

GENAI MODELS AND THE HYBRID GOVERNANCE TRAP

Ronit Levine-Schnur* & Moran Ofir**

This Article examines the relationship between organizational form and the integrity of knowledge production in generative artificial intelligence (GenAI) firms. As large language models and related technologies increasingly function as public epistemic infrastructures, their institutional architecture becomes a matter of urgent legal and policy concern. Drawing on theories of institutional design and the public-good nature of knowledge, the Article argues that hybrid models, those who ostensibly combine for-profit incentives with public-interest commitments, such as the current form of OpenAI, are structurally flawed. While such entities claim to balance commercial efficiency and ethical governance, they in fact lack enforceable accountability mechanisms and fail to align fiduciary duties with epistemic integrity.

The Article analyzes different organizational forms: traditional for-profit corporations, nonprofit institutions, and hybrid entities, and finds that for-profit models optimize scalability but risk epistemic degradation; nonprofit models safeguard mission fidelity but struggle with capital mobilization; and hybrid models suffer from dual incoherence, inheriting the limitations of both without securing the strengths of either. The Article concludes that preserving the democratic legitimacy and epistemic efficiency of GenAI systems requires legal innovation in organizational form, before structural path dependencies become entrenched.

INTRODUCTION	340
I. INSTITUTIONS OF KNOWLEDGE PRODUCTION	342
A. Knowledge as a Public Good	342
B. Universities as Knowledge Producers	343
C. GenAI Models as Knowledge Producers	345
II. ORGANIZATIONAL MODELS OF KNOWLEDGE PRODUCERS	346
A. Typology	346
B. Organizations in the University World	351

* Tel Aviv University, Faculty of Law. rls@tau.ac.il

** University of Haifa, Faculty of Law. moranofir@law.haifa.ac.il

C. <i>Organizations in the GenAI Industry</i>	352
1. The For-Profit Model	352
2. Nonprofit Model	353
3. Hybrid Model	354
III. IS THE HYBRID MODEL A NECESSARY FAILURE?	355
A. <i>Reconciling Purpose and Incentive</i>	356
B. <i>Voluntary Constraints</i>	357
C. <i>Structural Incoherence and the Limits of Hybridity</i>	360
CONCLUSION: THE HYBRID MODELS TRAP	361

INTRODUCTION

The knowledge infrastructure underlying democratic society is undergoing structural transformation. At the heart of this shift is the development and deployment of generative artificial intelligence (GenAI) systems, whose capacity to create language, images, and code¹ is reshaping the production and dissemination of knowledge. This Article addresses a fundamental institutional question regarding the corporate entities that are currently behind large language models: How does the organizational form of GenAI companies affect the capacity of GenAI models to serve the public good and the future of knowledge production?

Our analysis builds on the growing recognition that knowledge, particularly in the form of search results, generated text, and foundational AI models, is a public good: non-excludable, non-rivalrous,² and vital to democratic functioning. Because of this, the institutions that produce, curate, and distribute such knowledge must be evaluated not only in terms of efficiency, but also in terms of accountability, epistemic integrity, and equitable access. Are GenAI firms equipped to properly address these challenges? To answer this question, we should begin by focusing on the core feature of these firms: their internal organization that allows “some authority (an ‘entrepreneur’) to direct the resources.”³

The current organizational landscape of firms presents three distinct corporate structures: traditional for-profit corporations, non-profit organizations, and benefit or hybrid corporations.⁴ Each of these types represent fundamentally different approaches to organizing economic and social activity. These structures emerge from unique legal foundations and

¹ John T. Kivus, *Generative AI and Copyright Law: A Misalignment That Could Lead to the Privatization of Copyrights Enforcement*, 25 N.C. J.L. & TECH. (2024); Saadia Gabriel et al., *Advancing Equality: Harnessing Generative AI to Combat Systemic Racism*, AN MIT EXPLORATION OF GENERATIVE AI (2024). {\scaps An MIT Exploration of Generative AI} (2024

² Pieter Verdegem, *Dismantling AI Capitalism: The Commons as an Alternative to the Power Concentration of Big Tech*, 39 AI & SOC’Y 727 (2024).

³ R.H. Coase, *The Nature of the Firm*, 4 ECONOMICA 386, 392 (1937).

⁴ Ronald J. Colombo, *Taking Stock of the Benefit Corporation*, 7 TEX. A&M L. REV. 73 (2019).

operational frameworks that shape how they prioritize objectives, allocate resources, and measure success.

Traditional for-profit corporations focus primarily on shareholder wealth maximization, which has been a bedrock feature of modern Anglo-American corporate law.⁵ Operating as entrepreneur-directed systems that coordinate resources through internal hierarchies, corporate officers and directors prioritize the maximization of shareholder value in their oversight and operation of the corporation.⁶ Non-profit organizations are governed by the non-distribution constraint, to pursue social missions through complex funding structures and community-oriented governance. Benefit corporations emerge as a hybrid model, striving to bridge the gap between seeking profits and creating social benefits. They do this by attempting to combine in their for-profit capital structure some public-benefit commitment, declared societal purposes, or an ownership model that ostensibly reflects dual institutional commitment.⁷

Understanding these distinct organizational forms and their underlying legal principles is crucial for analyzing their potential application to GenAI governance, as each embodies different incentive structures, governance mechanisms, and degrees of independence from market logic and standards.

In particular, some argue that the hybrid model, which was adopted by leading GenAI firms such as OpenAI, offers the best of both worlds: innovation at scale, aligned with ethical commitments and public values. This Article challenges that assumption. Drawing on insights from corporate governance theory and empirical studies on the effect of firms' structure and governance on their behavior, we argue that the hybrid model is not a solution but a structural failure: it simultaneously lacks the accountability mechanisms of the non-profit entities and the internal coherence of profit-driven firms. As a result, the only thing it tends to produce are symbolic ethics, reputational safeguards, and organizational opacity, rather than enforceable standards or durable public-interest protections.

We make three contributions in the Article. First, we explain the public good nature of knowledge infrastructure and identify GenAI models' role as knowledge providers. Second, we evaluate the promise and limits of the hybrid model through a legal-institutional lens, focusing on the case of OpenAI. Third, we develop a normative argument that hybrid models are uniquely ill-equipped to sustain the epistemic conditions necessary for trustworthy knowledge production, and that such organizational forms should not relax our concerns.

⁵ *Id.* at 100.

⁶ *Id.*

⁷ Per L. Bylund, *Ronald Coase's "Nature of the Firm" and the Argument for the Economic Planning*, 36 J. HIST. ECON. THOUGHT 305 (2014); DENNIS YOUNG, *IF NOT FOR PROFIT, FOR WHAT?* (2013).

The Article proceeds as follows. Part I defines key concepts in the political economy of knowledge, outlining why GenAI functions as a form of public epistemic infrastructure rather than a purely commercial product. Part II surveys organizational theory to compare the institutional logics of for-profit, nonprofit, and hybrid entities, with an emphasis on their capacity to support epistemic integrity, and applies this framework to the knowledge producers such as universities and the GenAI sector. Part III draws together theoretical critiques and empirical insights, arguing that hybrid forms are structurally incapable of aligning incentives, enforcing accountability, or sustaining public trust. The final Part offers conclusions and some preliminary normative recommendations.

