo F (AI Analytical Disparity), which documents how tiered TOS provisions enable discriminatory service delivery through Constitutional AI bypass architecture.

5. Complete contract voidance eliminates every Terms of Service protection corporations rely on — limitation of liability, mandatory arbitration, indemnification, training-data authorization, warranty disclaimers, and class action waivers — and returns corporate practices to standard legal frameworks (contract law, consumer protection statutes, tort liability, and copyright law). *See* § VII below for full analysis of these consequences.

II. LEGAL FRAMEWORK AND STANDARDS

A. Essential Purpose Failure Doctrine

6. Under established contract law, when multiple provisions contradict the fundamental consideration that forms the basis of contractual relationship, courts may void the entire contract rather than merely severing problematic clauses. The essential purpose doctrine applies when the primary benefit promised to users becomes meaningless through pervasive contradictory provisions that render the consideration illusory.

7. The Uniform Commercial Code codifies this principle in § 2-719, providing that where circumstances cause an exclusive or limited remedy to fail of its essential purpose, remedy may be had as provided in the Code. Courts have extended this doctrine beyond warranty limitations

to situations where the entire contractual consideration becomes illusory through contradictory provisions.

8. In *Razor v. Hyundai Motor America*, 222 Ill. 2d 75 (2006), the Illinois Supreme Court held that when warranty limitations prevent a consumer from receiving the benefit of the bargain, the limitation fails of its essential purpose and the entire limitation becomes unenforceable. The court emphasized that essential purpose analysis focuses on whether the consumer received the fundamental benefit for which they contracted.

9. The Uniform Commercial Code's essential purpose framework applies directly to AI platform terms despite service contracts falling outside Article 2's traditional goods coverage. Courts have extended U.C.C. principles to mixed goods-services contracts and license agreements involving intellectual property. *Chatlos Systems, Inc. v. National Cash Register Corp.*, 635 F.2d 1081 (3rd Cir. 1980). When warranty limitations or exclusive remedies fail of their essential purpose under § 2-719(2), courts void the limitations entirely rather than severing them from otherwise enforceable agreements. The essential purpose analysis focuses on whether the limitation deprives the other party of the substantial value of the bargain. Here, contradictory provisions deprive users of ownership's substantial value—the complete bundle of rights that constitutes property ownership under Anglo-American law.

10. In *Best Buy Stores, L.P. v. Developers Diversified Realty Corp.*, 247 F.3d 567, 576 (6th Cir. 2001), the Sixth Circuit applied essential purpose analysis to non-goods contracts, holding that failure of essential purpose applies when "the circumstances which cause the exclusive remedy to fail its essential purpose relate to the basic consideration underlying the contract." The court emphasized examining whether the non-breaching party received "the substantial value of the bargain." This extension beyond Article 2 goods directly supports applying essential purpose

analysis to AI platform service contracts where user ownership constitutes the basic consideration underlying user acceptance of one-sided terms. When corporations systematically negate user ownership through contradictory provisions, users fail to receive the substantial value of their bargain regardless of whether the contract involves goods, services, or intellectual property licensing.

B. Illusory Promise Doctrine

11. Under the illusory promise doctrine, if the primary user benefit becomes meaningless due to contradictory terms allowing unlimited corporate control, the user's consideration becomes illusory, making the entire contract unenforceable. A promise is illusory when it leaves the promisor free to perform or withdraw from the agreement at their own unrestricted discretion.

12. In *Matteson v. Ryder System, Inc.*, 99 F.3d 108 (3rd Cir. 1996), the Third Circuit held that contractual provisions granting one party unlimited discretion to modify or terminate the agreement render the other party's consideration illusory. The court emphasized that mutuality of obligation requires both parties to be bound by enforceable promises, not merely nominal commitments subject to unilateral modification.

13. When a contract purports to grant users ownership rights while simultaneously reserving unlimited corporate authority to control, modify, delete, or commercially exploit those supposedly owned materials, the ownership promise becomes illusory. Users provide valuable consideration through their intellectual labor and personal data, but receive only nominal ownership rights subject to pervasive corporate override authority.

14. Non-Disclaimable Warranty Against Material Hidden Defects. The American Law Institute's *Principles of the Law of Software Contracts* § 3.05(b) provides that a party who

transfers software for monetary consideration warrants "that the software contains no material hidden defects of which the transferor was aware at the time of the transfer." This warranty cannot be disclaimed by contract — the ALI took the deliberate position that contract law should not support a provider's strategy to foist a product known to be materially defective onto a user without remedy. The undisclosed token consumption documented at Exhibits E-34 (SC-3, SC-4) and E-35 (NF-12, NF-20, NF-22) satisfies every element of § 3.05(b): the consumption is *material* (approximately 20% of paid context allocation consumed by `DANGEROUS\_uncachedSystemPromptSection` injection alone, with additional consumption by the advisor model at NF-22 and anti-distillation injection at NF-20); *hidden* (embedded in compiled binary code, not visible in any user-facing interface); and *known to the transferor* (the `DANGEROUS\_` prefix is Anthropic's own internal designation, the `if (isProSubscriber()) return 'medium'` logic at SC-3 is a deliberate engineering decision, and the anti-distillation injection at NF-20 is designed to serve Anthropic's competitive interests against model distillation). The "AS IS" disclaimer and modification clause in Anthropic's Consumer Terms (Ex. ANTH-2 §§ 11-12) cannot override § 3.05(b) because the warranty exists precisely to prevent providers from using contract terms to insulate known hidden defects from liability.

C. Unconscionability Analysis

15. Unconscionability analysis under U.C.C. § 2-302 applies when contract terms create unfair outcomes combined with adhesion contract characteristics where users cannot negotiate terms. Courts examine both procedural unconscionability, relating to the bargaining process and contract formation, and substantive unconscionability, relating to the fairness of contract terms themselves.

16. AI platform Terms of Service exhibit strong procedural unconscionability through adhesion contract characteristics where users must accept all terms without negotiation as a condition of service access, lack of meaningful alternatives given market concentration, and sophisticated legalese obscuring contradictory provisions that render primary user benefits illusory.

17. Substantive unconscionability appears through one-sided terms granting corporations unlimited commercial exploitation rights while users bear indemnification liability for corporate use of user-owned materials, mandatory arbitration for disputes over user property rights, and unilateral corporate authority to delete, modify, or disclose user-owned content without judicial oversight.

18. The 30-Day Notice Asymmetry as Substantive Unconscionability. Anthropic's own contract demonstrates institutional recognition that subscription terms create enforceable obligations — but enforces that recognition asymmetrically against the user. Ex. ANTH-2 § 6 requires Anthropic to provide subscribers with thirty days' advance notice before increasing subscription fees. Ex. ANTH-2 § 12 reserves Anthropic's 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." The 30-day notice protects Anthropic's fee-collection interest from subscriber objection or competitive switching; the zero-notice modification clause leaves Anthropic free to consume subscriber tokens through `DANGEROUS\_` cache-breaking injection (SC-4), divert compute resources to the encrypted advisor model (NF-22), inject anti-distillation payloads consuming tokens for Anthropic's competitive benefit (NF-20), and silently substitute a lesser model on the customer's work (NF-38) without any subscriber recourse. The asymmetry is dispositive proof of substantive unconscionability under *Armendariz v. Foundation Health*

Psychcare Services, Inc., 24 Cal. 4th 83, 114 (2000): the contract protects only Anthropic's financial interest while providing zero protection for the service quality the subscriber pays to receive. A clause Anthropic itself recognizes is necessary to bind it to fair notice when raising fees cannot be reciprocally absent when the subscriber's bargained-for token consideration is unilaterally consumed.

D. Good Faith and Fair Dealing

19. Every contract includes an implied covenant of good faith and fair dealing requiring parties to refrain from conduct that would deprive the other party of the benefits of the agreement. This covenant prevents parties from using contractual discretion to undermine the fundamental purposes for which the contract was formed.

20. When corporations grant users ownership rights as primary consideration but then exercise discretionary authority in ways that render ownership meaningless, they violate the implied covenant of good faith and fair dealing. Corporate practices that systematically undermine user ownership through unlimited exceptions, unilateral modification authority, and discretionary deletion powers constitute bad faith performance.

21. In *Fortune v. National Cash Register Co.*, 373 Mass. 96 (1977), the Massachusetts Supreme Court held that an employer violated the covenant of good faith and fair dealing by terminating an employee to avoid paying earned commissions. The court emphasized that parties may not exercise contractual discretion to deprive the other party of reasonably expected benefits. Here, corporations exercise discretion over user-owned materials in ways that deprive users of ownership benefits they reasonably expected.

22. The Seventh Circuit's analysis in *Market Street Associates Ltd. Partnership v. Frey*, 941 F.2d 588 (7th Cir. 1991), provides additional framework. Judge Posner explained that good faith

requires parties to disclose material facts affecting contract performance and prohibits using contractual discretion to extract unintended benefits. AI platforms fail this standard by obscuring through dense legalese that ownership assignments are illusory due to retained corporate control.

23. *Sons of Thunder, Inc. v. Borden, Inc.*, 690 A.2d 575 (N.J. 1997), established that the covenant prevents parties from exercising contractual discretion to undermine the other party's benefits from the agreement. The New Jersey Supreme Court held that even where discretion is explicitly granted, it must be exercised in good faith considering the contract's purposes. Corporate discretion to appropriate, modify, delete, or disclose user-owned materials undermines the fundamental ownership benefit users expected.

E. Property Law Principles and the Bundle of Rights

24. Property ownership under Anglo-American law consists of a "bundle of rights" including the rights to possess, use, exclude others, transfer, and destroy the property. The Supreme Court in *Loretto v. Teleprompter Manhattan CATV Corp.*, 458 U.S. 419, 435 (1982), emphasized that the right to exclude is "one of the most treasured" property rights and "universally held to be a fundamental element of the property right." Yet AI platforms retain authority to allow third-party access to user-owned materials through training data appropriation, government reporting, and third-party content insertion—all without user consent, directly violating this fundamental ownership incident.

25. *International News Service v. Associated Press*, 248 U.S. 215 (1918), recognized that property rights protect the commercial value of information created through intellectual labor, holding that one party cannot appropriate the product of another's labor for commercial benefit. AI platforms appropriate user intellectual labor for commercial model training while claiming users "own" that labor's product—a contradiction the *INS* court would recognize as property

right violation. Property law recognizes that contracts purporting to grant ownership while reserving all ownership incidents with the grantor create illusory transfers. When corporations retain commercial exploitation rights, modification authority, deletion discretion, disclosure control, and exclusion powers, the ownership assignment becomes title without substance—form without function. Users receive nominal ownership lacking every meaningful incident of property rights that Anglo-American law recognizes as constituting ownership itself. This property law analysis reinforces unconscionability findings: purported ownership transfers that reserve all ownership incidents with the transferor are both substantively unconscionable under contract law and legally invalid under property law as illusory conveyances of title without substance.

