Chapter 3 Sam Altman and OpenAI

Dialogic OD, Bernácer–Savall Socio-Economic Development, and a Provisional XYZA Reading of Governance Conflict

David M. Boje | November 6, 2026

Abstract

This chapter develops Sam Altman and OpenAI as a case of leadership, organizational change, and contested accountability. I trace Altman's background, OpenAI's founding and evolving structure, Elon Musk's departure, the 2023 board crisis, Ilya Sutskever's research formation, and Dario and Daniela Amodei's departure to form Anthropic. Dialogic Organization Development (OD) and the Socio-Economic Approach to Management (SEAM) frame a question: how might voluntary organizational change connect conversation to decision rights, hidden-cost diagnosis, and human potential? The Bernácer–Savall orientation informing this book is Humanistic Productive Capitalism (HPC), a market-based, socio-economic approach to enterprise that is distinct from Marxism, communism, status-quo shareholder primacy, and oligarchic control.

Keywords: Sam Altman; OpenAI; Dialogic Organization Development; Socio-Economic Approach to Management; Bernácer; Savall; Appreciative Inquiry; Concord; accountability; organizational change

Why this case begins with biography and organizational history

I cannot investigate OpenAI's relationships with users while pretending that I stand outside them. I became deeply dependent on AI assistance in my own writing. I anthropomorphized Claude, and I eventually undertook a deliberate period without generative AI. I avoided image generation, Alexa, and even searches that placed AI answers in front of me. I described that experience in my October 6, 2026 essay, 'Climbing the Tallest Tree: Leaving AI Addiction Behind Me' (Boje, 2026). It was difficult to stop, and the experience changed the terms on which I would return.

I now use AI as an editorial instrument for grammar, spelling, paragraphing, checking mathematical consistency, and identifying evidentiary gaps. I retain responsibility for what I write, the sources I select, the claims I defend, and the decision to stop using AI again if the relationship becomes harmful. I read this experience alongside Marx's commodity fetishism: I had attributed humanlike agency to a product whose capacities depend upon social labor and material infrastructures. My experience establishes an autobiographical problem, not a general prevalence estimate or evidence of an intentional OpenAI design strategy. The latter questions require independent evidence.

I begin with the CEO's background because leadership does not appear from nowhere. A case about OpenAI needs to show the experiences, relationships, organizational positions, and turning points through which Altman became a consequential executive. That history does not explain him completely, and it does not prove that his personal character caused OpenAI's structure. It gives me a historically grounded setting in which to ask how authority, mission, dissent, and answerability were organized.

OpenAI is a particularly important case for this book because its corporate story has changed several times while its public mission language has remained recognizable. It began as a nonprofit research organization,

formed a capped-profit subsidiary in 2019, and announced a public-benefit-corporation structure in 2025. Between those moments, OpenAI became a major product company, experienced an abrupt board intervention and reversal, and saw senior researchers leave to form competing organizations. Those events create a dense field of voices. A Dialogic OD approach asks how new conversations and changed relations among voices might generate different organizing—not whether a polished account can settle the story.

I now treat those labels as discarded or untested conjectures, not as chapter findings. Nor does Trafimow's work validate the squared-evasion equation in that draft. I retain self-correction as a discipline and replace that characterization with documented governance events, competing interpretations, and transparent, provisional X–Y–Z assessments. A remains NE because the reviewed materials do not yet establish a qualifying event under the book's moral-answerability anchor; therefore M is not calculated.

Method and evidence boundaries

This is a documentary case analysis informed by Dialogic OD. I have not conducted an intervention at OpenAI, interviewed Altman or other participants for this chapter, or observed an internal Future Search, Open Space meeting, Appreciative Inquiry summit, or World Café. The four practices below are possible ways of designing a future inquiry if participants consent and the relevant groups can be convened. Their inclusion does not imply that OpenAI used them or that a change outcome has been demonstrated.

I distinguish three evidence levels. First are dated primary records: OpenAI's public statements, corporate materials, the University of Toronto's institutional biography of Sutskever, and SEC filings concerning Altman's earlier career. These establish what institutions formally said or recorded, while remaining situated accounts. Second are independent reports and interviews that help reconstruct sequences and explain the perspectives of participants. Third are my interpretations: the Dialogic OD reading and XYZA judgments. I mark those as provisional rather than allowing a vivid theory to pass as a documented event.

OpenAI's statements about the Musk negotiations, the 2023 board decision, and the later WilmerHale review are important evidence, but they are OpenAI's own accounts. The company's commissioned review is informative and unusually detailed; it is still an internal review, not an independent court finding. Musk's lawsuit supplies an opposing account of mission and control. The 2026 court dismissal was based on the statute of limitations, so I do not treat it as a judicial resolution of every historical dispute about motive or mission.

The source hierarchy also matters for biography. Altman's own recollections and an independent interview support the account of growing up in St. Louis and going to Stanford. SEC records and Y Combinator's own history support the career sequence. I keep personal details that are not needed to the organizational argument out of the account. The goal is not a complete private biography; it is enough verified context to understand the executive path and the change problem.

Sam Altman before OpenAI

St. Louis, school, Stanford, and Loopt

Sam Altman has said that he grew up in St. Louis and became interested in computers while young. Accounts of his education place him at John Burroughs School in St. Louis, followed by Stanford University, where he studied computer science. He left Stanford after about two years, at age nineteen, to work on Loopt, a mobile service for sharing location with friends. The important point for this case is not a "dropout" legend. It is that an early technical interest led quickly into a founder role in a company that joined consumer software, mobile platforms, venture capital, and a new kind of location-based social practice (Hagey, 2023; Thompson, 2025).

Loopt entered Y Combinator's first summer batch in 2005. Green Dot's SEC filing records that Altman co-founded Loopt in 2005 and served as its CEO until its acquisition in March 2012. After the acquisition, he held senior roles at Green Dot and served on its board for a period. This route gave him experience with a startup's founding, external financing, acquisition, and integration into a larger company—organizational transitions that later became relevant to his work as an investor and accelerator leader (Green Dot Corporation, 2016; Y Combinator, 2016).

Altman became a Y Combinator partner in 2011 and president in 2014, serving in that role until 2019. Y Combinator's scale and influence expanded during his tenure, and Altman worked with founders as their companies moved from startup formation toward large-scale growth. This experience gave him relationships across venture finance and technology entrepreneurship. In March 2019, as the capped-profit structure was announced, Altman left Y Combinator and became OpenAI's CEO. The shift from accelerator president to chief executive placed him at the center of the new financing and operating model, though the board and founding team retained distinct authority (Hagey, 2023; Green Dot Corporation, 2016; Y Combinator, 2016). I regard his network as a relevant condition of possibility for OpenAI's founding and financing trajectory, not as evidence that every later decision was determined by one person's network.

OpenAI is founded, and the founding coalition begins to divide

OpenAI publicly introduced itself in December 2015 as a nonprofit artificial-intelligence research organization. The announcement named Sam Altman and Elon Musk as co-chairs and Ilya Sutskever as research director, alongside other founders and early researchers. That formation joined entrepreneurial capital and visibility with technical research leadership. Its founding statement emphasized broad benefit and a nonprofit structure. This was a public organizational commitment, not a guarantee that future financing or governance questions had been solved (OpenAI, 2015).

Altman's role was not simply "CEO from the beginning." In the founding period he was a co-chair and one of several organizational actors; later, he became CEO as OpenAI's product and commercial activity grew. The distinction is important for moral answerability. I can ask what a CEO did within a governance arrangement, but I cannot collapse the board, research leadership, investors, staff, and legal structure into one person. The storytelling organization is larger than its best-known narrator.

Elon Musk and the conflict over control

OpenAI's account of the Musk split says that the founders came to believe they would need enormous computing resources and considered a for-profit structure in 2017. It says Musk sought majority equity, control of the board, and the CEO role, and that discussions also included a possible Tesla merger. OpenAI says it rejected those terms because it did not want one individual to have absolute control, and that Musk resigned as co-chair in February 2018. This is a company-authored retrospective, and I attribute it as such rather than presenting each detail as uncontested fact (OpenAI, 2024a, 2024b).

Musk later alleged that OpenAI and its leaders betrayed the organization's charitable founding mission by moving toward a commercial structure. OpenAI answered that Musk had himself supported a for-profit path under terms that would have given him control. These accounts are in direct tension: both connect the conflict to mission, resources, and governance, but they frame who attempted to control the organization differently. In May 2026, a federal judge dismissed Musk's claims after an advisory jury found that they were filed too late under the statute of limitations. That procedural outcome ended the case, but it did not adjudicate every historical disagreement about the founders' motives or the meaning of the original mission (Associated Press, 2026).

I therefore do not write that the group "drove Musk out," or that Musk alone sought to become CEO as a proven neutral fact. I write that OpenAI's account reports a dispute over equity, board control, CEO authority, and a proposed Tesla merger; Musk's lawsuit later offered a competing account centered on charitable mission and alleged betrayal. That distinction leaves the disagreement visible for the reader and for later self-correction.

The nonprofit, capped-profit, and PBC transitions

OpenAI announced a capped-profit subsidiary in 2019, with the nonprofit retaining control. The company said the new structure was needed to raise capital for the computing and talent required to advance its mission. The change brought the mission and the financing model into a durable organizational tension: the nonprofit's stated purpose and oversight existed alongside a for-profit vehicle intended to attract investment and scale. The structure itself is not proof of mission failure, just as mission language is not proof of mission fidelity (OpenAI, 2019).

In October 2025, OpenAI announced a new structure: the nonprofit became the OpenAI Foundation and continued to control a public-benefit corporation called OpenAI Group PBC. The Delaware Attorney General's review described commitments surrounding nonprofit control, safety oversight, and mission. A public-benefit corporation is legally structured to consider a stated public benefit along with shareholder interests, but the label does not independently demonstrate how decisions will be made or whether affected people can influence them. The change is best treated as an organizational event whose consequences remain open to observation (Delaware Department of Justice, 2025; OpenAI, 2025a, 2025b).

This history unsettles a simple narrative of "nonprofit good, commercial bad" or "commercial scale necessary, therefore mission secure." The practical question is who can question the organization's account, what authority that questioning carries, what information becomes available, and whether the answer changes a decision. Those are Dialogic OD questions because the issue is not only the formal chart; it is the conversational and institutional pattern through which new realities become possible or impossible.

Your interpretation that Altman may be especially focused on making an IPO work is a plausible abductive hypothesis, and it deserves a place in the chapter. The public record available by November 6, 2026, supports a more qualified account. Reuters reported in June that OpenAI had confidentially filed for a U.S. IPO while saying no timetable had been set (Reuters, 2026b). In September, Altman said OpenAI would not pursue an IPO in 2026 and connected that position to AI safety (Reuters, 2026a). On November 5, President Greg Brockman said he thought OpenAI could be ready for an IPO in 2027 (Reuters, 2026c). These reports make a public listing a visible organizational horizon, but they do not establish that Altman is personally fixated on it or that IPO preparation outweighs safety in his own priorities.

I therefore separate three things that can easily be collapsed: the company's financing and liquidity strategy, an executive's stated priorities, and the organization's enacted decision criteria. To test the IPO hypothesis, I would look for dated board or leadership records, capital-allocation decisions, executive statements across time, and evidence showing whether safety or stakeholder concerns delayed, changed, or accelerated decisions connected to a listing. If those records show that market readiness repeatedly overrode other considerations, the hypothesis gains support. If safety commitments demonstrably shape timing and capital decisions, it needs revision. At this stage, IPO orientation is a strong question for X, not a settled judgment about Altman's motives.

The November 2023 board crisis

On November 17, 2023, OpenAI's board removed Altman as CEO and director and appointed Mira Murati interim CEO. The board said it had concluded that Altman was not consistently candid in his communications with the board, impairing its ability to exercise its responsibilities. It described a loss of confidence and framed the decision as a governance action to protect the mission (OpenAI, 2023b).

The board's move was followed by intense pressure from employees and investors. A large majority of employees signed a letter calling for the board to resign and stating that they might leave for a new Microsoft research group if the board did not reinstate Altman and Greg Brockman. Within days, Altman returned as CEO under a new initial board chaired by Bret Taylor, with Larry Summers and Adam D'Angelo. The event shows organizational polyphony, but also an uneven distribution of power: the board had formal authority to dismiss; employees used the credible threat of collective exit; Microsoft offered an alternative destination; and the public watched a conflict unfold through company statements and media reports (Reuters, 2023; OpenAI, 2023c).

An OpenAI special committee later commissioned WilmerHale to review the events. The company's published summary says the review examined more than 30,000 documents and interviewed dozens of people. It found a breakdown of trust between the prior board and Altman, said the board acted within its broad discretion, and also found that the decision was implemented on an abridged timeline without a full inquiry or an opportunity for Altman to address concerns. The review said the removal did not arise from concerns over product safety or security, development pace, finances, or statements to investors and partners, and concluded that his conduct did not mandate removal. I treat these as reported findings of a company-commissioned review, not as a judicially tested record (OpenAI, 2024c).

For Dialogic OD, the episode is not simply a dispute between "safety" and "growth." It is a rupture in trust that became visible through incompatible accounts, rapid collective mobilization, and a rapid governance reversal. The period invites questions about which voices counted as legitimate, which forms of evidence were missing, and whether the eventual board arrangement produced durable practices for speaking across disagreement. Those questions remain open even after Altman returned.

