

From Surveillance Capitalism to AI Imperialism

Omnibus Dossier Volume

1. Editorial Preface

This omnibus volume joins the full chapter sequence into a single dossier-length publication. It preserves the chapter-level evidence apparatus so the reader can move continuously through the argument while still inspecting the bounded evidence surface for each chapter.

1.1. How to Use This Volume

- Read each chapter as a standalone paper if needed; chapter-local appendices remain attached.
- Use the corpus companion for the full-source, full-claim, full-entity, full-event, and full-note audit surfaces.
- Treat **Documented Fact**, **Disputed Fact**, **Hypothesis / Interpretation**, and **Speculative Narrative Risk** as distinct evidentiary states throughout the volume.

1.2. Shared Citation and Evidence Rules

Body notes point to the real external documents first, then to the chapter-local appendix entries that explain why those sources matter in context. Chapter bibliographies remain local so each chapter still works when extracted from the omnibus.

2. Chapter 00. Opening framework

2.1. Abstract

The fight over artificial intelligence is not only a fight over products, model benchmarks, or corporate market share. By December 17, 2024, the Bipartisan House Task Force on Artificial Intelligence had already recorded the training-data conflict in terms overlapping with the dossier's own framework: reproduction, derivative-work theory, commercialization, consent, credit, and compensation. The dossier therefore does not end in a vague call for better ethics.

2.2. Keywords

Opening framework; surveillance capitalism; AI imperialism; AI weights; patrimony of humanity; evidence-based dossier

2.3. Research Question

How should the dossier’s opening framework connect surveillance capitalism, ownership of AI weights, and the three-reform program?

2.4. Scope and Framing Note

This paper examines opening framework 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 frames the civilizational ownership question and announces the three-reform destination. Adjacent reinforcement appears in Chapter 01, Chapter 04, Chapter 10, but those cross-references are supportive rather than required for basic comprehension.

2.5. Central Thesis

The fight over artificial intelligence is not only a fight over products, model benchmarks, or corporate market share. It is also a fight over who owns the condensed value of collective human cognition. In this dossier, the central question is whether captured human language, art, code, labor, culture, science, and emotion can be turned into privately governed strategic weights without triggering a deeper political challenge. ¹

This chapter also sets the repo’s two writing layers. The files in **chapters/** are the reference-corpus narrative: readable, source-traceable, and explicitly bounded. The files in **papers/** are the durable paper-format outputs: self-contained, numerically cited, linked to a global appendix, and reusable for later publication or adaptation. The two layers serve different audiences, but they defend the same argument. ²

2.6. Evidence and Method Note

This paper is derived from the chapter corpus for Chapter 00 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.

2.7. Introduction

This paper asks: How should the dossier’s opening framework connect surveillance capitalism, ownership of AI weights, and the three-reform program? Its working answer is clear from the start: The fight over artificial intelligence is not only a fight over products, model benchmarks, or corporate market share. It is also a fight over who owns the condensed value of collective human cognition.

¹Sources: TwinChat, AI Copyright Weights: A New Frontier... (n.d.; see Appendix A1); TwinChat, How AI Is Ruining Your Life to Build... (n.d.; see Appendix A2)

²Sources: U.S. House of Representatives, Bipartisan House Task Force Report on... (2024-12-17; see Appendix A3); TwinChat, AI Copyright Weights: A New Frontier... (n.d.; see Appendix A1)

In this dossier, the central question is whether captured human language, art, code, labor, culture, science, and emotion can be turned into privately governed strategic weights without triggering a deeper political challenge. ³

This chapter also sets the repo’s two writing layers. The files in **chapters/** are the reference-corpus narrative: readable, source-traceable, and explicitly bounded. The files in **papers/** are the durable paper-format outputs: self-contained, numerically cited, linked to a global appendix, and reusable for later publication or adaptation. The two layers serve different audiences, but they defend the same argument. ⁴ 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.

2.8. Historical and Institutional Context

By December 17, 2024, the Bipartisan House Task Force on Artificial Intelligence had already recorded the training-data conflict in terms overlapping with the dossier’s own framework: reproduction, derivative-work theory, commercialization, consent, credit, and compensation. Footnote 40 on physical page 116 cites Hassan Uriostegui’s essay directly. That does not prove the argument on the merits, but it does prove institutional uptake. ⁵

2.9. Main Analysis

Paul Lara’s published writing adds the dossier’s second framing line: technofeudalism, debt-heavy buildout, paper valuation, and the smoke-economy vocabulary needed to describe how extraction is financed and politically sold. The dossier uses that line as interpretive framing, not as a substitute for primary evidence. ⁶

Interpretation: the strongest defensible reading is not that every later chapter proves a single grand conspiracy. It is that a coherent line can be drawn from platform extraction, to social weakening, to pandemic acceleration, to model training, to labor correction, to access chokepoints, and finally to legitimacy and capture. The burden of the dossier is to keep the factual layer and the political interpretation distinct while still arguing forcefully from the evidence. ⁷

2.10. Position Within the Series

This paper sits inside a coordinated 11-paper architecture. Its nearest companions are Chapter 01 (Social networks as infrastructure for free data capture), Chapter 04 (AI arrives when the human

³Sources: TwinChat, AI Copyright Weights: A New Frontier... (n.d.; see Appendix A1); TwinChat, How AI Is Ruining Your Life to Build... (n.d.; see Appendix A2)

⁴Sources: U.S. House of Representatives, Bipartisan House Task Force Report on... (2024-12-17; see Appendix A3); TwinChat, AI Copyright Weights: A New Frontier... (n.d.; see Appendix A1)

⁵Sources: U.S. House of Representatives, Bipartisan House Task Force Report on... (2024-12-17; see Appendix A3); TwinChat, AI Copyright Weights: A New Frontier... (n.d.; see Appendix A1)

⁶Sources: Paul Lara, ¿Cómo hackear al tecnofeudalismo? (2026-06-29; see Appendix A4); Paul Lara, La gran deuda por la IA (2026-06-22; see Appendix A5); TwinChat, How AI Is Ruining Your Life to Build... (n.d.; see Appendix A2)

⁷Sources: TwinChat, AI Copyright Weights: A New Frontier... (n.d.; see Appendix A1); TwinChat, How AI Is Ruining Your Life to Build... (n.d.; see Appendix A2); U.S. House of Representatives, Bipartisan House Task Force Report on... (2024-12-17; see Appendix A3)

raw material already exists), Chapter 10 (What to do). In that series logic, Chapter 00 frames the civilizational ownership question and announces the three-reform destination.

2.11. Counterarguments

The opening can overreach if it treats the enclosure argument as already complete in legal doctrine rather than politically and institutionally emergent. The House citation proves relevance, not victory. The Lara framing sharpens the language of the dossier, but it does not replace the need for documentary support chapter by chapter.⁸

2.12. Limits of the Evidence

The opening can overreach if it treats the enclosure argument as already complete in legal doctrine rather than politically and institutionally emergent. The House citation proves relevance, not victory. The Lara framing sharpens the language of the dossier, but it does not replace the need for documentary support chapter by chapter.⁹

2.13. Reform Relevance

The dossier therefore does not end in a vague call for better ethics. It moves toward three explicit reforms:

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

Interpretation: those reforms are not presented as already enacted or legally settled. They are the dossier's normative answer to the documented structure of extraction, enclosure, and capture.¹⁰

2.14. Conclusion

The dossier therefore does not end in a vague call for better ethics. To test that argument, the dossier has to begin before frontier AI. In the architecture of this series, Chapter 00 therefore functions as a self-contained argument while also advancing the cumulative path toward the dossier's three reforms.

2.15. 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.

⁸Sources: U.S. House of Representatives, Bipartisan House Task Force Report on... (2024-12-17; see Appendix A3); Paul Lara, La gran deuda por la IA (2026-06-22; see Appendix A5)

⁹Sources: U.S. House of Representatives, Bipartisan House Task Force Report on... (2024-12-17; see Appendix A3); Paul Lara, La gran deuda por la IA (2026-06-22; see Appendix A5)

¹⁰Sources: TwinChat, AI Copyright Weights: A New Frontier... (n.d.; see Appendix A1); TwinChat, How AI Is Ruining Your Life to Build... (n.d.; see Appendix A2); U.S. House of Representatives, Bipartisan House Task Force Report on... (2024-12-17; see Appendix A3)

The underlying research OS distinguishes among **Documented Fact**, **Disputed Fact**, **Hypothesis / Interpretation**, and **Speculative Narrative Risk**. In Chapter 00, those boundaries remain visible in the prose, in the bibliography, and in the appendix sections that summarize claims and caution points.

2.16. References / Bibliography

1. TwinChat. AI Copyright Weights: A New Frontier in Intellectual Property Law. n.d.. T3 essay. Paper appendix: Appendix A1. Corpus companion entry: Source S-122. Audit ID: **twinchat-ai-copyright-weights**. Relevance: That the dossier has a coherent thesis about how weights condense previously captured human labor.
2. TwinChat. How AI Is Ruining Your Life to Build an Empire. n.d.. T3 essay. Paper appendix: Appendix A2. Corpus companion entry: Source S-123. Audit ID: **twinchat-ai-ruining-your-life**. Relevance: That there is a consistent argumentative line for reading surveillance, social deterioration, and AI within a single political framework.
3. U.S. House of Representatives. Bipartisan House Task Force Report on Artificial Intelligence. 2024-12-17. T1 report. Paper appendix: Appendix A3. Corpus companion entry: Source S-070. Audit ID: **house-ai-task-force-report-2024**. Relevance: That the dossier's copyright / weights / compensation framing entered a bipartisan U.S. congressional policy record by December 17, 2024.
4. Paul Lara. ¿Cómo hackear al tecnofeudalismo?. 2026-06-29. T3 essay. Paper appendix: Appendix A4. Corpus companion entry: Source S-104. Audit ID: **paul-lara-tecnofeudalismo-2026**. Relevance: That by late June 2026 the dossier's response layer had a documented co-framing voice organized around technofeudalism, democratic governance, and digital sovereignty.
5. Paul Lara. La gran deuda por la IA. 2026-06-22. T3 essay. Paper appendix: Appendix A5. Corpus companion entry: Source S-102. Audit ID: **paul-lara-great-ai-debt-2026**. Relevance: Supports a key factual or interpretive move inside this paper.

2.17. Appendix A. Source Register for This Paper

2.17.1. Appendix A1. 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 the dossier has a coherent thesis about how weights condense previously captured human labor. **Corpus companion entry:** Source S-122 **Audit ID:** **twinchat-ai-copyright-weights**

2.17.2. Appendix A2. 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:** That there is a consistent argumentative line for reading surveillance, social deterioration, and AI within a single political framework. **Corpus companion entry:** Source S-123 **Audit ID:** **twinchat-ai-ruining-your-life**

2.17.3. Appendix A3. 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 the dossier’s copyright / weights / compensation framing entered a bipartisan U.S. congressional policy record by December 17, 2024. **Corpus companion entry:** Source S-070 **Audit ID:** house-ai-task-force-report-2024

2.17.4. Appendix A4. ¿Cómo hackear al tecnofeudalismo?

Source citation: ¿Cómo hackear al tecnofeudalismo? **Institution / author:** Paul Lara **Date:** 2026-06-29 **Tier / type:** T3 / essay **Why it matters in this paper:** That by late June 2026 the dossier’s response layer had a documented co-framing voice organized around technofeudalism, democratic governance, and digital sovereignty. **Corpus companion entry:** Source S-104 **Audit ID:** paul-lara-tecnofeudalismo-2026

2.17.5. Appendix A5. La gran deuda por la IA

Source citation: La gran deuda por la IA **Institution / author:** Paul Lara **Date:** 2026-06-22 **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-102 **Audit ID:** paul-lara-great-ai-debt-2026

2.18. Appendix B. Claims Used in This Paper**2.18.1. Appendix B1. The dossier’s copyright / weights / compensation framing entered a bipartisan U.S...**

Claim statement: That the dossier’s copyright / weights / compensation framing entered a bipartisan U.S. congressional policy record by December 17, 2024. **Evidence status:** Documented Fact **Source support:** Appendix A3 **Corpus companion entry:** Claim C-219 **Audit Claim ID:** house-ai-task-force-report-2024-04-data-to-models-copyright-weights-claim-01

2.18.2. Appendix B2. A formal House report recorded the AI-training dispute in terms that directly ove...

Claim statement: That a formal House report recorded the AI-training dispute in terms that directly overlap with the dossier’s concern about reproduction, derivative-work theory, commercialization, consent, credit, and compensation. **Evidence status:** Documented Fact **Source support:** Appendix A3 **Corpus companion entry:** Claim C-220 **Audit Claim ID:** house-ai-task-force-report-2024-04-data-to-models-copyright-weights-claim-02

2.18.3. Appendix B3. Chapter 00 and Chapter 04 can cite institutional uptake of the argument without p...

Claim statement: That Chapter 00 and Chapter 04 can cite institutional uptake of the argument without pretending that congressional citation settles the underlying legal merits. **Evidence sta-**

tus: Documented Fact **Source support:** Appendix A3 **Corpus companion entry:** Claim C-221 **Audit Claim ID:** house-ai-task-force-report-2024-04-data-to-models-copyright-weights-claim-03

2.18.4. Appendix B4. By late June 2026 the dossier’s response layer had a documented co-framing voice...

Claim statement: That by late June 2026 the dossier’s response layer had a documented co-framing voice organized around technofeudalism, democratic governance, and digital sovereignty.

Evidence status: Documented Fact **Source support:** Appendix A4 **Corpus companion entry:** Claim C-346 **Audit Claim ID:** paul-lara-tecnofeudalismo-2026-10-counterstrategy-response-agenda-claim-02

2.18.5. Appendix B5. Chapter 10 can draw on a public interpretive vocabulary that connects critique of...

Claim statement: That Chapter 10 can draw on a public interpretive vocabulary that connects critique of extraction to concrete institutional countermeasures.

Evidence status: Documented Fact **Source support:** Appendix A4 **Corpus companion entry:** Claim C-347 **Audit Claim ID:** paul-lara-tecnofeudalismo-2026-10-counterstrategy-response-agenda-claim-02

2.18.6. Appendix B6. The dossier’s framing partnership with Paul Lara is textually grounded in an exis...

Claim statement: That the dossier’s framing partnership with Paul Lara is textually grounded in an existing public essay sequence rather than improvised after the fact.

Evidence status: Documented Fact **Source support:** Appendix A4 **Corpus companion entry:** Claim C-348 **Audit Claim ID:** paul-lara-tecnofeudalismo-2026-10-counterstrategy-response-agenda-claim-03

2.18.7. Appendix B7. The dossier has a coherent thesis about how weights condense previously captured...

Claim statement: That the dossier has a coherent thesis about how weights condense previously captured human labor.

Evidence status: Documented Fact **Source support:** Appendix A1 **Corpus companion entry:** Claim C-397 **Audit Claim ID:** twinchat-ai-copyright-weights-claim-01

2.18.8. Appendix B8. There is a clear argumentative line linking AI, property, and enclosure

Claim statement: That there is a clear argumentative line linking AI, property, and enclosure.

Evidence status: Documented Fact **Source support:** Appendix A1 **Corpus companion entry:** Claim C-398 **Audit Claim ID:** twinchat-ai-copyright-weights-claim-02

2.18.9. Appendix B9. There is a consistent argumentative line for reading surveillance, social deterio...

Claim statement: That there is a consistent argumentative line for reading surveillance, social deterioration, and AI within a single political framework.

Evidence status: Documented Fact

Source support: Appendix A2 **Corpus companion entry:** Claim C-399 **Audit Claim ID:** twinchat-ai-ruining-your-life-claim-01

2.18.10. Appendix B10. The dossier will not be a simple collection of scandals, but an integrated theory...

Claim statement: That the dossier will not be a simple collection of scandals, but an integrated theory of the process. **Evidence status:** Documented Fact **Source support:** Appendix A2 **Corpus companion entry:** Claim C-400 **Audit Claim ID:** twinchat-ai-ruining-your-life-claim-02

2.18.11. Appendix B11. 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 A1, Appendix A3 **Corpus companion entry:** Claim C-429 **Audit Claim ID:** vector04-weights-patrimony-governance-claim-01

2.18.12. Appendix B12. 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 A3 **Corpus companion entry:** Claim C-426 **Audit Claim ID:** vector04-patrimony-already-settled-overclaim-claim-01

2.19. Appendix C. Timeline Slice

2.19.1. Appendix C1. 2024-12-17: The Bipartisan House AI Task Force report cites Hassan Uriostegui's weights-and-copyr...

Event summary: The Bipartisan House AI Task Force report cites Hassan Uriostegui's weights-and-copyright argument **Event date:** 2024-12-17 **Source support:** Appendix A3 **Corpus companion entry:** Event T-062 **Audit Event ID:** event-house-ai-task-force-report-2024-12-17

2.20. Appendix D. Relevant Entities

2.20.1. Appendix D1. 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 A1, Appendix A2, Appendix A3 **Corpus companion entry:** Entity E-048 **Audit Entity ID:** person-hassan-uriostegui

2.21. Appendix E. Evidence Boundaries

2.21.1. Appendix E1. Disputed Matters

- None registered for this paper.

2.21.2. Appendix E2. Interpretive Boundaries

- Appendix B11: That frontier AI weights can be read as condensed claims over collectively produced human expression and therefore justify patrimony-oriented governance.

2.21.3. Appendix E3. Speculative Narrative Risks

- Appendix B12: That current legal uncertainty or congressional uptake already proves frontier AI weights are legally settled as patrimony, public trust property, or common heritage.

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

- Appendix A3. Companion note: Note N-031. Risk flag: review. Note focus: That the dossier's copyright / weights / compensation framing entered a bipartisan U.S. congressional policy record by December 17, 2024.
- Appendix A1. Companion note: Note N-035. Risk flag: medium. Note focus: That the dossier has a coherent thesis about how weights condense previously captured human labor.
- Appendix A2. Companion note: Note N-036. Risk flag: medium. Note focus: That there is a consistent argumentative line for reading surveillance, social deterioration, and AI within a single political framework.
- Appendix A4. Companion note: Note N-106. Risk flag: review. Note focus: That by late June 2026 the dossier's response layer had a documented co-framing voice organized around technofeudalism, democratic governance, and digital sovereignty.

3. Chapter 01. Social networks as infrastructure for free data capture

3.1. Abstract

Before AI could be marketed as a new intelligence order, social platforms had already built the behavioral raw material of that order. The FTC’s 2024 report on social media and video-streaming data practices gives this chapter its strongest regulatory anchor. Once social life is reorganized around measurable engagement, extractive targeting, and recommendation systems, platforms stop looking like passive infrastructure and start looking like privately governed civic systems.

3.2. Keywords

Social networks as infrastructure for free data capture; surveillance capitalism; AI imperialism; Section 230 reform; evidence-based dossier

3.3. Research Question

How does social networks as infrastructure for free data capture function within the dossier’s larger argument, and what does the documentary record allow this chapter to establish?

3.4. Scope and Framing Note

This paper examines social networks as infrastructure for free data capture 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 builds the structural case for refactoring Section 230 by documenting recommendation, extraction, and design power. Adjacent reinforcement appears in Chapter 02, Chapter 04, Chapter 10, but those cross-references are supportive rather than required for basic comprehension.

3.5. Central Thesis

Before AI could be marketed as a new intelligence order, social platforms had already built the behavioral raw material of that order. Their core innovation was not simply communication at scale. It was the normalization of a business model in which attention, relationships, preferences, and expression became persistent, monetizable data assets.¹¹

3.6. Evidence and Method Note

This paper is derived from the chapter corpus for Chapter 01 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

¹¹Sources: FTC, A Look Behind the Screens: Examining... (n.d.; see Appendix A1); FTC, FTC Imposes \$5 Billion Penalty and Sw... (2019-07-24; see Appendix A2)

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.

3.7. Introduction

This paper asks: How does social networks as infrastructure for free data capture function within the dossier’s larger argument, and what does the documentary record allow this chapter to establish? Its working answer is clear from the start: Before AI could be marketed as a new intelligence order, social platforms had already built the behavioral raw material of that order. Their core innovation was not simply communication at scale. It was the normalization of a business model in which attention, relationships, preferences, and expression became persistent, monetizable data assets.¹² 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.

3.8. Historical or Institutional Context

The FTC’s 2024 report on social media and video-streaming data practices gives this chapter its strongest regulatory anchor. The platform model is documented as one of extraction, retention, combination, and monetization, not as a neutral relay of user speech. That record is already enough to weaken the old fiction that sovereign-scale platforms are passive intermediaries.¹³

3.9. Main Analysis

The FTC’s 2019 Facebook action matters because Cambridge Analytica was not just a reputational scandal. It became a structural compliance event with board-level oversight duties and a formal federal record of deceptive and insufficiently controlled data governance.¹⁴

Facebook’s own S-1 and later Meta proxy filings document Peter Thiel as an early investor and a long-lived governance actor. That does not prove clandestine command. It does prove continuity between platform surveillance, elite venture governance, and later state-adjacent technology networks.¹⁵

Meta’s 2021 rebrand and 2023 10-K show that the company attached a metaverse-centered strategic narrative to a still-dominant advertising machine. The rebrand did not function as a data-governance reset. It functioned as narrative redirection while the extraction engine remained intact and financially central.¹⁶

The documented platform record therefore shows more than large-scale surveillance. It shows social participation being routed through privately governed interfaces that could rank visibility, shape reach, and condition economic opportunity. Interpretation: this is the earlier precedent for the

¹²Sources: FTC, A Look Behind the Screens: Examining... (n.d.; see Appendix A1); FTC, FTC Imposes \$5 Billion Penalty and Sw... (2019-07-24; see Appendix A2)

¹³Sources: FTC, A Look Behind the Screens: Examining... (n.d.; see Appendix A1)

¹⁴Sources: FTC, FTC Imposes \$5 Billion Penalty and Sw... (2019-07-24; see Appendix A2)

¹⁵Sources: SEC, Facebook, Inc. Registration Statement... (2012-02-01; see Appendix A3); SEC, Meta Platforms, Inc. 2022 Definitive... (2022-04-08; see Appendix A4)

¹⁶Sources: Meta, The Facebook Company Is Now Meta (2021-10-28; see Appendix A5); U.S. Securities and Exchange Commission, Meta Platforms, Inc. 2023 Annual Repo... (2024-02-02; see Appendix A6)

later AI story. Once an intermediary controls the interface layer, it can influence value without producing the underlying social life itself.¹⁷

Interpretation: Chapter 01 begins the case for **Refactor Section 230** by showing that reach, design, targeting, monetization, and recommendation are not incidental platform features. They are the product. A sovereign-scale platform framework should therefore track systemic design power rather than cling to passive-hosting logic built for an earlier internet.¹⁸

3.10. Position Within the Series

This paper sits inside a coordinated 11-paper architecture. Its nearest companions are Chapter 02 (Social erosion and moral reconfiguration), Chapter 04 (AI arrives when the human raw material already exists), Chapter 10 (What to do). In that series logic, Chapter 01 builds the structural case for refactoring Section 230 by documenting recommendation, extraction, and design power.

3.11. Counterarguments

This chapter is strongest on surveillance business models, regulatory intervention, governance continuity, and strategic redirection. It is weaker on any maximal claim that every platform dataset was fully mapped into every later foundation-model training stack. That bridge must be built more carefully in Chapter 04.¹⁹

3.12. Limits of the Evidence

This chapter is strongest on surveillance business models, regulatory intervention, governance continuity, and strategic redirection. It is weaker on any maximal claim that every platform dataset was fully mapped into every later foundation-model training stack. That bridge must be built more carefully in Chapter 04.²⁰

3.13. Reform Relevance

Once social life is reorganized around measurable engagement, extractive targeting, and recommendation systems, platforms stop looking like passive infrastructure and start looking like privately governed civic systems. That is why the dossier's Section 230 argument begins here: sovereign-scale platforms should be judged by the harms and dependencies created by their design, amplification, and monetization architecture.²¹

¹⁷Sources: FTC, A Look Behind the Screens: Examining... (n.d.; see Appendix A1); FTC, FTC Imposes \$5 Billion Penalty and Sw... (2019-07-24; see Appendix A2); U.S. Securities and Exchange Commission, Meta Platforms, Inc. 2023 Annual Repo... (2024-02-02; see Appendix A6)

¹⁸Sources: FTC, A Look Behind the Screens: Examining... (n.d.; see Appendix A1); FTC, FTC Imposes \$5 Billion Penalty and Sw... (2019-07-24; see Appendix A2); U.S. Securities and Exchange Commission, Meta Platforms, Inc. 2023 Annual Repo... (2024-02-02; see Appendix A6)

¹⁹Sources: FTC, A Look Behind the Screens: Examining... (n.d.; see Appendix A1); U.S. Securities and Exchange Commission, Meta Platforms, Inc. 2023 Annual Repo... (2024-02-02; see Appendix A6)

²⁰Sources: FTC, A Look Behind the Screens: Examining... (n.d.; see Appendix A1); U.S. Securities and Exchange Commission, Meta Platforms, Inc. 2023 Annual Repo... (2024-02-02; see Appendix A6)

²¹Sources: FTC, A Look Behind the Screens: Examining... (n.d.; see Appendix A1); FTC, FTC Imposes \$5 Billion Penalty and Sw... (2019-07-24; see Appendix A2)

3.14. Conclusion

Once social life is reorganized around measurable engagement, extractive targeting, and recommendation systems, platforms stop looking like passive infrastructure and start looking like privately governed civic systems. The next problem is not only that platforms captured behavior. In the architecture of this series, Chapter 01 therefore functions as a self-contained argument while also advancing the cumulative path toward the dossier's three reforms.

3.15. 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 01, those boundaries remain visible in the prose, in the bibliography, and in the appendix sections that summarize claims and caution points.

3.16. References / Bibliography

1. FTC. A Look Behind the Screens: Examining the Data Practices of Social Media and Video Streaming Services. n.d.. T1 report. Paper appendix: Appendix A1. Corpus companion entry: Source S-063. Audit ID: **ftc-2024-look-behind-screens**. Relevance: That the platform model relies on intensive extraction and processing of user data.
2. FTC. FTC Imposes \$5 Billion Penalty and Sweeping New Privacy Restrictions on Facebook. 2019-07-24. T1 press_release. Paper appendix: Appendix A2. Corpus companion entry: Source S-064. Audit ID: **ftc-facebook-privacy-2019**. Relevance: That there is formal regulatory backing for saying Facebook monetized and managed personal data within a framework of deceptive practices and insufficient oversight.
3. SEC. Facebook, Inc. Registration Statement on Form S-1. 2012-02-01. T1 filing. Paper appendix: Appendix A3. Corpus companion entry: Source S-052. Audit ID: **facebook-s1-2012**. Relevance: That Peter Thiel had a formal governance role inside Facebook at the time of the IPO filing and that the role dated back to 2005.
4. SEC. Meta Platforms, Inc. 2022 Definitive Proxy Statement. 2022-04-08. T1 proxy_statement. Paper appendix: Appendix A4. Corpus companion entry: Source S-080. Audit ID: **meta-proxy-2022-thiel-retirement**. Relevance: That Peter Thiel's formal role at Facebook / Meta lasted well beyond the early startup period and remained active into 2022.
5. Meta. The Facebook Company Is Now Meta. 2021-10-28. T1 company_statement. Paper appendix: Appendix A5. Corpus companion entry: Source S-081. Audit ID: **meta-rebrand-2021**. Relevance: That the rebrand was formally tied to a metaverse-centered strategic narrative, not merely a cosmetic name change.
6. U.S. Securities and Exchange Commission. Meta Platforms, Inc. 2023 Annual Report (Form 10-K). 2024-02-02. T1 sec_filing. Paper appendix: Appendix A6. Corpus companion entry: Source S-079. Audit ID: **meta-2023-10k**. Relevance: That the metaverse bet was not just branding language; it had a dedicated segment and a continuing loss structure visible in SEC reporting.

3.17. Appendix A. Source Register for This Paper

3.17.1. Appendix A1. 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 the platform model relies on intensive extraction and processing of user data. **Corpus companion entry:** Source S-063 **Audit ID:** ftc-2024-look-behind-screens

3.17.2. Appendix A2. FTC Imposes \$5 Billion Penalty and Sweeping New Privacy Restrictions on Facebook

Source citation: FTC Imposes \$5 Billion Penalty and Sweeping New Privacy Restrictions on Facebook **Institution / author:** FTC **Date:** 2019-07-24 **Tier / type:** T1 / press_release **Why it matters in this paper:** That there is formal regulatory backing for saying Facebook monetized and managed personal data within a framework of deceptive practices and insufficient oversight. **Corpus companion entry:** Source S-064 **Audit ID:** ftc-facebook-privacy-2019

3.17.3. Appendix A3. Facebook, Inc. Registration Statement on Form S-1

Source citation: Facebook, Inc. Registration Statement on Form S-1 **Institution / author:** SEC **Date:** 2012-02-01 **Tier / type:** T1 / filing **Why it matters in this paper:** That Peter Thiel had a formal governance role inside Facebook at the time of the IPO filing and that the role dated back to 2005. **Corpus companion entry:** Source S-052 **Audit ID:** facebook-s1-2012

3.17.4. Appendix A4. Meta Platforms, Inc. 2022 Definitive Proxy Statement

Source citation: Meta Platforms, Inc. 2022 Definitive Proxy Statement **Institution / author:** SEC **Date:** 2022-04-08 **Tier / type:** T1 / proxy_statement **Why it matters in this paper:** That Peter Thiel's formal role at Facebook / Meta lasted well beyond the early startup period and remained active into 2022. **Corpus companion entry:** Source S-080 **Audit ID:** meta-proxy-2022-thiel-retirement

3.17.5. Appendix A5. The Facebook Company Is Now Meta

Source citation: The Facebook Company Is Now Meta **Institution / author:** Meta **Date:** 2021-10-28 **Tier / type:** T1 / company_statement **Why it matters in this paper:** That the rebrand was formally tied to a metaverse-centered strategic narrative, not merely a cosmetic name change. **Corpus companion entry:** Source S-081 **Audit ID:** meta-rebrand-2021

3.17.6. Appendix A6. Meta Platforms, Inc. 2023 Annual Report (Form 10-K)

Source citation: Meta Platforms, Inc. 2023 Annual Report (Form 10-K) **Institution / author:** U.S. Securities and Exchange Commission **Date:** 2024-02-02 **Tier / type:** T1 / sec_filing **Why it matters in this paper:** That the metaverse bet was not just branding language; it had a

dedicated segment and a continuing loss structure visible in SEC reporting. **Corpus companion entry:** Source S-079 **Audit ID:** meta-2023-10k

3.18. Appendix B. Claims Used in This Paper

3.18.1. Appendix B1. Peter Thiel had a formal governance role inside Facebook at the time of the IPO f...

Claim statement: That Peter Thiel had a formal governance role inside Facebook at the time of the IPO filing and that the role dated back to 2005. **Evidence status:** Documented Fact **Source support:** Appendix A3 **Corpus companion entry:** Claim C-156 **Audit Claim ID:** facebook-s1-2012-01-surveillance-capitalism-claim-01

3.18.2. Appendix B2. Facebook itself described Thiel not merely as a commentator or associate, but as...

Claim statement: That Facebook itself described Thiel not merely as a commentator or associate, but as an early investor and standing director. **Evidence status:** Documented Fact **Source support:** Appendix A3 **Corpus companion entry:** Claim C-157 **Audit Claim ID:** facebook-s1-2012-01-surveillance-capitalism-claim-02

3.18.3. Appendix B3. Thiel occupied a structurally meaningful position in Facebook's governance archit...

Claim statement: That Thiel occupied a structurally meaningful position in Facebook's governance architecture, including a preferred-stock-designated seat and an audit-committee role. **Evidence status:** Documented Fact **Source support:** Appendix A3 **Corpus companion entry:** Claim C-158 **Audit Claim ID:** facebook-s1-2012-01-surveillance-capitalism-claim-03

3.18.4. Appendix B4. Facebook's public-market transition preserved concentrated internal control over...

Claim statement: That Facebook's public-market transition preserved concentrated internal control over directors and corporate decisions. **Evidence status:** Documented Fact **Source support:** Appendix A3 **Corpus companion entry:** Claim C-159 **Audit Claim ID:** facebook-s1-2012-01-surveillance-capitalism-claim-04

3.18.5. Appendix B5. Peter Thiel had a formal governance role inside Facebook at the time of the IPO f...

Claim statement: That Peter Thiel had a formal governance role inside Facebook at the time of the IPO filing and that the role dated back to 2005. **Evidence status:** Documented Fact **Source support:** Appendix A3 **Corpus companion entry:** Claim C-160 **Audit Claim ID:** facebook-s1-2012-11-facebook-thiel-intelligence-adjacency-claim-01

3.18.6. Appendix B6. Facebook itself described Thiel not merely as a commentator or associate, but as...

Claim statement: That Facebook itself described Thiel not merely as a commentator or associate, but as an early investor and standing director. **Evidence status:** Documented Fact **Source support:** Appendix A3 **Corpus companion entry:** Claim C-161 **Audit Claim ID:** facebook-s1-2012-11-facebook-thiel-intelligence-adjacency-claim-02

3.18.7. Appendix B7. Thiel occupied a structurally meaningful position in Facebook’s governance archit...

Claim statement: That Thiel occupied a structurally meaningful position in Facebook’s governance architecture, including a preferred-stock-designated seat and an audit-committee role. **Evidence status:** Documented Fact **Source support:** Appendix A3 **Corpus companion entry:** Claim C-162 **Audit Claim ID:** facebook-s1-2012-11-facebook-thiel-intelligence-adjacency-claim-03

3.18.8. Appendix B8. Facebook’s public-market transition preserved concentrated internal control over...

Claim statement: That Facebook’s public-market transition preserved concentrated internal control over directors and corporate decisions. **Evidence status:** Documented Fact **Source support:** Appendix A3 **Corpus companion entry:** Claim C-163 **Audit Claim ID:** facebook-s1-2012-11-facebook-thiel-intelligence-adjacency-claim-04

3.18.9. Appendix B9. The platform model relies on intensive extraction and processing of user data

Claim statement: That the platform model relies on intensive extraction and processing of user data. **Evidence status:** Documented Fact **Source support:** Appendix A1 **Corpus companion entry:** Claim C-196 **Audit Claim ID:** ftc-2024-look-behind-screens-claim-01

3.18.10. Appendix B10. Platforms cannot seriously be described as simple passive intermediaries

Claim statement: That platforms cannot seriously be described as simple passive intermediaries. **Evidence status:** Documented Fact **Source support:** Appendix A1 **Corpus companion entry:** Claim C-197 **Audit Claim ID:** ftc-2024-look-behind-screens-claim-02

3.18.11. Appendix B11. There is regulatory support for speaking of commercial surveillance at massive scale

Claim statement: That there is regulatory support for speaking of commercial surveillance at massive scale. **Evidence status:** Documented Fact **Source support:** Appendix A1 **Corpus companion entry:** Claim C-198 **Audit Claim ID:** ftc-2024-look-behind-screens-claim-03

3.18.12. Appendix B12. There is formal regulatory backing for saying Facebook monetized and managed pers...

Claim statement: That there is formal regulatory backing for saying Facebook monetized and managed personal data within a framework of deceptive practices and insufficient oversight. **Evidence status:** Documented Fact **Source support:** Appendix A2 **Corpus companion entry:** Claim C-199 **Audit Claim ID:** ftc-facebook-privacy-2019-01-surveillance-capitalism-claim-01

3.18.13. Appendix B13. The Cambridge Analytica scandal was not only a media episode, but a point of inst...

Claim statement: That the Cambridge Analytica scandal was not only a media episode, but a point of institutional intervention with structural compliance consequences. **Evidence status:** Documented Fact **Source support:** Appendix A2 **Corpus companion entry:** Claim C-200 **Audit Claim ID:** ftc-facebook-privacy-2019-01-surveillance-capitalism-claim-02

3.18.14. Appendix B14. Facebook's commercial-surveillance architecture generated enough risk to require...

Claim statement: That Facebook's commercial-surveillance architecture generated enough risk to require extraordinary board-level oversight mechanisms. **Evidence status:** Documented Fact **Source support:** Appendix A2 **Corpus companion entry:** Claim C-201 **Audit Claim ID:** ftc-facebook-privacy-2019-01-surveillance-capitalism-claim-03

3.18.15. Appendix B15. There is formal regulatory backing for saying Facebook monetized and managed pers...

Claim statement: That there is formal regulatory backing for saying Facebook monetized and managed personal data within a framework of deceptive practices and insufficient oversight. **Evidence status:** Documented Fact **Source support:** Appendix A2 **Corpus companion entry:** Claim C-202 **Audit Claim ID:** ftc-facebook-privacy-2019-claim-01

3.18.16. Appendix B16. The Cambridge Analytica scandal was not only a media episode, but a point of inst...

Claim statement: That the Cambridge Analytica scandal was not only a media episode, but a point of institutional intervention with structural compliance consequences. **Evidence status:** Documented Fact **Source support:** Appendix A2 **Corpus companion entry:** Claim C-203 **Audit Claim ID:** ftc-facebook-privacy-2019-claim-02

3.18.17. Appendix B17. Facebook's commercial-surveillance architecture generated enough risk to require...

Claim statement: That Facebook's commercial-surveillance architecture generated enough risk to require extraordinary board-level oversight mechanisms. **Evidence status:** Documented Fact **Source support:** Appendix A2 **Corpus companion entry:** Claim C-204 **Audit Claim ID:** ftc-facebook-privacy-2019-claim-03

3.18.18. Appendix B18. The metaverse bet was not just branding language; it had a dedicated segment and...

Claim statement: That the metaverse bet was not just branding language; it had a dedicated segment and a continuing loss structure visible in SEC reporting. **Evidence status:** Documented Fact **Source support:** Appendix A6 **Corpus companion entry:** Claim C-252 **Audit Claim ID:** meta-2023-10k-01-surveillance-capitalism-claim-01

3.18.19. Appendix B19. Meta’s core economic engine remained advertising, even while management invested...

Claim statement: That Meta’s core economic engine remained advertising, even while management invested heavily in AI and immersive technologies. **Evidence status:** Documented Fact **Source support:** Appendix A6 **Corpus companion entry:** Claim C-253 **Audit Claim ID:** meta-2023-10k-01-surveillance-capitalism-claim-02

3.18.20. Appendix B20. The rebrand period can be discussed in material financial terms rather than only...

Claim statement: That the rebrand period can be discussed in material financial terms rather than only as a symbolic communications move. **Evidence status:** Documented Fact **Source support:** Appendix A6 **Corpus companion entry:** Claim C-254 **Audit Claim ID:** meta-2023-10k-01-surveillance-capitalism-claim-03

3.18.21. Appendix B21. The metaverse bet was not just branding language; it had a dedicated segment and...

Claim statement: That the metaverse bet was not just branding language; it had a dedicated segment and a continuing loss structure visible in SEC reporting. **Evidence status:** Documented Fact **Source support:** Appendix A6 **Corpus companion entry:** Claim C-255 **Audit Claim ID:** meta-2023-10k-claim-01

3.18.22. Appendix B22. Meta’s core economic engine remained advertising, even while management invested...

Claim statement: That Meta’s core economic engine remained advertising, even while management invested heavily in AI and immersive technologies. **Evidence status:** Documented Fact **Source support:** Appendix A6 **Corpus companion entry:** Claim C-256 **Audit Claim ID:** meta-2023-10k-claim-02

3.18.23. Appendix B23. The rebrand period can be discussed in material financial terms rather than only...

Claim statement: That the rebrand period can be discussed in material financial terms rather than only as a symbolic communications move. **Evidence status:** Documented Fact **Source support:** Appendix A6 **Corpus companion entry:** Claim C-257 **Audit Claim ID:** meta-2023-10k-claim-03

3.18.24. Appendix B24. Peter Thiel’s formal role at Facebook / Meta lasted well beyond the early startup...

Claim statement: That Peter Thiel’s formal role at Facebook / Meta lasted well beyond the early startup period and remained active into 2022. **Evidence status:** Documented Fact **Source support:** Appendix A4 **Corpus companion entry:** Claim C-258 **Audit Claim ID:** meta-proxy-2022-thiel-retirement-01-surveillance-capitalism-claim-01

3.18.25. Appendix B25. The company itself presented Thiel’s board tenure as spanning from 2005 to the en...

Claim statement: That the company itself presented Thiel’s board tenure as spanning from 2005 to the end of his term at the 2022 annual meeting. **Evidence status:** Documented Fact **Source support:** Appendix A4 **Corpus companion entry:** Claim C-259 **Audit Claim ID:** meta-proxy-2022-thiel-retirement-01-surveillance-capitalism-claim-02

3.18.26. Appendix B26. Thiel still held committee-level governance authority at Meta immediately before...

Claim statement: That Thiel still held committee-level governance authority at Meta immediately before his departure from the board. **Evidence status:** Documented Fact **Source support:** Appendix A4 **Corpus companion entry:** Claim C-260 **Audit Claim ID:** meta-proxy-2022-thiel-retirement-01-surveillance-capitalism-claim-03

3.18.27. Appendix B27. Continuity between the Facebook surveillance era and the Meta rebrand era include...

Claim statement: That continuity between the Facebook surveillance era and the Meta rebrand era included continuity in elite board membership. **Evidence status:** Documented Fact **Source support:** Appendix A4 **Corpus companion entry:** Claim C-261 **Audit Claim ID:** meta-proxy-2022-thiel-retirement-01-surveillance-capitalism-claim-04

3.18.28. Appendix B28. Peter Thiel’s formal role at Facebook / Meta lasted well beyond the early startup...

Claim statement: That Peter Thiel’s formal role at Facebook / Meta lasted well beyond the early startup period and remained active into 2022. **Evidence status:** Documented Fact **Source support:** Appendix A4 **Corpus companion entry:** Claim C-262 **Audit Claim ID:** meta-proxy-2022-thiel-retirement-11-facebook-thiel-intelligence-adjacency-claim-01

3.18.29. Appendix B29. The company itself presented Thiel’s board tenure as spanning from 2005 to the en...

Claim statement: That the company itself presented Thiel’s board tenure as spanning from 2005 to the end of his term at the 2022 annual meeting. **Evidence status:** Documented Fact **Source support:** Appendix A4 **Corpus companion entry:** Claim C-263 **Audit Claim ID:** meta-proxy-2022-thiel-retirement-11-facebook-thiel-intelligence-adjacency-claim-02

3.18.30. Appendix B30. Thiel still held committee-level governance authority at Meta immediately before...

Claim statement: That Thiel still held committee-level governance authority at Meta immediately before his departure from the board. **Evidence status:** Documented Fact **Source support:** Appendix A4 **Corpus companion entry:** Claim C-264 **Audit Claim ID:** meta-proxy-2022-thiel-retirement-11-facebook-thiel-intelligence-adjacency-claim-03

3.18.31. Appendix B31. Continuity between the Facebook surveillance era and the Meta rebrand era include...

Claim statement: That continuity between the Facebook surveillance era and the Meta rebrand era included continuity in elite board membership. **Evidence status:** Documented Fact **Source support:** Appendix A4 **Corpus companion entry:** Claim C-265 **Audit Claim ID:** meta-proxy-2022-thiel-retirement-11-facebook-thiel-intelligence-adjacency-claim-04

3.18.32. Appendix B32. The rebrand was formally tied to a metaverse-centered strategic narrative, not me...

Claim statement: That the rebrand was formally tied to a metaverse-centered strategic narrative, not merely a cosmetic name change. **Evidence status:** Documented Fact **Source support:** Appendix A5 **Corpus companion entry:** Claim C-266 **Audit Claim ID:** meta-rebrand-2021-01-surveillance-capitalism-claim-01

3.18.33. Appendix B33. Meta itself created a reporting split capable of distinguishing its mature platfo...

Claim statement: That Meta itself created a reporting split capable of distinguishing its mature platform business from its speculative immersive-technology investment. **Evidence status:** Documented Fact **Source support:** Appendix A5 **Corpus companion entry:** Claim C-267 **Audit Claim ID:** meta-rebrand-2021-01-surveillance-capitalism-claim-02

3.18.34. Appendix B34. The company did not present the rebrand as a data-governance reset

Claim statement: That the company did not present the rebrand as a data-governance reset. **Evidence status:** Documented Fact **Source support:** Appendix A5 **Corpus companion entry:** Claim C-268 **Audit Claim ID:** meta-rebrand-2021-01-surveillance-capitalism-claim-03

3.18.35. Appendix B35. The rebrand was formally tied to a metaverse-centered strategic narrative, not me...

Claim statement: That the rebrand was formally tied to a metaverse-centered strategic narrative, not merely a cosmetic name change. **Evidence status:** Documented Fact **Source support:** Appendix A5 **Corpus companion entry:** Claim C-269 **Audit Claim ID:** meta-rebrand-2021-claim-01

3.18.36. Appendix B36. Meta itself created a reporting split capable of distinguishing its mature platfo...

Claim statement: That Meta itself created a reporting split capable of distinguishing its mature platform business from its speculative immersive-technology investment. **Evidence status:** Documented Fact **Source support:** Appendix A5 **Corpus companion entry:** Claim C-270 **Audit Claim ID:** meta-rebrand-2021-claim-02

3.18.37. Appendix B37. The company did not present the rebrand as a data-governance reset

Claim statement: That the company did not present the rebrand as a data-governance reset. **Evidence status:** Documented Fact **Source support:** Appendix A5 **Corpus companion entry:** Claim C-271 **Audit Claim ID:** meta-rebrand-2021-claim-03

3.18.38. Appendix B38. 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 A1, Appendix A6 **Corpus companion entry:** Claim C-418 **Audit Claim ID:** vector01-section230-sovereign-scale-reform-claim-01

3.18.39. Appendix B39. The Facebook-Thiel-Palantir-In-Q-Tel record supports a bounded argument about net...

Claim statement: That the Facebook-Thiel-Palantir-In-Q-Tel record supports a bounded argument about networked governance continuity between surveillance platforms, venture capital, and security-state technology ecosystems. **Evidence status:** Hypothesis / Interpretation **Source support:** Appendix A3, Appendix A4 **Corpus companion entry:** Claim C-459 **Audit Claim ID:** vector11-networked-adjacency-synthesis-claim-01

3.18.40. Appendix B40. The strongest reading of the Facebook-Thiel-Palantir-In-Q-Tel record is a structu...

Claim statement: That the strongest reading of the Facebook-Thiel-Palantir-In-Q-Tel record is a structured adjacency model built through board power, venture sponsorship, and government contracting rather than evidence of covert command. **Evidence status:** Hypothesis / Interpretation **Source support:** Appendix A3, Appendix A4 **Corpus companion entry:** Claim C-461 **Audit Claim ID:** vector11-structured-adjacency-not-command-claim-01

3.18.41. Appendix B41. The Meta rebrand and metaverse cycle can be read as a narrative and capital-alloc...

Claim statement: That the Meta rebrand and metaverse cycle can be read as a narrative and capital-allocation redirection after a reputational-regulatory rupture, even if the precise intent be-

hind the timing cannot be proved conclusively. **Evidence status:** Hypothesis / Interpretation **Source support:** Appendix A2, Appendix A5, Appendix A6 **Corpus companion entry:** Claim C-463 **Audit Claim ID:** vector12-rebrand-redirection-synthesis-claim-01

3.18.42. Appendix B42. 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 A1 **Corpus companion entry:** Claim C-417 **Audit Claim ID:** vector01-publisher-liability-overclaim-claim-01

3.18.43. Appendix B43. Intelligence adjacency or early-investor overlap alone proves direct CIA control...

Claim statement: That intelligence adjacency or early-investor overlap alone proves direct CIA control of Facebook. **Evidence status:** Speculative Narrative Risk **Source support:** Appendix A3 **Corpus companion entry:** Claim C-458 **Audit Claim ID:** vector11-direct-cia-control-facebook-over

3.18.44. Appendix B44. The Meta rebrand by itself proves a singular covert intent to erase Cambridge Ana...

Claim statement: That the Meta rebrand by itself proves a singular covert intent to erase Cambridge Analytica from public memory or that the metaverse pivot was primarily a cover-up operation. **Evidence status:** Speculative Narrative Risk **Source support:** Appendix A2, Appendix A5 **Corpus companion entry:** Claim C-462 **Audit Claim ID:** vector12-coverup-intent-overclaim-claim-01

3.19. Appendix C. Timeline Slice

3.19.1. Appendix C1. 2012-02-01: Facebook's S-1 publicly documents Peter Thiel's board seat and early-investor role

Event summary: Facebook's S-1 publicly documents Peter Thiel's board seat and early-investor role **Event date:** 2012-02-01 **Source support:** Appendix A3 **Corpus companion entry:** Event T-002 **Audit Event ID:** event-facebook-s1-thiel-2012-02-01

3.19.2. Appendix C2. 2019-07-24: The FTC imposes a record penalty and new privacy restrictions on Facebook

Event summary: The FTC imposes a record penalty and new privacy restrictions on Facebook **Event date:** 2019-07-24 **Source support:** Appendix A2 **Corpus companion entry:** Event T-006 **Audit Event ID:** event-ftc-facebook-2019-07-24

3.19.3. Appendix C3. 2021-10-28: Facebook formally rebrands its parent company as Meta

Event summary: Facebook formally rebrands its parent company as Meta **Event date:** 2021-10-28 **Source support:** Appendix A5 **Corpus companion entry:** Event T-015 **Audit Event ID:** event-meta-rebrand-2021-10-28

3.19.4. Appendix C4. 2022-04-08: Meta’s proxy states that Peter Thiel will retire from the board after the annual meeting

Event summary: Meta’s proxy states that Peter Thiel will retire from the board after the annual meeting **Event date:** 2022-04-08 **Source support:** Appendix A4 **Corpus companion entry:** Event T-017 **Audit Event ID:** event-meta-proxy-thiel-retirement-2022-04-08

3.19.5. Appendix C5. 2024-02-02: Meta files its 2023 annual report with continued Reality Labs losses

Event summary: Meta files its 2023 annual report with continued Reality Labs losses **Event date:** 2024-02-02 **Source support:** Appendix A6 **Corpus companion entry:** Event T-040 **Audit Event ID:** event-meta-10k-2024-02-02

3.20. Appendix D. Relevant Entities

3.20.1. Appendix D1. Reality Labs

Entity type: business_unit **Role in this paper:** Actor relevant to 12_cambridge_analytica_meta_metaverse_r05_talent_hoarding_section174_layoffs. **Relevant sources in this paper:** Appendix A5, Appendix A6 **Corpus companion entry:** Entity E-031 **Audit Entity ID:** org-reality-labs

3.20.2. Appendix D2. Facebook / Meta

Entity type: company **Role in this paper:** Actor relevant to 01_surveillance_capitalism, 03_covid_acceleration_and_geopolitics, 04_data_to_models_copyright_weights, 11_facebook_thiel_intelligence_adjacency and others. **Relevant sources in this paper:** Appendix A2, Appendix A3, Appendix A4, Appendix A5 **Corpus companion entry:** Entity E-015 **Audit Entity ID:** org-facebook

3.20.3. Appendix D3. 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_r05_talent_hoarding_section174_layoffs. **Relevant sources in this paper:** Appendix A5 **Corpus companion entry:** Entity E-054 **Audit Entity ID:** person-mark-zuckerberg

3.20.4. Appendix D4. 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 A3, Appendix A4 **Corpus companion entry:** Entity E-058 **Audit Entity ID:** person-peter-thiel

3.20.5. Appendix D5. Federal Trade Commission

Entity type: regulator **Role in this paper:** Actor relevant to 01_surveillance_capitalism. **Relevant sources in this paper:** Appendix A1, Appendix A2 **Corpus companion entry:** Entity E-017 **Audit Entity ID:** org-ftc

3.21. Appendix E. Evidence Boundaries

3.21.1. Appendix E1. Disputed Matters

- None registered for this paper.

