Chapter 4. Dario and Daniela Amodei, Anthropic, and the Conditions of Moral Answerability

The safety promise under organizational and market pressure | The Dark Side of AI Leadership

Abstract

This chapter follows Dario and Daniela Amodei from their research and organizational roles at OpenAI into the formation of Anthropic, a company whose public identity places safety and reliability near the center of its purpose. I examine that identity through dated company records, executive statements, court filings, public-policy proposals, and the company's own workforce and registration disclosures. The history does not support a simple contrast between a responsible Anthropic and an irresponsible industry. It shows a changing organization negotiating research ambition, commercial scale, governance, copyright claims, state power, labor, and public expectations. I use True Storytelling to go beneath corporate narratives and SEAM to ask which human and socioeconomic costs remain outside conventional reporting. The Concord Tesseract™ and multiplicative qualimetric procedure remain a demonstration: the evidence reviewed to date does not establish a qualifying event of Bakhtinian moral answerability, so I assign no A score and calculate no M value here. The chapter's argument is provisional and open to correction as the case record develops.

Keywords: Anthropic; Dario Amodei; Daniela Amodei; AI safety; True Storytelling; SEAM; Humanistic Productive Capitalism; moral answerability; Concord Tesseract™

1. Why Anthropic is a demanding case

I turn to Anthropic after examining NVIDIA and OpenAI because the organization presents a distinctive test of the argument I am developing. NVIDIA occupies an infrastructure chokepoint: its computing systems enable many AI products while material consequences extend through suppliers, data centers, energy, and water. OpenAI places mission, governance, commercialization, and public storytelling in sharp tension. Anthropic enters the book with an explicit safety-centered identity and a founding history connected to departures from OpenAI.

That identity makes the question of answerability especially consequential. When a company says that safety, reliability, and steerability are central to its work, what evidence would show that these commitments shape decisions when they conflict with speed, revenue, competition, or state demands?

I do not approach this case as a prosecution of Dario Amodei, Daniela Amodei, or Anthropic. The book asks how leaders and organizations learn to act within a culture of capitalism that makes some consequences visible and leaves others outside the decision frame. My own experience with AI is part of why I ask the question. In "Climbing the Tallest Tree: Leaving AI Addiction Behind Me," I describe a personal struggle with reliance on AI tools (Boje, 2026). That experience does not tell me what Anthropic's leaders intended. It does remind me that a system's language can invite human attachment, reliance, and trust, and that the organizations building such systems bear responsibilities for the conditions they create.

My provisional thesis is that Anthropic's safety commitments create a demanding, and therefore useful, case for examining the distance between declared principles and consequential organizational practice. The strongest rival interpretation is that Anthropic has made meaningful efforts to develop technical safeguards, governance mechanisms, and public boundaries, including decisions that may carry commercial or political costs. Both interpretations must remain open. A published policy is evidence of a commitment; it is not, by itself, proof of implementation. A lawsuit or conflict is evidence of a dispute; it is not, by itself, proof that a leader failed morally.

2. Dario and Daniela before Anthropic

Dario Amodei's route into AI leadership began in physics and biomedical research rather than in the familiar trajectory of a startup founder. Anthropic's SEC registration materials identify his undergraduate and doctoral training in physics, postdoctoral work at Stanford, research roles at Baidu and Google Brain, and his service as OpenAI's Vice President of Research from July 2016 through December 2020 (Anthropic PBC, 2026a, 2026b). This sequence matters because it places him inside a research organization as AI systems were scaling rapidly, before he became the chief executive of a company formed around a distinct account of how that scaling should be governed.

Daniela Amodei's path is equally important to the organizational story. The SEC filing identifies her as Anthropic's co-founder and President and records that she held several roles at

OpenAI from October 2018 to December 2020, including Engineering Manager, Vice President of People, and Vice President of Safety and Policy (Anthropic PBC, 2026a, 2026b). This answers a question I had earlier: Daniela did work at OpenAI before Anthropic. Her work connected technical teams, people operations, and safety policy. Dario and Daniela should therefore be considered as a leadership ensemble with different formal responsibilities, not collapsed into a single sibling narrative or treated as interchangeable representatives of the company.