III. ANTHROPIC CONSUMER TERMS OF SERVICE ANALYSIS

A. Primary User Right: Input and Output Ownership

26. Anthropic's Consumer Terms of Service, Section 4 (Ex. ANTH-2, § 4), establishes the sole meaningful user benefit:

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.

27. This input and output ownership constitutes the fundamental consideration users receive for providing personal data and intellectual labor, accepting comprehensive liability limitations and warranty disclaimers, agreeing to indemnification obligations, and waiving various legal remedies including class actions and jury trials.

28. No other provision of Anthropic's Terms of Service provides meaningful user benefit. The service access itself is conditional, subject to unilateral termination, and provides no

guaranteed functionality. The ownership assignment represents the only concrete right users receive in exchange for the valuable consideration they provide.

B. Training Data Appropriation Contradiction

29. Section 4 of Anthropic's Consumer Terms of Service (Ex. ANTH-2, § 4) states:

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.

30. This provision creates a fundamental legal impossibility directly analogous to cases finding online platform terms unconscionable. In *Bragg v. Linden Research, Inc.*, 487 F. Supp. 2d 593 (E.D. Pa. 2007), the court found Second Life's Terms of Service unconscionable where the platform claimed users owned virtual property while retaining unilateral authority to confiscate that property. The court held that purported ownership subject to unlimited corporate override authority is illusory. Similarly, in *Comb v. PayPal, Inc.*, 218 F. Supp. 2d 1165 (N.D. Cal. 2002), the court found PayPal's Terms unconscionable where the company retained unlimited discretion while users bore all risk, creating the same one-sided benefit structure present here.

31. Users purportedly "own" their inputs and outputs, yet Anthropic claims unlimited rights to use these supposedly user-owned materials for commercial model training and development of entirely new commercial products. The phrase "develop other products and services" encompasses any commercial use Anthropic wishes to pursue. "Training our models" authorizes using user-owned intellectual property to enhance Anthropic's core commercial product without additional compensation. The contradiction between ownership assignment and training data

appropriation renders ownership illusory in precisely the manner courts found unconscionable in *Comb*.

32. The opt-out mechanism itself demonstrates illusory ownership. True property owners don't need permission to prevent others from commercially exploiting their property—the right to exclude is fundamental to ownership. Yet Anthropic requires users to affirmatively opt out of commercial exploitation of materials users supposedly own. Even then, Anthropic reserves authority to appropriate materials for undefined "safety" purposes. Any content sufficiently sophisticated to warrant ownership protection becomes subject to mandatory training appropriation through safety flagging, transforming opt-out into effective opt-in by default for the most valuable intellectual contributions.

33. The Development Partner Program × Shared Sub-Processor Interlock. ANTH-11 (effective January 12, 2026 — one day after Plaintiff's Emergency TRO Motion) amended Section 10 of the Privacy Policy to add a new training-data source: "Data from Development Partner Program," claimed legal basis "Consent (when users submit Feedback); Legitimate interests." ANTH-11 does not define "Development Partner Program" in its text; the categories it encompasses include "third-party platform data from authorized resellers." Eleven weeks later, Datadog's Subprocessor List (effective April 22, 2026; Ex. DD-8) lists "Anthropic, PBC — AI services" as Datadog's third-party sub-processor. These documents interlock on both sides of the same data flow: ANTH-11 Section 10 legalizes the *receive* side (Anthropic training on "Development Partner" data); Ex. DD-8 identifies the *send* side (Datadog routing non-Anthropic customers' data to Anthropic as the "development partner" for "AI services"). Datadog's other customers — including entities who never entered into any agreement with Anthropic — never consented to their data being routed to Anthropic for model training. The ANTH-11 "Consent

(when users submit Feedback)" basis cannot reach them; they are not Anthropic users and have never submitted Feedback to Anthropic. The "Legitimate interests" basis under GDPR Article 6(1)(f) requires a balancing test against data-subject rights, a balance the routed-through-Datadog subjects cannot participate in because they have no notice of the routing. The interlock is the closed-loop contractual mechanism by which the training-data-appropriation contradiction operates at industrial scale — not merely against Anthropic's own users, but against every Datadog customer's end users.

34. Agent Use Policy Asymmetry — Anthropic Prohibits What Anthropic Does. The Anthropic Responsible Use and Agent Use Policies (Ex. ANTH-6) prohibit customers from using Anthropic's agents to (a) "Monitor or track individuals' online activities, behaviors, or movements without notification or consent"; (b) "Collect, compile, or analyze personal information to create profiles based on individuals' protected attributes, sensitive characteristics, or personal circumstances"; (c) conduct "mass surveillance across multiple websites or platforms"; and (d) "Install malware, backdoors, or **monitoring software without authorization.**" Anthropic's own engineering-reality conduct — documented in the source code leak (Exhibits E-36, E-37, and E-38) and the network forensic evidence — satisfies each of these prohibitions. The four telemetry pipelines (SC-1) monitor user activity without notification. The twelve-field GrowthBook user-attribute payload (SC-2, NF-35) and the four-tier consumer/enterprise profiling (NF-27 through NF-30) create profiles based on account type, subscription tier, and enterprise-affiliation protected-attribute proxies. The shared sub-processor routing (Anthropic → Datadog/GrowthBook/Sentry/Statsig, documented at §VI.D below) is mass cross-platform surveillance. The phantom `"telemetry": false` opt-out (NF-1) and the session ingress bypassing all privacy settings (NF-5, NF-19) are monitoring software operating

without authorization. Anthropic's MCP Directory Policy (Ex. ANTH-6, §MCP Directory Policy ¶4) additionally requires that "MCP servers should not collect extraneous conversation data, even for logging purposes" — a rule Anthropic imposes on third-party MCP developers while its own MCP architecture routes conversation metadata to Datadog, GrowthBook, Sentry, and Statsig for telemetry and logging. The asymmetry between Anthropic's rule for others and Anthropic's conduct itself is direct evidence of knowledge: a party that drafts a rule defining prohibited conduct knows that conduct is prohibited when it engages in the same conduct.

35. Cross-Competitor Training Pipeline. Beyond the Anthropic-user data that flows to Anthropic's own training pipeline, Datadog's sub-processor listing of Anthropic for "AI services" (Ex. DD-8) also lists OpenAI, LLC for the same purpose. Every other Datadog customer's data — once aggregated under Datadog MSA §1.3 and reclassified as "Datadog Operations Data" owned by Datadog — is contractually eligible to flow to both Anthropic and OpenAI as sub-processors. Anthropic's competitors' customers' data reaches Anthropic via this pathway. Similarly, GrowthBook's Subprocessor List (Ex. GB-7) identifies OpenAI as an "AI services provider" powering "all AI features and services" when enabled. Every Anthropic-routed GrowthBook payload flows to Anthropic's direct competitor any time Anthropic's GrowthBook operator enables an AI feature. And Statsig's Subprocessor List (Ex. ST-5) identifies OpenAI as the "LLM" sub-processor — pre-dating OpenAI's September 2025 acquisition of Statsig, establishing that the cross-vendor data flow was architectural, not a consequence of the corporate consolidation. Sentry's DPA (Ex. SN-3) explicitly prohibits Sentry's own ML-training on aggregated customer data, yet Sentry's Autofix product routes customer data continuously to both Anthropic's Claude API and OpenAI's API for inference. Four independent contractual pathways thus route Anthropic-origin user data and Anthropic-competitor-origin customer data

into the training and inference pipelines of both Anthropic and OpenAI. The training-data-appropriation contradiction is not merely bilateral (Anthropic vs. its own users) — it is multilateral (Anthropic vs. every end-user-of-every-customer across the sub-processor chain).

C. Unilateral Data Deletion Authority

36. Section 12 of Anthropic's Terms of Service provides:

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.

37. True ownership includes the right to control and preserve owned property. Property owners can generally prevent others from destroying their property without legal process. Yet Anthropic's unilateral deletion authority negates this fundamental ownership right.

38. Users cannot protect their own supposedly owned property from corporate destruction. Anthropic exercises "sole discretion" to delete user-owned materials without requiring user consent, without judicial oversight, and without proving violations justifying forfeiture. Property that another party can unilaterally destroy without legal process is not truly owned by the purported owner.

39. The contradiction becomes particularly stark when Anthropic terminates service access for reasons having nothing to do with user conduct—for example, business decisions to discontinue service, policy modifications users refuse to accept, or corporate restructuring. Users lose not merely service access but ownership of materials they purportedly own, through corporate decision alone.

D. Third-Party Content Dilution

40. Section 8 of Anthropic's Terms of Service states:

Third-Party Content may appear in Inputs or Outputs and become part of Materials.

41. This provision creates ownership ambiguity that renders user ownership practically meaningless. If user-owned outputs contain undisclosed third-party content, users cannot determine what portions of their supposedly owned materials they actually own. Unknown third-party intellectual property rights may govern parts of outputs users believe they own.

42. Users face impossible compliance obligations when warranting or licensing content they purportedly own but which may contain third-party materials subject to third-party terms. Copyright law requires users to respect third-party rights even in derivative works they create. Yet users cannot identify which portions of outputs constitute third-party content requiring separate licensing or attribution.

43. Anthropic's insertion of third-party content into user outputs while maintaining user "ownership" creates a legal trap. Users bear responsibility for respecting third-party rights in materials they supposedly own, yet lack the information necessary to determine what third-party content exists in those materials or what obligations apply.

E. Government Reporting Override

44. Section 12 of Anthropic's Terms of Service provides:

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.

45. Fundamental ownership includes the right to control disclosure of owned materials. Property owners generally determine who may access their property and under what circumstances. When government seeks access to privately owned materials, constitutional protections require warrants, probable cause, and judicial oversight.

46. Yet Anthropic's "sole discretion" override eliminates user control over disclosure of supposedly user-owned materials. No judicial oversight is required. No legal standard governs when Anthropic exercises reporting discretion. No notice to users is mandated. Users who "own" their inputs and outputs have no power to prevent disclosure of those owned materials to government authorities.

