european corporate governance institute

The Governance of AI Firms

Law Working Paper N° 976/2026
November 2026

Steen Thomsen
Copenhagen Business School and ECGI
Frederik Hovmark Pedersen
ESCP Business School Berlin

© Steen Thomsen and Frederik Hovmark Pedersen 2026. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

This paper can be downloaded without charge from:
https://doi.org/10.67929/ecgi-law-976-2026

european corporate governance institute

ECGI Working Paper Series in Law

The Governance of AI Firms

Working Paper N° 976/2026
November 2026

Steen Thomsen
Frederik Hovmark Pedersen

© Steen Thomsen and Frederik Hovmark Pedersen 2026. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Abstract

The leading developers of artificial intelligence (AI) are among the most valuable and influential companies in the world, yet their corporate governance has received little attention in the literature on AI governance. We contribute to filling this gap by comparing the governance of four leading AI firms (OpenAI, Anthropic, SpaceX and Alphabet/Google) and, more briefly, Mistral, Meta and DeepSeek. Two observations stand out. First, the leading AI firms are all closely held or controlled through dual-class shares, which reflects a broader revival of founder control in tech firms. Second, the firms differ sharply in whether and how their stated purpose is protected. OpenAI is controlled by a nonprofit foundation, Anthropic by founders under a purpose trust with the right to elect a majority of the board, Alphabet by founders with conventional checks and balances, and SpaceX by a single founder with few checks. Purpose-protecting devices - public benefit corporation status, foundation control and purpose trusts - are imperfect and must cope with enormous challenges, but from a governance perspective they are a great deal better than complete autarchy (non-governance). The paper concludes that private-law governance arrangements can be a valuable complement to public regulation.

Keywords: artificial intelligence, corporate governance, founder control, dual-class shares, public benefit corporations, enterprise foundations, OpenAI, Anthropic

Steen Thomsen

Copenhagen Business School and ECGI
ORCID 0000-0002-3709-3883

Frederik Hovmark Pedersen

ESCP Business School Berlin
ORCID 0009-0006-8394-8471

The Governance of AI Firms[1]

Steen Thomsen[2]
Center for Enterprise Foundations
Copenhagen Business School
and ECGI

Frederik Hovmark Pedersen[3]
ESCP Business School Berlin

Working paper 12 November 2026[4]

Abstract

The leading developers of artificial intelligence (AI) are among the most valuable and influential companies in the world, yet their corporate governance has received little attention in the literature on AI governance. We contribute to filling this gap by comparing the governance of four leading AI firms (OpenAI, Anthropic, SpaceX and Alphabet/Google) and, more briefly, Mistral, Meta and DeepSeek. Two observations stand out. First, the leading AI firms are all closely held or controlled through dual-class shares, which reflects a broader revival of founder control in tech firms. Second, the firms differ sharply in whether and how their stated purpose is protected. OpenAI is controlled by a nonprofit foundation, Anthropic by founders under a purpose trust with the right to elect a majority of the board, Alphabet by founders with conventional checks and balances, and SpaceX by a single founder with few checks. Purpose-protecting devices — public benefit corporation status, foundation control and purpose trusts — are imperfect and must cope with enormous challenges, but from a governance perspective they are a great deal better than complete autarchy (non-governance). The paper concludes that private-law governance arrangements can be a valuable complement to public regulation.

Keywords: artificial intelligence; corporate governance; founder control; dual-class shares; public benefit corporations; enterprise foundations; OpenAI; Anthropic.

[1] Comments by Patrick Bolton, Luigi Zingales and Henrik Lando at the Theory and Applications of Contracts Conference, Imperial College Business School, 29–30 June 2026 are gratefully acknowledged.

[2] Novo Nordisk Foundation Professor, Center for Enterprise Foundations, Copenhagen Business School, Denmark, e-mail: st.cefs@cbs.dk. Financial support from the Research Project on Long-Term Ownership and Value Creation in Enterprise Foundations is gratefully acknowledged. The project was funded by Copenhagen Business School and the Augustinus Foundation, the Carlsberg Foundation, the COWI Foundation, the Interogo Foundation, the Lauritzen Foundation, LEO Holding, the Lundbeck Foundation, the Rambøll Foundation, and the William Demant Foundation.

[3] Doctoral Candidate, ESCP Business School Berlin, Germany, e-mail: fhovmarkpedersen@escp.eu.

[4] AI acknowledgement. This paper was written with the assistance of Claude (Anthropic) and ChatGPT (OpenAI), generative AI models, under the direction of the authors. The authors defined the research questions, structure and main arguments, building on an earlier presentation and previous research, and reviewed and revised the text. Facts and references were checked against the sources cited. The authors bear sole responsibility for the content. Anthropic, the developer of Claude, and OpenAI, the developer of ChatGPT, are two of the firms analysed in the paper; but this played no role in the analysis, and the authors have no financial relationship with AI companies.