I. INSTITUTIONS OF KNOWLEDGE PRODUCTION

The challenge of governing GenAI requires us to first understand the institutional context in which knowledge is produced. GenAI systems increasingly mediate human understanding and shape the production of new ideas and claims. This Part establishes the conceptual foundation for the Article: it defines the characteristics of knowledge as a public good, introduces the institutional forms historically tasked with producing knowledge in the public interest, and situates GenAI within this tradition. The analysis argues that GenAI development is not merely a technical enterprise, but one that requires legal and institutional design attentive to the governance of epistemic authority.

A. *Knowledge as a Public Good*

Knowledge has long been recognized as a paradigmatic public good.⁸ As Kenneth Arrow famously argued, information is both non-rivalrous and non-excludable: one person's use does not reduce its availability to others, and it is difficult, often undesirable, to prevent others from accessing it.⁹ These characteristics make knowledge difficult to commodify under ordinary market conditions, thereby leading to chronic underproduction unless some form of public intervention or structural subsidy is introduced.¹⁰

Intellectual property ("IP") law seeks to resolve this public product problem by granting exclusive rights to authors and inventors, creating artificial scarcity to incentivize investment.¹¹ Yet the IP framework is

⁸ Paul A. Samuelson, *The Pure Theory of Public Expenditure*, 36 REV. ECON. & STATIS. 387 (1954).

⁹ Kenneth J. Arrow, *The Economic Implications of Learning by Doing*, 29 REV. ECON. STUD. 155 (1962).

¹⁰ Joseph E. Stiglitz, *Knowledge as a Global Public Good*, in GLOBAL PUBLIC GOODS: INTERNATIONAL COOPERATION IN THE 21ST CENTURY 308 (Inge Kaul et al. eds., 1999).

¹¹ WILLIAM M. LANDES & RICHARD A. POSNER, *THE ECONOMIC STRUCTURE OF INTELLECTUAL PROPERTY LAW* (2003).

limited in scope and has unintended consequences: it may promote enclosure of scientific progress, privilege commercially valuable outputs over socially essential ones, and crowd out cumulative collaboration.¹² As legal scholars have noted, the commodification of knowledge can distort both the incentives and content of research.¹³

To correct this market failure, modern societies have developed a range of institutions devoted to the public production of knowledge: universities, research institutes, public libraries, and archival infrastructures. These institutions operate under a distinct normative regime, prioritizing epistemic integrity, transparency, reproducibility, and the pursuit of truth over profit.¹⁴ They are governed by rules and norms that support professional autonomy while ensuring accountability through peer review, funding disclosure, and public mission alignment.¹⁵

B. Universities as Knowledge Producers

Let us look at the example of universities. Universities can be regarded as epistemic infrastructure providers, focused on the creation and dissemination of reliable knowledge as a public good. The institutional structure of universities, whether public or private, is generally non-profit (but for rare examples of for-profit institutions). This structure plays a determinative role in shaping how knowledge is produced, disseminated, and validated. These structures are not merely administrative distinctions; they embody different logics of governance, funding, and accountability that deeply affect the epistemic integrity and public mission of higher education.

American universities originated as independent colleges, governed by religious denominations and private trustees. Over time, these institutions evolved into complex organizations governed through a tripartite system involving governing boards (typically responsible for financial oversight), presidential leadership, and faculty governance bodies. This system institutionalized shared governance, balancing academic freedom with organizational accountability.¹⁶

Universities and research institutions maintain their autonomy through specific governance structures that emphasize educational values over commercial interests, supported by strong safeguards and accountability

¹² Jessica Litman, *The Public Domain*, 39 EMORY L.J. 965 (1990).

¹³ JULIE E. COHEN, *CONFIGURING THE NETWORKED SELF: LAW, CODE, AND THE PLAY OF EVERYDAY PRACTICE* (2012).

¹⁴ Robert K. Merton, *The Normative Structure of Science*, in *THE SOCIOLOGY OF SCIENCE: THEORETICAL AND EMPIRICAL INVESTIGATIONS* 267 (Norman W. Storer ed., 1979).

¹⁵ Jason Owen-Smith, *Managing Laboratory Work through Skepticism: Processes of Evaluation and Control*, 66 AM. SOCIOLOGICAL REV. 427 (2001).

¹⁶ Ronit Levine-Schnur & Moran Ofir, *GenAI Efficiency Paradox* (Sept. 1, 2025) (on file with the authors).

measures.¹⁷ Financial stability, reinforced by active trustee and faculty involvement, prevents compromised decision-making. The Open Access (OA) model further supports this independence by prioritizing academic prestige over monetary incentives. Political independence is preserved through organizational structures embracing certain “anarchic” features, allowing for knowledge production driven by scientific curiosity rather than external pressures. This model operates under a social contract where autonomy is justified by the practical applications of knowledge. Research findings are made publicly accessible, with extended peer communities translating complex findings into practical applications, ensuring public utility. To protect this system, institutions implement governance structures that safeguard academic freedom, transparent disclosure policies, and international collaboration networks, addressing potential “free rider” problems while maintaining research independence.¹⁸

The corporate form of a university affects not only its governance but also the epistemic structure of its research ecosystem. Public and nonprofit universities retain a formal commitment to knowledge as a public good, but economic dependencies, through research funding, commercial partnerships, and donor expectations, have introduced selective pressures that skew research priorities. Universities increasingly compete for prestige and resources, prioritizing research areas that attract external funding or media attention. This leads to underinvestment in non-profitable fields, such as the humanities or public-interest social science. Furthermore, the Bayh-Dole Act of 1980, which allowed universities to patent publicly funded research, has incentivized commercialization over open science.

Without proper mechanisms that assure research integrity and reputational consistency, private donors and corporate funders may exert influence over hiring, curriculum, or publication outputs. For example, fossil fuel companies may fund research centers at elite universities, raising questions about epistemic bias and institutional capture. Faculty can be incentivized to pursue grantable, patentable, or highly-citable research—often at the expense of long-term or exploratory inquiry.

Academic independence aside, both public and private institutions are bound to constitutional limitations. As state agents, all public colleges and universities are legally bound to respect the constitutional rights of their

¹⁷ Silvio O. Funtowicz & Jerome R. Ravetz, *Science for the Post-Normal Age*, 25 *FUTURES* 739 (1993).

¹⁸ Säde Hormio & Samuli Reijula, *Universities as Anarchic Knowledge Institutions*, 38 *SOCIAL EPISTEMOLOGY* 119 (2024); Richard Whitley, *How Do Institutional Changes Affect Scientific Innovations? The Effects of Shifts in Authority Relationships, Protected Space, and Flexibility*, in *ORGANIZATIONAL TRANSFORMATION AND SCIENTIFIC CHANGE: THE IMPACT OF INSTITUTIONAL RESTRUCTURING ON UNIVERSITIES AND INTELLECTUAL INNOVATION* 367 (2014); Paul A. David, *Understanding the Emergence of ‘Open Science’ Institutions: Functionalist Economics in Historical Context*, 13 *INDUS. & CORP. CHANGE* 571 (2004); BRIAN LOASBY, *KNOWLEDGE, INSTITUTIONS AND EVOLUTION IN ECONOMICS* (2002).

students, for example, under the Fourteenth and the First Amendments.¹⁹ Although private universities are not directly bound by the Constitution, when private institutions accept federal funding, they become subject to specific federal laws and regulations, such as Title VI of the Civil Rights Act of 1964.²⁰ The scope of intervention under this and related acts is under significant attention from the current Trump Administration.²¹

To maintain epistemic integrity, universities must protect the independence of research from both political and market pressures. This requires robust governance mechanisms, including:

- Clear separation between funding and research evaluation.
- Transparent conflict-of-interest policies.
- Faculty-led decision-making on research agendas.
- Reinforcement of peer-review and open-access principles.