The surrounding OpenAI network also matters, but the evidence requires careful distinctions. Dario's OpenAI research leadership overlapped institutionally with the organization's founding period and with Sam Altman's increasing executive role. Elon Musk was an OpenAI co-founder and co-chair in its earlier period; Dario joined OpenAI in 2016 and later held the Vice President of Research role. These positions establish overlapping organizational histories, not the details of personal relationships or private disagreements. The technical work that led to modern large language models likewise came from broad research teams, accumulated methods, data, and computing infrastructure. I do not attribute the LLM transition to one person or one dissertation. The chapter's task is to map relationships and organizational episodes without turning proximity into causation.

3. Departure, founding, and the promise of a different institution

Anthropic's registration statement records that a group of former OpenAI colleagues founded the company in January 2021, with Dario as chief executive and Daniela as President (Anthropic PBC, 2026a). Anthropic's announcement of its $124 million Series A in May 2021 described the purpose as building more reliable general AI systems and named the Amodeis in those leadership roles (Anthropic, 2021). This is the first public institutional expression in the record reviewed here of the company's alternative project: not simply a group of researchers moving between employers, but a new organization linking technical development with a claim about reliability and responsibility.

Accounts of why the group left OpenAI should be attributed to the people and sources making them. Dario has described disagreements about the pace and organization of AI development; Daniela has also described differences in the vision they wished to pursue. These statements are important speech acts and evidence about how the founders narrate the break. They do not give me complete access to internal deliberations, and they should not be converted into the claim that one side was expelled or that the later company resolved the original conflict. Leaving can represent principled dissent, strategic entrepreneurship, or both. The research

question is what kind of organization the departure made possible and whether its structures give dissent durable influence after the founding moment.

This is where the chapter connects with the preceding OpenAI case. The formation of Anthropic does not prove that OpenAI could not accommodate competing views, just as Anthropic's founding principles do not prove that its own organization can. A departure creates a new center of authority, new resource dependencies, and a new public story. The story can open a space for learning; it can also become a source of legitimacy that later decisions are expected to preserve. My task is to follow the story into organizational action.

4. From safety commitments to a public product

Anthropic publicly introduced Claude in March 2023 after earlier partner testing (Anthropic, 2023a). The launch moved the company from a research and governance promise into a relationship with users, developers, and commercial partners. A product release changes the ethical field: the organization's decisions become part of everyday practices, and users encounter the system through its responses, refusals, style, and limits. A public launch is not proof of safe deployment. It is an event that makes the company's claims testable against what systems do and how the company responds when they fail or cause harm.

In July 2023, Dario Amodei gave written testimony to the U.S. Senate Judiciary Subcommittee on Privacy, Technology, and the Law. One concise formulation in that testimony was that "the danger and the solution to the danger are often coupled" (Amodei, 2023, p. 1). I read this as an important speech act because it rejects the idea that capability and risk can be separated neatly. The phrase also raises a practical question: who gets to define the danger, who controls the proposed solution, and whose interests are included in the coupling? The full testimony—not a sentence excerpt—must be read in context, and the statement must be compared with decisions that followed.

In September 2023, Anthropic publicly described its Long-Term Benefit Trust. The company said the structure had been developed since its founding and that an earlier version of the idea had appeared in its Series A documents (Anthropic, 2023b). The Trust is a governance design intended to give a group of trustees a role in selecting board members over time. That design is relevant evidence, but a governance mechanism should be assessed by what it can do in a specific conflict: whether it has access to information, whether it can challenge management,

and whether its decision changes an outcome. Formal independence and enacted authority are related but not identical.

The available company records therefore show a sequence of commitments: a founding purpose tied to reliability, a public product, a congressional statement about coupling danger and solutions, and a governance mechanism presented as serving long-term benefit. This sequence does not settle the case. It creates the standards against which later conduct can be examined. It also creates the risk of a familiar organizational gap: strong public language may coexist with practical limits on whose concerns affect product and investment decisions.