1. Introduction

Who controls the companies that develop artificial intelligence (AI), and how are they governed? The question has moved from the margins to the centre of public debate. Frontier AI developers now raise capital on a scale previously seen only in sovereign finance: OpenAI closed a funding round of $122 billion at an $852 billion valuation in March 2026, Anthropic raised $65 billion at a $965 billion valuation in May 2026, and SpaceX, after absorbing Elon Musk's AI company xAI, listed on Nasdaq in June 2026 at a valuation well above $1 trillion (Anthropic, 2026e; OpenAI, 2026c; SpaceX, 2026a). At the same time, leading scientists warn that AI systems developed by companies motivated by short-term profits may not give safety the priority it deserves (Bengio et al., 2024), and two former OpenAI directors have argued publicly that AI firms cannot be trusted to govern themselves (Toner & McCauley, 2024). Most of the literature on "AI governance" is about something else: the public and organizational rules that are intended to make AI systems trustworthy, transparent, fair and safe (OECD, 2024; Taeihagh, 2021; Birkstedt, Minkkinen, Tandon, & Mäntymäki, 2023). This literature is valuable, and Section 2 summarizes what it has achieved, but it tends to treat the firms that build AI as black boxes that comply, or fail to comply, with external rules. In contrast, questions about who owns the firm, who appoints the board, what the firm is legally obliged to pursue, and who can change these arrangements, in other words, corporate governance, are typically left out. Yet the governance of AI firms determines who makes the decisions that matter most, such as when to release a model, how much to invest in safety, whom to sell to, and how to share the gains.

For a better understanding of AI corporate governance, this paper first proposes a governance framework that views governance as a hierarchy of connected mechanisms, from regulation through ownership and boards to management and operations (Thomsen & Conyon, 2019). The framework provides a wider menu of governance options for AI than the conventional AI governance literature and serves as a benchmark for the comparison that follows. The paper then documents the purpose, ownership, legal form, owner influence and board composition of the four leading AI firms — OpenAI, Anthropic, SpaceX and Alphabet — and more briefly Mistral, Meta and DeepSeek. The comparison is organized by governance mechanism and summarized in a set of comparative tables.

Two findings stand out. First, all of the leading AI firms are closely held or controlled through dual-class shares. None of them is a widely held corporation in the Berle and Means (1932) sense, with dispersed shareholders. This is not unique to AI. It reflects a broader shift in

corporate governance towards founder control, which has gone furthest in technology (Aggarwal, Eldar, Hochberg, & Litov, 2022; Goshen & Hamdani, 2016). Insider control is to be expected in knowledge-intensive firms whose strategic decisions are hard for outside investors to evaluate, so the relevant question is which checks are combined with it. Second, AI firms differ sharply in whether their stated purpose is protected by governance institutions. OpenAI is controlled by a nonprofit foundation, Anthropic by founders under a purpose trust with the right to elect a majority of its board, Alphabet by its founders under conventional checks and balances, and SpaceX by a single founder with few checks. The devices that protect purpose — public benefit corporation (PBC) status, foundation control and purpose trusts — are largely untested in a modern US context, although foundation ownership has a history of more than a hundred years in Denmark and elsewhere (Thomsen, 2017), and they are under pressure from capital needs and planned stock market listings.

The paper builds on the emerging literature on the governance of AI firms (e.g. Andhov & Murray, 2026; Fan & Nguyen, 2025; Fried & Reiter, 2026; Levine-Schnur & Ofir, 2025; Pedersen & Thomsen, 2026). It connects that literature to the corporate governance view of governance as a system of connected mechanisms and to the broader debate on founder control, and it applies research on enterprise foundations (Hansmann & Thomsen, 2021; Thomsen, 2017) to assess the foundation-like arrangements at OpenAI and Anthropic (see also Pedersen & Thomsen, 2026).

The rest of the paper is organized as follows. Section 2 summarizes the core functions of conventional AI governance, the evidence on how well they work and the emerging literature on the governance of AI firms. Section 3 proposes a governance framework for AI firms and discusses the industry characteristics that shape it. Section 4 places AI firms in the context of rising founder control and documents founder control in the AI industry. Section 5 introduces the firms. Sections 6 to 9 compare purpose, ownership and legal form, owner influence and boards. Section 10 benchmarks the firms against the framework. Section 11 discusses enterprise foundations and AI, and Section 12 concludes with implications for how AI firms should be governed.