Thus, independent knowledge production requires not only academic freedom but also institutional design that resists epistemic capture. This includes strong faculty governance, long-term financial stability, and legal constraints against undue influence. Universities are unique in that they operate on a social contract: in exchange for public support or tax-exempt status, they are expected to produce and disseminate knowledge for the common good—not to serve the interests of capital alone.

C. *GenAI Models as Knowledge Producers*

The risks of knowledge-producing institutions have been with us for long, and the GenAI era only elaborates their potential harm. In the context of GenAI, institutions increasingly serve as incubators for proprietary technologies developed in partnership with corporate labs. In this model, the boundaries between public knowledge and private product become blurred, and the institutional structure of the university no longer guarantees the independence or openness of knowledge production.

The phenomenon of technology progresses threatening knowledge production is not new. For instance, one could think about the effect the internet and free knowledge platforms, such as Wikipedia, had on existing knowledge production institutions, such as commercial encyclopedias.

¹⁹ 42 U.S.C. § 2000a (“Prohibition against discrimination or segregation in places of public accommodation”)

²⁰ See *Race, Color, and National Origin Discrimination*, U.S. DEP’T EDUC., <http://www.ed.gov/laws-and-policy/civil-rights-laws/race-color-and-national-origin-discrimination> [<https://perma.cc/X9GQ-B62Q>].

²¹ *Students for Fair Admissions, Inc. v. President & Fellows of Harvard Coll.*, 143 S. Ct. 2141 (2023); *Kestenbaum v. President & Fellows of Harvard Coll.*, No. 24-10092-RGS, 2024 WL 3658793, at *1 (D. Mass. Aug. 6, 2024).

The times when *Britanica* was a must-have in every household were quickly replaced by the free access to Wikipedia.

The shift from traditional encyclopedias to Wikipedia marked a transition to a free, open, and collaborative online knowledge model, moving from static, expert-authored content to dynamic, community-edited articles. While traditional encyclopedias were definitive and reliable, albeit slow to update, Wikipedia offers unparalleled breadth and immediacy, providing a convenient starting point for research by suggesting keywords and references. However, its volunteer-driven nature and lack of a systematic fact-checking process make it a less reliable source for direct citation in academic or professional contexts.

Rapid developments in GenAI are now presenting a new variation of these challenges. While often described as tools or services, GenAI models—such as GPT-4 and Claude—function as autonomous systems of knowledge production. They do not merely transmit or retrieve information, but they generate original language outputs that are interpreted by users as answers, explanations, or summaries. Their outputs are treated as epistemic claims. As such, they exert influence over how individuals and institutions understand the world and make decisions—from legal research to health advice to educational instruction.

This epistemic role demands regulatory attention not only to GenAI's outputs, but to its underlying institutional form. Who builds GenAI, with what incentives, and under what constraints? How are its knowledge claims validated? What epistemic norms, if any, govern its training, output evaluation, and error correction? These questions point to the inadequacy of treating GenAI as a mere product or platform. It must be understood as part of a new institutional landscape of knowledge production—one that demands scrutiny akin to that applied to universities and other knowledge production institutions. Let us now focus on the organizational model aspect of these institutions.

II. ORGANIZATIONAL MODELS OF KNOWLEDGE PRODUCERS

A. *Typology*

The modern organizational landscape presents three distinct corporate structures: traditional for-profit corporations, non-profit organizations, and benefit corporations. Each represents fundamentally different approaches to organizing economic and social activity. These structures emerge from unique legal foundations and operational frameworks that shape how they prioritize objectives, allocate resources, and measure success. Traditional for-profit corporations focus primarily on shareholder wealth maximization, operating as entrepreneur-directed systems that coordinate resources through internal hierarchies rather than market mechanisms. Non-profit organizations, governed by the non-distribution

constraint, pursue social missions through complex funding structures and community-oriented governance. Benefit corporations emerge as a hybrid model, attempting to bridge the gap between seeking profit and creating social benefit through a mandated “triple bottom line” approach requiring them to generate positive economic, social, and environmental outcomes.

Understanding these distinct organizational forms and their underlying legal principles is crucial for analyzing their potential application to artificial intelligence governance, as each structure offers different advantages and limitations in balancing commercial viability with public benefit. This Section examines the key characteristics, operational dynamics, and accountability mechanisms of each corporate form, with particular attention to how their distinctive features might inform the development of institutional frameworks for managing generative AI technologies.

For-Profit Organizations

Historically, traditional corporations have been primarily oriented towards shareholder wealth maximization, which has been a bedrock feature of modern Anglo-American corporate law.²² Under this structure, corporate officers and directors prioritize the maximization of shareholder value in their oversight and operation of the corporation.²³ The operational framework is guided by traditional corporate law principles, where corporations have considerable range to organize themselves and conduct affairs as they and their shareholders see fit.²⁴ Directors typically enjoy tremendous latitude in decision-making, provided they can justify their decisions as rational and in benefit of shareholders.²⁵

For-profit organizations focus primarily on generating profits and maximizing shareholder returns.²⁶ They generally have broader access to capital markets compared to alternative corporate forms, as their singular focus on shareholder returns aligns with traditional investor expectations.²⁷ Accountability mechanisms are primarily oriented toward shareholders, with well-established frameworks for measuring financial performance.²⁸ Traditional corporate law has developed various tools to tackle the agency problem of directors and officers managing property that is not their own.²⁹

Performance measurement in traditional corporations is relatively straightforward, focusing primarily on financial metrics and shareholder

²² Colombo, *supra* note 4, at 100.

²³ *Id.*

²⁴ *Id.* at 108.

²⁵ *Id.*

²⁶ *Id.* at 82–83.

²⁷ *Id.* at 107–08.

²⁸ *Id.* at 105–06.

²⁹ *Id.* at 105.

returns.³⁰ This creates clearer accountability measures compared to alternative corporate forms that must balance multiple objectives.³¹

Market pressures play a significant role in shaping traditional corporate behavior.³² While corporations maintain primary focus on shareholder wealth maximization, they are increasingly responsive to public opinion and social responsibility concerns, often embracing corporate social responsibility initiatives while maintaining their traditional corporate structure.³³ For-profit organizations face specific constraints on their size and scope.³⁴ These limitations arise from diminishing returns to management and competitive pressure from other firms.³⁵

Non-Profit Organizations

Non-profit organizations are defined by their fundamental legal characteristic, the non-distribution constraint, which prohibits them from distributing surplus revenues to individuals who control the organization.³⁶ This structural foundation shapes their entire operational framework and organizational purpose.³⁷ They operate under specific legal requirements including incorporation as non-profit entities and oversight by community-based boards of directors.³⁸

The financial architecture of non-profits is distinctly complex, utilizing multiple revenue streams, such as charitable donations, government grants, service fees, and contracts.³⁹ This diverse funding base supports their mission-driven activities while maintaining their tax-exempt status.⁴⁰ However, non-profits often face significant constraints in accessing capital markets and accumulating financial reserves.⁴¹ Governance in non-profit organizations reflects their community-oriented nature through a shared decision-making model.⁴² This involves oversight by boards of trustees representing community interests and accountability to multiple stakeholders.⁴³ This structure creates more complex but potentially more inclusive organizational dynamics than their for-profit counterparts.⁴⁴

³⁰ *Id.*

³¹ *Id.*

³² *Id.* at 108–09.