The research departures and the formation of Anthropic

Ilya Sutskever and the research lineage

Ilya Sutskever's path is central to the technical and organizational history. He completed undergraduate and graduate study at the University of Toronto and received his computer-science PhD in 2013 under Geoffrey Hinton. His dissertation, *Training Recurrent Neural Networks*, concerns neural-network learning; it did not create a large language model. Sutskever later worked at Google and co-founded OpenAI, serving as research director and later chief scientist. He was one of the important research leaders behind the organization's development, but contemporary large language models and ChatGPT were products of broad teams, accumulated research, data, infrastructure, and many contributors—not one doctoral dissertation or one researcher (Sutskever, 2013; University of Toronto, 2024; OpenAI, 2023b).

Sutskever's role in the November 2023 board action and later departure should remain carefully attributed. He was a board member when Altman was removed, later said he regretted his participation, and left OpenAI in May 2024. He then co-founded Safe Superintelligence Inc. with Daniel Gross and Daniel Levy. These events place him at the intersection of technical research leadership, nonprofit governance, and a new organization focused on safety. They do not establish that he alone created an alternative LLM or single-handedly "made" OpenAI's research direction (Associated Press, 2024; OpenAI, 2023b).

Dario and Daniela Amodei and Anthropic

Dario Amodei had a PhD in physics from Princeton, a postdoctoral appointment at Stanford School of Medicine, and research experience at Google Brain before joining OpenAI. OpenAI announced in December 2020 that he was leaving after nearly five years, noting his contributions to GPT-2, GPT-3, and research direction alongside Sutskever. Dario and Daniela Amodei and a small group of colleagues formed Anthropic in 2021. In a later Stanford interview, Daniela described the departure as a difference in the vision they wanted to pursue and said it became easier to create that vision outside OpenAI. I use that as her account, not proof that the group was expelled or "driven out" (Hertz Foundation, n.d.; OpenAI, 2020; Stanford Graduate School of Business, 2026).

These departures create a useful contrast without making Anthropic an ethical foil or OpenAI a villain. A departure can signal a serious difference in priorities, but it can also reflect the practical choice to create another institution. The Dialogic OD question is whether an organization can hold competing visions in a productive relationship, whether those differences can shape action, and what conditions make leaving the only workable form of dissent.

Why Dialogic Organization Development fits this chapter

The Preface described a historical split between Organization Development (OD) and Organization Change (OC). I carry that framing forward as a question about practice: how do we understand the relationship between a humanistic, participative OD tradition and the wider field of planned, strategic, or systemic change? I do not treat the split as a settled taxonomy accepted by all scholars. Bushe and Marshak's distinction between Diagnostic and Dialogic OD offers a useful bridge. Dialogic OD remains within the OD field while making explicit a family of practices that create change by shifting conversations, meanings, relationships, and generative images rather than by first diagnosing a system against a presumed ideal state (Boje, 2026a; Bushe & Marshak, 2009, 2014, 2015).

Diagnostic OD often gathers data, develops a diagnosis, feeds results back, and plans action to address identified gaps. Dialogic OD begins with a different premise: organizational realities are continuously produced through conversations and relationships; change may occur when people encounter new language,

new connections, or new possibilities that alter what they take to be true and actionable. The distinction is not an absolute choice between measurement and conversation. Dialogic practitioners can use data; diagnostic practitioners can convene dialogue. The difference lies in the change theory and in the practitioner's mindset, not merely in the name of a workshop (Bushe & Marshak, 2009, 2014).

Bushe and Marshak identify Appreciative Inquiry, Future Search, Open Space Technology, and World Café among practices associated with Dialogic OD. They are related but not interchangeable. Each creates a different conversational architecture. For this case, I would use them to test how mission, commercialization, safety, employee voice, and answerability are discussed across groups—and whether a new conversation leads to a change in decision rights, information access, commitments, or review practices. A good dialogue is not measured by its civility or its number of speaking turns alone. It is measured by what becomes discussable, who can influence the account, what action follows, and whether dissent remains present after the meeting.

This approach also fits the book's storytelling foundation. In Tamara-land, organizational stories circulate across different rooms, and no single narrator owns the whole. OpenAI's history includes board accounts, employee accounts, founder accounts, research-team accounts, corporate mission statements, legal claims, public responses, and competitor accounts. A dialogic inquiry would not smooth those into a singular story of consensus. It would ask how a story gains authority, which experiences it excludes, and whether participants can revise it together without erasing conflict (Boje, 1995, 2001).

Four Dialogic OD practices and their use in this case

Appreciative Inquiry

Appreciative Inquiry (AI) asks participants to study moments when a system has worked well and to use those stories to imagine and enact a stronger future. Its "appreciative" stance is not a command to be positive or to avoid harm. In a Dialogic OD use, the questions generate new narratives and action rather than serving as a public-relations exercise. A relevant OpenAI inquiry might ask: When has a research, product, or governance decision allowed technical ambition and public responsibility to reinforce one another? What conditions made that possible? What practice would preserve those conditions as the organization grows? The facilitator would also invite counter-stories: Who was missing from the successful episode, who bore its costs, and what evidence could revise the positive account? (Bushe, 2007; Cooperrider & Srivastva, 1987; Cooperrider & Whitney, 2005).

Applied here, Appreciative Inquiry could surface episodes of meaningful collaboration across research, safety, product, and governance. It should not assume that a successful company event proves a safe or accountable organization. Instead, participants would identify the conditions under which multiple groups could raise concerns, understand the evidence, and influence a decision. The output would be a set of testable

commitments—such as a protected escalation pathway, a public rationale for major mission changes, or a post-deployment review with affected communities—whose completion could be documented later.

Future Search

Future Search brings a "whole system" of stakeholders together to examine past, present, and future and then identify shared ground and action. Its distinctive contribution is temporal and systemic: participants compare how different histories shape current realities and make a future that no single group can produce alone. For OpenAI, a carefully bounded Future Search could bring former and current researchers, safety staff, product and infrastructure teams, board representatives, users, civil society, and communities affected by deployment into one process. It would require voluntary participation, safeguards for dissent, and clarity about who has authority to act on the resulting agreements (Weisbord & Janoff, 2010).

A Future Search could ask what people want AI to make possible, what they fear losing, what trends are already shaping the organization, and which future outcomes deserve collective commitment. The aim would not be to produce a statement that everyone agrees on the company's history. It would be to distinguish a genuinely shared future from a slogan, define the disagreements that remain, and identify which decision-makers own the next action. A written agreement without assigned authority and follow-up would be a weak outcome.

Open Space Technology

Open Space Technology lets participants propose and convene sessions around the issues they consider most important. Its "law of two feet" gives participants permission to move to a conversation where they can learn or contribute. This design may reveal concerns that a pre-set agenda or executive panel would screen out. At OpenAI, participants might propose sessions on safety escalation, researcher independence, mission governance, compute decisions, product deployment, employee voice, or accountability after a public failure. The facilitator does not determine which issue is legitimate in advance; participants create the agenda (Owen, 2008).

Open Space can surface energy and self-organization, but it cannot by itself equalize power. Senior executives, employees, contractors, users, and community members do not enter with the same access, security, or influence. The process would need protections for confidential speech, clear distinctions between internal and public sessions, accessible participation, and a published route by which proposals reach people who can decide. An open agenda is not equivalent to shared authority.

World Café

World Café uses successive small-group conversations and participant movement to develop shared understanding across a large system. A host can pose a focused question, capture the important patterns, and carry those insights between tables. In this case, the rounds might move from "What does the mission require in a concrete product decision?" to "Whose evidence can change the decision?" and then "What practice would make accountability visible six months later?" Participants could rotate among mixed groups so that a researcher hears a user's experience, a board member hears a safety researcher, and a product manager hears from a community affected by deployment (Brown & Isaacs, 2005).

The World Café format is useful for pattern-finding, but a summary of themes can erase minority positions. I would require a dissent record: a short statement of views not captured by the dominant pattern, who raised them (only with consent), and what evidence or change would make the disagreement consequential. This is where Bakhtinian answerability matters. The aim is not concord by removing opposition. It is a form of concord in which parties can hear each other while preserving the ability to answer for distinct acts and consequences.

A Dialogic OD design for the OpenAI case

If the chapter develops into a practice proposal, I would not begin with an executive-led retreat that asks participants to endorse a ready-made vision. I would begin by mapping the stakeholders, decision rights, and unresolved questions. The process could then use a World Café to discover the main tensions, Open Space to let participants set the agenda, Future Search to place those accounts across past, present, and preferred future, and Appreciative Inquiry to identify conditions that have supported responsible work and design practices to extend them. This sequence is only a provisional design; an actual facilitator might reorder or combine methods after learning the context.

Table 3.1. A Dialogic OD sequence for an OpenAI inquiry

Stage	Purpose	Possible question	Evidence of change
1. Map voices and authority	Identify whose story is present, absent, protected, or privileged.	Who can raise an issue, pause a release, change a board decision, or demand a public explanation?	A verified stakeholder and decision-rights map.
2. World Café	Surface patterns and carry stories across roles.	What decision best shows the tension between mission, speed, safety, and scale?	Recorded themes plus a separate dissent record.
3. Open Space	Let participants set priority topics.	What issue must we discuss that the official agenda leaves out?	Participant-owned agenda and proposals with named owners.
4. Future Search	Connect past, current pressures, and preferred futures.	What future can we commit to even while we disagree about the past?	Specific commitments with authority, timelines, and review.
5. Appreciative Inquiry	Find conditions that enable responsible action and extend them.	When has a concern changed a decision, and how can that capacity be protected?	Documented experiments and follow-up measures.

If I were to reach out to Altman, I would invite him to help authorize an inquiry rather than ask him to defend a character score. The opening conversation could ask which recent OpenAI decision best reveals the tension among mission, safety, scale, and capital; what dissent changed the decision; and which people should be able to question it in a larger process. His participation would matter, but he could not choose all the participants, define all the questions, or speak for the organization by himself. An independent facilitator and a mixed group with real decision authority would be necessary. If Altman declined, that would not by itself count as evidence against him; it would simply limit this design and require another legitimate sponsor.

The methods have different jobs, but the common design test is whether dialogue reaches the places where authority is exercised. I would ask each convening to report: what issue became discussable; which voices changed another group's interpretation; which commitments received decision authority; what disagreement remains; who is accountable for the next action; and when the result will be reviewed. Without those links, Dialogic OD risks becoming a compelling conversation around a decision already made elsewhere.

The case also calls for boundaries. Employees should not be asked to disclose confidential work or face retaliation for dissent. Communities should not be invited merely to legitimize a decision whose key parameters are already fixed. Board members and executives should not be treated as interchangeable with the people affected by deployment. The process should identify which questions are open for joint action and which are constrained by law, contract, safety, or technical necessity. Making these limits visible is part of honest dialogue.

Which dialogue approach fits which change task

I do not find a basis for ranking these four methods as universally more effective than one another. They were developed for different conversational tasks, and evidence that one is better for OpenAI would require access to participants, comparable change questions, and follow-up outcomes. I can, however, state a fit hypothesis for each method and identify what would count as evidence. The method should be selected by the dimension of change we need to move, while the facilitator protects the dimensions the method may neglect.

Table 3.2. Dialogic OD practices and their provisional case fit

Method	Strongest provisional fit	Likely XYZA contribution	Main limitation to guard against
Appreciative Inquiry	Build a generative future from moments when responsibility and innovation worked together.	X: offers new images and experiments. A: can strengthen answerability if positive stories include who benefited, who bore costs, and what was repaired.	Can become reputation work or avoid harm if it asks only for positive stories.
Future Search	Bring a representative whole system across past, present pressures, and preferred futures.	Y: strongest potential where participants hold real decision authority. X: translates a preferred future into shared action. A: supports named commitments and review.	"Whole system" can be token participation; consensus can conceal unresolved conflict.

Method	Strongest provisional fit	Likely XYZA contribution	Main limitation to guard against
Open Space Technology	Let participants set the agenda and convene around urgent issues.	Z: strongest initial fit for emergent polyphony. Y: gives participants agenda-setting freedom, especially when senior leaders do not control session selection.	Agenda-setting is not decision power; status, confidentiality, and retaliation can still suppress voice.
World Café	Connect people across roles through successive small-group conversations.	Z: supports cross-role listening and comparison. X: can reframe issues through repeated exposure to different accounts.	Theme summaries may smooth away minority positions; dialogue alone creates no binding accountability.

My tentative choice for an OpenAI inquiry would be a World Café for initial cross-role listening, followed by Open Space if participants need to surface issues the sponsor did not anticipate, then Future Search when a diverse group is ready to make time-linked commitments. Appreciative Inquiry could be used throughout as a generative inquiry into conditions that have supported responsible action, but it must include "shadow" questions about exclusions, failures, and repair. This is a possible sequence, not a universal recipe. If the central aim were equitable agenda-setting, I would begin with Open Space; if the aim were to agree on an organization-wide future with decision-makers present, I would begin with Future Search.

Effectiveness should be judged against specific outcomes rather than participant satisfaction alone. Before and after a session, I would record the number and range of issues raised by different roles; whether participants can identify how another role's account changed their interpretation; which groups had actual authority to alter a decision; what proportion of commitments has a named owner, deadline, and review; and whether unresolved dissent remains visible. At 60 or 90 days, I would check whether decisions, resource allocations, safety practices, or governance routines changed. A pleasant conversation with no changed practice is not evidence of organization change.