2. Conventional AI governance: core functions and how well they work

The OECD defines AI governance as the system of intergovernmental standards and policy frameworks designed to ensure that AI is innovative, trustworthy and respects human rights and democratic values (OECD, 2024). At the level of public policy, the most comprehensive

example is the European Union's AI Act (European Union, 2024a), which imposes risk-based obligations on providers and deployers of AI systems, including specific obligations for general-purpose AI models. At the organizational level, AI governance has been defined as the rules, practices, processes and tools that align an organization's use of AI with its strategy, legal requirements and ethical principles (Mäntymäki, Minkkinen, Birkstedt, & Viljanen, 2022; Papagiannidis, Mikalef, & Conboy, 2025). The research agenda that launched the field (Dafoe, 2018) and later proposals for frontier AI regulation (Anderljung et al., 2023) are framed in similar terms.

Since this body of rules is what most readers have in mind when they hear the term "AI governance", it is useful to summarize its core functions and what is known about how well they work. A review of 84 AI ethics guidelines issued by governments, firms and civil society found convergence on a small set of principles — transparency, justice and fairness, non-maleficence, responsibility and privacy — even if there is divergence in how the principles are interpreted and implemented (Jobin, Ienca, & Vayena, 2019). Building on this convergence, five core functions can be distinguished: accountability, transparency, fairness, privacy and data protection, and safety and security. Table 1 summarizes their aims, the main instruments that implement them, and the available evidence on their effectiveness.

Table 1. Core functions of conventional AI governance and evidence on effectiveness

Function	Aim	Main instruments	Evidence on effectiveness
Accountability	Identifiable people and bodies answerable for AI outcomes	Internal and external audits; ISO/IEC 42001; NIST AI Risk Management Framework; board oversight	Public audits can reduce bias; audit market lacks standards, independence and data access; policies decoupled from practice; few directors with AI expertise
Transparency	Understandable, documented AI systems	Model cards and datasheets; transparency indices; EU AI Act duties for general-purpose AI and Code of Practice	Foundation Model Transparency Index fell from 58 (2024) to 40 (2025); key documentation often left blank; Meta did not sign EU Code; EU enforcement began Aug. 2026
Fairness	No unjustified disparities across groups	Anti-discrimination law; bias audits; NYC Local Law 144	"Null compliance": 18 of 391 employers posted audits; two complaints in two years; weak enforcement
Privacy, data and IP	Lawful use of personal and copyrighted data	GDPR; EDPB guidance on AI models; copyright litigation	Italian fine on OpenAI annulled on jurisdiction (2026); $1.5bn Bartz v. Anthropic settlement; courts bite more than regulators
Safety and security	Prevent severe harms (misuse, loss of control, theft of models)	Red-teaming; frontier safety frameworks; Seoul commitments; California SB 53; New York RAISE Act	Pause commitments weakened under competition; no firm above C+ in independent safety index; evaluations fallible; laws require disclosure, not performance
Overall	From principles to practice	Ethics guidelines; EU AI Act; national AI institutes	Principles–practice gap; EU high-risk rules postponed (2026); US federal deregulation and challenges to state laws

Note. Authors' summary of the sources discussed in Section 2.

2.1 Accountability

Accountability means that identifiable people and bodies are answerable for what AI systems do. It is implemented through assigned roles, internal oversight structures and ex post checks such as audits, assurance and liability. Prominent instruments include end-to-end internal audit frameworks modelled on aerospace, medical-device and financial practice (Raji et al., 2020), layered audits of large language models (Mökander, Schuett, Kirk, & Floridi, 2024), the certifiable management-system standard ISO/IEC 42001 (ISO/IEC, 2023), and the US National Institute of Standards and Technology's AI Risk Management Framework (NIST, 2023). Boards are increasingly involved: about a third of S&P 500 companies disclosed some form of board oversight of AI in 2024 (ISS-Corporate, 2025), and by 2025 more than four out of five S&P 500 companies mentioned AI as a risk factor (Conference Board, 2026).

The evidence that these conventional AI governance structures change outcomes is thin. Public external audits can work. For example, after an audit publicly reported the error rates of commercial facial-analysis systems, the firms concerned substantially reduced their accuracy gaps for darker-skinned women (Raji & Buolamwini, 2019). But the audit market lacks

standards, clear independence requirements and access to data (Birhane, Steed, Ojewale, Vecchione, & Raji, 2024; Costanza-Chock, Raji, & Buolamwini, 2022). Inside firms, responsible-AI practitioners report that "policies, practices, and outcomes are decoupled" (Ali, Christin, Smart, & Katila, 2023). Specifically, ethics loses out to product launches, is hard to quantify, and depends on individual champions who bear personal risk (Ali, Christin, Smart, & Katila, 2023; Rakova, Yang, Cramer, & Chowdhury, 2021). Only 2.7% of S&P 500 directors disclosed AI expertise in 2025 (Conference Board, 2026). In institutional terms, formal accountability structures are spreading rapidly, but much of the adoption appears to be symbolic.