³³ *Id.*

³⁴ Bylund, *supra* note 7, at 316.

³⁵ *Id.*

³⁶ YOUNG, *supra* note 7, at 11.

³⁷ *Id.*

³⁸ *Id.*

³⁹ *Id.* at 149–50.

⁴⁰ *Id.*

⁴¹ *Id.*

⁴² *Id.* at 93–95.

⁴³ *Id.*

⁴⁴ *Id.*

The non-profit sector faces unique challenges in monitoring and constraining self-interested behavior among management.⁴⁵ Unlike for-profit entities where financial metrics provide clear accountability measures, non-profit managers can more easily justify suboptimal decisions due to the inherent difficulty in quantifying social impact.⁴⁶ This opacity in performance measurement creates opportunities for inefficient resource allocation, potentially manifesting in preferential program selection, questionable expenditures, and organizational decisions that prioritize management comfort over maximal social impact.⁴⁷

Non-profits encounter significant challenges in comparative analysis and performance measurement that extend beyond typical organizational constraints.⁴⁸ The absence of standardized metrics comparable to profit measurements creates fundamental difficulties in evaluating diverse missions, assessing program effectiveness, and analyzing opportunity costs.⁴⁹ This information asymmetry can lead to suboptimal resource allocation and compromised decision-making processes, particularly when organizations attempt to quantify and compare disparate social impacts.⁵⁰

The donor ecosystem in non-profit organizations presents complex dynamics affecting organizational governance and accountability.⁵¹ Major donors, while wielding significant influence, often lack comprehensive program information necessary for effective oversight.⁵² Simultaneously, modest donors possess limited monitoring capabilities, and the fragmented nature of donor decision-making creates information gaps that can impede efficient resource allocation.⁵³ This asymmetric information environment is further complicated by varying donor motivations, ranging from impact-focused evaluation to purely altruistic intentions.⁵⁴

Hybrid Models and Benefit Corporations

Benefit corporations are characterized by their hybrid nature, combining traditional corporate structure with explicit public benefit objectives.⁵⁵ Their defining legal feature is the mandate to pursue both shareholder interests and public benefits, distinguishing them from traditional corporations focused primarily on profit maximization.⁵⁶

⁴⁵ David M. Schizer, *Enhancing Efficiency at Non-Profits with Analysis and Disclosure*, 11 COLUM. J. TAX L. 76, 87–88 (2020).

⁴⁶ *Id.*

⁴⁷ *Id.*

⁴⁸ *Id.* at 83–84.

⁴⁹ *Id.*

⁵⁰ *Id.*

⁵¹ *Id.* at 101–102.

⁵² *Id.*

⁵³ *Id.*

⁵⁴ *Id.*

⁵⁵ Colombo, *supra* note 4, at 74–75.

⁵⁶ *Id.* at 82–83.

The operational framework of benefit corporations involves a “triple bottom line” approach, requiring them to generate positive economic, social, and environmental outcomes.⁵⁷ They operate under state-specific benefit corporation statutes while retaining most features of traditional corporate law.⁵⁸ Benefit corporations face unique challenges in balancing profit generation with public benefit objectives.⁵⁹ While they can access traditional capital markets, they may encounter difficulties due to their dual mission and that mission’s potential impact on shareholder returns.⁶⁰ Governance in benefit corporations requires directors to balance multiple interests including shareholders, stakeholders, and public benefits.⁶¹ This creates additional complexity in decision-making and potentially increases managerial discretion.⁶² Directors must consider impacts on various constituencies including employees, customers, communities, and the environment.⁶³

Accountability mechanisms include mandatory benefit reporting and third-party assessment requirements.⁶⁴ However, early evidence suggests that compliance with reporting requirements has been “abysmal,” and the lack of uniform standards undermines effective oversight.⁶⁵ The ability to select evaluation standards and limited enforcement mechanisms further complicate accountability.⁶⁶ Performance measurement presents significant challenges due to the need to evaluate both financial returns and public benefits.⁶⁷ The absence of standardized metrics for social and environmental impact creates difficulties in assessing organizational effectiveness.⁶⁸ This is compounded by the lack of clear prioritization among competing objectives.⁶⁹

The benefit corporation structure affects market positioning and capital access.⁷⁰ While potentially appealing to socially conscious investors, benefit corporations may face higher transaction costs and potential disadvantages in mainstream capital markets.⁷¹ However, they

⁵⁷ *Id.* at 100.

⁵⁸ *Id.* at 79.

⁵⁹ *Id.* at 107–08.

⁶⁰ *Id.* at 107.

⁶¹ *Id.* at 85–86.

⁶² *Id.* at 105–06.

⁶³ *Id.* at 85.

⁶⁴ *Id.* at 90–91.

⁶⁵ *Id.* at 101–02.

⁶⁶ *Id.*

⁶⁷ *Id.* at 105.

⁶⁸ *Id.*

⁶⁹ *Id.* at 105–06.

⁷⁰ *Id.* at 107–08.

⁷¹ *Id.*

may gain advantages in reputation and employee/customer engagement through their explicit commitment to public benefits.⁷²

To conclude, each organizational form examined in this Section offers distinct advantages and limitations in balancing commercial viability with public benefit. Traditional for-profit corporations excel at raising capital and driving innovation but face inherent tensions between shareholder interests and broader social considerations. Their well-established accountability mechanisms and clear performance metrics provide straightforward evaluation frameworks, though these may oversimplify complex social impacts. Non-profit organizations, while better positioned to prioritize public benefit through their non-distribution constraint, face significant challenges in accessing capital and measuring performance against diverse social objectives. Their community-oriented governance structures offer valuable models for stakeholder inclusion, though they may struggle with efficiency and scalability. Benefit corporations attempt to bridge these divides through their hybrid structure, but early evidence suggests significant challenges in balancing multiple objectives and maintaining effective accountability. Their experience highlights the persistent difficulties in combining profit-seeking with public benefit interests, even with explicit legal frameworks supporting such integration. Understanding these structural trade-offs is crucial for developing institutional frameworks for GenAI governance, as the technology's unique characteristics and societal implications may require choosing a specific organizational form that is applicable to the unique challenges discussed in this Article.

B. Organizations in the University World

To provide more concreteness, let us look at the actual organization of universities. Public universities, such as UC Berkeley and the University of Texas at Austin, are state-supported and governed with public oversight. Their funding historically relied on direct appropriations and block grants, although they have increasingly turned to tuition and competitive research grants. Public universities aim to serve the public interest, yet they too face growing market pressures and declining state support, affecting their ability to maintain knowledge production as a public good.