I treat these brief excerpts as speech acts to code, not as evidence that a promised practice occurred. Table 4.1 records the claim, its source, and the action that would need to be traced.

Table 4.1. Selected Amodei speech acts as researchable claims

Date and source	Short excerpt	What the speech act states	Question for the event record
July 25, 2023; Senate testimony (Amodei, 2023)	"the danger and the solution to the danger are often coupled."	Risk and mitigation are framed as connected.	Which proposed safeguards affected a documented decision?
October 2024; essay (Amodei, 2024)	"Fear is one kind of motivator, but it's not enough: we need hope as well."	A hopeful future is offered alongside risk.	Which promised benefits are observable, and who receives them?
February 26, 2026; company statement (Amodei, 2026a)	"we cannot in good conscience accede to their request."	The company states a boundary in a government dispute.	What action followed, and whose interests were protected or burdened?
September 12, 2026; policy essay (Amodei, 2026c)	"we owe it to humanity to try."	A public appeal for pacing and policy action.	Which proposal was enacted or adopted by the organization?

5. "Machines of Loving Grace" as an antenarrative

In October 2024, Amodei published "Machines of Loving Grace," an essay about how advanced AI might transform the world for the better (Amodei, 2024). Its title and scenarios make a strong wager on possible futures: accelerated scientific discovery, improvements in health, wider access to expertise, and material gains in human welfare. The essay is not a report of outcomes already achieved. It is a future-oriented narrative—a set of conditional expectations about what might happen if powerful AI systems are developed and used in certain ways.

I treat this essay as antenarrative: a story in advance that can organize investment, research, and public imagination before its predicted future arrives. Amodei writes that fear alone is insufficient and that hope is also needed. That is a serious rhetorical move. It makes room for aspiration without denying risk. At the same time, the hopeful future depends on organizational

choices about distribution, ownership, labor, safety, and access. The "country of geniuses in a datacenter" metaphor is powerful precisely because it condenses a political-economic wager into an image: extraordinary productive capacity concentrated in technical systems and the organizations that control them.

The question for this book is not whether hope is legitimate. The question is whether the institutions that promise broad benefit can make visible the people, places, and costs that may not share equally in that benefit. Which communities can contest the future being proposed? What happens to workers whose roles change? How are energy, water, and land counted? How do intellectual property disputes shape the knowledge commons? What kinds of evidence would make Amodei revise a forecast? A hopeful antenarrative becomes answerable when it permits those questions to alter action, not merely to appear in a footnote.

6. Copyright, training data, and the distribution of value

Anthropic's copyright disputes provide a concrete site for examining how AI companies describe the relation between training, source material, and value. In Bartz et al. v. Anthropic, the Northern District of California's June 2025 fair-use ruling distinguished between training on lawfully acquired books and Anthropic's acquisition and retention of pirated books to create a library. In July 2026, the court granted final approval to a $1.5 billion class settlement (Bartz et al. v. Anthropic PBC, 2025, 2026). These are separate legal events. The fair-use ruling should not be summarized as a finding that every method of obtaining or retaining training material was lawful, and the settlement should not be described as a finding of liability on every claim.

A second dispute, brought by music publishers in 2026, must also be kept separate from Bartz and from the earlier Concord Music Group litigation. In November 2026, the court granted Amodei's motion on the direct-infringement claims with leave to amend and directed further proceedings on Anthropic's motion (Concord Music Group, Inc. v. Anthropic PBC, 2026). This was a pleading-stage ruling, not a final resolution of the publishers' claims. The legal record will continue to develop; this chapter reports the status at the November 8, 2026 research cutoff.