2.2 Transparency and explainability

Transparency means disclosing sufficiently about data, models, capabilities, limitations and risks that users, regulators and the public can understand AI systems and hold them to account, while explainability is the narrower ability to explain how systems produce their outputs. Instruments include model cards (Mitchell et al., 2019) and datasheets for datasets (Gebru et al., 2021), the Stanford Foundation Model Transparency Index (Bommasani et al., 2023), and, since August 2025, the EU AI Act's obligations for providers of general-purpose AI models, supplemented by a voluntary Code of Practice. Most leading developers, including Anthropic, Google, Microsoft, Mistral and OpenAI, signed the Code, while xAI signed only its safety and security chapter, and Meta refused to sign (European Commission, n.d.).

Transparency at the frontier has declined. The average score on the Foundation Model Transparency Index rose from 37 out of 100 in 2023 to 58 in 2024, but fell back to 40 in 2025 (although the indicators and companies assessed changed in 2025, so the scores are not strictly comparable), with firms most opaque about training data, computing resources, usage and societal impact. Meta's score fell from 60 to 31 and Mistral's from 55 to 18 (Wan et al., 2025). An analysis of more than 32,000 model cards found that sections on limitations, evaluation and environmental impact are the ones most often left blank (Liang et al., 2024). Binding enforcement has only just begun: the EU AI Office's enforcement powers over general-purpose AI providers took effect in August 2026.

2.3 Fairness and non-discrimination

Fairness means preventing AI systems from producing unjustified disparities in outcomes or treatment across protected groups. The canonical failure is a widely used US health-care algorithm that used costs as a proxy for medical need and thereby systematically

underestimated the needs of Black patients (Obermeyer, Powers, Vogeli, & Mullainathan, 2019). The main instruments are anti-discrimination law, bias audits and disclosure mandates. New York City's Local Law 144, in force since 2023, requires employers using automated employment decision tools to commission independent bias audits and publish the results (Office of the New York State Comptroller, 2025).

The evidence on such mandates is discouraging. A study of 391 New York employers found that only 18 had posted audit reports and 13 had posted the required notices to candidates, a pattern the authors call "null compliance" (Wright et al., 2024). Audits that were published were limited in scope and of uncertain independence (Groves, Metcalf, Kennedy, Vecchione, & Strait, 2024), and the New York State Comptroller found that the city agency responsible had received only two complaints in two years and, in reviewing 32 companies, had found one instance of non-compliance where state auditors found at least 17 (Office of the New York State Comptroller, 2025). Disclosure-only mandates with weak enforcement have so far achieved little.

2.4 Privacy, data protection and intellectual property

Privacy and data protection concern the lawful basis, transparency and data-subject rights for personal data used to train and deploy AI systems. In practice, the lawful acquisition of copyrighted training data raises closely related issues. The main instruments are the EU's General Data Protection Regulation, guidance from the European Data Protection Board on AI models (EDPB, 2024), and copyright law. Enforcement has been uneven. The Italian data protection authority temporarily banned ChatGPT in 2023 (GPDP, 2023) and later fined OpenAI €15 million (GPDP, 2024), but the fine was annulled by a Rome court in March 2026 on the ground that the Irish authority had become OpenAI's lead regulator under the GDPR's one-stop-shop mechanism (Tribunale di Roma, 2026). Litigation, in contrast, appears to have more effect. In Bartz v. Anthropic, a US federal court held in 2025 that training models on lawfully purchased books was fair use but that building a library from pirated copies was not; the resulting $1.5 billion class settlement received final approval in July 2026 (Bartz v. Anthropic PBC, 2025, 2026). The New York Times' copyright suit against OpenAI and Microsoft, filed in 2023, had reached the summary judgment stage at the time of writing (In re OpenAI, Inc., Copyright Infringement Litigation, 2026). On balance, courts and private litigants appear to discipline AI firms more effectively than regulators do.

2.5 Safety and security

Safety and security concern the prevention of severe harms from capable models, including misuse for biological or cyber attacks, loss of human control, and theft of model weights. The main instruments are pre-deployment testing and red-teaming, firm-level frontier safety frameworks, such as Anthropic's Responsible Scaling Policy, OpenAI's Preparedness Framework and Google DeepMind's Frontier Safety Framework, and international voluntary commitments, such as those made by 16 firms at the Seoul AI Summit in 2024, and, increasingly, state laws. California's Transparency in Frontier Artificial Intelligence Act (SB 53), in force since January 2026, requires large frontier developers to publish safety frameworks and report critical safety incidents (State of California, 2025); New York's RAISE Act, with similar provisions, takes effect in 2027 (New York State, 2026).