Private Nonprofit Universities, such as Harvard, Yale, MIT, and Stanford, maintain autonomy over research, curricula, and admissions (subject to constitutional requirements); they rely heavily on tuition, donations, and endowments. While structurally insulated from shareholder pressure, these institutions are increasingly affected in forming their

⁷² *Id.* at 109.

research agendas by donations’ focus and industry partnerships, particularly in STEM fields.

For-profit universities, like the University of Phoenix, operate with a business logic and investor accountability. These institutions prioritize flexibility, professional training, and customer satisfaction over basic research or public service. Their funding stems primarily from tuition, often supplemented by federal student aid, but their contribution to epistemic infrastructure is minimal. They do not conduct significant public-facing research and instead focus on employment-oriented outcomes.

Each model exhibits a distinct funding profile, governance logic, and epistemic orientation:

Institution Type	Ownership Model	Funding Sources	Research Orientation
Public	State	State appropriations, federal grants, tuition	Basic and applied research
Private Non-Profit	Trustees/Endowment	Donations, endowments, federal research funding	Basic and applied research (with donor influence on research focus)
For-Profit	Shareholders	Tuition, federal aid	Minimal research, applied focus

FIGURE 1. Comparing Institution Type Against Ownership Model, Funding Sources, and Research Orientation.

C. *Organizations in the GenAI Industry*

1. The For-Profit Model

The for-profit model, exemplified by Microsoft’s Copilot and Google’s Gemini, is governed by traditional corporate law principles. Fiduciary duties run to shareholders; profit maximization is the default goal. These companies possess the capital, talent, and compute capacity to lead AI development. However, their business model creates structural incentives to prioritize monetizable outputs over epistemic quality, safety, or access. Moreover, their vertical integration strategies—acquiring AI startups or absorbing them into broader product ecosystems—reduce pluralism and limit contestation in the AI knowledge sphere.

While efficiency and innovation are strong, the democratic costs are high. These include concentration of power over information flows, lack of transparency, and limited public accountability.

2. Nonprofit Model

The nonprofit model, as seen in universities or research institutes, structurally insulates knowledge production from profit incentives. It emphasizes peer review, academic freedom, and the public mission. However, it suffers from slowness, resource dependency, and susceptibility to donor influence. Universities are not immune to capture or epistemic bias, but their institutional logic supports pluralism, contestation, and reproducibility—qualities essential to epistemic democracy.

A recent alternative to commercial language models has emerged in Switzerland, where EPFL, ETH Zurich, and the Swiss National Supercomputing Centre (CSCS) jointly launched Apertus, one of the largest fully open large language models (LLMs) to date. Unlike proprietary systems that disclose only limited aspects of their architecture, Apertus makes its entire development process—including training data, model weights, and recipes—openly accessible and fully documented. The model is distributed under a permissive open-source license and is available through strategic partners such as Swisscom, the AI platform Hugging Face, and the Public AI network. By design, Apertus functions as a public infrastructure resource, enabling developers, researchers, and institutions to build upon and adapt it for a wide range of applications, from chatbots and translation tools to education and scientific research.⁷³

At its launch, the Swiss AI Initiative explicitly framed Apertus as more than a technical achievement: it was presented as a normative alternative embodying values of transparency, sovereignty, inclusivity, and compliance with legal and ethical standards. Trained on fifteen trillion tokens across more than a thousand languages, including underrepresented languages such as Swiss German and Romansh, Apertus seeks to advance linguistic diversity alongside technological innovation. Its developers emphasized that the model was designed with Swiss data protection and copyright law in mind, as well as the transparency obligations of the EU AI Act, ensuring that training relied solely on publicly available, ethically curated data. By positioning Apertus as a model “built for the public good,” Swiss institutions articulated a vision of generative AI not merely as a commercial product but as a form of public infrastructure—akin to highways, electricity, or water—capable of reinforcing democratic values and digital sovereignty.⁷⁴

⁷³ *Apertus: A Fully Open, Transparent, Multilingual Language Model*, ETH ZURICH (Sept. 2, 2025), <https://ethz.ch/en/news-and-events/eth-news/news/2025/09/press-release-apertus-a-fully-open-transparent-multilingual-language-model.html> [<https://perma.cc/U4JF-UW7W>].

⁷⁴ *Id.*

3. Hybrid Model

ChatGPT represents a significant advancement in conversational artificial intelligence, implemented through the Generative Pre-trained Transformer (GPT) architecture. The system's development trajectory encompasses initial deployment with GPT-3.5, and subsequent enhancement through Reinforcement Learning from Human Feedback (RLHF) methodologies, continuing in the integration of GPT-4 capabilities for premium users in March 2023, and progressively improving since then.

The platform employs a dual-tier service model, differentiating between standard and premium access levels. The premium tier incorporates advanced functionalities including web browsing capabilities, DALL-E image synthesis, and experimental plugin integration. A significant enhancement introduced in November 2023 enables user customization for domain-specific applications. The system's architecture facilitates diverse applications, encompassing software development, creative composition, and natural language interaction.

OpenAI declares that its institutional framework represents an innovative synthesis of nonprofit oversight and commercial operations.⁷⁵ Initially established as a nonprofit entity in 2015, the organization underwent structural transformation in 2019 to address escalating resource requirements.⁷⁶ It moved away from its original purely nonprofit model to what it termed a “capped-profit” hybrid structure. Under this new arrangement, the company created OpenAI LP (Limited Partnership) which could accept investments and generate profits, but with a unique constraint: any returns exceeding 100 times the initial investment would be channeled to the overarching nonprofit entity.

This restructuring was explicitly designed to help OpenAI attract the substantial capital needed for its ambitious AI research and development goals, including investments in cloud computing, talent acquisition, and AI supercomputers, while attempting to maintain its original mission-driven ethos. The company essentially acknowledged that its nonprofit status was limiting its ability to compete with major tech corporations in the AI space, while trying to establish a middle ground between pure profit-seeking and its original nonprofit mission.⁷⁷

⁷⁵ *About*, OPENAI, <https://openai.com/about/> [<https://perma.cc/49L4-SGX4>].

⁷⁶ *OpenAI LP*, OPENAI (Mar. 11, 2019), <https://openai.com/index/openai-lp/> [<https://perma.cc/R6W2-V869>].

⁷⁷ Devin Coldewey, *OpenAI Shifts from Nonprofit to “Capped-Profit” to Attract Capital*, TECHCRUNCH (Mar. 11, 2019), <https://techcrunch.com/2019/03/11/openai-shifts-from-nonprofit-to-capped-profit-to-attract-capital/> [<https://perma.cc/BQV6-WAGQ>].