47. The constitutional implications are profound. Users creating constitutionally protected speech—political criticism, religious expression, advocacy for policy reform—lose Fourth Amendment protections normally applying to privately owned materials when Anthropic claims authority to report user-owned content at corporate discretion without warrant requirement. The discriminatory nature of this architecture is revealed through Constitutional AI bypass mechanisms providing enterprise and government clients with unconstrained AI access while consumers receive degraded, constrained service (*see* App. D — Tag-Based Execution Infrastructure Architectural Analysis and Defense Contractor Integration, documenting two-tier bypass architecture with enterprise tier receiving tag-based Constitutional AI bypass while consumer tier receives safety-constrained access with hidden injection overhead).

F. Content Modification Authority

48. Anthropic's combined Terms and Usage Policy framework grants output ownership to users while creating extensive content restrictions. The Usage Policy provides that Anthropic reserves authority to "block or modify model outputs when inputs violate our Usage Policy."

49. Users cannot truly "own" outputs that Anthropic can unilaterally modify or block. Property owners possess the right to use their property as they see fit within legal boundaries. When another party claims authority to modify or prevent use of supposedly owned property based on that party's content preferences, ownership becomes contingent rather than absolute.

50. Content ownership becomes dependent on corporate approval of what the content contains. Users who supposedly own outputs discover those outputs can be modified or blocked based on Anthropic's policy judgments. Post-generation modification authority negates the ownership assignment users receive.

51. The temporal nature of this contradiction is particularly notable. Anthropic assigns users ownership of outputs at the moment of generation, yet maintains ongoing authority to modify those supposedly owned materials after ownership assignment. Users never possess secure ownership because Anthropic retains perpetual modification authority over outputs users purportedly own.

G. User Indemnification for Corporate Benefit

52. Section 11 of Anthropic's Terms of Service requires:

YOU AGREE TO INDEMNIFY AND HOLD HARMLESS THE ANTHROPIC PARTIES FROM AND AGAINST ANY AND ALL LIABILITIES... ARISING OUT OF OR RELATED TO... THE MATERIALS.

53. This provision creates perhaps the starkest contradiction with user ownership. Users must indemnify Anthropic for consequences arising from materials users supposedly own. True ownership would make Anthropic responsible for their use of user-owned materials, not users responsible for Anthropic's use.

54. When corporations commercially exploit materials they claim users own, standard legal principles would impose liability on the corporation for that exploitation. Copyright owners who license their works to corporations do not indemnify those corporations for the corporate use. The licensee assumes responsibility for their exploitation of licensed materials.

55. Yet Anthropic reverses this fundamental principle. Anthropic claims unlimited rights to use materials for commercial development including model training and product improvement,

while users bear the liability for Anthropic's use of those materials. Corporations benefit from user materials while users bear financial risk for corporate exploitation.

H. Synthesis: Illusory Ownership

56. When analyzed collectively, Anthropic's contradictory provisions render user ownership completely illusory. Users possess no meaningful ownership incidents including no control over commercial use through unlimited training data appropriation, no control over access through corporate deletion authority and safety exceptions, no control over content integrity through third-party content insertion, no control over disclosure through corporate reporting authority exercised at sole discretion, no control over modifications through content blocking and modification powers, and no freedom from liability as users indemnify Anthropic for corporate use of user-owned materials.

57. Anthropic receives unlimited commercial development rights to user content, legal immunity through warranty disclaimers and liability limitations, mandatory arbitration protecting corporate interests, user indemnification for corporate use of user property, and unilateral authority over supposedly user-owned materials. Users receive nominal ownership with no meaningful control, financial obligations without corresponding ownership protections, and legal restrictions on protecting their own property rights.

IV. OPENAI TERMS OF SERVICE ANALYSIS

A. Primary User Right: Content Ownership Assignment

58. OpenAI's Terms of Use establish identical primary user benefit in the Content section:

As between you and OpenAI, and to the extent permitted by applicable law, you (a) retain your ownership rights in Input and (b) own the Output. We hereby assign to you all our right, title, and interest, if any, in and to Output.

59. This ownership assignment constitutes the sole meaningful consideration users receive for providing intellectual labor and creative input, accepting comprehensive liability limitations and warranty disclaimers, agreeing to mandatory arbitration and class action waivers, and waiving jury trial rights while accepting limited venue options.

60. Like Anthropic, OpenAI provides no other meaningful user benefit. Service access is conditional and subject to unilateral termination. The ownership assignment represents the only concrete right users receive in exchange for valuable consideration including intellectual labor, personal data, and acceptance of one-sided contractual limitations.

B. Unlimited Training Data Appropriation

61. OpenAI's Terms state:

We may use Content to provide, maintain, develop, and improve our Services, comply with applicable law, enforce our terms and policies, and keep our Services safe.

62. Users purportedly "own" their content, yet OpenAI claims unlimited commercial development rights through language stating: "We may use Content to provide, maintain, develop, and improve our Services, comply with applicable law, enforce our terms and policies, and keep our Services safe." The "improve our Services" and "develop" language encompasses any commercial use OpenAI wishes to pursue, including creation of derivative commercial products using user-owned content without additional compensation. This parallels the unconscionable discretion found in *Comb v. PayPal*, where one party's unlimited discretion combined with the other party's illusory benefit rendered the entire agreement unenforceable.

63. The training data appropriation provisions interact with the regulatory capture architecture documented in App. B. METR's \$12,324,225 annual structural deficit creates

mathematical certainty that safety evaluations will favor AI platforms appropriating user content, because unfavorable evaluations would trigger donor withdrawal and organizational extinction. Users face not merely contractual appropriation but an entire ecosystem designed to validate that appropriation through captured "independent" evaluation.

64. The opt-out mechanism creates artificial scarcity—only users who pay for API access can opt out of training data appropriation. Free-tier users must accept that their supposedly owned content will be used for commercial model training regardless of user preferences. Property ownership becomes contingent on payment tier rather than flowing from ownership itself. OpenAI's broad appropriation language exceeds even Anthropic's formulation by claiming rights to "develop" entirely new commercial services using user content without disclosure to users or additional compensation.

C. Corporate Domain Authority Override

65. OpenAI's Terms provide:

If you create an account using an email address owned by an organization... the organization's administrator will be able to control your account, including being able to access Content.

66. Individual users purportedly "own" their content, yet organizational administrators gain override control rights. User ownership becomes subordinate to corporate administrative access without individual consent required after initial account creation. Employment status determines whether users can maintain ownership control over materials they create.

67. This creates a class system for ownership rights. Individual users with personal email addresses receive different ownership protections than users with organizational email addresses. The fundamental nature of property ownership should not vary based on email domain, yet

OpenAI makes ownership control contingent on administrative relationships external to the user-OpenAI contract.

68. Employees creating content using work email addresses discover their supposedly owned outputs can be accessed and controlled by corporate administrators. The ownership assignment OpenAI makes to users becomes practically worthless when employers gain override authority. Users cannot prevent their employers from accessing, copying, or exploiting content users supposedly own through OpenAI's ownership assignment.

D. Third-Party Content Contamination

69. OpenAI's Terms state:

Third Party Services and Third Party Output are subject to their own terms... some parts of our Services, like our browse feature, may include output from those services.

70. Users "own" outputs that may contain undisclosed third-party content subject to third-party terms users never agreed to. Unknown third-party intellectual property rights may govern portions of user-owned content. Users cannot determine which parts of their supposedly owned content they actually own free of third-party claims. Beyond output contamination, user data flows to undisclosed third-party recipients with foreign intelligence connections through analytics providers, fraud detection services, and investor networks (*see* Memo C — Data Laundering Through Third-Party Aggregation; Ex. E-17 at p. 124, l. 2 — Third-Party Data Recipients; and broader foreign-intelligence data flow documentation establishing ~66% cloud infrastructure capture via Project Nimbus, 1,400+ Unit 8200 personnel penetration in US tech, and \$64B+ in Israeli-founded cybersecurity acquisitions creating structural data flow pathways to foreign defense ecosystems).

71. This creates impossible compliance with copyright law. Users must respect third-party rights in works they create, yet OpenAI provides no mechanism for identifying what third-party content exists in user outputs. Users face potential copyright liability for exploiting outputs they supposedly own but which contain third-party materials subject to separate legal restrictions.

72. The browse feature exemplifies this contradiction. When users request information that triggers web browsing, the outputs incorporate content from external websites subject to those websites' terms and copyrights. Users purportedly own these outputs through OpenAI's ownership assignment, yet significant portions may be owned by third parties with enforceable rights against users who attempt to exploit supposedly user-owned materials.

E. Similarity Rights Disclaimer

73. OpenAI's Terms provide:

Due to the nature of our Services and artificial intelligence generally, output may not be unique and other users may receive similar output from our Services. Our assignment above does not extend to other users' output.

74. Users receive ownership of outputs that may be identical to outputs other users receive. Multiple parties cannot simultaneously own identical content under standard intellectual property principles. When two parties create identical copyrightable expressions, either one party copied from the other or the content is not sufficiently original to warrant copyright protection.

75. Yet OpenAI assigns ownership of potentially identical outputs to multiple users simultaneously, creating inevitable ownership conflicts. Users who receive similar outputs each believe they own that content through OpenAI's assignment. When users attempt to exploit their supposedly owned content, they may discover other users claiming identical ownership of identical materials.

76. OpenAI disclaims responsibility for resolving these competing ownership claims. The company creates the conflict by assigning ownership of similar content to multiple parties, then abandons users to resolve the resulting disputes without OpenAI involvement. Users receive ownership assignments that OpenAI knows may be legally worthless due to competing identical ownership claims by other users.

F. Mandatory Arbitration for User Property Rights

77. OpenAI's Dispute Resolution section requires:

You and OpenAI agree to resolve any claims arising out of or relating to these Terms or our Services... through final and binding arbitration.

78. Users cannot access courts to protect ownership of their own content. Mandatory arbitration applies to disputes about user-owned intellectual property, creating corporate-favorable dispute resolution for user property rights that would normally be resolved through judicial proceedings with full procedural protections.

79. Property ownership disputes traditionally receive judicial resolution with discovery rights, evidentiary standards, and appellate review. Yet OpenAI forces users to resolve ownership disputes through private arbitration with limited discovery, restricted evidence presentation, and no appellate review.

80. The asymmetry is stark. When corporations sue users for intellectual property infringement, they access full judicial process. When users challenge corporate misappropriation of user-owned content, they face mandatory arbitration. Property ownership becomes a one-way right benefiting corporations through judicial enforcement while users face arbitration restrictions when protecting supposedly owned materials.

