

Chapter 10: What To Do

From Surveillance Capitalism to AI Imperialism | Chapter 10: What to do

0.1. Abstract

The response program is no longer an open-ended menu. The platform chapters already established that sovereign-scale systems shape exposure, amplification, targeting, monetization, and moderation conditions at mass scale. The essential political choice is whether intelligence infrastructure will remain a privately concentrated empire built on captured social life, or whether societies will recover sovereignty over the terms on which cognition itself is organized.

0.2. Keywords

What to do; surveillance capitalism; AI imperialism; Section 230 reform; AI weights; patrimony of humanity; government capture; evidence-based dossier

0.3. Research Question

What reform doctrine follows from the dossier's evidence on platform power, frontier-model enclosure, and state-company entanglement?

0.4. Scope and Framing Note

This paper examines what to do within the broader dossier *From Surveillance Capitalism to AI Imperialism*. It is intentionally written to circulate on its own: necessary context is restated, evidence boundaries are made explicit, and cited material is reproduced in paper-local form rather than delegated to repo navigation. It contributes to the series as a paper that synthesizes the entire dossier into the direct three-reform program. Adjacent reinforcement appears in Chapter 00, Chapter 01, Chapter 02, Chapter 04, Chapter 07, Chapter 08, Chapter 09, but those cross-references are supportive rather than required for basic comprehension.

0.5. Central Thesis

The response program is no longer an open-ended menu. The dossier now takes a direct three-part stance:

- Refactor Section 230
- AI Weights as Patrimony of Humanity
- Limit Government Capture by AI Companies

Cognitive literacy, digital hygiene, labor organization, competition policy, and public-interest infrastructure remain necessary, but they function as implementation arms of these three reforms rather than as disconnected alternatives.¹

0.6. Evidence and Method Note

This paper is derived from the chapter corpus for Chapter 10 and is grounded in the project's structured research OS. It relies on source notes, extracted claims, entity profiles, and timeline events already normalized under the dossier's evidence hierarchy. Documented facts are asserted directly where the record supports them; disputed matters are marked as such; interpretive claims are bounded explicitly; and speculative overreach is isolated in the evidence-boundary appendix.

0.7. Introduction

This paper asks: What reform doctrine follows from the dossier's evidence on platform power, frontier-model enclosure, and state-company entanglement? Its working answer is clear from the start: The response program is no longer an open-ended menu. The dossier now takes a direct three-part stance:

- Refactor Section 230
- AI Weights as Patrimony of Humanity
- Limit Government Capture by AI Companies

Cognitive literacy, digital hygiene, labor organization, competition policy, and public-interest infrastructure remain necessary, but they function as implementation arms of these three reforms rather than as disconnected alternatives.² The objective is not only to recount developments, but to show why the subject of this chapter belongs inside a structural argument about extraction, enclosure, legitimacy, and reform.

0.8. Historical or Institutional Context

The institutional context for what to do cannot be treated as an isolated debate. Each chapter in this series is part of a longer sequence in which platform design, data capture, labor correction, model governance, and state access accumulate into a single political structure. This section therefore identifies the historical or institutional setting needed for a standalone reader before moving deeper into the paper's own evidentiary claims.

0.9. Main Analysis

The platform chapters already established that sovereign-scale systems shape exposure, amplification, targeting, monetization, and moderation conditions at mass scale. The training chapters

¹Sources: U.S. Department of Education / ERIC, Artificial Intelligence and the Futur... (2023-05-24; see Appendix A1); arXiv, In Oxford Handbook on AI Governance:... (2021-08-05; see Appendix A2); arXiv, Interactions Between Artificial Intel... (2024-12-07; see Appendix A3)

²Sources: U.S. Department of Education / ERIC, Artificial Intelligence and the Futur... (2023-05-24; see Appendix A1); arXiv, In Oxford Handbook on AI Governance:... (2021-08-05; see Appendix A2); arXiv, Interactions Between Artificial Intel... (2024-12-07; see Appendix A3)

established that frontier capability is built from collectively produced expression under contested legal conditions. The access and legitimacy chapters established that frontier firms already move through policy formation, strategic access regimes, lobbying, procurement channels, and public-private legitimacy structures. This final chapter turns those upstream findings into doctrine.³

0.9.1. Refactor Section 230

Documented problem: sovereign-scale platforms do not merely host speech. They structure exposure, optimize amplification, monetize targeting, and govern the conditions under which civic, psychological, and political harm scales.⁴

Interpretation: the legal principle should shift from passive hosting to systemic design responsibility. Refactoring Section 230 does not require treating every user post as if it were authored by the platform. It requires treating sovereign-scale platforms as infrastructures whose reach, design, and monetized amplification create duties commensurate with their power.⁵

Policy direction: obligations should track reach, recommendation architecture, amplification, targeting, monetization, moderation architecture, and measurable systemic harm. Child-safety design duties, recommendation auditability, and sovereign-scale platform-risk obligations fit this framework more naturally than either the old passive-intermediary fiction or blanket publisher liability.⁶

0.9.2. AI Weights as Patrimony of Humanity

Documented problem: frontier-model capability is built from humanity’s accumulated language, art, code, labor, science, and socially filtered judgment under legal conditions that remain actively contested.⁷

Interpretation: once that structure is clear, frontier weights cannot be treated as ordinary private property by default. The dossier’s strongest normative direction is patrimony-oriented governance: public-trust logic, anti-monopoly limits, democratic oversight, interoperability duties, and restrictions on exclusive commercialization. In its strongest form, this doctrine pushes toward AI treated

³Sources: FTC, A Look Behind the Screens: Examining... (n.d.; see Appendix A4); HHS, Social Media and Youth Mental Health:... (n.d.; see Appendix A5); U.S. Copyright Office, Copyright and Artificial Intelligence... (2025-05-09; see Appendix A6); OpenAI, Introducing OpenAI for Government (2025-06-16; see Appendix A7); Anthropic, Claude Gov models for U.S. national s... (2025-06-06; see Appendix A8); U.S. Army Public Affairs, Army Launches Detachment 201: Executi... (2025-06-13; see Appendix A9)

⁴Sources: FTC, A Look Behind the Screens: Examining... (n.d.; see Appendix A4); HHS, Social Media and Youth Mental Health:... (n.d.; see Appendix A5); arXiv, When Kids Mode Isn’t For Kids: Invest... (2025-06-30; see Appendix A10); Pew Research Center, Social Media Seen as Mostly Good for... (2022-12-06; see Appendix A11)

⁵Sources: FTC, A Look Behind the Screens: Examining... (n.d.; see Appendix A4); HHS, Social Media and Youth Mental Health:... (n.d.; see Appendix A5); arXiv, Quantifying social organization and p... (2020-10-01; see Appendix A12)

⁶Sources: arXiv, When Kids Mode Isn’t For Kids: Invest... (2025-06-30; see Appendix A10); arXiv, Unfiltered: How Teens Engage in Body... (2025-04-02; see Appendix A13); FTC, A Look Behind the Screens: Examining... (n.d.; see Appendix A4)

⁷Sources: U.S. Copyright Office, Copyright and Artificial Intelligence... (2025-05-09; see Appendix A6); GovInfo, Copyright and Artificial Intelligence... (2023-08-30; see Appendix A14); OpenAI, Language Models are Unsupervised Mult... (2019-02-14; see Appendix A15); AP, Reddit strikes \$60M deal allowing Goo... (2024-02-22; see Appendix A16)

more like civic utility infrastructure: open and as broadly accessible as public-interest governance, safety constraints, and durable shared infrastructure can make it. This is advocated doctrine, not settled law.⁸

Policy direction: transparency around material training inputs, anti-enclosure limits for dominant model holders, democratic governance mechanisms, and restrictions on exclusive commercialization where public support is material belong inside this reform.⁹

0.9.3. Limit Government Capture by AI Companies

Documented problem: frontier firms already participate in policy formation, export-control protection, government-only offerings, classified-environment distribution channels, lobbying, electoral vehicles, and military-adjacent personnel structures.¹⁰

Interpretation: any frontier firm benefiting materially from public funding, procurement, compute support, export-control protection, regulatory protection, or national-security access should face heightened public obligations. Public support should increase transparency duties rather than deepen private enclosure. This is an anti-capture doctrine, not a generic appeal to ethics and not a claim of total hidden control.¹¹

Policy direction: procurement visibility, subsidy transparency, lobbying disclosure, conflict-of-interest rules, revolving-door limits, independent audits, model-risk reporting, and limits on exclusive private benefit from publicly supported AI infrastructure should be the baseline. Where firms are simultaneously building quasi-public intelligence infrastructure and private survivability hedges, the burden for public accountability should rise further: dependency-building should not be paired with private exemption from shared failure without stronger oversight and obligations.¹²

⁸Sources: TwinChat, AI Copyright Weights: A New Frontier... (n.d.; see Appendix A17); U.S. House of Representatives, Bipartisan House Task Force Report on... (2024-12-17; see Appendix A18); U.S. Copyright Office, Copyright and Artificial Intelligence... (2025-05-09; see Appendix A6); U.S. District Court for the Northern District of California via CourtListener, Bartz et al. v. Anthropic PBC, Order... (2025-06-23; see Appendix A19)

⁹Sources: U.S. Copyright Office, Copyright and Artificial Intelligence... (2025-05-09; see Appendix A6); OpenAI, OpenAI's proposals for the U.S. AI Ac... (2025-03-13; see Appendix A20); AP, Reddit strikes \$60M deal allowing Goo... (2024-02-22; see Appendix A16)

¹⁰Sources: OpenAI, OpenAI's proposals for the U.S. AI Ac... (2025-03-13; see Appendix A20); OpenAI, Introducing OpenAI for Government (2025-06-16; see Appendix A7); Anthropic, Claude Gov models for U.S. national s... (2025-06-06; see Appendix A8); Palantir Technologies, Anthropic and Palantir Partner to Bri... (2024-11-07; see Appendix A21); U.S. Army Public Affairs, Army Launches Detachment 201: Executi... (2025-06-13; see Appendix A9); Axios, Anthropic outspends OpenAI in biggest... (2026-04-21; see Appendix A22)

¹¹Sources: OpenAI, Introducing OpenAI for Government (2025-06-16; see Appendix A7); Anthropic, Claude Gov models for U.S. national s... (2025-06-06; see Appendix A8); Federal Election Commission, LEADING THE FUTURE - committee overview (2026-03-31; see Appendix A23); Dario Amodei, The Adolescence of Technology (2026-01-01; see Appendix A24)

¹²Sources: AP, FTC opens inquiry into Big Tech's par... (2024-01-25; see Appendix A25); OpenAI, Introducing OpenAI for Government (2025-06-16; see Appendix A7); Anthropic, Claude Gov models for U.S. national s... (2025-06-06; see Appendix A8); U.S. Army Public Affairs, Army Launches Detachment 201: Executi... (2025-06-13; see Appendix A9); Tad Friend, Sam Altman's Manifest Destiny (2016-10-03; see Appendix A26); Makena Kelly, Inside Mark Zuckerberg's Top-Secret H... (2023-12-14; see Appendix A27); Rebecca Rose, New Zealand woos the wealthy with 'go... (2025-09-24; see Appendix A28)

0.10. Position Within the Series

This paper sits inside a coordinated 11-paper architecture. Its nearest companions are Chapter 00 (Opening framework), Chapter 01 (Social networks as infrastructure for free data capture), Chapter 02 (Social erosion and moral reconfiguration), Chapter 04 (AI arrives when the human raw material already exists), Chapter 07 (Imperial control over access), Chapter 08 (Power network, legitimacy, and institutional capture), Chapter 09 (Future if no action is taken). In that series logic, Chapter 10 synthesizes the entire dossier into the direct three-reform program.