These cases make commodity fetishism a useful question for the book, though I do not adopt a Marxist conclusion in advance. The consumer often encounters a fluent answer as if it were an immaterial product, detached from the books, music, labor, data centers, and institutions that made it possible. The legal cases return attention to the conditions of production and the distribution of value. The qualimetric task is to record the episodes, count their recurrence and

consequences, and ask which costs can be traced to organizational decisions. The legal claim is not itself proof of a moral failure; a court ruling is not the same as an ethical verdict. But the dispute makes visible whose work and property are recognized, whose contribution is treated as available input, and how remedy is negotiated.

7. The Department of War dispute: boundaries, power, and answerability

In February and March 2026, Anthropic publicly described a dispute with the Department of War over the government's request for "any lawful use" of Claude. Amodei stated that the company would not accept uses involving fully autonomous weapons and mass domestic surveillance, while also distinguishing those boundaries from support for national-security work and existing military applications (Amodei, 2026a, 2026b). These are the company's own accounts of the negotiations and the limits it sought to maintain. They do not establish the government's full rationale, and they should not be summarized as a rejection of all military AI.

The dispute soon became a conflict over state authority and organizational autonomy. Anthropic said that it received a formal designation letter on March 4 and announced that it would challenge the government action. The company's public statements also addressed the anticipated effects on customers and included an apology from Amodei for the tone of a leaked internal message. These speech acts deserve attention both for their content and for their timing. A public refusal can represent an effort to maintain a boundary; it can also function as a high-stakes organizational story that builds trust with employees, customers, and the public. The evidence must examine the boundary's practical consequences and how the organization treated people affected by the dispute.

The litigation produced different outcomes under different statutory authorities. On August 27, 2026, the Northern District of California granted Anthropic partial summary judgment on constitutional and Administrative Procedure Act claims and set aside the challenged supply-chain designation under 10 U.S.C. § 3252, along with specified directives (Anthropic PBC v. U.S. Department of War, 2026). On September 25, a D.C. Circuit panel, by a 2–1 decision, upheld a separate designation under 41 U.S.C. § 4713, whose statutory definition differed from the one before the district court (United States v. Anthropic PBC, 2026). The decisions should not be compressed into a single victory or defeat. Each legal account needs its statute, parties,

procedural posture, and holding; neither ruling by itself settles the ethical evaluation of Anthropic's policy choice.

This episode is unusually relevant to the Y axis. It raises questions about who holds decision power over a model's use, whether a company's boundary protects people who may be subject to surveillance or autonomous weapons, and who bears the consequences when government and company positions conflict. The episode may also contain evidence of service to others. But a public statement of restraint is not enough to establish that affected people had voice, that their interests changed the decision, or that the organization acknowledged and repaired concrete harm. Those questions require the complete record and accounts from people beyond the executive team.

8. The frontier, employment, and the proposed public offering

In September 2026, Amodei published "We Must Pace the Frontier," a policy proposal calling for measures such as independent evaluation, domestic regulation, and international coordination (Amodei, 2026c). The essay continues the pairing of danger and possible solution visible in his 2023 testimony. It is a policy appeal, not enacted policy. The research question is whether the proposed safeguards become commitments the company accepts even when they constrain its own strategy, and whether the proposals include workers, users, and communities in the definition of the public interest.

Anthropic's October 2026 employment report said that the company had hired a smaller share of junior employees and that a smaller share of its workforce held technical roles (Anthropic, 2026a). Those findings deserve careful reading, but they do not establish mass layoffs, individual dismissals, or a decline in total headcount. Hiring mix, workforce composition, attrition, and layoffs are distinct measures. A humanistic productive analysis would ask how work is reorganized, what skills are valued, what opportunities are foreclosed, and whether workers have meaningful participation in changes affecting them. It would not turn a change in employee shares into a claim the report does not support.

Anthropic filed a public S-1 registration statement on October 26 and an amendment on November 2, 2026 (Anthropic PBC, 2026a, 2026b). These filings describe a proposed public offering; as of this chapter's November 8 cutoff, they do not establish a completed IPO, final offering price, or first trading date. The filings give the reader a view into the organization's financial position, risk disclosures, governance design, and commercial plans. A potential public

listing makes the question of Humanistic Productive Capitalism especially timely: how can productive capacity and entrepreneurial initiative be maintained while resisting the pressure to treat shareholder value, valuation, and quarterly performance as complete measures of organizational health?