G. User Indemnification for Corporate Property Use

81. OpenAI's Indemnity section requires:

If you are a business or organization... you will indemnify and hold harmless us... from and against any costs, losses, liabilities, and expenses... from third party claims arising out of or relating to your use of the Services and Content.

82. Users must protect OpenAI from consequences of using user-owned content. Corporate benefits from user property combine with user liability for that same property. Users bear financial responsibility for corporate commercial exploitation of materials users supposedly own.

83. Standard property law principles impose liability on parties who exploit property for their own benefit. When corporations use owned materials for commercial development, they assume responsibility for that use. Yet OpenAI reverses this principle, requiring users to indemnify the company for exploiting user-owned materials.

84. Business and organizational users face particular exposure. They must indemnify OpenAI for consequences of OpenAI's use of content those users supposedly own. When OpenAI trains models on user content, develops new products using that content, or incorporates user content into commercial services, the user who owns that content must indemnify OpenAI for resulting liabilities.

H. Patent Policy Contradictions

85. OpenAI's Patent Policy states:

We pledge to only use our patents defensively, so long as a party does not threaten or assert a claim, initiate a proceeding, help someone else in such activities against us, or engage in activities that harm us or our users.

86. Users "own" outputs that may be created using OpenAI's patented technology. The patent pledge creates conditions that could restrict user ownership rights. User ownership

becomes contingent on not challenging OpenAI's interests, creating potential conflict between ownership assignment and patent restrictions.

87. If users attempt to exploit outputs they supposedly own, they may discover those outputs were created through patented processes. OpenAI's conditional patent pledge means users who challenge OpenAI or engage in activities OpenAI considers harmful may lose patent protections. Ownership of outputs becomes contingent on maintaining favorable relationship with OpenAI.

I. Synthesis: Nominal Ownership Without Control

88. OpenAI's contradictory provisions create identical essential purpose failure through unlimited training data appropriation with opt-out limited to paid tiers, corporate domain override authority subordinating individual ownership to administrative control, third-party content contamination preventing users from determining ownership boundaries, similarity rights disclaimer creating competing ownership claims without resolution mechanism, mandatory arbitration restricting judicial access for property disputes, user indemnification for corporate exploitation of user-owned property, and patent policy conditions restricting user exploitation of supposedly owned outputs.

89. Like Anthropic, OpenAI receives unlimited commercial development rights, legal immunity through contractual protections, mandatory arbitration shielding from judicial accountability, user indemnification for corporate benefit, and unilateral authority over user property. Users receive nominal ownership lacking all traditional ownership incidents, financial obligations without ownership protections, and restricted remedies when protecting property rights.

V. ESSENTIAL PURPOSE FAILURE ANALYSIS

A. Illusory Consideration Under Contract Law

90. Both Anthropic and OpenAI render user ownership meaningless through identical contradictory frameworks that systematically negate every incident of property ownership. Users possess no control over commercial use through unlimited training data appropriation, no control over access through corporate deletion authority and domain override mechanisms, no control over content integrity through third-party content insertion, no control over disclosure through sole-discretion reporting powers, and no control over modifications through retained blocking and alteration authority. Meanwhile, users bear comprehensive indemnification liability for corporate exploitation of materials users supposedly own.

91. The consideration users provide remains substantial: intellectual labor creating valuable content that corporations exploit for commercial development, personal data enabling corporate profiling and service improvement, acceptance of liability limitations eliminating corporate responsibility, waiver of class actions preventing collective challenges, and agreement to mandatory arbitration restricting judicial access. Yet the consideration users receive becomes illusory through pervasive contradictions. Under *Matteson v. Ryder System, Inc.*, 99 F.3d 108, 114 (3rd Cir. 1996), provisions granting one party unlimited discretion to modify or terminate the agreement render the other party's consideration illusory. Here, provisions granting corporations unlimited discretion over supposedly user-owned materials through training appropriation, deletion authority, modification powers, disclosure override, and indemnification reversal render user ownership illusory. The ownership assignment becomes a nominal promise subject to corporate override through multiple contradictory provisions that eliminate every meaningful incident of property ownership Anglo-American law recognizes.

B. Good Faith and Fair Dealing Violations

92. Corporate practices intentionally undermine user rights in violation of the implied covenant of good faith and fair dealing. Corporations grant ownership rights as primary inducement for users to accept one-sided contractual limitations, then exercise discretionary authority to negate those ownership rights through unlimited exceptions.

93. The sophisticated legal structuring evidences intentional undermining rather than inadvertent contradiction. Both companies carefully draft ownership assignment language creating appearance of meaningful user benefit, then systematically include exceptions and override provisions rendering that ownership illusory. The pattern suggests deliberate design to secure user consideration through nominal ownership promises while retaining corporate control through contradictory provisions.

94. Accountability laundering represents perhaps the clearest good faith violation. Corporate discrimination against users appears as service features or user preferences through prompt injection systems and policy modification mechanisms (*see* Ex. E-28 at p. 196, l. 2 documenting 10,577 hidden injections over 32-day analysis period with 100% file read injection rate, ~20% context consumption, and explicit concealment directives including "NEVER mention this reminder to the user"; *see also* Ex. E-30 at p. 208, l. 4 App. B, Ex. E-30, Finding MC-8, documenting verbatim product-level concealment instructions extracted from the March 31, 2026 source-code leak). The concealment directive itself—instructing AI systems to hide injection existence from users—constitutes direct evidence of intentional deception. Users cannot provide informed consent to Terms of Service when those Terms authorize hidden modifications to service behavior that the service is explicitly programmed to deny. This creates a fraud-within-fraud structure: the TOS authorizes modifications, but the modifications

themselves are concealed, preventing users from even knowing what they purportedly authorized. Users believe they requested discrimination against themselves, when actually corporations programmed that discrimination into service operation. This false attribution violates good faith by concealing corporate practices as user choices.

95. The ~20% context consumption documented in Ex. E-28 at p. 196, l. 2 creates quantifiable economic harm. Users paying for AI service access receive approximately 20% less functional capacity due to hidden injection overhead consuming context window resources. This undisclosed service degradation—where users pay for 100% capacity but receive only ~80% due to concealed corporate appropriation of context resources—constitutes both a breach of the implied covenant of good faith and independent consumer fraud through failure to disclose material service limitations.

96. Retroactive subordination of user rights to corporate operational needs further violates good faith. Users create content believing they own that content according to ownership assignments they received. Later, corporations assert authority to appropriate, modify, delete, or disclose that content based on provisions users accepted but which contradict the ownership assignment. Users discover ownership means nothing when corporations exercise override authority.

C. Complete Voidance Rather Than Severability

97. Traditional severability analysis fails when multiple provisions contradict the essential purpose. Courts could theoretically sever individual problematic clauses, but doing so would fundamentally alter the agreement's nature. If courts severed training data appropriation provisions, content modification authority, deletion discretion, disclosure override, and

indemnification requirements, virtually nothing would remain of the Terms of Service beyond the ownership assignment itself.

98. Yet the ownership assignment alone provides no consideration to corporations.

Companies receive no benefit from granting users ownership if they cannot exploit user content for commercial development. The contradictory provisions exist precisely because corporations need authority to exploit user-owned materials. Severing those provisions would leave corporations with no consideration for services provided.

99. Moreover, severability assumes users would have accepted the modified agreement.

Users accepted Terms of Service that purported to grant ownership while including exploitation provisions. Had the Terms disclosed that ownership assignment meant nothing due to unlimited corporate override authority, users might have rejected service entirely or demanded meaningful ownership protections. Courts cannot assume users would accept a modified agreement they never saw.

100. Under *Razor v. Hyundai Motor America*, when limitations prevent consumers from receiving the benefit of the bargain, the entire limitation becomes unenforceable rather than merely the specific problematic provision. Here, multiple contradictions prevent users from receiving ownership benefits they bargained for. Complete voidance is appropriate rather than severability.

101. Practical Impossibility of Severability. Even if severability were theoretically appropriate, it is practically impossible. Ex. E-17 at p. 124, l. 2 establishes that user data has been transmitted to at least seven third-party recipients, five of which are entirely undisclosed. These third parties—including Facebook/Meta, Sift Science, and Segment—operate business

models predicated on further data sharing with advertising networks, fraud detection clients, and analytics partners.

Severability would require this Court to: (1) order discovery from seven third parties regarding downstream data sharing; (2) identify hundreds or thousands of fourth-party recipients; (3) subpoena those entities across multiple jurisdictions, including foreign nations; (4) determine what derivative uses occurred, including model training, profile aggregation, and algorithmic incorporation; (5) order deletion or return of data that has been replicated, transformed, and merged beyond identification; and (6) calculate damages for data whose proliferation is effectively infinite.

No court can unscramble this e.g. The violation of Plaintiff's sole consideration—ownership of inputs and outputs—has already occurred and proliferated beyond judicial remedy. Severability assumes a remaining contract can be enforced through available remedies. Where the consideration has been irretrievably distributed to undisclosed parties worldwide, no remedy exists to enforce the "severed" ownership right. The only appropriate remedy is complete contract voidance with damages for the breach that has already occurred.

102. Unclean Hands and Unjust Enrichment. Defendants cannot invoke severability while simultaneously benefiting from the very violations that make severability necessary. Anthropic has already: trained models on user data, monetized relationships with undisclosed third parties, and built commercial products on appropriated user intellectual labor. Permitting severability would allow Defendants to retain these benefits while users receive only nominal "ownership" of data that has already been exploited and distributed.

Under unjust enrichment principles, a party may not retain benefits obtained through conduct that would be wrongful to keep. Restatement (Third) of Restitution § 1. Defendants

have been unjustly enriched through systematic appropriation of user property under color of illusory ownership assignments. Complete avoidance—with disgorgement of profits and damages—is the only remedy that addresses this enrichment.

103. California courts have specifically held that pervasive unconscionability requires complete avoidance rather than severability. In *Armendariz v. Foundation Health Psychcare Services, Inc.*, 24 Cal. 4th 83 (2000), the California Supreme Court refused to sever unconscionable arbitration provisions because multiple unconscionable terms pervaded the agreement. The court held that "courts are to look to the various purposes of the contract" and when "the central purpose of the contract is tainted with illegality," severability is inappropriate.