0.11. Counterarguments

This chapter is partly normative by design. The current corpus supports the legitimacy of a three-reform agenda, but not yet a fully implementation-ready statutory blueprint in every domain. That is not a weakness if the distinction remains explicit.¹³

0.12. Limits of the Evidence

This chapter is partly normative by design. The current corpus supports the legitimacy of a three-reform agenda, but not yet a fully implementation-ready statutory blueprint in every domain. That is not a weakness if the distinction remains explicit.¹⁴

0.13. Reform Relevance

The essential political choice is whether intelligence infrastructure will remain a privately concentrated empire built on captured social life, or whether societies will recover sovereignty over the terms on which cognition itself is organized. The dossier's answer is explicit: sovereign-scale platform accountability, patrimony-oriented governance of frontier weights, and strong limits on opaque state-corporate capture.¹⁵

0.14. Conclusion

The essential political choice is whether intelligence infrastructure will remain a privately concentrated empire built on captured social life, or whether societies will recover sovereignty over the terms on which cognition itself is organized. The dossier closes here. In the architecture of this series, Chapter 10 therefore functions as a self-contained argument while also advancing the cumulative path toward the dossier's three reforms.

¹³Sources: arXiv, In Oxford Handbook on AI Governance:... (2021-08-05; see Appendix A2); arXiv, Interactions Between Artificial Intel... (2024-12-07; see Appendix A3); AP, FTC opens inquiry into Big Tech's par... (2024-01-25; see Appendix A25)

¹⁴Sources: arXiv, In Oxford Handbook on AI Governance:... (2021-08-05; see Appendix A2); arXiv, Interactions Between Artificial Intel... (2024-12-07; see Appendix A3); AP, FTC opens inquiry into Big Tech's par... (2024-01-25; see Appendix A25)

¹⁵Sources: FTC, A Look Behind the Screens: Examining... (n.d.; see Appendix A4); TwinChat, AI Copyright Weights: A New Frontier... (n.d.; see Appendix A17); OpenAI, Introducing OpenAI for Government (2025-06-16; see Appendix A7); Anthropic, Claude Gov models for U.S. national s... (2025-06-06; see Appendix A8)

0.15. Notes

This paper is designed to circulate independently. Public notes point readers to the real external documents first, then to the paper-local appendix entry that explains why each source matters.

The underlying research OS distinguishes among **Documented Fact**, **Disputed Fact**, **Hypothesis / Interpretation**, and **Speculative Narrative Risk**. In Chapter 10, those boundaries remain visible in the prose, in the bibliography, and in the appendix sections that summarize claims and caution points.

0.16. References / Bibliography

1. U.S. Department of Education / ERIC. Artificial Intelligence and the Future of Teaching and Learning: Insights and Recommendations. 2023-05-24. T1 report. Paper appendix: Appendix A1. Corpus companion entry: Source S-124. Audit ID: **us-ed-ai-future-teaching-learning-2023**. Relevance: That an institutional response agenda can legitimately include cognitive literacy, human oversight, and training rather than relying on consumer self-control alone.
2. arXiv. In Oxford Handbook on AI Governance: The Role of Workers in AI Ethics and Governance. 2021-08-05. T1 paper. Paper appendix: Appendix A2. Corpus companion entry: Source S-134. Audit ID: **workers-role-ai-ethics-governance-2021**. Relevance: That workers can be treated as governance actors in AI deployment rather than only as subjects of managerial adoption.
3. arXiv. Interactions Between Artificial Intelligence and Digital Public Infrastructure: Concepts, Benefits, and Challenges. 2024-12-07. T1 paper. Paper appendix: Appendix A3. Corpus companion entry: Source S-001. Audit ID: **ai-digital-public-infrastructure-2024**. Relevance: That a public-interest AI infrastructure layer is intellectually and technically legible within serious research, not an ideological fantasy.
4. FTC. A Look Behind the Screens: Examining the Data Practices of Social Media and Video Streaming Services. n.d.. T1 report. Paper appendix: Appendix A4. Corpus companion entry: Source S-063. Audit ID: **ftc-2024-look-behind-screens**. Relevance: That sovereign-scale platforms should be regulated on the basis of design, amplification, targeting, monetization, moderation architecture, and systemic reach rather than passiv...
5. HHS. Social Media and Youth Mental Health: The U.S. Surgeon General's Advisory. n.d.. T1 advisory. Paper appendix: Appendix A5. Corpus companion entry: Source S-069. Audit ID: **hhs-social-media-youth-advisory-2023**. Relevance: That documented youth exposure, recommendation amplification, trust erosion, and fragmentation support a sovereign-scale duty framework under a refactored Section 230 focused on...
6. U.S. Copyright Office. Copyright and Artificial Intelligence, Part 3: Generative AI Training (Pre-publication Version). 2025-05-09. T1 report. Paper appendix: Appendix A6. Corpus companion entry: Source S-126. Audit ID: **usco-ai-training-report-2025**. Relevance: That whether existing copyright and fair-use doctrine can absorb frontier-model training without consent or compensation remains actively contested across agencies, courts, and...
7. OpenAI. Introducing OpenAI for Government. 2025-06-16. T1 release. Paper appendix: Appendix A7. Corpus companion entry: Source S-090. Audit ID: **openai-for-government-2025**. Relevance: That by June 16, 2025 OpenAI had formally reorganized and expanded its U.S. public-sector work under an **OpenAI for Government** initiative rather than treating govern-

ment use as...

8. Anthropic. Claude Gov models for U.S. national security customers. 2025-06-06. T1 release. Paper appendix: Appendix A8. Corpus companion entry: Source S-004. Audit ID: **anthropic-claude-gov-2025**. Relevance: That by June 6, 2025 Anthropic had publicly launched a government-specific model line for U.S. national security customers.
9. U.S. Army Public Affairs. Army Launches Detachment 201: Executive Innovation Corps to Drive Tech Transformation. 2025-06-13. T1 press_release. Paper appendix: Appendix A9. Corpus companion entry: Source S-026. Audit ID: **army-detachment-201-2025**. Relevance: That by June 2025 official records documented frontier AI firms and adjacent partners operating government-only offerings, classified-environment distribution channels, and dire...
10. arXiv. When Kids Mode Isn't For Kids: Investigating TikTok's Under 13 Experience. 2025-06-30. T1 paper. Paper appendix: Appendix A10. Corpus companion entry: Source S-118. Audit ID: **tiktok-kids-mode-2025**. Relevance: Supports a key factual or interpretive move inside this paper.
11. Pew Research Center. Social Media Seen as Mostly Good for Democracy Across Many Nations, But U.S. is a Major Outlier. 2022-12-06. T2 report. Paper appendix: Appendix A11. Corpus companion entry: Source S-107. Audit ID: **pew-social-media-democracy-outlier-2022**. Relevance: That documented youth exposure, recommendation amplification, trust erosion, and fragmentation support a sovereign-scale duty framework under a refactored Section 230 focused on...
12. arXiv. Quantifying social organization and political polarization in online platforms. 2020-10-01. T1 paper. Paper appendix: Appendix A12. Corpus companion entry: Source S-128. Audit ID: **waller-anderson-online-polarization-2020**. Relevance: Supports a key factual or interpretive move inside this paper.
13. arXiv. Unfiltered: How Teens Engage in Body Image and Shaming Discussions via Instagram Direct Messages (DMs). 2025-04-02. T1 paper. Paper appendix: Appendix A13. Corpus companion entry: Source S-072. Audit ID: **instagram-body-image-dms-2025**. Relevance: Supports a key factual or interpretive move inside this paper.
14. GovInfo. Copyright and Artificial Intelligence; Notice of Inquiry. 2023-08-30. T1 federal_register_notice. Paper appendix: Appendix A14. Corpus companion entry: Source S-125. Audit ID: **usco-ai-noi-2023**. Relevance: Supports a key factual or interpretive move inside this paper.
15. OpenAI. Language Models are Unsupervised Multitask Learners. 2019-02-14. T1 paper. Paper appendix: Appendix A15. Corpus companion entry: Source S-108. Audit ID: **radford-gpt2-webtext-2019**. Relevance: Supports a key factual or interpretive move inside this paper.
16. AP. Reddit strikes \$60M deal allowing Google to train AI models on its posts, unveils IPO plans. 2024-02-22. T2 news_article. Paper appendix: Appendix A16. Corpus companion entry: Source S-024. Audit ID: **ap-reddit-google-ai-training-deal-2024**. Relevance: Supports a key factual or interpretive move inside this paper.
17. TwinChat. AI Copyright Weights: A New Frontier in Intellectual Property Law. n.d.. T3 essay. Paper appendix: Appendix A17. Corpus companion entry: Source S-122. Audit ID: **twinchat-ai-copyright-weights**. Relevance: That frontier AI weights can be read as condensed claims over collectively produced human expression and therefore justify patrimony-