I do not infer that an IPO automatically displaces safety, just as I do not infer that a public-benefit structure guarantees the public benefit. The change in financing and scale creates a new set of conditions to observe. Future chapters can test whether safety policies, stakeholder commitments, and governance arrangements shape resource allocation and product release decisions, or whether they remain parallel narratives beside commercial imperatives.

9. True Storytelling: going beneath the official account

I use True Storytelling as the chapter's Organization Development approach because it asks a company to work with what is true, sustainable, and energizing and to examine the relations beneath its official stories (Larsen et al., 2021). The approach is not a label for an intervention that Anthropic has conducted. I have not facilitated a True Storytelling process inside Anthropic, interviewed its leaders for this chapter, or observed internal deliberations. Here it serves as an analytic and prospective method: a way to identify the stories that organize action and specify what participatory work could make those stories more answerable.

The first principle, "going beneath," directs attention from the polished safety narrative toward the practices, routines, and material dependencies that sustain it. The chapter therefore follows the safety promise into model deployment, governance, copyright disputes, government negotiations, employment, and financial disclosures. The seventh principle, reflection, returns the organization to the consequences of its own decisions. Reflection is not simply a public apology or a new policy statement; it asks whether the organization learns from a specific event, changes a practice, and makes the change visible to affected stakeholders.

This case also requires polyphony. Researchers, product teams, users, artists, workers, policymakers, affected communities, and investors may tell different stories about the same event. A dialogic approach does not erase those differences in pursuit of a single official account. It asks whether opposing voices can affect decisions. Future research could use a True Storytelling circle, with consent and appropriate protections, to compare executive stories with worker, user, creator, and community accounts. That would be a proposed intervention, not evidence that such a process has occurred.

10. SEAM, HPC, and the socioeconomic ledger

SEAM adds an accounting discipline to this inquiry. Henri Savall's socioeconomic approach asks researchers to investigate dysfunctions, hidden costs, and unrealized human potential that ordinary financial reporting may not make decision-ready (Savall et al., 2024). Quarterly revenue, operating expense, losses, and valuation remain important evidence. They do not measure every consequence that matters to organizational decisions. A forensic socioeconomic ledger would examine such categories as rework, safety verification, labor and training burdens, conflict, turnover, intellectual-property disputes, energy and water use, regulatory costs, and missed opportunities for participation. Each entry needs a source, a stated estimation rule, and a clear distinction between observed expense, estimated consequence, and unresolved qualitative cost.

The chapter's working term for the broader economic orientation is Humanistic Productive Capitalism (HPC). I use the term in dialogue with Germán Bernácer's economic distinction and Henri Savall's socio-economic work. The Preface defines HPC's nine components and explains how the term differs from Marxism, socialism, and shareholder-centered capitalism. Here I apply the term as a research question: can a productive, entrepreneurial organization preserve innovation while taking human consequences, stakeholder participation, and non-speculative value seriously? This is not a claim that Anthropic has already achieved HPC. Nor does the label imply that every capitalist institution is identical or that conventional financial records are useless. It identifies what those records leave out when they stand alone.

For Anthropic, a socioeconomic inquiry would connect each episode to consequences rather than assigning dollar amounts by intuition. A court settlement is not automatically a measure of all the costs or benefits of training practices. A stated hiring mix is not a measure of worker displacement. A refusal of a government contract term is not a full account of who was protected or harmed. The ledger should collect the qualitative case, count comparable events where the record allows, and assess socioeconomic and financial consequences with explicit methods. Where the consequence cannot yet be priced, the record should say so rather than substitute a fabricated estimate.