The evidence suggests that self-regulation in this area can weaken under competitive pressure. Red-teaming is often vaguely scoped and can "verge on security theater" (Feffer, Sinha, Deng, Lipton, & Heidari, 2024). OpenAI's 2025 Preparedness Framework allows the company to adjust its safeguards if a competitor releases a high-risk system without comparable safeguards (OpenAI, 2025b), and the third version of Anthropic's Responsible Scaling Policy (2026) removed an implied unilateral commitment to pause development and made further safeguards largely conditional on the behaviour of competitors (Anthropic, 2026f; Pedersen & Thomsen, 2026). In the Future of Life Institute's independent AI Safety Index, no company scored higher than C+ overall in 2026, and none scored above D+ on existential safety (Future of Life Institute, 2026). The International AI Safety Report concluded that risk management techniques remain fallible and that some models can detect when they are being tested (Bengio et al., 2026). The new state laws require disclosure rather than specific safety performance.

2.6 How well does conventional AI governance work?

Three general conclusions emerge. First, there is a persistent gap between principles and practice. AI ethics guidelines still lack the shared aims, professional norms, proven implementation methods and accountability mechanisms that make ethical codes effective in fields such as medicine (Mittelstadt, 2019). Moreover, they are largely unenforced (Hagendorff, 2020), vulnerable to "ethics bluewashing" and "ethics shopping" (Floridi, 2019), and have been described as useless in their current form (Munn, 2023). Organizational research on decoupling predicts exactly this pattern (Meyer & Rowan, 1977; Westphal, 2023) of formal structures adopted for legitimacy while core practices remain unchanged (Ali et al., 2023).

Second, the regulatory regime is unstable. The EU's 2026 "digital omnibus" postponed the AI Act's obligations for high-risk systems by between 12 and 16 months, to 2 December 2027 for the systems listed in Annex III and to 2 August 2028 for those covered by Annex I (European Union, 2026). In the United States, the Biden administration's executive order on AI safety was revoked in January 2025, the federal AI Safety Institute was renamed the Center for AI Standards and Innovation (U.S. Department of Commerce, 2025), and a December 2025 executive order directed federal agencies to challenge state AI laws (Executive Order 14365, 2025); so far, the Department of Justice has intervened in one private challenge to Colorado's AI act rather than suing a state directly. The UK's AI Safety Institute was similarly renamed the AI Security Institute in 2025 (DSIT, 2025).

Third, and most important for this paper, conventional AI governance relies heavily on the firms themselves — on their internal accountability structures, voluntary disclosures and self-imposed safety frameworks. Where external enforcement is weak, the effectiveness of AI governance depends on who controls the firms and what they are committed to. This is why the governance of AI firms matters.

2.7 The governance of AI firms: an emerging literature

The governance of AI firms is a neglected topic in the current AI governance literature. It concerns the allocation of control rights over the firms that develop AI, including who owns them, who appoints and removes directors, what the firm's legal purpose is, how the board is composed, and how these arrangements can be changed. This is corporate governance in the ordinary sense (Shleifer & Vishny, 1997), applied to an unusual set of firms. However, the two kinds of governance are connected. Compliance with external rules depends on the internal incentives and authority structures of the firm (Cihon, Schuett, & Baum, 2021), and many proposals for AI risk management, such as ethics boards, internal audit and the "three lines of defence", are ultimately addressed to boards of directors (Schuett, 2025a, 2025b; Schuett, Reuel, & Carlier, 2025). But the governance of AI firms raises questions that compliance frameworks do not ask, such as whether founders or investors should control the firm, whether a nonprofit should own it, and whether its purpose can or should be protected against future owners.

A small but growing literature, mainly in law, addresses these questions. Much of it was prompted by OpenAI's governance crisis in November 2023, when the nonprofit board dismissed the chief executive and then reversed its decision within days under pressure from

employees and Microsoft (Tallarita, 2023; Trautman & Foster, 2025; Wasserman, 2023). Some authors see the crisis as evidence that conventional corporate governance is ill-equipped for firms whose activities pose existential risks (Tallarita, 2023), while others see it as evidence that the nonprofit structure made OpenAI unpredictable and dependent on individual directors (Wasserman, 2023). A Harvard Law Review note applies Hart and Zingales's (2017) concept of "amoral drift" to argue that the main threat to mission comes not from shareholders but from "superstakeholders" such as employees and strategic partners (Harvard Law Review, 2025). Others analyse the control of OpenAI's for-profit subsidiaries by its charity (Aprill, Chan Loui, & Horwitz, 2024), the risks of hybrid governance (Levine-Schnur & Ofir, 2025), and the 2025 recapitalization (Andhov & Murray, 2026; Goldblatt, 2026). Fan and Nguyen (2025) compare the structures of OpenAI, Anthropic and xAI and propose mandatory board-level safety committees. Fried and Reiter (2026) argue that "self-appointed mission guardians" at OpenAI and Anthropic could face the same problems that undermined the independent board of Ben & Jerry's. Oudin and Groza (2024) compare OpenAI's and Anthropic's arrangements with European corporate forms. Weiss (2026) criticizes the efficiency of mission-oriented control, profit caps, and alternative legal forms in AI startups. This paper builds on this literature but takes a broader comparative view.