- oriented governance.
18. U.S. House of Representatives. Bipartisan House Task Force Report on Artificial Intelligence. 2024-12-17. T1 report. Paper appendix: Appendix A18. Corpus companion entry: Source S-070. Audit ID: **house-ai-task-force-report-2024**. Relevance: That frontier AI weights can be read as condensed claims over collectively produced human expression and therefore justify patrimony-oriented governance.
 19. U.S. District Court for the Northern District of California via CourtListener. Bartz et al. v. Anthropic PBC, Order on Fair Use. 2025-06-23. T1 court_order. Paper appendix: Appendix A19. Corpus companion entry: Source S-035. Audit ID: **bartz-anthropic-fair-use-order-2025**. Relevance: That whether existing copyright and fair-use doctrine can absorb frontier-model training without consent or compensation remains actively contested across agencies, courts, and...
 20. OpenAI. OpenAI's proposals for the U.S. AI Action Plan. 2025-03-13. T1 policy_submission. Paper appendix: Appendix A20. Corpus companion entry: Source S-096. Audit ID: **openai-us-ai-action-plan-2025**. Relevance: That firms benefiting from export-control protection, strategic state access, or privileged public coordination should face stronger transparency and public-obligation duties th...
 21. Palantir Technologies. Anthropic and Palantir Partner to Bring Claude AI Models to AWS for U.S. Government Intelligence and Defense Operations. 2024-11-07. T1 press_release. Paper appendix: Appendix A21. Corpus companion entry: Source S-098. Audit ID: **palantir-anthropic-aws-gov-partnership-2024**. Relevance: That by November 7, 2024 Anthropic model access for U.S. intelligence and defense agencies was being publicly distributed through a Palantir-plus-AWS stack.
 22. Axios. Anthropic outspends OpenAI in biggest-ever lobbying quarter. 2026-04-21. T3 news_article. Paper appendix: Appendix A22. Corpus companion entry: Source S-027. Audit ID: **axios-anthropic-openai-lobbying-quarter-2026**. Relevance: That frontier legitimacy is increasingly stabilized through a combination of mission rhetoric, elite board composition, security credentials, policy access, and political spending.
 23. Federal Election Commission. LEADING THE FUTURE - committee overview. 2026-03-31. T1 committee_overview. Paper appendix: Appendix A23. Corpus companion entry: Source S-053. Audit ID: **fec-leading-the-future-committee-overview-2026**. Relevance: That frontier legitimacy is increasingly stabilized through a combination of mission rhetoric, elite board composition, security credentials, policy access, and political spending.
 24. Dario Amodei. The Adolescence of Technology. 2026-01-01. T3 essay. Paper appendix: Appendix A24. Corpus companion entry: Source S-047. Audit ID: **dario-amodei-adolescence-of-technology-2026**. Relevance: That by January 1, 2026, Anthropic's own CEO publicly described wealth concentration, backlash, and democratic erosion as serious AI-adjacent risks.
 25. AP. FTC opens inquiry into Big Tech's partnerships with leading AI startups. 2024-01-25. T2 news_article. Paper appendix: Appendix A25. Corpus companion entry: Source S-010. Audit ID: **ap-ftc-ai-partnerships-inquiry-2024**. Relevance: That formal competition scrutiny of frontier-AI partnerships was already underway in the United States by January 2024.
 26. Tad Friend. Sam Altman's Manifest Destiny. 2016-10-03. T2 profile. Paper appendix: Appendix A26. Corpus companion entry: Source S-083. Audit ID: **newyorker-sam-altman-manifest-destiny-2016**. Relevance: That by October 3, 2016 Sam Altman publicly described keeping guns, gold, potas-

sium iodide, antibiotics, batteries, water, gas masks, and land in Big Sur as part of survival pla...

27. Makena Kelly. Inside Mark Zuckerberg’s Top-Secret Hawaii Compound. 2023-12-14. T2 investigation. Paper appendix: Appendix A27. Corpus companion entry: Source S-133. Audit ID: `wired-zuckerberg-hawaii-compound-2023`. Relevance: That WIRED reported on December 14, 2023 that Mark Zuckerberg’s Kauai compound included a 5,000-square-foot underground shelter connected by tunnel, with living space, a mechani...
28. Rebecca Rose. New Zealand woos the wealthy with ‘golden visa’ changes. 2025-09-24. T2 news_article. Paper appendix: Appendix A28. Corpus companion entry: Source S-058. Audit ID: `ft-new-zealand-golden-visa-2025`. Relevance: That debate over wealthy foreigners seeking New Zealand bolt-holes remained live in serious financial press by 2025.
29. TwinChat. How AI Is Ruining Your Life to Build an Empire. n.d.. T3 essay. Paper appendix: Appendix A29. Corpus companion entry: Source S-123. Audit ID: `twinchat-ai-ruining-your-life`. Relevance: Supports a key factual or interpretive move inside this paper.
30. The Holy See. Encyclical Letter of His Holiness Leo XIV Magnifica Humanitas. 2026-05-15. T1 encyclical. Paper appendix: Appendix A30. Corpus companion entry: Source S-127. Audit ID: `vatican-magnifica-humanitas-2026`. Relevance: That by May 15, 2026, the AI legitimacy field had expanded into a formal Vatican encyclical rather than remaining confined to industry, policy, or technical discourse.

0.17. Appendix A. Source Register for This Paper

0.17.1. Appendix A1. Artificial Intelligence and the Future of Teaching and Learning: Insights and Recommendations

Source citation: Artificial Intelligence and the Future of Teaching and Learning: Insights and Recommendations **Institution / author:** U.S. Department of Education / ERIC **Date:** 2023-05-24 **Tier / type:** T1 / report **Why it matters in this paper:** That an institutional response agenda can legitimately include cognitive literacy, human oversight, and training rather than relying on consumer self-control alone. **Corpus companion entry:** Source S-124 **Audit ID:** `us-ed-ai-future-teaching-learning-2023`

0.17.2. Appendix A2. In Oxford Handbook on AI Governance: The Role of Workers in AI Ethics and Governance

Source citation: In Oxford Handbook on AI Governance: The Role of Workers in AI Ethics and Governance **Institution / author:** arXiv **Date:** 2021-08-05 **Tier / type:** T1 / paper **Why it matters in this paper:** That workers can be treated as governance actors in AI deployment rather than only as subjects of managerial adoption. **Corpus companion entry:** Source S-134 **Audit ID:** `workers-role-ai-ethics-governance-2021`

0.17.3. Appendix A3. Interactions Between Artificial Intelligence and Digital Public Infrastructure: Concepts, Benefits, and Challenges

Source citation: Interactions Between Artificial Intelligence and Digital Public Infrastructure: Concepts, Benefits, and Challenges **Institution / author:** arXiv **Date:** 2024-12-07 **Tier / type:** T1 / paper **Why it matters in this paper:** That a public-interest AI infrastructure layer is intellectually and technically legible within serious research, not an ideological fantasy. **Corpus companion entry:** Source S-001 **Audit ID:** ai-digital-public-infrastructure-2024

0.17.4. Appendix A4. A Look Behind the Screens: Examining the Data Practices of Social Media and Video Streaming Services

Source citation: A Look Behind the Screens: Examining the Data Practices of Social Media and Video Streaming Services **Institution / author:** FTC **Date:** n.d. **Tier / type:** T1 / report **Why it matters in this paper:** That sovereign-scale platforms should be regulated on the basis of design, amplification, targeting, monetization, moderation architecture, and systemic reach rather than passiv... **Corpus companion entry:** Source S-063 **Audit ID:** ftc-2024-look-behind-screens

0.17.5. Appendix A5. Social Media and Youth Mental Health: The U.S. Surgeon General’s Advisory

Source citation: Social Media and Youth Mental Health: The U.S. Surgeon General’s Advisory **Institution / author:** HHS **Date:** n.d. **Tier / type:** T1 / advisory **Why it matters in this paper:** That documented youth exposure, recommendation amplification, trust erosion, and fragmentation support a sovereign-scale duty framework under a refactored Section 230 focused on... **Corpus companion entry:** Source S-069 **Audit ID:** hhs-social-media-youth-advisory-2023

0.17.6. Appendix A6. Copyright and Artificial Intelligence, Part 3: Generative AI Training (Pre-publication Version)

Source citation: Copyright and Artificial Intelligence, Part 3: Generative AI Training (Pre-publication Version) **Institution / author:** U.S. Copyright Office **Date:** 2025-05-09 **Tier / type:** T1 / report **Why it matters in this paper:** That whether existing copyright and fair-use doctrine can absorb frontier-model training without consent or compensation remains actively contested across agencies, courts, and... **Corpus companion entry:** Source S-126 **Audit ID:** usco-ai-training-report-2025

0.17.7. Appendix A7. Introducing OpenAI for Government

Source citation: Introducing OpenAI for Government **Institution / author:** OpenAI **Date:** 2025-06-16 **Tier / type:** T1 / release **Why it matters in this paper:** That by June 16, 2025 OpenAI had formally reorganized and expanded its U.S. public-sector work under an OpenAI for Government initiative rather than treating government use as... **Corpus companion entry:** Source S-090 **Audit ID:** openai-for-government-2025

0.17.8. Appendix A8. Claude Gov models for U.S. national security customers

Source citation: Claude Gov models for U.S. national security customers **Institution / author:** Anthropic **Date:** 2025-06-06 **Tier / type:** T1 / release **Why it matters in this paper:** That by June 6, 2025 Anthropic had publicly launched a government-specific model line for U.S. national security customers. **Corpus companion entry:** Source S-004 **Audit ID:** anthropic-claude-gov-2025

0.17.9. Appendix A9. Army Launches Detachment 201: Executive Innovation Corps to Drive Tech Transformation

Source citation: Army Launches Detachment 201: Executive Innovation Corps to Drive Tech Transformation **Institution / author:** U.S. Army Public Affairs **Date:** 2025-06-13 **Tier / type:** T1 / press_release **Why it matters in this paper:** That by June 2025 official records documented frontier AI firms and adjacent partners operating government-only offerings, classified-environment distribution channels, and dire... **Corpus companion entry:** Source S-026 **Audit ID:** army-detachment-201-2025

0.17.10. Appendix A10. When Kids Mode Isn't For Kids: Investigating TikTok's Under 13 Experience

Source citation: When Kids Mode Isn't For Kids: Investigating TikTok's Under 13 Experience **Institution / author:** arXiv **Date:** 2025-06-30 **Tier / type:** T1 / paper **Why it matters in this paper:** Supports a key factual or interpretive move inside this paper. **Corpus companion entry:** Source S-118 **Audit ID:** tiktok-kids-mode-2025

0.17.11. Appendix A11. Social Media Seen as Mostly Good for Democracy Across Many Nations, But U.S. is a Major Outlier

Source citation: Social Media Seen as Mostly Good for Democracy Across Many Nations, But U.S. is a Major Outlier **Institution / author:** Pew Research Center **Date:** 2022-12-06 **Tier / type:** T2 / report **Why it matters in this paper:** That documented youth exposure, recommendation amplification, trust erosion, and fragmentation support a sovereign-scale duty framework under a refactored Section 230 focused on... **Corpus companion entry:** Source S-107 **Audit ID:** pew-social-media-democracy-outlier-2022

0.17.12. Appendix A12. Quantifying social organization and political polarization in online platforms

Source citation: Quantifying social organization and political polarization in online platforms **Institution / author:** arXiv **Date:** 2020-10-01 **Tier / type:** T1 / paper **Why it matters in this paper:** Supports a key factual or interpretive move inside this paper. **Corpus companion entry:** Source S-128 **Audit ID:** waller-anderson-online-polarization-2020

0.17.13. Appendix A13. Unfiltered: How Teens Engage in Body Image and Shaming Discussions via Instagram Direct Messages (DMs)

Source citation: Unfiltered: How Teens Engage in Body Image and Shaming Discussions via Instagram Direct Messages (DMs) **Institution / author:** arXiv **Date:** 2025-04-02 **Tier / type:** T1 / paper **Why it matters in this paper:** Supports a key factual or interpretive move inside this paper. **Corpus companion entry:** Source S-072 **Audit ID:** instagram-body-image-dms-2025

0.17.14. Appendix A14. Copyright and Artificial Intelligence; Notice of Inquiry

Source citation: Copyright and Artificial Intelligence; Notice of Inquiry **Institution / author:** GovInfo **Date:** 2023-08-30 **Tier / type:** T1 / federal_register_notice **Why it matters in this paper:** Supports a key factual or interpretive move inside this paper. **Corpus companion entry:** Source S-125 **Audit ID:** usco-ai-noi-2023