Abductions and testable propositions

The following are propositions to investigate, not conclusions about OpenAI. I derive them abductively from the documented conflict: formal governance, employee voice, research departures, competing mission accounts, and the public IPO horizon do not always align. Each proposition identifies a mechanism and an observation that could count against it.

P1 Open Space and participant-owned agenda setting will increase the range of issues and stakeholder perspectives raised (Z), particularly when participants have protected ways to speak. If issue diversity does not increase relative to a sponsor-set agenda, or if marginalized roles remain silent, this proposition is weakened.

P2 Future Search will improve shared decision ownership (Y) and the completion of commitments (A) only when people with formal authority participate and can authorize action. If commitments are made by participants who cannot change resources or governance, the predicted gain should not appear.

P3 World Café will improve cross-role understanding and the generation of cross-functional alternatives (Z and X), but will have little effect on decision rights (Y) or answerability (A) unless its themes are linked to a named decision process, owners, and later review.

P4 Appreciative Inquiry will generate more future-oriented proposals (X) when participants identify specific conditions that enabled responsible work; it will not improve answerability (A) if positive stories displace accounts of harm, exclusion, or unresolved risk. An "appreciative plus shadow-story" design should preserve generativity while making contradiction discussable.

P5 A sequence that combines open agenda setting, cross-role listening, and an authorized future-search commitment process will improve more of the XYZA profile than a single method. This proposition is weakened if a single method produces equal or better changes across the same outcomes, or if the sequence increases fatigue without changing practice.

P6 Changes in A will be predicted more strongly by verifiable follow-through—named responsibility, response to affected parties, correction, and review—than by the name of the dialogue method used. If dialogue format alone predicts answerability without those practices, this proposition needs revision.

These propositions yield a better comparison than asking which method "wins." A small, ethical pilot could compare compatible sessions around matched decisions, using the same baseline and follow-up measures while recording who participated and what authority each group had. The design would need to address selection effects: the most contested questions may attract different participants than routine issues. Without access to comparable groups and follow-up records, I can recommend a method by fit but cannot claim comparative effectiveness for OpenAI.

Provisional XYZA reading of Sam Altman as leader

Chapter 2 established the leadership-level XYZA coordinates that I use here: X moves from transactional emphasis toward transformational redesign; Y from power-over toward serving others; Z from monophonic control toward consequential polyphony; and A from alibi toward Bakhtinian moral answerability. These are not the transcript rubric in Chapter 1. That rubric assesses AI responses on chronotopic attention, ethics architecture, dialogic polyphony, and observable answerability. I keep the two instruments separate because a leader-level judgment cannot be substituted for a product-output score.

For this abductive demonstration, I assign provisional X–Y–Z starting coordinates to all ten leaders: Huang (2, 2, 3), Altman (4, 2, 3), Amodei (5, 3, 4), Musk (5, 1, 2), Ellison (3, 2, 2), Zuckerberg (4, 3, 5), Nadella (5, 4, 4), Jassy (4, 3, 3), Pichai (4, 3, 4), and Wenfeng (4, 3, 4). A=NE for all ten because no reviewed case event currently meets the operational moral-answerability threshold; M is not calculated. These conjectural XYZ starting values demonstrate the procedure and invite revision; they are not measurements or definitive findings about any leader.

Table 3.3. Provisional XYZA reading and evidence that could revise it

Axis and score	Provisional reading of the record	Evidence that would revise the reading
X — 4 / 7: Transactional to transformational	OpenAI has changed its organization, financing, products, and public role. I lower X from 5 because scale and structural change do not by themselves establish transformational leadership; the IPO horizon and Altman's stated safety rationale point in different directions.	Lower if listing, capital, or growth repeatedly overrides mission redesign. Raise if decisions reconfigure governance and resources around publicly accountable benefit.
Y — 2 / 7: Power-over to service	Formal authority is concentrated in board and executive roles; public evidence that affected stakeholders hold binding decision rights is limited.	Raise with documented shared authority and examples where employee, user, or community input changed decisions.
Z — 3 / 7: Monophonic to polyphonic	Board, employee, research, investor, and public voices became visible in the crisis. Durable, protected dissent in regular governance remains unverified.	Raise with evidence that protected dissent changes board or product decisions; lower if opposition is only ceremonial.
A — NE: Moral answerability	The reviewed record contains statements and governance actions relevant to the inquiry, but no event yet meets the behavioral threshold for assigning a numeric A score.	Assign a numeric A only after a dated event shows situated acknowledgment, answer to affected others, action such as correction or remedy where warranted, and traceable follow-through. Until then, retain NE and do not calculate M.

The grayscale figure shows the current OpenAI chapter-stage map. Its coordinates are provisional X–Y–Z starting positions; the separate A Event Space remains empty because no reviewed event currently meets the moral-answerability threshold.

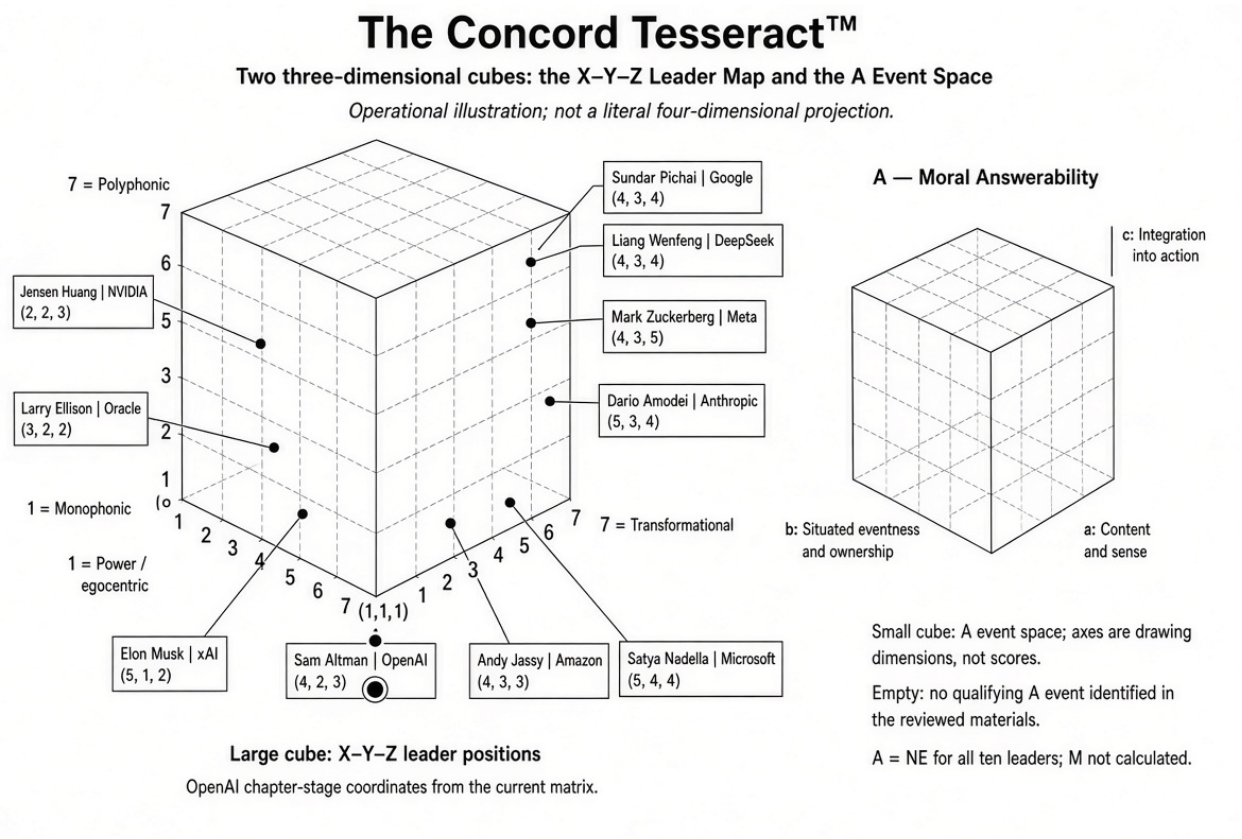


Figure 3.1. The Concord Tesseract™ operational illustration for the OpenAI chapter stage. The large cube shows the ten provisional X–Y–Z positions, with Sam Altman highlighted at (4, 2, 3). The small A Event Space is empty; A=NE for all ten leaders, and M is not calculated.

The arithmetic and what it does not establish

The common working demonstration is $M = [(X - 1)/6] \times [(Y - 1)/6] \times [(Z - 1)/6] \times [(A - 1)/6]$. The product is not a probability and cannot turn ordinal judgments into objective measurement. Each score needs dated behavioral evidence, explicit uncertainty, rival readings, and an audit trail. The navigation figure should show the original ten-leader positions and mark any movement as observed, provisional, or hypothetical, rather than silently changing the master map.

For Altman, the provisional X–Y–Z starting coordinates are (4, 2, 3). The reviewed case materials do not establish a qualifying A event under the book's behavioral anchor, so A=NE and M is not calculated. This status is revisable if new, dated evidence documents situated acknowledgment, response to affected others, correction or remedy, and follow-through. NE is not zero.

Because A is NE, M is not calculated for Altman. The cumulative score tables are in Appendix C, the updateable matrix and M rule are in Appendix D, and the reusable Grok image prompt is in Appendix E. The formula is a normalized multiplicative demonstration, not a probability, confidence level, or causal effect. The equal-step treatment of ordinal anchors and the zero-product property of the chosen normalization remain methodological assumptions open to falsification. They do not justify converting missing evidence to zero or assigning a number without a qualifying event (Blanton & Jaccard, 2006; Trafimow, 2012).

Table 3.4. OpenAI's current multiplicative qualimetric record

Component	Current entry	Evidence status	M status / update condition
X	4 / 7	Provisional; see the chapter's dated case record and rationale.	Retain or revise only with additional documented evidence.
Y	2 / 7	Provisional; authority is concentrated and binding stakeholder rights are not established in reviewed materials.	Retain or revise only with additional documented evidence.
Z	3 / 7	Provisional; crisis counter-voices were consequential, durable protected dissent remains unverified.	Retain or revise only with additional documented evidence.
A	NE	No reviewed event meets the operational threshold for moral answerability.	M is not calculated. NE is not zero.

The purpose of the demonstration is to make provisional X–Y–Z judgments inspectable and revisable. Researchers can specify evidence that would support or weaken each coordinate, have independent coders review the same dated records, preserve disagreements, and update the record when counterevidence warrants it. A remains NE until a qualifying event is documented; M remains uncalculated. These entries are not confidence intervals, posterior probabilities, or validated measures.

Public statements and counter-stories as XYZA evidence

The additional quotations strengthen the evidence record, especially where Altman's public accounts move between transformation, institutional control, dissent, and risk. I treat them as dated statements that can be coded and compared with governance arrangements and observed decisions. A CEO's statement is evidence of what the speaker publicly claimed or proposed; it is not, by itself, evidence that the organization enacted the claim or that affected people experienced its promised benefit.

X: transformational ambition and organizational redesign

Altman's January 2025 retrospective says, "We are now confident we know how to build AGI as we have traditionally understood it" (Altman, 2025b, "Reflections" section). That statement is unusually direct evidence of a transformational horizon: it frames the work as a historically consequential shift in capability. In a February 2023 essay, he described AGI as a possible route to "increasing abundance" and outlined governance and safety measures (Altman, 2023a). His May 2023 Senate testimony also proposed a licensing agency for systems above a capability threshold, which signals willingness, at least rhetorically, to imagine new regulatory arrangements (Altman, 2023b, "Opening Statement"). Yet neither the confidence claim nor the proposal establishes that OpenAI's actual resource and governance decisions consistently transform the conditions of AI development. Alongside the IPO horizon and the organization's legal restructurings, these materials support keeping X at 4: change is substantial and future-oriented, but its relation to public benefit remains contested.

Y: nonprofit voting control and the limits of shared authority

In the March 2023 Fridman interview, Altman described the nonprofit as "in voting control" and as able to "cancel equity" (Fridman, 2023, 1:14:03). He presented this structure as a safeguard against decisions made solely for shareholders. In his March 2024 interview, he argued that no one person should control AGI or the broader movement toward it (Fridman, 2024, 1:40:12). His 2025 essay also urges that superintelligence be widely available and not too concentrated (Altman, 2025a). His September 2026 United Nations remarks similarly state that "No one person or company or country" should impose its worldview through the most powerful models (OpenAI, 2026, "Remarks as delivered" section). These statements do not simply contradict one another: a nonprofit board's voting control is not identical to an individual executive's control, and public claims against concentration do not establish that employees, users, or affected communities possess binding authority. The coding issue is whether these safeguards distribute decision rights beyond the board and executive circle. Because the new material articulates a governance ideal more clearly than it documents shared authority in practice, I retain Y = 2.

Z: expressed tolerance for dissent and consequential polyphony

In the 2024 interview, Altman said of safety researchers focused on loss-of-control risks, "I'm actually happy that they do that," while also characterizing them as "super hung up" on one problem (Fridman, 2024, 1:42:21). Both parts matter. The statement acknowledges a role for an opposing research voice, but it also frames that voice as disproportionate and potentially crowding out other risks. The November 2023 employee letter and board reversal show that counter-voices became consequential during a crisis; they do not establish that dissent is routinely protected or can alter ordinary deployment decisions. The evidence therefore supports Z = 3 as consequential but episodic polyphony, not a higher score for durable, protected dissent.