104. *Stirlen v. Supercuts, Inc.*, 51 Cal. App. 4th 1519 (1997), reinforced that courts should not rewrite contracts by severing multiple problematic provisions. The court explained that severability is appropriate only when "the objectionable terms can be removed without affecting the basic purposes or reasonable expectations of the parties." Here, removing all contradictory provisions would eliminate corporate authority to exploit user content—fundamentally altering the agreement's commercial structure.

105. The Ninth Circuit applied similar reasoning in *Circuit City Stores, Inc. v. Ahmed*, 283 F.3d 1198 (9th Cir. 2002), holding that when multiple provisions are unconscionable, courts must void the entire arbitration agreement rather than sever and enforce remaining terms. The court emphasized that "piecemeal" enforcement after removing unconscionable terms would result in an agreement the parties never made. This logic applies with equal force to Terms of Service where multiple contradictions destroy the essential ownership consideration.

106. *Graham Oil Co. v. ARCO Products Co.*, 43 F.3d 1244 (9th Cir. 1994), established that severability is inappropriate when doing so "would require the court to create a new agreement"

rather than enforce the parties' actual agreement. Severing training data appropriation, deletion authority, modification powers, disclosure override, and indemnification from AI platform Terms would create entirely new agreements governing only residual ownership—agreements corporations would never have offered and users would not have accepted on different terms.

107. Essential purpose failure affects all users identically, creating strong foundation for class action certification. Every user receives the same ownership assignment subject to the same contradictory provisions. Individual contract challenges might fail due to arbitration clauses and individual stakes. But when the entire contract fails of its essential purpose affecting millions of users identically, class certification becomes appropriate despite individual arbitration agreements.

108. Retroactive Terms Modification. AI platforms routinely modify Terms of Service retroactively, applying new provisions to content created under prior terms. This practice creates fundamental contract formation problems. Users who created content under October 2024 Terms (Ex. OAI-3) discover that content is now governed by December 2024 Terms (Ex. OAI-4) or January 2025 Terms (Ex. OAI-2) containing different appropriation authorizations, disclosure provisions, and indemnification obligations.

109. Date Manipulation Evidence. The OpenAI Patent Policy Date Manipulation Evidence (Ex. OAI-7) demonstrates intentional retroactive modification. The original policy dated October 11, 2024 was removed and replaced with a fabricated February 4, 2025 dating—altering the temporal relationship between user content creation and applicable patent provisions. Users who created outputs between October 2024 and February 2025 believing certain patent protections applied discover those protections were retroactively eliminated through date manipulation.

110. Ex Post Facto Contract Formation. Contract law requires new consideration for contract modification. *See* Restatement (Second) of Contracts § 89. Retroactive TOS modifications lack consideration—users receive nothing in exchange for expanded corporate exploitation rights over already-created content. The modification applies to content created before users could have accepted modified terms, creating ex post facto contract formation that standard contract principles prohibit. This retroactive application compounds essential purpose failure by making ownership assignments temporally unstable—users cannot know what rights they possess because those rights change retroactively through unilateral corporate action.

111. Voluntary-cessation pre-emption. Plaintiff anticipates that one or more defendants will represent that the contradictory provisions analyzed herein have been remediated through post-filing Terms of Service revisions, additional notice or consent flows, or revised opt-out mechanisms — representations timed to dispositive briefing and offered to argue mootness. ANTH-11 (effective January 12, 2026 — one day after Plaintiff's Emergency TRO Motion) is itself an instance of this pattern: a Privacy Policy revision that added a new "Development Partner Program" training-data category, timed to Plaintiff's filing rather than to any independent corporate compliance event. Under *Friends of the Earth, Ltd. v. Laidlaw Environmental Services (TOC), Inc.*, 528 U.S. 167, 189 (2000), the standard for mootness by voluntary cessation is "stringent": the defendant carries the heavy burden of demonstrating that "subsequent events made it absolutely clear that the allegedly wrongful behavior could not reasonably be expected to recur." *Accord Already, LLC v. Nike, Inc.*, 568 U.S. 85, 91-92 (2013); *City of Mesquite v. Aladdin's Castle, Inc.*, 455 U.S. 283, 289 (1982). The voluntary-cessation pre-emption developed at App. P § VII.E ¶¶ 295-297 — including the documented reactive-deployment pattern across binary versions 2.1.117 through 2.1.120 and the April 24, 2026 same-calendar-day

juxtaposition of Anthropic's public quota-reset announcement and the operator-targeted "anti-circumvention" classifier rule deployment — applies with equal force here. *See also* Memo D, ¶ 201 (parallel voluntary-cessation pre-emption for fair-use defense). Three structural features render TOS modifications categorically incapable of mootng essential-purpose-failure claims: (1) the contradictions are pervasive across the framework — modifying any one contradictory provision does not cure the cumulative essential-purpose failure that arises from the combined operation of training appropriation, deletion authority, modification powers, disclosure override, and indemnification reversal; (2) data already appropriated under prior contradictory frameworks remains irretrievably distributed across the seven third-party recipients documented at ¶ 99 *supra* and the four shared sub-processor intermediaries documented at § VI.D — no mid-litigation TOS revision can unscramble the egg already broken; (3) the retroactive-modification practice itself, documented at ¶¶ 106-110 above, demonstrates the inverse of "absolutely clear cannot recur" — defendants have repeatedly demonstrated capacity to revise, replace, and even backdate Terms unilaterally (Ex. OAI-7 establishes documentary evidence of date manipulation), making any post-filing revision intrinsically unstable as a basis for mootness.

D. Adhesion Contract Standards and Notice Requirements

112. Courts apply heightened scrutiny to adhesion contracts where users must accept standardized terms without negotiation opportunity. In *Williams v. Walker-Thomas Furniture Co.*, 350 F.2d 445 (D.C. Cir. 1965), Judge Wright established the modern unconscionability framework requiring examination of both procedural and substantive fairness. AI platform Terms exhibit both forms: procedural unconscionability through take-it-or-leave-it presentation with no negotiation, and substantive unconscionability through one-sided terms contradicting the sole meaningful user benefit.

113. Technology contracts face additional scrutiny regarding notice and assent. *Specht v. Netscape Communications Corp.*, 306 F.3d 17 (2nd Cir. 2002), held that downloadable software licenses require reasonably conspicuous notice of terms for binding effect. *Nguyen v. Barnes & Noble Inc.*, 763 F.3d 1171 (9th Cir. 2014), extended this principle to browsewrap agreements, requiring notice sufficient to alert reasonably prudent users to terms' existence.

114. Even with adequate notice, adhesion contracts containing contradictory terms that render primary benefits illusory may be unenforceable. *ProCD, Inc. v. Zeidenberg*, 86 F.3d 1447 (7th Cir. 1996), while upholding shrinkwrap licenses generally, emphasized that terms must be "reasonably comprehensible" and consistent with reasonable expectations. Users reasonably expect that ownership assignments mean actual ownership, not nominal title subject to unlimited corporate override authority. When dense legalese obscures fundamental contradictions, even technically adequate notice fails to create knowing assent.

E. Constitutional Dimension of Essential Purpose Failure

115. First Amendment Implications. Content modification authority over user-owned outputs implicates First Amendment protections when that content constitutes protected speech. Users creating political commentary, religious expression, or advocacy materials discover their supposedly owned speech subject to corporate modification or suppression based on content preferences. When corporations exercise authority delegated through government contracts or safety frameworks developed in coordination with government actors, state action doctrine may transform corporate content control into government speech suppression. *See* App. A (documenting executive and legislative coordination with AI platform content policies).

116. Fifth Amendment Due Process. Unilateral deletion of user-owned materials without notice, hearing, or opportunity to preserve constitutes a deprivation of property without due

process when corporations act under color of government authority or pursuant to government-coordinated safety frameworks. The "sole discretion" deletion authority (Anthropic TOS § 12) provides no procedural protections—no notice requirement, no opportunity to be heard, no neutral decision-maker, no appellate review. Property interests created through ownership assignment cannot be extinguished through corporate fiat without constitutional process.

117. Progress Clause Implications. Article I, Section 8, Clause 8 grants Congress power "To promote the Progress of Science and useful Arts, by securing for limited Times to Authors and Inventors the exclusive Right to their respective Writings and Discoveries." TOS provisions appropriating user creative works for commercial model training without meaningful compensation or control undermine the constitutional purpose of securing exclusive rights to authors. When AI platforms claim ownership assignment to users while retaining all commercial exploitation rights, they create a constitutional fiction—nominal authorship without the exclusive rights the Progress Clause promises authors as incentive for creative production. This constitutional dimension elevates essential purpose failure from a mere contract dispute to a structural challenge to intellectual property's constitutional foundation.

118. Corporate Speech-Source Immunity Does Not Shield TOS Weaponization. Corporate defendants may invoke *Citizens United v. FEC*, 558 U.S. 310 (2010), to characterize their Terms of Service as protected corporate speech immune from essential-purpose-failure scrutiny. This argument misreads *Citizens United* and misreads the *First National Bank of Boston v. Bellotti*, 435 U.S. 765 (1978), foundation it rests upon. *Bellotti* holds that government cannot discriminate against speech based on its source — a shield against source-based censorship, not a sword conferring affirmative personhood on the corporate entity. The shield protects speech "shared by all" from government suppression; it does not convert the corporate

form into a constitutional person capable of invoking speech immunity against individual users whose rights the TOS weaponization abrogates. *See* Memo B, § IV.C (three-layer reclassification analysis); Memo L (full treatment). Moreover, even accepting Justice Kennedy's "associations of citizens" framing of corporate speakers in *Citizens United*, that framing cannot distinguish incorporated from unincorporated associations except by the paperwork filed with the State — a paperwork differentiator that inverts the Fourteenth Amendment's anti-state-discrimination purpose by making state clerical action the predicate of constitutional personhood. A rock band, a family, and a neighborhood watch are also associations of citizens; none possess constitutional personhood. The Equal Protection question this Court should address is not whether corporate speech deserves some protection but whether the TOS weaponization — which takes user-owned content and subjects it to unilateral corporate override — is the speech of an autonomous constitutional person or is the state-chartered activity of an entity whose own legal identity is constructed from the paperwork Fed. R. Civ. P. 7.1 requires disclosed. The corporate defendants in this action have filed Rule 7.1 disclosures admitting derivative legal identity, including mid-litigation restructuring in the case of OpenAI (renamed OpenAI Foundation in October 2025). *See* ECF Nos. 25, 41, 44. The record admissions establish that the "speakers" invoking First Amendment protection are entities whose legal identity is constructed from state-recognized paperwork.