0.17.15. Appendix A15. Language Models are Unsupervised Multitask Learners

Source citation: Language Models are Unsupervised Multitask Learners **Institution / author:** OpenAI **Date:** 2019-02-14 **Tier / type:** T1 / paper **Why it matters in this paper:** Supports a key factual or interpretive move inside this paper. **Corpus companion entry:** Source S-108 **Audit ID:** radford-gpt2-webtext-2019

0.17.16. Appendix A16. Reddit strikes \$60M deal allowing Google to train AI models on its posts, unveils IPO plans

Source citation: Reddit strikes \$60M deal allowing Google to train AI models on its posts, unveils IPO plans **Institution / author:** AP **Date:** 2024-02-22 **Tier / type:** T2 / news_article **Why it matters in this paper:** Supports a key factual or interpretive move inside this paper. **Corpus companion entry:** Source S-024 **Audit ID:** ap-reddit-google-ai-training-deal-2024

0.17.17. Appendix A17. AI Copyright Weights: A New Frontier in Intellectual Property Law

Source citation: AI Copyright Weights: A New Frontier in Intellectual Property Law **Institution / author:** TwinChat **Date:** n.d. **Tier / type:** T3 / essay **Why it matters in this paper:** That frontier AI weights can be read as condensed claims over collectively produced human expression and therefore justify patrimony-oriented governance. **Corpus companion entry:** Source S-122 **Audit ID:** twinchat-ai-copyright-weights

0.17.18. Appendix A18. Bipartisan House Task Force Report on Artificial Intelligence

Source citation: Bipartisan House Task Force Report on Artificial Intelligence **Institution / author:** U.S. House of Representatives **Date:** 2024-12-17 **Tier / type:** T1 / report **Why it matters in this paper:** That frontier AI weights can be read as condensed claims over collectively produced human expression and therefore justify patrimony-oriented governance. **Corpus companion entry:** Source S-070 **Audit ID:** house-ai-task-force-report-2024

0.17.19. Appendix A19. Bartz et al. v. Anthropic PBC, Order on Fair Use

Source citation: Bartz et al. v. Anthropic PBC, Order on Fair Use **Institution / author:** U.S. District Court for the Northern District of California via CourtListener **Date:** 2025-06-23 **Tier / type:** T1 / court_order **Why it matters in this paper:** That whether existing copyright and fair-use doctrine can absorb frontier-model training without consent or compensation remains actively contested across agencies, courts, and... **Corpus companion entry:** Source S-035 **Audit ID:** bartz-anthropic-fair-use-order-2025

0.17.20. Appendix A20. OpenAI's proposals for the U.S. AI Action Plan

Source citation: OpenAI's proposals for the U.S. AI Action Plan **Institution / author:** OpenAI **Date:** 2025-03-13 **Tier / type:** T1 / policy_submission **Why it matters in this paper:** That firms benefiting from export-control protection, strategic state access, or privileged public coordination should face stronger transparency and public-obligation duties th... **Corpus companion entry:** Source S-096 **Audit ID:** openai-us-ai-action-plan-2025

0.17.21. Appendix A21. Anthropic and Palantir Partner to Bring Claude AI Models to AWS for U.S. Government Intelligence and Defense Operations

Source citation: Anthropic and Palantir Partner to Bring Claude AI Models to AWS for U.S. Government Intelligence and Defense Operations **Institution / author:** Palantir Technologies **Date:** 2024-11-07 **Tier / type:** T1 / press_release **Why it matters in this paper:** That by November 7, 2024 Anthropic model access for U.S. intelligence and defense agencies was being publicly distributed through a Palantir-plus-AWS stack. **Corpus companion entry:** Source S-098 **Audit ID:** palantir-anthropic-aws-gov-partnership-2024

0.17.22. Appendix A22. Anthropic outspends OpenAI in biggest-ever lobbying quarter

Source citation: Anthropic outspends OpenAI in biggest-ever lobbying quarter **Institution / author:** Axios **Date:** 2026-04-21 **Tier / type:** T3 / news_article **Why it matters in this paper:** That frontier legitimacy is increasingly stabilized through a combination of mission rhetoric, elite board composition, security credentials, policy access, and political spending. **Corpus companion entry:** Source S-027 **Audit ID:** axios-anthropic-openai-lobbying-quarter-2026

0.17.23. Appendix A23. LEADING THE FUTURE - committee overview

Source citation: LEADING THE FUTURE - committee overview **Institution / author:** Federal Election Commission **Date:** 2026-03-31 **Tier / type:** T1 / committee_overview **Why it matters in this paper:** That frontier legitimacy is increasingly stabilized through a combination of mission rhetoric, elite board composition, security credentials, policy access, and political spending. **Corpus companion entry:** Source S-053 **Audit ID:** fec-leading-the-future-committee-overview-2026

0.17.24. Appendix A24. The Adolescence of Technology

Source citation: The Adolescence of Technology **Institution / author:** Dario Amodei **Date:** 2026-01-01 **Tier / type:** T3 / essay **Why it matters in this paper:** That by January 1, 2026, Anthropic’s own CEO publicly described wealth concentration, backlash, and democratic erosion as serious AI-adjacent risks. **Corpus companion entry:** Source S-047 **Audit ID:** dario-amodei-adolescence-of-technology-2026

0.17.25. Appendix A25. FTC opens inquiry into Big Tech’s partnerships with leading AI startups

Source citation: FTC opens inquiry into Big Tech’s partnerships with leading AI startups **Institution / author:** AP **Date:** 2024-01-25 **Tier / type:** T2 / news_article **Why it matters in this paper:** That formal competition scrutiny of frontier-AI partnerships was already underway in the United States by January 2024. **Corpus companion entry:** Source S-010 **Audit ID:** ap-ftc-ai-partnerships-inquiry-2024

0.17.26. Appendix A26. Sam Altman’s Manifest Destiny

Source citation: Sam Altman’s Manifest Destiny **Institution / author:** Tad Friend **Date:** 2016-10-03 **Tier / type:** T2 / profile **Why it matters in this paper:** That by October 3, 2016 Sam Altman publicly described keeping guns, gold, potassium iodide, antibiotics, batteries, water, gas masks, and land in Big Sur as part of survival plan... **Corpus companion entry:** Source S-083 **Audit ID:** newyorker-sam-altman-manifest-destiny-2016

0.17.27. Appendix A27. Inside Mark Zuckerberg’s Top-Secret Hawaii Compound

Source citation: Inside Mark Zuckerberg’s Top-Secret Hawaii Compound **Institution / author:** Makena Kelly **Date:** 2023-12-14 **Tier / type:** T2 / investigation **Why it matters in this paper:** That WIRED reported on December 14, 2023 that Mark Zuckerberg’s Kauai compound included a 5,000-square-foot underground shelter connected by tunnel, with living space, a mechanical... **Corpus companion entry:** Source S-133 **Audit ID:** wired-zuckerberg-hawaii-compound-2023

0.17.28. Appendix A28. New Zealand woos the wealthy with ‘golden visa’ changes

Source citation: New Zealand woos the wealthy with ‘golden visa’ changes **Institution / author:** Rebecca Rose **Date:** 2025-09-24 **Tier / type:** T2 / news_article **Why it matters in this paper:** That debate over wealthy foreigners seeking New Zealand bolt-holes remained live in serious financial press by 2025. **Corpus companion entry:** Source S-058 **Audit ID:** ft-new-zealand-golden-visa-2025

0.17.29. Appendix A29. How AI Is Ruining Your Life to Build an Empire

Source citation: How AI Is Ruining Your Life to Build an Empire **Institution / author:** TwinChat **Date:** n.d. **Tier / type:** T3 / essay **Why it matters in this paper:** Supports a key factual or interpretive move inside this paper. **Corpus companion entry:** Source S-123 **Audit ID:** twinchat-ai-ruining-your-life

0.17.30. Appendix A30. Encyclical Letter of His Holiness Leo XIV Magnifica Humanitas

Source citation: Encyclical Letter of His Holiness Leo XIV Magnifica Humanitas **Institution / author:** The Holy See **Date:** 2026-05-15 **Tier / type:** T1 / encyclical **Why it matters in this paper:** That by May 15, 2026, the AI legitimacy field had expanded into a formal Vatican encyclical rather than remaining confined to industry, policy, or technical discourse. **Corpus companion entry:** Source S-127 **Audit ID:** vatican-magnifica-humanitas-2026

0.18. Appendix B. Claims Used in This Paper

0.18.1. Appendix B1. A public-interest AI infrastructure layer is intellectually and technically legib...

Claim statement: That a public-interest AI infrastructure layer is intellectually and technically legible within serious research, not an ideological fantasy. **Evidence status:** Documented Fact **Source support:** Appendix A3 **Corpus companion entry:** Claim C-001 **Audit Claim ID:** ai-digital-public-infrastructure-2024-10-counterstrategy-response-agenda-claim-01

0.18.2. Appendix B2. Chapter 10 can discuss public or quasi-public AI infrastructure as a concrete res...

Claim statement: That Chapter 10 can discuss public or quasi-public AI infrastructure as a concrete response pathway. **Evidence status:** Documented Fact **Source support:** Appendix A3 **Corpus companion entry:** Claim C-002 **Audit Claim ID:** ai-digital-public-infrastructure-2024-10-counterstrategy-response-agenda-claim-02

0.18.3. Appendix B3. Interoperability, governance, and public-value design belong inside the response...

Claim statement: That interoperability, governance, and public-value design belong inside the response agenda alongside labor and competition policy. **Evidence status:** Documented Fact **Source support:** Appendix A3 **Corpus companion entry:** Claim C-003 **Audit Claim ID:** ai-digital-public-infrastructure-2024-10-counterstrategy-response-agenda-claim-03

0.18.4. Appendix B4. By June 6, 2025 Anthropic had publicly launched a government-specific model line...

Claim statement: That by June 6, 2025 Anthropic had publicly launched a government-specific model line for U.S. national security customers. **Evidence status:** Documented Fact **Source support:** Appendix A8 **Corpus companion entry:** Claim C-010 **Audit Claim ID:** anthropic-claude-gov-2025-07-export-controls-compute-defense-access-claim-01

0.18.5. Appendix B5. Anthropic publicly described those models as already deployed in high-level natio...

Claim statement: That Anthropic publicly described those models as already deployed in high-level national-security agencies and restricted to classified environments. **Evidence status:** Documented Fact

umented Fact **Source support:** Appendix A8 **Corpus companion entry:** Claim C-011 **Audit Claim ID:** anthropic-claude-gov-2025-07-export-controls-compute-defense-access-claim-02

0.18.6. Appendix B6. The company openly framed model adaptation for defense, intelligence, and cyberse...

Claim statement: That the company openly framed model adaptation for defense, intelligence, and cybersecurity use cases as part of its government offering. **Evidence status:** Documented Fact **Source support:** Appendix A8 **Corpus companion entry:** Claim C-012 **Audit Claim ID:** anthropic-claude-gov-2025-07-export-controls-compute-defense-access-claim-03

0.18.7. Appendix B7. Formal competition scrutiny of frontier-AI partnerships was already underway in t...

Claim statement: That formal competition scrutiny of frontier-AI partnerships was already underway in the United States by January 2024. **Evidence status:** Documented Fact **Source support:** Appendix A25 **Corpus companion entry:** Claim C-028 **Audit Claim ID:** ap-ftc-ai-partnerships-inquiry-2024-10-counterstrategy-response-agenda-claim-01