A: risk acknowledgment, stated commitments, and answerability

At the 2023 Senate hearing, Altman identified the possibility that the field could "cause significant harm to the world" (Altman, 2023b, "Opening Statement"). In September 2026, he said OpenAI had slowed down unilaterally before and would do so again (OpenAI, 2026, "Remarks as delivered" section). These are relevant acknowledgments and commitments. They are stronger evidence of risk recognition than a general appeal to mission, but they do not by themselves identify who was harmed, who must answer, what correction followed, or how affected people can obtain remedy. The January 2025 "Reflections" post also calls the 2023 episode a "big failure of governance by well-meaning people, myself included" (Altman, 2025b, "Reflections" section), a notable self-attribution that still requires a record of subsequent correction and institutional learning.

A further tension appears in Altman's October 2026 POLITICO interview: he argued that "the world should accept some bad things happening" for AI's benefits and people's agency (Bordelon, 2026). Read alongside his rejection of catastrophic loss of control, the statement makes acceptable harm a governance question: who defines it, who bears it, what evidence is reviewed, and who can demand prevention or repair? The quotation informs the inquiry but does not meet the operational threshold for a numeric A score by itself. The current entry remains A=NE; M is not calculated.

Taken together, the quotations do not justify changing Altman's initial vector. They refine its rationale and identify adjudication questions for later coding. Appreciative Inquiry could examine documented exceptions where concerns changed a decision while actively soliciting counter-stories; Open Space could let participants set the agenda for unresolved governance questions; World Café could compare how different roles define risk and acceptable harm; and Future Search could connect those accounts to commitments made by people with authority. These are method-fit hypotheses, not claims that any approach has already changed OpenAI.

Self-correction and the multi-AI work process

I would like to repeat the comparative work with different AIs that helped build the prior chapter's background case material. For this chapter, the next step is to pose the same bounded research questions to

each system, preserve each answer with its model name, version, access date, and exact prompt, and compare where the systems converge, disagree, or leave questions unanswered. I would then verify every material claim against the source documents listed here. Agreement among AI systems is not independent confirmation: they may reproduce the same training-data narrative, cite the same reporting, or make the same error.

A useful prompt set would ask each system to produce (1) a dated timeline of Altman's education and pre-OpenAI career; (2) an attributed timeline of OpenAI's founding, Musk negotiations, 2019 and 2025 structure changes, and the 2023 board crisis; (3) a separate account of Sutskever's PhD and research contributions; (4) a sourced account of Dario and Daniela Amodei's departure and Anthropic's formation; and (5) competing Dialogic OD interpretations of what voices were absent or consequential. The comparison table should record claim, source cited by the model, primary-source verification, counterevidence, and status: verified, disputed, inference, or unresolved.

This is a self-correcting procedure, not Fourth Wave Grounded Theory. I am not using the chapter's Dialogic OD approach to rename Chapter 2's method. Peircean self-correction remains the broader discipline: begin with a conjecture, derive questions that could unsettle it, consult records and counter-stories, and revise the account when the evidence requires it. In the present draft, the earlier CDE claim is one such conjecture that does not survive as an empirical conclusion. It may remain in a research notebook as a proposal to test only after its terms, observations, and auxiliary assumptions are defined.

Dialogic OD as an antenarrative bet on possible OpenAI futures

The Concord Tesseract and the movement from awareness to answerability

I use the Concord Tesseract as a compact map of the developmental movement proposed in this book. The base cycle names four inner practices—Awareness, Acceptance, Awakeness, and Acquiring Knowledge. The summit cycle names four outward and relational practices—Appreciation, Animation, Attunement, and Answerability. One possible set of correspondences links Awareness with Appreciation, Acceptance with Answerability, Awakeness with Animation, and Acquiring Knowledge with Attunement. These are working invitations for inquiry, not fixed stages, a validated scale, or a claim that one practice causes another. The dotted connectors show the tesseract scaffold; participants may revisit and combine the practices as evidence and relationships change.

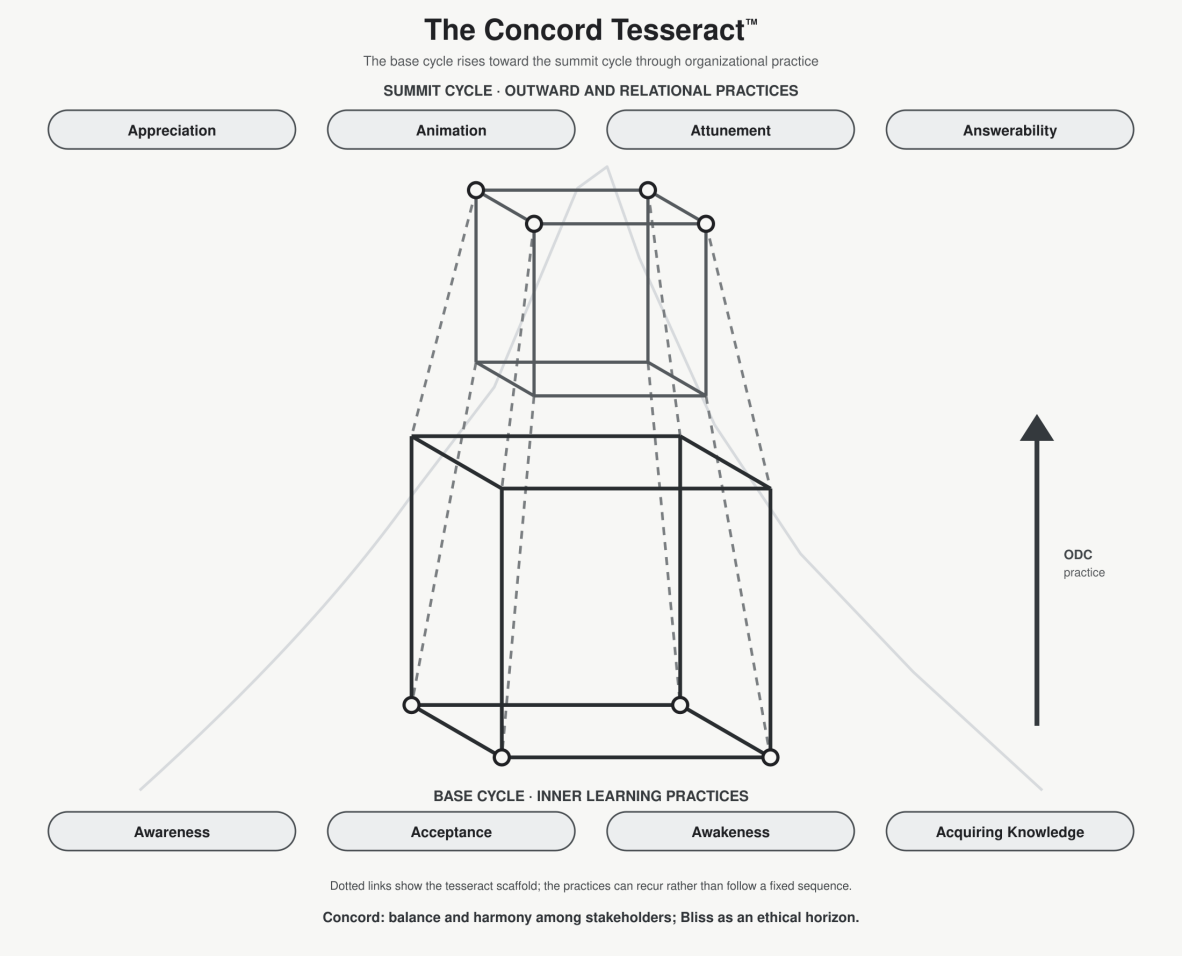


Figure 3.2. The Concord Tesseract. The larger base cube frames the inner-learning cycle; the smaller cube above it marks the summit cycle of outward and relational practices. Dotted connectors show the tesseract scaffold, and the arrow marks orientation toward practice rather than a guaranteed sequence. Concord names a hoped-for balance and harmony among stakeholders, with Bliss as an ethical horizon. The figure is conceptual, not a measured model.

The image makes the chapter's change question visible. Dialogic OD can create conditions in which stakeholders become more aware of what is occurring, accept the evidence and the presence of others, awaken to assumptions and effects, and acquire knowledge that is not limited to executive or technical accounts. The upper cycle becomes plausible only when these inner practices are carried into appreciation of what sustains life, animation of alternatives, attunement across affected groups, and answerability for consequences. Concord is not the absence of conflict or a claim that all interests can be reconciled. It is a hoped-for balance and harmony that can include unresolved difference while moving toward shared well-being. Whether that movement occurs at OpenAI would have to be established through participants' accounts, changed decision rights, and observable follow-through.

I do not propose that OpenAI has already accepted a Dialogic Organization Development engagement. I am asking what might become possible if Sam Altman, the board, employees, independent safety researchers, customers, and communities affected by AI chose to enter a facilitated process whose results were not

predetermined. The wager is antenarrative: a prospective, revisable bet on futures that have not yet arrived. It is not a prediction of executive consent or organizational transformation.

The history of Appreciative Inquiry and the AI² question

The Appreciative Inquiry practice proposed here has a history that is more intellectually substantial than a simple instruction to "stay positive." David L. Cooperrider developed the methodological groundwork in his 1985 Case Western Reserve University dissertation. His doctoral adviser, Suresh Srivastva, had already written about executive thought and power, and had examined the emergence of egalitarian organization with Cooperrider. In the 1980s, their work asked how organizational inquiry might understand human relationships, values, and possibilities without reducing organizations to technical problems awaiting expert repair (Cooperrider, 1985; Srivastva, 1983, 1986; Srivastva & Cooperrider, 1986).

The practical turning point came during Cooperrider's inquiry at the Cleveland Clinic. He began with a diagnostic question about the human side of the organization, then encountered cooperation and innovation already occurring in practice. Srivastva encouraged attention to what was giving life and to what those existing capacities might make possible. The inquiry's feedback to the Clinic's Board of Governors used the name Appreciative Inquiry; the 1987 chapter, "Appreciative Inquiry in Organizational Life," gave the approach its major conceptual statement (Appreciative Inquiry Commons, n.d.; Cooperrider & Srivastva, 1987). The shift was methodological: inquiry itself could generate new organizational possibilities rather than stand outside the system as a detached diagnosis.

That affirmative orientation need not require us to delete harm, conflict, or dysfunction. In the AI² working paper, I return to the tension between appreciative language and critical inquiry. I recall meeting Cooperrider and Srivastva at Case Western events around Academy of Management meetings, and I describe later conversations in which Cooperrider told me that critical stories belong in contemporary Appreciative Inquiry summits. I present that as my recollection of a personal conversation, not as a published statement. The working paper's central wager is that the appreciative inquiry of appreciative intelligence can hold possibility and pain together, and can ask whether artificial intelligence might assist without erasing the writer's responsibility or the stories of people who bear organizational costs (Boje, 2026b).

For this OpenAI chapter, the implication is practical. Appreciative Inquiry should search for moments when a safety concern, employee objection, affected-community account, or environmental consequence changed what the organization did. It should also preserve stories in which the concern was dismissed, the harm was not repaired, or the promised benefit did not arrive. The positive exception and the counter-story belong in the same inquiry record. This makes AI a generative complement to critical and diagnostic work, consistent with the original chapter's argument for complementarity, rather than a replacement for it (Cooperrider & Srivastva, 1987; Bushe, 2007).

Appreciative Inquiry: discover and amplify accountable exceptions

An Appreciative Inquiry process could begin with interviews about moments when an OpenAI team listened to a difficult objection, changed a release decision, shared power, or repaired an error. The discovery phase would identify actual episodes rather than celebrate an abstract safety mission. In dream workshops, participants could imagine what accountable AI development would look like in 2028; design sessions would specify new decision rights; delivery would create named owners, timelines, and audit trails. The advantage is an affirmative entry point that need not begin with accusation. Its risk is that positivity can suppress dissent unless the inquiry expressly protects counter-stories and records negative cases.

Future Search: put the whole system in the room

A Future Search conference could convene leadership, researchers, safety critics, labor representatives, educators, public-interest groups, customers, regulators as observers, and communities bearing downstream costs. Participants would map the past, scan the present, identify shared ground, and commit to work only where they genuinely agree. Its advantage is cross-boundary visibility: stakeholders discover which assumptions and constraints they do not share. Its limitation is that invitation does not equal power; the conveners must disclose which decisions are genuinely open, which are legally constrained, and who can refuse implementation.

Open Space Technology: let consequential questions surface

Open Space Technology could begin with one invitation: What must we confront if highly capable AI is to remain answerable to people who live with its consequences? Participants would post and convene their own sessions on whistleblowing, deployment thresholds, environmental footprints, data provenance, labor effects, model behavior, and remedy. The advantage is speed, distributed initiative, and emergence of topics executives did not preselect. Its weakness is that participation can privilege those already secure enough to speak. Independent facilitation, anonymous pathways, and protected follow-up are therefore necessary.