3.21.2. Appendix E2. Interpretive Boundaries

- Appendix B38: 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 B39: That the Facebook-Thiel-Palantir-In-Q-Tel record supports a bounded argument about networked governance continuity between surveillance platforms, venture capital, and security-state technology ecosystems.
- Appendix B40: That the strongest reading of the Facebook-Thiel-Palantir-In-Q-Tel record is a structured adjacency model built through board power, venture sponsorship, and government contracting rather than evidence of covert command.
- Appendix B41: That the Meta rebrand and metaverse cycle can be read as a narrative and capital-allocation redirection after a reputational-regulatory rupture, even if the precise intent behind the timing cannot be proved conclusively.

3.21.3. Appendix E3. Speculative Narrative Risks

- Appendix B42: 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 B43: That intelligence adjacency or early-investor overlap alone proves direct CIA control of Facebook.
- Appendix B44: That the Meta rebrand by itself proves a singular covert intent to erase Cambridge Analytica from public memory or that the metaverse pivot was primarily a cover-up operation.

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

- Appendix A3. Companion note: Note N-001. Risk flag: medium. Note focus: That Peter Thiel had a formal governance role inside Facebook at the time of the IPO filing and that the role dated back to 2005.

- Appendix A1. Companion note: Note N-002. Risk flag: low. Note focus: That the platform model relies on intensive extraction and processing of user data.
- Appendix A2. Companion note: Note N-003. Risk flag: low. Note focus: That there is formal regulatory backing for saying Facebook monetized and managed personal data within a framework of deceptive practices and insufficient oversight.
- Appendix A6. Companion note: Note N-005. Risk flag: low. Note focus: That the metaverse bet was not just branding language; it had a dedicated segment and a continuing loss structure visible in SEC reporting.
- Appendix A4. Companion note: Note N-006. Risk flag: medium. Note focus: That Peter Thiel's formal role at Facebook / Meta lasted well beyond the early startup period and remained active into 2022.
- Appendix A5. Companion note: Note N-007. Risk flag: medium. Note focus: That the rebrand was formally tied to a metaverse-centered strategic narrative, not merely a cosmetic name change.
- Appendix A3. Companion note: Note N-110. Risk flag: medium. Note focus: That Peter Thiel had a formal governance role inside Facebook at the time of the IPO filing and that the role dated back to 2005.
- Appendix A4. Companion note: Note N-114. Risk flag: medium. Note focus: That Peter Thiel's formal role at Facebook / Meta lasted well beyond the early startup period and remained active into 2022.
- Appendix A2. Companion note: Note N-123. Risk flag: low. Note focus: That there is formal regulatory backing for saying Facebook monetized and managed personal data within a framework of deceptive practices and insufficient oversight.
- Appendix A6. Companion note: Note N-126. Risk flag: low. Note focus: That the metaverse bet was not just branding language; it had a dedicated segment and a continuing loss structure visible in SEC reporting.
- Appendix A5. Companion note: Note N-127. Risk flag: medium. Note focus: That the rebrand was formally tied to a metaverse-centered strategic narrative, not merely a cosmetic name change.

4. Chapter 02. Social erosion and moral reconfiguration

4.1. Abstract

The platform system did not only capture behavior. The U.S. If attention is trained by systems optimized for engagement rather than flourishing, then the outcome is not only distraction.

4.2. Keywords

Social erosion and moral reconfiguration; surveillance capitalism; AI imperialism; Section 230 reform; evidence-based dossier

4.3. Research Question

How does social erosion and moral reconfiguration function within the dossier’s larger argument, and what does the documentary record allow this chapter to establish?

4.4. Scope and Framing Note

This paper examines social erosion and moral reconfiguration 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 deepens the Section 230 case by connecting amplification systems to youth harm, civic erosion, and moral reconfiguration. Adjacent reinforcement appears in Chapter 01, Chapter 03, Chapter 10, but those cross-references are supportive rather than required for basic comprehension.

4.5. Central Thesis

The platform system did not only capture behavior. It also reorganized norms. The strongest version of the claim is not nostalgic moral panic. It is that algorithmic optimization degraded attention, weakened trust, amplified social fragmentation, and altered what appears normal, desirable, and glamorous at scale.²²

4.6. Evidence and Method Note

This paper is derived from the chapter corpus for Chapter 02 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.

²²Sources: HHS, Social Media and Youth Mental Health:... (n.d.; see Appendix A1); arXiv, What Do We Mean When We Talk about Tr... (2023-02-07; see Appendix A2); arXiv, Quantifying social organization and p... (2020-10-01; see Appendix A3)

4.7. Introduction

This paper asks: How does social erosion and moral reconfiguration function within the dossier’s larger argument, and what does the documentary record allow this chapter to establish? Its working answer is clear from the start: The platform system did not only capture behavior. It also reorganized norms. The strongest version of the claim is not nostalgic moral panic. It is that algorithmic optimization degraded attention, weakened trust, amplified social fragmentation, and altered what appears normal, desirable, and glamorous at scale.²³ 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.

4.8. Historical or Institutional Context

The U.S. Surgeon General’s advisory provides serious institutional backing for discussing cumulative psychosocial risk among young people. That means the youth-harm argument does not rest only on intuition or generational panic.²⁴

4.9. Main Analysis

The Instagram body-image and TikTok under-13 records show that harms are not confined to public feeds. Private-message dynamics, recommendation mismatches, and child-directed product failures all matter. Platform design is therefore part of the harm structure, not just the venue in which the harm appears.²⁵

The trust and polarization literature strengthens the civic layer. Social media trust, content trust, and user trust are analytically distinct, but all are reshaped by platform architecture. Cross-national Pew reporting adds that large majorities across surveyed countries associate social media with manipulation, division, and reduced civility, while the United States appears as an unusually negative outlier.²⁶

The dossier’s more morally charged cultural claims also have bounded support. Serious reporting documents recommendation systems intensifying misogynistic content exposure. Research on OnlyFans documents mainstream entry by creators with no prior sex-work history and ties that normalization to platform visibility and pandemic-era adoption. These examples should remain illustrative, but they are not empty rhetoric.²⁷

Interpretation: Chapter 02 deepens the case for **Refactor Section 230** because the social harms here are not reducible to isolated user speech. They emerge from recommendation systems, exposure

²³Sources: HHS, Social Media and Youth Mental Health:... (n.d.; see Appendix A1); arXiv, What Do We Mean When We Talk about Tr... (2023-02-07; see Appendix A2); arXiv, Quantifying social organization and p... (2020-10-01; see Appendix A3)

²⁴Sources: HHS, Social Media and Youth Mental Health:... (n.d.; see Appendix A1)

²⁵Sources: arXiv, Unfiltered: How Teens Engage in Body... (2025-04-02; see Appendix A4); arXiv, When Kids Mode Isn’t For Kids: Invest... (2025-06-30; see Appendix A5)

²⁶Sources: arXiv, What Do We Mean When We Talk about Tr... (2023-02-07; see Appendix A2); arXiv, Quantifying social organization and p... (2020-10-01; see Appendix A3); Pew Research Center, Social Media Seen as Mostly Good for... (2022-12-06; see Appendix A6)

²⁷Sources: The Guardian, Social media algorithms ‘amplifying m... (2024-02-06; see Appendix A7); arXiv, Nudes? Shouldn’t I charge for these?:... (2022-05-20; see Appendix A8)

design, monetized aspiration, and systemic amplification. A sovereign-scale legal framework should therefore focus on design responsibility and measurable harm, not only on moderation after the fact.²⁸

4.10. Position Within the Series

This paper sits inside a coordinated 11-paper architecture. Its nearest companions are Chapter 01 (Social networks as infrastructure for free data capture), Chapter 03 (COVID as a historical accelerator), Chapter 10 (What to do). In that series logic, Chapter 02 deepens the Section 230 case by connecting amplification systems to youth harm, civic erosion, and moral reconfiguration.

4.11. Counterarguments

This chapter should not imply that every moral deformation of contemporary life is fully caused by platforms. The evidence is strongest on youth harm, fragmentation, misogynistic amplification, and monetized sexual self-presentation. Broader civilizational claims remain interpretive and should stay visibly bounded.²⁹

4.12. Limits of the Evidence

This chapter should not imply that every moral deformation of contemporary life is fully caused by platforms. The evidence is strongest on youth harm, fragmentation, misogynistic amplification, and monetized sexual self-presentation. Broader civilizational claims remain interpretive and should stay visibly bounded.³⁰

4.13. Reform Relevance

If attention is trained by systems optimized for engagement rather than flourishing, then the outcome is not only distraction. It is a weaker public with thinner trust, more manipulable norms, and reduced civic resilience. That matters because a society softened in this way is easier to govern through later layers of digital mediation and AI dependency.³¹

4.14. Conclusion

If attention is trained by systems optimized for engagement rather than flourishing, then the outcome is not only distraction. The next historical shock did not create this architecture. In the architecture of this series, Chapter 02 therefore functions as a self-contained argument while also advancing the cumulative path toward the dossier's three reforms.

²⁸Sources: HHS, Social Media and Youth Mental Health:... (n.d.; see Appendix A1); arXiv, When Kids Mode Isn't For Kids: Invest... (2025-06-30; see Appendix A5); Pew Research Center, Social Media Seen as Mostly Good for... (2022-12-06; see Appendix A6)

²⁹Sources: HHS, Social Media and Youth Mental Health:... (n.d.; see Appendix A1); The Guardian, Social media algorithms 'amplifying m...' (2024-02-06; see Appendix A7)

³⁰Sources: HHS, Social Media and Youth Mental Health:... (n.d.; see Appendix A1); The Guardian, Social media algorithms 'amplifying m...' (2024-02-06; see Appendix A7)

³¹Sources: arXiv, What Do We Mean When We Talk about Tr... (2023-02-07; see Appendix A2); arXiv, Quantifying social organization and p... (2020-10-01; see Appendix A3); Pew Research Center, Social Media Seen as Mostly Good for... (2022-12-06; see Appendix A6)

4.15. 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 02, those boundaries remain visible in the prose, in the bibliography, and in the appendix sections that summarize claims and caution points.

4.16. References / Bibliography

1. HHS. Social Media and Youth Mental Health: The U.S. Surgeon General’s Advisory. n.d.. T1 advisory. Paper appendix: Appendix A1. Corpus companion entry: Source S-069. Audit ID: `hhs-social-media-youth-advisory-2023`. Relevance: That there is serious institutional concern about the impact of social networks on minors.
2. arXiv. What Do We Mean When We Talk about Trust in Social Media? A Systematic Review. 2023-02-07. T1 paper. Paper appendix: Appendix A2. Corpus companion entry: Source S-120. Audit ID: `trust-social-media-systematic-review-2023`. Relevance: That trust in social media is a serious research problem with enough literature to support systematic synthesis.
3. arXiv. Quantifying social organization and political polarization in online platforms. 2020-10-01. T1 paper. Paper appendix: Appendix A3. Corpus companion entry: Source S-128. Audit ID: `waller-anderson-online-polarization-2020`. Relevance: That large-scale platform social organization can be measured as fragmented and politically polarized rather than discussed only impressionistically.
4. arXiv. Unfiltered: How Teens Engage in Body Image and Shaming Discussions via Instagram Direct Messages (DMs). 2025-04-02. T1 paper. Paper appendix: Appendix A4. Corpus companion entry: Source S-072. Audit ID: `instagram-body-image-dms-2025`. Relevance: That teen body-image discourse on Instagram is not confined to public feeds; it also unfolds in private channels that shape social reinforcement and harm.
5. arXiv. When Kids Mode Isn’t For Kids: Investigating TikTok’s Under 13 Experience. 2025-06-30. T1 paper. Paper appendix: Appendix A5. Corpus companion entry: Source S-118. Audit ID: `tiktok-kids-mode-2025`. Relevance: That a child-directed platform mode can still expose minors to content that the researchers did not classify as child-directed.
6. 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 A6. Corpus companion entry: Source S-107. Audit ID: `pew-social-media-democracy-outlier-2022`. Relevance: That cross-national evidence supports the claim that social media is widely associated with manipulation, division, and declining civility, even where some respondents still see...
7. The Guardian. Social media algorithms ‘amplifying misogynistic content’. 2024-02-06. T2 news_article. Paper appendix: Appendix A7. Corpus companion entry: Source S-067. Audit ID: `guardian-tiktok-misogyny-amplification-2024`. Relevance: That serious reporting documents a research finding that recommendation systems can intensify misogynistic content exposure over a short period.
8. arXiv. Nudes? Shouldn’t I charge for these?: Motivations of New Sexual Content Creators

on OnlyFans. 2022-05-20. T1 paper. Paper appendix: Appendix A8. Corpus companion entry: Source S-068. Audit ID: `hamilton-onlyfans-new-creators-2022`. Relevance: That mainstream visibility and platform design helped normalize monetized sexual content creation among people with no prior sex-work history.

4.17. Appendix A. Source Register for This Paper

4.17.1. Appendix A1. 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 there is serious institutional concern about the impact of social networks on minors.
Corpus companion entry: Source S-069 **Audit ID:** `hhs-social-media-youth-advisory-2023`

4.17.2. Appendix A2. What Do We Mean When We Talk about Trust in Social Media? A Systematic Review

Source citation: What Do We Mean When We Talk about Trust in Social Media? A Systematic Review
Institution / author: arXiv **Date:** 2023-02-07 **Tier / type:** T1 / paper **Why it matters in this paper:** That trust in social media is a serious research problem with enough literature to support systematic synthesis. **Corpus companion entry:** Source S-120 **Audit ID:** `trust-social-media-systematic-review-2023`

4.17.3. Appendix A3. 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:** That large-scale platform social organization can be measured as fragmented and politically polarized rather than discussed only impressionistically. **Corpus companion entry:** Source S-128 **Audit ID:** `waller-anderson-online-polarization-2020`

4.17.4. Appendix A4. 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:** That teen body-image discourse on Instagram is not confined to public feeds; it also unfolds in private channels that shape social reinforcement and harm. **Corpus companion entry:** Source S-072 **Audit ID:** `instagram-body-image-dms-2025`

4.17.5. Appendix A5. 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: That a child-directed platform mode can still expose minors to content that the researchers did not classify as child-directed. **Corpus companion entry:** Source S-118 **Audit ID:** tiktok-kids-mode-2025

4.17.6. Appendix A6. 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 cross-national evidence supports the claim that social media is widely associated with manipulation, division, and declining civility, even where some respondents still see... **Corpus companion entry:** Source S-107 **Audit ID:** pew-social-media-democracy-outlier-2022

4.17.7. Appendix A7. Social media algorithms ‘amplifying misogynistic content’

Source citation: Social media algorithms ‘amplifying misogynistic content’ **Institution / author:** The Guardian **Date:** 2024-02-06 **Tier / type:** T2 / news_article **Why it matters in this paper:** That serious reporting documents a research finding that recommendation systems can intensify misogynistic content exposure over a short period. **Corpus companion entry:** Source S-067 **Audit ID:** guardian-tiktok-misogyny-amplification-2024

4.17.8. Appendix A8. Nudes? Shouldn’t I charge for these?: Motivations of New Sexual Content Creators on OnlyFans

Source citation: Nudes? Shouldn’t I charge for these?: Motivations of New Sexual Content Creators on OnlyFans **Institution / author:** arXiv **Date:** 2022-05-20 **Tier / type:** T1 / paper **Why it matters in this paper:** That mainstream visibility and platform design helped normalize monetized sexual content creation among people with no prior sex-work history. **Corpus companion entry:** Source S-068 **Audit ID:** hamilton-onlyfans-new-creators-2022

4.18. Appendix B. Claims Used in This Paper

4.18.1. Appendix B1. Serious reporting documents a research finding that recommendation systems can in...

Claim statement: That serious reporting documents a research finding that recommendation systems can intensify misogynistic content exposure over a short period. **Evidence status:** Documented Fact **Source support:** Appendix A7 **Corpus companion entry:** Claim C-210 **Audit Claim ID:** guardian-tiktok-misogyny-amplification-2024-02-social-decay-youth-harm-hypersexualiza

4.18.2. Appendix B2. Chapter 2 can discuss algorithmic amplification of misogynistic norm formation wi...

Claim statement: That Chapter 2 can discuss algorithmic amplification of misogynistic norm formation with more than anecdotal support. **Evidence status:** Documented Fact

Source support: Appendix A7 **Corpus companion entry:** Claim C-211 **Audit Claim ID:** guardian-tiktok-misogyny-amplification-2024-02-social-decay-youth-harm-hypersexualization-claim

4.18.3. Appendix B3. The dossier can connect youth-facing platform design to moral and social reconfig...

Claim statement: That the dossier can connect youth-facing platform design to moral and social reconfiguration without reducing the argument to a generalized complaint about bad culture. **Evidence status:** Documented Fact **Source support:** Appendix A7 **Corpus companion entry:** Claim C-212 **Audit Claim ID:** guardian-tiktok-misogyny-amplification-2024-02-social-decay-youth-ha

4.18.4. Appendix B4. Mainstream visibility and platform design helped normalize monetized sexual conte...

Claim statement: That mainstream visibility and platform design helped normalize monetized sexual content creation among people with no prior sex-work history. **Evidence status:** Documented Fact **Source support:** Appendix A8 **Corpus companion entry:** Claim C-213 **Audit Claim ID:** hamilton-onlyfans-new-creators-2022-02-social-decay-youth-harm-hypersexualization-cla

4.18.5. Appendix B5. The pandemic was a documented adoption driver in this transition

Claim statement: That the pandemic was a documented adoption driver in this transition. **Evidence status:** Documented Fact **Source support:** Appendix A8 **Corpus companion entry:** Claim C-214 **Audit Claim ID:** hamilton-onlyfans-new-creators-2022-02-social-decay-youth-harm-hyper

4.18.6. Appendix B6. Chapter 2 can discuss digital prostitution glamour and normalization with a real...

Claim statement: That Chapter 2 can discuss digital prostitution glamour and normalization with a real research basis rather than relying only on celebrity examples or moral impression. **Evidence status:** Documented Fact **Source support:** Appendix A8 **Corpus companion entry:** Claim C-215 **Audit Claim ID:** hamilton-onlyfans-new-creators-2022-02-social-decay-youth-harm-hyper

4.18.7. Appendix B7. There is serious institutional concern about the impact of social networks on minors

Claim statement: That there is serious institutional concern about the impact of social networks on minors. **Evidence status:** Documented Fact **Source support:** Appendix A1 **Corpus companion entry:** Claim C-216 **Audit Claim ID:** hhs-social-media-youth-advisory-2023-claim-01

4.18.8. Appendix B8. There is enough basis to discuss psychosocial harm and cumulative risks in young...

Claim statement: That there is enough basis to discuss psychosocial harm and cumulative risks in young populations. **Evidence status:** Documented Fact **Source support:** Appendix A1 **Corpus companion entry:** Claim C-217 **Audit Claim ID:** hhs-social-media-youth-advisory-2023-claim-02

4.18.9. Appendix B9. The discussion about youth does not depend only on moral intuition or cultural no...

Claim statement: That the discussion about youth does not depend only on moral intuition or cultural nostalgia. **Evidence status:** Documented Fact **Source support:** Appendix A1 **Corpus companion entry:** Claim C-218 **Audit Claim ID:** hhs-social-media-youth-advisory-2023-claim-03

4.18.10. Appendix B10. Teen body-image discourse on Instagram is not confined to public feeds; it also u...

Claim statement: That teen body-image discourse on Instagram is not confined to public feeds; it also unfolds in private channels that shape social reinforcement and harm. **Evidence status:** Documented Fact **Source support:** Appendix A4 **Corpus companion entry:** Claim C-225 **Audit Claim ID:** instagram-body-image-dms-2025-02-social-decay-youth-harm-hypersexualization-claim-01

4.18.11. Appendix B11. Platform architecture matters: the difference between one-on-one and group settin...

Claim statement: That platform architecture matters: the difference between one-on-one and group settings changes the texture of support, criticism, and shaming. **Evidence status:** Documented Fact **Source support:** Appendix A4 **Corpus companion entry:** Claim C-226 **Audit Claim ID:** instagram-body-image-dms-2025-02-social-decay-youth-harm-hypersexualization-claim-02

4.18.12. Appendix B12. Chapter 2 can discuss mediated body-image norms with stronger empirical grounding...

Claim statement: That Chapter 2 can discuss mediated body-image norms with stronger empirical grounding than a general moral complaint. **Evidence status:** Documented Fact **Source support:** Appendix A4 **Corpus companion entry:** Claim C-227 **Audit Claim ID:** instagram-body-image-dms-2025-02-social-decay-youth-harm-hypersexualization-claim-03

4.18.13. Appendix B13. Cross-national evidence supports the claim that social media is widely associated...

Claim statement: That cross-national evidence supports the claim that social media is widely associated with manipulation, division, and declining civility, even where some respondents still see democratic benefits. **Evidence status:** Documented Fact **Source support:** Appendix A6 **Corpus companion entry:** Claim C-353 **Audit Claim ID:** pew-social-media-democracy-outlier-2022-02-social-d

4.18.14. Appendix B14. The United States can be treated as an unusually severe case in the trust-and-fra...

Claim statement: That the United States can be treated as an unusually severe case in the trust-and-fragmentation story rather than as just another generic example. **Evidence status:** Documented Fact **Source support:** Appendix A6 **Corpus companion entry:** Claim C-354 **Audit Claim ID:** pew-social-media-democracy-outlier-2022-02-social-decay-youth-harm-hypersexualization

4.18.15. Appendix B15. Chapter 2 can connect platform-mediated social erosion to democratic trust and ci...

Claim statement: That Chapter 2 can connect platform-mediated social erosion to democratic trust and civic coherence with stronger evidence than youth-harm studies alone provide. **Evidence status:** Documented Fact

Source support: Appendix A6 **Corpus companion entry:** Claim C-355 **Audit Claim ID:** pew-social-media-democracy-outlier-2022-02-social-decay-youth-harm-hypersexu

4.18.16. Appendix B16. A child-directed platform mode can still expose minors to content that the resear...

Claim statement: That a child-directed platform mode can still expose minors to content that the researchers did not classify as child-directed. **Evidence status:** Documented Fact

Source support: Appendix A5 **Corpus companion entry:** Claim C-385 **Audit Claim ID:** tiktok-kids-mode-2025-02-social-decay-youth-harm-hypersexualization-claim-01

4.18.17. Appendix B17. Chapter 2 can discuss platform design failures affecting minors with concrete evi...

Claim statement: That Chapter 2 can discuss platform design failures affecting minors with concrete evidence rather than only generalized concern. **Evidence status:** Documented Fact

Source support: Appendix A5 **Corpus companion entry:** Claim C-386 **Audit Claim ID:** tiktok-kids-mode-2025-02-social-decay-youth-harm-hypersexualization-claim-02

4.18.18. Appendix B18. Youth-risk arguments in the dossier can include the mismatch between stated safet...

Claim statement: That youth-risk arguments in the dossier can include the mismatch between stated safety framing and actual curation outcomes. **Evidence status:** Documented Fact

Source support: Appendix A5 **Corpus companion entry:** Claim C-387 **Audit Claim ID:** tiktok-kids-mode-2025-02-social-decay-youth-harm-hypersexualization-claim-03

4.18.19. Appendix B19. Trust in social media is a serious research problem with enough literature to sup...

Claim statement: That trust in social media is a serious research problem with enough literature to support systematic synthesis. **Evidence status:** Documented Fact

Source support: Appendix A2 **Corpus companion entry:** Claim C-391 **Audit Claim ID:** trust-social-media-systematic-review-2023-02-social-decay-youth-harm-hypersexualization-claim-0

4.18.20. Appendix B20. Trust in platforms, users, and content should not be treated as interchangeable c...

Claim statement: That trust in platforms, users, and content should not be treated as interchangeable concepts inside the dossier. **Evidence status:** Documented Fact

Source support: Appendix A2 **Corpus companion entry:** Claim C-392 **Audit Claim ID:** trust-social-media-systematic-review-2023-02-social-decay-youth-harm-hypersexualization-claim-0

4.18.21. Appendix B21. Chapter 2 can discuss erosion of trust and social interpretation with scholarly s...

Claim statement: That Chapter 2 can discuss erosion of trust and social interpretation with scholarly support, while still keeping conceptual precision. **Evidence status:** Documented Fact

Source support: Appendix A2 **Corpus companion entry:** Claim C-393 **Audit Claim ID:**

trust-social-media-systematic-review-2023-02-social-decay-youth-harm-hypersexualization-claim-0

4.18.22. Appendix B22. Large-scale platform social organization can be measured as fragmented and politi...

Claim statement: That large-scale platform social organization can be measured as fragmented and politically polarized rather than discussed only impressionistically. **Evidence status:** Documented Fact

Source support: Appendix A3 **Corpus companion entry:** Claim C-470 **Audit Claim ID:**

waller-anderson-online-polarization-2020-02-social-decay-youth-harm-hypersexualization-claim-0

4.18.23. Appendix B23. Chapter 2 can move beyond youth harm alone and document broader social fragmentat...

Claim statement: That Chapter 2 can move beyond youth harm alone and document broader social fragmentation as a platform-structured phenomenon. **Evidence status:** Documented Fact

Source support: Appendix A3 **Corpus companion entry:** Claim C-471 **Audit Claim ID:**

waller-anderson-online-polarization-2020-02-social-decay-youth-harm-hypersexualization-claim-0

4.18.24. Appendix B24. The strongest version of the argument should focus on system-level polarization d...

Claim statement: That the strongest version of the argument should focus on system-level polarization dynamics, not a simplistic claim that every individual user became equally radicalized. **Evidence status:** Documented Fact

Source support: Appendix A3 **Corpus companion entry:** Claim C-472 **Audit Claim ID:**

waller-anderson-online-polarization-2020-02-social-decay-youth-harm-hypersexualization-claim-0

4.18.25. Appendix B25. 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 A1, Appendix A6 **Corpus companion entry:** Claim C-419 **Audit Claim ID:**

vector02-design-duty-section230-claim-01

4.18.26. Appendix B26. Platform optimization alone fully explains contemporary moral decline, every yout...

Claim statement: That platform optimization alone fully explains contemporary moral decline, every youth pathology, or every downstream cultural change associated with digital life. **Evidence**

status: Speculative Narrative Risk **Source support:** Appendix A1 **Corpus companion entry:** Claim C-420 **Audit Claim ID:** vector02-total-moral-decline-overclaim-claim-01

4.19. Appendix C. Timeline Slice

4.19.1. Appendix C1. 2020-10-01: A large-scale study documents system-level polarization dynamics across Reddit commun...

Event summary: A large-scale study documents system-level polarization dynamics across Reddit communities **Event date:** 2020-10-01 **Source support:** Appendix A3 **Corpus companion entry:** Event T-008 **Audit Event ID:** event-reddit-polarization-study-2020-10-01

4.19.2. Appendix C2. 2022-05-20: A study documents mainstream acceptance as a driver of new OnlyFans sexual-content cr...

Event summary: A study documents mainstream acceptance as a driver of new OnlyFans sexual-content creators **Event date:** 2022-05-20 **Source support:** Appendix A8 **Corpus companion entry:** Event T-018 **Audit Event ID:** event-onlyfans-mainstreaming-study-2022-05-20

4.19.3. Appendix C3. 2022-12-06: Pew publishes cross-national findings on social media, democracy, manipulation, and d...

Event summary: Pew publishes cross-national findings on social media, democracy, manipulation, and division **Event date:** 2022-12-06 **Source support:** Appendix A6 **Corpus companion entry:** Event T-019 **Audit Event ID:** event-pew-social-democracy-outlier-2022-12-06

4.19.4. Appendix C4. 2024-02-06: Researchers report rapid amplification of misogynistic recommendations on TikTok

Event summary: Researchers report rapid amplification of misogynistic recommendations on TikTok **Event date:** 2024-02-06 **Source support:** Appendix A7 **Corpus companion entry:** Event T-041 **Audit Event ID:** event-tiktok-misogyny-amplification-2024-02-06

4.20. Appendix D. Relevant Entities

4.20.1. Appendix D1. Instagram

Entity type: platform **Role in this paper:** Actor relevant to 02_social_decay_youth_harm_hypersexualization **Relevant sources in this paper:** Appendix A4 **Corpus companion entry:** Entity E-021 **Audit Entity ID:** org-instagram

4.20.2. Appendix D2. OnlyFans

Entity type: platform **Role in this paper:** Actor relevant to 02_social_decay_youth_harm_hypersexualization **Relevant sources in this paper:** Appendix A8 **Corpus companion entry:** Entity E-026 **Audit Entity ID:** org-onlyfans

4.20.3. Appendix D3. 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 A3 **Corpus companion entry:** Entity E-032 **Audit Entity ID:** org-reddit

4.20.4. Appendix D4. TikTok

Entity type: platform **Role in this paper:** Actor relevant to 02_social_decay_youth_harm_hypersexualization **Relevant sources in this paper:** Appendix A5, Appendix A7 **Corpus companion entry:** Entity E-034 **Audit Entity ID:** org-tiktok

4.20.5. Appendix D5. 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 A1 **Corpus companion entry:** Entity E-019 **Audit Entity ID:** org-hhs-surgeon-general

4.20.6. Appendix D6. Pew Research Center

Entity type: research_institution **Role in this paper:** Actor relevant to 02_social_decay_youth_harm_hypersexualization **Relevant sources in this paper:** Appendix A6 **Corpus companion entry:** Entity E-030 **Audit Entity ID:** org-pew-research-center

4.21. Appendix E. Evidence Boundaries

4.21.1. Appendix E1. Disputed Matters

- None registered for this paper.

4.21.2. Appendix E2. Interpretive Boundaries

- Appendix B25: 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.

4.21.3. Appendix E3. Speculative Narrative Risks

- Appendix B26: That platform optimization alone fully explains contemporary moral decline, every youth pathology, or every downstream cultural change associated with digital life.

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

- Appendix A7. Companion note: Note N-008. Risk flag: medium. Note focus: That serious reporting documents a research finding that recommendation systems can intensify misogynistic content exposure over a short period.
- Appendix A8. Companion note: Note N-009. Risk flag: medium. Note focus: That mainstream visibility and platform design helped normalize monetized sexual content creation among people with no prior sex-work history.

- Appendix A1. Companion note: Note N-010. Risk flag: low. Note focus: That there is serious institutional concern about the impact of social networks on minors.
- Appendix A4. Companion note: Note N-011. Risk flag: review. Note focus: That teen body-image discourse on Instagram is not confined to public feeds; it also unfolds in private channels that shape social reinforcement and harm.
- Appendix A6. Companion note: Note N-012. Risk flag: review. Note focus: That cross-national evidence supports the claim that social media is widely associated with manipulation, division, and declining civility, even where some respondents still see democratic benefits.
- Appendix A5. Companion note: Note N-013. Risk flag: review. Note focus: That a child-directed platform mode can still expose minors to content that the researchers did not classify as child-directed.
- Appendix A2. Companion note: Note N-014. Risk flag: review. Note focus: That trust in social media is a serious research problem with enough literature to support systematic synthesis.
- Appendix A3. Companion note: Note N-015. Risk flag: medium. Note focus: That large-scale platform social organization can be measured as fragmented and politically polarized rather than discussed only impressionistically.

5. Chapter 03. COVID as a historical accelerator

5.1. Abstract

COVID should be treated as a historical accelerator, not as a proven master plot. The ODNI origin report remains the correct baseline for the most sensitive question. COVID widened the field in which screens became necessary, platforms became indispensable, and digitally mediated life became normalized.

5.2. Keywords

COVID as a historical accelerator; surveillance capitalism; AI imperialism; evidence-based dossier

5.3. Research Question

How does covid as a historical accelerator function within the dossier's larger argument, and what does the documentary record allow this chapter to establish?

5.4. Scope and Framing Note

This paper examines covid as a historical accelerator 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 shows how crisis conditions accelerated digital dependence and normalized the infrastructure later challenged by the reform agenda. Adjacent reinforcement appears in Chapter 01, Chapter 02, Chapter 04, but those cross-references are supportive rather than required for basic comprehension.

5.5. Central Thesis

COVID should be treated as a historical accelerator, not as a proven master plot. The strongest documentary case is that the pandemic consolidated digital dependency, normalized algorithmic mediation, and advantaged dominant platforms in ways that outlasted the emergency itself.³²

5.6. Evidence and Method Note

This paper is derived from the chapter corpus for Chapter 03 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.

³²Sources: ODNI, Report on Potential Links Between the... (n.d.; see Appendix A1); TIME, Big Tech Announces Striking Pandemic... (2021-08-02; see Appendix A2); Axios, Pandemic provides a techlash reprieve (2021-03-10; see Appendix A3)

5.7. Introduction

This paper asks: How does covid as a historical accelerator function within the dossier’s larger argument, and what does the documentary record allow this chapter to establish? Its working answer is clear from the start: COVID should be treated as a historical accelerator, not as a proven master plot. The strongest documentary case is that the pandemic consolidated digital dependency, normalized algorithmic mediation, and advantaged dominant platforms in ways that outlasted the emergency itself.³³ 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.

5.8. Historical or Institutional Context

The ODNI origin report remains the correct baseline for the most sensitive question. It supports a disciplined conclusion: the origin of COVID-19 remained officially disputed and unresolved. That means the dossier can discuss the pandemic’s geopolitical implications, but it cannot honestly present a totalized certainty claim about deliberate design or a proved Cold War operation.³⁴

5.9. Main Analysis

The evidence for acceleration is much stronger. Research on work-from-home durability, remote collaboration, and pandemic-era privacy risk shows that digital mediation was not merely a temporary improvisation. It reshaped institutional life and normalized new forms of oversight and dependence.³⁵

The social layer widened as well. COVID-era research and reporting show that digital platforms became central to health information, misinformation management, social interpretation, and everyday family life. Household and youth screen-time evidence makes that intensification measurable rather than anecdotal.³⁶

The concentration story also became clearer. Pandemic-era gains for large tech firms diverged sharply from small-business strain, and serious reporting documented a temporary reprieve from the pre-COVID techlash because the same firms under scrutiny had become indispensable infrastructure for work, school, commerce, and communication.³⁷

Interpretation: once the effects are separated from the unresolved origin question, a geopolitical reading becomes possible but must remain explicitly interpretive. The pandemic functioned like

³³Sources: ODNI, Report on Potential Links Between the... (n.d.; see Appendix A1); TIME, Big Tech Announces Striking Pandemic... (2021-08-02; see Appendix A2); Axios, Pandemic provides a techlash reprieve (2021-03-10; see Appendix A3)

³⁴Sources: ODNI, Report on Potential Links Between the... (n.d.; see Appendix A1)

³⁵Sources: arXiv, Work-From-Home is Here to Stay: Call... (2022-03-21; see Appendix A4); arXiv, How Work From Home Affects Collaborat... (2020-07-30; see Appendix A5); arXiv, Remote Working Pre- and Post-COVID-19... (2021-07-08; see Appendix A6)

³⁶Sources: arXiv, What social media told about us in th... (2021-01-05; see Appendix A7); arXiv, Excessive Screen Time is Associated w... (2025-08-13; see Appendix A8); Axios, Kids’ screen time up 50% during pandemic (2021-01-17; see Appendix A9)

³⁷Sources: TIME, Big Tech Announces Striking Pandemic... (2021-08-02; see Appendix A2); Axios, Pandemic provides a techlash reprieve (2021-03-10; see Appendix A3)

a force multiplier for mediation, dependence, and concentration even if the deeper origin story remains unresolved.³⁸

5.10. Position Within the Series

This paper sits inside a coordinated 11-paper architecture. Its nearest companions are 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). In that series logic, Chapter 03 shows how crisis conditions accelerated digital dependence and normalized the infrastructure later challenged by the reform agenda.

5.11. Counterarguments

The evidence strongly supports acceleration, dependence, and concentration. It does not support presenting the pandemic as a fully demonstrated act of coordinated strategic design. That stronger claim remains outside the documentary core unless much stronger evidence emerges.³⁹

5.12. Limits of the Evidence

The evidence strongly supports acceleration, dependence, and concentration. It does not support presenting the pandemic as a fully demonstrated act of coordinated strategic design. That stronger claim remains outside the documentary core unless much stronger evidence emerges.⁴⁰

5.13. Reform Relevance

COVID widened the field in which screens became necessary, platforms became indispensable, and digitally mediated life became normalized. That matters because the next phase of generative AI arrived not into a free population but into a population already habituated to living through interfaces.⁴¹

5.14. Conclusion

COVID widened the field in which screens became necessary, platforms became indispensable, and digitally mediated life became normalized. By the time frontier AI arrives at scale, the key precondition is already in place: the human raw material, and much of the dependence needed to exploit it, has already been captured. In the architecture of this series, Chapter 03 therefore functions as a self-contained argument while also advancing the cumulative path toward the dossier's three reforms.

³⁸Sources: ODNI, Report on Potential Links Between the... (n.d.; see Appendix A1); arXiv, Work-From-Home is Here to Stay: Call... (2022-03-21; see Appendix A4); TIME, Big Tech Announces Striking Pandemic... (2021-08-02; see Appendix A2)

³⁹Sources: ODNI, Report on Potential Links Between the... (n.d.; see Appendix A1)

⁴⁰Sources: ODNI, Report on Potential Links Between the... (n.d.; see Appendix A1)

⁴¹Sources: arXiv, Work-From-Home is Here to Stay: Call... (2022-03-21; see Appendix A4); arXiv, What social media told about us in th... (2021-01-05; see Appendix A7); Axios, Kids' screen time up 50% during pandemic (2021-01-17; see Appendix A9)

5.15. 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 03, those boundaries remain visible in the prose, in the bibliography, and in the appendix sections that summarize claims and caution points.

5.16. References / Bibliography

1. ODNI. Report on Potential Links Between the Wuhan Institute of Virology and the Origins of COVID-19. n.d.. T1 report. Paper appendix: Appendix A1. Corpus companion entry: Source S-086. Audit ID: `odni-covid-origin-report-2023`. Relevance: That the origin of COVID-19 remained a matter of official uncertainty and disputed assessment.
2. TIME. Big Tech Announces Striking Pandemic Gains as Small Businesses Strain to Find Their Footing. 2021-08-02. T3 news_analysis. Paper appendix: Appendix A2. Corpus companion entry: Source S-119. Audit ID: `time-big-tech-pandemic-gains-2021`. Relevance: That pandemic-era digital concentration can be documented through a concrete divergence between large tech earnings and small-business stress.
3. Axios. Pandemic provides a techlash reprieve. 2021-03-10. T3 news_article. Paper appendix: Appendix A3. Corpus companion entry: Source S-032. Audit ID: `axios-pandemic-techlash-reprieve-2021`. Relevance: That pandemic conditions altered the policy and legitimacy environment around dominant technology platforms, not only their earnings.
4. arXiv. Work-From-Home is Here to Stay: Call for Flexibility in Post-Pandemic Work Policies. 2022-03-21. T1 paper. Paper appendix: Appendix A4. Corpus companion entry: Source S-114. Audit ID: `smite-wfh-here-to-stay-2022`. Relevance: That pandemic-era remote work was not only temporary improvisation; it reshaped post-pandemic organizational policy and worker expectations.
5. arXiv. How Work From Home Affects Collaboration: A Large-Scale Study of Information Workers in a Natural Experiment During COVID-19. 2020-07-30. T1 paper. Paper appendix: Appendix A5. Corpus companion entry: Source S-135. Audit ID: `yang-wfh-collaboration-2020`. Relevance: That COVID forced a historically abrupt transition into remote, digitally mediated work for large numbers of information workers.
6. arXiv. Remote Working Pre- and Post-COVID-19: An Analysis of New Threats and Risks to Security and Privacy. 2021-07-08. T1 paper. Paper appendix: Appendix A6. Corpus companion entry: Source S-084. Audit ID: `nurse-remote-working-security-privacy-2021`. Relevance: That the pandemic-driven move to remote work created documented new privacy and surveillance risks, not just convenience or flexibility gains.
7. arXiv. What social media told about us in the time of COVID-19: a scoping review. 2021-01-05. T1 paper. Paper appendix: Appendix A7. Corpus companion entry: Source S-121. Audit ID: `tsao-social-media-covid-scoping-review-2021`. Relevance: That pandemic dependence on digital platforms was not limited to remote work.
8. arXiv. Excessive Screen Time is Associated with Mental Health Problems and ADHD in US Children and Adolescents: Physical Activity and Sleep as Parallel Mediators. 2025-08-13.

T1 paper. Paper appendix: Appendix A8. Corpus companion entry: Source S-046. Audit ID: dai-pandemic-screen-time-mental-health-2025. Relevance: That heavy youth screen exposure was a measurable feature of the pandemic period, not only a cultural impression.

9. Axios. Kids' screen time up 50% during pandemic. 2021-01-17. T2 news_article. Paper appendix: Appendix A9. Corpus companion entry: Source S-028. Audit ID: axios-kids-screen-time-pandemic-2021. Relevance: That pandemic-era digital dependence in family life can be described with a direct pre-pandemic versus pandemic comparison rather than only with pandemic-period association data.

5.17. Appendix A. Source Register for This Paper

5.17.1. Appendix A1. Report on Potential Links Between the Wuhan Institute of Virology and the Origins of COVID-19

Source citation: Report on Potential Links Between the Wuhan Institute of Virology and the Origins of COVID-19 **Institution / author:** ODNI **Date:** n.d. **Tier / type:** T1 / report **Why it matters in this paper:** That the origin of COVID-19 remained a matter of official uncertainty and disputed assessment. **Corpus companion entry:** Source S-086 **Audit ID:** odni-covid-origin-report-2023

5.17.2. Appendix A2. Big Tech Announces Striking Pandemic Gains as Small Businesses Strain to Find Their Footing

Source citation: Big Tech Announces Striking Pandemic Gains as Small Businesses Strain to Find Their Footing **Institution / author:** TIME **Date:** 2021-08-02 **Tier / type:** T3 / news_analysis **Why it matters in this paper:** That pandemic-era digital concentration can be documented through a concrete divergence between large tech earnings and small-business stress. **Corpus companion entry:** Source S-119 **Audit ID:** time-big-tech-pandemic-gains-2021

5.17.3. Appendix A3. Pandemic provides a techlash reprieve

Source citation: Pandemic provides a techlash reprieve **Institution / author:** Axios **Date:** 2021-03-10 **Tier / type:** T3 / news_article **Why it matters in this paper:** That pandemic conditions altered the policy and legitimacy environment around dominant technology platforms, not only their earnings. **Corpus companion entry:** Source S-032 **Audit ID:** axios-pandemic-techlash-reprieve-2021

5.17.4. Appendix A4. Work-From-Home is Here to Stay: Call for Flexibility in Post-Pandemic Work Policies

Source citation: Work-From-Home is Here to Stay: Call for Flexibility in Post-Pandemic Work Policies **Institution / author:** arXiv **Date:** 2022-03-21 **Tier / type:** T1 / paper **Why it matters in this paper:** That pandemic-era remote work was not only temporary improvisation; it reshaped post-pandemic organizational policy and worker expectations. **Corpus companion entry:** Source S-114 **Audit ID:** smite-wfh-here-to-stay-2022

5.17.5. Appendix A5. How Work From Home Affects Collaboration: A Large-Scale Study of Information Workers in a Natural Experiment During COVID-19

Source citation: How Work From Home Affects Collaboration: A Large-Scale Study of Information Workers in a Natural Experiment During COVID-19 **Institution / author:** arXiv **Date:** 2020-07-30 **Tier / type:** T1 / paper **Why it matters in this paper:** That COVID forced a historically abrupt transition into remote, digitally mediated work for large numbers of information workers. **Corpus companion entry:** Source S-135 **Audit ID:** yang-wfh-collaboration-2020

5.17.6. Appendix A6. Remote Working Pre- and Post-COVID-19: An Analysis of New Threats and Risks to Security and Privacy

Source citation: Remote Working Pre- and Post-COVID-19: An Analysis of New Threats and Risks to Security and Privacy **Institution / author:** arXiv **Date:** 2021-07-08 **Tier / type:** T1 / paper **Why it matters in this paper:** That the pandemic-driven move to remote work created documented new privacy and surveillance risks, not just convenience or flexibility gains. **Corpus companion entry:** Source S-084 **Audit ID:** nurse-remote-working-security-privacy-2021

5.17.7. Appendix A7. What social media told about us in the time of COVID-19: a scoping review

Source citation: What social media told about us in the time of COVID-19: a scoping review **Institution / author:** arXiv **Date:** 2021-01-05 **Tier / type:** T1 / paper **Why it matters in this paper:** That pandemic dependence on digital platforms was not limited to remote work. **Corpus companion entry:** Source S-121 **Audit ID:** tsao-social-media-covid-scoping-review-2021

5.17.8. Appendix A8. Excessive Screen Time is Associated with Mental Health Problems and ADHD in US Children and Adolescents: Physical Activity and Sleep as Parallel Mediators

Source citation: Excessive Screen Time is Associated with Mental Health Problems and ADHD in US Children and Adolescents: Physical Activity and Sleep as Parallel Mediators **Institution / author:** arXiv **Date:** 2025-08-13 **Tier / type:** T1 / paper **Why it matters in this paper:** That heavy youth screen exposure was a measurable feature of the pandemic period, not only a cultural impression. **Corpus companion entry:** Source S-046 **Audit ID:** dai-pandemic-screen-time-mental-health-2025

5.17.9. Appendix A9. Kids' screen time up 50% during pandemic

Source citation: Kids' screen time up 50% during pandemic **Institution / author:** Axios **Date:** 2021-01-17 **Tier / type:** T2 / news_article **Why it matters in this paper:** That pandemic-era digital dependence in family life can be described with a direct pre-pandemic versus pandemic comparison rather than only with pandemic-period association data. **Corpus companion entry:** Source S-028 **Audit ID:** axios-kids-screen-time-pandemic-2021

5.18. Appendix B. Claims Used in This Paper

5.18.1. Appendix B1. Pandemic-era digital dependence in family life can be described with a direct pre...

Claim statement: That pandemic-era digital dependence in family life can be described with a direct pre-pandemic versus pandemic comparison rather than only with pandemic-period association data. **Evidence status:** Documented Fact **Source support:** Appendix A9 **Corpus companion entry:** Claim C-084 **Audit Claim ID:** axios-kids-screen-time-pandemic-2021-03-covid-acceleration-and-

5.18.2. Appendix B2. Chapter 3 can now support the claim that COVID increased everyday child screen ex...

Claim statement: That Chapter 3 can now support the claim that COVID increased everyday child screen exposure at the household level. **Evidence status:** Documented Fact **Source support:** Appendix A9 **Corpus companion entry:** Claim C-085 **Audit Claim ID:** axios-kids-screen-time-pandemic-2021-03-covid-acceleration-and-geopolitics-claim-02

5.18.3. Appendix B3. Youth screen intensification during the pandemic was not only a retrospective cul...

Claim statement: That youth screen intensification during the pandemic was not only a retrospective cultural impression. **Evidence status:** Documented Fact **Source support:** Appendix A9 **Corpus companion entry:** Claim C-086 **Audit Claim ID:** axios-kids-screen-time-pandemic-2021-03-cov-

5.18.4. Appendix B4. Pandemic conditions altered the policy and legitimacy environment around dominant...

Claim statement: That pandemic conditions altered the policy and legitimacy environment around dominant technology platforms, not only their earnings. **Evidence status:** Documented Fact **Source support:** Appendix A3 **Corpus companion entry:** Claim C-096 **Audit Claim ID:** axios-pandemic-techlash-reprieve-2021-03-covid-acceleration-and-geopolitics-claim-01

5.18.5. Appendix B5. Chapter 3 can discuss a temporary softening of the pre-COVID techlash as part of...

Claim statement: That Chapter 3 can discuss a temporary softening of the pre-COVID techlash as part of the concentration story rather than relying solely on financial-gains reporting. **Evidence status:** Documented Fact **Source support:** Appendix A3 **Corpus companion entry:** Claim C-097 **Audit Claim ID:** axios-pandemic-techlash-reprieve-2021-03-covid-acceleration-and-geopolitics-

5.18.6. Appendix B6. The concentration narrative can include public dependence and political repriorit...

Claim statement: That the concentration narrative can include public dependence and political reprioritization, not just traditional antitrust metrics. **Evidence status:** Documented Fact **Source support:** Appendix A3 **Corpus companion entry:** Claim C-098 **Audit Claim ID:** axios-pandemic-techlash-reprieve-2021-03-covid-acceleration-and-geopolitics-claim-03

5.18.7. Appendix B7. Heavy youth screen exposure was a measurable feature of the pandemic period, not...

Claim statement: That heavy youth screen exposure was a measurable feature of the pandemic period, not only a cultural impression. **Evidence status:** Documented Fact **Source support:** Appendix A8 **Corpus companion entry:** Claim C-138 **Audit Claim ID:** dai-pandemic-screen-time-mental-health-2025-03-covid-acceleration-and-geopolitics-claim-01