0.18.8. Appendix B8. Chapter 10 can ground its competition-policy response layer in documented concern...

Claim statement: That Chapter 10 can ground its competition-policy response layer in documented concerns about structural dependence between model labs and major platforms. **Evidence status:** Documented Fact **Source support:** Appendix A25 **Corpus companion entry:** Claim C-029 **Audit Claim ID:** ap-ftc-ai-partnerships-inquiry-2024-10-counterstrategy-response-agenda-claim-02

0.18.9. Appendix B9. A response agenda focused on concentration, interoperability, and structural reme...

Claim statement: That a response agenda focused on concentration, interoperability, and structural remedies is anchored in real regulatory attention rather than in abstract anti-corporate rhetoric. **Evidence status:** Documented Fact **Source support:** Appendix A25 **Corpus companion entry:** Claim C-030 **Audit Claim ID:** ap-ftc-ai-partnerships-inquiry-2024-10-counterstrategy-response-agenda-claim-03

0.18.10. Appendix B10. By January 1, 2026, Anthropic's own CEO publicly described wealth concentration,...

Claim statement: That by January 1, 2026, Anthropic's own CEO publicly described wealth concentration, backlash, and democratic erosion as serious AI-adjacent risks. **Evidence status:** Documented Fact **Source support:** Appendix A24 **Corpus companion entry:** Claim C-141 **Audit Claim ID:** dario-amodei-adolescence-of-technology-2026-08-power-networks-legitimacy-capture-claim-04

0.18.11. Appendix B11. The 80%-of-wealth pledge angle can be grounded in a primary source rather than on...

Claim statement: That the 80%-of-wealth pledge angle can be grounded in a primary source rather than only in a secondary or tertiary retelling. **Evidence status:** Documented Fact **Source support:** Appendix A24 **Corpus companion entry:** Claim C-142 **Audit Claim ID:** dario-amodei-adolescence-of-technology-2026-08-power-networks-legitimacy-capture-claim-02

0.18.12. Appendix B12. Chapter 08 can document elite self-legitimation through philanthropy and public m...

Claim statement: That Chapter 08 can document elite self-legitimation through philanthropy and public moral positioning without pretending philanthropy resolves the structural problem. **Evidence status:** Documented Fact **Source support:** Appendix A24 **Corpus companion entry:** Claim C-143 **Audit Claim ID:** dario-amodei-adolescence-of-technology-2026-08-power-networks-legitimacy-capture-claim-03

0.18.13. Appendix B13. Chapter 09 can cite a frontier founder's own warning that extreme wealth concentr...

Claim statement: That Chapter 09 can cite a frontier founder's own warning that extreme wealth concentration under AI could destabilize democratic society. **Evidence status:** Documented Fact **Source support:** Appendix A24 **Corpus companion entry:** Claim C-144 **Audit Claim ID:** dario-amodei-adolescence-of-technology-2026-08-power-networks-legitimacy-capture-claim-04

0.18.14. Appendix B14. Chapter 10 can explicitly distinguish between voluntary philanthropy and the stro...

Claim statement: That Chapter 10 can explicitly distinguish between voluntary philanthropy and the stronger structural responses the dossier argues are necessary. **Evidence status:** Documented Fact **Source support:** Appendix A24 **Corpus companion entry:** Claim C-145 **Audit Claim ID:** dario-amodei-adolescence-of-technology-2026-08-power-networks-legitimacy-capture-claim-05

0.18.15. Appendix B15. Debate over wealthy foreigners seeking New Zealand bolt-holes remained live in se...

Claim statement: That debate over wealthy foreigners seeking New Zealand bolt-holes remained live in serious financial press by 2025. **Evidence status:** Documented Fact **Source support:** Appendix A28 **Corpus companion entry:** Claim C-182 **Audit Claim ID:** ft-new-zealand-golden-visa-2025-09-future-if-unchecked-claim-01

0.18.16. Appendix B16. The Financial Times reported on September 24, 2025 that Peter Thiel's New Zealand...

Claim statement: That the Financial Times reported on September 24, 2025 that Peter Thiel's New Zealand citizenship had been granted under 'exceptional circumstances' despite reportedly spending 12 days in the country, and that his Lake Wānaka lodge plans were later rejected in 2022.

Evidence status: Documented Fact **Source support:** Appendix A28 **Corpus companion entry:** Claim C-183 **Audit Claim ID:** ft-new-zealand-golden-visa-2025-11-facebook-thiel-intelligence-a

0.18.17. Appendix B17. By October 3, 2016 Sam Altman publicly described keeping guns, gold, potassium io...

Claim statement: That by October 3, 2016 Sam Altman publicly described keeping guns, gold, potassium iodide, antibiotics, batteries, water, gas masks, and land in Big Sur as part of survival planning, and named Peter Thiel's house in New Zealand as a backup plan. **Evidence status:** Documented Fact **Source support:** Appendix A26 **Corpus companion entry:** Claim C-275 **Audit Claim ID:** newyorker-sam-altman-manifest-destiny-2016-08-power-networks-legitimacy-capture-claim

0.18.18. Appendix B18. The public record already contained elite retreat and private-fallback planning b...

Claim statement: That the public record already contained elite retreat and private-fallback planning by a future central AI executive years before the current frontier-AI dependency regime matured. **Evidence status:** Documented Fact **Source support:** Appendix A26 **Corpus companion entry:** Claim C-276 **Audit Claim ID:** newyorker-sam-altman-manifest-destiny-2016-09-future-if-un

0.18.19. Appendix B19. By June 16, 2025 OpenAI had formally reorganized and expanded its U.S. public-sec...

Claim statement: That by June 16, 2025 OpenAI had formally reorganized and expanded its U.S. public-sector work under an OpenAI for Government initiative rather than treating government use as incidental. **Evidence status:** Documented Fact **Source support:** Appendix A7 **Corpus companion entry:** Claim C-298 **Audit Claim ID:** openai-for-government-2025-07-export-controls-compu

0.18.20. Appendix B20. OpenAI publicly tied government access to secure environments, limited national-s...

Claim statement: That OpenAI publicly tied government access to secure environments, limited national-security custom models, and structured forward visibility into upcoming capabilities. **Evidence status:** Documented Fact **Source support:** Appendix A7 **Corpus companion entry:** Claim C-299 **Audit Claim ID:** openai-for-government-2025-07-export-controls-compute-defense-access

0.18.21. Appendix B21. OpenAI's own record documents a formal DoD CDAO pilot with a \$200 million ceiling...

Claim statement: That OpenAI's own record documents a formal DoD CDAO pilot with a \$200 million ceiling under this initiative. **Evidence status:** Documented Fact **Source support:** Appendix A7 **Corpus companion entry:** Claim C-300 **Audit Claim ID:** openai-for-government-2025-07-export-controls-compute-defense-access-claim-03

0.18.22. Appendix B22. By November 7, 2024 Anthropic model access for U.S. intelligence and defense agen...

Claim statement: That by November 7, 2024 Anthropic model access for U.S. intelligence and defense agencies was being publicly distributed through a Palantir-plus-AWS stack. **Evidence**

status: Documented Fact **Source support:** Appendix A21 **Corpus companion entry:** Claim C-

328 **Audit Claim ID:** palantir-anthropic-aws-gov-partnership-2024-07-export-controls-compute-defer

0.18.23. Appendix B23. Classified-environment frontier-model access was tied to specific accredited infr...

Claim statement: That classified-environment frontier-model access was tied to specific accredited infrastructure and defense-cleared platform arrangements rather than open public availability.

Evidence status: Documented Fact **Source support:** Appendix A21 **Corpus companion entry:** Claim C-329 **Audit Claim ID:** palantir-anthropic-aws-gov-partnership-2024-07-export-controls-c

0.18.24. Appendix B24. Chapter 07 can document a concrete distribution channel where model access, cloud...

Claim statement: That Chapter 07 can document a concrete distribution channel where model access, cloud concentration, accreditation, and defense use cases are bundled together. **Evidence**

status: Documented Fact **Source support:** Appendix A21 **Corpus companion entry:** Claim C-

330 **Audit Claim ID:** palantir-anthropic-aws-gov-partnership-2024-07-export-controls-compute-defer

0.18.25. Appendix B25. An institutional response agenda can legitimately include cognitive literacy, hum...

Claim statement: That an institutional response agenda can legitimately include cognitive literacy, human oversight, and training rather than relying on consumer self-control alone. **Evidence**

status: Documented Fact **Source support:** Appendix A1 **Corpus companion entry:** Claim C-

404 **Audit Claim ID:** us-ed-ai-future-teaching-learning-2023-10-counterstrategy-response-agenda-c

0.18.26. Appendix B26. The response layer in Chapter 10 can draw on primary public guidance for educatio...

Claim statement: That the response layer in Chapter 10 can draw on primary public guidance for educational governance under AI. **Evidence status:** Documented Fact **Source**

support: Appendix A1 **Corpus companion entry:** Claim C-405 **Audit Claim ID:**

us-ed-ai-future-teaching-learning-2023-10-counterstrategy-response-agenda-claim-02

0.18.27. Appendix B27. Defensive measures around attention, judgment, and delegated cognition are compat...

Claim statement: That defensive measures around attention, judgment, and delegated cognition are compatible with mainstream institutional guidance. **Evidence status:** Documented Fact

Source support: Appendix A1 **Corpus companion entry:** Claim C-406 **Audit Claim ID:**

us-ed-ai-future-teaching-learning-2023-10-counterstrategy-response-agenda-claim-03

0.18.28. Appendix B28. By May 15, 2026, the AI legitimacy field had expanded into a formal Vatican encyc...

Claim statement: That by May 15, 2026, the AI legitimacy field had expanded into a formal Vatican encyclical rather than remaining confined to industry, policy, or technical discourse. **Evidence**

status: Documented Fact **Source support:** Appendix A30 **Corpus companion entry:** Claim C-

413 **Audit Claim ID:** vatican-magnifica-humanitas-2026-08-power-networks-legitimacy-capture-claim-

0.18.29. Appendix B29. A primary moral-religious authority explicitly treated concentration of digital p...

Claim statement: That a primary moral-religious authority explicitly treated concentration of digital power, hidden labor, and AI-enabled war as matters of public ethics and institutional respon-

sibility. **Evidence status:** Documented Fact **Source support:** Appendix A30 **Corpus compan-**

ion entry: Claim C-414 **Audit Claim ID:** vatican-magnifica-humanitas-2026-08-power-networks-legitim

0.18.30. Appendix B30. Chapter 08 can document legitimacy formation through religious and civilizational...

Claim statement: That Chapter 08 can document legitimacy formation through religious and civilizational authority without falling into conspiratorial overreach. **Evidence status:** Documented

Fact **Source support:** Appendix A30 **Corpus companion entry:** Claim C-415 **Audit Claim**

ID: vatican-magnifica-humanitas-2026-08-power-networks-legitimacy-capture-claim-03