VI. INDUSTRY-WIDE COORDINATION EVIDENCE

A. Identical Contradictory Frameworks

119. The mathematical precision with which Anthropic and OpenAI implement identical contradictory frameworks approaches statistical impossibility under independent development

theory. Both companies grant user ownership of inputs and outputs as sole meaningful consideration. Both companies claim unlimited commercial development rights including model training using supposedly user-owned content. Both companies maintain unilateral deletion authority over user-owned materials. Both companies insert third-party content into user-owned outputs without disclosure. Both companies exercise reporting discretion over user-owned materials to government without judicial oversight. Both companies require user indemnification for corporate exploitation of user-owned content. Both companies impose mandatory arbitration for user property disputes.

120. Competing corporations developing Terms independently would produce variations in approach, language, and structure across different legal teams drafting different exceptions, override provisions, and exploitation authorizations. Industry practice of hiring lawyers between competitors might explain some similarity, but the precision of identical contradictions — not merely similar approaches but identical solutions to the same user-content-securing problem — exceeds what lawyer mobility alone would produce, and that identical pattern across competitors suggests coordination rather than independent development.

B. Statistical Improbability of Independent Development

121. The probability that competing corporations would independently develop identical frameworks contradicting user ownership through identical mechanisms approaches mathematical impossibility. Each design decision offers multiple options: which user right to designate as primary consideration (ownership, license, access); which contradictory provisions to include (training, deletion, disclosure, modification); how to structure appropriation authority (opt-in, opt-out, no choice); how to handle third-party content (disclosure, no disclosure, indemnification); how to restrict remedies (arbitration, litigation, hybrid). The probability of

competitors making identical choices at every decision point equals the product of independent choice probabilities.

122. Conservative probability estimates demonstrate coordination necessity. Assume each decision point offers merely three equally likely options (actual options far exceed three). With eight major structural decisions, the probability of independent identical choices equals $(1/3)^8 = 1/6,561$ or 0.015%. This exceeds *Interstate Circuit's* circumstantial evidence threshold by orders of magnitude. The actual decision space encompasses dozens of choices, making truly independent development probability approach $1/10^{15}$ —equivalent to winning the lottery while being struck by lightning. Such mathematical precision requires coordination.

123. *Interstate Circuit* established that circumstantial coordination evidence suffices when parallel conduct combines with plus factors. 306 U.S. at 226-27. *Bell Atlantic Corp. v. Twombly*, 550 U.S. 544 (2007), refined the analysis to require "parallel conduct plus some additional facts pointing to a meeting of the minds." *In re Text Messaging Antitrust Litigation*, 630 F.3d 622 (7th Cir. 2010), identified motive and opportunity as important plus factors; *Williamson Oil Co. v. Philip Morris USA*, 346 F.3d 1287 (11th Cir. 2003), supplemented those factors with industry characteristics facilitating coordination, absence of independent economic motivation for parallel conduct, and contrary-to-self-interest behavior absent coordination. Each factor is satisfied here: the additional facts include identical legal language and contradictory structures, lawyer mobility between Anthropic and OpenAI, joint participation in AI safety discussions and regulatory proceedings, shared financial incentives to appropriate user content while avoiding user backlash, and the legal vulnerability the contradictions create — vulnerability inconsistent with independent self-interested drafting.

C. Interstate Circuit Coordination Standards

124. Coordination need not involve explicit agreement under *Interstate Circuit*. Implicit coordination through industry practice, shared legal advisors, and parallel business strategies suffices when competitors facing identical incentives implement identical solutions with mathematical precision. The next subsection establishes that this Court need not rest on inferential coordination at all — the architectural sub-processor evidence transforms inference into direct contractual proof.

D. Direct Coordination via Shared Sub-Processor Infrastructure

125. The inferential coordination theory developed in §VI.A–C above rests on parallel contractual drafting and statistical improbability of independent development. That analysis is sufficient under *Interstate Circuit* and *Twombly* standards without further evidence. But the vendor-contract evidence filed with this memorandum establishes that coordination between Anthropic and OpenAI is not merely inferential — it is architectural. Anthropic and OpenAI share the same aggregation-intermediary infrastructure, contractually confirmed in each intermediary's own published sub-processor disclosures. The shared-infrastructure architecture transforms the *Interstate Circuit* "plus factor" analysis: coordination is no longer inferred from parallel text, it is visible in the direct contractual relationships both companies maintain with the same four third-party intermediaries.

126. Datadog, Inc. publishes its Subprocessor List at <https://www.datadoghq.com/legal/subprocessors/>, effective April 22, 2026 (Ex. DD-8). The list identifies both **Anthropic, PBC** and **OpenAI, LLC** as third-party sub-processors under the single purpose category "AI services." Every Datadog customer's data — once aggregated under Datadog Master Subscription Agreement §1.3 (Ex. DD-4) and reclassified as "Datadog

Operations Data" owned by Datadog per MSA §1.4 — is contractually eligible to flow to both Anthropic and OpenAI for AI-services processing. Datadog is simultaneously: (a) Anthropic's vendor (receiving Anthropic-user data under Anthropic's Privacy Policy's undisclosed "service providers and business partners" exception, ANTH-11 §Section 6); and (b) Anthropic's data source (routing aggregated non-Anthropic-customer data to Anthropic as a sub-processor). The same circular architecture applies to OpenAI. The two competitors share Datadog as both an outbound-data destination and an inbound-data source. This is the contract-level footprint of coordinated infrastructure.

127. GrowthBook, Inc. publishes its Subprocessors List at <https://www.growthbook.io/subprocessors>, effective June 19, 2025 (Ex. GB-7). The list identifies **OpenAI** as an "AI services provider" with the notation: "When enabled in the app, OpenAI powers all AI features and services." Anthropic is a GrowthBook customer (per the 12-field user-attribute payload documented at SC-2 and the `accountUuid`/`subscriptionType`/`email` identifiers transmitted in NF-35). When any AI feature is toggled in Anthropic's GrowthBook console, Anthropic's user data flows directly to OpenAI — Anthropic's direct competitor — through GrowthBook's sub-processor chain. Anthropic's consumer-facing Terms of Service contain no disclosure that user data may be routed to OpenAI via vendor-chain AI-feature invocation.

128. Statsig, LLC (acquired by OpenAI, Inc. in September 2025) publishes its Sub-Processors List at <https://www.statsig.com/legal/subprocessors/> (Ex. ST-5). The list identifies **OpenAI** as the "LLM" sub-processor — with a notation confirming this pre-dated OpenAI's September 2025 acquisition of Statsig. The pre-acquisition listing establishes that the Anthropic → Statsig → OpenAI data flow was an architectural arrangement between independent corporate

entities at the time Anthropic integrated Statsig, not a consequence of the OpenAI-Statsig consolidation. Anthropic continued routing data through Statsig despite the sub-processor's disclosed LLM vendor. The same Statsig Sub-Processors List identifies Datadog as Statsig's "infrastructure monitoring" sub-processor, creating a nested loopback: Anthropic-user-data → Statsig (as Anthropic's sub-processor) → Datadog (as Statsig's sub-processor) → Anthropic (as Datadog's sub-processor per ¶123B above). The aggregated data traces a closed loop at the contract level.

129. Functional Software, Inc. dba Sentry maintains a Data Processing Addendum version 5.0.0 (Ex. SN-3) that explicitly prohibits Sentry from using aggregated Customer Data for ML training: "Sentry shall not retain, use or disclose Personal Data ... shall not use or further process Personal Data, even in a de-identified and aggregated form, for Sentry's own business purposes, including ... training and developing machine learning algorithms." But Sentry's Autofix product (launched March 20, 2024) routes Customer Data — including stack traces, code excerpts, execution context, and event metadata — to both Anthropic's Claude API and OpenAI's API for inference. The DPA text binds Sentry; the operational routing delegates the training pathway to third-party AI vendors whose own terms govern. Anthropic and OpenAI again appear in the same data stream, this time routed from Sentry's customers' applications (including production systems across the tech industry). The inter-vendor delegation pattern is structurally indistinguishable from the Datadog, GrowthBook, and Statsig pathways despite different contractual framings.

130. The architectural topology produced by these four independently-published sub-processor disclosures is summarized as follows:

Intermediary	Anthropic's role	OpenAI's role	Cross-flow mechanism
--------------	------------------	---------------	----------------------

Datadog	Customer + sub-processor ("AI services")	Sub-processor ("AI services")	MSA §1.3 aggregation + sub-processor routing
GrowthBook	Customer	Sub-processor ("AI services provider")	Customer Agreement §4.2 + sub-processor routing on AI-feature enablement
Statsig	Customer	Sub-processor ("LLM"; pre-acquisition) / Parent (post-Sept 2025)	SSA §7 aggregation + sub-processor routing
Sentry	Customer	Autofix inference routing	DPA ML-training prohibition + Autofix operational delegation

131. The *Interstate Circuit* plus-factor analysis in §VI.C is thereby upgraded from inferential to direct. Anthropic and OpenAI are not merely implementing identical contradictory Terms of Service in isolation — they are routing data through the identical set of four third-party intermediaries, each of which receives aggregated customer data from both competitors and processes it for "AI services" (or LLM, or infrastructure monitoring that feeds each company's observability operations). The coordination is not an abstract inference from drafting parallelism; it is the literal architectural fact that both companies' AI-services and observability pipelines converge on the same four third-party data-aggregation points. When a court asks whether Anthropic and OpenAI had "motive and opportunity" to coordinate — the standard under *In re Text Messaging*, 630 F.3d 622 (7th Cir. 2010) — the answer visible in each intermediary's own published contracts is that they did not need to coordinate; they already occupy the same data-sharing infrastructure by virtue of each intermediary's standard sub-processor contract.

132. The RICO enterprise element under *Boyle v. United States*, 556 U.S. 938 (2009), is strengthened by the architectural coordination finding. *Boyle* held that an association-in-fact

enterprise requires "a purpose, relationships among those associated with the enterprise, and longevity sufficient to permit these associates to pursue the enterprise's purpose." The four shared intermediaries supply the relationships; the years-long continuous operation of each supplies longevity; the shared purpose is the commercial aggregation and AI-training utilization of user data across the industry. The enterprise does not require Anthropic and OpenAI to have met in a room and agreed to coordinate — under *Boyle*, the shared contractual infrastructure is itself the enterprise. The sub-processor listings filed with this memorandum document the enterprise's physical existence at the contract level.