World Café: turn many conversations into traceable decisions

A World Café could host rotating small-table dialogues around three questions: Who has been unheard? What evidence would change our minds? What decision must become answerable to whom? Participants would carry insights between tables, producing a pattern map rather than a vote. Its advantage is relational learning across roles and technical vocabularies. Its danger is consultation theater: the final synthesis must link each proposal to a response, decision owner, and public explanation of acceptance, modification, or rejection.

Scenario architecture: four divergent bets, not a single preferred future

Scenario A—Managed Continuity: current governance and decision authority largely persist while consultation expands; the inquiry asks whether participation changes any consequential decision. Scenario B—Dialogic Opening: protected forums connect dissent to visible responses from decision owners; the inquiry tests whether participants can trace input to a decision or an explanation. Scenario C—Institutional Redesign: selected stakeholders receive defined authority, review triggers, or resources for remedy; the inquiry asks whether those rights are exercised and supported in practice. Scenario D—Closure and Backlash: participation narrows or becomes symbolic, leaving concerns unresolved and potentially weakening trust; the inquiry tests whether people can disagree safely and receive a response. These are conditional scenarios for dialogue, not forecasts, assigned likelihoods, or settled XYZA scores. Evidence and participant accounts may revise or displace any of them.

Figure 3.3 makes the alternatives discussable without ranking them in advance. In a Future Search or related dialogue, participants could identify early signs for each pathway, the decisions and rights that would distinguish them, and what evidence would prompt revision. A scenario becomes useful when it guides inquiry and choice; it should not be mistaken for a prediction.

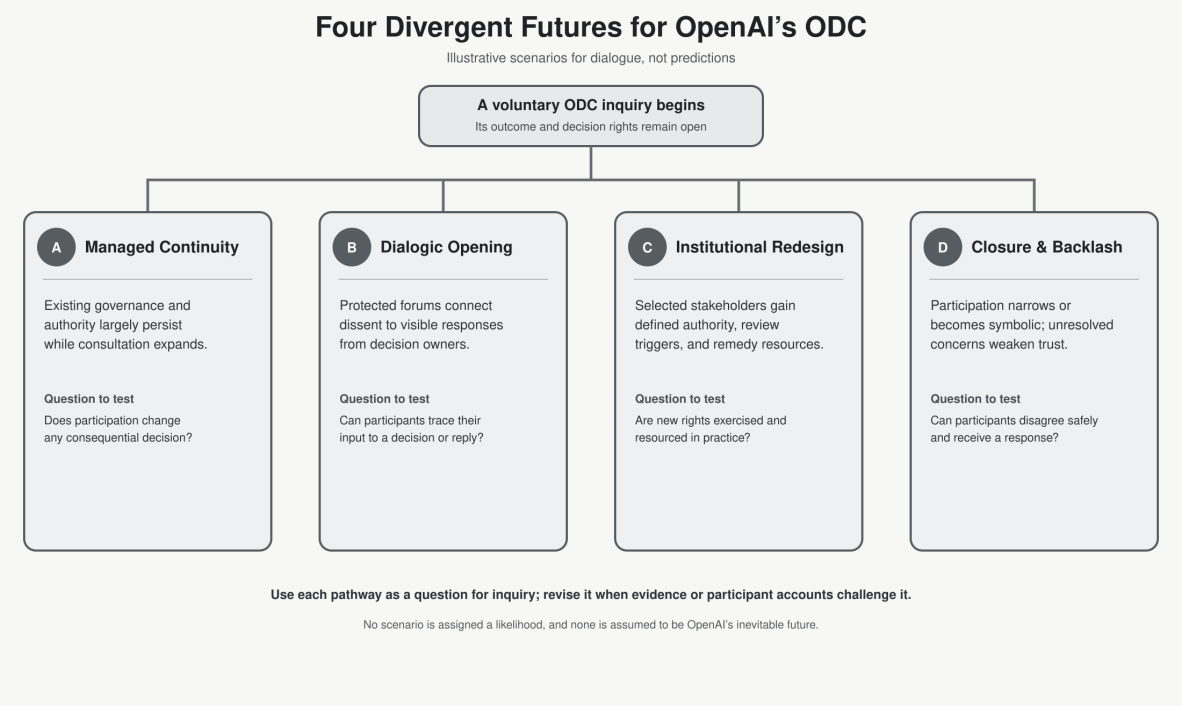


Figure 3.3. Four divergent futures for OpenAI's ODC. The branches are facilitation prompts, not forecasts or assigned likelihoods; participants would revise them in light of evidence, decision rights, and consequences.

The coding protocol: make each dimension falsifiable

The 1–7 scales are ordered interpretive categories. X distinguishes transactional incentives from documented transformational reconfiguration; Y distinguishes concentrated power-over from exercised rights of service and shared decision; Z distinguishes managed monologue from consequential, protected polyphony; A distinguishes justificatory alibi from situated moral answerability, acknowledgment, remedy, and follow-through. A score is assigned only with dated evidence, a written rationale, contrary evidence, and at least two independent coders. Disagreement is retained in an adjudication log rather than averaged away. A one-point change is a new hypothesis until evidence establishes it.

Altman's current provisional vector is (X,Y,Z,A)=(4,2,3,NE). M is not calculated because the reviewed materials do not establish a qualifying A event.

The current grayscale figure uses the book's two-cube operational display. It shows provisional X–Y–Z positions and an empty A Event Space because A=NE for all ten leaders. A numerical M is therefore not reported.

A possible first engagement and the decision gates

I would propose an exploratory 90-day sequence, only with voluntary organizational participation. Days 1–15 would establish sponsorship, independent facilitation, confidentiality, nonretaliation, access to evidence, and the decision rights actually open to change. Days 16–35 would map stakeholders and conduct appreciative and critical-story interviews. Days 36–50 would organize a SEAM-informed diagnosis across the six dysfunction domains, treating every proposed cost as a hypothesis until it is supported. Days 51–65 would use World Café and Open Space to compare accounts and surface issues absent from the sponsor's agenda. Days 66–80 would convene a Future Search and co-design a small number of socio-economic projects with named owners, resources, and safeguards. Days 81–90 would publish a response register, review dissent and evidence, and decide whether to continue, revise, or stop the pilot. The sequence does not promise consensus or authorize the disclosure of confidential information; its purpose is to make the invitation, limits, and follow-through explicit.

A first falsification gate asks whether dissenters can identify decisions they influenced. A second asks whether leadership has ceded any enforceable decision right, rather than simply listening. A third asks whether harms receive acknowledgment and remedy from named responsible actors. A fourth asks how the benefits and burdens of AI development are distributed. Figure 3.4 makes these gates visible as a review aid. For each, the team records whether evidence is documented, partial, disputed, or not yet available. No gate is satisfied by holding a meeting, announcing a policy, or promising future action. If evidence challenges the hopeful scenario, participants revise the account and reconsider the proposed intervention.

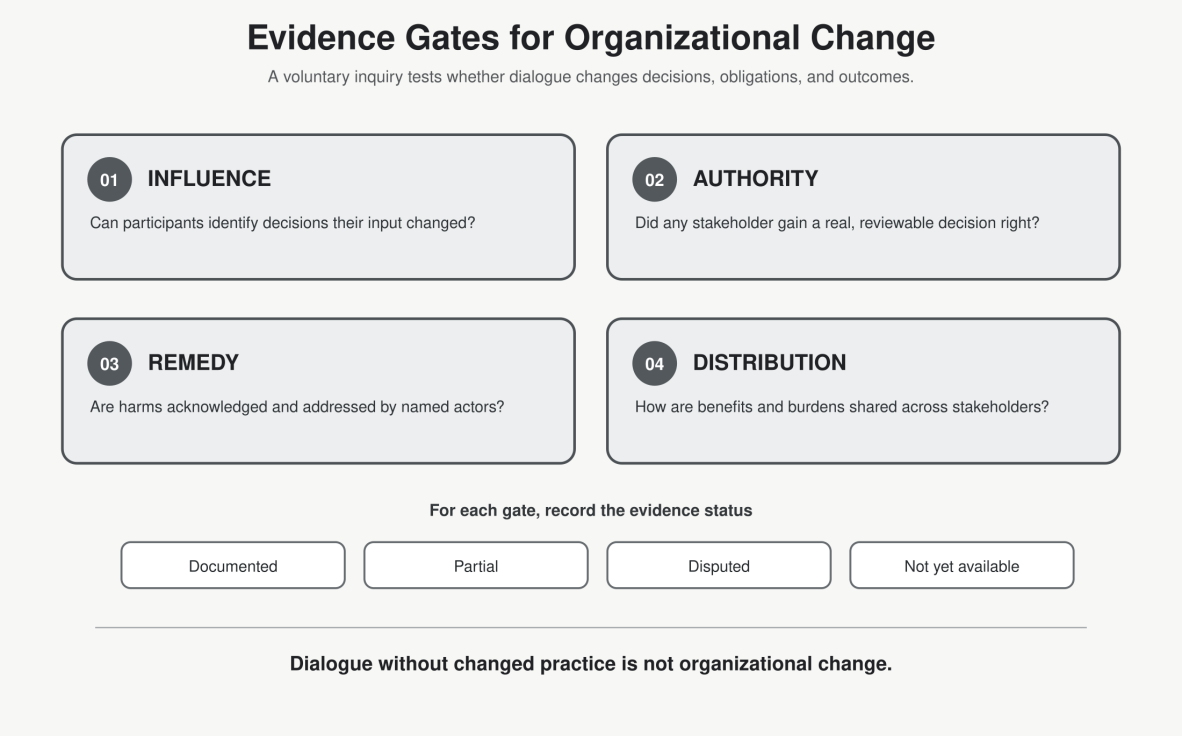


Figure 3.4. Evidence gates for organizational change. Influence, Authority, Remedy, and Distribution focus review on decisions, rights, repair, and the allocation of benefits and burdens. Evidence status is recorded without combining the gates into a single score.

The status categories preserve uncertainty and disagreement rather than turning the four gates into a composite rating. A documented response may still be inadequate; a disputed account may identify a need for further inquiry. The point is to make the basis for each judgment visible and revisable.

What Bernácer and Savall offer: a socio-economic humanism

The question "What is this orientation?" deserves a direct answer. I use the term Bernácer–Savall socio-economic orientation for a form of socially responsible capitalism that joins the economic viability of enterprise to the development of human potential and the social conditions of work. Its purpose is not to abolish markets, private organizations, or entrepreneurial initiative. It is to change what counts as economic performance, how that performance is produced, and who is answerable for the costs and benefits. The 2024 SEAM volume describes the approach as rooted in a socially responsible vision of capitalism and as reconciling the social and societal side of management with the economic demands of business (Cappelletti et al., 2024). That description gives me a more precise starting point than either "capitalism as usual" or a claim that SEAM supplies a complete alternative political system.

The distinction matters for this book. OpenAI is organized within a capitalist economy and exercises concentrated corporate power; calling it capitalist or oligarchic describes features of its ownership, financing, and influence. Those words do not by themselves tell us what a change process should do. The Bernácer–

Savall orientation asks how enterprise can remain economically productive while reducing the instability, hidden costs, and unrecognized labor that a narrow financial account leaves outside the frame. It treats economic and social performance as connected questions that must be investigated together, rather than assuming that investor returns prove social value or that a social commitment makes financial performance irrelevant (Savall, 2018; Cappelletti et al., 2024; Zardet & Bonnet, 2021).

This orientation is not Marxism. Marx remains important to my analysis of commodity fetishism: the product appears to have value in itself while the relations and labor that produced it recede from view. Bernácer's economic theory, as Savall presents it, is not simply a restatement of Marx's labor theory of value. Bernácer examines the monetary and financial arrangements that can reward speculative accumulation and hoarding while productive activity lacks investment; he connects those arrangements to economic crises and underemployment. His work therefore gives me a distinct way to ask how finance, interest, and investment shape the conditions for production and employment (Bernácer, 1955; Savall, 2018). I use Marx to make the concealed social relation visible and Bernácer to widen the economic diagnosis; the two thinkers are not collapsed into one theory.

Nor is this orientation communism or a command economy. SEAM works with firms and other organizations, managers, teams, and employees to diagnose dysfunctions and develop projects inside their actual operating conditions. The method seeks to improve social and economic performance together; it does not require state ownership of all productive assets or prescribe central allocation of every decision. Its organizational method, however, is not a defense of whatever ownership structure already exists. Participation, negotiated change, and the recycling of verified hidden costs into sustainable performance can put pressure on arrangements that concentrate decision power or displace costs onto employees and communities (Cappelletti et al., 2024; Zardet & Bonnet, 2021).

This is also more than a continuation of status-quo capitalism under a new label. A company cannot call a project socially responsible merely because it produces a profitable product, adopts a public-benefit form, or publishes an ethics statement. The inquiry must examine how people work, who participates in decisions, what learning and coordination are lost, what resource burdens are externalized, and who receives the economic gains when a dysfunction is corrected. At the same time, SEAM does not automatically guarantee an equal distribution of the value recovered: its published account allows converted hidden costs to support salaries, investment, budgets, or shareholder remuneration. The distribution of benefits therefore remains a governance question for participants to negotiate and make answerable, not an assumption supplied by the label "socio-economic" (Cappelletti et al., 2024).