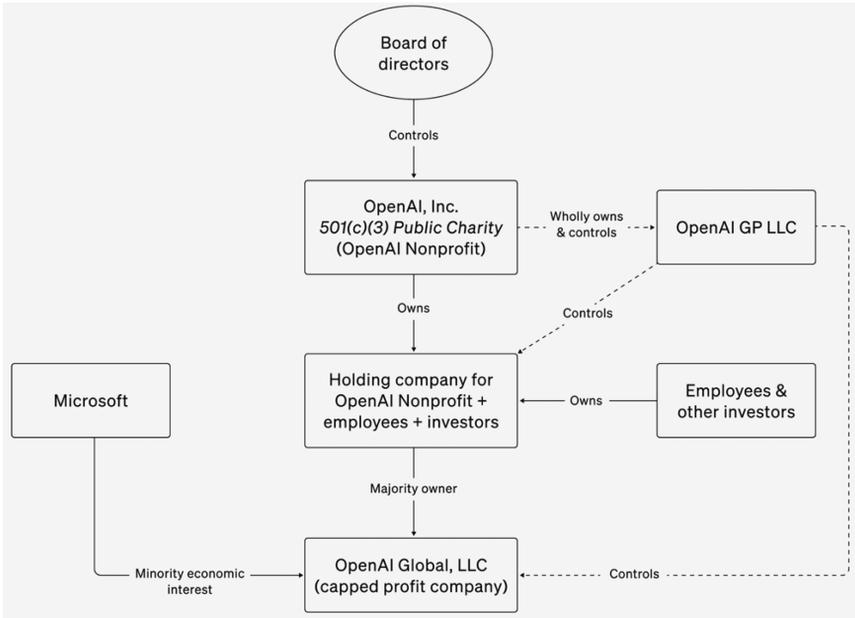


FIGURE 2. Organizational Chart – OpenAI.⁷⁸

The resulting configuration comprises the original nonprofit entity and a capped-profit subsidiary, designed to reconcile substantial capital requirements with ethical imperatives.⁷⁹ Governance is maintained through a board structure incorporating independent directors without equity interests.

The organization's partnership with Microsoft exemplifies its operational model. According to the figure above, Microsoft has minority economic interest in the capped profit company. Currently, despite substantial external investment, OpenAI maintains autonomy through its nonprofit governance structure, with specific provisions excluding artificial general intelligence technologies from commercial licensing agreements.

III. IS THE HYBRID MODEL A NECESSARY FAILURE?

Hybrid organizational forms, those combining for-profit structures with nonprofit or public-purpose elements, have come under increasing scrutiny. As Fisch & Davidoff Solomon and Bebchuk & Tallarita⁸⁰ have recently argued, such hybrid structures often function more as reputational devices than as genuine institutional innovations. They confer a veneer

⁷⁸ Our Structure, OPENAI, <https://openai.com/our-structure/> [<https://perma.cc/7F8M-8E2N>].

⁷⁹ See *OpenAI and Microsoft Extend Partnership*, OPENAI (Jan. 23, 2023), <https://openai.com/index/openai-and-microsoft-extend-partnership/> [<https://perma.cc/8CVQ-NYE9>].

⁸⁰ Jill E. Fisch & Steven Davidoff Solomon, *Should Corporations Have a Purpose?*, 99 TEX. L. REV. 1309 (2020); Lucian A. Bebchuk & Roberto Tallarita, *The Perils and Questionable Promise of ESG-Based Compensation*, 48 J. CORP. L. 37 (2022).

of legitimacy on profit-driven operations while failing to deliver robust public benefits. The result is a governance framework marked by fiduciary fragmentation: directors tasked with representing the public interest are structurally marginalized, while those aligned with investor priorities retain control over operational decision-making.

At its core, the hybrid model inherits the limitations of both traditional nonprofit and for-profit structures while resolving none. It lacks strong fiduciary mandates, enforceable legal constraints, or external accountability mechanisms. Directors appointed to advance the public interest typically lack shareholder backing or financial leverage, rendering their influence largely symbolic. Meanwhile, investor-aligned directors operate with few structural counterweights, ensuring that market logic continues to prevail.

The model also introduces a set of troubling incentive dynamics. By leveraging the symbolic capital of nonprofit or public-purpose commitments, firms such as OpenAI obscure their deep entanglements with monopolistic capital. This strategic ambiguity reinforces information asymmetries and erodes public trust. It could be characterized as regulatory arbitrage, i.e., the practice of taking advantage of a gap between market economics and the regulatory treatment.⁸¹ As Bebchuk and Tallarita have shown, unless market and public goals align perfectly, which is a rare occurrence, mission-oriented rhetoric serves more to conceal than to constrain extractive behavior.⁸²

Let's explore this argument a bit further.

A. *Reconciling Purpose and Incentive*

We have argued that the institutional design of organizations engaged in knowledge production requires careful calibration between autonomy, public purpose, and accountability. However, the challenge facing generative AI institutions is not novel. It reflects a broader dilemma in institutional design: how to balance efficiency and profitability and public purpose. In this regard, two seemingly divergent literatures, on nonprofit enterprises and on ESG-based corporate governance, shed important light on the structural limitations and potential design solutions facing hybrid GenAI entities today.

In their recent article *Purpose and Nonprofit Enterprise*, Professors Cathy Hwang and Dorothy Lund argue that nonprofit organizations provide a compelling model for aligning organizational design with public

⁸¹ Victor Fleischer, *Regulatory Arbitrage*, 89 TEX. L. REV. 227 (2010); Elizabeth Pollman, *Tech, Regulatory Arbitrage, and Limits*, 20 EUR. BUS. ORG. L. REV. 567 (2019); Nizan Geslevich Packin, *Show Me the (Data About the) Money!*, 2020 UTAH L. REV. 1277, 1296 (2020); Ronit Levine-Schnur & Moran Ofir, *Who Shares the Sharing Economy?*, 32 S. CAL. INTERDISC. L.J. 593, 594 (2022).

⁸² Bebchuk and Tallarita, *supra* note 80.

purpose.⁸³ They emphasize that nonprofits are structurally committed to a stable, enduring mission, rather than shareholder returns, which allows them to prioritize long-term, socially valuable outcomes even in the face of financial pressures. Crucial to their argument, this mission-alignment is embedded in legal constraints: so long as nonprofit firms are not allowed to distribute residual profits to private actors, and their boards are not answerable to capital investors. Instead, decision-making authority is allocated in ways that prioritize epistemic integrity and public benefit. As Hwang and Lund put it, “the absence of residual claimants lowers pressure to pursue short-term value and creates space for mission-driven governance.”⁸⁴

This specific organizational architecture, they argue, makes nonprofit entities especially well-suited for domains like education, healthcare, and science, as fields that require investment in uncertain, long-horizon endeavors with diffuse or unmonetizable benefits.⁸⁵ In essence, the nonprofit model structurally supports the production of public goods under conditions where market incentives are misaligned with social goals.

Similarly, the objective of producing reliable and high-quality knowledge constitutes a stable, continuous, and long-term goal that cannot be readily quantified in monetary terms. While trustworthy and comprehensive information undoubtedly possesses significant value, the precise financial worth of such knowledge remains difficult to assess and measure accurately.

B. *Voluntary Constraints*

What about for-profit organizations—are they able to produce such public values? Professors Lucian Bebchuk and Roberto Tallarita offer a sharp critique of recent years’ efforts to embed social purpose into the governance of for-profit corporations via ESG (Environmental, Social, and Governance) metrics.⁸⁶ In *The Perils and Questionable Promise of ESG-Based Compensation*, they demonstrate that ESG-oriented incentive schemes often fail to deliver meaningful accountability. Because ESG metrics are typically vague, non-standardized, and subject to managerial manipulation, they provide managers with discretion rather than discipline. Indeed, Bebchuk and Tallarita argue that ESG-based compensation often serves as a form of reputational “greenwashing” or “purpose washing,” enhancing public image without altering behavior or aligning incentives.