0.18.31. Appendix B31. Chapters 09 and 10 can cite a major normative institution when discussing concent...

Claim statement: That Chapters 09 and 10 can cite a major normative institution when discussing concentration, dehumanization, and the need for response beyond consumer choice. **Ev-**

idence status: Documented Fact **Source support:** Appendix A30 **Corpus companion entry:**

Claim C-416 **Audit Claim ID:** vatican-magnifica-humanitas-2026-08-power-networks-legitimacy-captur

0.18.32. Appendix B32. By June 2025 official records documented frontier AI firms and adjacent partners...

Claim statement: That by June 2025 official records documented frontier AI firms and adjacent partners operating government-only offerings, classified-environment distribution channels,

and direct military personnel interfaces rather than engaging the state only through lobbying or public rhetoric. **Evidence status:** Documented Fact **Source support:** Appendix A7, Appendix

A8, Appendix A21, Appendix A9 **Corpus companion entry:** Claim C-441 **Audit Claim ID:**

vector08-formal-state-interface-documented-claim-01

0.18.33. Appendix B33. WIRED reported on December 14, 2023 that Mark Zuckerberg's Kauai compound include...

Claim statement: That WIRED reported on December 14, 2023 that Mark Zuckerberg's Kauai compound included a 5,000-square-foot underground shelter connected by tunnel, with

living space, a mechanical room, and an escape hatch. **Evidence status:** Documented Fact **Source support:** Appendix A27 **Corpus companion entry:** Claim C-489 **Audit Claim ID:** wired-zuckerberg-hawaii-compound-2023-08-power-networks-legitimacy-capture-claim-01

0.18.34. Appendix B34. WIRED reported strict NDAs, crew separation, self-sufficiency features, and secre...

Claim statement: That WIRED reported strict NDAs, crew separation, self-sufficiency features, and secrecy controls around the Kauai project that resembled secure-installation practices more than ordinary private construction. **Evidence status:** Documented Fact **Source support:** Appendix A27 **Corpus companion entry:** Claim C-490 **Audit Claim ID:** wired-zuckerberg-hawaii-compound-2023-08-power-networks-legitimacy-capture-claim-02

0.18.35. Appendix B35. Workers can be treated as governance actors in AI deployment rather than only as...

Claim statement: That workers can be treated as governance actors in AI deployment rather than only as subjects of managerial adoption. **Evidence status:** Documented Fact **Source support:** Appendix A2 **Corpus companion entry:** Claim C-491 **Audit Claim ID:** workers-role-ai-ethics-governance-2021-10-counterstrategy-response-agenda-claim-01

0.18.36. Appendix B36. Chapter 10 can legitimately include bargaining, protected reporting, and workplac...

Claim statement: That Chapter 10 can legitimately include bargaining, protected reporting, and workplace consultation as part of a serious response agenda. **Evidence status:** Documented Fact **Source support:** Appendix A2 **Corpus companion entry:** Claim C-492 **Audit Claim ID:** workers-role-ai-ethics-governance-2021-10-counterstrategy-response-agenda-claim-02

0.18.37. Appendix B37. Labor organization belongs inside AI governance discussions, not outside them

Claim statement: That labor organization belongs inside AI governance discussions, not outside them. **Evidence status:** Documented Fact **Source support:** Appendix A2 **Corpus companion entry:** Claim C-493 **Audit Claim ID:** workers-role-ai-ethics-governance-2021-10-counterstrategy-resp

0.18.38. Appendix B38. Whether existing copyright and fair-use doctrine can absorb frontier-model traini...

Claim statement: That whether existing copyright and fair-use doctrine can absorb frontier-model training without consent or compensation remains actively contested across agencies, courts, and litigants. **Evidence status:** Disputed Fact **Source support:** Appendix A6, Appendix A19 **Corpus companion entry:** Claim C-428 **Audit Claim ID:** vector04-training-legality-contested-claim-0

0.18.39. Appendix B39. Sovereign-scale platforms should be regulated on the basis of design, amplificati...

Claim statement: That sovereign-scale platforms should be regulated on the basis of design, amplification, targeting, monetization, moderation architecture, and systemic reach rather than passive-hosting logic. **Evidence status:** Hypothesis / Interpretation **Source support:** Appendix A4 **Corpus companion entry:** Claim C-418 **Audit Claim ID:** vector01-section230-sovereign-scale-reform-claim-01

0.18.40. Appendix B40. Documented youth exposure, recommendation amplification, trust erosion, and fragm...

Claim statement: That documented youth exposure, recommendation amplification, trust erosion, and fragmentation support a sovereign-scale duty framework under a refactored Section 230 focused on design, exposure, and systemic harm. **Evidence status:** Hypothesis / Interpretation **Source support:** Appendix A5, Appendix A11 **Corpus companion entry:** Claim C-419 **Audit Claim ID:** vector02-design-duty-section230-claim-01

0.18.41. Appendix B41. Frontier AI weights can be read as condensed claims over collectively produced hu...

Claim statement: That frontier AI weights can be read as condensed claims over collectively produced human expression and therefore justify patrimony-oriented governance. **Evidence status:** Hypothesis / Interpretation **Source support:** Appendix A17, Appendix A6, Appendix A18 **Corpus companion entry:** Claim C-429 **Audit Claim ID:** vector04-weights-patrimony-governance-claim-01

0.18.42. Appendix B42. Current official records show frontier labs and adjacent platform firms building...

Claim statement: That current official records show frontier labs and adjacent platform firms building formal government-facing channels through secure offerings, classified-environment access, accredited cloud stacks, and even military personnel embedding rather than engaging the state only through lobbying or abstract policy rhetoric. **Evidence status:** Hypothesis / Interpretation **Source support:** Appendix A7, Appendix A8, Appendix A21, Appendix A9 **Corpus companion entry:** Claim C-437 **Audit Claim ID:** vector07-formal-government-interface-synthesis-claim-01

0.18.43. Appendix B43. Firms benefiting from export-control protection, strategic state access, or privi...

Claim statement: That firms benefiting from export-control protection, strategic state access, or privileged public coordination should face stronger transparency and public-obligation duties than ordinary software vendors. **Evidence status:** Hypothesis / Interpretation **Source support:** Appendix A20 **Corpus companion entry:** Claim C-438 **Audit Claim ID:** vector07-public-obligations-for-protected-firms-claim-01

0.18.44. Appendix B44. Frontier legitimacy is increasingly stabilized through a combination of mission r...

Claim statement: That frontier legitimacy is increasingly stabilized through a combination of mission rhetoric, elite board composition, security credentials, policy access, and political spending. **Evidence status:** Hypothesis / Interpretation **Source support:** Appendix A22, Appendix A23 **Corpus companion entry:** Claim C-444 **Audit Claim ID:** vector08-legitimacy-stabilized-through-institutions-claim-01

0.18.45. Appendix B45. The legitimacy field around frontier AI is now reinforced not only by boards, pol...

Claim statement: That the legitimacy field around frontier AI is now reinforced not only by boards, policy essays, and lobbying, but by operational state interfaces that normalize privileged access and public-private dependence. **Evidence status:** Hypothesis / Interpretation **Source support:** Appendix A7, Appendix A8, Appendix A21, Appendix A9, Appendix A20 **Corpus companion entry:** Claim C-445 **Audit Claim ID:** vector08-legitimacy-through-operational-state-interface

0.18.46. Appendix B46. The strongest bounded reading of elite fallback architecture is a revealed-prefer...

Claim statement: That the strongest bounded reading of elite fallback architecture is a revealed-preference gap between public stewardship language and private survivability hedging. **Evidence status:** Hypothesis / Interpretation **Source support:** Appendix A26, Appendix A27, Appendix A28 **Corpus companion entry:** Claim C-447 **Audit Claim ID:** vector08-revealed-preference-gap-claim-01

0.18.47. Appendix B47. The continuity scenario includes asymmetry of fallback and exit, with elite actor...

Claim statement: That the continuity scenario includes asymmetry of fallback and exit, with elite actors purchasing redundancy while ordinary users are asked to normalize dependence. **Evidence status:** Hypothesis / Interpretation **Source support:** Appendix A26, Appendix A27, Appendix A28 **Corpus companion entry:** Claim C-449 **Audit Claim ID:** vector09-asymmetric-fallback-continuity-claim-01

0.18.48. Appendix B48. If current platform habits, labor structures, and compute chokepoints continue re...

Claim statement: That if current platform habits, labor structures, and compute chokepoints continue reinforcing one another, the most credible continuity scenario is dependency, access hierarchy, and weaker cognitive sovereignty. **Evidence status:** Hypothesis / Interpretation **Source support:** Appendix A24 **Corpus companion entry:** Claim C-450 **Audit Claim ID:** vector09-continuity-scenario-dependency-hierarchy-claim-01

0.18.49. Appendix B49. Frontier firms receiving material public support, state coordination, export-cont...

Claim statement: That frontier firms receiving material public support, state coordination, export-control protection, or national-security access should face heightened transparency, audit, conflict-of-interest, procurement-visibility, and exclusive-benefit limits. **Evidence status:** Hypothesis / Interpretation **Source support:** Appendix A20, Appendix A25 **Corpus companion entry:** Claim C-452 **Audit Claim ID:** vector10-anti-capture-public-obligations-claim-01

0.18.50. Appendix B50. Sovereign-scale platforms should face a Section 230 framework keyed to systemic r...

Claim statement: That sovereign-scale platforms should face a Section 230 framework keyed to systemic reach, design, amplification, targeting, monetization, moderation architecture, and measurable harm rather than passive-hosting logic alone. **Evidence status:** Hypothesis / Interpretation **Source support:** Appendix A4, Appendix A5 **Corpus companion entry:** Claim C-455 **Audit Claim ID:** vector10-section230-reform-doctrine-claim-01

0.18.51. Appendix B51. Cognitive literacy, labor organization, competition policy, and public-interest i...

Claim statement: That cognitive literacy, labor organization, competition policy, and public-interest infrastructure should function as implementation arms of the dossier's three reforms rather than as disconnected options. **Evidence status:** Hypothesis / Interpretation **Source support:** Appendix A1, Appendix A2, Appendix A3, Appendix A4 **Corpus companion entry:** Claim C-456 **Audit Claim ID:** vector10-three-reforms-implementation-arms-claim-01

0.18.52. Appendix B52. Frontier AI weights trained on collectively produced human expression justify a p...

Claim statement: That frontier AI weights trained on collectively produced human expression justify a patrimony-oriented governance response including public-trust logic, anti-monopoly limits, democratic oversight, and restrictions on exclusive commercialization. **Evidence status:** Hypothesis / Interpretation **Source support:** Appendix A6, Appendix A18, Appendix A17 **Corpus companion entry:** Claim C-457 **Audit Claim ID:** vector10-weights-patrimony-response-claim-01

0.18.53. Appendix B53. Refactoring Section 230 for sovereign-scale platforms requires treating platforms...

Claim statement: That refactoring Section 230 for sovereign-scale platforms requires treating platforms as ordinary publishers for every user utterance regardless of design, amplification, targeting, or monetization. **Evidence status:** Speculative Narrative Risk **Source support:** Appendix A4 **Corpus companion entry:** Claim C-417 **Audit Claim ID:** vector01-publisher-liability-overclaim-claim-01