5.18.8. Appendix B8. Chapter 3 can connect COVID-era digital dependence to youth mental-health strain...

Claim statement: That Chapter 3 can connect COVID-era digital dependence to youth mental-health strain with population-scale evidence rather than only anecdote. **Evidence status:** Documented Fact **Source support:** Appendix A8 **Corpus companion entry:** Claim C-139 **Audit Claim ID:** dai-pandemic-screen-time-mental-health-2025-03-covid-acceleration-and-geopolitics-claim-02

5.18.9. Appendix B9. Pandemic mediation extended into everyday child and adolescent life through screen...

Claim statement: That pandemic mediation extended into everyday child and adolescent life through screen-intensive routines linked to sleep and activity disruption. **Evidence status:** Documented Fact **Source support:** Appendix A8 **Corpus companion entry:** Claim C-140 **Audit Claim ID:** dai-pandemic-screen-time-mental-health-2025-03-covid-acceleration-and-geopolitics-claim-03

5.18.10. Appendix B10. The pandemic-driven move to remote work created documented new privacy and surveillance...

Claim statement: That the pandemic-driven move to remote work created documented new privacy and surveillance risks, not just convenience or flexibility gains. **Evidence status:** Documented Fact **Source support:** Appendix A6 **Corpus companion entry:** Claim C-277 **Audit Claim ID:** nurse-remote-working-security-privacy-2021-03-covid-acceleration-and-geopolitics-claim-04

5.18.11. Appendix B11. COVID-era digital dependence can be tied to intensified workplace monitoring and...

Claim statement: That COVID-era digital dependence can be tied to intensified workplace monitoring and governance, not only to consumer screen time. **Evidence status:** Documented Fact **Source support:** Appendix A6 **Corpus companion entry:** Claim C-278 **Audit Claim ID:** nurse-remote-working-security-privacy-2021-03-covid-acceleration-and-geopolitics-claim-02

5.18.12. Appendix B12. Chapter 3 can connect pandemic mediation to surveillance expansion with academic...

Claim statement: That Chapter 3 can connect pandemic mediation to surveillance expansion with academic support. **Evidence status:** Documented Fact **Source support:** Appendix A6 **Corpus companion entry:** Claim C-279 **Audit Claim ID:** nurse-remote-working-security-privacy-2021-03-covid-acceleration-and-geopolitics-claim-05

5.18.13. Appendix B13. The origin of COVID-19 remained a matter of official uncertainty and disputed ass...

Claim statement: That the origin of COVID-19 remained a matter of official uncertainty and disputed assessment. **Evidence status:** Documented Fact **Source support:** Appendix A1 **Corpus companion entry:** Claim C-283 **Audit Claim ID:** odni-covid-origin-report-2023-claim-01

5.18.14. Appendix B14. It is legitimate to treat the origin question as unresolved within the dossier

Claim statement: That it is legitimate to treat the origin question as unresolved within the dossier. **Evidence status:** Documented Fact **Source support:** Appendix A1 **Corpus companion entry:** Claim C-284 **Audit Claim ID:** odni-covid-origin-report-2023-claim-02

5.18.15. Appendix B15. Any closed narrative of total intent lacks sufficient support from this source

Claim statement: That any closed narrative of total intent lacks sufficient support from this source. **Evidence status:** Documented Fact **Source support:** Appendix A1 **Corpus companion entry:** Claim C-285 **Audit Claim ID:** odni-covid-origin-report-2023-claim-03

5.18.16. Appendix B16. Pandemic-era remote work was not only temporary improvisation; it reshaped post-p...

Claim statement: That pandemic-era remote work was not only temporary improvisation; it reshaped post-pandemic organizational policy and worker expectations. **Evidence status:** Documented Fact **Source support:** Appendix A4 **Corpus companion entry:** Claim C-373 **Audit Claim ID:** smite-wfh-here-to-stay-2022-03-covid-acceleration-and-geopolitics-claim-01

5.18.17. Appendix B17. COVID accelerated a durable normalization of digitally mediated work arrangements

Claim statement: That COVID accelerated a durable normalization of digitally mediated work arrangements. **Evidence status:** Documented Fact **Source support:** Appendix A4 **Corpus companion entry:** Claim C-374 **Audit Claim ID:** smite-wfh-here-to-stay-2022-03-covid-acceleration-a

5.18.18. Appendix B18. Chapter 3 can discuss the pandemic as a long-tail institutional accelerator rather...

Claim statement: That Chapter 3 can discuss the pandemic as a long-tail institutional accelerator rather than only as a short emergency episode. **Evidence status:** Documented Fact **Source support:** Appendix A4 **Corpus companion entry:** Claim C-375 **Audit Claim ID:** smite-wfh-here-to-stay-2022-03-covid-acceleration-and-geopolitics-claim-03

5.18.19. Appendix B19. Pandemic-era digital concentration can be documented through a concrete divergenc...

Claim statement: That pandemic-era digital concentration can be documented through a concrete divergence between large tech earnings and small-business stress. **Evidence status:** Documented Fact **Source support:** Appendix A2 **Corpus companion entry:** Claim C-388 **Audit Claim ID:** time-big-tech-pandemic-gains-2021-03-covid-acceleration-and-geopolitics-claim-01

5.18.20. Appendix B20. Chapter 3 can discuss COVID as an accelerator of already dominant digital interme...

Claim statement: That Chapter 3 can discuss COVID as an accelerator of already dominant digital intermediaries rather than only as a public-health disruption. **Evidence status:** Documented Fact **Source support:** Appendix A2 **Corpus companion entry:** Claim C-389 **Audit Claim ID:** time-big-tech-pandemic-gains-2021-03-covid-acceleration-and-geopolitics-claim-02

5.18.21. Appendix B21. The metaverse growth narrative emerged in a moment when platform incumbents were...

Claim statement: That the metaverse growth narrative emerged in a moment when platform incumbents were already benefiting from pandemic-era asymmetry. **Evidence status:** Documented Fact **Source support:** Appendix A2 **Corpus companion entry:** Claim C-390 **Audit Claim ID:** time-big-tech-pandemic-gains-2021-03-covid-acceleration-and-geopolitics-claim-03

5.18.22. Appendix B22. Pandemic dependence on digital platforms was not limited to remote work

Claim statement: That pandemic dependence on digital platforms was not limited to remote work. **Evidence status:** Documented Fact **Source support:** Appendix A7 **Corpus companion entry:** Claim C-394 **Audit Claim ID:** tsao-social-media-covid-scoping-review-2021-03-covid-acceleration-and-geopolitics-claim-04

5.18.23. Appendix B23. Social media became a critical layer of mediation for health information, public...

Claim statement: That social media became a critical layer of mediation for health information, public interpretation, mental-health discourse, and misinformation management during the first COVID wave. **Evidence status:** Documented Fact **Source support:** Appendix A7 **Corpus companion entry:** Claim C-395 **Audit Claim ID:** tsao-social-media-covid-scoping-review-2021-03-covid-acceleration-and-geopolitics-claim-05

5.18.24. Appendix B24. Chapter 3 can support claims about platform-mediated pandemic life with evidence...

Claim statement: That Chapter 3 can support claims about platform-mediated pandemic life with evidence broader than workplace digitization alone. **Evidence status:** Documented Fact **Source support:** Appendix A7 **Corpus companion entry:** Claim C-396 **Audit Claim ID:** tsao-social-media-covid-scoping-review-2021-03-covid-acceleration-and-geopolitics-claim-03

5.18.25. Appendix B25. COVID forced a historically abrupt transition into remote, digitally mediated wor...

Claim statement: That COVID forced a historically abrupt transition into remote, digitally mediated work for large numbers of information workers. **Evidence status:** Documented Fact **Source support:** Appendix A5 **Corpus companion entry:** Claim C-494 **Audit Claim ID:** yang-wfh-collaboration-2020-03-covid-acceleration-and-geopolitics-claim-01

5.18.26. Appendix B26. Pandemic-era mediation cannot be reduced to one simple effect; lockdown condition...

Claim statement: That pandemic-era mediation cannot be reduced to one simple effect; lockdown conditions and work-from-home conditions must be analytically separated. **Evidence status:** Documented Fact **Source support:** Appendix A5 **Corpus companion entry:** Claim C-495 **Audit Claim ID:** yang-wfh-collaboration-2020-03-covid-acceleration-and-geopolitics-claim-02

5.18.27. Appendix B27. Chapter 3 can discuss a documented reorganization of communication patterns durin...

Claim statement: That Chapter 3 can discuss a documented reorganization of communication patterns during the pandemic rather than only making a general claim about “more online life.” **Evidence status:** Documented Fact **Source support:** Appendix A5 **Corpus companion entry:** Claim C-496 **Audit Claim ID:** yang-wfh-collaboration-2020-03-covid-acceleration-and-geopolitics-cl

5.18.28. Appendix B28. The balance of probability between zoonotic spillover and laboratory-related orig...

Claim statement: That the balance of probability between zoonotic spillover and laboratory-related origin remains contested across official and expert communities. **Evidence status:** Disputed Fact **Source support:** Appendix A1 **Corpus companion entry:** Claim C-423 **Audit Claim ID:** vector03-origin-balance-contested-claim-01

5.18.29. Appendix B29. COVID can be treated as a historic accelerator of digital dependence, mediation,...

Claim statement: That COVID can be treated as a historic accelerator of digital dependence, mediation, and concentration without claiming that the pandemic sequence itself was intentionally engineered. **Evidence status:** Hypothesis / Interpretation **Source support:** Appendix A4, Appendix A2, Appendix A7 **Corpus companion entry:** Claim C-421 **Audit Claim ID:** vector03-acceleration-without-intent-claim-01

5.18.30. Appendix B30. Documented pandemic acceleration effects alone prove COVID-19 was a deliberate ge...

Claim statement: That documented pandemic acceleration effects alone prove COVID-19 was a deliberate geopolitical operation or “Cold War act”. **Evidence status:** Speculative Narrative

Risk Source support: Appendix A1 **Corpus companion entry:** Claim C-422 **Audit Claim ID:** vector03-covid-cold-war-act-overclaim-claim-01

5.19. Appendix C. Timeline Slice

5.19.1. Appendix C1. 2020-07-30: A large-scale COVID work-from-home study documents changed collaboration and communic...

Event summary: A large-scale COVID work-from-home study documents changed collaboration and communication patterns **Event date:** 2020-07-30 **Source support:** Appendix A5 **Corpus companion entry:** Event T-007 **Audit Event ID:** event-pandemic-wfh-collaboration-2020-07-30

5.19.2. Appendix C2. 2021-01-05: A scoping review documents social media as core pandemic communication infrastructure

Event summary: A scoping review documents social media as core pandemic communication infrastructure **Event date:** 2021-01-05 **Source support:** Appendix A7 **Corpus companion entry:** Event T-009 **Audit Event ID:** event-covid-social-media-mediation-2021-01-05

5.19.3. Appendix C3. 2021-01-17: Axios reports a sharp pre-pandemic versus pandemic jump in children's screen time

Event summary: Axios reports a sharp pre-pandemic versus pandemic jump in children's screen time **Event date:** 2021-01-17 **Source support:** Appendix A9 **Corpus companion entry:** Event T-010 **Audit Event ID:** event-pandemic-kids-screen-migration-2021-01-17

5.19.4. Appendix C4. 2021-03-10: Pandemic-era dependence gives dominant tech firms a temporary reprieve from the pre-C...

Event summary: Pandemic-era dependence gives dominant tech firms a temporary reprieve from the pre-COVID techlash **Event date:** 2021-03-10 **Source support:** Appendix A3 **Corpus companion entry:** Event T-011 **Audit Event ID:** event-pandemic-techlash-reprieve-2021-03-10

5.19.5. Appendix C5. 2021-07-08: Research highlights new surveillance and privacy risks in pandemic remote work

Event summary: Research highlights new surveillance and privacy risks in pandemic remote work **Event date:** 2021-07-08 **Source support:** Appendix A6 **Corpus companion entry:** Event T-012 **Audit Event ID:** event-pandemic-remote-surveillance-2021-07-08

5.19.6. Appendix C6. 2021-08-02: A pandemic-era earnings cycle highlights the divergence between big tech gains and sm...

Event summary: A pandemic-era earnings cycle highlights the divergence between big tech gains and small-business strain **Event date:** 2021-08-02 **Source support:** Appendix A2 **Corpus companion entry:** Event T-013 **Audit Event ID:** event-big-tech-pandemic-gains-2021-08-02

5.19.7. Appendix C7. 2022-03-21: Cross-company research argues work from home became a normalized post-pandemic expect...

Event summary: Cross-company research argues work from home became a normalized post-pandemic expectation **Event date:** 2022-03-21 **Source support:** Appendix A4 **Corpus companion entry:** Event T-016 **Audit Event ID:** event-post-pandemic-wfh-entitlement-2022-03-21

5.19.8. Appendix C8. 2023-06-23: ODNI publishes its report on possible links between the WIV and the origins of COVID-19

Event summary: ODNI publishes its report on possible links between the WIV and the origins of COVID-19 **Event date:** 2023-06-23 **Source support:** Appendix A1 **Corpus companion entry:** Event T-025 **Audit Event ID:** event-odni-report-2023-06-23

5.19.9. Appendix C9. 2025-08-13: A national study documents heavy youth screen exposure during the pandemic period

Event summary: A national study documents heavy youth screen exposure during the pandemic period **Event date:** 2025-08-13 **Source support:** Appendix A8 **Corpus companion entry:** Event T-083 **Audit Event ID:** event-pandemic-youth-screen-time-nsch-2025-08-13

5.20. Appendix D. Relevant Entities

5.20.1. Appendix D1. Facebook / Meta

Entity type: company **Role in this paper:** Actor relevant to 01_surveillance_capitalism, 03_covid_acceleration_and_geopolitics, 04_data_to_models_copyright_weights, 11_facebook_thiel_intelligence_adjacency and others. **Relevant sources in this paper:** Appendix A2 **Corpus companion entry:** Entity E-015 **Audit Entity ID:** org-facebook

5.20.2. Appendix D2. Office of the Director of National Intelligence

Entity type: government **Role in this paper:** Actor relevant to 03_covid_acceleration_and_geopolitics. **Relevant sources in this paper:** Appendix A1 **Corpus companion entry:** Entity E-025 **Audit Entity ID:** org-odni

5.20.3. Appendix D3. Big Tech

Entity type: industry_block **Role in this paper:** Actor relevant to 03_covid_acceleration_and_geopolitics. **Relevant sources in this paper:** Appendix A3, Appendix A2 **Corpus companion entry:** Entity E-062 **Audit Entity ID:** sector-big-tech

5.21. Appendix E. Evidence Boundaries

5.21.1. Appendix E1. Disputed Matters

- Appendix B28: That the balance of probability between zoonotic spillover and laboratory-related origin remains contested across official and expert communities.

5.21.2. Appendix E2. Interpretive Boundaries

- Appendix B29: That COVID can be treated as a historic accelerator of digital dependence, mediation, and concentration without claiming that the pandemic sequence itself was intentionally engineered.

5.21.3. Appendix E3. Speculative Narrative Risks

- Appendix B30: That documented pandemic acceleration effects alone prove COVID-19 was a deliberate geopolitical operation or “Cold War act”.

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

- Appendix A9. Companion note: Note N-016. Risk flag: medium. Note focus: That pandemic-era digital dependence in family life can be described with a direct pre-pandemic versus pandemic comparison rather than only with pandemic-period association data.
- Appendix A3. Companion note: Note N-017. Risk flag: review. Note focus: That pandemic conditions altered the policy and legitimacy environment around dominant technology platforms, not only their earnings.
- Appendix A8. Companion note: Note N-018. Risk flag: low. Note focus: That heavy youth screen exposure was a measurable feature of the pandemic period, not only a cultural impression.
- Appendix A6. Companion note: Note N-019. Risk flag: review. Note focus: That the pandemic-driven move to remote work created documented new privacy and surveillance risks, not just convenience or flexibility gains.
- Appendix A1. Companion note: Note N-020. Risk flag: medium. Note focus: That the origin of COVID-19 remained a matter of official uncertainty and disputed assessment.
- Appendix A4. Companion note: Note N-021. Risk flag: review. Note focus: That pandemic-era remote work was not only temporary improvisation; it reshaped post-pandemic organizational policy and worker expectations.
- Appendix A2. Companion note: Note N-022. Risk flag: medium. Note focus: That pandemic-era digital concentration can be documented through a concrete divergence between large tech earnings and small-business stress.
- Appendix A7. Companion note: Note N-023. Risk flag: low. Note focus: That pandemic dependence on digital platforms was not limited to remote work.
- Appendix A5. Companion note: Note N-024. Risk flag: review. Note focus: That COVID forced a historically abrupt transition into remote, digitally mediated work for large numbers of information workers.

6. Chapter 04. AI arrives when the human raw material already exists

6.1. Abstract

Foundation models do not emerge from nowhere. The U.S. Once captured human production can be recombined into proprietary weights, the surveillance economy acquires a new layer.

6.2. Keywords

AI arrives when the human raw material already exists; surveillance capitalism; AI imperialism; AI weights; patrimony of humanity; evidence-based dossier

6.3. Research Question

How does ai arrives when the human raw material already exists function within the dossier's larger argument, and what does the documentary record allow this chapter to establish?

6.4. Scope and Framing Note

This paper examines ai arrives when the human raw material already exists 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 provides the central evidentiary basis for treating frontier weights as a patrimony and governance problem rather than a neutral market outcome. Adjacent reinforcement appears in Chapter 00, Chapter 05, Chapter 06, Chapter 10, but those cross-references are supportive rather than required for basic comprehension.

6.5. Central Thesis

Foundation models do not emerge from nowhere. They emerge after a long period in which human expression, cultural production, and socially filtered judgment were captured at scale, weakly compensated if at all, and converted into machine-trainable inputs. ⁴²

6.6. Evidence and Method Note

This paper is derived from the chapter corpus for Chapter 04 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.

⁴²Sources: U.S. Copyright Office, Copyright and Artificial Intelligence... (2025-05-09; see Appendix A1); OpenAI, Language Models are Unsupervised Mult... (2019-02-14; see Appendix A2); AP, Reddit strikes \$60M deal allowing Goo... (2024-02-22; see Appendix A3)

6.7. Introduction

This paper asks: How does ai arrives when the human raw material already exists function within the dossier’s larger argument, and what does the documentary record allow this chapter to establish? Its working answer is clear from the start: Foundation models do not emerge from nowhere. They emerge after a long period in which human expression, cultural production, and socially filtered judgment were captured at scale, weakly compensated if at all, and converted into machine-trainable inputs.⁴³ 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.

6.8. Historical or Institutional Context

The U.S. Copyright Office has already treated AI training on copyrighted works as a major live policy problem rather than a settled technical inevitability. The 2023 inquiry and 2025 report together show that training-input legality, disclosure, and compensation were contested near the start of the generative-AI boom and remain contested now.⁴⁴

6.9. Main Analysis

The technical and legal record also undermines the industry’s preferred abstraction. Research on memorization shows that large language models can reproduce identifiable material from training data, while court records and legal scholarship now separate at least two layers of conflict: training-use doctrine and input-acquisition legality.⁴⁵

The social-to-model bridge is also explicit. GPT-2’s WebText was built from socially filtered outbound links shared on Reddit, and by February 2024 Reddit had moved from diffuse public conversation to direct licensing, striking a major agreement that allowed Google to use Reddit posts for AI training and related product improvement. The extraction economy did not end with unpaid capture; it evolved into direct licensing over accumulated human expression.⁴⁶

The labor layer matters too. Research on annotation and outsourced human labor shows that model building depends not only on scraped creative works but also on significant human verification, correction, and operational support that is often geographically externalized and weakly valued.⁴⁷

Taken together, the record shows a sequence in which collectively produced expression becomes training input, archived conversation becomes licensable model fuel, and additional human labor

⁴³Sources: U.S. Copyright Office, Copyright and Artificial Intelligence... (2025-05-09; see Appendix A1); OpenAI, Language Models are Unsupervised Mult... (2019-02-14; see Appendix A2); AP, Reddit strikes \$60M deal allowing Goo... (2024-02-22; see Appendix A3)

⁴⁴Sources: GovInfo, Copyright and Artificial Intelligence... (2023-08-30; see Appendix A4); U.S. Copyright Office, Copyright and Artificial Intelligence... (2025-05-09; see Appendix A1)

⁴⁵Sources: arXiv, Extracting Training Data from Large L... (2020-12-14; see Appendix A5); arXiv, Generative AI Training and Copyright Law (2025-02-21; 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 A7); U.S. District Court for the Southern District of New York via Justia, The New York Times Company v. Microso... (2025-04-04; see Appendix A8)

⁴⁶Sources: OpenAI, Language Models are Unsupervised Mult... (2019-02-14; see Appendix A2); AP, Reddit strikes \$60M deal allowing Goo... (2024-02-22; see Appendix A3)

⁴⁷Sources: Big Data & Society, The problem with annotation. Human la... (2023-07-25; see Appendix A9)

is used to clean, annotate, and operationalize the resulting systems. Value that first appeared as diffuse public production is progressively reorganized inside proprietary model pipelines.⁴⁸

Interpretation: the strategic move is not only model creation but control of the interface through which machine-packaged intelligence is accessed. In that reading, the shelf is the interface: the prompt box, chatbot, or subscription layer becomes the first point of contact between a buyer and capability distilled from human work. Once that first contact is captured, value can be repriced downward and dependence can deepen even where total technical replacement remains incomplete.⁴⁹

The House AI Task Force report gives this chapter an institutional bridge beyond agencies and courts. By December 17, 2024, a bipartisan congressional report was already recording the training conflict in language overlapping with the dossier's core concerns about commercialization, consent, credit, compensation, and derivative-work theory.⁵⁰

Interpretation: this is the dossier's primary evidentiary base for **AI Weights as Patrimony of Humanity**. If frontier weights are built from collectively produced human expression and labor, then they cannot be treated as ordinary private assets by default. The strongest normative direction pushes toward broad, open, utility-like availability under public-interest rules rather than permanent private scarcity. That is advocated doctrine, not settled law.⁵¹

6.10. Position Within the Series

This paper sits inside a coordinated 11-paper architecture. Its nearest companions are Chapter 00 (Opening framework), Chapter 05 (Talent, overhiring, valuation, and layoffs), Chapter 06 (AI as a dual-use rhetorical weapon), Chapter 10 (What to do). In that series logic, Chapter 04 provides the central evidentiary basis for treating frontier weights as a patrimony and governance problem rather than a neutral market outcome.

6.11. Counterarguments

The current record is strong on conflict, appropriation structure, memorization, licensing, and the social-to-model bridge. It is still weaker on full training-stack provenance across every frontier lab. This chapter should therefore challenge enclosure forcefully without pretending that every corpus is already fully mapped.⁵²

⁴⁸Sources: OpenAI, Language Models are Unsupervised Mult... (2019-02-14; see Appendix A2); AP, Reddit strikes \$60M deal allowing Goo... (2024-02-22; see Appendix A3); Big Data & Society, The problem with annotation. Human la... (2023-07-25; see Appendix A9)

⁴⁹Sources: OpenAI, Language Models are Unsupervised Mult... (2019-02-14; see Appendix A2); AP, Reddit strikes \$60M deal allowing Goo... (2024-02-22; see Appendix A3); Big Data & Society, The problem with annotation. Human la... (2023-07-25; see Appendix A9)

⁵⁰Sources: U.S. House of Representatives, Bipartisan House Task Force Report on... (2024-12-17; see Appendix A10); TwinChat, AI Copyright Weights: A New Frontier... (n.d.; see Appendix A11)

⁵¹Sources: TwinChat, AI Copyright Weights: A New Frontier... (n.d.; see Appendix A11); U.S. Copyright Office, Copyright and Artificial Intelligence... (2025-05-09; see Appendix A1); U.S. House of Representatives, Bipartisan House Task Force Report on... (2024-12-17; see Appendix A10)

⁵²Sources: U.S. Copyright Office, Copyright and Artificial Intelligence... (2025-05-09; see Appendix A1); U.S. District Court for the Southern District of New York via Justia, The New York Times Company v. Microso... (2025-04-04; see Appendix A8)

6.12. Limits of the Evidence

The current record is strong on conflict, appropriation structure, memorization, licensing, and the social-to-model bridge. It is still weaker on full training-stack provenance across every frontier lab. This chapter should therefore challenge enclosure forcefully without pretending that every corpus is already fully mapped.⁵³

6.13. Reform Relevance

Once captured human production can be recombined into proprietary weights, the surveillance economy acquires a new layer. It no longer only predicts behavior. It begins to substitute for the producers whose accumulated traces made the system possible in the first place. That is why the ownership question has to expand into a governance question.⁵⁴

6.14. Conclusion

Once captured human production can be recombined into proprietary weights, the surveillance economy acquires a new layer. The next question is how that enclosure story meets labor, finance, and layoffs. In the architecture of this series, Chapter 04 therefore functions as a self-contained argument while also advancing the cumulative path toward the dossier's three reforms.

6.15. 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 04, those boundaries remain visible in the prose, in the bibliography, and in the appendix sections that summarize claims and caution points.

6.16. References / Bibliography

1. U.S. Copyright Office. Copyright and Artificial Intelligence, Part 3: Generative AI Training (Pre-publication Version). 2025-05-09. T1 report. Paper appendix: Appendix A1. Corpus companion entry: Source S-126. Audit ID: **usco-ai-training-report-2025**. Relevance: That the U.S. Copyright Office treats AI training on copyrighted works as a major live legal and policy problem rather than a settled matter.
2. OpenAI. Language Models are Unsupervised Multitask Learners. 2019-02-14. T1 paper. Paper appendix: Appendix A2. Corpus companion entry: Source S-108. Audit ID: **radford-gpt2-webtext-2019**. Relevance: That a major OpenAI training corpus was explicitly built from socially filtered outbound links shared on Reddit.

⁵³Sources: U.S. Copyright Office, Copyright and Artificial Intelligence... (2025-05-09; see Appendix A1); U.S. District Court for the Southern District of New York via Justia, The New York Times Company v. Microso... (2025-04-04; see Appendix A8)

⁵⁴Sources: OpenAI, Language Models are Unsupervised Mult... (2019-02-14; see Appendix A2); AP, Reddit strikes \$60M deal allowing Goo... (2024-02-22; see Appendix A3); TwinChat, AI Copyright Weights: A New Frontier... (n.d.; see Appendix A11)

3. 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 A3. Corpus companion entry: Source S-024. Audit ID: **ap-reddit-google-ai-training-deal-2024**. Relevance: That by February 2024 a major social platform was explicitly licensing accumulated user conversation data for frontier-AI training.
4. GovInfo. Copyright and Artificial Intelligence; Notice of Inquiry. 2023-08-30. T1 federal_register_notice. Paper appendix: Appendix A4. Corpus companion entry: Source S-125. Audit ID: **usco-ai-noi-2023**. Relevance: That by August 2023 the U.S. Copyright Office had already elevated AI training inputs and disclosure questions into a formal federal inquiry.
5. arXiv. Extracting Training Data from Large Language Models. 2020-12-14. T1 paper. Paper appendix: Appendix A5. Corpus companion entry: Source S-038. Audit ID: **carlini-extracting-training-data-2020**. Relevance: That large language models can sometimes reproduce identifiable verbatim material from their training data rather than only abstract statistical patterns.
6. arXiv. Generative AI Training and Copyright Law. 2025-02-21. T1 paper. Paper appendix: Appendix A6. Corpus companion entry: Source S-051. Audit ID: **dornis-generative-ai-training-copyright**. Relevance: That there is serious academic legal scholarship directly challenging the idea that generative AI training automatically falls under familiar exceptions.
7. 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 A7. Corpus companion entry: Source S-035. Audit ID: **bartz-anthropic-fair-use-order-2025**. Relevance: That the strongest primary legal record already separates the question of model training from the question of how the training corpus was assembled.
8. U.S. District Court for the Southern District of New York via Justia. The New York Times Company v. Microsoft Corporation et al, Opinion on Motions to Dismiss. 2025-04-04. T1 court_order. Paper appendix: Appendix A8. Corpus companion entry: Source S-085. Audit ID: **nyt-openai-motion-dismiss-order-2025**. Relevance: That another federal court outside the Anthropic case treated training-stage copying and dataset assembly as central issues serious enough to survive dismissal.
9. Big Data & Society. The problem with annotation. Human labour and outsourcing between France and Madagascar. 2023-07-25. T1 paper. Paper appendix: Appendix A9. Corpus companion entry: Source S-077. Audit ID: **le-ludec-cornet-casilli-annotation-madagascar-2023**. Relevance: That AI systems depend on significant human labor for data generation, annotation, verification, and operational correction.
10. U.S. House of Representatives. Bipartisan House Task Force Report on Artificial Intelligence. 2024-12-17. T1 report. Paper appendix: Appendix A10. Corpus companion entry: Source S-070. Audit ID: **house-ai-task-force-report-2024**. Relevance: That the dossier's copyright / weights / compensation framing entered a bipartisan U.S. congressional policy record by December 17, 2024.
11. TwinChat. AI Copyright Weights: A New Frontier in Intellectual Property Law. n.d.. T3 essay. Paper appendix: Appendix A11. Corpus companion entry: Source S-122. Audit ID: **twinchat-ai-copyright-weights**. Relevance: That the dossier has a coherent thesis about how weights condense previously captured human labor.

6.17. Appendix A. Source Register for This Paper

6.17.1. Appendix A1. 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 the U.S. Copyright Office treats AI training on copyrighted works as a major live legal and policy problem rather than a settled matter. **Corpus companion entry:** Source S-126 **Audit ID:** usco-ai-training-report-2025

6.17.2. Appendix A2. 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:** That a major OpenAI training corpus was explicitly built from socially filtered outbound links shared on Reddit. **Corpus companion entry:** Source S-108 **Audit ID:** radford-gpt2-webtext-2019

6.17.3. Appendix A3. 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:** That by February 2024 a major social platform was explicitly licensing accumulated user conversation data for frontier-AI training. **Corpus companion entry:** Source S-024 **Audit ID:** ap-reddit-google-ai-training-deal-2024

6.17.4. Appendix A4. 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:** That by August 2023 the U.S. Copyright Office had already elevated AI training inputs and disclosure questions into a formal federal inquiry. **Corpus companion entry:** Source S-125 **Audit ID:** usco-ai-noi-2023

6.17.5. Appendix A5. Extracting Training Data from Large Language Models

Source citation: Extracting Training Data from Large Language Models **Institution / author:** arXiv **Date:** 2020-12-14 **Tier / type:** T1 / paper **Why it matters in this paper:** That large language models can sometimes reproduce identifiable verbatim material from their training data rather than only abstract statistical patterns. **Corpus companion entry:** Source S-038 **Audit ID:** carlini-extracting-training-data-2020

6.17.6. Appendix A6. Generative AI Training and Copyright Law

Source citation: Generative AI Training and Copyright Law **Institution / author:** arXiv **Date:** 2025-02-21 **Tier / type:** T1 / paper **Why it matters in this paper:** That there is

serious academic legal scholarship directly challenging the idea that generative AI training automatically falls under familiar exceptions. **Corpus companion entry:** Source S-051 **Audit ID:** dornis-generative-ai-training-copyright-2025

6.17.7. Appendix A7. 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 the strongest primary legal record already separates the question of model training from the question of how the training corpus was assembled. **Corpus companion entry:** Source S-035 **Audit ID:** bartz-anthropic-fair-use-order-2025

6.17.8. Appendix A8. The New York Times Company v. Microsoft Corporation et al, Opinion on Motions to Dismiss

Source citation: The New York Times Company v. Microsoft Corporation et al, Opinion on Motions to Dismiss **Institution / author:** U.S. District Court for the Southern District of New York via Justia **Date:** 2025-04-04 **Tier / type:** T1 / court_order **Why it matters in this paper:** That another federal court outside the Anthropic case treated training-stage copying and dataset assembly as central issues serious enough to survive dismissal. **Corpus companion entry:** Source S-085 **Audit ID:** nyt-openai-motion-dismiss-order-2025

6.17.9. Appendix A9. The problem with annotation. Human labour and outsourcing between France and Madagascar

Source citation: The problem with annotation. Human labour and outsourcing between France and Madagascar **Institution / author:** Big Data & Society **Date:** 2023-07-25 **Tier / type:** T1 / paper **Why it matters in this paper:** That AI systems depend on significant human labor for data generation, annotation, verification, and operational correction. **Corpus companion entry:** Source S-077 **Audit ID:** le-ludec-cornet-casilli-annotation-madagascar-2023

6.17.10. Appendix A10. 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 the dossier's copyright / weights / compensation framing entered a bipartisan U.S. congressional policy record by December 17, 2024. **Corpus companion entry:** Source S-070 **Audit ID:** house-ai-task-force-report-2024

6.17.11. Appendix A11. 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 the dossier has a coherent thesis about how weights condense previously captured human labor. **Corpus companion entry:** Source S-122 **Audit ID:** twinchat-ai-copyright-weights

6.18. Appendix B. Claims Used in This Paper

6.18.1. Appendix B1. By February 2024 a major social platform was explicitly licensing accumulated use...

Claim statement: That by February 2024 a major social platform was explicitly licensing accumulated user conversation data for frontier-AI training. **Evidence status:** Documented Fact **Source support:** Appendix A3 **Corpus companion entry:** Claim C-073 **Audit Claim ID:** ap-reddit-google-ai-training-deal-2024-04-data-to-models-copyright-weights-claim-01

6.18.2. Appendix B2. Chapter 4 can extend the social-data-to-models bridge beyond GPT-2-era WebText in...

Claim statement: That Chapter 4 can extend the social-data-to-models bridge beyond GPT-2-era WebText into the later commercial frontier-model period. **Evidence status:** Documented Fact **Source support:** Appendix A3 **Corpus companion entry:** Claim C-074 **Audit Claim ID:** ap-reddit-google-ai-training-deal-2024-04-data-to-models-copyright-weights-claim-02

6.18.3. Appendix B3. The extraction economy did not stop at unpaid capture; it evolved into direct lic...

Claim statement: That the extraction economy did not stop at unpaid capture; it evolved into direct licensing arrangements over archived human expression. **Evidence status:** Documented Fact **Source support:** Appendix A3 **Corpus companion entry:** Claim C-075 **Audit Claim ID:** ap-reddit-google-ai-training-deal-2024-04-data-to-models-copyright-weights-claim-03

6.18.4. Appendix B4. The strongest primary legal record already separates the question of model traini...

Claim statement: That the strongest primary legal record already separates the question of model training from the question of how the training corpus was assembled. **Evidence status:** Documented Fact **Source support:** Appendix A7 **Corpus companion entry:** Claim C-104 **Audit Claim ID:** bartz-anthropic-fair-use-order-2025-04-data-to-models-copyright-weights-claim-01

6.18.5. Appendix B5. By June 23, 2025, a federal court had accepted a fair-use defense for Anthropic's...

Claim statement: That by June 23, 2025, a federal court had accepted a fair-use defense for Anthropic's training on lawfully acquired books while preserving exposure for piracy-related conduct. **Evidence status:** Documented Fact **Source support:** Appendix A7 **Corpus companion entry:** Claim C-105 **Audit Claim ID:** bartz-anthropic-fair-use-order-2025-04-data-to-models-copyright-

6.18.6. Appendix B6. Chapter 4 can rely on a primary court order, not only press summaries, when argui...

Claim statement: That Chapter 4 can rely on a primary court order, not only press summaries, when arguing that AI copyright conflict has split into doctrinal and supply-chain layers. **Evidence**

status: Documented Fact **Source support:** Appendix A7 **Corpus companion entry:** Claim C-106 **Audit Claim ID:** bartz-anthropic-fair-use-order-2025-04-data-to-models-copyright-weights-cla

6.18.7. Appendix B7. Large language models can sometimes reproduce identifiable verbatim material from...

Claim statement: That large language models can sometimes reproduce identifiable verbatim material from their training data rather than only abstract statistical patterns. **Evidence status:** Documented Fact **Source support:** Appendix A5 **Corpus companion entry:** Claim C-114 **Audit Claim ID:** carlini-extracting-training-data-2020-04-data-to-models-copyright-weights-claim-0

6.18.8. Appendix B8. Chapter 4 can connect training-data appropriation to memorization and leakage wit...

Claim statement: That Chapter 4 can connect training-data appropriation to memorization and leakage with direct technical evidence. **Evidence status:** Documented Fact **Source support:** Appendix A5 **Corpus companion entry:** Claim C-115 **Audit Claim ID:** carlini-extracting-training-data-2020-04-data-to-models-copyright-weights-claim-02

6.18.9. Appendix B9. The notion of model weights as a harmlessly abstract transformation of source mat...

Claim statement: That the notion of model weights as a harmlessly abstract transformation of source material is too simple for the full technical record. **Evidence status:** Documented Fact **Source support:** Appendix A5 **Corpus companion entry:** Claim C-116 **Audit Claim ID:** carlini-extracting-training-data-2020-04-data-to-models-copyright-weights-claim-03

6.18.10. Appendix B10. There is serious academic legal scholarship directly challenging the idea that ge...

Claim statement: That there is serious academic legal scholarship directly challenging the idea that generative AI training automatically falls under familiar exceptions. **Evidence status:** Documented Fact **Source support:** Appendix A6 **Corpus companion entry:** Claim C-153 **Audit Claim ID:** dornis-generative-ai-training-copyright-2025-04-data-to-models-copyright-weights-clai

6.18.11. Appendix B11. Memorization is not only a technical curiosity but also part of the copyright debate

Claim statement: That memorization is not only a technical curiosity but also part of the copyright debate. **Evidence status:** Documented Fact **Source support:** Appendix A6 **Corpus companion entry:** Claim C-154 **Audit Claim ID:** dornis-generative-ai-training-copyright-2025-04-data-to-models-copyright-weights-clai

6.18.12. Appendix B12. Chapter 4 can present the legal status of training as contested and actively theo...

Claim statement: That Chapter 4 can present the legal status of training as contested and actively theorized rather than settled by industry analogy. **Evidence status:** Documented Fact

Source support: Appendix A6 **Corpus companion entry:** Claim C-155 **Audit Claim ID:** dornis-generative-ai-training-copyright-2025-04-data-to-models-copyright-weights-claim-03

6.18.13. Appendix B13. The dossier’s copyright / weights / compensation framing entered a bipartisan U.S...

Claim statement: That the dossier’s copyright / weights / compensation framing entered a bipartisan U.S. congressional policy record by December 17, 2024. **Evidence status:** Documented Fact **Source support:** Appendix A10 **Corpus companion entry:** Claim C-219 **Audit Claim ID:** house-ai-task-force-report-2024-04-data-to-models-copyright-weights-claim-01

6.18.14. Appendix B14. A formal House report recorded the AI-training dispute in terms that directly ove...

Claim statement: That a formal House report recorded the AI-training dispute in terms that directly overlap with the dossier’s concern about reproduction, derivative-work theory, commercialization, consent, credit, and compensation. **Evidence status:** Documented Fact **Source support:** Appendix A10 **Corpus companion entry:** Claim C-220 **Audit Claim ID:** house-ai-task-force-report-2024-04-data-to-models-copyright-weights-claim-02

6.18.15. Appendix B15. Chapter 00 and Chapter 04 can cite institutional uptake of the argument without p...

Claim statement: That Chapter 00 and Chapter 04 can cite institutional uptake of the argument without pretending that congressional citation settles the underlying legal merits. **Evidence status:** Documented Fact **Source support:** Appendix A10 **Corpus companion entry:** Claim C-221 **Audit Claim ID:** house-ai-task-force-report-2024-04-data-to-models-copyright-weights-claim-03

6.18.16. Appendix B16. AI systems depend on significant human labor for data generation, annotation, ver...

Claim statement: That AI systems depend on significant human labor for data generation, annotation, verification, and operational correction. **Evidence status:** Documented Fact **Source support:** Appendix A9 **Corpus companion entry:** Claim C-246 **Audit Claim ID:** le-ludec-cornet-casilli-annotation-madagascar-2023-04-data-to-models-copyright-weights-claim-04

6.18.17. Appendix B17. Chapter 4 can connect model building to undervalued and geographically externaliz...

Claim statement: That Chapter 4 can connect model building to undervalued and geographically externalized human work, not only to scraped copyrighted material. **Evidence status:** Documented Fact **Source support:** Appendix A9 **Corpus companion entry:** Claim C-247 **Audit Claim ID:** le-ludec-cornet-casilli-annotation-madagascar-2023-04-data-to-models-copyright-weights-claim-05

6.18.18. Appendix B18. A serious research literature already frames AI production as a labor-reorganization...

Claim statement: That a serious research literature already frames AI production as a labor-reorganization process rather than a purely autonomous technical breakthrough. **Evidence status:** Documented Fact **Source support:** Appendix A9 **Corpus companion entry:** Claim C-248 **Audit Claim ID:** le-ludec-cornet-casilli-annotation-madagascar-2023-04-data-to-models-copyright-weights-claim-01

6.18.19. Appendix B19. Another federal court outside the Anthropic case treated training-stage copying as...

Claim statement: That another federal court outside the Anthropic case treated training-stage copying and dataset assembly as central issues serious enough to survive dismissal. **Evidence status:** Documented Fact **Source support:** Appendix A8 **Corpus companion entry:** Claim C-280 **Audit Claim ID:** nyt-openai-motion-dismiss-order-2025-04-data-to-models-copyright-weights-claim-02

6.18.20. Appendix B20. Chapter 4 can rely on more than one primary court record when arguing that the AI...

Claim statement: That Chapter 4 can rely on more than one primary court record when arguing that the AI copyright conflict has split into doctrinal and supply-chain layers. **Evidence status:** Documented Fact **Source support:** Appendix A8 **Corpus companion entry:** Claim C-281 **Audit Claim ID:** nyt-openai-motion-dismiss-order-2025-04-data-to-models-copyright-weights-claim-02

6.18.21. Appendix B21. The legal record does not treat training as a purely abstract statistical process...

Claim statement: That the legal record does not treat training as a purely abstract statistical process detached from source acquisition and storage practices. **Evidence status:** Documented Fact **Source support:** Appendix A8 **Corpus companion entry:** Claim C-282 **Audit Claim ID:** nyt-openai-motion-dismiss-order-2025-04-data-to-models-copyright-weights-claim-03

6.18.22. Appendix B22. A major OpenAI training corpus was explicitly built from socially filtered outbound...

Claim statement: That a major OpenAI training corpus was explicitly built from socially filtered outbound links shared on Reddit. **Evidence status:** Documented Fact **Source support:** Appendix A2 **Corpus companion entry:** Claim C-356 **Audit Claim ID:** radford-gpt2-webtext-2019-04-data-to-models-copyright-weights-claim-01

6.18.23. Appendix B23. Chapter 4 can make a direct bridge from social-platform activity to model-training...

Claim statement: That Chapter 4 can make a direct bridge from social-platform activity to model-training inputs instead of speaking only in broad abstractions about “the internet.” **Evidence status:** Documented Fact **Source support:** Appendix A2 **Corpus companion entry:** Claim C-357 **Audit Claim ID:** radford-gpt2-webtext-2019-04-data-to-models-copyright-weights-claim-02

6.18.24. Appendix B24. Training-data pipelines can incorporate unpaid human curation signals even when t...

Claim statement: That training-data pipelines can incorporate unpaid human curation signals even when the final corpus is presented as a neutral technical dataset. **Evidence status:** Documented Fact **Source support:** Appendix A2 **Corpus companion entry:** Claim C-358 **Audit Claim ID:** radford-gpt2-webtext-2019-04-data-to-models-copyright-weights-claim-03

6.18.25. Appendix B25. The dossier has a coherent thesis about how weights condense previously captured...

Claim statement: That the dossier has a coherent thesis about how weights condense previously captured human labor. **Evidence status:** Documented Fact **Source support:** Appendix A11 **Corpus companion entry:** Claim C-397 **Audit Claim ID:** twinchat-ai-copyright-weights-claim-01

6.18.26. Appendix B26. There is a clear argumentative line linking AI, property, and enclosure

Claim statement: That there is a clear argumentative line linking AI, property, and enclosure. **Evidence status:** Documented Fact **Source support:** Appendix A11 **Corpus companion entry:** Claim C-398 **Audit Claim ID:** twinchat-ai-copyright-weights-claim-02

6.18.27. Appendix B27. By August 2023 the U.S. Copyright Office had already elevated AI training inputs...

Claim statement: That by August 2023 the U.S. Copyright Office had already elevated AI training inputs and disclosure questions into a formal federal inquiry. **Evidence status:** Documented Fact **Source support:** Appendix A4 **Corpus companion entry:** Claim C-407 **Audit Claim ID:** usco-ai-noi-2023-04-data-to-models-copyright-weights-claim-01

6.18.28. Appendix B28. The legal and policy debate around training data was not retrofitted later; it wa...

Claim statement: That the legal and policy debate around training data was not retrofitted later; it was a live official issue near the start of the generative AI boom. **Evidence status:** Documented Fact **Source support:** Appendix A4 **Corpus companion entry:** Claim C-408 **Audit Claim ID:** usco-ai-noi-2023-04-data-to-models-copyright-weights-claim-02

6.18.29. Appendix B29. Chapter 4 can document a formal institutional path from open questions about AI t...

Claim statement: That Chapter 4 can document a formal institutional path from open questions about AI training to later official analysis and policy conflict. **Evidence status:** Documented Fact **Source support:** Appendix A4 **Corpus companion entry:** Claim C-409 **Audit Claim ID:** usco-ai-noi-2023-04-data-to-models-copyright-weights-claim-03

6.18.30. Appendix B30. The U.S. Copyright Office treats AI training on copyrighted works as a major live...

Claim statement: That the U.S. Copyright Office treats AI training on copyrighted works as a major live legal and policy problem rather than a settled matter. **Evidence status:** Documented Fact **Source support:** Appendix A1 **Corpus companion entry:** Claim C-410 **Audit Claim ID:** usco-ai-training-report-2025-04-data-to-models-copyright-weights-claim-01

6.18.31. Appendix B31. Official analysis recognizes copyrighted creative works as central inputs to gene...

Claim statement: That official analysis recognizes copyrighted creative works as central inputs to generative AI development. **Evidence status:** Documented Fact **Source support:** Appendix A1 **Corpus companion entry:** Claim C-411 **Audit Claim ID:** usco-ai-training-report-2025-04-data-to-models-copyright-weights-claim-02

6.18.32. Appendix B32. Chapter 4 can ground its argument about training-data appropriation and licensing...

Claim statement: That Chapter 4 can ground its argument about training-data appropriation and licensing conflict in a primary U.S. government source. **Evidence status:** Documented Fact **Source support:** Appendix A1 **Corpus companion entry:** Claim C-412 **Audit Claim ID:** usco-ai-training-report-2025-04-data-to-models-copyright-weights-claim-03

6.18.33. Appendix B33. 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 A1, Appendix A7 **Corpus companion entry:** Claim C-428 **Audit Claim ID:** vector04-training-legality-contested-claim-04

6.18.34. Appendix B34. Dominant intermediaries can consolidate value and pricing power by controlling th...

Claim statement: That dominant intermediaries can consolidate value and pricing power by controlling the access layer through which users reach collectively produced goods, even when they did not produce the underlying material. **Evidence status:** Hypothesis / Interpretation **Source support:** Appendix A2, Appendix A3, Appendix A9 **Corpus companion entry:** Claim C-424 **Audit Claim ID:** vector04-access-layer-power-claim-01

6.18.35. Appendix B35. The AI prompt box, chatbot, or subscription layer can be read as the new shelf or...

Claim statement: That the AI prompt box, chatbot, or subscription layer can be read as the new shelf or point of sale of intelligence because it mediates access to capabilities condensed from collective human production. **Evidence status:** Hypothesis / Interpretation **Source support:**

Appendix A2, Appendix A3, Appendix A9 **Corpus companion entry:** Claim C-425 **Audit Claim ID:** vector04-interface-as-shelf-claim-01

6.18.36. Appendix B36. 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 A11, Appendix A1, Appendix A10 **Corpus companion entry:** Claim C-429 **Audit Claim ID:** vector04-weights-patrimony-governance-claim-01

6.18.37. Appendix B37. 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 A10, Appendix A1 **Corpus companion entry:** Claim C-426 **Audit Claim ID:** vector04-patrimony-already-settled-overclaim-claim-01

6.18.38. Appendix B38. The supermarket / shelf analogy is itself established as a documented fact by the...

Claim statement: That the supermarket / shelf analogy is itself established as a documented fact by the legal or regulatory record rather than as a bounded interpretation built on that record. **Evidence status:** Speculative Narrative Risk **Source support:** Appendix A1, Appendix A3, Appendix A9 **Corpus companion entry:** Claim C-427 **Audit Claim ID:** vector04-shelf-analogy-overproof-claim-01

6.19. Appendix C. Timeline Slice

6.19.1. Appendix C1. 2019-02-14: OpenAI documents a Reddit-filtered corpus for GPT-2 training

Event summary: OpenAI documents a Reddit-filtered corpus for GPT-2 training **Event date:** 2019-02-14 **Source support:** Appendix A2 **Corpus companion entry:** Event T-005 **Audit Event ID:** event-gpt2-webtext-reddit-2019-02-14

6.19.2. Appendix C2. 2023-07-25: A peer-reviewed study maps outsourced annotation labor in AI production chains

Event summary: A peer-reviewed study maps outsourced annotation labor in AI production chains **Event date:** 2023-07-25 **Source support:** Appendix A9 **Corpus companion entry:** Event T-027 **Audit Event ID:** event-annotation-madagascar-paper-2023-07-25

6.19.3. Appendix C3. 2023-08-30: The U.S. Copyright Office opens a formal inquiry into AI training and copyright

Event summary: The U.S. Copyright Office opens a formal inquiry into AI training and copyright

Event date: 2023-08-30 **Source support:** Appendix A4 **Corpus companion entry:** Event T-028 **Audit Event ID:** event-usco-ai-noi-2023-08-30

6.19.4. Appendix C4. 2024-02-22: Reddit announces a \$60 million agreement allowing Google to train AI models on Reddit...

Event summary: Reddit announces a \$60 million agreement allowing Google to train AI models on Reddit posts

Event date: 2024-02-22 **Source support:** Appendix A3 **Corpus companion entry:** Event T-043 **Audit Event ID:** event-reddit-google-training-deal-2024-02-22

6.19.5. Appendix C5. 2024-12-17: The Bipartisan House AI Task Force report cites Hassan Uriostegui's weights-and-copyr...