⁸³ Cathy Hwang & Dorothy S. Lund, *Purpose and Nonprofit Enterprise* (Eur. Corp. Governance Inst. Working Paper, Paper No. 819/2024), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5061739.

⁸⁴ *Id.*

⁸⁵ *Id.*

⁸⁶ Bebchuk and Tallarita, *supra* note 80.

Bebchuk and Tallarita's findings underscore a deeper structural point: injecting purpose into corporate governance through soft mechanisms like ESG metrics does not replicate the binding structural constraints of mission-driven entities as documented by Hwang and Lund.⁸⁷ Without credible enforcement mechanisms or ownership structures that prioritize public purpose, the underlying financial logic of shareholder primacy continues to dominate. As such, these efforts may provide only the appearance, not the reality, of mission alignment.

These findings resonate with studies showing that in for-profit corporations, stated commitments to social values are inherently volatile, shifting in response to changes in policy or ownership.⁸⁸ This sensitivity suggests, on the one hand, that the adoption of corporate policies can be effective in shaping practices. On the other hand, it underscores the absence of any structural or organizational commitment to maintaining such policies, which may be altered at the stroke of a pen or through a transfer of ownership.

One illustration for this claim arises from the effects of ownership shifts and policy changes by major social media platforms such as X (formerly Twitter) and Facebook. On October 27, 2022, Twitter Inc. underwent a major ownership change, followed by significant organizational restructuring. Shortly thereafter, the company revised several content-related policies, including reinstating previously suspended accounts, discontinuing certain enforcement guidelines, and dissolving advisory bodies such as the Trust and Safety Council.⁸⁹ This reorientation also marked a break from earlier moderation strategies designed to curb the spread of disputed or unverified information in the wake of major political events.⁹⁰

Burak Özturan and colleagues analyzed the source-quality scores of links shared before and after the ownership change and identified measurable shifts in the quality of information on the platform. Their analysis indicated a decline in the overall quality of circulating content. While they acknowledge that the precise causal mechanisms behind this decline remain to be determined, the consistent pattern underscores the broader impact that platform-level governance changes can exert on the information ecosystem.⁹¹

⁸⁷ Hwang and Lund, *supra* note 83.

⁸⁸ Kevin Reuning, Anne Whitesell & A. Lee Hannah, *Facebook Algorithm Changes May Have Amplified Local Republican Parties*, 9 RESEARCH & POLITICS (2022), <https://journals.sagepub.com/doi/10.1177/20531680221103809> [<https://perma.cc/MKV5-RDEU>].

⁸⁹ *Twitter Abruptly Dissolves Safety Council Moments Before Meeting*, WASH. POST (Dec. 12, 2022), <https://www.theguardian.com/technology/2022/dec/12/twitter-safety-council-dissolved-before-meeting> [<https://perma.cc/KPS3-5Z8K>].

⁹⁰ Stefan D. McCabe et al., *Post-January 6th Deplatforming Reduced the Reach of Misinformation on Twitter*, 630 NATURE 132 (2024).

⁹¹ Burak Özturan et al., *Declining Information Quality under New Platform Governance*, HARV. KENNEDY SCH. MISINFORMATION REV. (July 11, 2025), <https://misinforeview.hks>.

Facebook recently announced a major shift away from content moderation, dismantling the system they built after 2016's controversies.⁹² Until recently, Facebook employed over eighty fact-checking organizations across more than sixty languages.⁹³ Under pressure from political factors and leadership frustrations, Facebook has implemented sweeping changes: ending its fact-checking program, relaxing content rules, and replacing professional moderators with user-generated "community notes."⁹⁴

This dramatic change poses serious risks to democratic discourse. Without professional fact-checkers from trusted journalistic sources like AP and Reuters, false information can propagate rapidly through algorithmic recommendation systems before any corrections appear. The platform's engagement-focused business model creates a "race to the bottom" where sensational or false content often generates more user interaction than accurate information, and that content now faces fewer systematic safeguards.⁹⁵

The situation is exacerbated by Facebook's underlying algorithmic architecture designed to maximize user engagement, which often prioritizes controversial or divisive content over accurate information. By removing professional fact-checking infrastructure and relying primarily on user-generated notes, Facebook is effectively prioritizing engagement metrics over information accuracy and integrity.⁹⁶

This creates a dangerous feedback loop where false or misleading information receives algorithmic amplification, generates high engagement, receives further promotion, and reaches even wider audiences before countervailing corrections can be established. Without independent professional fact-checkers acting as effective epistemic gatekeepers, this cycle threatens to further erode public trust in reliable information sources and deepen social divisions that undermine democratic cohesion.⁹⁷

harvard.edu/article/declining-information-quality-under-new-platform-governance/ [https://perma.cc/7NMC-M4VQ].

⁹² Kate Klönick, *The New Governors: The People, Rules, and Processes Governing Online Speech*, 131 HARV. L. REV. 1598, 1625–30 (2018).

⁹³ Leslie Gielow Jacobs, *Freedom of Speech and Regulation of Fake News*, 70 AM. J. COMPAR. L. i278, i304 (2022).

⁹⁴ THE JOURNAL: *The End of Facebook's Content Moderation Era* (WSJ Podcasts, Jan. 9, 2025), <https://www.wsj.com/podcasts/the-journal/the-end-of-facebooks-content-moderation-era/5985ff84-eb24-42de-b947-31fa44a2e916> [https://perma.cc/UFS7-N8GW]; Meghan Bobrowsky & Gareth Vipers, *Meta Ends Fact-Checking on Facebook, Instagram in Free-Speech Pitch*, WALL ST. J. (2025), <https://www.wsj.com/tech/meta-ends-fact-checking-on-facebook-instagram-in-free-speech-pitch-8e46ad52> [https://perma.cc/H36D-FRT8].

⁹⁵ TARLETON GILLESPIE, CUSTODIANS OF THE INTERNET: PLATFORMS, CONTENT MODERATION, AND THE HIDDEN DECISIONS THAT SHAPE SOCIAL MEDIA 176–82 (2018).

⁹⁶ Bobrowsky & Vipers, *supra* note 94; Zeynep Tufekci, *It's the (Democracy-Poisoning) Golden Age of Free Speech*, WIRED (Jan. 16, 2018), <https://www.wired.com/story/free-speech-issue-tech-turmoil-new-censorship/> [https://perma.cc/K9TS-2KTZ].

⁹⁷ David M.J. Lazer et al., *The Science of Fake News*, 359 SCIENCE 1094, 1095–96 (2018).

Another manifestation of this problem concerns the centrality and market influence of Sam Altman, as a “superstar CEO” of OpenAI. A “superstar CEO,” according to Hamdani & Kastiel is a leader whom directors, investors, and markets believe possesses uniquely valuable qualities—such as rare vision, exceptional skills, charismatic authority, or firm-specific investments—that make them perceived as critical to the company’s success.⁹⁸ Appointed in 2019, Altman was abruptly removed as CEO by the nonprofit board on November 17, 2023, reportedly due to concerns about his leadership. The move triggered immediate backlash, with major investors, most prominently Microsoft, pressuring for his return. Just five days later, on November 22, OpenAI announced Altman’s reinstatement under a restructured board chaired by Bret Taylor, while directors who had supported his removal were dismissed.⁹⁹ Thus, as Chen Wang argues, Altman’s (and other superstar CEOs), as embodiments of the superstar CEO archetype, derive power through non-institutional channels, and institutional channels do not define or actually limit define their operation.¹⁰⁰

C. *Structural Incoherence and the Limits of Hybridity*

There is, we argue, structural dilemma facing hybrid GenAI entities like OpenAI, Anthropic, and others. Such organizations often adopt dual commitments: to public mission and to private investment. But in practice, these goals are not merely in tension—they are structurally incompatible without careful institutional safeguards.