0.18.54. Appendix B54. Current legal uncertainty or congressional uptake already proves frontier AI weig...

Claim statement: That current legal uncertainty or congressional uptake already proves frontier AI weights are legally settled as patrimony, public trust property, or common heritage. **Evidence status:** Speculative Narrative Risk **Source support:** Appendix A18, Appendix A6 **Corpus companion entry:** Claim C-426 **Audit Claim ID:** vector04-patrimony-already-settled-overclaim-claim-01

0.18.55. Appendix B55. Bunker ownership, survival planning, or retreat architecture prove elite foreknow...

Claim statement: That bunker ownership, survival planning, or retreat architecture prove elite foreknowledge of imminent collapse or that frontier firms are knowingly engineering it. **Evidence status:** Speculative Narrative Risk **Source support:** Appendix A26, Appendix A27, Appendix A28 **Corpus companion entry:** Claim C-440 **Audit Claim ID:** vector08-bunker-foreknowledge-overclaim-claim-01

0.18.56. Appendix B56. Founder philanthropy or public warnings about concentration prove elite foreknowl...

Claim statement: That founder philanthropy or public warnings about concentration prove elite foreknowledge of imminent collapse or absolve frontier firms of responsibility. **Evidence status:** Speculative Narrative Risk **Source support:** Appendix A24 **Corpus companion entry:** Claim C-446 **Audit Claim ID:** vector08-philanthropy-foreknowledge-overclaim-claim-01

0.18.57. Appendix B57. Formal government-only model channels, classified distribution stacks, or militar...

Claim statement: That formal government-only model channels, classified distribution stacks, or military personnel embedding by themselves prove unlawful capture or total hidden state control over frontier AI firms. **Evidence status:** Speculative Narrative Risk **Source support:** Appendix A7, Appendix A8, Appendix A21, Appendix A9 **Corpus companion entry:** Claim C-448 **Audit Claim ID:** vector08-state-interface-hidden-control-overclaim-claim-01

0.18.58. Appendix B58. The continuity scenario is a fixed prophecy of inevitable collapse rather than a...

Claim statement: That the continuity scenario is a fixed prophecy of inevitable collapse rather than a political trajectory that could still be altered by institutional reform. **Evidence status:** Speculative Narrative Risk **Source support:** Appendix A30 **Corpus companion entry:** Claim C-451 **Audit Claim ID:** vector09-inevitable-collapse-overclaim-claim-01

0.18.59. Appendix B59. Individual digital hygiene, consumer choice, or personal restraint alone can solv...

Claim statement: That individual digital hygiene, consumer choice, or personal restraint alone can solve sovereign-scale platform harm, enclosure of frontier weights, or opaque state-company

entanglement without structural reform. **Evidence status:** Speculative Narrative Risk **Source support:** Appendix A1, Appendix A3 **Corpus companion entry:** Claim C-453 **Audit Claim ID:** vector10-personal-hygiene-alone-overclaim-claim-01

0.18.60. Appendix B60. Voluntary elite philanthropy or self-limitation can substitute for structural ref...

Claim statement: That voluntary elite philanthropy or self-limitation can substitute for structural reform of platform power, frontier-weight ownership, and state-company capture. **Evidence status:** Speculative Narrative Risk **Source support:** Appendix A24, Appendix A30 **Corpus companion entry:** Claim C-454 **Audit Claim ID:** vector10-philanthropy-substitute-overclaim-claim-01

0.19. Appendix C. Timeline Slice

0.19.1. Appendix C1. 2016-10-03: The New Yorker profile records Sam Altman's survival-prep measures and New Zealand fa...

Event summary: The New Yorker profile records Sam Altman's survival-prep measures and New Zealand fallback plan with Peter Thiel **Event date:** 2016-10-03 **Source support:** Appendix A26 **Corpus companion entry:** Event T-004 **Audit Event ID:** event-altman-prep-new-zealand-2016-10-03

0.19.2. Appendix C2. 2021-08-05: The role of workers in AI ethics and governance is posted to arXiv

Event summary: The role of workers in AI ethics and governance is posted to arXiv **Event date:** 2021-08-05 **Source support:** Appendix A2 **Corpus companion entry:** Event T-014 **Audit Event ID:** event-workers-ai-governance-2021-08-05

0.19.3. Appendix C3. 2023-05-24: The U.S. Department of Education releases its AI teaching-and-learning recommendations

Event summary: The U.S. Department of Education releases its AI teaching-and-learning recommendations **Event date:** 2023-05-24 **Source support:** Appendix A1 **Corpus companion entry:** Event T-024 **Audit Event ID:** event-us-ed-ai-report-2023-05-24

0.19.4. Appendix C4. 2023-12-14: WIRED documents bunker-like features, secrecy controls, and self-sufficiency planning...

Event summary: WIRED documents bunker-like features, secrecy controls, and self-sufficiency planning at Zuckerberg's Kauai compound **Event date:** 2023-12-14 **Source support:** Appendix A27 **Corpus companion entry:** Event T-035 **Audit Event ID:** event-wired-zuckerberg-kauai-compound-2023-12-14

0.19.5. Appendix C5. 2024-01-25: The FTC opens an inquiry into major AI investments and partnerships

Event summary: The FTC opens an inquiry into major AI investments and partnerships **Event date:** 2024-01-25 **Source support:** Appendix A25 **Corpus companion entry:** Event T-037 **Audit Event ID:** event-ftc-ai-partnerships-inquiry-2024-01-25

0.19.6. Appendix C6. 2024-11-07: Palantir, Anthropic, and AWS announce a classified-environment channel for Claude in...

Event summary: Palantir, Anthropic, and AWS announce a classified-environment channel for Claude in U.S. defense and intelligence work **Event date:** 2024-11-07 **Source support:** Appendix A21 **Corpus companion entry:** Event T-059 **Audit Event ID:** event-palantir-anthropic-aws-government-

0.19.7. Appendix C7. 2024-12-07: Researchers publish a paper on AI and digital public infrastructure

Event summary: Researchers publish a paper on AI and digital public infrastructure **Event date:** 2024-12-07 **Source support:** Appendix A3 **Corpus companion entry:** Event T-061 **Audit Event ID:** event-ai-dpi-paper-2024-12-07

0.19.8. Appendix C8. 2025-06-06: Anthropic launches Claude Gov for U.S. national security customers

Event summary: Anthropic launches Claude Gov for U.S. national security customers **Event date:** 2025-06-06 **Source support:** Appendix A8 **Corpus companion entry:** Event T-077 **Audit Event ID:** event-anthropic-claude-gov-2025-06-06

0.19.9. Appendix C9. 2025-06-16: OpenAI launches OpenAI for Government and announces a DoD CDAO pilot

Event summary: OpenAI launches OpenAI for Government and announces a DoD CDAO pilot **Event date:** 2025-06-16 **Source support:** Appendix A7 **Corpus companion entry:** Event T-079 **Audit Event ID:** event-openai-for-government-2025-06-16

0.19.10. Appendix C10. 2026-01-01: Dario Amodei publicly warns of AI-driven wealth concentration and describes Anthropic...

Event summary: Dario Amodei publicly warns of AI-driven wealth concentration and describes Anthropic's 80% wealth pledge **Event date:** 2026-01-01 **Source support:** Appendix A24 **Corpus companion entry:** Event T-086 **Audit Event ID:** event-dario-adolescence-technology-2026-01-01

0.20. Appendix D. Relevant Entities

0.20.1. Appendix D1. Anthropic

Entity type: company **Role in this paper:** Actor relevant to 04_data_to_models_copyright_weights, 06_ai_double_use_reasoning_deepseek, 07_export_controls_compute_defense_access, 08_power_networks_le and others. **Relevant sources in this paper:** Appendix A19, Appendix A22, Appendix A8, Appendix A21 **Corpus companion entry:** Entity E-008 **Audit Entity ID:** org-anthropic

0.20.2. Appendix D2. Google

Entity type: company **Role in this paper:** Actor relevant to 04_data_to_models_copyright_weights. **Relevant sources in this paper:** Appendix A16 **Corpus companion entry:** Entity E-018

Audit Entity ID: org-google

0.20.3. Appendix D3. OpenAI

Entity type: company **Role in this paper:** Actor relevant to 04_data_to_models_copyright_weights, 06_ai_double_use_reasoning_deepseek, 07_export_controls_compute_defense_access, 08_power_networks_legitimacy_capture, and others. **Relevant sources in this paper:** Appendix A20, Appendix A15, Appendix A22, Appendix A7 **Corpus companion entry:** Entity E-027 **Audit Entity ID:** org-openai

0.20.4. Appendix D4. Palantir Technologies

Entity type: company **Role in this paper:** Actor relevant to 07_export_controls_compute_defense_access, 08_power_networks_legitimacy_capture, 11_facebook_thiel_intelligence_adjacency. **Relevant sources in this paper:** Appendix A21, Appendix A9 **Corpus companion entry:** Entity E-029 **Audit Entity ID:** org-palantir

0.20.5. Appendix D5. Digital public infrastructure

Entity type: governance_framework **Role in this paper:** Actor relevant to 10_counterstrategy_response_agenda. **Relevant sources in this paper:** Appendix A3 **Corpus companion entry:** Entity E-003 **Audit Entity ID:** concept-digital-public-infrastructure

0.20.6. Appendix D6. Federal Election Commission

Entity type: government **Role in this paper:** Actor relevant to 08_power_networks_legitimacy_capture. **Relevant sources in this paper:** Appendix A23 **Corpus companion entry:** Entity E-016 **Audit Entity ID:** org-fec

0.20.7. Appendix D7. U.S. Department of Defense Chief Digital and Artificial Intelligence Office

Entity type: government **Role in this paper:** Actor relevant to 07_export_controls_compute_defense_access, 08_power_networks_legitimacy_capture. **Relevant sources in this paper:** Appendix A7 **Corpus companion entry:** Entity E-014 **Audit Entity ID:** org-dod-cdao

0.20.8. Appendix D8. U.S. Department of Education, Office of Educational Technology

Entity type: government **Role in this paper:** Actor relevant to 09_future_if_unchecked, 10_counterstrategy_response_agenda. **Relevant sources in this paper:** Appendix A1 **Corpus companion entry:** Entity E-038 **Audit Entity ID:** org-us-ed-oet

0.20.9. Appendix D9. United States Army

Entity type: government **Role in this paper:** Actor relevant to 07_export_controls_compute_defense_access, 08_power_networks_legitimacy_capture. **Relevant sources in this paper:** Appendix A9 **Corpus companion entry:** Entity E-035 **Audit Entity ID:** org-us-army

0.20.10. Appendix D10. U.S. Copyright Office

Entity type: government_agency **Role in this paper:** Actor relevant to 04_data_to_models_copyright_weights

Relevant sources in this paper: Appendix A14, Appendix A6 **Corpus companion entry:**

Entity E-037 **Audit Entity ID:** org-us-copyright-office

0.20.11. Appendix D11. Workers as AI governance actors

Entity type: labor **Role in this paper:** Actor relevant to 10_counterstrategy_response_agenda.