Event summary: The Bipartisan House AI Task Force report cites Hassan Uriostegui's weights-and-copyright argument

Event date: 2024-12-17 **Source support:** Appendix A10 **Corpus companion entry:** Event T-062 **Audit Event ID:** event-house-ai-task-force-report-2024-12-17

6.19.6. Appendix C6. 2025-04-04: A federal court lets core New York Times copyright claims against OpenAI and Microsof...

Event summary: A federal court lets core New York Times copyright claims against OpenAI and Microsoft proceed

Event date: 2025-04-04 **Source support:** Appendix A8 **Corpus companion entry:** Event T-073 **Audit Event ID:** event-nyt-openai-dismissal-order-2025-04-04

6.19.7. Appendix C7. 2025-05-09: The U.S. Copyright Office releases its pre-publication report on generative AI training

Event summary: The U.S. Copyright Office releases its pre-publication report on generative AI training

Event date: 2025-05-09 **Source support:** Appendix A1 **Corpus companion entry:** Event T-075 **Audit Event ID:** event-usco-ai-training-report-2025-05-09

6.19.8. Appendix C8. 2025-06-23: Judge Alsup's order separates AI training fair use from pirate-library exposure

Event summary: Judge Alsup's order separates AI training fair use from pirate-library exposure

Event date: 2025-06-23 **Source support:** Appendix A7 **Corpus companion entry:** Event T-080 **Audit Event ID:** event-bartz-anthropic-order-2025-06-23

6.20. Appendix D. Relevant Entities

6.20.1. 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 A7 **Corpus companion entry:** Entity E-008 **Audit Entity ID:** org-anthropic

6.20.2. 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 A3 **Corpus companion entry:** Entity E-018 **Audit Entity ID:** org-google

6.20.3. 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 A8, Appendix A2 **Corpus companion entry:** Entity E-027 **Audit Entity ID:** org-openai

6.20.4. Appendix D4. 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 A4, Appendix A1 **Corpus companion entry:** Entity E-037 **Audit Entity ID:** org-us-copyright-office

6.20.5. Appendix D5. 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 A11, Appendix A10 **Corpus companion entry:** Entity E-048 **Audit Entity ID:** person-hassan-uriostegui

6.20.6. Appendix D6. 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 A3, Appendix A2 **Corpus companion entry:** Entity E-032 **Audit Entity ID:** org-reddit

6.21. Appendix E. Evidence Boundaries

6.21.1. Appendix E1. Disputed Matters

- Appendix B33: 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.

6.21.2. Appendix E2. Interpretive Boundaries

- Appendix B34: That dominant intermediaries can consolidate value and pricing power by controlling the access layer through which users reach collectively produced goods, even when they did not produce the underlying material.

- Appendix B35: That the AI prompt box, chatbot, or subscription layer can be read as the new shelf or point of sale of intelligence because it mediates access to capabilities condensed from collective human production.
- Appendix B36: That frontier AI weights can be read as condensed claims over collectively produced human expression and therefore justify patrimony-oriented governance.

6.21.3. Appendix E3. Speculative Narrative Risks

- Appendix B37: That current legal uncertainty or congressional uptake already proves frontier AI weights are legally settled as patrimony, public trust property, or common heritage.
- Appendix B38: That the supermarket / shelf analogy is itself established as a documented fact by the legal or regulatory record rather than as a bounded interpretation built on that record.

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

- Appendix A3. Companion note: Note N-027. Risk flag: review. Note focus: That by February 2024 a major social platform was explicitly licensing accumulated user conversation data for frontier-AI training.
- Appendix A7. Companion note: Note N-028. Risk flag: review. Note focus: That the strongest primary legal record already separates the question of model training from the question of how the training corpus was assembled.
- Appendix A5. Companion note: Note N-029. Risk flag: review. Note focus: That large language models can sometimes reproduce identifiable verbatim material from their training data rather than only abstract statistical patterns.
- Appendix A6. Companion note: Note N-030. Risk flag: review. Note focus: That there is serious academic legal scholarship directly challenging the idea that generative AI training automatically falls under familiar exceptions.
- Appendix A10. Companion note: Note N-031. Risk flag: review. Note focus: That the dossier's copyright / weights / compensation framing entered a bipartisan U.S. congressional policy record by December 17, 2024.
- Appendix A9. Companion note: Note N-032. Risk flag: low. Note focus: That AI systems depend on significant human labor for data generation, annotation, verification, and operational correction.
- Appendix A8. Companion note: Note N-033. Risk flag: medium. Note focus: That another federal court outside the Anthropic case treated training-stage copying and dataset assembly as central issues serious enough to survive dismissal.
- Appendix A2. Companion note: Note N-034. Risk flag: low. Note focus: That a major OpenAI training corpus was explicitly built from socially filtered outbound links shared on Reddit.
- Appendix A11. Companion note: Note N-035. Risk flag: medium. Note focus: That the dossier has a coherent thesis about how weights condense previously captured human labor.
- Appendix A4. Companion note: Note N-037. Risk flag: review. Note focus: That by August 2023 the U.S. Copyright Office had already elevated AI training inputs and disclosure questions into a formal federal inquiry.
- Appendix A1. Companion note: Note N-038. Risk flag: review. Note focus: That the U.S. Copyright Office treats AI training on copyrighted works as a major live legal and policy problem

rather than a settled matter.

7. Chapter 05. Talent, overhiring, valuation, and layoffs

7.1. Abstract

The cleanest version of the layoffs story is not that AI suddenly made technical labor obsolete. Section 174 is a real structural factor. If this reading holds, then the AI boom is also a class story.

7.2. Keywords

Talent, overhiring, valuation, and layoffs; surveillance capitalism; AI imperialism; evidence-based dossier

7.3. Research Question

How does talent, overhiring, valuation, and layoffs function within the dossier's larger argument, and what does the documentary record allow this chapter to establish?

7.4. Scope and Framing Note

This paper examines talent, overhiring, valuation, and layoffs 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 connects the enclosure story to labor correction, balance-sheet pressure, and the financial narratives that stabilized layoffs. Adjacent reinforcement appears in Chapter 04, Chapter 06, Chapter 09, but those cross-references are supportive rather than required for basic comprehension.

7.5. Central Thesis

The cleanest version of the layoffs story is not that AI suddenly made technical labor obsolete. It is that AI arrived as the legitimizing language for a correction that already had deeper causes: pandemic-era overhiring, higher rates, margin pressure, weaker demand, and the tax treatment of research and software spending. The cuts were not born from one machine. They were narrated through one. ⁵⁵

7.6. Evidence and Method Note

This paper is derived from the chapter corpus for Chapter 05 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.

⁵⁵Sources: Internal Revenue Service, Notice 2023-63: Guidance on Section 1... (2023-09-08; see Appendix A1); Associated Press, Amazon, Salesforce jettison jobs in 1... (2023-01-04; see Appendix A2); Meta, Mark Zuckerberg's Message to Meta Emp... (2023-03-14; see Appendix A3); Financial Times, AI exuberance masks broad weakness in... (2024-09-09; see Appendix A4)

7.7. Introduction

This paper asks: How does talent, overhiring, valuation, and layoffs function within the dossier's larger argument, and what does the documentary record allow this chapter to establish? Its working answer is clear from the start: The cleanest version of the layoffs story is not that AI suddenly made technical labor obsolete. It is that AI arrived as the legitimizing language for a correction that already had deeper causes: pandemic-era overhiring, higher rates, margin pressure, weaker demand, and the tax treatment of research and software spending. The cuts were not born from one machine. They were narrated through one.⁵⁶ 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.

7.8. Historical or Institutional Context

Section 174 is a real structural factor. IRS guidance and congressional repair efforts show that domestic research and software expenses were caught in a more punitive amortization regime after 2021, and Congress itself treated the burden as significant enough to attempt reversal. This is not the whole story, but it is not rhetorical decoration either.⁵⁷

7.9. Main Analysis

The post-pandemic overhiring record is also strong. Executive statements and major reporting across Meta, Amazon, and Salesforce show a later turn toward efficiency, focus, and cost discipline after an earlier period of aggressive expansion.⁵⁸

Top-tier reporting ties the broader correction to weaker demand, margin pressure, and investor impatience, not only to AI. That matters because it prevents a monocausal story. The financial correction was already underway before AI became the universal explanation for it.⁵⁹

The BIS annual report adds the macro-financial layer. By June 2026, a major multilateral institution was openly warning that the AI capex race was outrunning earnings and free cash flow, with leverage and financing structures becoming part of the risk profile. The system was promising abundance while leaning on increasingly fragile financial scaffolding.⁶⁰

Interpretation: this is where the dossier's economic thesis sharpens. AI did not invent the correction,

⁵⁶Sources: Internal Revenue Service, Notice 2023-63: Guidance on Section 1... (2023-09-08; see Appendix A1); Associated Press, Amazon, Salesforce jettison jobs in l... (2023-01-04; see Appendix A2); Meta, Mark Zuckerberg's Message to Meta Emp... (2023-03-14; see Appendix A3); Financial Times, AI exuberance masks broad weakness in... (2024-09-09; see Appendix A4)

⁵⁷Sources: Internal Revenue Service, Notice 2023-63: Guidance on Section 1... (2023-09-08; see Appendix A1); Congress.gov, H.R. 7024: Tax Relief for American Fa... (2024-01-16; see Appendix A5); Congress.gov / CRS, H.R. 7024: Tax Relief for American Fa... (2024-01-31; see Appendix A6); Congress.gov / Library of Congress, H.R. 7024 text: Deduction for domesti... (2024-03-21; see Appendix A7)

⁵⁸Sources: Meta, Mark Zuckerberg's Message to Meta Emp... (2023-03-14; see Appendix A3); Associated Press, Meta posts lower Q4 profit, announces... (2023-02-01; see Appendix A8); Associated Press, Meta posts sharp profit, revenue incr... (2024-02-01; see Appendix A9); Associated Press, Amazon, Salesforce jettison jobs in l... (2023-01-04; see Appendix A2)

⁵⁹Sources: Financial Times, AI exuberance masks broad weakness in... (2024-09-09; see Appendix A4); Axios, Big Tech layoffs shattered industry,... (2024-01-03; see Appendix A10)

⁶⁰Sources: Bank for International Settlements, Annual Economic Report 2026 (2026-06-28; see Appendix A11)

but it became the language through which it was justified. The same story that encloses captured human production also helps reclassify labor cuts as technological destiny instead of the outcome of financial structure, policy, and managerial excess. That interpretation supports Chapter 04 by showing how enclosure and labor subordination move together.⁶¹

7.10. Position Within the Series

This paper sits inside a coordinated 11-paper architecture. Its nearest companions are Chapter 04 (AI arrives when the human raw material already exists), Chapter 06 (AI as a dual-use rhetorical weapon), Chapter 09 (Future if no action is taken). In that series logic, Chapter 05 connects the enclosure story to labor correction, balance-sheet pressure, and the financial narratives that stabilized layoffs.

7.11. Counterarguments

This chapter should not imply that executives always lied when they invoked AI. Some reallocation toward AI was clearly real. The stronger, more defensible claim is narrower: AI provided a legitimizing narrative layered on top of a structural correction with multiple preexisting causes.⁶²

7.12. Limits of the Evidence

This chapter should not imply that executives always lied when they invoked AI. Some reallocation toward AI was clearly real. The stronger, more defensible claim is narrower: AI provided a legitimizing narrative layered on top of a structural correction with multiple preexisting causes.⁶³

7.13. Reform Relevance

If this reading holds, then the AI boom is also a class story. It narrows the protected tier, weakens the technical middle class, and presents downward pressure on labor as the rational price of the future. That strengthens the patrimony argument: a system built from collective human production is being enclosed while many of the humans who made it possible are told the future no longer needs them on equal terms.⁶⁴

7.14. Conclusion

If this reading holds, then the AI boom is also a class story. That contradiction sharpens in the next chapter. In the architecture of this series, Chapter 05 therefore functions as a self-contained argument while also advancing the cumulative path toward the dossier's three reforms.

⁶¹Sources: Meta, Mark Zuckerberg's Message to Meta Emp... (2023-03-14; see Appendix A3); Financial Times, AI exuberance masks broad weakness in... (2024-09-09; see Appendix A4); Bank for International Settlements, Annual Economic Report 2026 (2026-06-28; see Appendix A11)

⁶²Sources: Meta, Mark Zuckerberg's Message to Meta Emp... (2023-03-14; see Appendix A3); Associated Press, Amazon, Salesforce jettison jobs in l... (2023-01-04; see Appendix A2)

⁶³Sources: Meta, Mark Zuckerberg's Message to Meta Emp... (2023-03-14; see Appendix A3); Associated Press, Amazon, Salesforce jettison jobs in l... (2023-01-04; see Appendix A2)

⁶⁴Sources: Internal Revenue Service, Notice 2023-63: Guidance on Section 1... (2023-09-08; see Appendix A1); Bank for International Settlements, Annual Economic Report 2026 (2026-06-28; see Appendix A11); TwinChat, AI Copyright Weights: A New Frontier... (n.d.; see Appendix A12)

7.15. 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 05, those boundaries remain visible in the prose, in the bibliography, and in the appendix sections that summarize claims and caution points.

7.16. References / Bibliography

1. Internal Revenue Service. Notice 2023-63: Guidance on Section 174 and specified research or experimental expenditures. 2023-09-08. T1 official_notice. Paper appendix: Appendix A1. Corpus companion entry: Source S-075. Audit ID: `irs-notice-2023-63`. Relevance: That Section 174 amortization was not a rumor or niche accounting theory; it was the subject of active IRS guidance.
2. Associated Press. Amazon, Salesforce jettison jobs in latest tech worker purge. 2023-01-04. T2 news_article. Paper appendix: Appendix A2. Corpus companion entry: Source S-008. Audit ID: `ap-amazon-salesforce-tech-purge-2023`. Relevance: That major tech layoffs were publicly tied to pandemic overexpansion and deteriorating macro conditions.
3. Meta. Mark Zuckerberg's Message to Meta Employees. 2023-03-14. T1 company_statement. Paper appendix: Appendix A3. Corpus companion entry: Source S-082. Audit ID: `meta-year-of-efficiency-2023`. Relevance: That Meta itself framed layoffs as part of a structural response to financial conditions and organizational inefficiency.
4. Financial Times. AI exuberance masks broad weakness in tech sector, say investors. 2024-09-09. T2 news_analysis. Paper appendix: Appendix A4. Corpus companion entry: Source S-055. Audit ID: `ft-ai-exuberance-tech-weakness-2024`. Relevance: That AI enthusiasm did not eliminate the broader structural weakness of the sector.
5. Congress.gov. H.R. 7024: Tax Relief for American Families and Workers Act of 2024. 2024-01-16. T1 legislation. Paper appendix: Appendix A5. Corpus companion entry: Source S-042. Audit ID: `congress-hr7024-tax-relief-2024`. Relevance: That Congress formally moved to soften or reverse the immediate impact of Section 174 on domestic research expensing.
6. Congress.gov / CRS. H.R. 7024: Tax Relief for American Families and Workers Act of 2024 (CRS summary). 2024-01-31. T1 legislative_summary. Paper appendix: Appendix A6. Corpus companion entry: Source S-041. Audit ID: `congress-hr7024-summary-2024`. Relevance: That Congress formally treated the post-2021 Section 174 regime as burdensome enough to warrant legislative reversal or postponement.
7. Congress.gov / Library of Congress. H.R. 7024 text: Deduction for domestic research and experimental expenditures. 2024-03-21. T1 bill_text. Paper appendix: Appendix A7. Corpus companion entry: Source S-043. Audit ID: `congress-hr7024-text-2024`. Relevance: That the practical Section 174 fight directly implicated software development, not only laboratory research.
8. Associated Press. Meta posts lower Q4 profit, announces huge stock buyback. 2023-02-01. T2 news_article. Paper appendix: Appendix A8. Corpus companion entry: Source S-013. Audit ID: `ap-meta-q4-2022-overhiring`. Relevance: That a major tech company publicly linked

layoffs to pandemic overhiring and slowing post-lockdown growth.

9. Associated Press. Meta posts sharp profit, revenue increase in Q4 thanks to cost cuts and advertising rebound. 2024-02-01. T2 news_article. Paper appendix: Appendix A9. Corpus companion entry: Source S-014. Audit ID: `ap-meta-q4-2023-year-of-efficiency`. Relevance: That Meta's improved financial results were publicly tied to both cost cutting and layoffs, not just to a magical AI turnaround.
10. Axios. Big Tech layoffs shattered industry, worker confidence. 2024-01-03. T3 news_analysis. Paper appendix: Appendix A10. Corpus companion entry: Source S-034. Audit ID: `axios-tech-layoffs-worker-confidence-2024`. Relevance: That the layoff cycle was large enough to materially change worker expectations across the sector.
11. Bank for International Settlements. Annual Economic Report 2026. 2026-06-28. T1 report. Paper appendix: Appendix A11. Corpus companion entry: Source S-036. Audit ID: `bis-annual-economic-report-2026`. Relevance: That a primary multilateral financial institution explicitly warned in June 2026 that the AI capex race was outrunning earnings and free cash flow.
12. TwinChat. AI Copyright Weights: A New Frontier in Intellectual Property Law. n.d.. T3 essay. Paper appendix: Appendix A12. Corpus companion entry: Source S-122. Audit ID: `twinchat-ai-copyright-weights`. Relevance: Supports a key factual or interpretive move inside this paper.

7.17. Appendix A. Source Register for This Paper

7.17.1. Appendix A1. Notice 2023-63: Guidance on Section 174 and specified research or experimental expenditures

Source citation: Notice 2023-63: Guidance on Section 174 and specified research or experimental expenditures **Institution / author:** Internal Revenue Service **Date:** 2023-09-08 **Tier / type:** T1 / official_notice **Why it matters in this paper:** That Section 174 amortization was not a rumor or niche accounting theory; it was the subject of active IRS guidance. **Corpus companion entry:** Source S-075 **Audit ID:** `irs-notice-2023-63`

7.17.2. Appendix A2. Amazon, Salesforce jettison jobs in latest tech worker purge

Source citation: Amazon, Salesforce jettison jobs in latest tech worker purge **Institution / author:** Associated Press **Date:** 2023-01-04 **Tier / type:** T2 / news_article **Why it matters in this paper:** That major tech layoffs were publicly tied to pandemic overexpansion and deteriorating macro conditions. **Corpus companion entry:** Source S-008 **Audit ID:** `ap-amazon-salesforce-tech-purge-2023`

7.17.3. Appendix A3. Mark Zuckerberg's Message to Meta Employees

Source citation: Mark Zuckerberg's Message to Meta Employees **Institution / author:** Meta **Date:** 2023-03-14 **Tier / type:** T1 / company_statement **Why it matters in this paper:** That Meta itself framed layoffs as part of a structural response to financial conditions and organizational inefficiency. **Corpus companion entry:** Source S-082 **Audit ID:** `meta-year-of-efficiency-2023`

7.17.4. Appendix A4. AI exuberance masks broad weakness in tech sector, say investors

Source citation: AI exuberance masks broad weakness in tech sector, say investors **Institution / author:** Financial Times **Date:** 2024-09-09 **Tier / type:** T2 / news_analysis **Why it matters in this paper:** That AI enthusiasm did not eliminate the broader structural weakness of the sector. **Corpus companion entry:** Source S-055 **Audit ID:** ft-ai-exuberance-tech-weakness-2024

7.17.5. Appendix A5. H.R. 7024: Tax Relief for American Families and Workers Act of 2024

Source citation: H.R. 7024: Tax Relief for American Families and Workers Act of 2024 **Institution / author:** Congress.gov **Date:** 2024-01-16 **Tier / type:** T1 / legislation **Why it matters in this paper:** That Congress formally moved to soften or reverse the immediate impact of Section 174 on domestic research expensing. **Corpus companion entry:** Source S-042 **Audit ID:** congress-hr7024-tax-relief-2024

7.17.6. Appendix A6. H.R. 7024: Tax Relief for American Families and Workers Act of 2024 (CRS summary)

Source citation: H.R. 7024: Tax Relief for American Families and Workers Act of 2024 (CRS summary) **Institution / author:** Congress.gov / CRS **Date:** 2024-01-31 **Tier / type:** T1 / legislative_summary **Why it matters in this paper:** That Congress formally treated the post-2021 Section 174 regime as burdensome enough to warrant legislative reversal or postponement. **Corpus companion entry:** Source S-041 **Audit ID:** congress-hr7024-summary-2024

7.17.7. Appendix A7. H.R. 7024 text: Deduction for domestic research and experimental expenditures

Source citation: H.R. 7024 text: Deduction for domestic research and experimental expenditures **Institution / author:** Congress.gov / Library of Congress **Date:** 2024-03-21 **Tier / type:** T1 / bill_text **Why it matters in this paper:** That the practical Section 174 fight directly implicated software development, not only laboratory research. **Corpus companion entry:** Source S-043 **Audit ID:** congress-hr7024-text-2024

7.17.8. Appendix A8. Meta posts lower Q4 profit, announces huge stock buyback

Source citation: Meta posts lower Q4 profit, announces huge stock buyback **Institution / author:** Associated Press **Date:** 2023-02-01 **Tier / type:** T2 / news_article **Why it matters in this paper:** That a major tech company publicly linked layoffs to pandemic overhiring and slowing post-lockdown growth. **Corpus companion entry:** Source S-013 **Audit ID:** ap-meta-q4-2022-overhiring

7.17.9. Appendix A9. Meta posts sharp profit, revenue increase in Q4 thanks to cost cuts and advertising rebound

Source citation: Meta posts sharp profit, revenue increase in Q4 thanks to cost cuts and advertising rebound **Institution / author:** Associated Press **Date:** 2024-02-01 **Tier / type:** T2 / news_article **Why it matters in this paper:** That Meta's improved financial results were publicly tied to both cost cutting and layoffs, not just to a magical AI turnaround. **Corpus companion entry:** Source S-014 **Audit ID:** ap-meta-q4-2023-year-of-efficiency

7.17.10. Appendix A10. Big Tech layoffs shattered industry, worker confidence

Source citation: Big Tech layoffs shattered industry, worker confidence **Institution / author:** Axios **Date:** 2024-01-03 **Tier / type:** T3 / news_analysis **Why it matters in this paper:** That the layoff cycle was large enough to materially change worker expectations across the sector. **Corpus companion entry:** Source S-034 **Audit ID:** axios-tech-layoffs-worker-confidence-2024

7.17.11. Appendix A11. Annual Economic Report 2026

Source citation: Annual Economic Report 2026 **Institution / author:** Bank for International Settlements **Date:** 2026-06-28 **Tier / type:** T1 / report **Why it matters in this paper:** That a primary multilateral financial institution explicitly warned in June 2026 that the AI capex race was outrunning earnings and free cash flow. **Corpus companion entry:** Source S-036 **Audit ID:** bis-annual-economic-report-2026

7.17.12. Appendix A12. 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:** Supports a key factual or interpretive move inside this paper. **Corpus companion entry:** Source S-122 **Audit ID:** twinchat-ai-copyright-weights

7.18. Appendix B. Claims Used in This Paper

7.18.1. Appendix B1. Major tech layoffs were publicly tied to pandemic overexpansion and deteriorating...

Claim statement: That major tech layoffs were publicly tied to pandemic overexpansion and deteriorating macro conditions. **Evidence status:** Documented Fact **Source support:** Appendix A2 **Corpus companion entry:** Claim C-022 **Audit Claim ID:** ap-amazon-salesforce-tech-purge-2023-05-talent-hoarding-section174-layoffs-claim-01

7.18.2. Appendix B2. Efficiency language was already being used by top executives in early 2023, befor...

Claim statement: That efficiency language was already being used by top executives in early 2023, before later AI-centered narratives became even more prominent. **Evidence status:** Documented

Fact Source support: Appendix A2 **Corpus companion entry:** Claim C-023 **Audit Claim ID:** ap-amazon-salesforce-tech-purge-2023-05-talent-hoarding-section174-layoffs-claim-02

7.18.3. Appendix B3. Chapter 5 can frame layoffs as part of a structural correction involving cost dis...

Claim statement: That Chapter 5 can frame layoffs as part of a structural correction involving cost discipline and margin protection, not only as a pure automation story. **Evidence status:** Documented **Fact Source support:** Appendix A2 **Corpus companion entry:** Claim C-024 **Audit Claim ID:** ap-amazon-salesforce-tech-purge-2023-05-talent-hoarding-section174-layoffs-claim-03

7.18.4. Appendix B4. A major tech company publicly linked layoffs to pandemic overhiring and slowing p...

Claim statement: That a major tech company publicly linked layoffs to pandemic overhiring and slowing post-lockdown growth. **Evidence status:** Documented **Fact Source support:** Appendix A8 **Corpus companion entry:** Claim C-037 **Audit Claim ID:** ap-meta-q4-2022-overhiring-claim-01

7.18.5. Appendix B5. Efficiency rhetoric was already being layered onto a revenue slowdown and investo...

Claim statement: That efficiency rhetoric was already being layered onto a revenue slowdown and investor pressure over unprofitable strategic bets. **Evidence status:** Documented **Fact Source support:** Appendix A8 **Corpus companion entry:** Claim C-038 **Audit Claim ID:** ap-meta-q4-2022-overhiring-claim-02

7.18.6. Appendix B6. Metaverse losses were part of the financial backdrop to Meta's workforce correction

Claim statement: That metaverse losses were part of the financial backdrop to Meta's workforce correction. **Evidence status:** Documented **Fact Source support:** Appendix A8 **Corpus companion entry:** Claim C-039 **Audit Claim ID:** ap-meta-q4-2022-overhiring-claim-03

7.18.7. Appendix B7. Meta's improved financial results were publicly tied to both cost cutting and lay...

Claim statement: That Meta's improved financial results were publicly tied to both cost cutting and layoffs, not just to a magical AI turnaround. **Evidence status:** Documented **Fact Source support:** Appendix A9 **Corpus companion entry:** Claim C-040 **Audit Claim ID:** ap-meta-q4-2023-year-of-efficiency-claim-01

7.18.8. Appendix B8. AI hiring resumed selectively even after a broad labor correction, reinforcing th...

Claim statement: That AI hiring resumed selectively even after a broad labor correction, reinforcing the idea of reallocation rather than simple technological replacement. **Evidence status:**

Documented Fact **Source support:** Appendix A9 **Corpus companion entry:** Claim C-041 **Audit Claim ID:** ap-meta-q4-2023-year-of-efficiency-claim-02

7.18.9. Appendix B9. Metaverse-related losses remained large even as the company regained investor con...

Claim statement: That metaverse-related losses remained large even as the company regained investor confidence. **Evidence status:** Documented Fact **Source support:** Appendix A9 **Corpus companion entry:** Claim C-042 **Audit Claim ID:** ap-meta-q4-2023-year-of-efficiency-claim-03

7.18.10. Appendix B10. The layoff cycle was large enough to materially change worker expectations across...

Claim statement: That the layoff cycle was large enough to materially change worker expectations across the sector. **Evidence status:** Documented Fact **Source support:** Appendix A10 **Corpus companion entry:** Claim C-102 **Audit Claim ID:** axios-tech-layoffs-worker-confidence-2024-05-talent

7.18.11. Appendix B11. Chapter 5 can discuss the correction as a broad labor-market event rather than a...

Claim statement: That Chapter 5 can discuss the correction as a broad labor-market event rather than a collection of isolated firm stories. **Evidence status:** Documented Fact **Source support:** Appendix A10 **Corpus companion entry:** Claim C-103 **Audit Claim ID:** axios-tech-layoffs-worker-confidence-2024-05-talent-hoarding-section174-layoffs-claim-02

7.18.12. Appendix B12. A primary multilateral financial institution explicitly warned in June 2026 that...

Claim statement: That a primary multilateral financial institution explicitly warned in June 2026 that the AI capex race was outrunning earnings and free cash flow. **Evidence status:** Documented Fact **Source support:** Appendix A11 **Corpus companion entry:** Claim C-107 **Audit Claim ID:** bis-annual-economic-report-2026-05-talent-hoarding-section174-layoffs-claim-01

7.18.13. Appendix B13. Debt issuance, leverage, opaque financing structures, and circular financing belo...

Claim statement: That debt issuance, leverage, opaque financing structures, and circular financing belong to the documented risk structure of the AI buildout. **Evidence status:** Documented Fact **Source support:** Appendix A11 **Corpus companion entry:** Claim C-108 **Audit Claim ID:** bis-annual-economic-report-2026-05-talent-hoarding-section174-layoffs-claim-02

7.18.14. Appendix B14. Chapter 05 can support the thesis of structural correction and financial overhang...

Claim statement: That Chapter 05 can support the thesis of structural correction and financial overhang without relying only on opinion pieces or market chatter. **Evidence status:** Documented

Fact Source support: Appendix A11 **Corpus companion entry:** Claim C-109 **Audit Claim ID:** bis-annual-economic-report-2026-05-talent-hoarding-section174-layoffs-claim-03

7.18.15. Appendix B15. Chapters 06, 09, and 10 can discuss the “cheap AI” narrative against a documented...

Claim statement: That Chapters 06, 09, and 10 can discuss the “cheap AI” narrative against a documented backdrop of subsidy-like financing pressure and fragile commercial returns. **Evidence status:** Documented **Fact Source support:** Appendix A11 **Corpus companion entry:** Claim C-110 **Audit Claim ID:** bis-annual-economic-report-2026-05-talent-hoarding-section174-layoffs-claim-

7.18.16. Appendix B16. Congress formally treated the post-2021 Section 174 regime as burdensome enough t...

Claim statement: That Congress formally treated the post-2021 Section 174 regime as burdensome enough to warrant legislative reversal or postponement. **Evidence status:** Documented **Fact Source support:** Appendix A6 **Corpus companion entry:** Claim C-123 **Audit Claim ID:** congress-hr7024-summary-2024-05-talent-hoarding-section174-layoffs-claim-01

7.18.17. Appendix B17. The practical effect of the existing rule can be described in plain English: firm...

Claim statement: That the practical effect of the existing rule can be described in plain English: firms were being pushed into five-year deduction treatment for domestic research costs until the proposed relief would restore expensing through 2025. **Evidence status:** Documented **Fact Source support:** Appendix A6 **Corpus companion entry:** Claim C-124 **Audit Claim ID:** congress-hr7024-summary-2024-05-talent-hoarding-section174-layoffs-claim-02

7.18.18. Appendix B18. Chapter 5 can treat Section 174 as a material structural factor in the post-pande...

Claim statement: That Chapter 5 can treat Section 174 as a material structural factor in the post-pandemic correction without pretending it was the only cause of layoffs. **Evidence status:** Documented **Fact Source support:** Appendix A6 **Corpus companion entry:** Claim C-125 **Audit Claim ID:** congress-hr7024-summary-2024-05-talent-hoarding-section174-layoffs-claim-03

7.18.19. Appendix B19. Congress formally moved to soften or reverse the immediate impact of Section 174...

Claim statement: That Congress formally moved to soften or reverse the immediate impact of Section 174 on domestic research expensing. **Evidence status:** Documented **Fact Source support:** Appendix A5 **Corpus companion entry:** Claim C-126 **Audit Claim ID:** congress-hr7024-tax-relief-2024-claim-01

7.18.20. Appendix B20. The practical consequences of Section 174 were concrete enough to generate a seri...

Claim statement: That the practical consequences of Section 174 were concrete enough to generate a serious legislative response. **Evidence status:** Documented Fact **Source support:** Appendix A5 **Corpus companion entry:** Claim C-127 **Audit Claim ID:** congress-hr7024-tax-relief-2024-claim-02

7.18.21. Appendix B21. Chapter 5 can frame Section 174 as a live business issue rather than an after-the...

Claim statement: That Chapter 5 can frame Section 174 as a live business issue rather than an after-the-fact excuse invented by commentators. **Evidence status:** Documented Fact **Source support:** Appendix A5 **Corpus companion entry:** Claim C-128 **Audit Claim ID:** congress-hr7024-tax-relief-2024-claim-03

7.18.22. Appendix B22. The practical Section 174 fight directly implicated software development, not onl...

Claim statement: That the practical Section 174 fight directly implicated software development, not only laboratory research. **Evidence status:** Documented Fact **Source support:** Appendix A7 **Corpus companion entry:** Claim C-129 **Audit Claim ID:** congress-hr7024-text-2024-05-talent-hoarding-section174-layoffs-claim-01

7.18.23. Appendix B23. The legislative response was not vague rhetoric; it proposed specific statutory l...

Claim statement: That the legislative response was not vague rhetoric; it proposed specific statutory language restoring domestic deductions and naming software explicitly. **Evidence status:** Documented Fact **Source support:** Appendix A7 **Corpus companion entry:** Claim C-130 **Audit Claim ID:** congress-hr7024-text-2024-05-talent-hoarding-section174-layoffs-claim-02

7.18.24. Appendix B24. Chapter 5 can connect the accounting treatment of software and R&D spending to th...

Claim statement: That Chapter 5 can connect the accounting treatment of software and R&D spending to the financial environment of large tech employers. **Evidence status:** Documented Fact **Source support:** Appendix A7 **Corpus companion entry:** Claim C-131 **Audit Claim ID:** congress-hr7024-text-2024-05-talent-hoarding-section174-layoffs-claim-03

7.18.25. Appendix B25. AI enthusiasm did not eliminate the broader structural weakness of the sector

Claim statement: That AI enthusiasm did not eliminate the broader structural weakness of the sector. **Evidence status:** Documented Fact **Source support:** Appendix A4 **Corpus companion entry:** Claim C-173 **Audit Claim ID:** ft-ai-exuberance-tech-weakness-2024-05-talent-hoarding-sec

7.18.26. Appendix B26. A serious T2 source explicitly connects post-pandemic overhiring, weak demand, an...

Claim statement: That a serious T2 source explicitly connects post-pandemic overhiring, weak demand, and investment reallocation to the industry's ongoing malaise. **Evidence status:** Documented Fact **Source support:** Appendix A4 **Corpus companion entry:** Claim C-174 **Audit Claim ID:** ft-ai-exuberance-tech-weakness-2024-05-talent-hoarding-section174-layoffs-claim-02

7.18.27. Appendix B27. Chapter 5 can argue AI was layered onto an existing correction rather than causin...

Claim statement: That Chapter 5 can argue AI was layered onto an existing correction rather than causing it from scratch. **Evidence status:** Documented Fact **Source support:** Appendix A4 **Corpus companion entry:** Claim C-175 **Audit Claim ID:** ft-ai-exuberance-tech-weakness-2024-05-talent-hoarding-section174-layoffs-claim-03

7.18.28. Appendix B28. Section 174 amortization was not a rumor or niche accounting theory; it was the s...

Claim statement: That Section 174 amortization was not a rumor or niche accounting theory; it was the subject of active IRS guidance. **Evidence status:** Documented Fact **Source support:** Appendix A1 **Corpus companion entry:** Claim C-240 **Audit Claim ID:** irs-notice-2023-63-claim-01

7.18.29. Appendix B29. U.S. firms with research-heavy cost structures had to account for a changed tax t...

Claim statement: That U.S. firms with research-heavy cost structures had to account for a changed tax treatment of those expenditures. **Evidence status:** Documented Fact **Source support:** Appendix A1 **Corpus companion entry:** Claim C-241 **Audit Claim ID:** irs-notice-2023-63-claim-02

7.18.30. Appendix B30. Chapter 5 can legitimately treat Section 174 as one documented structural factor...

Claim statement: That Chapter 5 can legitimately treat Section 174 as one documented structural factor in the financial environment facing software and R&D-heavy companies. **Evidence status:** Documented Fact **Source support:** Appendix A1 **Corpus companion entry:** Claim C-242 **Audit Claim ID:** irs-notice-2023-63-claim-03

7.18.31. Appendix B31. Meta itself framed layoffs as part of a structural response to financial conditio...

Claim statement: That Meta itself framed layoffs as part of a structural response to financial conditions and organizational inefficiency. **Evidence status:** Documented Fact **Source support:** Appendix A3 **Corpus companion entry:** Claim C-272 **Audit Claim ID:** meta-year-of-efficiency-2023-claim-01

7.18.32. Appendix B32. AI investment was presented alongside workforce reduction, not as a later justifi...

Claim statement: That AI investment was presented alongside workforce reduction, not as a later justification invented by outsiders. **Evidence status:** Documented Fact **Source support:** Appendix A3 **Corpus companion entry:** Claim C-273 **Audit Claim ID:** meta-year-of-efficiency-2023-claim-02

7.18.33. Appendix B33. Interest rates, regulation, and margin discipline were explicit parts of the comp...

Claim statement: That interest rates, regulation, and margin discipline were explicit parts of the company's own explanation. **Evidence status:** Documented Fact **Source support:** Appendix A3 **Corpus companion entry:** Claim C-274 **Audit Claim ID:** meta-year-of-efficiency-2023-claim-03

7.18.34. Appendix B34. AI was often used as a legitimizing narrative layered on top of an existing struc...

Claim statement: That AI was often used as a legitimizing narrative layered on top of an existing structural correction driven by overhiring, rates, margin pressure, and tax-treatment shifts rather than as the sole cause of the cuts. **Evidence status:** Hypothesis / Interpretation **Source support:** Appendix A1, Appendix A3, Appendix A4, Appendix A11 **Corpus companion entry:** Claim C-431 **Audit Claim ID:** vector05-legitimizing-narrative-synthesis-claim-01

7.18.35. Appendix B35. AI alone caused the post-pandemic tech layoff cycle or fully explains wage compre...

Claim statement: That AI alone caused the post-pandemic tech layoff cycle or fully explains wage compression and labor insecurity across the sector. **Evidence status:** Speculative Narrative Risk **Source support:** Appendix A1, Appendix A8, Appendix A11 **Corpus companion entry:** Claim C-430 **Audit Claim ID:** vector05-ai-alone-caused-layoffs-overclaim-claim-01

7.19. Appendix C. Timeline Slice

7.19.1. Appendix C1. 2023-01-04: Amazon and Salesforce tie layoffs to overhiring and cost structure

Event summary: Amazon and Salesforce tie layoffs to overhiring and cost structure **Event date:** 2023-01-04 **Source support:** Appendix A2 **Corpus companion entry:** Event T-020 **Audit Event ID:** event-amazon-salesforce-purge-2023-01-04

7.19.2. Appendix C2. 2023-03-14: Meta announces deeper cuts in its Year of Efficiency

Event summary: Meta announces deeper cuts in its Year of Efficiency **Event date:** 2023-03-14 **Source support:** Appendix A3 **Corpus companion entry:** Event T-022 **Audit Event ID:** event-meta-efficiency-2023-03-14

7.19.3. Appendix C3. 2023-09-08: The IRS issues Notice 2023-63 on Section 174 expenditures

Event summary: The IRS issues Notice 2023-63 on Section 174 expenditures **Event date:** 2023-09-08 **Source support:** Appendix A1 **Corpus companion entry:** Event T-029 **Audit Event ID:** event-irs-section174-2023-09-08

7.19.4. Appendix C4. 2024-01-31: The House passes H.R. 7024 with Section 174 relief language

Event summary: The House passes H.R. 7024 with Section 174 relief language **Event date:** 2024-01-31 **Source support:** Appendix A5 **Corpus companion entry:** Event T-038 **Audit Event ID:** event-hr7024-passed-house-2024-01-31

7.19.5. Appendix C5. 2024-01-31: Congressional summary explains temporary domestic R&D expensing relief

Event summary: Congressional summary explains temporary domestic R&D expensing relief **Event date:** 2024-01-31 **Source support:** Appendix A6 **Corpus companion entry:** Event T-039 **Audit Event ID:** event-hr7024-summary-2024-01-31

7.19.6. Appendix C6. 2024-03-21: H.R. 7024 text explicitly covers domestic R&D deductions and software

Event summary: H.R. 7024 text explicitly covers domestic R&D deductions and software **Event date:** 2024-03-21 **Source support:** Appendix A7 **Corpus companion entry:** Event T-046 **Audit Event ID:** event-hr7024-text-2024-03-21

7.19.7. Appendix C7. 2024-09-09: Financial Times says AI exuberance masks broader tech weakness

Event summary: Financial Times says AI exuberance masks broader tech weakness **Event date:** 2024-09-09 **Source support:** Appendix A4 **Corpus companion entry:** Event T-054 **Audit Event ID:** event-ft-tech-weakness-ai-2024-09-09

7.20. Appendix D. Relevant Entities

7.20.1. Appendix D1. Reality Labs

Entity type: business_unit **Role in this paper:** Actor relevant to 12_cambridge_analytica_meta_metaverse_r
05_talent_hoarding_section174_layoffs. **Relevant sources in this paper:** Appendix A8, Appendix A9 **Corpus companion entry:** Entity E-031 **Audit Entity ID:** org-reality-labs

7.20.2. Appendix D2. Salesforce

Entity type: company **Role in this paper:** Actor relevant to 05_talent_hoarding_section174_layoffs. **Relevant sources in this paper:** Appendix A2, Appendix A10 **Corpus companion entry:** Entity E-033 **Audit Entity ID:** org-salesforce

7.20.3. Appendix D3. Internal Revenue Service**Entity type:** government **Role in this paper:** Actor relevant to 05_talent_hoarding_section174_layoffs.**Relevant sources in this paper:** Appendix A1 **Corpus companion entry:** Entity E-023**Audit Entity ID:** org-irs**7.20.4. Appendix D4. U.S. Congress / Congress.gov / CRS****Entity type:** government **Role in this paper:** Actor relevant to 05_talent_hoarding_section174_layoffs.**Relevant sources in this paper:** Appendix A5, Appendix A6, Appendix A7 **Corpus companion entry:** Entity E-036 **Audit Entity ID:** org-us-congress**7.20.5. Appendix D5. Andy Jassy****Entity type:** person **Role in this paper:** Actor relevant to 05_talent_hoarding_section174_layoffs.**Relevant sources in this paper:** Appendix A2 **Corpus companion entry:** Entity E-042**Audit Entity ID:** person-andy-jassy**7.20.6. Appendix D6. Marc Benioff****Entity type:** person **Role in this paper:** Actor relevant to 05_talent_hoarding_section174_layoffs.**Relevant sources in this paper:** Appendix A2 **Corpus companion entry:** Entity E-053**Audit Entity ID:** person-marc-benioff**7.20.7. Appendix D7. 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 A3 **Corpus companion entry:** Entity E-054**Audit Entity ID:** person-mark-zuckerberg**7.21. Appendix E. Evidence Boundaries****7.21.1. Appendix E1. Disputed Matters**

- None registered for this paper.

7.21.2. Appendix E2. Interpretive Boundaries

- Appendix B34: That AI was often used as a legitimizing narrative layered on top of an existing structural correction driven by overhiring, rates, margin pressure, and tax-treatment shifts rather than as the sole cause of the cuts.

7.21.3. Appendix E3. Speculative Narrative Risks

- Appendix B35: That AI alone caused the post-pandemic tech layoff cycle or fully explains wage compression and labor insecurity across the sector.

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

- Appendix A2. Companion note: Note N-039. Risk flag: low. Note focus: That major tech layoffs were publicly tied to pandemic overexpansion and deteriorating macro conditions.
- Appendix A8. Companion note: Note N-040. Risk flag: low. Note focus: That a major tech company publicly linked layoffs to pandemic overhiring and slowing post-lockdown growth.
- Appendix A9. Companion note: Note N-041. Risk flag: low. Note focus: That Meta's improved financial results were publicly tied to both cost cutting and layoffs, not just to a magical AI turnaround.
- Appendix A10. Companion note: Note N-042. Risk flag: medium. Note focus: That the layoff cycle was large enough to materially change worker expectations across the sector.
- Appendix A11. Companion note: Note N-043. Risk flag: review. Note focus: That a primary multilateral financial institution explicitly warned in June 2026 that the AI capex race was outrunning earnings and free cash flow.
- Appendix A6. Companion note: Note N-044. Risk flag: low. Note focus: That Congress formally treated the post-2021 Section 174 regime as burdensome enough to warrant legislative reversal or postponement.
- Appendix A5. Companion note: Note N-045. Risk flag: low. Note focus: That Congress formally moved to soften or reverse the immediate impact of Section 174 on domestic research expensing.
- Appendix A7. Companion note: Note N-046. Risk flag: low. Note focus: That the practical Section 174 fight directly implicated software development, not only laboratory research.
- Appendix A4. Companion note: Note N-047. Risk flag: low. Note focus: That AI enthusiasm did not eliminate the broader structural weakness of the sector.
- Appendix A1. Companion note: Note N-048. Risk flag: medium. Note focus: That Section 174 amortization was not a rumor or niche accounting theory; it was the subject of active IRS guidance.
- Appendix A3. Companion note: Note N-049. Risk flag: medium. Note focus: That Meta itself framed layoffs as part of a structural response to financial conditions and organizational inefficiency.

8. Chapter 06. AI as a dual-use rhetorical weapon

8.1. Abstract

One of the clearest contradictions in the current AI regime is that systems are marketed simultaneously as dangerous enough to justify elite governance and safe enough to be rolled out as broad dependence infrastructure. OpenAI publicly framed reasoning as a distinct capability layer tied to additional train-time and test-time compute, and paired it with preparedness, safety, and misuse-risk language. This is how AI becomes a dual-use rhetorical weapon.

8.2. Keywords

AI as a dual-use rhetorical weapon; surveillance capitalism; AI imperialism; evidence-based dossier

8.3. Research Question

How does ai as a dual-use rhetorical weapon function within the dossier's larger argument, and what does the documentary record allow this chapter to establish?

8.4. Scope and Framing Note

This paper examines ai as a dual-use rhetorical weapon 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 shows how danger rhetoric, cheapness rhetoric, and reasoning rhetoric together normalize frontier enclosure and strategic dependence. Adjacent reinforcement appears in Chapter 04, Chapter 05, Chapter 07, Chapter 10, but those cross-references are supportive rather than required for basic comprehension.

8.5. Central Thesis

One of the clearest contradictions in the current AI regime is that systems are marketed simultaneously as dangerous enough to justify elite governance and safe enough to be rolled out as broad dependence infrastructure. The point is not that reasoning advances are fake. The point is that real advances are being packaged inside a politics of fear, prestige, and access management at the same time. ⁶⁵

8.6. Evidence and Method Note

This paper is derived from the chapter corpus for Chapter 06 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.

⁶⁵Sources: OpenAI, Learning to reason with LLMs (2024-09-12; see Appendix A1); OpenAI, OpenAI o1 System Card (2024-12-05; see Appendix A2); Financial Times, OpenAI acknowledges new models increa... (2024-09-13; see Appendix A3)

8.7. Introduction

This paper asks: How does ai as a dual-use rhetorical weapon function within the dossier’s larger argument, and what does the documentary record allow this chapter to establish? Its working answer is clear from the start: One of the clearest contradictions in the current AI regime is that systems are marketed simultaneously as dangerous enough to justify elite governance and safe enough to be rolled out as broad dependence infrastructure. The point is not that reasoning advances are fake. The point is that real advances are being packaged inside a politics of fear, prestige, and access management at the same time.⁶⁶ 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.

8.8. Historical or Institutional Context

OpenAI publicly framed reasoning as a distinct capability layer tied to additional train-time and test-time compute, and paired it with preparedness, safety, and misuse-risk language. The product story and the security story arrived together.⁶⁷

8.9. Main Analysis

The Reuters Q* and **Strawberry** reporting shows that this reasoning story has a serious public prehistory rather than being a later marketing improvisation. That matters because it confirms both sides of the contradiction: the capability arc may be real, but so is the later struggle over how it is narrated, priced, and gated.⁶⁸

Primary research strengthens the middle claim. Test-time compute is not just a corporate slogan. A 2024 paper on compute-optimal inference-time scaling showed meaningful efficiency improvements over simpler baselines, while later research and open-replication efforts showed both rapid diffusion pressure and an active technical dispute over what those replications actually reproduced. In other words, the dossier no longer needs to pretend that reasoning scarcity is fake in order to criticize how it is governed.⁶⁹

The pricing layer is also direct now. OpenAI’s own plan and API pages, along with Anthropic’s pricing and **Claude 4** launch material, show that stronger reasoning access is tiered across premium plans, higher usage classes, and higher-priced API lanes. Reasoning is not only a capability. It is

⁶⁶Sources: OpenAI, Learning to reason with LLMs (2024-09-12; see Appendix A1); OpenAI, OpenAI o1 System Card (2024-12-05; see Appendix A2); Financial Times, OpenAI acknowledges new models increa... (2024-09-13; see Appendix A3)

⁶⁷Sources: OpenAI, Learning to reason with LLMs (2024-09-12; see Appendix A1); OpenAI, OpenAI o1 System Card (2024-12-05; see Appendix A2); Financial Times, OpenAI acknowledges new models increa... (2024-09-13; see Appendix A3)

⁶⁸Sources: Reuters, OpenAI researchers warned board of AI... (2023-11-23; see Appendix A4); Reuters, Exclusive: OpenAI working on new reas... (2024-07-15; see Appendix A5)

⁶⁹Sources: Charlie Snell; Jaehoon Lee; Kelvin Xu; Aviral Kumar, Scaling LLM Test-Time Compute Optimal... (2024-08-06; see Appendix A6); Niklas Muennighoff; Zitong Yang; Weijia Shi; Xiang Lisa Li; Li Fei-Fei; Hannaneh Hajishirzi; Luke Zettlemoyer; Percy Liang; Emmanuel Candès; Tatsunori Hashimoto, s1: Simple test-time scaling (2025-01-31; see Appendix A7); Guojun Wu, It’s Not That Simple. An Analysis of... (2025-07-19; see Appendix A8)

also a gated commercial product.⁷⁰

DeepSeek complicates total exclusivity claims without eliminating concentration. Its releases show that parts of reasoning capability diffuse faster than absolutist frontier rhetoric suggests. The real story is not zero diffusion versus perfect monopoly. It is diffusion at the edges and political curation at the core.⁷¹

Interpretation: this chapter supports the patrimony argument by showing how a real compute-intensive technical mechanism is turned into a managed access regime. Fear legitimizes the gate, convenience fills the funnel, and premium reasoning tiers convert collectively grounded capability into selectively governed scarcity.⁷²

8.10. Position Within the Series

This paper sits inside a coordinated 11-paper architecture. Its nearest companions are Chapter 04 (AI arrives when the human raw material already exists), Chapter 05 (Talent, overhiring, valuation, and layoffs), Chapter 07 (Imperial control over access), Chapter 10 (What to do). In that series logic, Chapter 06 shows how danger rhetoric, cheapness rhetoric, and reasoning rhetoric together normalize frontier enclosure and strategic dependence.