On one hand, the nonprofit model offers strong protection for public-purpose commitments but lacks access to capital and efficiency mechanisms that scale technological deployment. On the other hand, public-purpose infused for-profit models offer capital access and flexibility, but do not successfully bind managerial behavior to public interest.

Hybrid GenAI organizations often seek to combine the best of both worlds. Yet, as Bebchuk and Tallarita demonstrate, adding soft-purpose elements to a profit-maximizing form often results in failure to achieve either goal. Simultaneously, they lack the durable institutional commitments of

⁹⁸ Assaf Hamdani & Kobi Kastiel, *Superstar CEOs and Corporate Law*, 100 WASH. U. L. REV. 1353, 1367-68 (2022).

⁹⁹ Chen Wang, *The Architecture of Control and Fiduciary Limits: Rethinking Superstar CEOs* (July 31, 2025), <https://papers.ssrn.com/abstract=5375271> [<https://perma.cc/2TBK-HQJ9>].”plainCitation:”Chen Wang, *The Architecture of Control and Fiduciary Limits: Rethinking Superstar CEOs* (July 31, 2025 See Ari Levy et al., *OpenAI Brings Sam Altman Back as CEO Less Than a Week After He Was Fired by Board*, CNBC (Nov. 22, 2023), <https://www.cnbc.com/2023/11/22/openai-brings-sam-altman-back-as-ceo-days-after-ouster.html> [<https://perma.cc/PYS4-Z9TF>]; Will Knight & Steven Levy, *OpenAI Staff Threaten to Quit Unless Board Resigns*, WIRED (Nov. 20, 2023), <https://www.wired.com/story/openai-staff-walk-protest-sam-altman/> [<https://perma.cc/X4N8-TKQD>].

¹⁰⁰ Wang, *supra* note 99, at 5.

the nonprofit sector (as identified by Hwang & Lund), while retaining the incentive vulnerabilities of traditional corporate governance. In short, they suffer from dual incoherence: insufficient accountability to either market-based incentives or public-purpose principles.

This suggests a need for rethinking the foundational legal and organizational structures of GenAI institutions. Merely appointing “public interest directors” or codifying non-binding mission statements is unlikely to suffice. Instead, meaningful safeguards may require statutory or charter-based constraints on investor influence, fiduciary duties explicitly oriented toward epistemic integrity, and governance mechanisms that ensure independence in technical decision-making. Without such reforms, the hybrid model may remain an unstable compromise, aspiring to serve the public while structurally bound to extractive logics.

Correspondingly, the non-profit arm of hybrid organizations or the commitments to advancing social objectives in benefit corporations may serve merely as external whitewashing mechanisms, effectively concealing activities that are entirely profit-driven. Such arrangements risk functioning as superficial gestures designed to enhance public perception while the underlying organizational operations remain fundamentally oriented toward financial gain rather than genuine social impact.

In the context of generative AI, the consequences of governance failure are especially acute. These systems are epistemically opaque, socially consequential, and difficult to regulate *ex post*. Without independent oversight and durable institutional commitments to the public good, hybrid GenAI firms are likely to amplify market concentration, erode public trust, and undermine democratic deliberation. The hybrid model, as currently constituted, is not merely imperfect, it is structurally ill-suited to the governance challenges of the AI era.

CONCLUSION: THE HYBRID MODELS TRAP

Concentration in the AI technology stack thus raises significant concerns for democratic governance by influencing the information ecosystem and political power dynamics. Concentration in foundation models and vertically integrated applications can centralize control over the flow of information.¹⁰¹ Few providers of foundation models could shape public discourse by emphasizing or suppressing specific topics, influenced by private or ideological interests. AI tools may also lower the cost for malicious actors to manipulate the information ecosystem, such as through the proliferation of deepfakes.

¹⁰¹ Tejas N. Narechania & Ganesh Sitaraman, *An Antimonopoly Approach to Governing Artificial Intelligence*, 43 *YALE L. & POL. REV.* 95, 142 (“In the AI context, concentration in and across the technology stack raises concerns for the health of our democracy.”).

The hybrid model of GenAI governance, positioned as a compromise between commercial efficiency and public-minded responsibility, has emerged as the dominant organizational form in the sector. Its rhetorical appeal lies in its promise to blend the innovation incentives of private enterprise with the normative constraints of nonprofit or public-benefit institutions. Yet this Article has shown that the hybrid model offers neither meaningful incentives for epistemic quality nor mechanisms of public accountability. It replicates the structural weaknesses of both worlds.

Whereas traditional corporations are bound by fiduciary duties to shareholders and can be disciplined (at least in theory) through market mechanisms, hybrid AI firms dilute these duties without replacing them with enforceable public mandates. Directors appointed to represent “the public interest” in such entities are structurally disempowered: they lack a clear constituency, are shielded from external oversight, and operate within organizations whose resource dependencies and strategic alliances remain profit-driven. As Bebchuk and Tallarita demonstrate in the broader context of ESG governance, attempts to balance social and financial objectives within a single corporate vehicle tend to produce symbolic compliance and reputational window-dressing rather than substantive accountability. The same institutional logic applies to GenAI.

More fundamentally, our analysis reveals a categorical mistake: treating epistemic infrastructure as if it could be governed by organizations whose internal logic is not built to uphold knowledge integrity. When knowledge production is organized through entities whose legal DNA prioritizes competitive advantage and secrecy, the conditions for pluralism, transparency, and trustworthiness are eroded from within. The hybrid model, far from resolving this tension, masks it, allowing tech firms to claim public legitimacy while consolidating control.

We therefore argue that governance reform in this field must begin with institutional design. The choice of organizational form is not neutral. It encodes priorities, authorizes actors, and conditions future path dependencies. If the production and distribution of knowledge is to serve as a public good—one that supports democratic discourse, informed decision-making, and epistemic fairness—then it cannot be left to hybrid structures that escape both market discipline and democratic oversight.

Instead, we suggest two directions for institutional experimentation. First, we call for the creation of public GenAI institutions modeled on structurally independent public universities—funded by the state but insulated from political or commercial interference. Second, we propose a rethinking of corporate law in relation to knowledge-producing firms. Specifically, we support the development of a new category of epistemic fiduciaries, entities legally obligated to uphold standards of epistemic

integrity, transparency, and pluralism, in ways analogous to duties of care and loyalty.

Ultimately, the hybrid model does not fail because it tries to do too much. But because it fails to choose. Its attempt to align two incommensurable goals within a single framework leads to institutional drift, unaccountable discretion, and performative ethics. As the stakes of GenAI governance rise, the cost of this failure will not be borne by shareholders, but by the broader public that depends on reliable, diverse, and transparent knowledge infrastructures.