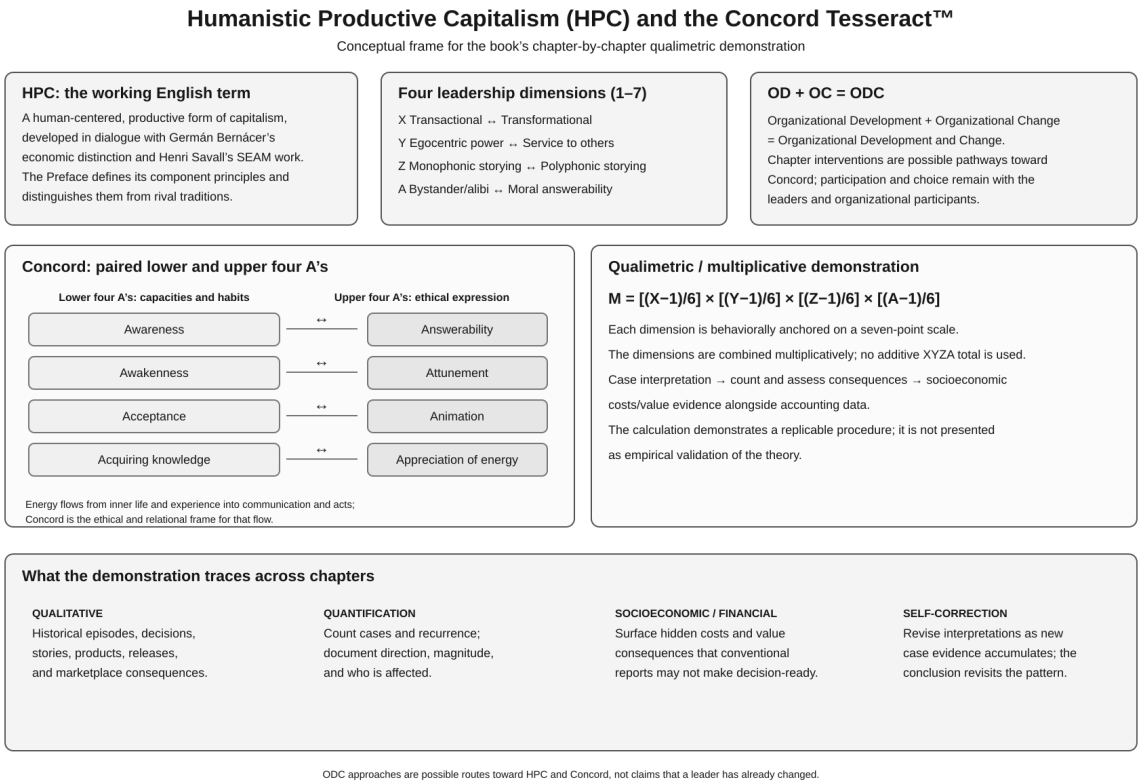


Figure 4.1. Humanistic Productive Capitalism and the Concord Tesseract™: conceptual frame for the chapter-by-chapter qualimetric demonstration. Working grayscale illustration.

11. Concord Tesseract™ and the multiplicative demonstration

The Concord Tesseract™ organizes this chapter's inquiry through X, Y, Z, and A. X ranges from transactional to transformational practice; Y ranges from egocentric power toward observable service to others; Z ranges from monophonic toward polyphonic organizing; and A concerns event-specific moral answerability in Bakhtin's sense (Bakhtin, 1993). The first three dimensions are leadership and organizing coordinates. A asks how a person responds to a concrete, once-occurring event with acknowledgment, ownership, answerability, and action. A cannot be inferred from a safety identity, a legal victory, or an eloquent public statement alone.

For visual purposes, the A cube uses three drawing dimensions: content and sense; situated eventness and ownership; and integration into action. These axes help locate and examine an event in the image. They are not three additional numerical ratings. The event receives one behaviorally anchored A rating only when the evidence supports a code. Appendix A gives the operational anchors, the minimum event record, and the procedure for updating the chapter-by-chapter display.

I have not identified a qualifying A event in the materials reviewed for this draft. This is a record-status statement: no event in the present evidence set has yet met the operational threshold. It is not a claim that Dario Amodei or any other leader has never acted answerably. It is also not a reason to convert missing evidence into a numeric score. Some episodes may inform Y by raising questions of power, service, and whose interests were represented. Those events should not be promoted to A unless they also document the response to a concrete event and the integration of that response into action.