8.11. Counterarguments

The chapter should not collapse three separate issues into one: the reality of test-time compute as a technical mechanism, the open challenge to frontier mystique, and the unresolved dispute over what simple replications actually prove. The strongest claim is about rhetoric, gating, and governance, not about proving that every frontier capability turned out to be hollow.⁷³

8.12. Limits of the Evidence

The chapter should not collapse three separate issues into one: the reality of test-time compute as a technical mechanism, the open challenge to frontier mystique, and the unresolved dispute over what simple replications actually prove. The strongest claim is about rhetoric, gating, and governance,

⁷⁰Sources: OpenAI, ChatGPT Plans | Free, Go, Plus, Pro,... (n.d.; see Appendix A9); OpenAI, Pricing | OpenAI API (n.d.; see Appendix A10); Anthropic, Plans & Pricing | Claude by Anthropic (n.d.; see Appendix A11); Anthropic, Introducing Claude 4 (2025-05-22; see Appendix A12)

⁷¹Sources: arXiv, DeepSeek-V3 Technical Report (2024-12-27; see Appendix A13); arXiv, DeepSeek-R1: Incentivizing Reasoning... (2025-01-22; see Appendix A14); Niklas Muennighoff; Zitong Yang; Weijia Shi; Xiang Lisa Li; Li Fei-Fei; Hannaneh Hajishirzi; Luke Zettlemoyer; Percy Liang; Emmanuel Candès; Tatsunori Hashimoto, s1: Simple test-time scaling (2025-01-31; see Appendix A7)

⁷²Sources: OpenAI, OpenAI o1 System Card (2024-12-05; see Appendix A2); OpenAI, ChatGPT Plans | Free, Go, Plus, Pro,... (n.d.; see Appendix A9); Anthropic, Plans & Pricing | Claude by Anthropic (n.d.; see Appendix A11); Charlie Snell; Jaehoon Lee; Kelvin Xu; Aviral Kumar, Scaling LLM Test-Time Compute Optimal... (2024-08-06; see Appendix A6)

⁷³Sources: Charlie Snell; Jaehoon Lee; Kelvin Xu; Aviral Kumar, Scaling LLM Test-Time Compute Optimal... (2024-08-06; see Appendix A6); Niklas Muennighoff; Zitong Yang; Weijia Shi; Xiang Lisa Li; Li Fei-Fei; Hannaneh Hajishirzi; Luke Zettlemoyer; Percy Liang; Emmanuel Candès; Tatsunori Hashimoto, s1: Simple test-time scaling (2025-01-31; see Appendix A7); Guojun Wu, It's Not That Simple. An Analysis of... (2025-07-19; see Appendix A8)

not about proving that every frontier capability turned out to be hollow.⁷⁴

8.13. Reform Relevance

This is how AI becomes a dual-use rhetorical weapon. Danger justifies concentration, while convenience justifies mass adoption. A public told that the system is too dangerous to democratize and too useful to refuse is being prepared to accept hierarchy as prudence. That pushes Chapter 04's ownership question toward a sharper answer: capability built from collective human production is being enclosed as a governed premium.⁷⁵

8.14. Conclusion

This is how AI becomes a dual-use rhetorical weapon. Once managed dependency exists, the next question is who controls the most consequential access lanes and under what geopolitical terms. In the architecture of this series, Chapter 06 therefore functions as a self-contained argument while also advancing the cumulative path toward the dossier's three reforms.

8.15. 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 06, those boundaries remain visible in the prose, in the bibliography, and in the appendix sections that summarize claims and caution points.

8.16. References / Bibliography

1. OpenAI. Learning to reason with LLMs. 2024-09-12. T1 release. Paper appendix: Appendix A1. Corpus companion entry: Source S-093. Audit ID: `openai-learning-to-reason-with-llms-2024`. Relevance: That OpenAI publicly framed reasoning as a new scaling layer based on additional train-time and test-time compute.
2. OpenAI. OpenAI o1 System Card. 2024-12-05. T1 system_card. Paper appendix: Appendix A2. Corpus companion entry: Source S-095. Audit ID: `openai-o1-system-card-2024`. Relevance: That OpenAI itself paired stronger reasoning capability with formal safety, preparedness, and misuse-risk framing.
3. Financial Times. OpenAI acknowledges new models increase risk of misuse to create bioweapons. 2024-09-13. T2 news_article. Paper appendix: Appendix A3. Corpus companion entry: Source S-060. Audit ID: `ft-openai-o1-bioweapon-risk-2024`. Relevance: That serious T2 reporting documented OpenAI publicly acknowledging higher misuse risk at the same moment it was releasing stronger reasoning models.

⁷⁴Sources: Charlie Snell; Jaehoon Lee; Kelvin Xu; Aviral Kumar, Scaling LLM Test-Time Compute Optimal... (2024-08-06; see Appendix A6); Niklas Muennighoff; Zitong Yang; Weijia Shi; Xiang Lisa Li; Li Fei-Fei; Hannaneh Hajishirzi; Luke Zettlemoyer; Percy Liang; Emmanuel Candès; Tatsunori Hashimoto, s1: Simple test-time scaling (2025-01-31; see Appendix A7); Guojun Wu, It's Not That Simple. An Analysis of... (2025-07-19; see Appendix A8)

⁷⁵Sources: OpenAI, OpenAI o1 System Card (2024-12-05; see Appendix A2); OpenAI, ChatGPT Plans | Free, Go, Plus, Pro,... (n.d.; see Appendix A9); Anthropic, Introducing Claude 4 (2025-05-22; see Appendix A12)

4. Reuters. OpenAI researchers warned board of AI breakthrough ahead of CEO ouster, sources say. 2023-11-23. T2 news_article. Paper appendix: Appendix A4. Corpus companion entry: Source S-110. Audit ID: **reuters-qstar-board-warning-2023**. Relevance: That by November 23, 2023 there was serious mainstream reporting explicitly linking the board crisis to a letter about an AI breakthrough.
5. Reuters. Exclusive: OpenAI working on new reasoning technology under code name ‘Strawberry’. 2024-07-15. T2 news_article. Paper appendix: Appendix A5. Corpus companion entry: Source S-111. Audit ID: **reuters-strawberry-reasoning-2024**. Relevance: That by mid-July 2024 Reuters had published serious reporting about an OpenAI reasoning project under the codename **Strawberry**.
6. Charlie Snell; Jaehoon Lee; Kelvin Xu; Aviral Kumar. Scaling LLM Test-Time Compute Optimally can be More Effective than Scaling Model Parameters. 2024-08-06. T1 paper. Paper appendix: Appendix A6. Corpus companion entry: Source S-115. Audit ID: **snell-test-time-compute-2024**. Relevance: That test-time compute scaling is a documented technical mechanism rather than only a commercial branding story.
7. Niklas Muennighoff; Zitong Yang; Weijia Shi; Xiang Lisa Li; Li Fei-Fei; Hannaneh Hajishirzi; Luke Zettlemoyer; Percy Liang; Emmanuel Candès; Tatsunori Hashimoto. s1: Simple test-time scaling. 2025-01-31. T1 paper. Paper appendix: Appendix A7. Corpus companion entry: Source S-112. Audit ID: **s1-simple-test-time-scaling-2025**. Relevance: That OpenAI’s reasoning launch quickly triggered open replication efforts focused on test-time scaling.
8. Guojun Wu. It’s Not That Simple. An Analysis of Simple Test-Time Scaling. 2025-07-19. T1 paper. Paper appendix: Appendix A8. Corpus companion entry: Source S-076. Audit ID: **its-not-that-simple-test-time-scaling-2025**. Relevance: That there is direct technical pushback against simplistic claims that frontier reasoning was trivially replicated.
9. OpenAI. ChatGPT Plans | Free, Go, Plus, Pro, Business, and Enterprise. n.d.. T1 pricing_page. Paper appendix: Appendix A9. Corpus companion entry: Source S-039. Audit ID: **chatgpt-pricing-2026**. Relevance: That as of capture on July 3, 2026, OpenAI’s official ChatGPT plan structure reserved the highest-end reasoning access for upper paid tiers rather than offering it uniformly acr...
10. OpenAI. Pricing | OpenAI API. n.d.. T1 pricing_page. Paper appendix: Appendix A10. Corpus companion entry: Source S-088. Audit ID: **openai-api-pricing-2026**. Relevance: That as of July 3, 2026, OpenAI’s official API pricing imposed a large price premium on GPT-5.5 Pro relative to GPT-5.5.
11. Anthropic. Plans & Pricing | Claude by Anthropic. n.d.. T1 pricing_page. Paper appendix: Appendix A11. Corpus companion entry: Source S-005. Audit ID: **anthropic-claude-pricing-2026**. Relevance: That as of July 3, 2026, Anthropic’s official current pricing structure tied higher usage, priority access, and stronger feature access to premium paid plans.
12. Anthropic. Introducing Claude 4. 2025-05-22. T1 release. Paper appendix: Appendix A12. Corpus companion entry: Source S-003. Audit ID: **anthropic-claude-4-release-2025**. Relevance: That on May 22, 2025 Anthropic launched reasoning-oriented hybrid frontier models together with immediate commercial distribution across subscription plans, API access, and majo...
13. arXiv. DeepSeek-V3 Technical Report. 2024-12-27. T1 paper. Paper appendix: Appendix A13.

Corpus companion entry: Source S-049. Audit ID: `deepseek-v3-technical-report-2024`. Relevance: That by late 2024 there was already a competitor publicly claiming frontier-adjacent performance through a technical report.

14. arXiv. DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning. 2025-01-22. T1 paper. Paper appendix: Appendix A14. Corpus companion entry: Source S-048. Audit ID: `deepseek-r1-reinforcement-learning-2025`. Relevance: That by January 2025 a non-U.S. lab was publicly claiming frontier-adjacent reasoning performance in a primary technical paper.

8.17. Appendix A. Source Register for This Paper

8.17.1. Appendix A1. Learning to reason with LLMs

Source citation: Learning to reason with LLMs **Institution / author:** OpenAI **Date:** 2024-09-12 **Tier / type:** T1 / release **Why it matters in this paper:** That OpenAI publicly framed reasoning as a new scaling layer based on additional train-time and test-time compute. **Corpus companion entry:** Source S-093 **Audit ID:** `openai-learning-to-reason-with-llms-2024`

8.17.2. Appendix A2. OpenAI o1 System Card

Source citation: OpenAI o1 System Card **Institution / author:** OpenAI **Date:** 2024-12-05 **Tier / type:** T1 / system_card **Why it matters in this paper:** That OpenAI itself paired stronger reasoning capability with formal safety, preparedness, and misuse-risk framing. **Corpus companion entry:** Source S-095 **Audit ID:** `openai-o1-system-card-2024`

8.17.3. Appendix A3. OpenAI acknowledges new models increase risk of misuse to create bioweapons

Source citation: OpenAI acknowledges new models increase risk of misuse to create bioweapons **Institution / author:** Financial Times **Date:** 2024-09-13 **Tier / type:** T2 / news_article **Why it matters in this paper:** That serious T2 reporting documented OpenAI publicly acknowledging higher misuse risk at the same moment it was releasing stronger reasoning models. **Corpus companion entry:** Source S-060 **Audit ID:** `ft-openai-o1-bioweapon-risk-2024`

8.17.4. Appendix A4. OpenAI researchers warned board of AI breakthrough ahead of CEO ouster, sources say

Source citation: OpenAI researchers warned board of AI breakthrough ahead of CEO ouster, sources say **Institution / author:** Reuters **Date:** 2023-11-23 **Tier / type:** T2 / news_article **Why it matters in this paper:** That by November 23, 2023 there was serious mainstream reporting explicitly linking the board crisis to a letter about an AI breakthrough. **Corpus companion entry:** Source S-110 **Audit ID:** `reuters-qstar-board-warning-2023`

8.17.5. Appendix A5. Exclusive: OpenAI working on new reasoning technology under code name ‘Strawberry’

Source citation: Exclusive: OpenAI working on new reasoning technology under code name ‘Strawberry’ **Institution / author:** Reuters **Date:** 2024-07-15 **Tier / type:** T2 / news_article **Why it matters in this paper:** That by mid-July 2024 Reuters had published serious reporting about an OpenAI reasoning project under the codename Strawberry. **Corpus companion entry:** Source S-111 **Audit ID:** reuters-strawberry-reasoning-2024

8.17.6. Appendix A6. Scaling LLM Test-Time Compute Optimally can be More Effective than Scaling Model Parameters

Source citation: Scaling LLM Test-Time Compute Optimally can be More Effective than Scaling Model Parameters **Institution / author:** Charlie Snell; Jaehoon Lee; Kelvin Xu; Aviral Kumar **Date:** 2024-08-06 **Tier / type:** T1 / paper **Why it matters in this paper:** That test-time compute scaling is a documented technical mechanism rather than only a commercial branding story. **Corpus companion entry:** Source S-115 **Audit ID:** snell-test-time-compute-2024

8.17.7. Appendix A7. s1: Simple test-time scaling

Source citation: s1: Simple test-time scaling **Institution / author:** Niklas Muennighoff; Zitong Yang; Weijia Shi; Xiang Lisa Li; Li Fei-Fei; Hannaneh Hajishirzi; Luke Zettlemoyer; Percy Liang; Emmanuel Candès; Tatsunori Hashimoto **Date:** 2025-01-31 **Tier / type:** T1 / paper **Why it matters in this paper:** That OpenAI’s reasoning launch quickly triggered open replication efforts focused on test-time scaling. **Corpus companion entry:** Source S-112 **Audit ID:** s1-simple-test-time-scaling-2025

8.17.8. Appendix A8. It’s Not That Simple. An Analysis of Simple Test-Time Scaling

Source citation: It’s Not That Simple. An Analysis of Simple Test-Time Scaling **Institution / author:** Guojun Wu **Date:** 2025-07-19 **Tier / type:** T1 / paper **Why it matters in this paper:** That there is direct technical pushback against simplistic claims that frontier reasoning was trivially replicated. **Corpus companion entry:** Source S-076 **Audit ID:** its-not-that-simple-test-time-scaling-2025

8.17.9. Appendix A9. ChatGPT Plans | Free, Go, Plus, Pro, Business, and Enterprise

Source citation: ChatGPT Plans | Free, Go, Plus, Pro, Business, and Enterprise **Institution / author:** OpenAI **Date:** n.d. **Tier / type:** T1 / pricing_page **Why it matters in this paper:** That as of capture on July 3, 2026, OpenAI’s official ChatGPT plan structure reserved the highest-end reasoning access for upper paid tiers rather than offering it uniformly across tiers. **Corpus companion entry:** Source S-039 **Audit ID:** chatgpt-pricing-2026

8.17.10. Appendix A10. Pricing | OpenAI API

Source citation: Pricing | OpenAI API **Institution / author:** OpenAI **Date:** n.d. **Tier / type:** T1 / pricing_page **Why it matters in this paper:** That as of July 3, 2026, OpenAI’s

official API pricing imposed a large price premium on GPT-5.5 Pro relative to GPT-5.5. **Corpus companion entry:** Source S-088 **Audit ID:** openai-api-pricing-2026

8.17.11. Appendix A11. Plans & Pricing | Claude by Anthropic

Source citation: Plans & Pricing | Claude by Anthropic **Institution / author:** Anthropic **Date:** n.d. **Tier / type:** T1 / pricing_page **Why it matters in this paper:** That as of July 3, 2026, Anthropic's official current pricing structure tied higher usage, priority access, and stronger feature access to premium paid plans. **Corpus companion entry:** Source S-005 **Audit ID:** anthropic-claude-pricing-2026

8.17.12. Appendix A12. Introducing Claude 4

Source citation: Introducing Claude 4 **Institution / author:** Anthropic **Date:** 2025-05-22 **Tier / type:** T1 / release **Why it matters in this paper:** That on May 22, 2025 Anthropic launched reasoning-oriented hybrid frontier models together with immediate commercial distribution across subscription plans, API access, and major cloud platforms. **Corpus companion entry:** Source S-003 **Audit ID:** anthropic-claude-4-release-2025

8.17.13. Appendix A13. DeepSeek-V3 Technical Report

Source citation: DeepSeek-V3 Technical Report **Institution / author:** arXiv **Date:** 2024-12-27 **Tier / type:** T1 / paper **Why it matters in this paper:** That by late 2024 there was already a competitor publicly claiming frontier-adjacent performance through a technical report. **Corpus companion entry:** Source S-049 **Audit ID:** deepseek-v3-technical-report-2024

8.17.14. Appendix A14. DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning

Source citation: DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning **Institution / author:** arXiv **Date:** 2025-01-22 **Tier / type:** T1 / paper **Why it matters in this paper:** That by January 2025 a non-U.S. lab was publicly claiming frontier-adjacent reasoning performance in a primary technical paper. **Corpus companion entry:** Source S-048 **Audit ID:** deepseek-r1-reinforcement-learning-2025

8.18. Appendix B. Claims Used in This Paper

8.18.1. Appendix B1. On May 22, 2025 Anthropic launched reasoning-oriented hybrid frontier models together with immediate commercial distribution across subscription plans, API access, and major cloud platforms.

Claim statement: That on May 22, 2025 Anthropic launched reasoning-oriented hybrid frontier models together with immediate commercial distribution across subscription plans, API access, and major cloud platforms. **Evidence status:** Documented Fact **Source support:** Appendix A12 **Corpus companion entry:** Claim C-007 **Audit Claim ID:** anthropic-claude-4-release-2025-06-ai-double-use-reasoning-deepseek-claim-01

8.18.2. Appendix B2. Anthropic itself tied extended-thinking capability to differentiated plan availab...

Claim statement: That Anthropic itself tied extended-thinking capability to differentiated plan availability, with full-model access concentrated in paid tiers even while Sonnet 4 also reached free users. **Evidence status:** Documented Fact **Source support:** Appendix A12 **Corpus companion entry:** Claim C-008 **Audit Claim ID:** anthropic-claude-4-release-2025-06-ai-double-use-reasoning

8.18.3. Appendix B3. Reasoning-capable frontier models were commercialized with explicit API price dif...

Claim statement: That reasoning-capable frontier models were commercialized with explicit API price differentiation at launch rather than released only as a research artifact. **Evidence status:** Documented Fact **Source support:** Appendix A12 **Corpus companion entry:** Claim C-009 **Audit Claim ID:** anthropic-claude-4-release-2025-06-ai-double-use-reasoning-deepseek-claim-03

8.18.4. Appendix B4. As of July 3, 2026, Anthropic's official current pricing structure tied higher us...

Claim statement: That as of July 3, 2026, Anthropic's official current pricing structure tied higher usage, priority access, and stronger feature access to premium paid plans. **Evidence status:** Documented Fact **Source support:** Appendix A11 **Corpus companion entry:** Claim C-013 **Audit Claim ID:** anthropic-claude-pricing-2026-06-ai-double-use-reasoning-deepseek-claim-01

8.18.5. Appendix B5. Anthropic's own current pricing also differentiates materially across model famil...

Claim statement: That Anthropic's own current pricing also differentiates materially across model families at the API level rather than presenting frontier access as uniform. **Evidence status:** Documented Fact **Source support:** Appendix A11 **Corpus companion entry:** Claim C-014 **Audit Claim ID:** anthropic-claude-pricing-2026-06-ai-double-use-reasoning-deepseek-claim-02

8.18.6. Appendix B6. Chapter 06 can compare more than one frontier lab when arguing that reasoning-era...

Claim statement: That Chapter 06 can compare more than one frontier lab when arguing that reasoning-era capability is stratified by availability and price. **Evidence status:** Documented Fact **Source support:** Appendix A11 **Corpus companion entry:** Claim C-015 **Audit Claim ID:** anthropic-claude-pricing-2026-06-ai-double-use-reasoning-deepseek-claim-03

8.18.7. Appendix B7. As of capture on July 3, 2026, OpenAI's official ChatGPT plan structure reserved...

Claim statement: That as of capture on July 3, 2026, OpenAI's official ChatGPT plan structure reserved the highest-end reasoning access for upper paid tiers rather than offering it uniformly across all users. **Evidence status:** Documented Fact **Source support:** Appendix A9 **Corpus companion entry:** Claim C-117 **Audit Claim ID:** chatgpt-pricing-2026-06-ai-double-use-reasoning-deepseek-c

8.18.8. Appendix B8. Current official consumer access to reasoning-capable models is stratified by pla...

Claim statement: That current official consumer access to reasoning-capable models is stratified by plan, usage envelope, and feature set. **Evidence status:** Documented Fact **Source support:** Appendix A9 **Corpus companion entry:** Claim C-118 **Audit Claim ID:** chatgpt-pricing-2026-06-ai-double-use-reasoning-deepseek-claim-02

8.18.9. Appendix B9. Chapter 06 can ground its reasoning-premium argument in a primary access-policy s...

Claim statement: That Chapter 06 can ground its reasoning-premium argument in a primary access-policy source rather than relying only on commentary about hype or control. **Evidence status:** Documented Fact **Source support:** Appendix A9 **Corpus companion entry:** Claim C-119 **Audit Claim ID:** chatgpt-pricing-2026-06-ai-double-use-reasoning-deepseek-claim-03

8.18.10. Appendix B10. By January 2025 a non-U.S. lab was publicly claiming frontier-adjacent reasoning...

Claim statement: That by January 2025 a non-U.S. lab was publicly claiming frontier-adjacent reasoning performance in a primary technical paper. **Evidence status:** Documented Fact **Source support:** Appendix A14 **Corpus companion entry:** Claim C-146 **Audit Claim ID:** deepseek-r1-reinforcement-learning-2025-06-ai-double-use-reasoning-deepseek-claim-01

8.18.11. Appendix B11. Reasoning capability was not confined entirely to closed U.S. product stacks and...

Claim statement: That reasoning capability was not confined entirely to closed U.S. product stacks and was partly diffusing through open publication and model release. **Evidence status:** Documented Fact **Source support:** Appendix A14 **Corpus companion entry:** Claim C-147 **Audit Claim ID:** deepseek-r1-reinforcement-learning-2025-06-ai-double-use-reasoning-deepseek-cla

8.18.12. Appendix B12. Chapters 06 and 07 can contrast frontier-access politics with a documented patter...

Claim statement: That Chapters 06 and 07 can contrast frontier-access politics with a documented pattern of partial technical diffusion. **Evidence status:** Documented Fact **Source support:** Appendix A14 **Corpus companion entry:** Claim C-148 **Audit Claim ID:** deepseek-r1-reinforcement-learning-2025-06-ai-double-use-reasoning-deepseek-claim-03

8.18.13. Appendix B13. By late 2024 there was already a competitor publicly claiming frontier-adjacent p...

Claim statement: That by late 2024 there was already a competitor publicly claiming frontier-adjacent performance through a technical report. **Evidence status:** Documented Fact **Source support:** Appendix A13 **Corpus companion entry:** Claim C-149 **Audit Claim ID:** deepseek-v3-technical-report-2024-claim-01

8.18.14. Appendix B14. The debate over reasoning, efficiency, and cost cannot be reduced to the marketin...

Claim statement: That the debate over reasoning, efficiency, and cost cannot be reduced to the marketing of a single Western actor. **Evidence status:** Documented Fact **Source support:** Appendix A13 **Corpus companion entry:** Claim C-150 **Audit Claim ID:** deepseek-v3-technical-report-2024-claim-02

8.18.15. Appendix B15. Narratives of absolute scarcity around advanced capabilities must be contrasted w...

Claim statement: That narratives of absolute scarcity around advanced capabilities must be contrasted with open technical reports and partial replicability. **Evidence status:** Documented Fact **Source support:** Appendix A13 **Corpus companion entry:** Claim C-151 **Audit Claim ID:** deepseek-v3-technical-report-2024-claim-03

8.18.16. Appendix B16. Serious T2 reporting documented OpenAI publicly acknowledging higher misuse risk...

Claim statement: That serious T2 reporting documented OpenAI publicly acknowledging higher misuse risk at the same moment it was releasing stronger reasoning models. **Evidence status:** Documented Fact **Source support:** Appendix A3 **Corpus companion entry:** Claim C-187 **Audit Claim ID:** ft-openai-o1-bioweapon-risk-2024-06-ai-double-use-reasoning-deepseek-claim-01

8.18.17. Appendix B17. Chapter 06 can ground the dual-use argument in contemporaneous mainstream reporti...

Claim statement: That Chapter 06 can ground the dual-use argument in contemporaneous mainstream reporting, not only in retrospective critique. **Evidence status:** Documented Fact **Source support:** Appendix A3 **Corpus companion entry:** Claim C-188 **Audit Claim ID:** ft-openai-o1-bioweapon-risk-2024-06-ai-double-use-reasoning-deepseek-claim-02

8.18.18. Appendix B18. Danger rhetoric around reasoning models accompanied commercialization rather than...

Claim statement: That danger rhetoric around reasoning models accompanied commercialization rather than halting deployment altogether. **Evidence status:** Documented Fact **Source support:** Appendix A3 **Corpus companion entry:** Claim C-189 **Audit Claim ID:** ft-openai-o1-bioweapon-risk-2024-06-ai-double-use-reasoning-deepseek-claim-03

8.18.19. Appendix B19. There is direct technical pushback against simplistic claims that frontier reason...

Claim statement: That there is direct technical pushback against simplistic claims that frontier reasoning was trivially replicated. **Evidence status:** Documented Fact **Source support:** Appendix A8 **Corpus companion entry:** Claim C-243 **Audit Claim ID:** its-not-that-simple-test-time-scaling-2025-06-ai-double-use-reasoning-deepseek-claim-01

8.18.20. Appendix B20. The paper argued simple test-time scaling often reflects scaling down through len...

Claim statement: That the paper argued simple test-time scaling often reflects scaling down through length constraints rather than true scaling up. **Evidence status:** Documented Fact **Source support:** Appendix A8 **Corpus companion entry:** Claim C-244 **Audit Claim ID:** its-not-that-simple-test-time-scaling-2025-06-ai-double-use-reasoning-deepseek-claim-02

8.18.21. Appendix B21. The paper argued o1-like and DeepSeek-R1-like systems differ from simple replicat...

Claim statement: That the paper argued o1-like and DeepSeek-R1-like systems differ from simple replications because they learn to scale up test-time compute through reinforcement learning. **Evidence status:** Documented Fact **Source support:** Appendix A8 **Corpus companion entry:** Claim C-245 **Audit Claim ID:** its-not-that-simple-test-time-scaling-2025-06-ai-double-use-reasoning-deepseek-claim-03

8.18.22. Appendix B22. As of July 3, 2026, OpenAI's official API pricing imposed a large price premium o...

Claim statement: That as of July 3, 2026, OpenAI's official API pricing imposed a large price premium on GPT-5.5 Pro relative to GPT-5.5. **Evidence status:** Documented Fact **Source support:** Appendix A10 **Corpus companion entry:** Claim C-292 **Audit Claim ID:** openai-api-pricing-2026-06-ai-double-use-reasoning-deepseek-claim-01

8.18.23. Appendix B23. OpenAI's current API record treats top-tier reasoning-capable access as materiall...

Claim statement: That OpenAI's current API record treats top-tier reasoning-capable access as materially higher-cost infrastructure rather than as a negligible increment over ordinary flagship use. **Evidence status:** Documented Fact **Source support:** Appendix A10 **Corpus companion entry:** Claim C-293 **Audit Claim ID:** openai-api-pricing-2026-06-ai-double-use-reasoning-deepseek-claim-02

8.18.24. Appendix B24. Chapter 06 can document reasoning-premium monetization through a primary pricing...

Claim statement: That Chapter 06 can document reasoning-premium monetization through a primary pricing source rather than only through media summaries. **Evidence status:** Documented Fact **Source support:** Appendix A10 **Corpus companion entry:** Claim C-294 **Audit Claim ID:** openai-api-pricing-2026-06-ai-double-use-reasoning-deepseek-claim-03

8.18.25. Appendix B25. OpenAI publicly framed reasoning as a new scaling layer based on additional train...

Claim statement: That OpenAI publicly framed reasoning as a new scaling layer based on additional train-time and test-time compute. **Evidence status:** Documented Fact **Source support:** Appendix A1 **Corpus companion entry:** Claim C-310 **Audit Claim ID:** openai-learning-to-reason-with-llms-2024-06-ai-double-use-reasoning-deepseek-claim-01

8.18.26. Appendix B26. The company explicitly marketed o1 as a major reasoning advance over GPT-4o in ma...

Claim statement: That the company explicitly marketed o1 as a major reasoning advance over GPT-4o in math, coding, and science-heavy evaluations. **Evidence status:** Documented Fact **Source support:** Appendix A1 **Corpus companion entry:** Claim C-311 **Audit Claim ID:** openai-learning-to-reason-with-llms-2024-06-ai-double-use-reasoning-deepseek-claim-02

8.18.27. Appendix B27. Chapter 06 can document reasoning as a commercial and strategic product category,...

Claim statement: That Chapter 06 can document reasoning as a commercial and strategic product category, not just a later media label. **Evidence status:** Documented Fact **Source support:** Appendix A1 **Corpus companion entry:** Claim C-312 **Audit Claim ID:** openai-learning-to-reason-with-llms-2024-06-ai-double-use-reasoning-deepseek-claim-03

8.18.28. Appendix B28. OpenAI itself paired stronger reasoning capability with formal safety, preparedne...

Claim statement: That OpenAI itself paired stronger reasoning capability with formal safety, preparedness, and misuse-risk framing. **Evidence status:** Documented Fact **Source support:** Appendix A2 **Corpus companion entry:** Claim C-316 **Audit Claim ID:** openai-o1-system-card-2024-06-ai-double-use-reasoning-deepseek-claim-01

8.18.29. Appendix B29. The company documented a training-data mix combining public web-scale material, p...

Claim statement: That the company documented a training-data mix combining public web-scale material, proprietary partnership data, and internal datasets. **Evidence status:** Documented Fact **Source support:** Appendix A2 **Corpus companion entry:** Claim C-317 **Audit Claim ID:** openai-o1-system-card-2024-06-ai-double-use-reasoning-deepseek-claim-02

8.18.30. Appendix B30. Chapter 06 can ground the coexistence of capability escalation and safety-governa...

Claim statement: That Chapter 06 can ground the coexistence of capability escalation and safety-governance packaging in a primary source. **Evidence status:** Documented Fact **Source support:** Appendix A2 **Corpus companion entry:** Claim C-318 **Audit Claim ID:** openai-o1-system-card-2024-06-ai-double-use-reasoning-deepseek-claim-03

8.18.31. Appendix B31. By November 23, 2023 there was serious mainstream reporting explicitly linking th...

Claim statement: That by November 23, 2023 there was serious mainstream reporting explicitly linking the board crisis to a letter about an AI breakthrough. **Evidence status:** Documented Fact **Source support:** Appendix A4 **Corpus companion entry:** Claim C-362 **Audit Claim ID:** reuters-qstar-board-warning-2023-13-openai-governance-qstar-talent-exodus-claim-01

8.18.32. Appendix B32. Chapter 06 can describe the Q* thread as a real reporting trail in the public rec...

Claim statement: That Chapter 06 can describe the Q* thread as a real reporting trail in the public record rather than a purely fringe rumor. **Evidence status:** Documented Fact **Source support:** Appendix A4 **Corpus companion entry:** Claim C-363 **Audit Claim ID:** reuters-qstar-board-warning-2023-13-openai-governance-qstar-talent-exodus-claim-02

8.18.33. Appendix B33. The dossier should treat Q*-related causal claims as reported and contested rather... {#ch-06-appendix-b33-the-dossier-should-treat-q-related-causal-claims-as-reported-and-contested-rathe}

Claim statement: That the dossier should treat Q*-related causal claims as reported and contested rather than as settled fact. **Evidence status:** Documented Fact **Source support:** Appendix A4 **Corpus companion entry:** Claim C-364 **Audit Claim ID:** reuters-qstar-board-warning-2023-13-openai-governance-qstar-talent-exodus-claim-03

8.18.34. Appendix B34. By mid-July 2024 Reuters had published serious reporting about an OpenAI reasonin...

Claim statement: That by mid-July 2024 Reuters had published serious reporting about an OpenAI reasoning project under the codename Strawberry. **Evidence status:** Documented Fact **Source support:** Appendix A5 **Corpus companion entry:** Claim C-365 **Audit Claim ID:** reuters-strawberry-reasoning-2024-13-openai-governance-qstar-talent-exodus-claim-01

8.18.35. Appendix B35. The dossier can document a pre-release public reporting trail for reasoning-model...

Claim statement: That the dossier can document a pre-release public reporting trail for reasoning-model development before o1 entered the market. **Evidence status:** Documented Fact **Source support:** Appendix A5 **Corpus companion entry:** Claim C-366 **Audit Claim ID:** reuters-strawberry-reasoning-2024-13-openai-governance-qstar-talent-exodus-claim-02

8.18.36. Appendix B36. Chapter 06 can connect later reasoning-model branding to an earlier, serious repo...

Claim statement: That Chapter 06 can connect later reasoning-model branding to an earlier, serious reporting record without overstating certainty about every internal detail. **Evidence status:** Documented Fact **Source support:** Appendix A5 **Corpus companion entry:** Claim C-367 **Audit Claim ID:** reuters-strawberry-reasoning-2024-13-openai-governance-qstar-talent-exodus-claim-03

8.18.37. Appendix B37. OpenAI's reasoning launch quickly triggered open replication efforts focused on t...

Claim statement: That OpenAI's reasoning launch quickly triggered open replication efforts focused on test-time scaling. **Evidence status:** Documented Fact **Source support:** Appendix A7 **Corpus companion entry:** Claim C-368 **Audit Claim ID:** s1-simple-test-time-scaling-2025-06-ai-doubt

8.18.38. Appendix B38. The paper proposed budget forcing as a simple way to control test-time compute by...

Claim statement: That the paper proposed budget forcing as a simple way to control test-time compute by truncating or extending model thinking. **Evidence status:** Documented Fact **Source support:** Appendix A7 **Corpus companion entry:** Claim C-369 **Audit Claim ID:** s1-simple-test-time-scaling-2025-06-ai-double-use-reasoning-deepseek-claim-02

8.18.39. Appendix B39. The paper reported benchmark gains substantial enough to challenge simplistic fro...

Claim statement: That the paper reported benchmark gains substantial enough to challenge simplistic frontier-exclusivity narratives. **Evidence status:** Documented Fact **Source support:** Appendix A7 **Corpus companion entry:** Claim C-370 **Audit Claim ID:** s1-simple-test-time-scaling-2025-06-ai-double-use-reasoning-deepseek-claim-03

8.18.40. Appendix B40. Test-time compute scaling is a documented technical mechanism rather than only a...

Claim statement: That test-time compute scaling is a documented technical mechanism rather than only a commercial branding story. **Evidence status:** Documented Fact **Source support:** Appendix A6 **Corpus companion entry:** Claim C-376 **Audit Claim ID:** snell-test-time-compute-2024-06-ai-double-use-reasoning-deepseek-claim-01

8.18.41. Appendix B41. The paper reported a compute-optimal strategy improving test-time compute efficie...

Claim statement: That the paper reported a compute-optimal strategy improving test-time compute efficiency by more than 4x over a best-of-N baseline. **Evidence status:** Documented Fact **Source support:** Appendix A6 **Corpus companion entry:** Claim C-377 **Audit Claim ID:** snell-test-time-compute-2024-06-ai-double-use-reasoning-deepseek-claim-02

8.18.42. Appendix B42. Under a FLOPs-matched evaluation the paper found cases where a smaller model usin...

Claim statement: That under a FLOPs-matched evaluation the paper found cases where a smaller model using test-time compute outperformed a model 14x larger. **Evidence status:** Documented Fact **Source support:** Appendix A6 **Corpus companion entry:** Claim C-378 **Audit Claim ID:** snell-test-time-compute-2024-06-ai-double-use-reasoning-deepseek-claim-03

8.18.43. Appendix B43. Whether simple budget-forcing style replications capture the same scaling mechani...

Claim statement: That whether simple budget-forcing style replications capture the same scaling mechanism as o1-like or R1-like reasoning remains contested. **Evidence status:** Disputed Fact **Source support:** Appendix A7, Appendix A8 **Corpus companion entry:** Claim C-434 **Audit Claim ID:** vector06-simple-replication-contested-claim-01

8.18.44. Appendix B44. The public record does not settle whether Q* or any specific capability breakthro...

Claim statement: That the public record does not settle whether Q* or any specific capability breakthrough directly triggered the November 2023 OpenAI board rupture. **Evidence status:** Disputed Fact **Source support:** Appendix A4 **Corpus companion entry:** Claim C-465 **Audit Claim ID:** vector13-qstar-causality-contested-claim-01

8.18.45. Appendix B45. Frontier labs are turning reasoning into a premium political-economic category by...

Claim statement: That frontier labs are turning reasoning into a premium political-economic category by coupling genuine capability gains to danger rhetoric, controlled release, and tiered strategic positioning. **Evidence status:** Hypothesis / Interpretation **Source support:** Appendix A2, Appendix A3 **Corpus companion entry:** Claim C-433 **Audit Claim ID:** vector06-reasoning-premium-control-claim-01

8.18.46. Appendix B46. Current official pricing and launch records from OpenAI and Anthropic show reason...

Claim statement: That current official pricing and launch records from OpenAI and Anthropic show reasoning capability being commercialized through tiered plans, priority access, and premium API pricing rather than released as an equal-access utility. **Evidence status:** Hypothesis / Interpretation **Source support:** Appendix A9, Appendix A10, Appendix A11, Appendix A12 **Corpus companion entry:** Claim C-435 **Audit Claim ID:** vector06-tiered-reasoning-access-synthesis-claim-01

8.18.47. Appendix B47. Chapter 06 is strongest when it argues that reasoning tiers package a real comput...

Claim statement: That Chapter 06 is strongest when it argues that reasoning tiers package a real compute-intensive technical mechanism inside a commercial and political access regime rather than treating the whole phenomenon as fake. **Evidence status:** Hypothesis / Interpretation **Source support:** Appendix A6, Appendix A2, Appendix A9, Appendix A10 **Corpus companion entry:** Claim C-436 **Audit Claim ID:** vector06-ttc-real-but-packaged-claim-01

8.18.48. Appendix B48. OpenAI's governance rupture, safety rhetoric, leadership churn, and reasoning-mod...

Claim statement: That OpenAI's governance rupture, safety rhetoric, leadership churn, and reasoning-model rollout are best read as parts of one reorganization of authority rather than as isolated episodes. **Evidence status:** Hypothesis / Interpretation **Source support:** Appendix A5 **Corpus companion entry:** Claim C-464 **Audit Claim ID:** vector13-governance-productization-synthesis

8.18.49. Appendix B49. Open reasoning papers or public model releases eliminate the politics of compute...

Claim statement: That open reasoning papers or public model releases eliminate the politics of compute concentration, distribution control, or institutional gatekeeping around frontier AI. **Evidence status:** Speculative Narrative Risk **Source support:** Appendix A13, Appendix A14 **Corpus companion entry:** Claim C-432 **Audit Claim ID:** vector06-open-publication-solves-control-overcl

8.18.50. Appendix B50. One secret breakthrough, letter, or hidden internal discovery fully explains the...

Claim statement: That one secret breakthrough, letter, or hidden internal discovery fully explains the OpenAI governance crisis, later product strategy, and subsequent talent exodus by itself. **Evidence status:** Speculative Narrative Risk **Source support:** Appendix A4, Appendix A5 **Corpus companion entry:** Claim C-466 **Audit Claim ID:** vector13-secret-trigger-overclaim-claim-01

8.19. Appendix C. Timeline Slice

8.19.1. Appendix C1. 2023-11-23: Reuters publishes a Q*-linked reporting trail tied to the OpenAI board crisis {#ch-06-appendix-c1-2023-11-23-reuters-publishes-a-q-linked-reporting-trail-tied-to-the-openai-board-crisis}

Event summary: Reuters publishes a Q*-linked reporting trail tied to the OpenAI board crisis **Event date:** 2023-11-23 **Source support:** Appendix A4 **Corpus companion entry:** Event T-033 **Audit Event ID:** event-reuters-qstar-report-2023-11-23

8.19.2. Appendix C2. 2024-07-15: Reuters reports on OpenAI's Strawberry reasoning project before o1 release

Event summary: Reuters reports on OpenAI's Strawberry reasoning project before o1 release **Event date:** 2024-07-15 **Source support:** Appendix A5 **Corpus companion entry:** Event T-050 **Audit Event ID:** event-reuters-strawberry-2024-07-15

8.19.3. Appendix C3. 2024-08-06: Researchers publish a compute-optimal test-time scaling paper for LLMs

Event summary: Researchers publish a compute-optimal test-time scaling paper for LLMs **Event date:** 2024-08-06 **Source support:** Appendix A6 **Corpus companion entry:** Event T-053 **Audit Event ID:** event-snell-test-time-compute-2024-08-06

8.19.4. Appendix C4. 2024-09-12: OpenAI launches o1-preview as a reasoning-focused model line

Event summary: OpenAI launches o1-preview as a reasoning-focused model line **Event date:** 2024-09-12 **Source support:** Appendix A1 **Corpus companion entry:** Event T-055 **Audit Event ID:** event-openai-o1-preview-2024-09-12

8.19.5. Appendix C5. 2024-09-13: Financial Times reports OpenAI acknowledging higher bioweapon misuse risk for o1

Event summary: Financial Times reports OpenAI acknowledging higher bioweapon misuse risk for o1 **Event date:** 2024-09-13 **Source support:** Appendix A3 **Corpus companion entry:** Event T-056 **Audit Event ID:** event-ft-openai-bioweapon-risk-2024-09-13

8.19.6. Appendix C6. 2024-12-05: OpenAI updates the o1 system card with preparedness and training-data details

Event summary: OpenAI updates the o1 system card with preparedness and training-data details **Event date:** 2024-12-05 **Source support:** Appendix A2 **Corpus companion entry:** Event T-060 **Audit Event ID:** event-openai-o1-system-card-2024-12-05

8.19.7. Appendix C7. 2024-12-27: DeepSeek presents the DeepSeek-V3 technical report on arXiv

Event summary: DeepSeek presents the DeepSeek-V3 technical report on arXiv **Event date:** 2024-12-27 **Source support:** Appendix A13 **Corpus companion entry:** Event T-063 **Audit Event ID:** event-deepseek-v3-2024-12-27

8.19.8. Appendix C8. 2025-01-22: DeepSeek publishes its R1 reasoning paper and open release claims

Event summary: DeepSeek publishes its R1 reasoning paper and open release claims **Event date:** 2025-01-22 **Source support:** Appendix A14 **Corpus companion entry:** Event T-065 **Audit Event ID:** event-deepseek-r1-2025-01-22

8.19.9. Appendix C9. 2025-01-31: Researchers release s1 as an open attempt to replicate test-time scaling

Event summary: Researchers release s1 as an open attempt to replicate test-time scaling **Event date:** 2025-01-31 **Source support:** Appendix A7 **Corpus companion entry:** Event T-067 **Audit Event ID:** event-s1-simple-test-time-scaling-2025-01-31

8.19.10. Appendix C10. 2025-05-22: Anthropic launches Claude 4 with extended-thinking commercial availability

Event summary: Anthropic launches Claude 4 with extended-thinking commercial availability **Event date:** 2025-05-22 **Source support:** Appendix A12 **Corpus companion entry:** Event T-076 **Audit Event ID:** event-anthropic-claude-4-launch-2025-05-22

8.19.11. Appendix C11. 2025-07-19: A follow-up paper challenges simplistic readings of simple test-time scaling

Event summary: A follow-up paper challenges simplistic readings of simple test-time scaling **Event date:** 2025-07-19 **Source support:** Appendix A8 **Corpus companion entry:** Event T-082 **Audit Event ID:** event-its-not-that-simple-test-time-scaling-2025-07-19

8.20. Appendix D. Relevant Entities

8.20.1. 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 A11, Appendix A12 **Corpus companion entry:** Entity E-008 **Audit Entity ID:** org-anthropic

8.20.2. Appendix D2. DeepSeek

Entity type: company **Role in this paper:** Actor relevant to 06_ai_double_use_reasoning_deepseek, 07_export_controls_compute_defense_access. **Relevant sources in this paper:** Appendix A13, Appendix A14 **Corpus companion entry:** Entity E-013 **Audit Entity ID:** org-deepseek

8.20.3. 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_le and others. **Relevant sources in this paper:** Appendix A1, Appendix A2, Appendix A3, Appendix A4, Appendix A5, Appendix A9, Appendix A10 **Corpus companion entry:** Entity E-027 **Audit Entity ID:** org-openai

8.20.4. Appendix D4. Test-time compute scaling

Entity type: technical_mechanism **Role in this paper:** Actor relevant to 06_ai_double_use_reasoning_deepseek **Relevant sources in this paper:** Appendix A6, Appendix A7, Appendix A8 **Corpus companion entry:** Entity E-004 **Audit Entity ID:** concept-test-time-compute-scaling

8.21. Appendix E. Evidence Boundaries

8.21.1. Appendix E1. Disputed Matters

- Appendix B43: That whether simple budget-forcing style replications capture the same scaling mechanism as o1-like or R1-like reasoning remains contested.
- Appendix B44: That the public record does not settle whether Q* or any specific capability breakthrough directly triggered the November 2023 OpenAI board rupture.

8.21.2. Appendix E2. Interpretive Boundaries

- Appendix B45: That frontier labs are turning reasoning into a premium political-economic category by coupling genuine capability gains to danger rhetoric, controlled release, and tiered strategic positioning.
- Appendix B46: That current official pricing and launch records from OpenAI and Anthropic show reasoning capability being commercialized through tiered plans, priority access, and premium API pricing rather than released as an equal-access utility.

- Appendix B47: That Chapter 06 is strongest when it argues that reasoning tiers package a real compute-intensive technical mechanism inside a commercial and political access regime rather than treating the whole phenomenon as fake.
- Appendix B48: That OpenAI's governance rupture, safety rhetoric, leadership churn, and reasoning-model rollout are best read as parts of one reorganization of authority rather than as isolated episodes.

8.21.3. Appendix E3. Speculative Narrative Risks

- Appendix B49: That open reasoning papers or public model releases eliminate the politics of compute concentration, distribution control, or institutional gatekeeping around frontier AI.
- Appendix B50: That one secret breakthrough, letter, or hidden internal discovery fully explains the OpenAI governance crisis, later product strategy, and subsequent talent exodus by itself.

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

- Appendix A12. Companion note: Note N-054. Risk flag: low. Note focus: That on May 22, 2025 Anthropic launched reasoning-oriented hybrid frontier models together with immediate commercial distribution across subscription plans, API access, and major cloud platforms.
- Appendix A11. Companion note: Note N-055. Risk flag: low. Note focus: That as of July 3, 2026, Anthropic's official current pricing structure tied higher usage, priority access, and stronger feature access to premium paid plans.
- Appendix A9. Companion note: Note N-056. Risk flag: low. Note focus: That as of capture on July 3, 2026, OpenAI's official ChatGPT plan structure reserved the highest-end reasoning access for upper paid tiers rather than offering it uniformly across all users.
- Appendix A14. Companion note: Note N-057. Risk flag: medium. Note focus: That by January 2025 a non-U.S. lab was publicly claiming frontier-adjacent reasoning performance in a primary technical paper.
- Appendix A13. Companion note: Note N-058. Risk flag: medium. Note focus: That by late 2024 there was already a competitor publicly claiming frontier-adjacent performance through a technical report.
- Appendix A3. Companion note: Note N-059. Risk flag: medium. Note focus: That serious T2 reporting documented OpenAI publicly acknowledging higher misuse risk at the same moment it was releasing stronger reasoning models.
- Appendix A8. Companion note: Note N-060. Risk flag: low. Note focus: That there is direct technical pushback against simplistic claims that frontier reasoning was trivially replicated.
- Appendix A10. Companion note: Note N-061. Risk flag: low. Note focus: That as of July 3, 2026, OpenAI's official API pricing imposed a large price premium on GPT-5.5 Pro relative to GPT-5.5.
- Appendix A1. Companion note: Note N-062. Risk flag: medium. Note focus: That OpenAI publicly framed reasoning as a new scaling layer based on additional train-time and test-time compute.
- Appendix A2. Companion note: Note N-063. Risk flag: medium. Note focus: That OpenAI itself paired stronger reasoning capability with formal safety, preparedness, and misuse-risk framing.

- Appendix A7. Companion note: Note N-064. Risk flag: medium. Note focus: That OpenAI's reasoning launch quickly triggered open replication efforts focused on test-time scaling.
- Appendix A6. Companion note: Note N-065. Risk flag: low. Note focus: That test-time compute scaling is a documented technical mechanism rather than only a commercial branding story.
- Appendix A4. Companion note: Note N-141. Risk flag: medium. Note focus: That by November 23, 2023 there was serious mainstream reporting explicitly linking the board crisis to a letter about an AI breakthrough.
- Appendix A5. Companion note: Note N-142. Risk flag: medium. Note focus: That by mid-July 2024 Reuters had published serious reporting about an OpenAI reasoning project under the codename **Strawberry**.

9. Chapter 07. Imperial control over access

9.1. Abstract

The AI system is not simply a market of competing products. The United States has already moved from broad leadership rhetoric into formal doctrine, consultation processes, and implementation guidance. If frontier AI is governed through chokepoints, then intelligence itself becomes a stratified geopolitical infrastructure.

9.2. Keywords

Imperial control over access; surveillance capitalism; AI imperialism; government capture; evidence-based dossier

9.3. Research Question

How does imperial control over access function within the dossier's larger argument, and what does the documentary record allow this chapter to establish?

9.4. Scope and Framing Note

This paper examines imperial control over access 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 establishes chokepoints, strategic access, and state-company interfaces that justify anti-capture obligations. Adjacent reinforcement appears in Chapter 06, Chapter 08, Chapter 10, but those cross-references are supportive rather than required for basic comprehension.

9.5. Central Thesis

The AI system is not simply a market of competing products. It is also an architecture of chokepoints. Export controls, cloud dependence, chip access, and defense-state adjacency increasingly define who can build, scale, or even reach frontier intelligence. The question is no longer only who invents. It is who gets permission.⁷⁶

9.6. Evidence and Method Note

This paper is derived from the chapter corpus for Chapter 07 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.