3. A governance framework for AI firms

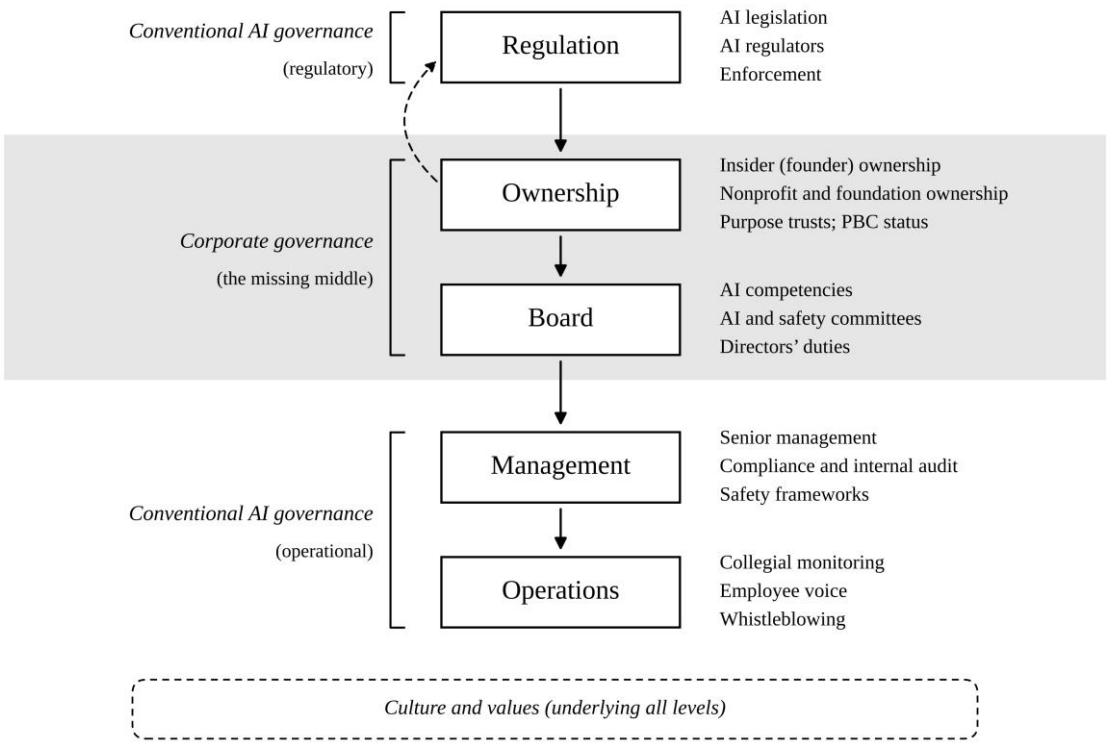
The previous section showed that conventional AI governance relies heavily on the firms themselves. To analyse how those firms are governed, and to ask how they *should* be governed, we need a broader framework than the one used in the AI governance literature. This section proposes one, drawing on an understanding of corporate governance as a bundle of connected governance mechanisms (Thomsen & Conyon, 2019). It then uses the framework to identify what can be expected in AI firms and to set out a menu of governance options against which the leading firms can be compared in the rest of the paper.

3.1 Governance as a hierarchy of mechanisms

Governance can be understood as a set of mechanisms that address agency problems (Jensen & Meckling, 1976; Shleifer & Vishny, 1997). The mechanisms are not independent of one another. Some are logically prior to others (Figure 1). Owners elect the board, and the board hires and fires the chief executive and sets her pay, so ownership is logically prior to the board, and the board is logically prior to management. Managers in turn delegate to lower levels of the organization, where operational decisions are made. All of this takes place within a

framework of law and regulation which firms must largely take as given, at least in the short run, and which is itself shaped by politics. Underlying the whole system are culture and values, which govern much behaviour without anyone having to think about it. This ordering resembles Williamson's (2000) levels of institutional analysis, which run from informal institutions through the formal rules of the game to governance structures and day-to-day resource allocation.

Figure 1. A hierarchy of governance mechanisms for AI firms



Note. Based on Thomsen and Conyon (2019, ch. 3). Solid arrows show the direction of formal authority; the dashed arrow shows intervention from the top of the hierarchy when owners fail (Section 3.1). Shading marks the missing middle (Section 3.4). Corporate governance is used here in the narrow sense of ownership and boards; in a broader sense it also includes executive incentives and internal control (Thomsen & Conyon, 2019).