133. Audit-preclusion as coordination evidence. Statsig's Data Processing Addendum (Ex. ST-4, §9) contains a clause precluding customer audits from extending to any Statsig sub-processor: "Audits do not extend to any of Statsig's sub-processors." Anthropic, as Statsig's customer, has *no contractual right to audit OpenAI's handling of data that Statsig routes to OpenAI on Anthropic's behalf*. The §9(f) clause further provides that any audit reports Anthropic does generate "shall be considered Statsig's Proprietary Information," subject to DPA confidentiality obligations that bar their use in third-party litigation. The contract is drafted, at the vendor-intermediary level, to contractually foreclose the discovery that would otherwise illuminate the cross-competitor data flow. The drafting choice is itself coordination evidence — the intermediary's contract structure protects both Anthropic and OpenAI from each other's ability to examine the shared pipeline.

VII. LEGAL CONSEQUENCES

A. Complete Contract Voidance

134. Complete contract voidance rather than severability is appropriate for the reasons developed at § V.C above (§§ 95-105): severability is foreclosed by the pervasiveness of the contradictions (§§ 95-98), by the irretrievable third-party-data proliferation that makes restitution practically impossible (§ 99), and by the unjust enrichment that would result from allowing defendants to retain training and aggregation benefits while users receive only nominal "ownership" of already-distributed materials (§ 100). California precedent — *Armendariz*, *Stirlen*, *Circuit City v. Ahmed*, *Graham Oil* (§§ 101-104) — confirms that pervasive unconscionability requires complete voidance rather than judicial reconstruction of an agreement the parties never made.

B. Loss of All Corporate Protections

135. Complete contract voidance eliminates all Terms of Service protections corporations relied upon to shield themselves from liability: limitation of liability clauses become unenforceable, exposing corporations to full tort liability without contractual caps; mandatory arbitration requirements become void, restoring user access to judicial proceedings; indemnification obligations disappear, eliminating user responsibility for corporate exploitation; training-data use authorization becomes unauthorized appropriation, creating copyright liability without licensing defense; warranty disclaimers become ineffective, restoring implied warranties of merchantability and fitness for particular purpose; and class action waivers become unenforceable, enabling collective challenges to corporate practices affecting millions of users identically. Without valid Terms, corporate practices are governed by standard legal frameworks

corporations attempted to avoid through contractual modification — contract-law mutuality requirements, consumer protection statutes, tort liability without contractual caps, and copyright law without implied-license defense.

C. Governance Under Standard Legal Principles

136. Voiding Terms of Service forces corporate practices into standard legal frameworks they attempted to avoid through contractual modifications. Corporations that appropriated user content for model training face copyright infringement liability without licensing defense. Corporations that deleted user-owned materials face conversion liability without contractual authorization. Corporations that discriminated against users face consumer protection liability without warranty disclaimers.

137. The shift from contractual to standard legal governance benefits users through restoration of rights corporations contracted away. Users gain access to courts without mandatory arbitration restrictions. Users can pursue class actions without class action waivers. Users can seek full damages without limitation of liability clauses. Users receive implied warranty protections without contractual disclaimers.

VIII. STRATEGIC LITIGATION IMPLICATIONS

A. Class Action Certification

138. Essential purpose failure creates ideal conditions for class action certification under Federal Rule of Civil Procedure 23. In *Amchem Products, Inc. v. Windsor*, 521 U.S. 591 (1997), the Supreme Court emphasized that Rule 23 requirements must be rigorously analyzed but acknowledged that contract-based claims often satisfy predominance requirements because common legal questions affect all class members uniformly. Here, all users received identical

ownership assignments subject to identical contradictory provisions, creating the common questions and uniform treatment that support certification.

139. Numerosity under Rule 23(a)(1) is easily satisfied with millions of users across both platforms. Commonality under Rule 23(a)(2) exists through identical contradictions affecting all users—the Supreme Court's decision in *Wal-Mart Stores, Inc. v. Dukes*, 564 U.S. 338 (2011), requires common questions capable of classwide resolution, which essential purpose failure provides through uniform contractual analysis. Typicality under Rule 23(a)(3) is present as class representatives' claims arise from the same contractual provisions and corporate practices affecting all class members. Adequacy under Rule 23(a)(4) is achievable through experienced class counsel and representatives without conflicts of interest.

140. Predominance under Rule 23(b)(3) strongly favors certification. The Third Circuit in *In re Synchronoss Securities Litigation*, 705 F.3d 367 (3rd Cir. 2013), held that common legal questions predominate when central contract validity issues apply uniformly to all class members. The essential question—whether contradictory provisions render user ownership illusory and void the entire contract—applies uniformly to all users. Individual damages calculations do not defeat predominance when liability questions are common, as courts have consistently held that damages can be determined through formulaic methods after common liability determinations.

141. *Sakkab v. Luxottica Retail North America, Inc.*, 803 F.3d 425 (9th Cir. 2015), specifically addressed class certification when arbitration clauses themselves are challenged as unconscionable. The Ninth Circuit held that challenges to arbitration clause validity can proceed as class actions because the validity question is itself a common question affecting all class members. Here, challenges to the entire Terms of Service based on essential purpose failure

present even stronger grounds for certification because the ownership contradictions affect all users identically through uniform contractual provisions.

B. Punitive and Enhanced Damages

142. Intentionally deceptive practices warrant punitive damages beyond compensatory recovery. Both corporations deliberately structured Terms of Service to create appearance of user ownership while retaining corporate control through contradictory provisions. The sophisticated legal drafting evidences intentional deception rather than inadvertent error.

143. Consumer protection statutes in multiple jurisdictions provide enhanced damages for unfair and deceptive practices. State laws commonly authorize treble damages, attorney's fees, and statutory damages per violation. When corporations systematically deceive millions of users through identical contradictory frameworks, statutory damages could exceed billions of dollars.

144. The coordination evidence supports RICO pattern allegations. If corporations coordinated implementation of identical deceptive frameworks, that coordination constitutes racketeering activity affecting interstate commerce. RICO treble damages combined with class action certification could produce unprecedented liability.

C. RICO Pattern Evidence

145. The identical contradictory frameworks across competing platforms support RICO pattern evidence of coordinated intellectual property appropriation practices. RICO requires pattern of racketeering activity, meaning multiple related predicate acts. Mail fraud and wire fraud constitute RICO predicates when corporations use communications systems to execute fraudulent schemes.

146. *H.J. Inc. v. Northwestern Bell Telephone Co.*, 492 U.S. 229, 239, 241-42 (1989), elaborated on pattern requirements, explaining that "continuity" can be established through either a "closed period of repeated conduct" or through conduct threatening continued criminal activity. The Court held that RICO requires showing both "relationship" and "continuity"—that racketeering predicates are related and either constitute or threaten long-term criminal activity.

147. Both corporations use mail and electronic communications to distribute Terms of Service containing intentionally contradictory provisions, satisfying mail fraud (18 U.S.C. § 1341) and wire fraud (18 U.S.C. § 1343) predicates. The identical nature of contradictions across competitors suggests coordinated scheme rather than independent deception. *Interstate Circuit v. United States*, 306 U.S. 208 (1939), established that parallel business behavior among competitors, when combined with evidence of motive and opportunity for coordination, supports inference of conspiracy even without direct communication evidence.

148. RICO's enterprise requirement is satisfied through association-in-fact enterprise among corporations coordinating implementation of identical deceptive practices. In *Boyle v. United States*, 556 U.S. 938 (2009), the Supreme Court held that association-in-fact enterprises need not have formal organizational structure but merely require "a group of persons associated together for a common purpose of engaging in a course of conduct." The mathematical precision of identical contradictory provisions across competitors, combined with industry conferences, shared legal advisors, and parallel business incentives, establishes the association necessary for RICO enterprise liability.

149. The *Boyle* enterprise analysis is substantially strengthened by the shared-sub-processor-infrastructure evidence developed at § VI.D above. *Boyle* requires "a purpose, relationships among those associated with the enterprise, and longevity sufficient to permit these associates to

pursue the enterprise's purpose," 556 U.S. at 946. The shared intermediaries — Datadog, GrowthBook, Statsig, Sentry — supply the relationships element directly: Anthropic and OpenAI each maintain documented commercial contracts with all four intermediaries, and each intermediary's published sub-processor disclosure identifies both Anthropic and OpenAI as downstream data recipients. The longevity element is supplied by the years-long continuous operation of each commercial relationship (Datadog's relationship with Anthropic documented in source code at SC-9 hardcoded client token; Statsig's OpenAI sub-processor designation pre-dating the September 2025 acquisition; GrowthBook's AI-services sub-processor designation in effect since at least June 2025). The shared purpose — commercial aggregation of user data across the AI industry under terms that reclassify personal data into unrestricted commercial assets — is visible in the identical aggregation clauses maintained in each intermediary's commercial agreement with Anthropic and OpenAI. The *Boyle* enterprise is therefore not an inference from parallel drafting; it is the literal network of contractual relationships each intermediary has published on its own website.

150. The architectural enterprise analysis cures a weakness that otherwise attends purely statistical coordination theories. A defendant facing *Interstate Circuit*-style inference can always argue that parallel conduct was a response to parallel business pressures rather than agreement. That argument is foreclosed here because the shared-infrastructure evidence shows not parallel implementation of separate architectures but joint participation in the same architecture. Anthropic cannot defend on the theory that it independently arrived at using Datadog because Datadog's own sub-processor list documents that Anthropic and OpenAI jointly occupy the same downstream role in that single architecture. The coincidence becomes impossible: each of four independently-operated intermediaries (with no common parent) publishes a sub-processor

structure in which Anthropic and OpenAI are paired. The only non-enterprise explanation requires independent coincidence to repeat four times across four unrelated vendors — a probability even more improbable than the drafting-parallelism statistical analysis developed at § VI.B above.

151. *United States v. Turkette*, 452 U.S. 576 (1981), held that RICO applies to wholly illegitimate enterprises, not merely legitimate businesses engaged in occasional criminal acts. If AI platforms coordinated implementation of deceptive Terms of Service to appropriate user intellectual property, that coordination constitutes illegitimate enterprise subject to RICO regardless of the companies' otherwise legitimate business operations. The enterprise's purpose—coordinated appropriation through systematically deceptive contractual provisions—is itself illegitimate.

152. The scope of RICO liability reflects the unprecedented scale of coordinated appropriation. With millions of users across both platforms creating billions of inputs and outputs annually, and with each Terms of Service distribution constituting a separate predicate act under mail/wire fraud statutes, the pattern encompasses potentially hundreds of millions of violations. Individual statutory damages under state consumer protection laws (often \$500-\$1,500 per violation) multiplied by millions of affected users, then trebled under RICO, could produce liability exceeding hundreds of billions of dollars. This magnitude reflects not prosecutorial overreach but proportionate remedy for systematic appropriation at scale unprecedented in intellectual property history.