⁷⁶Sources: The White House, Removing Barriers to American Leaders... (2025-01-23; see Appendix A1); Bureau of Industry and Security, Guidance Regarding Enforcement of Lic... (2026-05-31; see Appendix A2); Amazon, Amazon and Anthropic Announce Strateg... (2023-09-25; see Appendix A3)

9.7. Introduction

This paper asks: How does imperial control over access function within the dossier’s larger argument, and what does the documentary record allow this chapter to establish? Its working answer is clear from the start: The AI system is not simply a market of competing products. It is also an architecture of chokepoints. Export controls, cloud dependence, chip access, and defense-state adjacency increasingly define who can build, scale, or even reach frontier intelligence. The question is no longer only who invents. It is who gets permission.⁷⁷ 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.

9.8. Historical or Institutional Context

The United States has already moved from broad leadership rhetoric into formal doctrine, consultation processes, and implementation guidance. Executive action, federal requests for information, and corporate positioning now form a visible policy sequence rather than a vague background condition.⁷⁸

9.9. Main Analysis

BIS guidance gives the advanced-computing export regime an operational enforcement layer. The politics of AI access is therefore not merely metaphorical. It is already administrative, technical, and executable.⁷⁹

Hyperscaler dependence is also documented. At least one major frontier lab publicly tied the majority of its workloads and future development to a single cloud giant, which means frontier capability is already routed through a small number of infrastructural landlords.⁸⁰

The formal government-facing layer is now public as well. OpenAI’s **OpenAI for Government**, Anthropic’s **Claude Gov**, and the Palantir-AWS-Anthropic classified distribution lane show that frontier access is organized not only through markets but through accredited, restricted, and government-priority channels.⁸¹

The personnel interface is explicit too. The Army’s **Detachment 201** demonstrates that the boundary between senior frontier-tech leadership and military institutions has become porous in public view. That does not prove improper command. It does prove public institutional integration.⁸²

⁷⁷Sources: The White House, Removing Barriers to American Leaders... (2025-01-23; see Appendix A1); Bureau of Industry and Security, Guidance Regarding Enforcement of Lic... (2026-05-31; see Appendix A2); Amazon, Amazon and Anthropic Announce Strateg... (2023-09-25; see Appendix A3)

⁷⁸Sources: The White House, Removing Barriers to American Leaders... (2025-01-23; see Appendix A1); Federal Register, Request for Information on the Develo... (2025-02-06; see Appendix A4); OpenAI, OpenAI’s proposals for the U.S. AI Ac... (2025-03-13; see Appendix A5)

⁷⁹Sources: Bureau of Industry and Security, Guidance Regarding Enforcement of Lic... (2026-05-31; see Appendix A2)

⁸⁰Sources: Amazon, Amazon and Anthropic Announce Strateg... (2023-09-25; see Appendix A3)

⁸¹Sources: OpenAI, Introducing OpenAI for Government (2025-06-16; see Appendix A6); Anthropic, Claude Gov models for U.S. national s... (2025-06-06; see Appendix A7); Palantir Technologies, Anthropic and Palantir Partner to Bri... (2024-11-07; see Appendix A8)

⁸²Sources: U.S. Army Public Affairs, Army Launches Detachment 201: Executi... (2025-06-13; see Appendix A9)

In-Q-Tel, Palantir, and the Thiel-linked network provide a serious documentary basis for talking about venture-state-security adjacency without drifting into mythology. The structure is the point: chips, clouds, accredited environments, contracts, and strategic interfaces remain highly concentrated even if some model knowledge diffuses.⁸³

Interpretation: this is where **Limit Government Capture by AI Companies** begins to look structurally necessary. Once public support, export-control protection, privileged state access, and defense channels become material to frontier advantage, the resulting intelligence cannot be defended as a purely ordinary private market outcome.⁸⁴

9.10. Position Within the Series

This paper sits inside a coordinated 11-paper architecture. Its nearest companions are Chapter 06 (AI as a dual-use rhetorical weapon), Chapter 08 (Power network, legitimacy, and institutional capture), Chapter 10 (What to do). In that series logic, Chapter 07 establishes chokepoints, strategic access, and state-company interfaces that justify anti-capture obligations.

9.11. Counterarguments

This chapter must not confuse chokepoints with total closure. The evidence supports concentrated leverage, privileged lanes, and formal state-company interfaces, not omnipotent control over all capability diffusion.⁸⁵

9.12. Limits of the Evidence

This chapter must not confuse chokepoints with total closure. The evidence supports concentrated leverage, privileged lanes, and formal state-company interfaces, not omnipotent control over all capability diffusion.⁸⁶

9.13. Reform Relevance

If frontier AI is governed through chokepoints, then intelligence itself becomes a stratified geopolitical infrastructure. Access is no longer merely a subscription question. It becomes a question of empire, supply chains, protected lanes, and permission.⁸⁷

⁸³Sources: In-Q-Tel, In-Q-Tel (n.d.; see Appendix A10); Palantir, Gotham (n.d.; see Appendix A11); SEC, Meta Platforms, Inc. 2022 Definitive... (2022-04-08; see Appendix A12); arXiv, DeepSeek-R1: Incentivizing Reasoning... (2025-01-22; see Appendix A13)

⁸⁴Sources: OpenAI, Introducing OpenAI for Government (2025-06-16; see Appendix A6); Anthropic, Claude Gov models for U.S. national s... (2025-06-06; see Appendix A7); Bureau of Industry and Security, Guidance Regarding Enforcement of Lic... (2026-05-31; see Appendix A2)

⁸⁵Sources: arXiv, DeepSeek-R1: Incentivizing Reasoning... (2025-01-22; see Appendix A13); Amazon, Amazon and Anthropic Announce Strateg... (2023-09-25; see Appendix A3)

⁸⁶Sources: arXiv, DeepSeek-R1: Incentivizing Reasoning... (2025-01-22; see Appendix A13); Amazon, Amazon and Anthropic Announce Strateg... (2023-09-25; see Appendix A3)

⁸⁷Sources: The White House, Removing Barriers to American Leaders... (2025-01-23; see Appendix A1); Bureau of Industry and Security, Guidance Regarding Enforcement of Lic... (2026-05-31; see Appendix A2); OpenAI, Introducing OpenAI for Government (2025-06-16; see Appendix A6)

9.14. Conclusion

If frontier AI is governed through chokepoints, then intelligence itself becomes a stratified geopolitical infrastructure. But chokepoints alone do not explain public acceptance. In the architecture of this series, Chapter 07 therefore functions as a self-contained argument while also advancing the cumulative path toward the dossier's three reforms.

9.15. 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 07, those boundaries remain visible in the prose, in the bibliography, and in the appendix sections that summarize claims and caution points.

9.16. References / Bibliography

1. The White House. Removing Barriers to American Leadership in Artificial Intelligence. 2025-01-23. T1 executive_order. Paper appendix: Appendix A1. Corpus companion entry: Source S-130. Audit ID: **whitehouse-eo-14179-2025**. Relevance: That AI dominance was elevated to explicit state doctrine in 2025.
2. Bureau of Industry and Security. Guidance Regarding Enforcement of License Requirements for Advanced Computing Items for Entities Headquartered in Country Group D:5 and Macau. 2026-05-31. T1 guidance_pdf. Paper appendix: Appendix A2. Corpus companion entry: Source S-037. Audit ID: **bis-guidance-advanced-computing-d5-macau-2026**. Relevance: That there is a primary U.S. government record directly describing continued license requirements for advanced computing exports linked to Country Group D:5 and Macau.
3. Amazon. Amazon and Anthropic Announce Strategic Collaboration to Advance Generative AI. 2023-09-25. T1 press_release. Paper appendix: Appendix A3. Corpus companion entry: Source S-002. Audit ID: **amazon-anthropic-strategic-collaboration-2023**. Relevance: That at least one major frontier lab publicly committed the majority of its workloads and future model development to a hyperscaler cloud provider.
4. Federal Register. Request for Information on the Development of an Artificial Intelligence (AI) Action Plan. 2025-02-06. T1 federal_register_notice. Paper appendix: Appendix A4. Corpus companion entry: Source S-054. Audit ID: **fr-ai-action-plan-rfi-2025**. Relevance: That the state created an official consultation pathway for shaping the 2025 AI doctrine.
5. OpenAI. OpenAI's proposals for the U.S. AI Action Plan. 2025-03-13. T1 policy_submission. Paper appendix: Appendix A5. Corpus companion entry: Source S-096. Audit ID: **openai-us-ai-action-plan-2025**. Relevance: That OpenAI presents itself as a policy actor seeking to co-design national AI strategy.
6. OpenAI. Introducing OpenAI for Government. 2025-06-16. T1 release. Paper appendix: Appendix A6. 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...

7. Anthropic. Claude Gov models for U.S. national security customers. 2025-06-06. T1 release. Paper appendix: Appendix A7. 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.
8. 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 A8. 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.
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 13, 2025 the U.S. Army had created a formal structure to embed senior commercial-tech executives in a military innovation role.
10. In-Q-Tel. In-Q-Tel. n.d.. T1 organization_homepage. Paper appendix: Appendix A10. Corpus companion entry: Source S-074. Audit ID: **iqt-homepage-2026**. Relevance: That there is a formal institutional channel connecting strategic technology development with intelligence and defense needs.
11. Palantir. Gotham. n.d.. T1 product_page. Paper appendix: Appendix A11. Corpus companion entry: Source S-099. Audit ID: **palantir-gotham-2026**. Relevance: That Palantir openly markets Gotham to defense agencies and the intelligence community.
12. SEC. Meta Platforms, Inc. 2022 Definitive Proxy Statement. 2022-04-08. T1 proxy_statement. Paper appendix: Appendix A12. Corpus companion entry: Source S-080. Audit ID: **meta-proxy-2022-thiel-retirement**. Relevance: That Peter Thiel's formal role at Facebook / Meta lasted well beyond the early startup period and remained active into 2022.
13. arXiv. DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning. 2025-01-22. T1 paper. Paper appendix: Appendix A13. Corpus companion entry: Source S-048. Audit ID: **deepseek-r1-reinforcement-learning-2025**. Relevance: That by January 2025 a non-U.S. lab was publicly claiming frontier-adjacent reasoning performance in a primary technical paper.

9.17. Appendix A. Source Register for This Paper

9.17.1. Appendix A1. Removing Barriers to American Leadership in Artificial Intelligence

Source citation: Removing Barriers to American Leadership in Artificial Intelligence **Institution** / **author:** The White House **Date:** 2025-01-23 **Tier** / **type:** T1 / executive_order **Why it matters in this paper:** That AI dominance was elevated to explicit state doctrine in 2025. **Corpus companion entry:** Source S-130 **Audit ID:** whitehouse-eo-14179-2025

9.17.2. Appendix A2. Guidance Regarding Enforcement of License Requirements for Advanced Computing Items for Entities Headquartered in Country Group D:5 and Macau

Source citation: Guidance Regarding Enforcement of License Requirements for Advanced Computing Items for Entities Headquartered in Country Group D:5 and Macau **Institution / author:** Bureau of Industry and Security **Date:** 2026-05-31 **Tier / type:** T1 / guidance_pdf **Why it matters in this paper:** That there is a primary U.S. government record directly describing continued license requirements for advanced computing exports linked to Country Group D:5 and Macau. **Corpus companion entry:** Source S-037 **Audit ID:** bis-guidance-advanced-computing-d5-macau-2026

9.17.3. Appendix A3. Amazon and Anthropic Announce Strategic Collaboration to Advance Generative AI

Source citation: Amazon and Anthropic Announce Strategic Collaboration to Advance Generative AI **Institution / author:** Amazon **Date:** 2023-09-25 **Tier / type:** T1 / press_release **Why it matters in this paper:** That at least one major frontier lab publicly committed the majority of its workloads and future model development to a hyperscaler cloud provider. **Corpus companion entry:** Source S-002 **Audit ID:** amazon-anthropic-strategic-collaboration-2023

9.17.4. Appendix A4. Request for Information on the Development of an Artificial Intelligence (AI) Action Plan

Source citation: Request for Information on the Development of an Artificial Intelligence (AI) Action Plan **Institution / author:** Federal Register **Date:** 2025-02-06 **Tier / type:** T1 / federal_register_notice **Why it matters in this paper:** That the state created an official consultation pathway for shaping the 2025 AI doctrine. **Corpus companion entry:** Source S-054 **Audit ID:** fr-ai-action-plan-rfi-2025

9.17.5. Appendix A5. 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 OpenAI presents itself as a policy actor seeking to co-design national AI strategy. **Corpus companion entry:** Source S-096 **Audit ID:** openai-us-ai-action-plan-2025

9.17.6. Appendix A6. 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

9.17.7. Appendix A7. 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

9.17.8. Appendix A8. 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

9.17.9. 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 13, 2025 the U.S. Army had created a formal structure to embed senior commercial-tech executives in a military innovation role. **Corpus companion entry:** Source S-026 **Audit ID:** army-detachment-201-2025

9.17.10. Appendix A10. In-Q-Tel

Source citation: In-Q-Tel **Institution / author:** In-Q-Tel **Date:** n.d. **Tier / type:** T1 / organization_homepage **Why it matters in this paper:** That there is a formal institutional channel connecting strategic technology development with intelligence and defense needs. **Corpus companion entry:** Source S-074 **Audit ID:** iqt-homepage-2026

9.17.11. Appendix A11. Gotham

Source citation: Gotham **Institution / author:** Palantir **Date:** n.d. **Tier / type:** T1 / product_page **Why it matters in this paper:** That Palantir openly markets Gotham to defense agencies and the intelligence community. **Corpus companion entry:** Source S-099 **Audit ID:** palantir-gotham-2026

9.17.12. Appendix A12. Meta Platforms, Inc. 2022 Definitive Proxy Statement

Source citation: Meta Platforms, Inc. 2022 Definitive Proxy Statement **Institution / author:** SEC **Date:** 2022-04-08 **Tier / type:** T1 / proxy_statement **Why it matters in this paper:** That Peter Thiel's formal role at Facebook / Meta lasted well beyond the early startup

period and remained active into 2022. **Corpus companion entry:** Source S-080 **Audit ID:** meta-proxy-2022-thiel-retirement

9.17.13. Appendix A13. DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning

Source citation: DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning **Institution / author:** arXiv **Date:** 2025-01-22 **Tier / type:** T1 / paper **Why it matters in this paper:** That by January 2025 a non-U.S. lab was publicly claiming frontier-adjacent reasoning performance in a primary technical paper. **Corpus companion entry:** Source S-048 **Audit ID:** deepseek-r1-reinforcement-learning-2025

9.18. Appendix B. Claims Used in This Paper

9.18.1. Appendix B1. At least one major frontier lab publicly committed the majority of its workloads...

Claim statement: That at least one major frontier lab publicly committed the majority of its workloads and future model development to a hyperscaler cloud provider. **Evidence status:** Documented Fact **Source support:** Appendix A3 **Corpus companion entry:** Claim C-004 **Audit Claim ID:** amazon-anthropic-strategic-collaboration-2023-07-export-controls-compute-defense-access-claim-01

9.18.2. Appendix B2. Chapter 07 can document compute dependence and cloud concentration with a primary...

Claim statement: That Chapter 07 can document compute dependence and cloud concentration with a primary source rather than only with abstract market language. **Evidence status:** Documented Fact **Source support:** Appendix A3 **Corpus companion entry:** Claim C-005 **Audit Claim ID:** amazon-anthropic-strategic-collaboration-2023-07-export-controls-compute-defense-access-claim-01

9.18.3. Appendix B3. Chip access, cloud infrastructure, model distribution, and capital investment were...

Claim statement: That chip access, cloud infrastructure, model distribution, and capital investment were being bundled together in a single strategic relationship. **Evidence status:** Documented Fact **Source support:** Appendix A3 **Corpus companion entry:** Claim C-006 **Audit Claim ID:** amazon-anthropic-strategic-collaboration-2023-07-export-controls-compute-defense-access-claim-01

9.18.4. 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 A7 **Corpus companion entry:** Claim C-010 **Audit Claim ID:** anthropic-claude-gov-2025-07-export-controls-compute-defense-access-claim-01

9.18.5. 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 **Source support:** Appendix A7 **Corpus companion entry:** Claim C-011 **Audit Claim ID:** anthropic-claude-gov-2025-07-export-controls-compute-defense-access-claim-02

9.18.6. 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 A7 **Corpus companion entry:** Claim C-012 **Audit Claim ID:** anthropic-claude-gov-2025-07-export-controls-compute-defense-access-claim-03

9.18.7. Appendix B7. By June 13, 2025 the U.S. Army had created a formal structure to embed senior com...

Claim statement: That by June 13, 2025 the U.S. Army had created a formal structure to embed senior commercial-tech executives in a military innovation role. **Evidence status:** Documented Fact **Source support:** Appendix A9 **Corpus companion entry:** Claim C-078 **Audit Claim ID:** army-detachment-201-2025-07-export-controls-compute-defense-access-claim-01

9.18.8. Appendix B8. The initial cohort explicitly included executives from Palantir and OpenAI, not o...

Claim statement: That the initial cohort explicitly included executives from Palantir and OpenAI, not only traditional defense contractors. **Evidence status:** Documented Fact **Source support:** Appendix A9 **Corpus companion entry:** Claim C-079 **Audit Claim ID:** army-detachment-201-2025-07-export-controls-compute-defense-access-claim-02

9.18.9. Appendix B9. Chapter 07 and Chapter 08 can document a direct institutional bridge between fron...

Claim statement: That Chapter 07 and Chapter 08 can document a direct institutional bridge between frontier/commercial AI leadership and military structures without relying on rumor or guilt by association. **Evidence status:** Documented Fact **Source support:** Appendix A9 **Corpus companion entry:** Claim C-080 **Audit Claim ID:** army-detachment-201-2025-07-export-controls-compute

9.18.10. Appendix B10. There is a primary U.S. government record directly describing continued license r...

Claim statement: That there is a primary U.S. government record directly describing continued license requirements for advanced computing exports linked to Country Group D:5 and Macau. **Evidence status:** Documented Fact **Source support:** Appendix A2 **Corpus companion entry:** Claim C-111 **Audit Claim ID:** bis-guidance-advanced-computing-d5-macau-2026-07-export-controls-com

9.18.11. Appendix B11. Chapter 07 can document a concrete implementation layer beneath abstract talk of...

Claim statement: That Chapter 07 can document a concrete implementation layer beneath abstract talk of export controls and access politics. **Evidence status:** Documented Fact

Source support: Appendix A2 **Corpus companion entry:** Claim C-112 **Audit Claim ID:**

bis-guidance-advanced-computing-d5-macau-2026-07-export-controls-compute-defense-access-claim-0

9.18.12. Appendix B12. The AI Diffusion Rule and the earlier November 2023 controls should be treated as...

Claim statement: That the AI Diffusion Rule and the earlier November 2023 controls should be treated as connected parts of a continuing export-control regime rather than isolated headlines. **Evidence status:** Documented Fact

Source support: Appendix A2 **Corpus companion entry:**

Claim C-113 **Audit Claim ID:** bis-guidance-advanced-computing-d5-macau-2026-07-export-controls-com

9.18.13. Appendix B13. By January 2025 a non-U.S. lab was publicly claiming frontier-adjacent reasoning...

Claim statement: That by January 2025 a non-U.S. lab was publicly claiming frontier-adjacent reasoning performance in a primary technical paper. **Evidence status:** Documented Fact

Source support: Appendix A13 **Corpus companion entry:** Claim C-146 **Audit Claim ID:**

deepseek-r1-reinforcement-learning-2025-06-ai-double-use-reasoning-deepseek-claim-01

9.18.14. Appendix B14. Reasoning capability was not confined entirely to closed U.S. product stacks and...

Claim statement: That reasoning capability was not confined entirely to closed U.S. product stacks and was partly diffusing through open publication and model release. **Evidence status:** Documented Fact

Source support: Appendix A13 **Corpus companion entry:** Claim C-147

Audit Claim ID: deepseek-r1-reinforcement-learning-2025-06-ai-double-use-reasoning-deepseek-cla

9.18.15. Appendix B15. Chapters 06 and 07 can contrast frontier-access politics with a documented patter...

Claim statement: That Chapters 06 and 07 can contrast frontier-access politics with a documented pattern of partial technical diffusion. **Evidence status:** Documented Fact **Source**

support: Appendix A13 **Corpus companion entry:** Claim C-148 **Audit Claim ID:**

deepseek-r1-reinforcement-learning-2025-06-ai-double-use-reasoning-deepseek-claim-03

9.18.16. Appendix B16. The state created an official consultation pathway for shaping the 2025 AI doctrine

Claim statement: That the state created an official consultation pathway for shaping the 2025 AI doctrine. **Evidence status:** Documented Fact **Source support:** Appendix A4 **Corpus compan-**

ion entry: Claim C-167 **Audit Claim ID:** fr-ai-action-plan-rfi-2025-08-power-networks-legitimacy-ca

9.18.17. Appendix B17. Legitimacy in AI was being organized through formal administrative processes, not...

Claim statement: That legitimacy in AI was being organized through formal administrative processes, not only informal lobbying. **Evidence status:** Documented Fact **Source support:** Appendix A4 **Corpus companion entry:** Claim C-168 **Audit Claim ID:** fr-ai-action-plan-rfi-2025-08-power-networks-legitimacy-capture-claim-02

9.18.18. Appendix B18. Chapter 8 can map institutional capture or influence as a documented procedural f...

Claim statement: That Chapter 8 can map institutional capture or influence as a documented procedural field. **Evidence status:** Documented Fact **Source support:** Appendix A4 **Corpus companion entry:** Claim C-169 **Audit Claim ID:** fr-ai-action-plan-rfi-2025-08-power-networks-legitimacy

9.18.19. Appendix B19. The language of AI dominance did not remain only in a political order, but moved...

Claim statement: That the language of AI dominance did not remain only in a political order, but moved into an official federal consultation process. **Evidence status:** Documented Fact **Source support:** Appendix A4 **Corpus companion entry:** Claim C-170 **Audit Claim ID:** fr-ai-action-plan-rfi-2025-claim-01

9.18.20. Appendix B20. OSTP and federal agencies activated a formal channel to gather input from industr...

Claim statement: That OSTP and federal agencies activated a formal channel to gather input from industry, academia, and other actors on U.S. AI strategy. **Evidence status:** Documented Fact **Source support:** Appendix A4 **Corpus companion entry:** Claim C-171 **Audit Claim ID:** fr-ai-action-plan-rfi-2025-claim-02

9.18.21. Appendix B21. The dossier can document a clear institutional sequence between the presidential...

Claim statement: That the dossier can document a clear institutional sequence between the presidential order, the RFI, and later corporate policy positioning. **Evidence status:** Documented Fact **Source support:** Appendix A4 **Corpus companion entry:** Claim C-172 **Audit Claim ID:** fr-ai-action-plan-rfi-2025-claim-03

9.18.22. Appendix B22. There is a formal institutional channel connecting strategic technology developme...

Claim statement: That there is a formal institutional channel connecting strategic technology development with intelligence and defense needs. **Evidence status:** Documented Fact **Source support:** Appendix A10 **Corpus companion entry:** Claim C-234 **Audit Claim ID:** iqt-homepage-2026-08-power-networks-legitimacy-capture-claim-01

9.18.23. Appendix B23. The adjacency between frontier innovation and the state apparatus can be document...

Claim statement: That the adjacency between frontier innovation and the state apparatus can be documented from primary sources. **Evidence status:** Documented Fact **Source support:** Appendix A10 **Corpus companion entry:** Claim C-235 **Audit Claim ID:** iqt-homepage-2026-08-power-networks-legitimacy-capture-claim-02

9.18.24. Appendix B24. Chapter 8 can discuss legitimacy networks without relying on conspiratorial infer...

Claim statement: That Chapter 8 can discuss legitimacy networks without relying on conspiratorial inference. **Evidence status:** Documented Fact **Source support:** Appendix A10 **Corpus companion entry:** Claim C-236 **Audit Claim ID:** iqt-homepage-2026-08-power-networks-legitimacy-capt

9.18.25. Appendix B25. There is an organization officially dedicated to connecting technological innovat...

Claim statement: That there is an organization officially dedicated to connecting technological innovation with intelligence, defense, and national-security needs. **Evidence status:** Documented Fact **Source support:** Appendix A10 **Corpus companion entry:** Claim C-237 **Audit Claim ID:** iqt-homepage-2026-claim-01

9.18.26. Appendix B26. Adjacency between venture capital, frontier tech, and the state apparatus is not...

Claim statement: That adjacency between venture capital, frontier tech, and the state apparatus is not speculation: in the case of IQT, it is expressed by the organization itself. **Evidence status:** Documented Fact **Source support:** Appendix A10 **Corpus companion entry:** Claim C-238 **Audit Claim ID:** iqt-homepage-2026-claim-02

9.18.27. Appendix B27. The dossier can discuss a technology-transfer infrastructure toward the state wit...

Claim statement: That the dossier can discuss a technology-transfer infrastructure toward the state without depending on conspiratorial sources. **Evidence status:** Documented Fact **Source support:** Appendix A10 **Corpus companion entry:** Claim C-239 **Audit Claim ID:** iqt-homepage-2026-claim-03

9.18.28. Appendix B28. Peter Thiel's formal role at Facebook / Meta lasted well beyond the early startup...

Claim statement: That Peter Thiel's formal role at Facebook / Meta lasted well beyond the early startup period and remained active into 2022. **Evidence status:** Documented Fact **Source support:** Appendix A12 **Corpus companion entry:** Claim C-262 **Audit Claim ID:** meta-proxy-2022-thiel-retirement-11-facebook-thiel-intelligence-adjacency-claim-01

9.18.29. Appendix B29. The company itself presented Thiel’s board tenure as spanning from 2005 to the en...

Claim statement: That the company itself presented Thiel’s board tenure as spanning from 2005 to the end of his term at the 2022 annual meeting. **Evidence status:** Documented Fact **Source support:** Appendix A12 **Corpus companion entry:** Claim C-263 **Audit Claim ID:** meta-proxy-2022-thiel-retirement-11-facebook-thiel-intelligence-adjacency-claim-02

9.18.30. Appendix B30. Thiel still held committee-level governance authority at Meta immediately before...

Claim statement: That Thiel still held committee-level governance authority at Meta immediately before his departure from the board. **Evidence status:** Documented Fact **Source support:** Appendix A12 **Corpus companion entry:** Claim C-264 **Audit Claim ID:** meta-proxy-2022-thiel-retirement-11-facebook-thiel-intelligence-adjacency-claim-03

9.18.31. Appendix B31. Continuity between the Facebook surveillance era and the Meta rebrand era include...

Claim statement: That continuity between the Facebook surveillance era and the Meta rebrand era included continuity in elite board membership. **Evidence status:** Documented Fact **Source support:** Appendix A12 **Corpus companion entry:** Claim C-265 **Audit Claim ID:** meta-proxy-2022-thiel-retirement-11-facebook-thiel-intelligence-adjacency-claim-04

9.18.32. Appendix B32. 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 A6 **Corpus companion entry:** Claim C-298 **Audit Claim ID:** openai-for-government-2025-07-export-controls-compu

9.18.33. Appendix B33. 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 A6 **Corpus companion entry:** Claim C-299 **Audit Claim ID:** openai-for-government-2025-07-export-controls-compute-defense-access

9.18.34. Appendix B34. 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 A6 **Corpus companion entry:** Claim C-300 **Audit Claim ID:** openai-for-government-2025-07-export-controls-compute-defense-access-claim-03

9.18.35. Appendix B35. OpenAI presents itself as a policy actor seeking to co-design national AI strategy

Claim statement: That OpenAI presents itself as a policy actor seeking to co-design national AI strategy. **Evidence status:** Documented Fact **Source support:** Appendix A5 **Corpus companion entry:** Claim C-319 **Audit Claim ID:** openai-us-ai-action-plan-2025-08-power-networks-legitimacy

9.18.36. Appendix B36. The company publicly links its commercial and technical agenda to national-securi...

Claim statement: That the company publicly links its commercial and technical agenda to national-security and industrial-policy language. **Evidence status:** Documented Fact **Source support:** Appendix A5 **Corpus companion entry:** Claim C-320 **Audit Claim ID:** openai-us-ai-action-plan-2025-08-power-networks-legitimacy-capture-claim-02

9.18.37. Appendix B37. Chapter 8 can treat lobbying and doctrinal positioning as part of the legitimacy...

Claim statement: That Chapter 8 can treat lobbying and doctrinal positioning as part of the legitimacy architecture of frontier firms. **Evidence status:** Documented Fact **Source support:** Appendix A5 **Corpus companion entry:** Claim C-321 **Audit Claim ID:** openai-us-ai-action-plan-2025-08-power-networks-legitimacy-capture-claim-03

9.18.38. Appendix B38. OpenAI officially defended a public-policy agenda combining global commercial exp...

Claim statement: That OpenAI officially defended a public-policy agenda combining global commercial expansion, export controls, and preservation of access to copyrighted material for training. **Evidence status:** Documented Fact **Source support:** Appendix A5 **Corpus companion entry:** Claim C-322 **Audit Claim ID:** openai-us-ai-action-plan-2025-claim-01

9.18.39. Appendix B39. The company presents itself not only as a model provider, but as an actor seeking...

Claim statement: That the company presents itself not only as a model provider, but as an actor seeking to co-design U.S. national AI strategy. **Evidence status:** Documented Fact **Source support:** Appendix A5 **Corpus companion entry:** Claim C-323 **Audit Claim ID:** openai-us-ai-action-plan-2025-claim-02

9.18.40. Appendix B40. The language of national security, competition with the PRC, and government adopt...

Claim statement: That the language of national security, competition with the PRC, and government adoption is already integrated into its official 2025 discourse. **Evidence status:** Documented Fact **Source support:** Appendix A5 **Corpus companion entry:** Claim C-324 **Audit Claim ID:** openai-us-ai-action-plan-2025-claim-03

9.18.41. Appendix B41. 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 sta-**

tus: Documented Fact **Source support:** Appendix A8 **Corpus companion entry:** Claim C-328

Audit Claim ID: palantir-anthropic-aws-gov-partnership-2024-07-export-controls-compute-defense-

9.18.42. Appendix B42. 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 A8 **Corpus companion en-**

try: Claim C-329 **Audit Claim ID:** palantir-anthropic-aws-gov-partnership-2024-07-export-controls-

9.18.43. Appendix B43. 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 sta-**

tus: Documented Fact **Source support:** Appendix A8 **Corpus companion entry:** Claim C-330

Audit Claim ID: palantir-anthropic-aws-gov-partnership-2024-07-export-controls-compute-defense-

9.18.44. Appendix B44. Palantir openly markets Gotham to defense agencies and the intelligence community

Claim statement: That Palantir openly markets Gotham to defense agencies and the intelligence community. **Evidence status:** Documented Fact **Source support:** Appendix A11 **Corpus com-**

panion entry: Claim C-331 **Audit Claim ID:** palantir-gotham-2026-11-facebook-thiel-intelligence-adj-

9.18.45. Appendix B45. Palantir's public product framing already includes operational AI, deployment at...

Claim statement: That Palantir's public product framing already includes operational AI, deployment at the edge, and mission-critical decision support. **Evidence status:** Documented Fact

Source support: Appendix A11 **Corpus companion entry:** Claim C-332 **Audit Claim ID:**

palantir-gotham-2026-11-facebook-thiel-intelligence-adjacency-claim-02

9.18.46. Appendix B46. Palantir's position inside security-state and defense workflows does not need to...

Claim statement: That Palantir's position inside security-state and defense workflows does not need to be inferred from hostile or conspiratorial sources because the company states it directly. **Ev-**

idence status: Documented Fact **Source support:** Appendix A11 **Corpus companion entry:**

Claim C-333 **Audit Claim ID:** palantir-gotham-2026-11-facebook-thiel-intelligence-adjacency-claim-

9.18.47. Appendix B47. AI dominance was elevated to explicit state doctrine in 2025

Claim statement: That AI dominance was elevated to explicit state doctrine in 2025. **Evidence status:** Documented Fact **Source support:** Appendix A1 **Corpus companion entry:** Claim C-477 **Audit Claim ID:** whitehouse-eo-14179-2025-08-power-networks-legitimacy-capture-claim-01

9.18.48. Appendix B48. Deregulatory and dominance-oriented state framing can serve as an external source...

Claim statement: That deregulatory and dominance-oriented state framing can serve as an external source of legitimacy for frontier firms. **Evidence status:** Documented Fact **Source support:** Appendix A1 **Corpus companion entry:** Claim C-478 **Audit Claim ID:** whitehouse-eo-14179-2025-08-power-networks-legitimacy-capture-claim-02

9.18.49. Appendix B49. Chapter 8 can ground the idea of institutional capture or mutual reinforcement in...

Claim statement: That Chapter 8 can ground the idea of institutional capture or mutual reinforcement in official state language rather than metaphor alone. **Evidence status:** Documented Fact **Source support:** Appendix A1 **Corpus companion entry:** Claim C-479 **Audit Claim ID:** whitehouse-eo-14179-2025-08-power-networks-legitimacy-capture-claim-03

9.18.50. Appendix B50. Competition over AI was officially framed by the U.S. government in terms of glob...

Claim statement: That competition over AI was officially framed by the U.S. government in terms of global dominance, economic competitiveness, and national security. **Evidence status:** Documented Fact **Source support:** Appendix A1 **Corpus companion entry:** Claim C-480 **Audit Claim ID:** whitehouse-eo-14179-2025-claim-01

9.18.51. Appendix B51. There was an explicit turn toward a more deregulatory line aimed at removing inte...

Claim statement: That there was an explicit turn toward a more deregulatory line aimed at removing internal barriers for U.S. actors. **Evidence status:** Documented Fact **Source support:** Appendix A1 **Corpus companion entry:** Claim C-481 **Audit Claim ID:** whitehouse-eo-14179-2025-claim-02

9.18.52. Appendix B52. The 2025 state framework directly connected AI policy, national security, and adm...

Claim statement: That the 2025 state framework directly connected AI policy, national security, and administrative reordering. **Evidence status:** Documented Fact **Source support:** Appendix A1 **Corpus companion entry:** Claim C-482 **Audit Claim ID:** whitehouse-eo-14179-2025-claim-03

9.18.53. Appendix B53. 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 A6, Appendix A7, Appendix A8, Appendix A9 **Corpus companion entry:** Claim C-437 **Audit Claim ID:** vector07-formal-government-interface-synthesis-claim-01

9.18.54. Appendix B54. 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 A1, Appendix A5, Appendix A2 **Corpus companion entry:** Claim C-438 **Audit Claim ID:** vector07-public-obligations-for-protected-firms-claim-01

9.18.55. Appendix B55. The Facebook-Thiel-Palantir-In-Q-Tel record supports a bounded argument about net...

Claim statement: That the Facebook-Thiel-Palantir-In-Q-Tel record supports a bounded argument about networked governance continuity between surveillance platforms, venture capital, and security-state technology ecosystems. **Evidence status:** Hypothesis / Interpretation **Source support:** Appendix A12, Appendix A10 **Corpus companion entry:** Claim C-459 **Audit Claim ID:** vector11-networked-adjacency-synthesis-claim-01

9.18.56. Appendix B56. The strongest reading of the Facebook-Thiel-Palantir-In-Q-Tel record is a structu...

Claim statement: That the strongest reading of the Facebook-Thiel-Palantir-In-Q-Tel record is a structured adjacency model built through board power, venture sponsorship, and government contracting rather than evidence of covert command. **Evidence status:** Hypothesis / Interpretation **Source support:** Appendix A12, Appendix A10, Appendix A11 **Corpus companion entry:** Claim C-461 **Audit Claim ID:** vector11-structured-adjacency-not-command-claim-01

9.18.57. Appendix B57. Open reasoning papers or public model releases eliminate the politics of compute...

Claim statement: That open reasoning papers or public model releases eliminate the politics of compute concentration, distribution control, or institutional gatekeeping around frontier AI. **Evidence status:** Speculative Narrative Risk **Source support:** Appendix A13 **Corpus companion entry:** Claim C-432 **Audit Claim ID:** vector06-open-publication-solves-control-overclaim-claim-01

9.18.58. Appendix B58. Export controls, hyperscaler dependence, and defense adjacency prove total state...

Claim statement: That export controls, hyperscaler dependence, and defense adjacency prove total state ownership or seamless command over every frontier-model training, deployment, or access pathway. **Evidence status:** Speculative Narrative Risk **Source support:** Appendix A2, Appendix A3 **Corpus companion entry:** Claim C-439 **Audit Claim ID:** vector07-total-state-command-overclaim-claim-01

9.19. Appendix C. Timeline Slice

9.19.1. Appendix C1. 2022-04-08: Meta's proxy states that Peter Thiel will retire from the board after the annual meeting

Event summary: Meta's proxy states that Peter Thiel will retire from the board after the annual meeting **Event date:** 2022-04-08 **Source support:** Appendix A12 **Corpus companion entry:** Event T-017 **Audit Event ID:** event-meta-proxy-thiel-retirement-2022-04-08

9.19.2. Appendix C2. 2023-09-25: Amazon and Anthropic tie frontier-model development to AWS cloud and chips

Event summary: Amazon and Anthropic tie frontier-model development to AWS cloud and chips **Event date:** 2023-09-25 **Source support:** Appendix A3 **Corpus companion entry:** Event T-030 **Audit Event ID:** event-amazon-anthropic-collaboration-2023-09-25

9.19.3. Appendix C3. 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 A8 **Corpus companion entry:** Event T-059 **Audit Event ID:** event-palantir-anthropic-aws-government-c

9.19.4. Appendix C4. 2025-01-22: DeepSeek publishes its R1 reasoning paper and open release claims

Event summary: DeepSeek publishes its R1 reasoning paper and open release claims **Event date:** 2025-01-22 **Source support:** Appendix A13 **Corpus companion entry:** Event T-065 **Audit Event ID:** event-deepseek-r1-2025-01-22

9.19.5. Appendix C5. 2025-01-23: The White House makes sustaining and enhancing U.S. global AI dominance official policy

Event summary: The White House makes sustaining and enhancing U.S. global AI dominance official policy **Event date:** 2025-01-23 **Source support:** Appendix A1 **Corpus companion entry:** Event T-066 **Audit Event ID:** event-eo-14179-2025-01-23

9.19.6. Appendix C6. 2025-02-06: The AI Action Plan RFI is published in the Federal Register

Event summary: The AI Action Plan RFI is published in the Federal Register **Event date:** 2025-02-06 **Source support:** Appendix A4 **Corpus companion entry:** Event T-068 **Audit Event ID:** event-fr-rfi-2025-02-06

9.19.7. Appendix C7. 2025-03-13: OpenAI publishes its proposals for the U.S. AI Action Plan

Event summary: OpenAI publishes its proposals for the U.S. AI Action Plan **Event date:** 2025-03-13 **Source support:** Appendix A5 **Corpus companion entry:** Event T-072 **Audit Event ID:** event-openai-action-plan-2025-03-13

9.19.8. 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 A7 **Corpus companion entry:** Event T-077 **Audit Event ID:** event-anthropic-claude-gov-2025-06-06

9.19.9. Appendix C9. 2025-06-13: The U.S. Army launches Detachment 201 and commissions senior tech executives into a r...

Event summary: The U.S. Army launches Detachment 201 and commissions senior tech executives into a reserve innovation corps **Event date:** 2025-06-13 **Source support:** Appendix A9 **Corpus companion entry:** Event T-078 **Audit Event ID:** event-army-detachment-201-2025-06-13

9.19.10. Appendix C10. 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 A6 **Corpus companion entry:** Event T-079 **Audit Event ID:** event-openai-for-government-2025-06-16

9.19.11. Appendix C11. 2026-05-31: BIS reaffirms license requirements for advanced computing items tied to D:5 and Macau...

Event summary: BIS reaffirms license requirements for advanced computing items tied to D:5 and Macau entities **Event date:** 2026-05-31 **Source support:** Appendix A2 **Corpus companion entry:** Event T-092 **Audit Event ID:** event-bis-advanced-computing-guidance-2026-05-31

9.20. Appendix D. Relevant Entities

9.20.1. Appendix D1. Amazon Web Services

Entity type: cloud_provider **Role in this paper:** Actor relevant to 07_export_controls_compute_defense_acc... **Relevant sources in this paper:** Appendix A3 **Corpus companion entry:** Entity E-009

Audit Entity ID: org-aws

9.20.2. Appendix D2. Amazon

Entity type: company **Role in this paper:** Actor relevant to 07_export_controls_compute_defense_access.

Relevant sources in this paper: Appendix A3 **Corpus companion entry:** Entity E-006

Audit Entity ID: org-amazon

9.20.3. Appendix D3. 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_legitimacy_capture and others. **Relevant sources in this paper:** Appendix A3, Appendix A7, Appendix A8 **Corpus companion entry:** Entity E-008 **Audit Entity ID:** org-anthropic

9.20.4. Appendix D4. DeepSeek

Entity type: company **Role in this paper:** Actor relevant to 06_ai_double_use_reasoning_deepseek, 07_export_controls_compute_defense_access. **Relevant sources in this paper:** Appendix A13 **Corpus companion entry:** Entity E-013 **Audit Entity ID:** org-deepseek

9.20.5. Appendix D5. Facebook / Meta

Entity type: company **Role in this paper:** Actor relevant to 01_surveillance_capitalism, 03_covid_acceleration_and_geopolitics, 04_data_to_models_copyright_weights, 11_facebook_thiel_intelligence_adjacency and others. **Relevant sources in this paper:** Appendix A12 **Corpus companion entry:** Entity E-015 **Audit Entity ID:** org-facebook

9.20.6. Appendix D6. 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 A5, Appendix A6 **Corpus companion entry:** Entity E-027 **Audit Entity ID:** org-openai

9.20.7. Appendix D7. 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 A11, Appendix A8, Appendix A9 **Corpus companion entry:** Entity E-029 **Audit Entity ID:** org-palantir

9.20.8. Appendix D8. Bureau of Industry and Security

Entity type: government **Role in this paper:** Actor relevant to 07_export_controls_compute_defense_access. **Relevant sources in this paper:** Appendix A2 **Corpus companion entry:** Entity E-010 **Audit Entity ID:** org-bis

9.20.9. Appendix D9. The White House / OSTP process

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 A1, Appendix A4 **Corpus companion entry:** Entity E-039 **Audit Entity ID:** org-white-house

9.20.10. Appendix D10. 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 A6 **Corpus companion entry:** Entity E-014 **Audit Entity ID:** org-dod-cdao

9.20.11. Appendix D11. 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

9.20.12. Appendix D12. Central Intelligence Agency

Entity type: government_agency **Role in this paper:** Actor relevant to 11_facebook_thiel_intelligence_adjacency. **Relevant sources in this paper:** Appendix A10 **Corpus companion entry:** Entity E-012 **Audit Entity ID:** org-cia

9.20.13. Appendix D13. In-Q-Tel

Entity type: organization **Role in this paper:** Actor relevant to 01_surveillance_capitalism, 08_power_networks_legitimacy_capture, 11_facebook_thiel_intelligence_adjacency. **Relevant sources in this paper:** Appendix A10 **Corpus companion entry:** Entity E-022 **Audit Entity ID:** org-iqt

9.20.14. Appendix D14. 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 A12 **Corpus companion entry:** Entity E-058 **Audit Entity ID:** person-peter-thiel

9.20.15. Appendix D15. 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 A5 **Corpus companion entry:** Entity E-059 **Audit Entity ID:** person-sam-altman

9.21. Appendix E. Evidence Boundaries

9.21.1. Appendix E1. Disputed Matters

- None registered for this paper.

9.21.2. Appendix E2. Interpretive Boundaries

- Appendix B53: 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 B54: 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 B55: That the Facebook-Thiel-Palantir-In-Q-Tel record supports a bounded argument about networked governance continuity between surveillance platforms, venture capital, and security-state technology ecosystems.
- Appendix B56: That the strongest reading of the Facebook-Thiel-Palantir-In-Q-Tel record is a structured adjacency model built through board power, venture sponsorship, and government contracting rather than evidence of covert command.

9.21.3. Appendix E3. Speculative Narrative Risks

- Appendix B57: That open reasoning papers or public model releases eliminate the politics of compute concentration, distribution control, or institutional gatekeeping around frontier AI.
- Appendix B58: That export controls, hyperscaler dependence, and defense adjacency prove total state ownership or seamless command over every frontier-model training, deployment, or access pathway.

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

- Appendix A13. Companion note: Note N-057. Risk flag: medium. Note focus: That by January 2025 a non-U.S. lab was publicly claiming frontier-adjacent reasoning performance in a primary technical paper.
- Appendix A3. Companion note: Note N-067. Risk flag: low. Note focus: That at least one major frontier lab publicly committed the majority of its workloads and future model development to a hyperscaler cloud provider.
- Appendix A7. 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 A9. Companion note: Note N-069. Risk flag: low. Note focus: That by June 13, 2025 the U.S. Army had created a formal structure to embed senior commercial-tech executives in a military innovation role.
- Appendix A2. Companion note: Note N-070. Risk flag: low. Note focus: That there is a primary U.S. government record directly describing continued license requirements for advanced computing exports linked to Country Group D:5 and Macau.

- Appendix A4. Companion note: Note N-071. Risk flag: low. Note focus: That the language of AI dominance did not remain only in a political order, but moved into an official federal consultation process.
- Appendix A6. 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 A5. Companion note: Note N-073. Risk flag: medium. Note focus: That OpenAI officially defended a public-policy agenda combining global commercial expansion, export controls, and preservation of access to copyrighted material for training.
- Appendix A8. 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 A1. Companion note: Note N-076. Risk flag: low. Note focus: That competition over AI was officially framed by the U.S. government in terms of global dominance, economic competitiveness, and national security.
- Appendix A4. Companion note: Note N-086. Risk flag: low. Note focus: That the state created an official consultation pathway for shaping the 2025 AI doctrine.
- Appendix A10. Companion note: Note N-088. Risk flag: medium. Note focus: That there is a formal institutional channel connecting strategic technology development with intelligence and defense needs.
- Appendix A5. Companion note: Note N-093. Risk flag: medium. Note focus: That OpenAI presents itself as a policy actor seeking to co-design national AI strategy.
- Appendix A1. Companion note: Note N-096. Risk flag: low. Note focus: That AI dominance was elevated to explicit state doctrine in 2025.
- Appendix A10. Companion note: Note N-113. Risk flag: medium. Note focus: That there is an organization officially dedicated to connecting technological innovation with intelligence, defense, and national-security needs.
- Appendix A12. Companion note: Note N-114. Risk flag: medium. Note focus: That Peter Thiel's formal role at Facebook / Meta lasted well beyond the early startup period and remained active into 2022.
- Appendix A11. Companion note: Note N-116. Risk flag: medium. Note focus: That Palantir openly markets Gotham to defense agencies and the intelligence community.

10. Chapter 08. Power network, legitimacy, and institutional capture

10.1. Abstract

Frontier firms do not consolidate power by product alone. OpenAI publicly treated the November 2023 crisis as a governance rupture and responded with a visible legitimacy-repair process that included board reset, formal review, and later restructuring. Once legitimacy is institutionally assembled, frontier AI stops looking like just another industry and begins to look like a state-adjacent governing infrastructure.

10.2. Keywords

Power network, legitimacy, and institutional capture; surveillance capitalism; AI imperialism; government capture; evidence-based dossier

10.3. Research Question

How does power network, legitimacy, and institutional capture function within the dossier's larger argument, and what does the documentary record allow this chapter to establish?

10.4. Scope and Framing Note

This paper examines power network, legitimacy, and institutional capture 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 shows how access is legitimized, stabilized, and normalized through governance, lobbying, elite networks, and institutional prestige. Adjacent reinforcement appears in Chapter 07, Chapter 09, Chapter 10, but those cross-references are supportive rather than required for basic comprehension.

10.5. Central Thesis

Frontier firms do not consolidate power by product alone. They do so by building legitimacy around the product. That legitimacy is assembled through governance form, elite board composition, security credentials, policy access, operational state interfaces, philanthropy, and political spending. The result is not necessarily a secret cabal. It is a documented field of mutual reinforcement through ordinary institutions. ⁸⁸

10.6. Evidence and Method Note

This paper is derived from the chapter corpus for Chapter 08 and is grounded in the project's structured research OS. It relies on source notes, extracted claims, entity profiles, and timeline

⁸⁸Sources: OpenAI, OpenAI announces leadership transition (2023-11-17; see Appendix A1); Associated Press, OpenAI appoints former top US cyberwa... (2024-06-14; see Appendix A2); Axios, Anthropic outspends OpenAI in biggest... (2026-04-21; see Appendix A3)

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.

10.7. Introduction

This paper asks: How does power network, legitimacy, and institutional capture function within the dossier's larger argument, and what does the documentary record allow this chapter to establish? Its working answer is clear from the start: Frontier firms do not consolidate power by product alone. They do so by building legitimacy around the product. That legitimacy is assembled through governance form, elite board composition, security credentials, policy access, operational state interfaces, philanthropy, and political spending. The result is not necessarily a secret cabal. It is a documented field of mutual reinforcement through ordinary institutions.⁸⁹ 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.

10.8. Historical or Institutional Context

OpenAI publicly treated the November 2023 crisis as a governance rupture and responded with a visible legitimacy-repair process that included board reset, formal review, and later restructuring. That shows legitimacy was not incidental to frontier status. It had to be actively rebuilt.⁹⁰

10.9. Main Analysis

The company's later appointments and committee structures show legitimacy being rebuilt through establishment credentials as much as technical authority. Security and defense credibility became part of the governance story.⁹¹

The policy and state-interface layers described in Chapter 07 matter differently here. **OpenAI for Government**, **Claude Gov**, classified-distribution channels, and **Detachment 201** show how privileged state access can be normalized as routine institutional partnership rather than treated as an exceptional concentration of power.⁹²

The influence layer widened by 2026. Lobbying disclosures and PAC records show that frontier legitimacy expanded beyond public messaging into formal Washington and electoral infrastructure.