Each level has a distinct governance function (Thomsen & Conyon, 2019). Political governance and regulation set the rules of the game: company and securities law, investor protection, competition law and sector regulation. They determine which ownership and board structures are available to firms and how well outside investors are protected, and they are themselves shaped by politics (La Porta, Lopez-de-Silanes, Shleifer, & Vishny, 1998; Roe, 2003). Ownership allocates residual control rights and residual claims (Hansmann, 1996). Large owners have incentives to monitor that dispersed shareholders lack (Shleifer & Vishny,

1986), and ownership structures vary systematically with the characteristics of firms and their environment (Demsetz & Lehn, 1985). The board separates decision control from decision management: management initiates and implements decisions, while the board ratifies and monitors the most important of them (Fama & Jensen, 1983). Boards also advise management, and their members are chosen by those who control the firm, so board structure is itself an outcome of ownership (Adams, Hermalin, & Weisbach, 2010). Management makes and implements strategic decisions and designs the internal control system, including risk management, compliance and internal audit. At the operational level, decisions are made by employees whose specific knowledge cannot easily be transmitted upwards, so decision rights are delegated to where the knowledge is (Jensen & Meckling, 1992). Here formal controls reach only part of behaviour, and corporate culture and shared values do much of the governing (Guiso, Sapienza, & Zingales, 2015).

The view of governance as a bundle of interdependent mechanisms has a long history. Jensen (1993) identified four control forces acting on the corporation: capital markets; the legal, political and regulatory system; product and factor markets; and the internal control system headed by the board. Walsh and Seward (1990) distinguished internal from external control mechanisms, and empirical work has shown that mechanisms act as substitutes or complements, so that the effect of any one of them depends on the others (Agrawal & Knoeber, 1996; Rediker & Seth, 1995; Ward, Brown, & Rodriguez, 2009). Organizational approaches add that the effectiveness of a mechanism depends on its fit with the firm's environment and with other mechanisms (Aguilera, Filatotchev, Gospel, & Jackson, 2008), and Aguilera, Desender, Bednar and Lee (2015) argue that external mechanisms such as the legal system, markets and stakeholders must be analysed together with internal mechanisms such as ownership and boards. The hierarchy in Figure 1 builds on this view but adds an ordering: it specifies which mechanisms hold formal authority over which. This matters for AI, where the debate has tended to treat public regulation and firm-level safety practices as separate topics rather than as parts of one connected system.

The hierarchy can be read as successive lines of defence. If social norms fail, law and regulation are a second line. Since much of what firms do cannot be judged by courts or regulators, boards are a third. Board failure is in turn checked by owners, who can replace the board, which makes ownership a fourth line. Finally, if owners fail, the problem returns to the top of the hierarchy, where legislators, regulators and enforcement agencies can change the rules or intervene directly. The "three lines of defence" model, which comes from financial regulation and has

been proposed for AI firms (Schuett, 2025a), covers only the bottom of this hierarchy and depends on the layers above it.

3.2 Industry characteristics and their governance implications

The framework is general, but which governance problems are most severe, and which mechanisms are likely to be adopted, depends on the characteristics of the industry. Governance research has long recognized that one size does not fit all. Firms with more growth opportunities rely more on equity-based pay and less on debt (Smith & Watts, 1992), firms in R&D-intensive industries benefit from boards with more insiders, whose firm-specific knowledge is valuable (Coles, Daniel, & Naveen, 2008), and the value of governance mechanisms depends on the competitive environment (Giroud & Mueller, 2010). Six characteristics of the frontier AI industry stand out: extreme knowledge intensity, powerful externalities, exponential growth, intense competition, enormous capital requirements and rapid technological change. Table 2 summarizes them together with the governance challenges they create and the responses they invite.

Extreme knowledge intensity. Frontier AI development depends on a small pool of specialized researchers and engineers. Research has moved from universities to industry: almost 70 percent of new AI PhDs took jobs in industry in 2020, compared with 21 percent in 2004, and industry produced 96 percent of the largest AI models in 2021 (Ahmed, Wahed, & Thompson, 2023). The movement of AI faculty to industry has had measurable effects on the formation and funding of AI start-ups (Gofman & Jin, 2024). Knowledge of this kind is specific and costly to transfer (Jensen & Meckling, 1992), so outside investors, directors and regulators find it hard to evaluate the firm's strategic decisions. This is the classic setting for insider control: the costs of ownership are lowest for those who understand the technology (Hansmann, 1996), and founders with an idiosyncratic vision may rationally be given control by investors who cannot verify it (Goshen & Hamdani, 2016). Knowledge intensity also shifts power within the firm. When critical resources are embodied in people rather than in physical assets, those people gain bargaining power over the firm and its owners (Rajan & Zingales, 1998; Zingales, 2000), which is the background for the employee influence discussed in Section 3.3.