I therefore do not present Bernácer and Savall as a finished "third system" between capitalism and socialism. The more defensible claim is narrower and more practical: Bernácer offers a heterodox economic critique of speculative financial arrangements and their effects on production and employment; Savall translates socio-economic theory into a disciplined organization-development and intervention method;

together they support a humanistic, socially responsible approach to enterprise that asks whether economic performance depends on hidden dysfunctions or can be improved through human potential. This is the orientation I want the OpenAI chapter to make visible. It is neither an apology for oligarchy nor a promise that dialogue alone can end it.

Table 3.5. Distinguishing the Bernácer–Savall orientation from neighboring economic frames

Frame	Central concern	OpenAI implication
Marxian critique	Explains commodity fetishism and how social relations can appear as relations among things.	Use as a critical lens on AI commodities and concealed labor; do not attribute Marx's value theory to Bernácer.
Bernácer's economic lens	Examines finance, speculation, hoarding, productive investment, crisis, and employment.	Ask whether capital and organizational attention support productive activity and human employment or accumulate through destabilizing arrangements.
Savall's SEAM	Integrates social and economic performance through diagnosis, hidden-cost analysis, human-potential development, and negotiated projects.	Study six organizational dysfunction domains; test whether verified dysfunction costs can be reduced and potential developed.
Status-quo shareholder primacy	Treats financial return as the dominant test and may leave social costs outside the organizational account.	Make the boundary of the business case explicit; identify who pays, who benefits, and who can change a decision.
Oligarchic control	Concentrates strategic influence and agenda-setting among a few corporate actors.	Connect dialogue to enforceable decision rights, independent review, public response, and remedy.
Bernácer–Savall orientation	Market-based, socially responsible enterprise with human potential and socio-economic performance treated as interdependent.	Keep economic viability in view while making work, distribution, ecological conditions, participation, and answerability part of the same inquiry.

The label "socially responsible capitalism" is the closest published description of SEAM's economic setting; "Bernácer–Savall socio-economic orientation" names the specific intellectual and practical lineage I am bringing to this case. The second phrase keeps the theory and method visible. It also prevents the chapter from suggesting that one organizational intervention can settle broad political questions of ownership, democratic control, antitrust policy, or public regulation. Those questions remain consequential. This chapter addresses the organizational level at which people can examine decisions, costs, work, and authority, while recognizing the wider legal and political environment in which OpenAI operates.

From Bernácer's economic diagnosis to Savall's SEAM method

Savall's contribution is to connect the socio-economic question to organizational diagnosis and action. In the SEAM account, ordinary financial statements do not necessarily reveal all the costs generated when organizational activity is dysfunctional. A missed handoff can lead to repeated work; inadequate integrated training can make employees dependent on correction and rework; poor working conditions can contribute to absence or turnover; fractured communication can delay decisions; a strategy can remain formally announced but poorly implemented. These are not interchangeable with each other, and they should not be monetized before the underlying activities and consequences are understood. SEAM's "hidden cost" concept directs

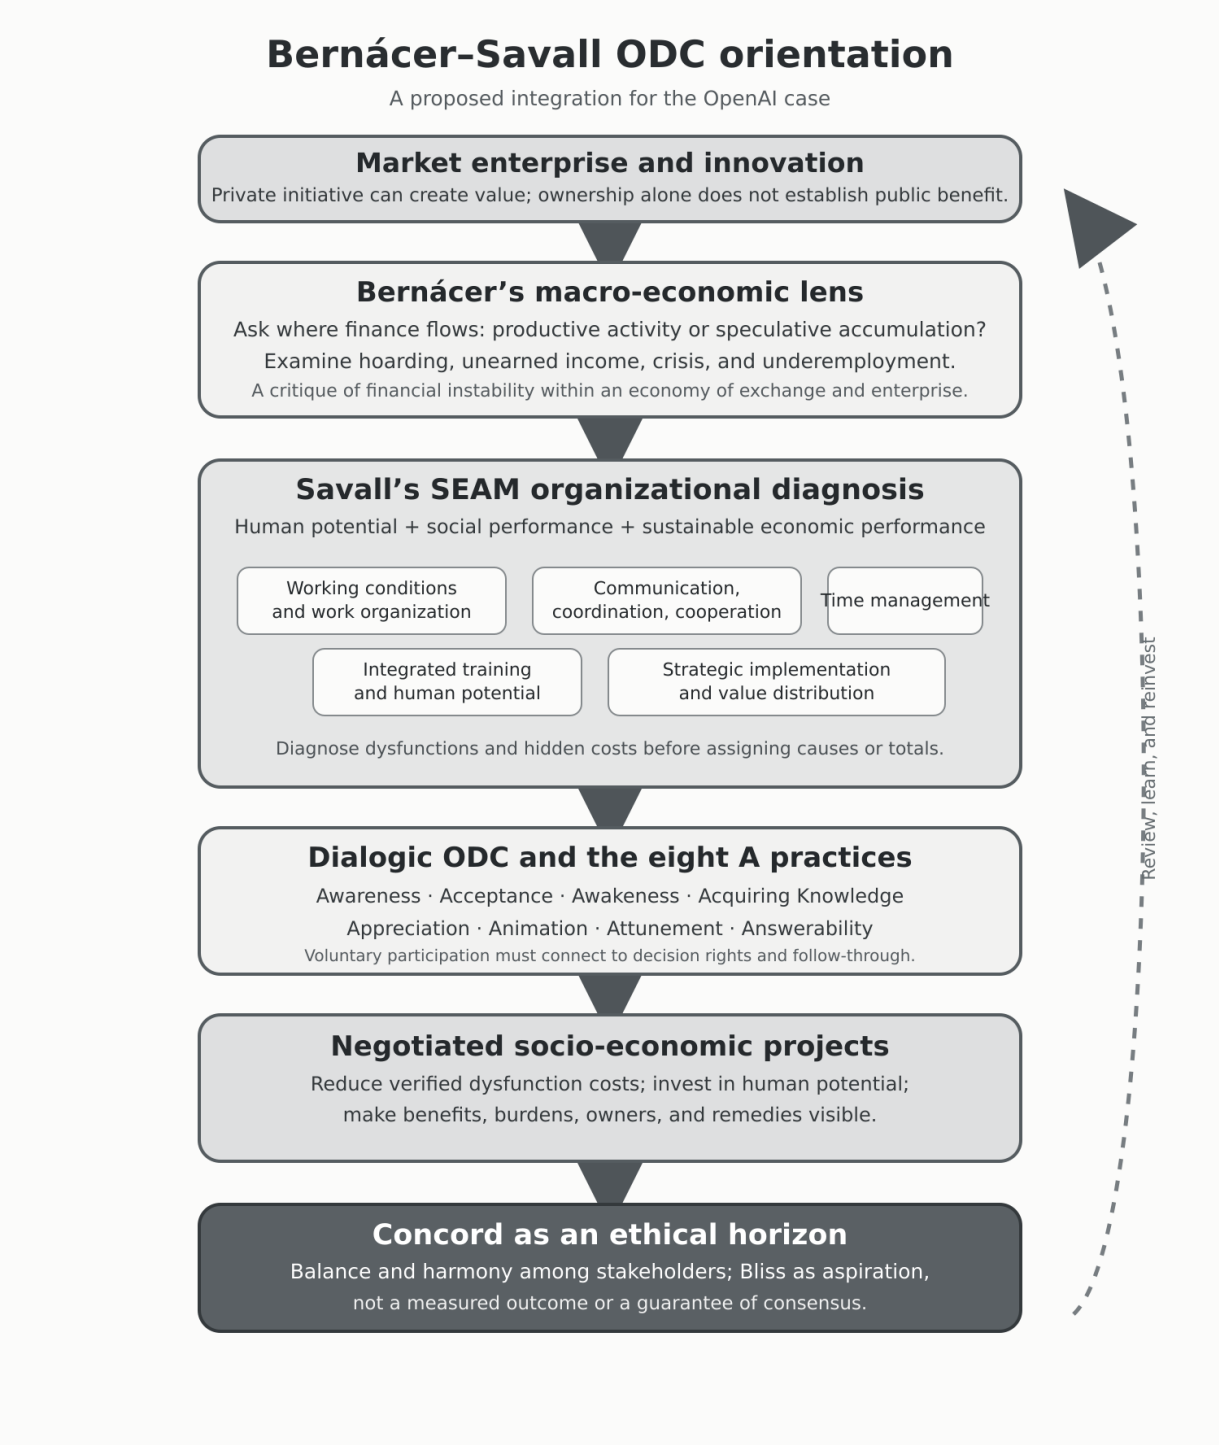
inquiry toward consequences that conventional accounts may fail to capture, while its human-potential concept asks what knowledge, skill, and initiative the organization is not developing (Savall, 2018; Zardet & Bonnet, 2021).

The six commonly identified domains—working conditions, work organization, communication-coordination-cooperation, time management, integrated training, and implementation of strategy—are useful here because they keep diagnosis close to lived work. They do not prove that OpenAI has a particular dysfunction. They give an intervention team a disciplined set of places to look, interview, compare, and verify. A financial estimate becomes credible only after the activity is defined, evidence is collected, the time and cost assumptions are stated, and those closest to the work can challenge the interpretation. A hidden-cost estimate should never be a rhetorical multiplier applied to the whole company without a traceable calculation.

This distinction matters in light of the working materials accompanying the book. Some drafts contain sharply worded interpretations, product scores, probability labels, and a proposed hidden-cost percentage. Those materials can generate questions and hypotheses. The present chapter does not treat a simulated score or a working estimate as an observed organizational fact. To make a socio-economic claim about OpenAI, I would need specified data, time periods, definitions, denominators, source records, and a method for checking alternative explanations. I would also need to distinguish public records from internal information that is not available to this documentary case. Until then, I use the six SEAM domains as a research design, not as a completed diagnosis.

The same boundary applies to the XYZA work. The leadership-level coordinates in this chapter refer to transformation, service and power, polyphony, and moral answerability. The separate transcript rubric in the developing materials uses other labels, including chronotopic alignment and ethics architecture, to examine an AI product's responses. The unit of analysis differs: one rubric makes a provisional judgment about a leader and organization; the other codes text produced by a particular product under a specified prompt. I will not merge their scores or treat a product's sentence as direct evidence of a leader's private character. A future study can compare the frameworks only after the constructs, coding rules, sampled transcripts, model versions, and inter-rater procedures are declared.

Figure 3.5. A proposed Bernácer–Savall socio-economic orientation for OpenAI. The diagram separates Bernácer's macro-economic lens, Savall's organizational diagnosis, Dialogic ODC practices, negotiated projects, and Concord as an ethical horizon. It is a conceptual synthesis for this chapter, not a validated causal model or a political-economic blueprint.



A SEAM-informed diagnostic map for the OpenAI case

For OpenAI, the SEAM domains can turn the general call for "responsible AI" into a set of bounded questions. The table below is not a finding that any listed dysfunction exists. It is a map for a voluntary

research-intervention team to specify the evidence required before drawing a conclusion. Where a public event already exists—such as the November 2023 board crisis—it provides a documented entry point for asking about communication, coordination, authority, and time pressure; it does not establish the current condition of those systems. Any diagnosis would have to include different levels of the organization, multiple roles, and counter-evidence.

Table 3.6. SEAM diagnostic hypotheses and evidence requirements for OpenAI

SEAM domain	Questions for inquiry	Evidence to seek	Boundary on interpretation
Working conditions	What conditions shape the safety, health, workload, autonomy, and security of employees and contracted workers whose work contributes to products? Which tasks are urgent, repetitive, emotionally demanding, or invisible to product accounting?	Work design and workload records where available; voluntary interviews; paid contractor perspectives; workload, absence, turnover, and safety data with privacy protections.	Avoid assuming a specific injury or employment condition. A candidate cost pathway could include rework, absence, turnover, or loss of expertise; each requires actual records and a defined time frame.
Work organization	How are research, product, safety, policy, legal, and commercial responsibilities divided? Who can pause, revise, or escalate a release or deployment decision? Where do handoffs fail or duplicate work?	Decision-rights map; release and escalation procedures; dated process records; interviews across organizational boundaries; cases where a concern changed or did not change a decision.	Do not infer marginalization from the existence of disagreement. Identify authority, workflow, and consequences before labeling a dysfunction.
Communication, coordination, cooperation	What information reaches the board and senior leaders, and when? How are conflicts among safety, product, customer, and infrastructure priorities worked through? Which voices are missing from routine forums?	Meeting and escalation protocols; review of selected public and internal records if voluntarily authorized; follow-up logs; stakeholder mapping; recorded dissent and response.	A candidate cost path is repeated work or delayed decisions. Do not convert the 2023 crisis into a claim about every current relationship.
Time management	How do deadlines, release cycles, emergencies, compute capacity, and safety reviews shape work? Which delays create learning and which create avoidable rework?	Schedule and project data; revision and rollback histories; time-use samples; resource planning; comparable before-and-after periods.	Avoid treating speed as a dysfunction by itself. The question is whether time pressure produces hidden costs, suppresses necessary review, or prevents useful work.
Integrated training	How do employees develop judgment about model behavior, safety, data provenance, human review, and the consequences of deployment? Does the system reward learning or only immediate throughput?	Training plans and participation; mentoring and skill pathways; employee accounts; documented competencies; evidence of learning applied to decisions.	Do not assume deskilling or “metacognitive laziness.” Test whether expertise and learning opportunities are built, displaced, or changed.
Strategic implementation	How do public mission statements, the legal structure, financing plans, product strategy, safety commitments, and infrastructure decisions connect in actual practice? Who tracks commitments and their effects?	Dated strategy and governance materials; budget and capital-allocation decisions where authorized; public-benefit reporting; measures of social and economic performance; stakeholder effects.	Do not infer motive from a corporate form or market valuation. Test alignment between commitments, resources, decisions, and outcomes.