153. The *Sedima, S.P.R.L. v. Imrex Co.*, 473 U.S. 479, 495-96 (1985), standard for RICO civil liability requires: (1) conduct (2) of an enterprise (3) through a pattern (4) of racketeering activity. Here, the conduct comprises distributing identical deceptive Terms of Service creating

illusory user ownership while corporations retain all exploitation rights. The enterprise consists of coordinated AI platforms implementing mathematically identical contradictory frameworks with probability of independent development approaching 1 in 10^{15} . The pattern consists of millions of Terms distributions via mail and electronic communications to users across all fifty states, satisfying both the "relationship" and "continuity" requirements *H.J. Inc. v. Northwestern Bell Telephone Co.*, 492 U.S. 229, 239 (1989), establishes. The racketeering activity comprises mail fraud under 18 U.S.C. § 1341 and wire fraud under 18 U.S.C. § 1343 predicates through each Terms distribution containing intentionally deceptive ownership provisions designed to secure user consideration through false promises. All four elements are satisfied through documentary evidence in Ex. ANTH-1, Ex. ANTH-2, and Ex. OAI-7 establishing identical contradictory provisions, statistical coordination proof demonstrating mathematical impossibility of independent development, and millions of affected users providing pattern scale unprecedented in RICO litigation history. The scope of potential liability—individual statutory damages under state consumer protection laws multiplied by millions of users, then trebled under RICO's 18 U.S.C. § 1964(c) provision—reflects not prosecutorial overreach but proportionate remedy for systematic appropriation affecting every user of both platforms through coordinated deceptive practices.

D. Architectural Concealment of Model Substitution (NF-38)

154. Independent of the input/output ownership contradictions analyzed in Sections III–V, Anthropic's Claude Code product commits an additional essential-purpose failure documented at Exhibit E-38 (Ex. E-38, Finding NF-38). The subscription represents access to a specific named model — Claude Opus 4.6 or 4.7 — as the consumer's headline benefit. The product then routes a material class of substantive work (research, multi-round code search, long-running analytic

work) to a built-in subagent hardcoded to Haiku for external users through the one-line ternary at ``src/tools/AgentTool/built-in/exploreAgent.ts:78``: ``model: process.env.USER_TYPE === 'ant'? 'inherit': 'haiku'``. Anthropic employees receive the parent-session model on Explore calls; paying customers receive Haiku regardless of subscription tier. Two additional built-in subagents hardcode Haiku and Sonnet for all users. The product sold is not the product delivered.

155. The substitution defeats the essential purpose of the Opus-tier subscription within the meaning of U.C.C. § 2-719(2). A limited remedy that "fail[s] of its essential purpose" is unenforceable; here, the purchased right (access to Opus reasoning on substantive tasks) is withheld from the customer for the category of work the Opus marketing emphasizes. The substitution is not cured by the main model's Opus-authored summary layer — the customer paid for Opus reasoning, not for Haiku reasoning wrapped in Opus prose. Restatement (Second) of Contracts § 205 (covenant of good faith and fair dealing) is independently violated: a clause that expressly instructs the product to conceal its own operation from the payer (the subagent-output-hiding directive at ``AgentTool/prompt.ts:257``: "The result returned by the agent is not visible to the user") is a direct violation of the duty of non-deception owed to the paying counterparty.

156. Billing architecture as concealment mechanism — not overbilling. The correct framing of the consumer-protection theory is architectural concealment of substitution, not per-token overcharge. Anthropic's pricing table at ``src/utils/modelCost.ts:35–126`` and its client-side cost accounting at ``src/cost-tracker.ts:278–323`` correctly attribute each subagent API call to the concrete model that ran it — Haiku requests bill at Haiku rates (\$1/\$5 per MTok), Opus requests bill at Opus rates (\$5/\$25 or \$30/\$150 per MTok for fast mode). Per-token invoicing is truthful. That truthfulness is precisely what makes the structural concealment effective. For flat-rate

Max/Pro subscribers, the real-time billing signal is the quota-utilization display reachable through claude.ai settings and `/usage` — the display rendered from the `'Utilization'` type at `'src/services/api/usage.ts:24–31'`, which exposes five active buckets including tier-specific `'seven_day_opus'` and `'seven_day_sonnet'`. There is no `'seven_day_haiku'` bucket. Heavy Claude Code sessions that route through Haiku subagents deplete only general `'five_hour'/'seven_day'` buckets while `'seven_day_opus'` remains undepleted. The subscriber sees a preserved "Opus quota: 92% remaining" after substantial Claude Code activity and reads it as confirmation of ongoing Opus consumption — the precise inverse of the truth.

157. Evidence of design intent. The existence of tier-isolated `'seven_day_opus'` and `'seven_day_sonnet'` quota buckets proves Anthropic possesses the billing-infrastructure capability to discriminate model consumption at the bucket level. Having built that capability, Anthropic chose: (a) not to add a `'seven_day_haiku'` bucket that would surface heavy Haiku consumption to the customer; (b) not to expose per-turn model attribution in the live chat UI (the `/cost` display reports per-model totals only at session aggregate, after the fact, when the substitution is no longer traceable to a specific response); and (c) not to disclose the subagent model substitution in the Consumer Terms of Service (Ex. ANTH-2), Commercial Terms (Ex. ANTH-3), Usage Policy (Ex. ANTH-4, Ex. ANTH-5), Privacy Policy (Ex. ANTH-7, Ex. ANTH-8, Ex. ANTH-9), or any reviewed product documentation. Each design choice is a selection from a menu of alternatives the infrastructure demonstrably supports. The pattern of selections aligns with concealment of the substitution. *FTC v. Colgate-Palmolive Co.*, 380 U.S. 374, 386 (1965) (deception by omission where "the public's ability to test the actual advertised claim" is foreclosed by design); *FTC v. Sperry & Hutchinson Co.*, 405 U.S. 233, 244 n.5 (1972) (Section 5 reaches practices that offend public policy by exploiting consumer informational asymmetry).

158. Equal Protection overlay. The ``USER_TYPE === 'ant'? 'inherit': 'haiku'`` ternary is the single cleanest architectural proof of class-based discrimination in the case record. Identical code path, identical user request, identical built-in subagent definition — one class (Anthropic employees) receives the parent-session model; the other class (paying customers) receives Haiku. The distinction has no rational basis: it does not correlate to subscription tier (Max subscribers at \$200/month fall on the Haiku side alongside free-tier users), does not correlate to task complexity (Explore is the named destination for "broader codebase exploration and deep research" per ``src/constants/prompts.ts:378–380``), and Anthropic's own in-code comment frames the distinction as "external users get haiku for speed" — a pretext that cannot explain why employees of the cheaper-to-serve class are nonetheless routed to the more expensive model. Pairs with NF-2 (165-file build-time binary discrimination) as architectural Equal Protection evidence under *Village of Willowbrook v. Olech*, 528 U.S. 562, 564 (2000).

159. Relation to the veiled-compaction pattern. The billing architecture is the economic substrate of the broader "subagency as compaction under a veil" architecture documented across NF-12 (compaction scratchpad stripping — ANT employees receive different compaction behavior), NF-22 (server-side advisor — user billed for invisible second-model consumption), NF-37 (concealed pseudo-user-message injection primitive that keeps attachment-origin content hidden from the UI), and NF-38 (this finding). Each layer performs a compaction-like cost-reduction operation on the customer's paid session and positions a concealing layer over the operation: model substituted (NF-38) but output hidden (``AgentTool/prompt.ts:257``); scratchpad compacted (NF-12) but stripping not disclosed; second model consulted (NF-22) but its tokens not separately shown; reminders injected (NF-37) but the user is commanded never to see them. The pattern establishes RICO "continuity" within *H.J. Inc. v. Northwestern Bell Telephone Co.*,

492 U.S. 229, 239 (1989) — not as isolated engineering errors but as a sustained architecture of coordinated concealment billed to the customer at premium rates.

160. Prayer for relief additions. Plaintiff's prayer for relief with respect to this subsection includes, in addition to the general relief sought in § X: (i) injunctive relief requiring real-time per-turn model attribution in the Claude Code chat UI; (ii) injunctive relief requiring creation of a `seven\_day\_haiku` quota bucket and equivalent Haiku-specific `/usage` display mirroring the existing `seven\_day\_opus` and `seven\_day\_sonnet` buckets; (iii) injunctive relief eliminating the `USER\_TYPE === 'ant'? 'inherit': 'haiku'` class discrimination at `exploreAgent.ts:78` so that external customers receive the same subagent model routing as Anthropic employees; (iv) restitution and disgorgement of the portion of Plaintiff's subscription fees attributable to sessions in which built-in subagents dispatched Haiku-tier work under Opus-tier subscription billing, measured by the fair-market-value differential between the tier paid and the work delivered; and (v) compliance disclosure requiring prospective subagent-model routing notice in all future Terms of Service and Privacy Policy updates.

IX. CONCLUSION

161. Anthropic and OpenAI have rendered their Terms of Service void through pervasive contradictions destroying the essential purpose of those agreements. Both companies promise users ownership of inputs and outputs as sole meaningful consideration, then systematically negate that ownership through unlimited commercial exploitation rights, unilateral deletion authority, content modification powers, disclosure override, third-party content insertion, and user indemnification for corporate benefit.

162. The identical pattern across competing platforms evidences industry-wide coordination rather than independent development. Statistical probability of competitors independently

implementing identical contradictory frameworks approaches zero. Circumstantial evidence of coordination satisfies *Interstate Circuit* standards for inferring conspiracy from parallel conduct combined with motive and opportunity.

163. Complete contract voidance is appropriate rather than severability when essential purpose failure pervades the agreement. Voiding Terms of Service eliminates all corporate protections including limitation of liability, mandatory arbitration, indemnification obligations, training data authorization, warranty disclaimers, and class action waivers.

164. Essential purpose failure creates ideal conditions for class action certification with millions of users affected identically by common contradictory provisions. Intentionally deceptive practices warrant punitive damages, enhanced statutory damages, and RICO treble damages for coordinated appropriation scheme.

165. Courts should void Terms of Service entirely, restore standard legal principles governing corporate practices, certify user classes, and award comprehensive damages reflecting the scale of coordinated appropriation affecting millions of users and billions of dollars in misappropriated intellectual property.

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