The qualimetric/multiplicative demonstration is $M=[(X-1)/6]\times[(Y-1)/6]\times[(Z-1)/6]\times[(A-1)/6]$, calculated only when all four ratings are supported by reviewed evidence. For this chapter, A is NE because the materials reviewed do not establish a qualifying event under the behavioral anchor, so M is not calculated. Appendix C presents the current cumulative score tables, Appendix D specifies the updateable matrix and multiplicative rule, and Appendix E provides the reusable Grok prompt for the chapter figure. Dario Amodei's provisional X–Y–Z coordinates remain a priori starting placements; any revision must cite the new evidence, preserve the prior entry, and explain the change. This is a transparent demonstration for future empirical inquiry, not a validated measure of character or a probability.

12. OD + OC = ODC: possible practice, not an accomplished intervention

The book's formula, OD + OC = ODC, names a possible relationship between Organization Development and Organization Change. In this chapter, True Storytelling is the proposed OD approach. Organizational change concerns the actual consequences for governance, product development, employment, stakeholder relations, and resource allocation. Combining them is not a claim that Anthropic has adopted the method. It is a proposal for how an organization could invite affected participants to examine a story, test it against experience, and alter its practices when evidence warrants.

A future process would require voluntary participation, protection from retaliation, clarity about confidentiality, and a way for participant accounts to influence decisions. It should include voices beyond executive leadership and make the limits of participation explicit. A circle that collects stories but gives participants no influence can become another instrument of legitimacy. A circle that permits disagreement to change decisions might become a step toward concord. In

this book, concord does not mean consensus or the elimination of conflict. It means a relational and ethical capacity to hold differences while remaining answerable for the effects of one's acts.

13. Self-correction and provisional conclusion

The inquiry proceeds through abduction, deduction, and induction. Abduction proposes rival explanations: for example, that Anthropic's public safety boundaries reflect a durable organizational commitment, a strategic response to market and regulatory conditions, or an interaction between the two. Deduction specifies what evidence each explanation would predict. Induction assembles dated records, compares actions with statements, counts recurring episode types, and revises the explanation when the record contradicts it. The three paired A cycles—capacities, outward ethical practices, and the event-level answering process—can be examined alongside this larger self-correcting cycle.

My present conclusion is deliberately limited. Anthropic's record contains public safety commitments, governance design, policy proposals, product releases, legal disputes, workforce disclosures, and a proposed public offering. Those events make the company's commitments available for examination. They do not yet justify a final verdict about whether Anthropic has achieved moral answerability or Humanistic Productive Capitalism. The strongest next move is to deepen the case: authenticate the complete source for each speech act; compare policy versions with actual deployment and remediation; gather stakeholder accounts; and build the socioeconomic ledger without pretending that quarterly reports answer every question.

I remain open to the possibility that these leaders and their organizations can choose different practices. The purpose of the chapter is not blame or shame. It is to demonstrate a way to trace how organizational development and change might move an industry from concentrated, speculative, and instrumental forms of capitalism toward a more humanistic and productive orientation. If later evidence changes the interpretation, the map should change with it. A researcher who falsifies my theory is helping the work become more answerable.

References

Amodei, D. (2023, July 25). Written testimony of Dario Amodei, Ph.D. U.S. Senate Committee on the Judiciary, Subcommittee on Privacy, Technology, and the Law. https://www.judiciary.senate.gov/imo/media/doc/2023-07-26_-_testimony_-_amodei.pdf

Amodei, D. (2024, October). Machines of loving grace: How AI could transform the world for the better. https://www.darioamodei.com/essay/machines-of-loving-grace

Amodei, D. (2026a, February 26). Statement from Dario Amodei on our discussions with the Department of War. Anthropic. https://www.anthropic.com/news/statement-department-of-war