Relevant sources in this paper: Appendix A2 **Corpus companion entry:** Entity E-005 **Audit**

Entity ID: group-workers-ai-governance

0.20.12. Appendix D12. Dario Amodei

Entity type: person **Role in this paper:** Actor relevant to 08_power_networks_legitimacy_capture, 09_future_if_unchecked, 10_counterstrategy_response_agenda. **Relevant sources in this**

paper: Appendix A24 **Corpus companion entry:** Entity E-044 **Audit Entity ID:**

person-dario-amodei

0.20.13. Appendix D13. Hassan Uriostegui

Entity type: person **Role in this paper:** Actor relevant to 04_data_to_models_copyright_weights, 08_power_networks_legitimacy_capture. **Relevant sources in this paper:** Appendix A17,

Appendix A29, Appendix A18 **Corpus companion entry:** Entity E-048 **Audit Entity ID:**

person-hassan-uriostegui

0.20.14. Appendix D14. Mark Zuckerberg

Entity type: person **Role in this paper:** Actor relevant to 05_talent_hoarding_section174_layoffs, 08_power_networks_legitimacy_capture, 09_future_if_unchecked, 12_cambridge_analytica_meta_metaverse

Relevant sources in this paper: Appendix A27 **Corpus companion entry:** Entity E-054

Audit Entity ID: person-mark-zuckerberg

0.20.15. Appendix D15. Peter Thiel

Entity type: person **Role in this paper:** Actor relevant to 01_surveillance_capitalism, 08_power_networks_legitimacy_capture, 09_future_if_unchecked, 11_facebook_thiel_intelligence_adjacency.

Relevant sources in this paper: Appendix A28 **Corpus companion entry:** Entity E-058

Audit Entity ID: person-peter-thiel

0.20.16. Appendix D16. Sam Altman

Entity type: person **Role in this paper:** Actor relevant to 07_export_controls_compute_defense_access, 08_power_networks_legitimacy_capture, 09_future_if_unchecked, 10_counterstrategy_response_agenda

and others. **Relevant sources in this paper:** Appendix A20, Appendix A26 **Corpus companion**

entry: Entity E-059 **Audit Entity ID:** person-sam-altman

0.20.17. Appendix D17. Instagram

Entity type: platform **Role in this paper:** Actor relevant to 02_social_decay_youth_harm_hypersexualization

Relevant sources in this paper: Appendix A13 **Corpus companion entry:** Entity E-021

Audit Entity ID: org-instagram

0.20.18. Appendix D18. Reddit

Entity type: platform **Role in this paper:** Actor relevant to 02_social_decay_youth_harm_hypersexualization
04_data_to_models_copyright_weights. **Relevant sources in this paper:** Appendix A16,

Appendix A15, Appendix A12 **Corpus companion entry:** Entity E-032 **Audit Entity ID:**
org-reddit

0.20.19. Appendix D19. TikTok

Entity type: platform **Role in this paper:** Actor relevant to 02_social_decay_youth_harm_hypersexualization

Relevant sources in this paper: Appendix A10 **Corpus companion entry:** Entity E-034

Audit Entity ID: org-tiktok

0.20.20. Appendix D20. Leading the Future

Entity type: political_committee **Role in this paper:** Actor relevant to 08_power_networks_legitimacy_capture

Relevant sources in this paper: Appendix A23 **Corpus companion entry:** Entity E-024

Audit Entity ID: org-leading-the-future

0.20.21. Appendix D21. U.S. Surgeon General / HHS

Entity type: public_health **Role in this paper:** Actor relevant to 02_social_decay_youth_harm_hypersexualization

Relevant sources in this paper: Appendix A5 **Corpus companion entry:** Entity E-019

Audit Entity ID: org-hhs-surgeon-general

0.20.22. Appendix D22. Federal Trade Commission

Entity type: regulator **Role in this paper:** Actor relevant to 01_surveillance_capitalism. **Relevant**

sources in this paper: Appendix A4 **Corpus companion entry:** Entity E-017 **Audit**

Entity ID: org-ftc

0.20.23. Appendix D23. FTC inquiry into AI investments and partnerships

Entity type: regulatory_action **Role in this paper:** Actor relevant to 10_counterstrategy_response_agenda. **Relevant sources in this paper:** Appendix A25 **Corpus companion**

entry: Entity E-061 **Audit Entity ID:** program-ftc-ai-partnerships-inquiry

0.20.24. Appendix D24. Pew Research Center

Entity type: research_institution **Role in this paper:** Actor relevant to 02_social_decay_youth_harm_hypers

Relevant sources in this paper: Appendix A11 **Corpus companion entry:** Entity E-030

Audit Entity ID: org-pew-research-center

0.21. Appendix E. Evidence Boundaries

0.21.1. Appendix E1. Disputed Matters

- Appendix B38: That whether existing copyright and fair-use doctrine can absorb frontier-model training without consent or compensation remains actively contested across agencies, courts, and litigants.

0.21.2. Appendix E2. Interpretive Boundaries

- Appendix B39: That sovereign-scale platforms should be regulated on the basis of design, amplification, targeting, monetization, moderation architecture, and systemic reach rather than passive-hosting logic.
- Appendix B40: That documented youth exposure, recommendation amplification, trust erosion, and fragmentation support a sovereign-scale duty framework under a refactored Section 230 focused on design, exposure, and systemic harm.
- Appendix B41: That frontier AI weights can be read as condensed claims over collectively produced human expression and therefore justify patrimony-oriented governance.
- Appendix B42: That current official records show frontier labs and adjacent platform firms building formal government-facing channels through secure offerings, classified-environment access, accredited cloud stacks, and even military personnel embedding rather than engaging the state only through lobbying or abstract policy rhetoric.
- Appendix B43: That firms benefiting from export-control protection, strategic state access, or privileged public coordination should face stronger transparency and public-obligation duties than ordinary software vendors.
- Appendix B44: That frontier legitimacy is increasingly stabilized through a combination of mission rhetoric, elite board composition, security credentials, policy access, and political spending.
- Appendix B45: That the legitimacy field around frontier AI is now reinforced not only by boards, policy essays, and lobbying, but by operational state interfaces that normalize privileged access and public-private dependence.
- Appendix B46: That the strongest bounded reading of elite fallback architecture is a revealed-preference gap between public stewardship language and private survivability hedging.
- Appendix B47: That the continuity scenario includes asymmetry of fallback and exit, with elite actors purchasing redundancy while ordinary users are asked to normalize dependence.
- Appendix B48: That if current platform habits, labor structures, and compute chokepoints continue reinforcing one another, the most credible continuity scenario is dependency, access hierarchy, and weaker cognitive sovereignty.
- Appendix B49: That frontier firms receiving material public support, state coordination, export-control protection, or national-security access should face heightened transparency, audit, conflict-of-interest, procurement-visibility, and exclusive-benefit limits.
- Appendix B50: That sovereign-scale platforms should face a Section 230 framework keyed to systemic reach, design, amplification, targeting, monetization, moderation architecture, and measurable harm rather than passive-hosting logic alone.
- Appendix B51: That cognitive literacy, labor organization, competition policy, and public-interest infrastructure should function as implementation arms of the dossier's three reforms rather than as disconnected options.

- Appendix B52: That frontier AI weights trained on collectively produced human expression justify a patrimony-oriented governance response including public-trust logic, anti-monopoly limits, democratic oversight, and restrictions on exclusive commercialization.

0.21.3. Appendix E3. Speculative Narrative Risks

- Appendix B53: That refactoring Section 230 for sovereign-scale platforms requires treating platforms as ordinary publishers for every user utterance regardless of design, amplification, targeting, or monetization.
- Appendix B54: That current legal uncertainty or congressional uptake already proves frontier AI weights are legally settled as patrimony, public trust property, or common heritage.
- Appendix B55: That bunker ownership, survival planning, or retreat architecture prove elite foreknowledge of imminent collapse or that frontier firms are knowingly engineering it.
- Appendix B56: That founder philanthropy or public warnings about concentration prove elite foreknowledge of imminent collapse or absolve frontier firms of responsibility.
- Appendix B57: That formal government-only model channels, classified distribution stacks, or military personnel embedding by themselves prove unlawful capture or total hidden state control over frontier AI firms.
- Appendix B58: That the continuity scenario is a fixed prophecy of inevitable collapse rather than a political trajectory that could still be altered by institutional reform.
- Appendix B59: That individual digital hygiene, consumer choice, or personal restraint alone can solve sovereign-scale platform harm, enclosure of frontier weights, or opaque state-company entanglement without structural reform.
- Appendix B60: That voluntary elite philanthropy or self-limitation can substitute for structural reform of platform power, frontier-weight ownership, and state-company capture.

0.21.4. Appendix E4. Sensitive or Review-Worthy Note Flags

- Appendix A8. Companion note: Note N-068. Risk flag: low. Note focus: That by June 6, 2025 Anthropic had publicly launched a government-specific model line for U.S. national security customers.
- Appendix A7. Companion note: Note N-072. Risk flag: low. Note focus: That by June 16, 2025 OpenAI had formally reorganized and expanded its U.S. public-sector work under an **OpenAI for Government** initiative rather than treating government use as incidental.
- Appendix A21. Companion note: Note N-074. Risk flag: low. Note focus: That by November 7, 2024 Anthropic model access for U.S. intelligence and defense agencies was being publicly distributed through a Palantir-plus-AWS stack.
- Appendix A24. Companion note: Note N-084. Risk flag: review. Note focus: That by January 1, 2026, Anthropic's own CEO publicly described wealth concentration, backlash, and democratic erosion as serious AI-adjacent risks.
- Appendix A26. Companion note: Note N-089. Risk flag: review. Note focus: That by October 3, 2016 Sam Altman publicly described maintaining a survival-prep cache and a New Zealand fallback with Peter Thiel.
- Appendix A30. Companion note: Note N-095. Risk flag: review. Note focus: That by May 15, 2026, the AI legitimacy field had expanded into a formal Vatican encyclical rather than

remaining confined to industry, policy, or technical discourse.

- Appendix A27. Companion note: Note N-097. Risk flag: review. Note focus: That a serious T2 investigation documented bunker-like features, strong secrecy controls, and private redundancy architecture at Zuckerberg’s Kauai compound by December 2023.
- Appendix A3. Companion note: Note N-103. Risk flag: review. Note focus: That a public-interest AI infrastructure layer is intellectually and technically legible within serious research, not an ideological fantasy.
- Appendix A25. Companion note: Note N-104. Risk flag: review. Note focus: That formal competition scrutiny of frontier-AI partnerships was already underway in the United States by January 2024.
- Appendix A1. Companion note: Note N-107. Risk flag: review. Note focus: That an institutional response agenda can legitimately include cognitive literacy, human oversight, and training rather than relying on consumer self-control alone.
- Appendix A2. Companion note: Note N-108. Risk flag: review. Note focus: That workers can be treated as governance actors in AI deployment rather than only as subjects of managerial adoption.
- Appendix A28. Companion note: Note N-111. Risk flag: review. Note focus: That the Thiel/New Zealand strand can be grounded in a serious T2 business-paper source, not only in contextual essays.