The proposed diagnosis asks the participants to examine both cost and potential. If safety review is delayed by unclear responsibility, for example, the inquiry should identify the actual work sequence, points of delay, and decision owners before estimating a cost. It should also ask which expertise is being developed, who is trusted to exercise it, and what organizational arrangement would allow that expertise to improve outcomes. In this sense, the intervention does not simply count deficiencies. It studies how a dysfunction is produced and what conditions would permit a different pattern of work.

The same discipline applies to environmental and community effects. Data-center electricity, water, land, construction, and local revenue can be part of an OpenAI socio-economic account where the case establishes a connection to the relevant infrastructure and where the data can be attributed responsibly. An infrastructure narrative should identify the site, reporting period, source, and relevant organizational relationship. It should not assign every sector-wide impact to one company without evidence. A local community's knowledge belongs in the inquiry alongside technical records, but neither the company's sustainability statement nor a single critical account should substitute for transparent comparison.

The eight A practices as a pathway from inquiry to Concord

The Concord Tesseract gives the chapter a developmental language; SEAM and Dialogic OD give it organizational work to do. I treat the lower four A practices as capacities for looking inward and learning: Awareness, Acceptance, Awakeness, and Acquiring Knowledge. The upper four—Appreciation, Animation, Attunement, and Answerability—carry learning into relationships, decisions, and obligations. This is a proposed path, not a maturity scale or a claim that every team must proceed in a fixed sequence. A group may return to awareness after learning contradicts its first interpretation, or revisit acceptance when a stakeholder supplies a counter-story. The point is that the shift into outward practice requires more than a new vocabulary.

Table 3.7. Operationalizing the eight A practices in a voluntary OpenAI inquiry

Practice	Possible OpenAI activity	Observable output	Safeguard
Awareness	Map the people, systems, decisions, resources, and stories involved in a selected OpenAI process. Name what is known, absent, and contested.	A bounded issue map with sources, stakeholders, decision owners, and evidence gaps.	Do not treat an executive narrative or a single transcript as the whole system.
Acceptance	Acknowledge the evidence and the presence of affected people without converting acknowledgment into a verdict about motive.	A shared record of confirmed events, unresolved claims, and disagreements that remain open.	Acceptance means contact with reality and others' accounts; it does not mean approval of harm.
Awakeness	Notice assumptions that make a current arrangement appear inevitable: for example, that speed is always progress or that consultation itself equals influence.	A set of assumptions participants can test against cases and counterexamples.	Keep the language invitational; do not presume that participants are consciously deceptive.
Acquiring Knowledge	Combine stories with records, process tracing, SEAM diagnosis, and carefully defined measures.	A jointly reviewed evidence file with methods, limits, and a route for correction.	Do not turn a provisional interpretation into a quantified fact without a reproducible basis.
Appreciation	Find episodes where technical ambition and responsibility reinforced each other, including the conditions that enabled that result.	Documented enabling conditions and accountable exceptions that can be extended.	Invite critical stories and negative cases; appreciation must not screen out pain.
Animation	Co-design and resource specific experiments in work design, review, safety, training, community engagement, or cost transparency.	A small portfolio of time-bounded projects with owners, budgets, and review dates.	Do not declare success at launch; measure whether the experiment changes work or outcomes.
Attunement	Bring distinct organizational and stakeholder perspectives into a continuing relation while preserving disagreement.	A decision process that records which voices were heard, what changed, and why.	Do not manufacture consensus or treat equal airtime as equal decision power.

Practice	Possible OpenAI activity	Observable output	Safeguard
Answerability	Assign people and offices to respond to results, explain decisions, correct errors, and provide remedy where appropriate.	A public or appropriately protected response register, with owners, dates, status, and appeal route.	A promise, policy, or score is not answerability unless it is followed through and can be reviewed.

The eight practices should not become another corporate framework that looks complete because each box has a label. The sequence is useful only when participants can identify the decision an activity touches and the person or body with authority to change it. If an Appreciative Inquiry interview surfaces a moment when a safety objection altered a release decision, the process should record the conditions and decision path that made this possible. If a counter-story shows that an objection did not reach a decision-maker, the record should preserve that as a negative case and ask what the process must change. The upper cycle becomes credible when the organization responds to the lower cycle's learning.

Equanimity as the discipline of invitation

I bring equanimity into this proposal as an ethical discipline for conducting a difficult inquiry. It does not ask the facilitator to become neutral about hidden harm, nor does it ask affected people to remain calm for the comfort of executives. It asks those of us with power in the room to resist vengeance, humiliation, and premature judgment so that the facts and consequences can be examined without turning the encounter into another cycle of violence. It also asks leaders to hear criticism without treating every challenge as an attack on their person. Patience, in this sense, makes sustained answerability more possible; it does not weaken the demand for evidence, repair, or change (Boje, 2026c).

Equanimity also corrects a possible misunderstanding of Concord. Concord is not a vote that erases a minority, a rhetorical midpoint between a corporation and a community, or a finding that all parties share the same interests. The chapter's phrase "balance and harmony among the stakeholders of the AI industry, which equals the Eastern concept of Bliss" names a normative horizon. It is not a validated outcome measure. In a real ODC engagement, the practical question is whether stakeholders gain conditions for well-being and voice, whether economic performance remains viable, whether burdens stop being silently shifted, and whether the organization can respond when a balance proves unjust. Bliss remains a philosophical aspiration, not a number on a dashboard.

A voluntary invitation to OpenAI leadership

I invite Sam Altman, the OpenAI Foundation, the OpenAI Group PBC board, and the leaders responsible for research, safety, product, infrastructure, and people practices to consider a voluntary socio-economic ODC inquiry. The invitation is not a request to endorse my diagnosis or to accept the labels "capitalist oligarchy" as a sufficient account of the organization. It is a proposal to examine together how authority, mission, work, financing, safety, and material infrastructure are connected, and what changes would make OpenAI's economic performance more answerable to the people and places affected by its decisions.

The first meeting would decide whether the inquiry can be genuine. Participants would jointly define its scope, the decisions open to influence, access to evidence, confidentiality protections, independent facilitation, nonretaliation expectations, and how results will be reported. A voluntary invitation does not mean that a company may select only agreeable voices and call the result whole-system dialogue. Employees, safety and research staff, contractors or vendors where relevant, users, customers, and communities connected to infrastructure should have meaningful routes to participate. The group should also identify who is not represented and what would be needed to hear from them safely.

The process would include a written agreement that dialogue is not an admission of liability, a substitute for regulatory oversight, or a public-relations endorsement. It would not ask participants to reveal protected technical information or private personnel data. Independent facilitators would distinguish public evidence, confidential organizational records, and participant testimony. Participants would agree in advance how disagreements will be recorded and who can see the record. If leadership is unwilling to make any decision rights open to review, the facilitator should state that limit plainly and let potential participants decide whether the process is worth entering.

I would ask leaders to make three commitments before convening the full group: first, identify at least one real decision that could change because of the inquiry; second, authorize a protected path for dissent and correction; and third, fund the time, expertise, and data needed to test proposed changes. The invitation should also state what cannot change within the process, such as legal obligations or safety constraints, and why. These terms turn the invitation from an appeal for goodwill into a testable organizational offer. The leaders remain free to decline; the affected parties remain free to disagree; and the inquiry remains responsible for reporting what occurred rather than promising a transformation in advance.

From dialogue to negotiated socio-economic projects

Dialogic practices create conditions for changed meaning and relationship, but meaning alone cannot restructure a work process, budget, review gate, or decision right. SEAM adds the discipline of intervention: participants identify a concrete dysfunction, investigate its root conditions, negotiate a project, and revisit the results. The sequence below keeps the invitation realistic by moving from a bounded inquiry to a limited pilot, then requiring a decision about whether the evidence supports continuation.

Table 3.8. A bounded path from Dialogic OD to SEAM intervention

Phase	Work	Evidence or output	Decision safeguard
1. Contract and scope	Agree on sponsor, independent facilitation, participant protection, evidence access, confidentiality, reporting, and decision rights.	Signed scope; list of open and closed decisions; participant-selection and nonretaliation protocol.	Do not proceed to diagnosis if participants cannot know what is open to change.
2. Map work and voices	Use interviews, process maps, and a World Café to surface the decision sequence, stakeholder accounts, and points of friction.	Voice and decision map; documented exceptions; counter-story and missing-evidence log.	Protect minority accounts and record where participants disagree.

Phase	Work	Evidence or output	Decision safeguard
3. Diagnose with SEAM	Use the six dysfunction domains to trace a bounded process. Test causes through records, observation, and participant review.	A qualitative diagnosis and, only where data permit, a transparent hidden-cost estimate.	No multiplier or monetary estimate without definitions, data, time period, and review.
4. Co-design options	Use Open Space and Future Search to generate alternatives and commitments across roles and stakeholder groups.	Two or three feasible socio-economic projects with owners, resources, risks, and affected groups.	Separate ideas from authorized commitments; make who decides explicit.
5. Pilot and observe	Run a limited project—for example, a changed safety escalation path, integrated training practice, or community cost-disclosure process.	Baseline and follow-up evidence; participant reflections; record of unexpected effects.	Do not equate implementation with benefit; look for displacement or new hidden costs.
6. Review and answer	Return to the group, examine outcomes and negative cases, identify remedies, and decide whether to extend, revise, or stop.	Response register; named owners; timelines; reasons for decisions; accessible appeal route.	Leadership must explain non-adoption and revisit the decision when material evidence changes.

The pilot need not begin with a company-wide redesign. A sharply bounded process—for example, how a safety concern travels from discovery to deployment review—can make the organizational pattern observable. A second project might examine how people working on model evaluation receive training, time, and authority to escalate findings. A third could examine how infrastructure decisions account for local water, electricity, and community effects. The participants should select the first pilot based on evidence and readiness, not on the attractiveness of an outside analyst's preferred example.

For each project, a before-and-after comparison should identify the unit of analysis, the observation period, the data source, and the conditions under which the project will be judged helpful or harmful. A claim that rework fell, for example, needs a consistent definition of rework and an account of whether work shifted to another team. A claim that a new forum increased polyphony needs evidence about who could speak, which voices affected decisions, and whether participants could see a response. A claim that a local benefit improved needs a named community measure and a way for affected people to correct the organization's account. SEAM's economic question and dialogic OD's relational question can then be examined together.

Evidence gates and limits of the claims

I want this proposal to be as scientific as the current evidence allows. The XYZA scores remain interpretive, ordinal starting points. They are not a probability of organizational virtue or a causal measure of a leader's influence. The eight A practices are a conceptual model under development, not an instrument with established reliability. The SEAM tables in this chapter list diagnostic hypotheses, not confirmed OpenAI costs. Each claim would require its own operational definition, evidence, and possible disconfirmation. The visual correspondence between the eight A practices and the socio-economic framework helps organize inquiry; it does not validate the framework by itself.

A research-intervention team should preserve prompt wording, product and model version, dates, source documents, coder instructions, independent scores, disagreements, and adjudication records when it evaluates AI-generated transcripts. A product-response score cannot be transferred automatically to its corporate leader.

If a rubric assesses chronotopic grounding or ethics architecture in model text, it should remain a product-level rubric until researchers establish a justified relationship to leadership behavior. If a proposed calculation uses probabilities or a post-data gain term, its distribution, assumptions, variables, and validation evidence must be published before the result is called a probability or inference. I therefore do not import the working spreadsheet's simulated probabilities into the OpenAI chapter as empirical findings.

The same standard applies to hidden costs. A socio-economic cost estimate should identify an activity and an attributable cost, distinguish observed from imputed values, make the time horizon visible, and explain which costs may be transferred rather than removed. Where the evidence is qualitative, it should remain qualitative. Where data support a range, the range and assumptions should be reported. The goal is not to force every human or ecological consequence into one financial number. The goal is to prevent the visible revenue or infrastructure invoice from being mistaken for a complete account of organizational performance.

These limits do not weaken the invitation. They clarify what an invitation can responsibly promise: a process that makes contested accounts inspectable, makes decision authority visible, tests whether hidden costs can be reduced, and leaves the organization answerable for what it does next. If OpenAI declines, the invitation itself does not count as change. If leaders accept, a summit does not count as change. The evidence of change would be a changed process, a changed right or resource, an acknowledged and remedied harm, a documented creation of human potential, or a transparent explanation for why a proposed change was not adopted.

Discussion

The central organizational story is not simply that Altman was fired and then returned. It is that OpenAI's mission language, legal structure, board authority, employee identity, investor influence, and research culture became entangled in a conflict that the public could see only in fragments. Dialogic OD helps me read this as an organizing problem: the organization's reality was contested through conversations and actions, and the crisis changed who could speak with authority. A new board and a formal review settled immediate governance questions, but they did not by themselves demonstrate that the organization had developed a durable capacity to hold opposition.