⁸⁹Sources: OpenAI, OpenAI announces leadership transition (2023-11-17; see Appendix A1); Associated Press, OpenAI appoints former top US cyberwa... (2024-06-14; see Appendix A2); Axios, Anthropic outspends OpenAI in biggest... (2026-04-21; see Appendix A3)

⁹⁰Sources: OpenAI, OpenAI announces leadership transition (2023-11-17; see Appendix A1); OpenAI, OpenAI announces new members to board... (2024-03-08; see Appendix A4); OpenAI, OpenAI board forms safety and securit... (2024-05-28; see Appendix A5)

⁹¹Sources: Associated Press, OpenAI appoints former top US cyberwa... (2024-06-14; see Appendix A2); OpenAI, OpenAI board forms safety and securit... (2024-05-28; see Appendix A5); Financial Times, OpenAI appoints one of Wall Street's... (2025-01-14; see Appendix A6)

⁹²Sources: 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 A9); U.S. Army Public Affairs, Army Launches Detachment 201: Executi... (2025-06-13; see Appendix A10)

That means legitimacy was not only reputational. It was procedural and political. ⁹³

The moral-authority layer widened too. Dario Amodei publicly warned that AI-driven wealth concentration could destabilize democracy while also describing the founders' 80% wealth pledge. The Vatican's 2026 encyclical shows that the legitimacy field had expanded beyond industry and policy into formal civilizational and religious authority. These sources do not absolve frontier firms. They show how moral positioning becomes part of the legitimacy field itself. ⁹⁴

Another part of the legitimacy problem is private survivability hedging. In 2016, Sam Altman publicly described storing guns, gold, potassium iodide, antibiotics, gas masks, and having a New Zealand fallback plan tied to Peter Thiel. In 2023, Wired documented Mark Zuckerberg's Kauai compound as including an underground shelter, secrecy controls on construction crews, and self-sufficiency features; in 2024, Zuckerberg publicly minimized the structure as "just a little shelter." None of this proves that these actors expect imminent collapse or are engineering it. But it does sharpen a bounded political question: legitimacy weakens when elites ask publics and states to normalize dependency on systems they govern while privately investing in exemption from shared risk. ⁹⁵

Named-individual material belongs here only when it clarifies institutional control, mission conflict, or legal architecture. The Musk-OpenAI litigation matters because it shows that nonprofit mission, restructuring, and governance control became a formal legal battleground. The AP injunction-denial reporting further shows that the dispute remained live even after emergency relief was denied. ⁹⁶

Interpretation: this chapter builds the structural case for **Limit Government Capture by AI Companies**. The strongest reading is not hidden sovereign command. It is that governance form, security credentials, policy access, lobbying, philanthropy, and operational state interfaces together normalize extraordinary concentration before democratic controls catch up. ⁹⁷

10.10. Position Within the Series

This paper sits inside a coordinated 11-paper architecture. Its nearest companions are Chapter 07 (Imperial control over access), Chapter 09 (Future if no action is taken), Chapter 10 (What to do). In that series logic, Chapter 08 shows how access is legitimized, stabilized, and normalized through governance, lobbying, elite networks, and institutional prestige.

⁹³Sources: Axios, Anthropic outspends OpenAI in biggest... (2026-04-21; see Appendix A3); Federal Election Commission, LEADING THE FUTURE - committee overview (2026-03-31; see Appendix A11)

⁹⁴Sources: Dario Amodei, The Adolescence of Technology (2026-01-01; see Appendix A12); The Holy See, Encyclical Letter of His Holiness Leo... (2026-05-15; see Appendix A13)

⁹⁵Sources: Tad Friend, Sam Altman's Manifest Destiny (2016-10-03; see Appendix A14); Makena Kelly, Inside Mark Zuckerberg's Top-Secret H... (2023-12-14; see Appendix A15); Charlotte Phillipp, Mark Zuckerberg Downplays His 5,000-S... (2024-12-20; see Appendix A16)

⁹⁶Sources: CourtListener / PACER, Musk v. Altman, 4:24-cv-04722 (2024-08-05; see Appendix A17); Matt O'Brien; Barbara Ortutay, Judge denies Elon Musk's request to b... (2025-03-05; see Appendix A18); Financial Times, Judge denies Elon Musk's attempt to i... (2025-03-05; see Appendix A19)

⁹⁷Sources: 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); Axios, Anthropic outspends OpenAI in biggest... (2026-04-21; see Appendix A3); Dario Amodei, The Adolescence of Technology (2026-01-01; see Appendix A12)

10.11. Counterarguments

The chapter must preserve a crucial distinction: elite overlap and influence channels do not by themselves prove unified hidden control. The record is strongest on formal influence, legitimacy repair, disclosed interfaces, and institutional reinforcement, not on secret ownership of every outcome.
98

10.12. Limits of the Evidence

The chapter must preserve a crucial distinction: elite overlap and influence channels do not by themselves prove unified hidden control. The record is strongest on formal influence, legitimacy repair, disclosed interfaces, and institutional reinforcement, not on secret ownership of every outcome.
99

10.13. Reform Relevance

Once legitimacy is institutionally assembled, frontier AI stops looking like just another industry and begins to look like a state-adjacent governing infrastructure. That is why any material government relationship with frontier firms should trigger transparency, conflict-of-interest controls, auditability, procurement visibility, and limits on exclusive private benefit.
100

10.14. Conclusion

Once legitimacy is institutionally assembled, frontier AI stops looking like just another industry and begins to look like a state-adjacent governing infrastructure. At that point the dossier can ask its final scenario question: if the present trajectory continues, what kind of social order does it produce? In the architecture of this series, Chapter 08 therefore functions as a self-contained argument while also advancing the cumulative path toward the dossier's three reforms.

10.15. 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 08, those boundaries remain visible in the prose, in the bibliography, and in the appendix sections that summarize claims and caution points.

⁹⁸Sources: CourtListener / PACER, Musk v. Altman, 4:24-cv-04722 (2024-08-05; see Appendix A17); Federal Election Commission, LEADING THE FUTURE - committee overview (2026-03-31; see Appendix A11)

⁹⁹Sources: CourtListener / PACER, Musk v. Altman, 4:24-cv-04722 (2024-08-05; see Appendix A17); Federal Election Commission, LEADING THE FUTURE - committee overview (2026-03-31; see Appendix A11)

¹⁰⁰Sources: 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 A10); Axios, Anthropic outspends OpenAI in biggest... (2026-04-21; see Appendix A3)

10.16. References / Bibliography

1. OpenAI. OpenAI announces leadership transition. 2023-11-17. T1 company_statement. Paper appendix: Appendix A1. Corpus companion entry: Source S-092. Audit ID: **openai-leadership-transition-2023**. Relevance: That OpenAI publicly treated its crisis as a governance and legitimacy rupture, not as a simple executive reshuffle.
2. Associated Press. OpenAI appoints former top US cyberwarrior Paul Nakasone to its board of directors. 2024-06-14. T2 news_article. Paper appendix: Appendix A2. Corpus companion entry: Source S-021. Audit ID: **ap-openai-nakasone-board-2024**. Relevance: That OpenAI brought a former top U.S. military and intelligence official directly into its governance structure.
3. Axios. Anthropic outspends OpenAI in biggest-ever lobbying quarter. 2026-04-21. T3 news_article. Paper appendix: Appendix A3. Corpus companion entry: Source S-027. Audit ID: **axios-anthropic-openai-lobbying-quarter-2026**. Relevance: That frontier AI firms were spending meaningful and rising sums on federal lobbying by early 2026.
4. OpenAI. OpenAI announces new members to board of directors. 2024-03-08. T1 company_statement. Paper appendix: Appendix A4. Corpus companion entry: Source S-094. Audit ID: **openai-new-board-members-2024**. Relevance: That OpenAI deliberately rebuilt its board with figures carrying regulatory, legal, nonprofit, and major-platform credentials.
5. OpenAI. OpenAI board forms safety and security committee. 2024-05-28. T1 company_statement. Paper appendix: Appendix A5. Corpus companion entry: Source S-089. Audit ID: **openai-board-forms-safety-security-committee-2024**. Relevance: That OpenAI publicly moved safety and security oversight to a named board-level structure after the governance crisis.
6. Financial Times. OpenAI appoints one of Wall Street's most powerful dealmakers to its board. 2025-01-14. T2 news_article. Paper appendix: Appendix A6. Corpus companion entry: Source S-061. Audit ID: **ft-openai-ogunlesi-board-2025**. Relevance: That OpenAI continued to load its board with elite finance and infrastructure figures during its restructuring phase.
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 government 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. 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 A9. 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.
10. 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 A10. Corpus companion entry: Source S-026. Audit ID: **army-detachment-201-2025**. Relevance:

- vance: That by June 13, 2025 the U.S. Army had created a formal structure to embed senior commercial-tech executives in a military innovation role.
11. Federal Election Commission. **LEADING THE FUTURE** - committee overview. 2026-03-31. T1 committee_overview. Paper appendix: Appendix A11. Corpus companion entry: Source S-053. Audit ID: **fec-leading-the-future-committee-overview-2026**. Relevance: That there is an official federal record for **LEADING THE FUTURE** as an active Super PAC in the 2025-2026 election cycle.
 12. Dario Amodei. The Adolescence of Technology. 2026-01-01. T3 essay. Paper appendix: Appendix A12. 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.
 13. The Holy See. Encyclical Letter of His Holiness Leo XIV Magnifica Humanitas. 2026-05-15. T1 encyclical. Paper appendix: Appendix A13. 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.
 14. Tad Friend. Sam Altman's Manifest Destiny. 2016-10-03. T2 profile. Paper appendix: Appendix A14. 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, potassium iodide, antibiotics, batteries, water, gas masks, and land in Big Sur as part of survival pla...
 15. Makena Kelly. Inside Mark Zuckerberg's Top-Secret Hawaii Compound. 2023-12-14. T2 investigation. Paper appendix: Appendix A15. 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...
 16. Charlotte Phillipp. Mark Zuckerberg Downplays His 5,000-Square-Foot Underground Bunker in Hawaii: 'Just a Little Shelter'. 2024-12-20. T3 news_article. Paper appendix: Appendix A16. Corpus companion entry: Source S-106. Audit ID: **people-zuckerberg-little-shelter-2024**. Relevance: That by December 20, 2024 Mark Zuckerberg publicly minimized the Kauai bunker as 'just a little shelter,' creating a documented contrast between the estate's described features...
 17. CourtListener / PACER. Musk v. Altman, 4:24-cv-04722. 2024-08-05. T1 court_docket. Paper appendix: Appendix A17. Corpus companion entry: Source S-045. Audit ID: **courtlistener-musk-v-altman-docket-2024**. Relevance: That the Musk-OpenAI dispute had a concrete federal-court vehicle with named defendants and an active procedural history.
 18. Matt O'Brien; Barbara Ortutay. Judge denies Elon Musk's request to block OpenAI for-profit conversion but welcomes trial. 2025-03-05. T2 news_article. Paper appendix: Appendix A18. Corpus companion entry: Source S-015. Audit ID: **ap-musk-openai-injunction-denied-2025**. Relevance: That by March 5, 2025 Musk had failed to obtain immediate judicial relief blocking OpenAI's conversion path.
 19. Financial Times. Judge denies Elon Musk's attempt to immediately block OpenAI's conversion to for-profit entity. 2025-03-05. T2 news_article. Paper appendix: Appendix A19. Corpus companion entry: Source S-057. Audit ID: **ft-musk-openai-injunction-denied-2025**. Rel-

evance: That by March 5, 2025 Musk had failed to obtain immediate judicial relief blocking OpenAI's conversion path.

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 OpenAI presents itself as a policy actor seeking to co-design national AI strategy.

10.17. Appendix A. Source Register for This Paper

10.17.1. Appendix A1. OpenAI announces leadership transition

Source citation: OpenAI announces leadership transition **Institution / author:** OpenAI **Date:** 2023-11-17 **Tier / type:** T1 / company_statement **Why it matters in this paper:** That OpenAI publicly treated its crisis as a governance and legitimacy rupture, not as a simple executive reshuffle. **Corpus companion entry:** Source S-092 **Audit ID:** openai-leadership-transition-2023

10.17.2. Appendix A2. OpenAI appoints former top US cyberwarrior Paul Nakasone to its board of directors

Source citation: OpenAI appoints former top US cyberwarrior Paul Nakasone to its board of directors **Institution / author:** Associated Press **Date:** 2024-06-14 **Tier / type:** T2 / news_article **Why it matters in this paper:** That OpenAI brought a former top U.S. military and intelligence official directly into its governance structure. **Corpus companion entry:** Source S-021 **Audit ID:** ap-openai-nakasone-board-2024

10.17.3. Appendix A3. 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 AI firms were spending meaningful and rising sums on federal lobbying by early 2026. **Corpus companion entry:** Source S-027 **Audit ID:** axios-anthropic-openai-lobbying-quarter-2026

10.17.4. Appendix A4. OpenAI announces new members to board of directors

Source citation: OpenAI announces new members to board of directors **Institution / author:** OpenAI **Date:** 2024-03-08 **Tier / type:** T1 / company_statement **Why it matters in this paper:** That OpenAI deliberately rebuilt its board with figures carrying regulatory, legal, non-profit, and major-platform credentials. **Corpus companion entry:** Source S-094 **Audit ID:** openai-new-board-members-2024

10.17.5. Appendix A5. OpenAI board forms safety and security committee

Source citation: OpenAI board forms safety and security committee **Institution / author:** OpenAI **Date:** 2024-05-28 **Tier / type:** T1 / company_statement **Why it matters in this paper:** That OpenAI publicly moved safety and security oversight to a named board-level

structure after the governance crisis. **Corpus companion entry:** Source S-089 **Audit ID:** openai-board-forms-safety-security-committee-2024

10.17.6. Appendix A6. OpenAI appoints one of Wall Street’s most powerful dealmakers to its board

Source citation: OpenAI appoints one of Wall Street’s most powerful dealmakers to its board **Institution / author:** Financial Times **Date:** 2025-01-14 **Tier / type:** T2 / news_article **Why it matters in this paper:** That OpenAI continued to load its board with elite finance and infrastructure figures during its restructuring phase. **Corpus companion entry:** Source S-061 **Audit ID:** ft-openai-ogunlesi-board-2025

10.17.7. 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

10.17.8. 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

10.17.9. Appendix A9. 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

10.17.10. Appendix A10. 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 13, 2025 the U.S. Army had created a formal structure to embed senior commercial-tech executives in a military innovation role. **Corpus companion entry:** Source S-026 **Audit ID:** army-detachment-201-2025

10.17.11. Appendix A11. 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 there is an official federal record for LEADING THE FUTURE as an active Super PAC in the 2025-2026 election cycle. **Corpus companion entry:** Source S-053 **Audit ID:** fec-leading-the-future-committee-overview-2026

10.17.12. Appendix A12. 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

10.17.13. Appendix A13. 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

10.17.14. Appendix A14. 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

10.17.15. Appendix A15. 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 mechani... **Corpus companion entry:** Source S-133 **Audit ID:** wired-zuckerberg-hawaii-compound-2023

10.17.16. Appendix A16. Mark Zuckerberg Downplays His 5,000-Square-Foot Underground Bunker in Hawaii: 'Just a Little Shelter'

Source citation: Mark Zuckerberg Downplays His 5,000-Square-Foot Underground Bunker in Hawaii: 'Just a Little Shelter' **Institution / author:** Charlotte Phillipp **Date:** 2024-12-20 **Tier / type:** T3 / news_article **Why it matters in this paper:** That by December 20, 2024 Mark

Zuckerberg publicly minimized the Kauai bunker as ‘just a little shelter,’ creating a documented contrast between the estate’s described features... **Corpus companion entry:** Source S-106 **Audit ID:** people-zuckerberg-little-shelter-2024

10.17.17. Appendix A17. Musk v. Altman, 4:24-cv-04722

Source citation: Musk v. Altman, 4:24-cv-04722 **Institution / author:** CourtListener / PACER **Date:** 2024-08-05 **Tier / type:** T1 / court_docket **Why it matters in this paper:** That the Musk-OpenAI dispute had a concrete federal-court vehicle with named defendants and an active procedural history. **Corpus companion entry:** Source S-045 **Audit ID:** courtlistener-musk-v-altman-docket-2024

10.17.18. Appendix A18. Judge denies Elon Musk’s request to block OpenAI for-profit conversion but welcomes trial

Source citation: Judge denies Elon Musk’s request to block OpenAI for-profit conversion but welcomes trial **Institution / author:** Matt O’Brien; Barbara Ortutay **Date:** 2025-03-05 **Tier / type:** T2 / news_article **Why it matters in this paper:** That by March 5, 2025 Musk had failed to obtain immediate judicial relief blocking OpenAI’s conversion path. **Corpus companion entry:** Source S-015 **Audit ID:** ap-musk-openai-injunction-denied-2025

10.17.19. Appendix A19. Judge denies Elon Musk’s attempt to immediately block OpenAI’s conversion to for-profit entity

Source citation: Judge denies Elon Musk’s attempt to immediately block OpenAI’s conversion to for-profit entity **Institution / author:** Financial Times **Date:** 2025-03-05 **Tier / type:** T2 / news_article **Why it matters in this paper:** That by March 5, 2025 Musk had failed to obtain immediate judicial relief blocking OpenAI’s conversion path. **Corpus companion entry:** Source S-057 **Audit ID:** ft-musk-openai-injunction-denied-2025

10.17.20. 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 OpenAI presents itself as a policy actor seeking to co-design national AI strategy. **Corpus companion entry:** Source S-096 **Audit ID:** openai-us-ai-action-plan-2025

10.18. Appendix B. Claims Used in This Paper

10.18.1. Appendix B1. 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

10.18.2. Appendix B2. 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 **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

10.18.3. Appendix B3. 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

10.18.4. Appendix B4. By March 5, 2025 Musk had failed to obtain immediate judicial relief blocking Ope...

Claim statement: That by March 5, 2025 Musk had failed to obtain immediate judicial relief blocking OpenAI's conversion path. **Evidence status:** Documented Fact **Source support:** Appendix A18 **Corpus companion entry:** Claim C-043 **Audit Claim ID:** ap-musk-openai-injunction-denied-2025-14-named-individuals-disputes-and-risk-claim-01

10.18.5. Appendix B5. The dispute over OpenAI's nonprofit mission and restructuring remained procedural...

Claim statement: That the dispute over OpenAI's nonprofit mission and restructuring remained procedurally live even after that denial because the court signaled willingness to expedite trial. **Evidence status:** Documented Fact **Source support:** Appendix A18 **Corpus companion entry:** Claim C-044 **Audit Claim ID:** ap-musk-openai-injunction-denied-2025-14-named-individuals-disputes-

10.18.6. Appendix B6. Chapter 08 can distinguish between a live governance challenge, a denied prelimin...

Claim statement: That Chapter 08 can distinguish between a live governance challenge, a denied preliminary injunction, and a full merits finding. **Evidence status:** Documented Fact **Source support:** Appendix A18 **Corpus companion entry:** Claim C-045 **Audit Claim ID:** ap-musk-openai-injunction-denied-2025-14-named-individuals-disputes-and-risk-claim-03

10.18.7. Appendix B7. OpenAI brought a former top U.S. military and intelligence official directly into...

Claim statement: That OpenAI brought a former top U.S. military and intelligence official directly into its governance structure. **Evidence status:** Documented Fact **Source support:** Appendix A2 **Corpus companion entry:** Claim C-064 **Audit Claim ID:** ap-openai-nakasone-board-2024-08-power-networks-legitimacy-capture-claim-01

10.18.8. Appendix B8. Defense and cybersecurity credentials became part of OpenAI's legitimacy and over...

Claim statement: That defense and cybersecurity credentials became part of OpenAI's legitimacy and oversight architecture. **Evidence status:** Documented Fact **Source support:** Appendix A2 **Corpus companion entry:** Claim C-065 **Audit Claim ID:** ap-openai-nakasone-board-2024-08-power-netw

10.18.9. Appendix B9. The company publicly linked security governance with its post-crisis institutiona...

Claim statement: That the company publicly linked security governance with its post-crisis institutional rebuild. **Evidence status:** Documented Fact **Source support:** Appendix A2 **Corpus companion entry:** Claim C-066 **Audit Claim ID:** ap-openai-nakasone-board-2024-08-power-networks-leg

10.18.10. Appendix B10. By June 13, 2025 the U.S. Army had created a formal structure to embed senior com...

Claim statement: That by June 13, 2025 the U.S. Army had created a formal structure to embed senior commercial-tech executives in a military innovation role. **Evidence status:** Documented Fact **Source support:** Appendix A10 **Corpus companion entry:** Claim C-078 **Audit Claim ID:** army-detachment-201-2025-07-export-controls-compute-defense-access-claim-01

10.18.11. Appendix B11. The initial cohort explicitly included executives from Palantir and OpenAI, not o...

Claim statement: That the initial cohort explicitly included executives from Palantir and OpenAI, not only traditional defense contractors. **Evidence status:** Documented Fact **Source support:** Appendix A10 **Corpus companion entry:** Claim C-079 **Audit Claim ID:** army-detachment-201-2025-07-export-controls-compute-defense-access-claim-02

10.18.12. Appendix B12. Chapter 07 and Chapter 08 can document a direct institutional bridge between fron...

Claim statement: That Chapter 07 and Chapter 08 can document a direct institutional bridge between frontier/commercial AI leadership and military structures without relying on rumor or guilt by association. **Evidence status:** Documented Fact **Source support:** Appendix A10 **Corpus companion entry:** Claim C-080 **Audit Claim ID:** army-detachment-201-2025-07-export-controls-compute

10.18.13. Appendix B13. Frontier AI firms were spending meaningful and rising sums on federal lobbying by...

Claim statement: That frontier AI firms were spending meaningful and rising sums on federal lobbying by early 2026. **Evidence status:** Documented Fact **Source support:** Appendix A3 **Corpus companion entry:** Claim C-081 **Audit Claim ID:** axios-anthropic-openai-lobbying-quarter-2026-08

10.18.14. Appendix B14. OpenAI and Anthropic were not only releasing models and policy essays, but also p...

Claim statement: That OpenAI and Anthropic were not only releasing models and policy essays, but also participating in Washington through formal lobbying channels. **Evidence status:** Documented Fact **Source support:** Appendix A3 **Corpus companion entry:** Claim C-082 **Audit Claim ID:** axios-anthropic-openai-lobbying-quarter-2026-08-power-networks-legitimacy-capture-cla

10.18.15. Appendix B15. Chapter 08 can discuss institutional influence with evidence tied to disclosed sp...

Claim statement: That Chapter 08 can discuss institutional influence with evidence tied to disclosed spending and issue categories, not only inferred proximity. **Evidence status:** Documented Fact **Source support:** Appendix A3 **Corpus companion entry:** Claim C-083 **Audit Claim ID:** axios-anthropic-openai-lobbying-quarter-2026-08-power-networks-legitimacy-capture-claim-03

10.18.16. Appendix B16. The Musk-OpenAI dispute had a concrete federal-court vehicle with named defendant...

Claim statement: That the Musk-OpenAI dispute had a concrete federal-court vehicle with named defendants and an active procedural history. **Evidence status:** Documented Fact **Source support:** Appendix A17 **Corpus companion entry:** Claim C-135 **Audit Claim ID:** courtlistener-musk-v-altman-docket-2024-14-named-individuals-disputes-and-risk-claim-01

10.18.17. Appendix B17. Chapter 8 can describe the conflict as a formal institutional struggle rather tha...

Claim statement: That Chapter 8 can describe the conflict as a formal institutional struggle rather than a social-media quarrel. **Evidence status:** Documented Fact **Source support:** Appendix A17 **Corpus companion entry:** Claim C-136 **Audit Claim ID:** courtlistener-musk-v-altman-docket-2024-14-named-individuals-disputes-and-risk-claim-02

10.18.18. Appendix B18. The dossier can anchor named-individual discussion in court records when discussi...

Claim statement: That the dossier can anchor named-individual discussion in court records when discussing governance conflict. **Evidence status:** Documented Fact **Source support:** Appendix A17 **Corpus companion entry:** Claim C-137 **Audit Claim ID:** courtlistener-musk-v-altman-docket-2024-

10.18.19. Appendix B19. 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 A12 **Corpus companion entry:** Claim C-141 **Audit Claim ID:** dario-amodei-adolescence-of-technology-2026-08-power-networks-legitimacy-capture-clai

10.18.20. Appendix B20. 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 A12 **Corpus companion entry:** Claim C-142 **Audit Claim ID:** dario-amodei-adolescence-of-technology-2026-08-power-networks-legitimacy-capture-claim-02

10.18.21. Appendix B21. 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 A12 **Corpus companion entry:** Claim C-143 **Audit Claim ID:** dario-amodei-adolescence-of-technology-2026-08-power-networks-legiti

10.18.22. Appendix B22. 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 A12 **Corpus companion entry:** Claim C-144 **Audit Claim ID:** dario-amodei-adolescence-of-technology-2026-08-power-networks-legitimacy-capture-claim-04

10.18.23. Appendix B23. 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 A12 **Corpus companion entry:** Claim C-145 **Audit Claim ID:** dario-amodei-adolescence-of-technology-2026-08-power-networks-legitimacy-capture-claim-05

10.18.24. Appendix B24. There is an official federal record for LEADING THE FUTURE as an active Super PAC...

Claim statement: That there is an official federal record for LEADING THE FUTURE as an active Super PAC in the 2025-2026 election cycle. **Evidence status:** Documented Fact **Source support:** Appendix A11 **Corpus companion entry:** Claim C-164 **Audit Claim ID:** fec-leading-the-future-committee-overview-2026-08-power-networks-legitimacy-capture-claim-01

10.18.25. Appendix B25. Chapter 08 can support the electoral-influence layer with a primary filing-backed...

Claim statement: That Chapter 08 can support the electoral-influence layer with a primary filing-backed source rather than relying only on secondary reporting. **Evidence status:** Documented Fact **Source support:** Appendix A11 **Corpus companion entry:** Claim C-165 **Audit Claim ID:** fec-leading-the-future-committee-overview-2026-08-power-networks-legitimacy-capture-claim-0

10.18.26. Appendix B26. The political vehicle discussed in later secondary reporting had real, large-scal...

Claim statement: That the political vehicle discussed in later secondary reporting had real, large-scale financial activity documented by the FEC. **Evidence status:** Documented Fact **Source support:** Appendix A11 **Corpus companion entry:** Claim C-166 **Audit Claim ID:** fec-leading-the-future-committee-overview-2026-08-power-networks-legitimacy-capture-claim-03

10.18.27. Appendix B27. By March 5, 2025 Musk had failed to obtain immediate judicial relief blocking Ope...

Claim statement: That by March 5, 2025 Musk had failed to obtain immediate judicial relief blocking OpenAI's conversion path. **Evidence status:** Documented Fact **Source support:** Appendix A19 **Corpus companion entry:** Claim C-179 **Audit Claim ID:** ft-musk-openai-injunction-denied-2025-14-named-individuals-disputes-and-risk-claim-01

10.18.28. Appendix B28. The dispute over OpenAI's nonprofit mission and restructuring remained procedural...

Claim statement: That the dispute over OpenAI's nonprofit mission and restructuring remained procedurally live even after that denial. **Evidence status:** Documented Fact **Source support:** Appendix A19 **Corpus companion entry:** Claim C-180 **Audit Claim ID:** ft-musk-openai-injunction-denied-2025-14-named-individuals-disputes-and-risk-claim-02

10.18.29. Appendix B29. Chapter 08 can distinguish between a live governance challenge and a successful m...

Claim statement: That Chapter 08 can distinguish between a live governance challenge and a successful merits finding. **Evidence status:** Documented Fact **Source support:** Appendix A19 **Corpus companion entry:** Claim C-181 **Audit Claim ID:** ft-musk-openai-injunction-denied-2025-14-na

10.18.30. Appendix B30. OpenAI continued to load its board with elite finance and infrastructure figures...

Claim statement: That OpenAI continued to load its board with elite finance and infrastructure figures during its restructuring phase. **Evidence status:** Documented Fact **Source support:** Appendix A6 **Corpus companion entry:** Claim C-190 **Audit Claim ID:** ft-openai-ogunlesi-board-2025-08-power-networks-legitimacy-capture-claim-01

10.18.31. Appendix B31. Chapter 08 can treat board composition as part of a broader legitimacy and capita...

Claim statement: That Chapter 08 can treat board composition as part of a broader legitimacy and capital-alignment architecture, not just mission oversight. **Evidence status:** Documented Fact **Source support:** Appendix A6 **Corpus companion entry:** Claim C-191 **Audit Claim ID:** ft-openai-ogunlesi-board-2025-08-power-networks-legitimacy-capture-claim-02

10.18.32. Appendix B32. The company’s governance network was expanding toward actors useful for large-sca...

Claim statement: That the company’s governance network was expanding toward actors useful for large-scale infrastructure buildout and institutional credibility. **Evidence status:** Documented Fact **Source support:** Appendix A6 **Corpus companion entry:** Claim C-192 **Audit Claim ID:** ft-openai-ogunlesi-board-2025-08-power-networks-legitimacy-capture-claim-03

10.18.33. Appendix B33. 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 A14 **Corpus companion entry:** Claim C-275 **Audit Claim ID:** newyorker-sam-altman-manifest-destiny-2016-08-power-networks-legitimacy-capture-claim

10.18.34. Appendix B34. OpenAI publicly moved safety and security oversight to a named board-level struct...

Claim statement: That OpenAI publicly moved safety and security oversight to a named board-level structure after the governance crisis. **Evidence status:** Documented Fact **Source support:** Appendix A5 **Corpus companion entry:** Claim C-295 **Audit Claim ID:** openai-board-forms-safety-security-committee-2024-13-openai-governance-qstar-talent-exodus-cla

10.18.35. Appendix B35. The company explicitly linked committee formation to a new frontier-model trainin...

Claim statement: That the company explicitly linked committee formation to a new frontier-model training cycle. **Evidence status:** Documented Fact **Source support:** Appendix A5 **Corpus companion entry:** Claim C-296 **Audit Claim ID:** openai-board-forms-safety-security-committee-20

10.18.36. Appendix B36. Board governance, safety rhetoric, and external security expertise were being int...

Claim statement: That board governance, safety rhetoric, and external security expertise were being integrated into the same institutional response. **Evidence status:** Documented Fact **Source support:** Appendix A5 **Corpus companion entry:** Claim C-297 **Audit Claim ID:** openai-board-forms-safety-security-committee-2024-13-openai-governance-qstar-talent-exodus-cla

10.18.37. Appendix B37. 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

10.18.38. Appendix B38. 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

10.18.39. Appendix B39. 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

10.18.40. Appendix B40. OpenAI publicly treated its crisis as a governance and legitimacy rupture, not as...

Claim statement: That OpenAI publicly treated its crisis as a governance and legitimacy rupture, not as a simple executive reshuffle. **Evidence status:** Documented Fact **Source support:** Appendix A1 **Corpus companion entry:** Claim C-304 **Audit Claim ID:** openai-leadership-transition-2023-08-power-networks-legitimacy-capture-claim-01

10.18.41. Appendix B41. The nonprofit-mission structure remained central to how the company justified pow...

Claim statement: That the nonprofit-mission structure remained central to how the company justified power over leadership. **Evidence status:** Documented Fact **Source support:** Appendix A1 **Corpus companion entry:** Claim C-305 **Audit Claim ID:** openai-leadership-transition-2023-08-power-

10.18.42. Appendix B42. Chapter 8 can discuss institutional legitimacy as embedded in corporate form, boa...

Claim statement: That Chapter 8 can discuss institutional legitimacy as embedded in corporate form, board authority, and mission rhetoric. **Evidence status:** Documented Fact **Source support:** Appendix A1 **Corpus companion entry:** Claim C-306 **Audit Claim ID:** openai-leadership-transition-2023-08-power-networks-legitimacy-capture-claim-03

10.18.43. Appendix B43. The November 2023 crisis was recognized by OpenAI itself as a formal governance r...

Claim statement: That the November 2023 crisis was recognized by OpenAI itself as a formal governance rupture. **Evidence status:** Documented Fact **Source support:** Appendix A1 **Corpus companion entry:** Claim C-307 **Audit Claim ID:** openai-leadership-transition-2023-claim-01

10.18.44. Appendix B44. The board justified Altman’s removal in terms of candor, oversight capacity, and...

Claim statement: That the board justified Altman’s removal in terms of candor, oversight capacity, and institutional mission. **Evidence status:** Documented Fact **Source support:** Appendix A1 **Corpus companion entry:** Claim C-308 **Audit Claim ID:** openai-leadership-transition-2023-claim-02

10.18.45. Appendix B45. OpenAI still publicly presented itself as an exceptional structure where nonprofi...

Claim statement: That OpenAI still publicly presented itself as an exceptional structure where nonprofit mission and oversight were meant to prevail over ordinary corporate logic. **Evidence status:** Documented Fact **Source support:** Appendix A1 **Corpus companion entry:** Claim C-309 **Audit Claim ID:** openai-leadership-transition-2023-claim-03

10.18.46. Appendix B46. OpenAI deliberately rebuilt its board with figures carrying regulatory, legal, no...

Claim statement: That OpenAI deliberately rebuilt its board with figures carrying regulatory, legal, nonprofit, and major-platform credentials. **Evidence status:** Documented Fact **Source support:** Appendix A4 **Corpus companion entry:** Claim C-313 **Audit Claim ID:** openai-new-board-members-2024-08-power-networks-legitimacy-capture-claim-01

10.18.47. Appendix B47. Legitimacy and governance repair were explicit parts of OpenAI’s public response...

Claim statement: That legitimacy and governance repair were explicit parts of OpenAI’s public response to the November 2023 crisis. **Evidence status:** Documented Fact **Source support:** Appendix A4 **Corpus companion entry:** Claim C-314 **Audit Claim ID:** openai-new-board-members-2024-08-power-networks-legitimacy-capture-claim-02

10.18.48. Appendix B48. Chapter 8 can discuss institutional board composition as a structural mechanism o...

Claim statement: That Chapter 8 can discuss institutional board composition as a structural mechanism of legitimacy, not just a public-relations detail. **Evidence status:** Documented Fact **Source support:** Appendix A4 **Corpus companion entry:** Claim C-315 **Audit Claim ID:** openai-new-board-members-2024-08-power-networks-legitimacy-capture-claim-03

10.18.49. Appendix B49. OpenAI presents itself as a policy actor seeking to co-design national AI strategy

Claim statement: That OpenAI presents itself as a policy actor seeking to co-design national AI strategy. **Evidence status:** Documented Fact **Source support:** Appendix A20 **Corpus companion entry:** Claim C-319 **Audit Claim ID:** openai-us-ai-action-plan-2025-08-power-networks-legitimacy-capture-claim-04

10.18.50. Appendix B50. The company publicly links its commercial and technical agenda to national-security...

Claim statement: That the company publicly links its commercial and technical agenda to national-security and industrial-policy language. **Evidence status:** Documented Fact **Source support:** Appendix A20 **Corpus companion entry:** Claim C-320 **Audit Claim ID:** openai-us-ai-action-plan-2025-08-power-networks-legitimacy-capture-claim-02

10.18.51. Appendix B51. Chapter 8 can treat lobbying and doctrinal positioning as part of the legitimacy...

Claim statement: That Chapter 8 can treat lobbying and doctrinal positioning as part of the legitimacy architecture of frontier firms. **Evidence status:** Documented Fact **Source support:** Appendix A20 **Corpus companion entry:** Claim C-321 **Audit Claim ID:** openai-us-ai-action-plan-2025-08-power-networks-legitimacy-capture-claim-03

10.18.52. Appendix B52. OpenAI officially defended a public-policy agenda combining global commercial exp...

Claim statement: That OpenAI officially defended a public-policy agenda combining global commercial expansion, export controls, and preservation of access to copyrighted material for training. **Evidence status:** Documented Fact **Source support:** Appendix A20 **Corpus companion entry:** Claim C-322 **Audit Claim ID:** openai-us-ai-action-plan-2025-claim-01

10.18.53. Appendix B53. The company presents itself not only as a model provider, but as an actor seeking...

Claim statement: That the company presents itself not only as a model provider, but as an actor seeking to co-design U.S. national AI strategy. **Evidence status:** Documented Fact **Source support:** Appendix A20 **Corpus companion entry:** Claim C-323 **Audit Claim ID:** openai-us-ai-action-plan-2025-claim-02

10.18.54. Appendix B54. The language of national security, competition with the PRC, and government adopt...

Claim statement: That the language of national security, competition with the PRC, and government adoption is already integrated into its official 2025 discourse. **Evidence status:** Documented Fact **Source support:** Appendix A20 **Corpus companion entry:** Claim C-324 **Audit Claim ID:** openai-us-ai-action-plan-2025-claim-03

10.18.55. Appendix B55. 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 A9 **Corpus companion entry:** Claim C-328 **Audit Claim ID:** palantir-anthropic-aws-gov-partnership-2024-07-export-controls-compute-defense-

10.18.56. Appendix B56. 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 A9 **Corpus companion entry:** Claim C-329 **Audit Claim ID:** palantir-anthropic-aws-gov-partnership-2024-07-export-controls-c

10.18.57. Appendix B57. 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 A9 **Corpus companion entry:** Claim C-330

Audit Claim ID: palantir-anthropic-aws-gov-partnership-2024-07-export-controls-compute-defense-

10.18.58. Appendix B58. By December 20, 2024 Mark Zuckerberg publicly minimized the Kauai bunker as 'just...

Claim statement: That by December 20, 2024 Mark Zuckerberg publicly minimized the Kauai bunker as 'just a little shelter,' creating a documented contrast between the estate's described features and his own framing of them.

Evidence status: Documented Fact **Source support:** Appendix A16 **Corpus companion entry:** Claim C-352 **Audit Claim ID:** people-zuckerberg-little-shelter-2024-08-power-networks-legitimacy-capture-claim-01

10.18.59. Appendix B59. 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 A13 **Corpus companion entry:** Claim C-413 **Audit Claim ID:** vatican-magnifica-humanitas-2026-08-power-networks-legitimacy-capture-claim-

10.18.60. Appendix B60. 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 responsibility.

Evidence status: Documented Fact **Source support:** Appendix A13 **Corpus companion entry:** Claim C-414 **Audit Claim ID:** vatican-magnifica-humanitas-2026-08-power-networks-legitim

10.18.61. Appendix B61. 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 A13 **Corpus companion entry:** Claim C-415 **Audit Claim ID:** vatican-magnifica-humanitas-2026-08-power-networks-legitimacy-capture-claim-03

10.18.62. Appendix B62. 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. **Evidence status:** Documented Fact **Source support:** Appendix A13 **Corpus companion entry:** Claim C-416 **Audit Claim ID:** vatican-magnifica-humanitas-2026-08-power-networks-legitimacy-capture-claim-04

10.18.63. Appendix B63. 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 A9, Appendix A10 **Corpus companion entry:** Claim C-441 **Audit Claim ID:** vector08-formal-state-interface-documented-claim-01

10.18.64. Appendix B64. 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 A15 **Corpus companion entry:** Claim C-489 **Audit Claim ID:** wired-zuckerberg-hawaii-compound-2023-08-power-networks-legitimacy-capture-claim-01

10.18.65. Appendix B65. 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 A15 **Corpus companion entry:** Claim C-490 **Audit Claim ID:** wired-zuckerberg-hawaii-compound-2023-08-power-networks-legitimacy-capture-claim-02

10.18.66. Appendix B66. The precise degree of direct OpenAI alignment behind LEADING THE FUTURE remains c...

Claim statement: That the precise degree of direct OpenAI alignment behind LEADING THE FUTURE remains contested beyond the current FEC filing record and should remain bounded to the Axios reporting layer. **Evidence status:** Disputed Fact **Source support:** Appendix A11 **Corpus companion entry:** Claim C-443 **Audit Claim ID:** vector08-leading-the-future-linkage-bounded-claim-03

10.18.67. Appendix B67. The public record does not settle whether Q* or any specific capability breakthro...

Claim statement: That the public record does not settle whether Q* or any specific capability breakthrough directly triggered the November 2023 OpenAI board rupture. **Evidence status:** Disputed Fact **Source support:** Appendix A1 **Corpus companion entry:** Claim C-465 **Audit Claim ID:** vector13-qstar-causality-contested-claim-01

10.18.68. Appendix B68. Allegations advanced by named actors in OpenAI governance disputes are procedural...

Claim statement: That allegations advanced by named actors in OpenAI governance disputes are procedurally live and politically relevant without being established as true. **Evidence status:** Disputed Fact **Source support:** Appendix A17 **Corpus companion entry:** Claim C-468 **Audit Claim ID:** vector14-live-allegations-not-established-claim-01

10.18.69. Appendix B69. 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 A9, Appendix A10 **Corpus companion entry:** Claim C-437 **Audit Claim ID:** vector07-formal-government-interface-synthesis-claim-01

10.18.70. Appendix B70. 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 A4, Appendix A2, Appendix A3, Appendix A11 **Corpus companion entry:** Claim C-444 **Audit Claim ID:** vector08-legitimacy-stabilized-through-institutions-claim-01

10.18.71. Appendix B71. 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 A9, Appendix A10, Appendix A20 **Corpus companion entry:** Claim C-445 **Audit Claim ID:** vector08-legitimacy-through-operational-state-interfaces-claim-01

10.18.72. Appendix B72. 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 A14, Appendix A15, Appendix A16 **Corpus companion entry:** Claim C-447 **Audit Claim ID:** vector08-revealed-preference-gap-claim-01

10.18.73. Appendix B73. OpenAI's governance rupture, safety rhetoric, leadership churn, and reasoning-mod...

Claim statement: That OpenAI's governance rupture, safety rhetoric, leadership churn, and reasoning-model rollout are best read as parts of one reorganization of authority rather than as isolated episodes. **Evidence status:** Hypothesis / Interpretation **Source support:** Appendix A1, Appendix A5 **Corpus companion entry:** Claim C-464 **Audit Claim ID:** vector13-governance-productization-synthesis-claim-01

10.18.74. Appendix B74. Named-individual disputes should enter the dossier only when they clarify governa...

Claim statement: That named-individual disputes should enter the dossier only when they clarify governance structure, institutional control, or legal architecture rather than functioning as scandal collage. **Evidence status:** Hypothesis / Interpretation **Source support:** Appendix A17 **Corpus companion entry:** Claim C-469 **Audit Claim ID:** vector14-structural-relevance-synthesis-claim-01

10.18.75. Appendix B75. 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 A14, Appendix A15, Appendix A16 **Corpus companion entry:** Claim C-440 **Audit Claim ID:** vector08-bunker-foreknowledge-overclaim-claim-01

10.18.76. Appendix B76. Elite overlap among frontier labs, venture capital, lobbying networks, and defens...

Claim statement: That elite overlap among frontier labs, venture capital, lobbying networks, and defense-adjacent institutions proves a single hidden command chain. **Evidence status:** Speculative Narrative Risk **Source support:** Appendix A2, Appendix A3 **Corpus companion entry:** Claim C-442 **Audit Claim ID:** vector08-hidden-command-chain-overclaim-claim-01

10.18.77. Appendix B77. 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 A12 **Corpus companion entry:** Claim C-446 **Audit Claim ID:** vector08-philanthropy-foreknowledge-overclaim-claim-01

10.18.78. Appendix B78. 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 A9, Appendix A10 **Corpus companion entry:** Claim C-448 **Audit Claim ID:** vector08-state-interface-hidden-control-overclaim-claim-01

10.18.79. Appendix B79. Elite overlap, litigation conflict, or reputational controversy around named indi...

Claim statement: That elite overlap, litigation conflict, or reputational controversy around named individuals justifies guilt by association without direct structural proof. **Evidence status:** Speculative Narrative Risk **Source support:** Appendix A17 **Corpus companion entry:** Claim C-467 **Audit Claim ID:** vector14-guilt-by-association-overclaim-claim-01

10.19. Appendix C. Timeline Slice

10.19.1. 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 A14 **Corpus companion entry:** Event T-004 **Audit Event ID:** event-altman-prep-new-zealand-2016-10-03

10.19.2. Appendix C2. 2023-11-17: OpenAI announces Sam Altman’s departure and names Mira Murati interim CEO

Event summary: OpenAI announces Sam Altman’s departure and names Mira Murati interim CEO **Event date:** 2023-11-17 **Source support:** Appendix A1 **Corpus companion entry:** Event T-032 **Audit Event ID:** event-openai-board-2023-11-17

10.19.3. Appendix C3. 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 A15 **Corpus companion entry:** Event T-035 **Audit Event ID:** event-wired-zuckerberg-kauai-compound-20

10.19.4. Appendix C4. 2024-03-08: OpenAI expands its board with new institutional figures

Event summary: OpenAI expands its board with new institutional figures **Event date:** 2024-03-08 **Source support:** Appendix A4 **Corpus companion entry:** Event T-045 **Audit Event ID:** event-openai-board-expansion-2024-03-08

10.19.5. Appendix C5. 2024-05-28: OpenAI forms a board-level Safety and Security Committee

Event summary: OpenAI forms a board-level Safety and Security Committee **Event date:** 2024-05-28 **Source support:** Appendix A5 **Corpus companion entry:** Event T-048 **Audit Event ID:** event-openai-safety-committee-2024-05-28

10.19.6. Appendix C6. 2024-06-14: OpenAI appoints Paul Nakasone to its board and safety committee

Event summary: OpenAI appoints Paul Nakasone to its board and safety committee **Event date:** 2024-06-14 **Source support:** Appendix A2 **Corpus companion entry:** Event T-049 **Audit Event ID:** event-openai-nakasone-board-2024-06-14

10.19.7. Appendix C7. 2024-08-05: Musk files a federal complaint against Altman and OpenAI entities

Event summary: Musk files a federal complaint against Altman and OpenAI entities **Event date:** 2024-08-05 **Source support:** Appendix A17 **Corpus companion entry:** Event T-051 **Audit Event ID:** event-musk-v-altman-docket-2024-08-05

10.19.8. Appendix C8. 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 A9 **Corpus companion entry:** Event T-059 **Audit Event ID:** event-palantir-anthropic-aws-government-c

10.19.9. Appendix C9. 2025-01-14: OpenAI adds Adebayo Ogunlesi to its board during restructuring

Event summary: OpenAI adds Adebayo Ogunlesi to its board during restructuring **Event date:** 2025-01-14 **Source support:** Appendix A6 **Corpus companion entry:** Event T-064 **Audit Event ID:** event-openai-ogunlesi-board-2025-01-14

10.19.10. Appendix C10. 2025-03-05: Judge denies Musk's immediate bid to block OpenAI's conversion path and offers an exp...

Event summary: Judge denies Musk's immediate bid to block OpenAI's conversion path and offers an expedited trial track **Event date:** 2025-03-05 **Source support:** Appen-

dix A18, Appendix A19 **Corpus companion entry:** Event T-071 **Audit Event ID:** event-musk-openai-injunction-denied-2025-03-05

10.19.11. Appendix C11. 2025-03-13: OpenAI publishes its proposals for the U.S. AI Action Plan

Event summary: OpenAI publishes its proposals for the U.S. AI Action Plan **Event date:** 2025-03-13 **Source support:** Appendix A20 **Corpus companion entry:** Event T-072 **Audit Event ID:** event-openai-action-plan-2025-03-13

10.19.12. Appendix C12. 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

10.19.13. Appendix C13. 2025-06-13: The U.S. Army launches Detachment 201 and commissions senior tech executives into a r...