Amodei, D. (2026b, March 5). Where things stand with the Department of War. Anthropic. https://www.anthropic.com/news/where-stand-department-war

Amodei, D. (2026c, September 12). We must pace the frontier. https://www.darioamodei.com/post/we-must-pace-the-frontier

Anthropic. (2021, May 28). Anthropic raises $124 million to build more reliable, general AI systems. https://www.anthropic.com/news/anthropic-raises-124-million-to-build-more-reliable-general-ai-systems

Anthropic. (2023a, March 14). Introducing Claude. https://www.anthropic.com/news/introducing-claude

Anthropic. (2023b, September 19). The long-term benefit trust. https://www.anthropic.com/news/the-long-term-benefit-trust

Anthropic. (2026, October 16). How employment is changing at Anthropic. https://www.anthropic.com/research/employment-at-anthropic

Anthropic PBC. (2026, October 26). Registration statement on Form S-1. U.S. Securities and Exchange Commission. https://www.sec.gov/Archives/edgar/data/2130357/000162828026067168/anthropicpbc-sx1.htm

Anthropic PBC. (2026, November 2). Registration statement amendment on Form S-1/A. U.S. Securities and Exchange Commission. https://www.sec.gov/Archives/edgar/data/2130357/000162828026068676/anthropicpbc-sx1ax1.htm

Bakhtin, M. M. (1993). Toward a philosophy of the act (V. Liapunov, Trans.; V. Liapunov & M. Holquist, Eds.). University of Texas Press.

Bartz et al. v. Anthropic PBC, No. 4:24-cv-05417-AMO (N.D. Cal. June 23, 2025) (Order on fair use).

Bartz et al. v. Anthropic PBC, No. 4:24-cv-05417-AMO (N.D. Cal. July 20, 2026) (Order granting final approval of class settlement and judgment). https://law.justia.com/cases/federal/district-courts/california/candce/4%3A2024cv05417/434709/680/

Bernácer, G. (1916). Sociedad y felicidad: Ensayo de mecánica social. Francisco Beltrán.

Bernácer, G. (1955). Una economía libre sin crisis y sin paro. Aguilar.

Boje, D. M. (2026, October 6). Climbing the tallest tree: Leaving AI addiction behind me [Working paper]. https://davidboje.com/ws/Climbing_Tallest_Tree_Storytelling_Mountain_2026.pdf

Concord Music Group, Inc. v. Anthropic PBC, No. 5:26-cv-00880-EKL (N.D. Cal. Nov. 6, 2026) (Order regarding motion to dismiss hearing and meet and confer). https://cases.justia.com/federal/district-courts/california/candce/5%3A2026cv00880/463362/141/0.pdf

Kahn, J. (2023, September 26). Anthropic's CEO says why he quit his job at OpenAI to start a competitor that just received billions from Amazon and Google. Fortune. https://fortune.com/2023/09/26/anthropic-ceo-interview-quit-open-ai-amazon-investment/

Kantrowitz, A. (2025, July 29). The making of Dario Amodei. Big Technology. https://www.bigtechnology.com/p/the-making-of-dario-amodei

Larsen, J., Boje, D. M., & Bruun, L. (2021). True storytelling: Seven principles for an ethical and sustainable change-management strategy. Routledge.

Savall, H., Zardet, V., Bonnet, M., & Savall, A. (2024). Presentation of SEAM and opening up [Business] storytelling. World Scientific Publishing Company, Series 1. [Verify the final APA 7 record against the author's reference copy.]

Anthropic PBC v. U.S. Department of War, No. 3:26-cv-01996-RFL (N.D. Cal. Aug. 27, 2026) (Order on cross-motions for summary judgment, Doc. 250). https://www.govinfo.gov/app/details/USCOURTS-cand-3_26-cv-01996/USCOURTS-cand-3_26-cv-01996-5

United States v. Anthropic PBC, Nos. 26-1049 & 26-1162 (D.C. Cir. Sept. 25, 2026). https://www.govinfo.gov/app/details/USCOURTS-caDC-26-01049