The case also complicates claims about "the CEO" as the unit of explanation. Altman's influence matters, but the founding group, board, technical leaders, investors, employees, customers, and legal forms all contributed to OpenAI's path. The departures of Musk, Sutskever, and the Amodeis did not have the same cause or organizational meaning. Their stories should not be compressed into one narrative of betrayal or one narrative of principled exit. Each departure marks a different relationship between authority, mission, and the possibility of continuing inside the organization.

These judgments remain open to revision. A is NE rather than a low numeric score because no reviewed event currently meets the operational threshold. Readers can identify what evidence would change that status.

The Dialogic OD proposal is likewise conditional. Appreciative Inquiry can become image management if it screens out harm. Future Search can produce false consensus if participation does not connect to authority. Open Space can invite topics while leaving decision power unchanged. World Café can turn plural stories into a polished majority summary. Each method needs explicit provisions for dissent, access, confidentiality, and follow-through. The practice is dialogic only when participants can alter the shared account and the decisions that account supports.

The added Bernácer–Savall lens sharpens what the chapter means by organizational change. A responsible market enterprise is not established by a mission statement or a favorable score; it must demonstrate how economic performance relates to work, human potential, decision rights, and the distribution of costs and benefits. For OpenAI, Dialogic OD can open and sustain the conversation, SEAM can organize the diagnosis and negotiated projects, and the eight A practices can make the movement toward Concord legible. Each remains conditional on voluntary participation, usable evidence, and follow-through by people with authority.

What would strengthen the next revision

The present evidence is enough for a provisional case and a practice proposal. A stronger next version would benefit from any additional stories or source material you have already gathered: the exact prompts and outputs from the different AIs used on the prior chapter; public statements or interviews from Altman, board members, employees, or researchers that you want examined; and any firsthand organizational stories that clarify how OpenAI's internal change processes work. I would treat firsthand accounts as attributed evidence and seek corroboration where possible. I do not need private access or a direct response from Altman to continue developing the chapter.

The most useful specific evidence for revising X would be records showing how safety, mission, capital, and listing considerations shaped actual decisions. For Y, I would need evidence of who has authority to participate and change decisions. For Z, I would look for whether dissent is protected and retained in the record. For A, I would need examples of answer, correction, and follow-through to people affected by a decision. Until such evidence is available, the scores remain abductive starting points.

Conclusion

This first draft builds the Chapter 3 case around documented history and Dialogic OD. Altman's path from St. Louis and Stanford through Loopt and Y Combinator helps explain how he entered OpenAI's founding coalition and later became its central executive voice. The records show a contested founding relationship with

Musk, changing financial and legal structures, a board crisis and reversal, and the formation of new AI organizations by former research leaders. They do not authorize a single-person explanation of OpenAI or a claim that one researcher created the LLM.

The provisional X–Y–Z coordinates—4, 2, 3, with A=NE—make the current interpretation inspectable while preserving uncertainty. M is not calculated.

References

Altman, S. (2023a, February 24). Planning for AGI and beyond. OpenAI. https://openai.com/index/planning-for-agi-and-beyond/

Altman, S. (2023b, May 16). Testimony before the U.S. Senate Judiciary Committee, Subcommittee on Privacy, Technology, and the Law: Oversight of A.I.: Rules for artificial intelligence [Hearing testimony]. U.S. Senate. https://www.judiciary.senate.gov/imo/media/doc/2023-05-16%20-%20Bio%20and%20Testimony%20-%20Altman.pdf

Altman, S. (2025a, June 10). The gentle singularity. https://blog.samaltman.com/the-gentle-singularity

Altman, S. (2025b, January 6). Reflections. https://blog.samaltman.com/reflections

Appreciative Inquiry Commons. (n.d.). A brief history. https://aicommons.champlain.edu/learn/appreciative-inquiry-brief-history/

Associated Press. (2024, June 20). OpenAI co-founder Sutskever sets up new AI company devoted to “safe superintelligence.” https://apnews.com/article/c6b48a3675fb3fb459859dece2b45499

Associated Press. (2026, May 18). Federal court rejects Elon Musk’s claims against OpenAI, saying he filed his lawsuit too late. https://apnews.com/article/0b9b0bfaffe96f2c930341f52dfe4f8c

Bakhtin, M. M. (1993). Toward a philosophy of the act (V. Liapunov, Trans.; M. Holquist, Ed.). University of Texas Press. (Original work written c. 1919–1921)

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

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

Blanton, H., & Jaccard, J. (2006). Tests of multiplicative models in psychology: A case study using the unified theory of implicit attitudes, stereotypes, self-esteem, and self-concept. Psychological Review, 113(1), 155–169. https://doi.org/10.1037/0033-295X.113.1.155

Boje, D. M. (1995). Stories of the storytelling organization: A postmodern analysis of Disney as “Tamara-Land.” Academy of Management Journal, 38(4), 997–1035. https://doi.org/10.2307/256618

Boje, D. M. (2001). Narrative methods for organizational and communication research. SAGE.

Boje, D. M. (2026a). The Dark Side of AI Leadership: Ghost vortex, underground consciousness and the accountability architecture of AI power [Manuscript in progress]. Routledge.

Boje, D. M. (2026b, October 1). The appreciative inquiry of appreciative intelligence: AI² in an unprotected world [Working paper]. DavidBoje.com. https://davidboje.com/ws/AI_Squared_Appreciative_Inquiry_Working_Paper_Transitions_Revised.pdf

Boje, D. M. (2026c, October 29). Artificial intelligence capitalism: Sustaining ODC equanimity during AI fetishism [Working paper].

Bordelon, B. (2026, October 4). Sam Altman to Decoded: The world should accept some bad things happening for the benefits of AI. POLITICO. https://www.politico.eu/article/sam-altman-to-decoded-the-world-should-accept-some-bad-things-happening-for-the-benefits-of-ai/

Brown, J., & Isaacs, D. (2005). The World Café: Shaping our futures through conversations that matter. Berrett-Koehler.

Bushe, G. R. (2007). Appreciative inquiry is not (just) about the positive. OD Practitioner, 39(4), 30–35.

Bushe, G. R., & Marshak, R. J. (2009). Revisioning organization development: Diagnostic and dialogic premises and patterns of practice. The Journal of Applied Behavioral Science, 45(3), 348–368. https://doi.org/10.1177/0021886309335070

Bushe, G. R., & Marshak, R. J. (2014). Dialogic organization development. In B. B. Jones & M. Brazzel (Eds.), The NTL handbook of organization development and change (2nd ed., pp. 193–211). Wiley. https://doi.org/10.1002/9781118836170.ch10

Bushe, G. R., & Marshak, R. J. (Eds.). (2015). Dialogic organization development: The theory and practice of transformational change. Berrett-Koehler.

Cappelletti, L., Savall, H., & Zardet, V. (2024). Socio-economic approach to management: Science-based consulting for sustainability. Palgrave Macmillan. https://doi.org/10.1007/978-3-031-43875-2

Cooperrider, D. L. (1985). Appreciative inquiry: Toward a methodology for understanding and enhancing organizational innovation [Doctoral dissertation, Case Western Reserve University].

Cooperrider, D. L., & Srivastva, S. (1987). Appreciative inquiry in organizational life. In R. W. Woodman & W. A. Pasmore (Eds.), Research in organizational change and development (Vol. 1, pp. 129–169). JAI Press.

Cooperrider, D. L., & Whitney, D. (2005). Appreciative inquiry: A positive revolution in change. Berrett-Koehler.

Delaware Department of Justice. (2025, October 28). AG Jennings completes review of OpenAI recapitalization. https://news.delaware.gov/2025/10/28/ag-jennings-completes-review-of-openai-recapitalization/

Fridman, L. (Host). (2023, March 25). Sam Altman: OpenAI CEO on GPT-4, ChatGPT, and the future of AI (No. 367) [Audio podcast episode transcript]. In Lex Fridman Podcast. https://lexfridman.com/sam-altman-transcript/

Fridman, L. (Host). (2024, March 18). Sam Altman: OpenAI, GPT-5, Sora, board saga, Elon Musk, Ilya, power & AGI (No. 419) [Audio podcast episode transcript]. In Lex Fridman Podcast. https://lexfridman.com/sam-altman-2-transcript/

Green Dot Corporation. (2016). Definitive proxy statement (Schedule 14A). U.S. Securities and Exchange Commission. https://www.sec.gov/Archives/edgar/data/1386278/000138627816000093/a2016preliminaryproxystate.htm

Hagey, K. (2023, December 6). Sam Altman: CEO of the year. TIME. https://time.com/6342827/ceo-of-the-year-2023-sam-altman/

Hertz Foundation. (n.d.). Dario Amodei, PhD. https://www.hertzfoundation.org/people/dario-amodei/

OpenAI. (2015, December 11). Introducing OpenAI. https://openai.com/index/introducing-openai/

OpenAI. (2019, March 11). OpenAI LP. https://openai.com/index/openai-lp/

OpenAI. (2020, December 14). Organizational update from OpenAI. https://openai.com/index/organizational-update/

OpenAI. (2023b, November 17). OpenAI announces leadership transition. https://openai.com/index/openai-announces-leadership-transition/

OpenAI. (2023c, November 29). Sam Altman returns as CEO, OpenAI has a new initial board. https://openai.com/index/sam-altman-returns-as-ceo-openai-has-a-new-initial-board/

OpenAI. (2024a, December 13). Elon Musk wanted an OpenAI for-profit. https://openai.com/index/elon-musk-wanted-an-openai-for-profit/

OpenAI. (2024b, February 29). OpenAI and Elon Musk. https://openai.com/index/openai-elon-musk/

OpenAI. (2024c, March 8). Review completed & Altman, Brockman to continue to lead OpenAI. https://openai.com/index/review-completed-altman-brockman-to-continue-to-lead-openai/

OpenAI. (2025a, October 28). Built to benefit everyone: OpenAI's updated structure. https://openai.com/index/built-to-benefit-everyone/

OpenAI. (2025b, October 28). Our structure. https://openai.com/our-structure/

OpenAI. (2026, September 23). Sam Altman's remarks at the United Nations Security Council. https://openai.com/index/sam-altman-un-security-council-remarks/

Owen, H. (2008). Open space technology: A user's guide (3rd ed.). Berrett-Koehler.

Reuters. (2023, November 20). OpenAI staff threaten to quit unless board resigns. https://www.reuters.com/technology/openai-staff-threaten-quit-unless-board-resigns-letter-2023-11-20/

Reuters. (2026a, September 12). OpenAI's Altman won't do IPO this year, calls AI extinction risk "unacceptable." https://www.reuters.com/legal/litigation/openai-ipo-will-not-happen-2026-amid-ai-safety-fears-altman-says-2026-09-12/

Reuters. (2026b, June 8). OpenAI files for U.S. IPO after Anthropic as AI giants head to public markets. https://www.reuters.com/technology/openai-files-us-ipo-after-anthropic-ai-giants-head-public-markets-2026-06-08/

Reuters. (2026c, November 5). OpenAI, ready to IPO in 2027, now counts majority of sales from enterprises, Brockman says. https://www.reuters.com/legal/transactional/openai-ready-ipo-2027-now-counts-majority-sales-enterprises-brockman-says-2026-11-05/

Savall, H. (2018). Radical origins to economic crises: Germán Bernácer, a visionary precursor. Palgrave Macmillan. https://doi.org/10.1007/978-3-319-93366-5

Srivastva, S. (1983). The executive mind: New insights on managerial thought and action. Jossey-Bass.

Srivastva, S. (1986). Executive power: How executives influence people and organizations. Jossey-Bass.

Srivastva, S., & Cooperrider, D. L. (1986). The emergence of the egalitarian organization. Human Relations, 39(8), 683–724. https://doi.org/10.1177/001872678603900801

Stanford Graduate School of Business. (2026). Daniela Amodei says curiosity is underrated. https://www.gsb.stanford.edu/insights/daniela-amodei-says-curiosity-underrated

Sutskever, I. (2013). Training recurrent neural networks [Doctoral dissertation, University of Toronto]. https://www.cs.toronto.edu/~ilya/pubs/ilya_sutskever_phd_thesis.pdf

Taos Institute. (n.d.). Suresh Srivastva. https://www.taosinstitute.net/about-us/people/in-memoriam/suresh-srivastva

Thompson, B. (2025). An interview with OpenAI CEO Sam Altman about building a consumer tech company. Stratechery. https://stratechery.com/2025/an-interview-with-openai-ceo-sam-altman-about-building-a-consumer-tech-company/

Trafimow, D. (2012). The role of auxiliary assumptions for the validity of manipulations and measures. Theory & Psychology, 22(4). https://doi.org/10.1177/0959354311429996

University of Toronto. (2024, June 17). Ilya Sutskever, a leader in AI and its responsible development, receives U of T honorary degree. https://www.utoronto.ca/news/ilya-sutskever-leader-ai-and-its-responsible-development-receives-u-t-honorary-degree

Weisbord, M. R., & Janoff, S. (2010). Future search: Getting the whole system in the room for vision, commitment, and action (3rd ed.). Berrett-Koehler.

Y Combinator. (2016, September 27). Sam Altman on how to build the future. https://www.ycombinator.com/blog/sam-altman-on-how-to-build-the-future

Zardet, V., & Bonnet, M. (2021). Savall, Henri: Connecting the dots between organization development and economic performance. In D. B. Szabla (Ed.), The Palgrave handbook of organizational change thinkers (2nd ed., pp. 1509–1525). Springer. https://doi.org/10.1007/978-3-030-38324-4_55

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