Event summary: The U.S. Army launches Detachment 201 and commissions senior tech executives into a reserve innovation corps **Event date:** 2025-06-13 **Source support:** Appendix A10 **Corpus companion entry:** Event T-078 **Audit Event ID:** event-army-detachment-201-2025-06-13

10.19.14. Appendix C14. 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

10.19.15. Appendix C15. 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 A12 **Corpus companion entry:** Event T-086 **Audit Event ID:** event-dario-adolescence-technology-2026-01-01

10.19.16. Appendix C16. 2026-03-31: FEC data shows large-scale activity for Leading the Future in the 2025-2026 cycle

Event summary: FEC data shows large-scale activity for Leading the Future in the 2025-2026 cycle **Event date:** 2026-03-31 **Source support:** Appendix A11 **Corpus companion entry:** Event T-088 **Audit Event ID:** event-fec-leading-the-future-overview-2026-03-31

10.19.17. Appendix C17. 2026-04-21: Axios reports record lobbying quarters for Anthropic and OpenAI

Event summary: Axios reports record lobbying quarters for Anthropic and OpenAI **Event date:** 2026-04-21 **Source support:** Appendix A3 **Corpus companion entry:** Event T-089 **Audit Event ID:** event-anthropic-openai-lobbying-2026-04-21

10.20. Appendix D. Relevant Entities

10.20.1. 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_legitimacy_capture and others. **Relevant sources in this paper:** Appendix A3, Appendix A8, Appendix A9 **Corpus companion entry:** Entity E-008 **Audit Entity ID:** org-anthropic

10.20.2. Appendix D2. 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 A1, Appendix A20, Appendix A5, Appendix A2, Appendix A17, Appendix A19, Appendix A18, Appendix A6, Appendix A3, Appendix A7 **Corpus companion entry:** Entity E-027 **Audit Entity ID:** org-openai

10.20.3. Appendix D3. 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 A9, Appendix A10 **Corpus companion entry:** Entity E-029 **Audit Entity ID:** org-palantir

10.20.4. Appendix D4. Federal Election Commission

Entity type: government **Role in this paper:** Actor relevant to 08_power_networks_legitimacy_capture. **Relevant sources in this paper:** Appendix A11 **Corpus companion entry:** Entity E-016 **Audit Entity ID:** org-fec

10.20.5. Appendix D5. 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

10.20.6. Appendix D6. 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 A10

Corpus companion entry: Entity E-035 **Audit Entity ID:** org-us-army

10.20.7. Appendix D7. Adebayo Ogunlesi

Entity type: person **Role in this paper:** Actor relevant to 08_power_networks_legitimacy_capture.
Relevant sources in this paper: Appendix A6 **Corpus companion entry:** Entity E-041
Audit Entity ID: person-adebayo-ogunlesi

10.20.8. Appendix D8. Bret Taylor

Entity type: person **Role in this paper:** Actor relevant to 08_power_networks_legitimacy_capture,
13_openai_governance_qstar_talent_exodus, 14_named_individuals_disputes_and_risk. **Rel-**
evant sources in this paper: Appendix A4 **Corpus companion entry:** Entity E-043 **Audit**
Entity ID: person-bret-taylor

10.20.9. Appendix D9. 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 A12 **Corpus companion entry:** Entity E-044 **Audit Entity ID:**
person-dario-amodei

10.20.10. Appendix D10. Elon Musk

Entity type: person **Role in this paper:** Actor relevant to 14_named_individuals_disputes_and_risk,
08_power_networks_legitimacy_capture. **Relevant sources in this paper:** Appendix A17,
Appendix A19, Appendix A18 **Corpus companion entry:** Entity E-045 **Audit Entity ID:**
person-elon-musk

10.20.11. Appendix D11. Fidji Simo

Entity type: person **Role in this paper:** Actor relevant to 08_power_networks_legitimacy_capture.
Relevant sources in this paper: Appendix A4 **Corpus companion entry:** Entity E-046
Audit Entity ID: person-fidji-simo

10.20.12. Appendix D12. Greg Brockman

Entity type: person **Role in this paper:** Actor relevant to 13_openai_governance_qstar_talent_exodus.
Relevant sources in this paper: Appendix A17 **Corpus companion entry:** Entity E-047
Audit Entity ID: person-greg-brockman

10.20.13. Appendix D13. Jakub Pachocki

Entity type: person **Role in this paper:** Actor relevant to 13_openai_governance_qstar_talent_exodus.
Relevant sources in this paper: Appendix A5 **Corpus companion entry:** Entity E-050
Audit Entity ID: person-jakub-pachocki

10.20.14. Appendix D14. John Schulman

Entity type: person **Role in this paper:** Actor relevant to 13_openai_governance_qstar_talent_exodus.
Relevant sources in this paper: Appendix A5 **Corpus companion entry:** Entity E-051
Audit Entity ID: person-john-schulman

10.20.15. Appendix D15. Larry Summers

Entity type: person **Role in this paper:** Actor relevant to 08_power_networks_legitimacy_capture,
13_openai_governance_qstar_talent_exodus. **Relevant sources in this paper:** Appendix A4
Corpus companion entry: Entity E-052 **Audit Entity ID:** person-larry-summers

10.20.16. Appendix D16. 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 A15, Appendix A16 **Corpus companion entry:**
Entity E-054 **Audit Entity ID:** person-mark-zuckerberg

10.20.17. Appendix D17. Mira Murati

Entity type: person **Role in this paper:** Actor relevant to 13_openai_governance_qstar_talent_exodus.
Relevant sources in this paper: Appendix A1 **Corpus companion entry:** Entity E-055
Audit Entity ID: person-mira-murati

10.20.18. Appendix D18. Nicole Seligman

Entity type: person **Role in this paper:** Actor relevant to 08_power_networks_legitimacy_capture.
Relevant sources in this paper: Appendix A4 **Corpus companion entry:** Entity E-056
Audit Entity ID: person-nicole-seligman

10.20.19. Appendix D19. Paul Nakasone

Entity type: person **Role in this paper:** Actor relevant to 07_export_controls_compute_defense_access,
08_power_networks_legitimacy_capture, 13_openai_governance_qstar_talent_exodus. **Rele-**
vant sources in this paper: Appendix A2, Appendix A5 **Corpus companion entry:** Entity
E-057 **Audit Entity ID:** person-paul-nakasone

10.20.20. Appendix D20. 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: This entity is contextual in this paper rather than source-
leading. **Corpus companion entry:** Entity E-058 **Audit Entity ID:** person-peter-thiel

10.20.21. Appendix D21. 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 A1, Appendix A20, Appendix A17, Appendix A19, Appendix A18, Appendix A14 **Corpus companion entry:** Entity E-059 **Audit Entity ID:** person-sam-altman

10.20.22. Appendix D22. Sue Desmond-Hellmann

Entity type: person **Role in this paper:** Actor relevant to 08_power_networks_legitimacy_capture. **Relevant sources in this paper:** Appendix A4 **Corpus companion entry:** Entity E-060 **Audit Entity ID:** person-sue-desmond-hellmann

10.20.23. Appendix D23. 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 A11 **Corpus companion entry:** Entity E-024 **Audit Entity ID:** org-leading-the-future

10.21. Appendix E. Evidence Boundaries

10.21.1. Appendix E1. Disputed Matters

- Appendix B66: That the precise degree of direct OpenAI alignment behind LEADING THE FUTURE remains contested beyond the current FEC filing record and should remain bounded to the Axios reporting layer.
- Appendix B67: That the public record does not settle whether Q* or any specific capability breakthrough directly triggered the November 2023 OpenAI board rupture.
- Appendix B68: That allegations advanced by named actors in OpenAI governance disputes are procedurally live and politically relevant without being established as true.

10.21.2. Appendix E2. Interpretive Boundaries

- Appendix B69: 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 B70: That frontier legitimacy is increasingly stabilized through a combination of mission rhetoric, elite board composition, security credentials, policy access, and political spending.
- Appendix B71: 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 B72: That the strongest bounded reading of elite fallback architecture is a revealed-preference gap between public stewardship language and private survivability hedging.
- Appendix B73: That OpenAI's governance rupture, safety rhetoric, leadership churn, and reasoning-model rollout are best read as parts of one reorganization of authority rather than as

isolated episodes.

- Appendix B74: That named-individual disputes should enter the dossier only when they clarify governance structure, institutional control, or legal architecture rather than functioning as scandal collage.

10.21.3. Appendix E3. Speculative Narrative Risks

- Appendix B75: That bunker ownership, survival planning, or retreat architecture prove elite foreknowledge of imminent collapse or that frontier firms are knowingly engineering it.
- Appendix B76: That elite overlap among frontier labs, venture capital, lobbying networks, and defense-adjacent institutions proves a single hidden command chain.
- Appendix B77: That founder philanthropy or public warnings about concentration prove elite foreknowledge of imminent collapse or absolve frontier firms of responsibility.
- Appendix B78: 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 B79: That elite overlap, litigation conflict, or reputational controversy around named individuals justifies guilt by association without direct structural proof.

10.21.4. 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 A10. Companion note: Note N-069. Risk flag: low. Note focus: That by June 13, 2025 the U.S. Army had created a formal structure to embed senior commercial-tech executives in a military innovation role.
- 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 A20. Companion note: Note N-073. Risk flag: medium. Note focus: That OpenAI officially defended a public-policy agenda combining global commercial expansion, export controls, and preservation of access to copyrighted material for training.
- Appendix A9. 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 A2. Companion note: Note N-079. Risk flag: review. Note focus: That OpenAI brought a former top U.S. military and intelligence official directly into its governance structure.
- Appendix A3. Companion note: Note N-081. Risk flag: medium. Note focus: That frontier AI firms were spending meaningful and rising sums on federal lobbying by early 2026.
- Appendix A12. 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 A11. Companion note: Note N-085. Risk flag: low. Note focus: That there is an official federal record for **LEADING THE FUTURE** as an active Super PAC in the 2025-2026 election

cycle.

- Appendix A6. Companion note: Note N-087. Risk flag: medium. Note focus: That OpenAI continued to load its board with elite finance and infrastructure figures during its restructuring phase.
- Appendix A14. 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 A1. Companion note: Note N-091. Risk flag: medium. Note focus: That OpenAI publicly treated its crisis as a governance and legitimacy rupture, not as a simple executive reshuffle.
- Appendix A4. Companion note: Note N-092. Risk flag: low. Note focus: That OpenAI deliberately rebuilt its board with figures carrying regulatory, legal, nonprofit, and major-platform credentials.
- Appendix A20. Companion note: Note N-093. Risk flag: medium. Note focus: That OpenAI presents itself as a policy actor seeking to co-design national AI strategy.
- Appendix A16. Companion note: Note N-094. Risk flag: review. Note focus: That Zuckerberg publicly minimized the compound's bunker framing after the earlier investigation.
- Appendix A13. 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 A15. 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 A5. Companion note: Note N-138. Risk flag: review. Note focus: That OpenAI publicly moved safety and security oversight to a named board-level structure after the governance crisis.
- Appendix A1. Companion note: Note N-140. Risk flag: medium. Note focus: That the November 2023 crisis was recognized by OpenAI itself as a formal governance rupture.
- Appendix A18. Companion note: Note N-144. Risk flag: medium. Note focus: That by March 5, 2025 Musk had failed to obtain immediate judicial relief blocking OpenAI's conversion path.
- Appendix A17. Companion note: Note N-148. Risk flag: medium. Note focus: That the Musk-OpenAI dispute had a concrete federal-court vehicle with named defendants and an active procedural history.
- Appendix A19. Companion note: Note N-149. Risk flag: medium. Note focus: That by March 5, 2025 Musk had failed to obtain immediate judicial relief blocking OpenAI's conversion path.

11. Chapter 09. Future if no action is taken

11.1. Abstract

The most credible continuity scenario is not science-fiction apocalypse. The U.S. If no action is taken, platform immunity, private control over weights, and opaque state-company entanglement stop looking like separate policy issues and start looking like one reinforced system.

11.2. Keywords

Future if no action is taken; surveillance capitalism; AI imperialism; evidence-based dossier

11.3. Research Question

How does future if no action is taken function within the dossier's larger argument, and what does the documentary record allow this chapter to establish?

11.4. Scope and Framing Note

This paper examines future if no action is taken 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 explains why inaction would compound dependency, hierarchy, and cognitive enclosure across the whole system. Adjacent reinforcement appears in Chapter 05, Chapter 07, Chapter 08, Chapter 10, but those cross-references are supportive rather than required for basic comprehension.

11.5. Central Thesis

The most credible continuity scenario is not science-fiction apocalypse. It is a slower regime of deepening dependence, shrinking cognitive sovereignty, weakened middle-class bargaining power, and increasingly hierarchical access to machine intelligence.¹⁰¹

11.6. Evidence and Method Note

This paper is derived from the chapter corpus for Chapter 09 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.

11.7. Introduction

This paper asks: How does future if no action is taken function within the dossier's larger argument, and what does the documentary record allow this chapter to establish? Its working answer is clear

¹⁰¹Sources: U.S. Department of Education / ERIC, Artificial Intelligence and the Futur... (2023-05-24; see Appendix A1); arXiv, GPTs are GPTs: An Early Look at the L... (2023-03-17; see Appendix A2); arXiv, Computing Power and the Governance of... (2024-02-13; see Appendix A3)

from the start: The most credible continuity scenario is not science-fiction apocalypse. It is a slower regime of deepening dependence, shrinking cognitive sovereignty, weakened middle-class bargaining power, and increasingly hierarchical access to machine intelligence.¹⁰² 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.

11.8. Historical or Institutional Context

The U.S. Department of Education has already framed AI in learning environments as a governance issue involving automated decisions, bias, privacy, and human oversight. Dependency concerns are therefore not merely cultural panic.¹⁰³

11.9. Main Analysis

The GPTs paper and the IMF note both support an uneven future in which exposure to generative AI is distributed asymmetrically across occupations, income positions, and labor systems. The future problem is not uniform extinction. It is differential vulnerability.¹⁰⁴

Compute-governance research supports the claim that high-end capability is likely to remain infrastructure-concentrated unless policy changes the terms of access. That means unequal access to machine intelligence is already an infrastructure question, not only a talent question.¹⁰⁵

The continuity scenario also includes unequal fallback. Some of the same actors asking institutions and publics to normalize deeper dependence on AI and platform infrastructure have also publicly described or been credibly reported as building private redundancy, jurisdictional exit, and survivability options: Altman's 2016 prep disclosures and New Zealand fallback, Zuckerberg's Kauai shelter complex, and Thiel's long association with New Zealand retreat logic and elite-exit politics. That does not prove foreknowledge, coordinated collapse engineering, or a hidden command chain. It does show how continuity can stratify not only access to intelligence, but access to backup, redundancy, and escape when shared systems fail or become politically untrustworthy.¹⁰⁶

The financial and moral-warning layers also matter. The BIS warned of fragile capex dynamics and weak near-term returns in the AI buildout, while Dario Amodi publicly warned that AI-era wealth concentration could strain democratic society. These are not identical claims, but together they strengthen the dossier's continuity scenario of hierarchy rather than broad emancipation.¹⁰⁷

¹⁰²Sources: U.S. Department of Education / ERIC, Artificial Intelligence and the Futur... (2023-05-24; see Appendix A1); arXiv, GPTs are GPTs: An Early Look at the L... (2023-03-17; see Appendix A2); arXiv, Computing Power and the Governance of... (2024-02-13; see Appendix A3)

¹⁰³Sources: U.S. Department of Education / ERIC, Artificial Intelligence and the Futur... (2023-05-24; see Appendix A1)

¹⁰⁴Sources: arXiv, GPTs are GPTs: An Early Look at the L... (2023-03-17; see Appendix A2); IMF, Gen-AI: Artificial Intelligence and t... (2024-01-14; see Appendix A4)

¹⁰⁵Sources: arXiv, Computing Power and the Governance of... (2024-02-13; see Appendix A3)

¹⁰⁶Sources: Tad Friend, Sam Altman's Manifest Destiny (2016-10-03; see Appendix A5); Makena Kelly, Inside Mark Zuckerberg's Top-Secret H... (2023-12-14; see Appendix A6); Charlotte Phillipp, Mark Zuckerberg Downplays His 5,000-S... (2024-12-20; see Appendix A7); Mark O'Connell, Why Silicon Valley billionaires are p... (2018-02-15; see Appendix A8); Rebecca Rose, New Zealand woos the wealthy with 'go... (2025-09-24; see Appendix A9)

¹⁰⁷Sources: Bank for International Settlements, Annual Economic Report 2026 (2026-06-28; see Appendix A10); Dario Amodi, The Adolescence of Technology (2026-01-01; see Appendix A11)

Interpretation: if current platform habits, labor structures, weights enclosure, compute chokepoints, and asymmetric fallback systems continue reinforcing one another, the likely future is not total machine takeover but a regime of dependency, access hierarchy, and weaker cognitive sovereignty. A contextual formulation from Dion Wiggins is useful here: continuity can disguise dependency when systems still appear to function while effective authority migrates elsewhere. That is precisely why the dossier’s three reforms belong together rather than as isolated policy tweaks.¹⁰⁸

11.10. Position Within the Series

This paper sits inside a coordinated 11-paper architecture. Its nearest companions are Chapter 05 (Talent, overhiring, valuation, and layoffs), Chapter 07 (Imperial control over access), Chapter 08 (Power network, legitimacy, and institutional capture), Chapter 10 (What to do). In that series logic, Chapter 09 explains why inaction would compound dependency, hierarchy, and cognitive enclosure across the whole system.

11.11. Counterarguments

This chapter must remain explicit that it is a projection built from documented trend lines, not a prophecy. The evidence is stronger on labor exposure, governance risk, and access concentration than on long-run learning outcomes or a single inevitable end state.¹⁰⁹

11.12. Limits of the Evidence

This chapter must remain explicit that it is a projection built from documented trend lines, not a prophecy. The evidence is stronger on labor exposure, governance risk, and access concentration than on long-run learning outcomes or a single inevitable end state.¹¹⁰

11.13. Reform Relevance

If no action is taken, platform immunity, private control over weights, and opaque state-company entanglement stop looking like separate policy issues and start looking like one reinforced system. That is why **Refactor Section 230**, **AI Weights as Patrimony of Humanity**, and **Limit Government Capture by AI Companies** must be read as one program of democratic reassertion.

¹¹¹

¹⁰⁸Sources: FTC, A Look Behind the Screens: Examining... (n.d.; see Appendix A12); Meta, Mark Zuckerberg’s Message to Meta Emp... (2023-03-14; see Appendix A13); arXiv, Computing Power and the Governance of... (2024-02-13; see Appendix A3); Dario Amodei, The Adolescence of Technology (2026-01-01; see Appendix A11); Dion Wiggins, The Digital Sovereignty Imperative (n.d.; see Appendix A14)

¹⁰⁹Sources: U.S. Department of Education / ERIC, Artificial Intelligence and the Futur... (2023-05-24; see Appendix A1); IMF, Gen-AI: Artificial Intelligence and t... (2024-01-14; see Appendix A4)

¹¹⁰Sources: U.S. Department of Education / ERIC, Artificial Intelligence and the Futur... (2023-05-24; see Appendix A1); IMF, Gen-AI: Artificial Intelligence and t... (2024-01-14; see Appendix A4)

¹¹¹Sources: FTC, A Look Behind the Screens: Examining... (n.d.; see Appendix A12); OpenAI, OpenAI’s proposals for the U.S. AI Ac... (2025-03-13; see Appendix A15); arXiv, Computing Power and the Governance of... (2024-02-13; see Appendix A3)

11.14. Conclusion

If no action is taken, platform immunity, private control over weights, and opaque state-company entanglement stop looking like separate policy issues and start looking like one reinforced system. The final chapter therefore has to answer the question the whole dossier implies: what does resistance look like once the solution space is no longer open-ended but doctrinally explicit? In the architecture of this series, Chapter 09 therefore functions as a self-contained argument while also advancing the cumulative path toward the dossier's three reforms.

11.15. 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 09, those boundaries remain visible in the prose, in the bibliography, and in the appendix sections that summarize claims and caution points.

11.16. 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 a primary U.S. education authority already treated AI-mediated learning as a governance question rather than as a neutral productivity upgrade.
2. arXiv. GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models. 2023-03-17. T1 paper. Paper appendix: Appendix A2. Corpus companion entry: Source S-065. Audit ID: **gpts-are-gpts-labor-impact-2023**. Relevance: That generative AI exposure reaches well beyond narrow routine clerical work and extends into professional and white-collar task bundles.
3. arXiv. Computing Power and the Governance of Artificial Intelligence. 2024-02-13. T1 paper. Paper appendix: Appendix A3. Corpus companion entry: Source S-040. Audit ID: **compute-governance-ai-2024**. Relevance: That access hierarchy in frontier AI is materially tied to compute concentration and not only to abstract innovation differences.
4. IMF. Gen-AI: Artificial Intelligence and the Future of Work. 2024-01-14. T1 report. Paper appendix: Appendix A4. Corpus companion entry: Source S-071. Audit ID: **imf-gen-ai-future-of-work-2024**. Relevance: That labor-market exposure to generative AI is systematically uneven across geography, occupation, and class position.
5. Tad Friend. Sam Altman's Manifest Destiny. 2016-10-03. T2 profile. Paper appendix: Appendix A5. 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, potassium iodide, antibiotics, batteries, water, gas masks, and land in Big Sur as part of survival pla...
6. Makena Kelly. Inside Mark Zuckerberg's Top-Secret Hawaii Compound. 2023-12-14. T2 investigation. Paper appendix: Appendix A6. 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...
7. Charlotte Phillipp. Mark Zuckerberg Downplays His 5,000-Square-Foot Underground Bunker in Hawaii: ‘Just a Little Shelter’. 2024-12-20. T3 news_article. Paper appendix: Appendix A7. Corpus companion entry: Source S-106. Audit ID: **people-zuckerberg-little-shelter-2024**. Relevance: That by December 20, 2024 Mark Zuckerberg publicly minimized the Kauai bunker as ‘just a little shelter,’ creating a documented contrast between the estate’s described features...
 8. Mark O’Connell. Why Silicon Valley billionaires are prepping for the apocalypse in New Zealand. 2018-02-15. T3 longform_feature. Paper appendix: Appendix A8. Corpus companion entry: Source S-066. Audit ID: **guardian-billionaires-new-zealand-2018**. Relevance: That by February 15, 2018 New Zealand had already become a recognizable elite-safe-haven trope in public collapse-prep discourse around Silicon Valley.
 9. Rebecca Rose. New Zealand woos the wealthy with ‘golden visa’ changes. 2025-09-24. T2 news_article. Paper appendix: Appendix A9. 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.
 10. Bank for International Settlements. Annual Economic Report 2026. 2026-06-28. T1 report. Paper appendix: Appendix A10. Corpus companion entry: Source S-036. Audit ID: **bis-annual-economic-report-2026**. Relevance: That a primary multilateral financial institution explicitly warned in June 2026 that the AI capex race was outrunning earnings and free cash flow.
 11. Dario Amodei. The Adolescence of Technology. 2026-01-01. T3 essay. Paper appendix: Appendix A11. 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.
 12. FTC. A Look Behind the Screens: Examining the Data Practices of Social Media and Video Streaming Services. n.d.. T1 report. Paper appendix: Appendix A12. Corpus companion entry: Source S-063. Audit ID: **ftc-2024-look-behind-screens**. Relevance: Supports a key factual or interpretive move inside this paper.
 13. Meta. Mark Zuckerberg’s Message to Meta Employees. 2023-03-14. T1 company_statement. Paper appendix: Appendix A13. Corpus companion entry: Source S-082. Audit ID: **meta-year-of-efficiency-2023**. Relevance: Supports a key factual or interpretive move inside this paper.
 14. Dion Wiggins. The Digital Sovereignty Imperative. n.d.. T3 website_homepage. Paper appendix: Appendix A14. Corpus companion entry: Source S-050. Audit ID: **digital-sovereignty-imperative-homepage-2026**. Relevance: That Dion Wiggins publicly defines digital sovereignty in survivability terms: whether systems still run under your authority when cooperation is withdrawn.
 15. OpenAI. OpenAI’s proposals for the U.S. AI Action Plan. 2025-03-13. T1 policy_submission. Paper appendix: Appendix A15. Corpus companion entry: Source S-096. Audit ID: **openai-us-ai-action-plan-2025**. Relevance: 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 acc...

11.17. Appendix A. Source Register for This Paper

11.17.1. 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 a primary U.S. education authority already treated AI-mediated learning as a governance question rather than as a neutral productivity upgrade. **Corpus companion entry:** Source S-124 **Audit ID:** us-ed-ai-future-teaching-learning-2023

11.17.2. Appendix A2. GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models

Source citation: GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models **Institution / author:** arXiv **Date:** 2023-03-17 **Tier / type:** T1 / paper **Why it matters in this paper:** That generative AI exposure reaches well beyond narrow routine clerical work and extends into professional and white-collar task bundles. **Corpus companion entry:** Source S-065 **Audit ID:** gpts-are-gpts-labor-impact-2023

11.17.3. Appendix A3. Computing Power and the Governance of Artificial Intelligence

Source citation: Computing Power and the Governance of Artificial Intelligence **Institution / author:** arXiv **Date:** 2024-02-13 **Tier / type:** T1 / paper **Why it matters in this paper:** That access hierarchy in frontier AI is materially tied to compute concentration and not only to abstract innovation differences. **Corpus companion entry:** Source S-040 **Audit ID:** compute-governance-ai-2024

11.17.4. Appendix A4. Gen-AI: Artificial Intelligence and the Future of Work

Source citation: Gen-AI: Artificial Intelligence and the Future of Work **Institution / author:** IMF **Date:** 2024-01-14 **Tier / type:** T1 / report **Why it matters in this paper:** That labor-market exposure to generative AI is systematically uneven across geography, occupation, and class position. **Corpus companion entry:** Source S-071 **Audit ID:** imf-gen-ai-future-of-work-2024

11.17.5. Appendix A5. 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 pla... **Corpus companion entry:** Source S-083 **Audit ID:** newyorker-sam-altman-manifest-destiny-2016

11.17.6. Appendix A6. 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 mechani... **Corpus companion entry:** Source S-133 **Audit ID:** wired-zuckerberg-hawaii-compound-2023

11.17.7. Appendix A7. Mark Zuckerberg Downplays His 5,000-Square-Foot Underground Bunker in Hawaii: ‘Just a Little Shelter’

Source citation: Mark Zuckerberg Downplays His 5,000-Square-Foot Underground Bunker in Hawaii: ‘Just a Little Shelter’ **Institution / author:** Charlotte Phillipp **Date:** 2024-12-20 **Tier / type:** T3 / news_article **Why it matters in this paper:** That by December 20, 2024 Mark Zuckerberg publicly minimized the Kauai bunker as ‘just a little shelter,’ creating a documented contrast between the estate’s described features... **Corpus companion entry:** Source S-106 **Audit ID:** people-zuckerberg-little-shelter-2024

11.17.8. Appendix A8. Why Silicon Valley billionaires are prepping for the apocalypse in New Zealand

Source citation: Why Silicon Valley billionaires are prepping for the apocalypse in New Zealand **Institution / author:** Mark O’Connell **Date:** 2018-02-15 **Tier / type:** T3 / longform_feature **Why it matters in this paper:** That by February 15, 2018 New Zealand had already become a recognizable elite-safe-haven trope in public collapse-prep discourse around Silicon Valley. **Corpus companion entry:** Source S-066 **Audit ID:** guardian-billionaires-new-zealand-2018

11.17.9. Appendix A9. 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

11.17.10. Appendix A10. Annual Economic Report 2026

Source citation: Annual Economic Report 2026 **Institution / author:** Bank for International Settlements **Date:** 2026-06-28 **Tier / type:** T1 / report **Why it matters in this paper:** That a primary multilateral financial institution explicitly warned in June 2026 that the AI capex race was outrunning earnings and free cash flow. **Corpus companion entry:** Source S-036 **Audit ID:** bis-annual-economic-report-2026

11.17.11. Appendix A11. 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

11.17.12. Appendix A12. 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:** Supports a key factual or interpretive move inside this paper. **Corpus companion entry:** Source S-063 **Audit ID:** ftc-2024-look-behind-screens

11.17.13. Appendix A13. Mark Zuckerberg’s Message to Meta Employees

Source citation: Mark Zuckerberg’s Message to Meta Employees **Institution / author:** Meta **Date:** 2023-03-14 **Tier / type:** T1 / company_statement **Why it matters in this paper:** Supports a key factual or interpretive move inside this paper. **Corpus companion entry:** Source S-082 **Audit ID:** meta-year-of-efficiency-2023

11.17.14. Appendix A14. The Digital Sovereignty Imperative

Source citation: The Digital Sovereignty Imperative **Institution / author:** Dion Wiggins **Date:** n.d. **Tier / type:** T3 / website_homepage **Why it matters in this paper:** That Dion Wiggins publicly defines digital sovereignty in survivability terms: whether systems still run under your authority when cooperation is withdrawn. **Corpus companion entry:** Source S-050 **Audit ID:** digital-sovereignty-imperative-homepage-2026

11.17.15. Appendix A15. 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 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 acc... **Corpus companion entry:** Source S-096 **Audit ID:** openai-us-ai-action-plan-2025

11.18. Appendix B. Claims Used in This Paper

11.18.1. Appendix B1. A primary multilateral financial institution explicitly warned in June 2026 that...

Claim statement: That a primary multilateral financial institution explicitly warned in June 2026 that the AI capex race was outrunning earnings and free cash flow. **Evidence status:** Documented Fact **Source support:** Appendix A10 **Corpus companion entry:** Claim C-107 **Audit Claim ID:** bis-annual-economic-report-2026-05-talent-hoarding-section174-layoffs-claim-01

11.18.2. Appendix B2. Debt issuance, leverage, opaque financing structures, and circular financing below...

Claim statement: That debt issuance, leverage, opaque financing structures, and circular financing belong to the documented risk structure of the AI buildout. **Evidence status:** Documented Fact **Source support:** Appendix A10 **Corpus companion entry:** Claim C-108 **Audit Claim ID:** bis-annual-economic-report-2026-05-talent-hoarding-section174-layoffs-claim-02

11.18.3. Appendix B3. Chapter 05 can support the thesis of structural correction and financial overhang...

Claim statement: That Chapter 05 can support the thesis of structural correction and financial overhang without relying only on opinion pieces or market chatter. **Evidence status:** Documented Fact **Source support:** Appendix A10 **Corpus companion entry:** Claim C-109 **Audit Claim ID:** bis-annual-economic-report-2026-05-talent-hoarding-section174-layoffs-claim-03

11.18.4. Appendix B4. Chapters 06, 09, and 10 can discuss the “cheap AI” narrative against a documented...

Claim statement: That Chapters 06, 09, and 10 can discuss the “cheap AI” narrative against a documented backdrop of subsidy-like financing pressure and fragile commercial returns. **Evidence status:** Documented Fact **Source support:** Appendix A10 **Corpus companion entry:** Claim C-110 **Audit Claim ID:** bis-annual-economic-report-2026-05-talent-hoarding-section174-layoffs-claim-

11.18.5. Appendix B5. Access hierarchy in frontier AI is materially tied to compute concentration and n...

Claim statement: That access hierarchy in frontier AI is materially tied to compute concentration and not only to abstract innovation differences. **Evidence status:** Documented Fact **Source support:** Appendix A3 **Corpus companion entry:** Claim C-120 **Audit Claim ID:** compute-governance-ai-2024-09-future-if-unchecked-claim-01

11.18.6. Appendix B6. Chapter 09 can speak about unequal access to high-end intelligence as a documente...

Claim statement: That Chapter 09 can speak about unequal access to high-end intelligence as a documented infrastructure problem. **Evidence status:** Documented Fact **Source support:** Appendix A3 **Corpus companion entry:** Claim C-121 **Audit Claim ID:** compute-governance-ai-2024-09-future-if-unchecked-claim-02

11.18.7. Appendix B7. The project’s language of chokepoints and access stratification has direct suppor...

Claim statement: That the project’s language of chokepoints and access stratification has direct support in primary research on compute governance. **Evidence status:** Documented Fact **Source support:** Appendix A3 **Corpus companion entry:** Claim C-122 **Audit Claim ID:** compute-governance-ai-2024-09-future-if-unchecked-claim-03

11.18.8. Appendix B8. 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 A11 **Corpus companion entry:** Claim C-141 **Audit**

Claim ID: dario-amodei-adolescence-of-technology-2026-08-power-networks-legitimacy-capture-claim-01

11.18.9. Appendix B9. 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 A11 **Corpus companion entry:** Claim C-142 **Audit**

Claim ID: dario-amodei-adolescence-of-technology-2026-08-power-networks-legitimacy-capture-claim-02

11.18.10. Appendix B10. 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 A11 **Corpus companion entry:** Claim C-143 **Audit**

Claim ID: dario-amodei-adolescence-of-technology-2026-08-power-networks-legitimacy-capture-claim-03

11.18.11. Appendix B11. 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 A11 **Corpus companion entry:** Claim C-144 **Audit**

Claim ID: dario-amodei-adolescence-of-technology-2026-08-power-networks-legitimacy-capture-claim-04

11.18.12. Appendix B12. 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 A11 **Corpus companion entry:** Claim C-145 **Audit**

Claim ID: dario-amodei-adolescence-of-technology-2026-08-power-networks-legitimacy-capture-claim-05

11.18.13. Appendix B13. Dion Wiggins publicly defines digital sovereignty in survivability terms: whether...

Claim statement: That Dion Wiggins publicly defines digital sovereignty in survivability terms: whether systems still run under your authority when cooperation is withdrawn. **Evidence status:** Documented Fact

Source support: Appendix A14 **Corpus companion entry:** Claim C-152 **Audit**

Claim ID: digital-sovereignty-imperative-homepage-2026-09-future-if-unchecked-claim-01

11.18.14. Appendix B14. 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 A9 **Corpus companion entry:** Claim C-182 **Audit Claim ID:** ft-new-zealand-golden-visa-2025-09-future-if-unchecked-claim-01

11.18.15. Appendix B15. 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 A9 **Corpus companion entry:** Claim C-183 **Audit Claim ID:** ft-new-zealand-golden-visa-2025-11-facebook-thiel-intelligence-a

11.18.16. Appendix B16. Generative AI exposure reaches well beyond narrow routine clerical work and exten...

Claim statement: That generative AI exposure reaches well beyond narrow routine clerical work and extends into professional and white-collar task bundles. **Evidence status:** Documented Fact **Source support:** Appendix A2 **Corpus companion entry:** Claim C-205 **Audit Claim ID:** gpts-are-gpts-labor-impact-2023-09-future-if-unchecked-claim-01

11.18.17. Appendix B17. A future scenario of labor stratification can be grounded in documented different...

Claim statement: That a future scenario of labor stratification can be grounded in documented differential exposure rather than in pure speculation. **Evidence status:** Documented Fact **Source support:** Appendix A2 **Corpus companion entry:** Claim C-206 **Audit Claim ID:** gpts-are-gpts-labor-impact-2023-09-future-if-unchecked-claim-02

11.18.18. Appendix B18. Chapter 09 can discuss weakening bargaining power for parts of the technical and...

Claim statement: That Chapter 09 can discuss weakening bargaining power for parts of the technical and professional middle class without claiming universal job extinction. **Evidence status:** Documented Fact **Source support:** Appendix A2 **Corpus companion entry:** Claim C-207 **Audit Claim ID:** gpts-are-gpts-labor-impact-2023-09-future-if-unchecked-claim-03

11.18.19. Appendix B19. By February 15, 2018 New Zealand had already become a recognizable elite-safe-hav...

Claim statement: That by February 15, 2018 New Zealand had already become a recognizable elite-safe-haven trope in public collapse-prep discourse around Silicon Valley. **Evidence status:**

Documented Fact **Source support:** Appendix A8 **Corpus companion entry:** Claim C-208 **Audit Claim ID:** guardian-billionaires-new-zealand-2018-09-future-if-unchecked-claim-01

11.18.20. Appendix B20. A serious 2018 longform outlet publicly connected Peter Thiel’s New Zealand citiz...

Claim statement: That a serious 2018 longform outlet publicly connected Peter Thiel’s New Zealand citizenship controversy, Sam Altman’s retreat plan, and The Sovereign Individual within a single elite-safe-haven narrative. **Evidence status:** Documented Fact **Source support:** Appendix A8 **Corpus companion entry:** Claim C-209 **Audit Claim ID:** guardian-billionaires-new-zealand-2018-11-facebook-thiel-intelligence-adjacency-claim-01

11.18.21. Appendix B21. Labor-market exposure to generative AI is systematically uneven across geography,...

Claim statement: That labor-market exposure to generative AI is systematically uneven across geography, occupation, and class position. **Evidence status:** Documented Fact **Source support:** Appendix A4 **Corpus companion entry:** Claim C-222 **Audit Claim ID:** imf-gen-ai-future-of-work-2024-09-future-if-unchecked-claim-01

11.18.22. Appendix B22. A future of stratified labor outcomes can be documented from mainstream instituti...

Claim statement: That a future of stratified labor outcomes can be documented from mainstream institutional analysis rather than inferred only from anecdote. **Evidence status:** Documented Fact **Source support:** Appendix A4 **Corpus companion entry:** Claim C-223 **Audit Claim ID:** imf-gen-ai-future-of-work-2024-09-future-if-unchecked-claim-02

11.18.23. Appendix B23. Chapter 09 can connect AI deployment to differentiated middle-class risk without...

Claim statement: That Chapter 09 can connect AI deployment to differentiated middle-class risk without reducing the issue to a single universal displacement story. **Evidence status:** Documented Fact **Source support:** Appendix A4 **Corpus companion entry:** Claim C-224 **Audit Claim ID:** imf-gen-ai-future-of-work-2024-09-future-if-unchecked-claim-03

11.18.24. Appendix B24. 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 A5 **Corpus companion entry:** Claim C-275 **Audit Claim ID:** newyorker-sam-altman-manifest-destiny-2016-08-power-networks-legitimacy-capture-claim

11.18.25. Appendix B25. 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 A5 **Corpus companion entry:** Claim C-276

Audit Claim ID: newyorker-sam-altman-manifest-destiny-2016-09-future-if-unchecked

11.18.26. Appendix B26. By December 20, 2024 Mark Zuckerberg publicly minimized the Kauai bunker as 'just...

Claim statement: That by December 20, 2024 Mark Zuckerberg publicly minimized the Kauai bunker as 'just a little shelter,' creating a documented contrast between the estate's described features and his own framing of them. **Evidence status:** Documented Fact

Source support: Appendix A7 **Corpus companion entry:** Claim C-352

Audit Claim ID: people-zuckerberg-little-shelter-2024-08-power-networks-legitimacy-capture-claim-01

11.18.27. Appendix B27. A primary U.S. education authority already treated AI-mediated learning as a gove...

Claim statement: That a primary U.S. education authority already treated AI-mediated learning as a governance question rather than as a neutral productivity upgrade. **Evidence status:** Documented Fact

Source support: Appendix A1 **Corpus companion entry:** Claim C-401

Audit Claim ID: us-ed-ai-future-teaching-learning-2023-09-future-if-unchecked-claim-01

11.18.28. Appendix B28. Institutional guidance had already identified bias, privacy, safety, and over-del...

Claim statement: That institutional guidance had already identified bias, privacy, safety, and over-delegation as central risks in educational AI adoption. **Evidence status:** Documented Fact

Source support: Appendix A1 **Corpus companion entry:** Claim C-402

Audit Claim ID: us-ed-ai-future-teaching-learning-2023-09-future-if-unchecked-claim-02

11.18.29. Appendix B29. Chapter 09 can ground dependency and cognitive-mediation concerns in official edu...

Claim statement: That Chapter 09 can ground dependency and cognitive-mediation concerns in official educational governance language. **Evidence status:** Documented Fact

Source support: Appendix A1 **Corpus companion entry:** Claim C-403

Audit Claim ID: us-ed-ai-future-teaching-learning-2023-09-future-if-unchecked-claim-03

11.18.30. Appendix B30. 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 A6 **Corpus companion entry:** Claim C-489 **Audit Claim ID:** wired-zuckerberg-hawaii-compound-2023-08-power-networks-legitimacy-capture-claim-01

11.18.31. Appendix B31. 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 A6 **Corpus companion entry:** Claim C-490 **Audit Claim ID:** wired-zuckerberg-hawaii-compound-2023-08-power-networks-legitimacy-capture-claim-02

11.18.32. Appendix B32. 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 A15 **Corpus companion entry:** Claim C-445 **Audit Claim ID:** vector08-legitimacy-through-operational-state-interfaces-claim-01

11.18.33. Appendix B33. 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 A5, Appendix A6, Appendix A7, Appendix A8, Appendix A9, Appendix A14 **Corpus companion entry:** Claim C-449 **Audit Claim ID:** vector09-asymmetric-fallback-continuity-claim-01

11.18.34. Appendix B34. 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 A2, Appendix A4, Appendix A3, Appendix A11 **Corpus companion entry:** Claim C-450 **Audit Claim ID:** vector09-continuity-scenario-dependency-hierarchy-claim-01

11.18.35. Appendix B35. 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 A5, Appendix A6, Appendix A7, Appendix A8, Appendix A9 **Corpus companion entry:** Claim C-440 **Audit Claim ID:** vector08-bunker-foreknowledge-overclaim-claim-01

11.18.36. Appendix B36. 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 A11 **Corpus companion entry:** Claim C-446 **Audit Claim ID:** vector08-philanthropy-foreknowledge-overclaim-claim-01

11.18.37. Appendix B37. 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 A10 **Corpus companion entry:** Claim C-451 **Audit Claim ID:** vector09-inevitable-collapse-overclaim-claim-01

11.18.38. Appendix B38. Peter Thiel's New Zealand retreat logic, ideological associations, and political...

Claim statement: That Peter Thiel's New Zealand retreat logic, ideological associations, and political ties prove a unified hidden command chain linking platform governance, Palantir, and civilizational collapse planning. **Evidence status:** Speculative Narrative Risk **Source support:** Appendix A8, Appendix A9 **Corpus companion entry:** Claim C-460 **Audit Claim ID:** vector11-prepper-command-overclaim-claim-01

11.19. Appendix C. Timeline Slice

11.19.1. 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 A5 **Corpus companion entry:** Event T-004 **Audit Event ID:** event-altman-prep-new-zealand-2016-10-03

11.19.2. Appendix C2. 2023-03-17: OpenAI and OpenResearch authors publish GPTs are GPTs on occupational exposure to LLMs

Event summary: OpenAI and OpenResearch authors publish GPTs are GPTs on occupational exposure to LLMs **Event date:** 2023-03-17 **Source support:** Appendix A2 **Corpus companion entry:** Event T-023 **Audit Event ID:** event-gpts-are-gpts-2023-03-17

11.19.3. 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

11.19.4. 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 A6 **Corpus companion entry:** Event T-035 **Audit Event ID:** event-wired-zuckerberg-kauai-compound-2023-12-14

11.19.5. Appendix C5. 2024-01-14: The IMF publishes Gen-AI: Artificial Intelligence and the Future of Work

Event summary: The IMF publishes Gen-AI: Artificial Intelligence and the Future of Work **Event date:** 2024-01-14 **Source support:** Appendix A4 **Corpus companion entry:** Event T-036 **Audit Event ID:** event-imf-gen-ai-work-2024-01-14

11.19.6. Appendix C6. 2024-02-13: Researchers publish Computing Power and the Governance of Artificial Intelligence on...

Event summary: Researchers publish Computing Power and the Governance of Artificial Intelligence on arXiv **Event date:** 2024-02-13 **Source support:** Appendix A3 **Corpus companion entry:** Event T-042 **Audit Event ID:** event-compute-governance-paper-2024-02-13

11.19.7. Appendix C7. 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 A11 **Corpus companion entry:** Event T-086 **Audit Event ID:** event-dario-adolescence-technology-2026-01-01

11.20. Appendix D. Relevant Entities

11.20.1. Appendix D1. 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

11.20.2. Appendix D2. AI compute chokepoints

Entity type: infrastructure **Role in this paper:** Actor relevant to 09_future_if_unchecked, 10_counterstrategy_response_agenda. **Relevant sources in this paper:** Appendix A3 **Corpus companion entry:** Entity E-002 **Audit Entity ID:** concept-ai-compute-chokepoint

11.20.3. Appendix D3. International Monetary Fund

Entity type: multilateral **Role in this paper:** Actor relevant to 09_future_if_unchecked. **Relevant sources in this paper:** Appendix A4 **Corpus companion entry:** Entity E-020 **Audit Entity ID:** org-imf

11.20.4. Appendix D4. 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 A11 **Corpus companion entry:** Entity E-044 **Audit Entity ID:** person-dario-amodei

11.20.5. Appendix D5. 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 A13, Appendix A6, Appendix A7 **Corpus companion entry:** Entity E-054 **Audit Entity ID:** person-mark-zuckerberg

11.20.6. Appendix D6. 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 A8, Appendix A9 **Corpus companion entry:** Entity E-058 **Audit Entity ID:** person-peter-thiel

11.20.7. Appendix D7. 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 A15, Appendix A5 **Corpus companion entry:** Entity E-059 **Audit Entity ID:** person-sam-altman

11.21. Appendix E. Evidence Boundaries**11.21.1. Appendix E1. Disputed Matters**

- None registered for this paper.

11.21.2. Appendix E2. Interpretive Boundaries

- Appendix B32: 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 B33: That the continuity scenario includes asymmetry of fallback and exit, with elite actors purchasing redundancy while ordinary users are asked to normalize dependence.
- Appendix B34: 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.

11.21.3. Appendix E3. Speculative Narrative Risks

- Appendix B35: That bunker ownership, survival planning, or retreat architecture prove elite foreknowledge of imminent collapse or that frontier firms are knowingly engineering it.
- Appendix B36: That founder philanthropy or public warnings about concentration prove elite foreknowledge of imminent collapse or absolve frontier firms of responsibility.
- Appendix B37: 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 B38: That Peter Thiel’s New Zealand retreat logic, ideological associations, and political ties prove a unified hidden command chain linking platform governance, Palantir, and civilizational collapse planning.

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

- Appendix A10. Companion note: Note N-043. Risk flag: review. Note focus: That a primary multilateral financial institution explicitly warned in June 2026 that the AI capex race was outrunning earnings and free cash flow.
- Appendix A11. 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 A5. 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 A7. Companion note: Note N-094. Risk flag: review. Note focus: That Zuckerberg publicly minimized the compound’s bunker framing after the earlier investigation.
- Appendix A6. 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-098. Risk flag: review. Note focus: That access hierarchy in frontier AI is materially tied to compute concentration and not only to abstract innovation differences.
- Appendix A14. Companion note: Note N-099. Risk flag: review. Note focus: That Dion Wiggins publicly frames sovereignty in survivability terms and links continuity to disguised dependency.
- Appendix A2. Companion note: Note N-100. Risk flag: review. Note focus: That generative AI exposure reaches well beyond narrow routine clerical work and extends into professional and white-collar task bundles.
- Appendix A4. Companion note: Note N-101. Risk flag: review. Note focus: That labor-market exposure to generative AI is systematically uneven across geography, occupation, and class position.
- Appendix A1. Companion note: Note N-102. Risk flag: review. Note focus: That a primary U.S. education authority already treated AI-mediated learning as a governance question rather than as a neutral productivity upgrade.
- Appendix A9. 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.

- Appendix A8. Companion note: Note N-112. Risk flag: review. Note focus: That by early 2018 New Zealand was already a publicly legible elite-safe-haven trope in Silicon Valley collapse-prep discourse.

12. Chapter 10. What to do

12.1. 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.

12.2. Keywords

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

12.3. Research Question

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

12.4. 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.

12.5. 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. ¹¹²

¹¹²Sources: 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)

12.6. 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.

12.7. 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.¹¹³ 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.

12.8. 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.

12.9. 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 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-

¹¹³Sources: 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)

private legitimacy structures. This final chapter turns those upstream findings into doctrine. ¹¹⁴

12.9.1. 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. ¹¹⁵

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. ¹¹⁶

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. ¹¹⁷

12.9.2. 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. ¹¹⁸

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 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

¹¹⁴Sources: 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)

¹¹⁵Sources: 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)

¹¹⁶Sources: 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)

¹¹⁷Sources: 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)

¹¹⁸Sources: 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)

settled law. ¹¹⁹

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. ¹²⁰

12.9.3. 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. ¹²¹

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. ¹²²

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. ¹²³

12.10. 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

¹¹⁹Sources: 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)

¹²⁰Sources: 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)

¹²¹Sources: 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)

¹²²Sources: 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)

¹²³Sources: 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)

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.

12.11. 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.¹²⁴

12.12. 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.¹²⁵

12.13. 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.¹²⁶

12.14. 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.

12.15. 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

¹²⁴Sources: 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)

¹²⁵Sources: 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)

¹²⁶Sources: 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)

visible in the prose, in the bibliography, and in the appendix sections that summarize claims and caution points.

12.16. 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 government 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, potassium 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.

12.17. Appendix A. Source Register for This Paper

12.17.1. 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

12.17.2. 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

12.17.3. 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

12.17.4. 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

12.17.5. 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

12.17.6. 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

12.17.7. 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

12.17.8. 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

12.17.9. 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

12.17.10. 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

12.17.11. 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

12.17.12. 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

12.17.13. 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

12.17.14. 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

12.17.15. 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

12.17.16. 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

12.17.17. 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

12.17.18. 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

12.17.19. 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

12.17.20. 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

12.17.21. 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

12.17.22. 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

12.17.23. 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

12.17.24. 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

12.17.25. 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

12.17.26. 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 pla... **Corpus companion entry:** Source S-083 **Audit ID:** newyorker-sam-altman-manifest-destiny-2016

12.17.27. 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 mechani... **Corpus companion entry:** Source S-133 **Audit ID:** wired-zuckerberg-hawaii-compound-2023

12.17.28. 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

12.17.29. 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

12.17.30. 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

12.18. Appendix B. Claims Used in This Paper

12.18.1. 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

12.18.2. 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-co

12.18.3. 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

12.18.4. 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

12.18.5. 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 **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

12.18.6. 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

12.18.7. 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

12.18.8. 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

12.18.9. 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

12.18.10. 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

12.18.11. 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-05

12.18.12. 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-legiti

12.18.13. 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

12.18.14. 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

12.18.15. 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

12.18.16. 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

12.18.17. 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

12.18.18. 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

12.18.19. 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

12.18.20. 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

12.18.21. 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

12.18.22. 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

12.18.23. 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

12.18.24. 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

12.18.25. 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

12.18.26. 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

12.18.27. 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

12.18.28. 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

12.18.29. 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 responsibility.

Evidence status: Documented Fact **Source support:** Appendix A30 **Corpus companion entry:** Claim C-414 **Audit Claim ID:** vatican-magnifica-humanitas-2026-08-power-networks-legitim

12.18.30. 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

12.18.31. 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.

Evidence 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

12.18.32. 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

12.18.33. 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

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

12.18.34. 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

12.18.35. 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

12.18.36. 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

12.18.37. 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

12.18.38. 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

12.18.39. 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

12.18.40. 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

12.18.41. 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

12.18.42. 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

12.18.43. 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

12.18.44. 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

12.18.45. 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

12.18.46. 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

12.18.47. 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

12.18.48. 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

12.18.49. 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

12.18.50. 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

12.18.51. 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

12.18.52. 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

12.18.53. 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

12.18.54. 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

12.18.55. 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

12.18.56. 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

12.18.57. 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

12.18.58. 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

12.18.59. 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

12.18.60. 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

12.19. Appendix C. Timeline Slice

12.19.1. 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

12.19.2. 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

12.19.3. 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

12.19.4. 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-20

12.19.5. 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

12.19.6. 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-

12.19.7. 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

12.19.8. 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

12.19.9. 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

12.19.10. 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

12.20. Appendix D. Relevant Entities

12.20.1. 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

12.20.2. 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

12.20.3. 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

12.20.4. 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

12.20.5. 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

12.20.6. 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

12.20.7. 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

12.20.8. 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

12.20.9. 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

12.20.10. 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**12.20.11. 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**12.20.12. 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

12.20.13. 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

12.20.14. 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**12.20.15. 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**12.20.16. 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_agendaand others. **Relevant sources in this paper:** Appendix A20, Appendix A26 **Corpus companion****entry:** Entity E-059 **Audit Entity ID:** person-sam-altman

12.20.17. 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**12.20.18. 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**12.20.19. 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**12.20.20. Appendix D20. Leading the Future****Entity type:** political_committee **Role in this paper:** Actor relevant to 08_power_networks_legitimacy_captu**Relevant sources in this paper:** Appendix A23 **Corpus companion entry:** Entity E-024**Audit Entity ID:** org-leading-the-future**12.20.21. Appendix D21. U.S. Surgeon General / HHS****Entity type:** public_health **Role in this paper:** Actor relevant to 02_social_decay_youth_harm_hypersexuali**Relevant sources in this paper:** Appendix A5 **Corpus companion entry:** Entity E-019**Audit Entity ID:** org-hhs-surgeon-general**12.20.22. Appendix D22. Federal Trade Commission****Entity type:** regulator **Role in this paper:** Actor relevant to 01_surveillance_capitalism. **Rel-****evant sources in this paper:** Appendix A4 **Corpus companion entry:** Entity E-017 **Audit****Entity ID:** org-ftc**12.20.23. Appendix D23. FTC inquiry into AI investments and partnerships****Entity type:** regulatory_action **Role in this paper:** Actor relevant to 10_counterstrat-
egy_response_agenda. **Relevant sources in this paper:** Appendix A25 **Corpus companion****entry:** Entity E-061 **Audit Entity ID:** program-ftc-ai-partnerships-inquiry**12.20.24. 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

12.21. Appendix E. Evidence Boundaries

12.21.1. 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.

12.21.2. 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.

12.21.3. 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.

12.21.4. 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.