Powerful externalities. The consequences of frontier AI reach far beyond the firms' customers and investors. Negative externalities include malicious use in cyber or biological attacks, accidents and possible loss of control over advanced systems, effects on employment and income distribution, and effects on information ecosystems and democratic processes

(Acemoglu, 2021; Bengio et al., 2024, 2026; Hendrycks, Mazeika, & Woodside, 2023; Korinek & Stiglitz, 2019). Positive externalities may also be large, since AI is a general-purpose technology whose benefits spread across the economy (Bresnahan & Trajtenberg, 1995). Externalities are the standard case for public regulation, but where regulation is incomplete they raise the question of whether firms should take account of wider interests than those of their shareholders (Hart & Zingales, 2017). This is the economic rationale for purpose-protecting devices such as PBC status, foundation ownership and purpose trusts. It also creates an agency problem between the firm's controllers and the rest of society that is unusually severe in AI (Section 3.3).

Exponential growth. The compute used to train frontier models has doubled roughly every six months since 2010, far faster than Moore's law (Sevilla et al., 2022), and the revenues and valuations of the leading firms have grown at similar rates (Section 1). Rapid growth means that governance must scale. Arrangements designed for a research laboratory with a few hundred employees must work for a firm worth around a trillion dollars, ownership is diluted with each funding round, and informal founder governance gives way to the demands of outside investors and eventually public markets. The benefits of founder control may also change over the life cycle: evidence from dual-class firms suggests that the valuation premium associated with founder control declines as firms mature (Cremers, Lauterbach, & Pajuste, 2024; Kim & Michaely, 2019).

Intense competition. The frontier is contested by a handful of firms backed by the largest technology companies and, increasingly, by states. Economies of scale in compute and data push towards concentration (Korinek & Vipra, 2025), while the rapid diffusion of techniques and open-weight models keeps rivalry intense. Competition cuts both ways for governance. Product market competition disciplines managers and can substitute for governance (Giroud & Mueller, 2010), but it can also erode ethical restraint when restraint is costly (Shleifer, 2004). In a race to develop a powerful technology, competitors have incentives to cut back on safety, and the incentive grows with the number of competitors and the value of winning (Armstrong, Bostrom, & Shulman, 2016; Hendrycks et al., 2023). Unilateral safety commitments are therefore fragile. This increases the value of precommitment devices that are hard to reverse, of industry coordination and of international regulation, although coordination among competitors is itself limited by competition law.

Enormous capital requirements. The amortized cost of training the largest models has grown by about 2.4 times per year since 2016, and training runs costing more than a billion dollars were projected for 2027 (Cottier et al., 2024). Spending on data centres, chips and energy is larger still, and the leading firms now raise single funding rounds of tens of billions of dollars (Section 1). Capital needs of this size cannot be met by founders or by nonprofits. Nonprofits cannot issue equity, so their access to capital is limited to donations, debt and retained earnings (Glaeser & Shleifer, 2001; Hansmann, 1980), which helps explain why OpenAI created first a capped-profit and later a PBC subsidiary. Outside capital dilutes founder control, brings in strategic investors that also supply computing power and distribution, and ultimately points towards a stock market listing. Dual-class shares, founder share classes, purpose trusts and foundation control can all be seen as devices to raise capital without giving up control of the firm or its purpose (Section 4.2).

Rapid technological change. AI capabilities and business models change within months, while legislation takes years. This "pacing problem" (Marchant, Allenby, & Herkert, 2011) is a version of the Collingridge (1980) dilemma: when a technology is young, its effects are hard to predict, and when they have become clear, the technology is entrenched and hard to control. Regulators therefore risk regulating the previous generation of technology. This increases the weight of private governance, but it also means that fixed rules, within firms as well as in law, quickly become obsolete. Typical responses are adaptive instruments, such as safety frameworks that are revised as capabilities develop and principles-based regulation supplemented by codes of practice, together with technical competence on boards. High uncertainty also raises the value of fast decisions and of founders' judgement (Goshen & Hamdani, 2016).

These characteristics interact and pull in different directions. Knowledge intensity, rapid growth and fast technological change favour insider control and fast decisions. Externalities and race dynamics favour purpose protection and external checks. Capital requirements put pressure on both, because outside capital dilutes founder control and outside investors may resist constraints on profit. Much of the governance of AI firms can be read as an attempt to reconcile these pressures, and the leading firms reconcile them in different ways (Sections 4 to 10).

Table 2. Industry characteristics of frontier AI and governance responses

Characteristic	Evidence	Governance challenge	Governance responses	Examples in leading AI firms
Extreme knowledge intensity	Almost 70% of new AI PhDs go to industry; industry produces nearly all of the largest models (Ahmed et al., 2023)	Outsiders cannot evaluate strategy; key employees gain bargaining power	Founder control and dual-class shares; insider-heavy or technically expert boards; equity pay and retention	Founder control at all four firms; employee pressure in the OpenAI 2023 crisis; employee Class E shares at Anthropic
Powerful externalities	Risks of misuse, accidents and loss of control; labour market and democratic effects; large potential benefits (Acemoglu, 2021; Bengio et al., 2026)	Private and social returns diverge; conflict between profit and safety	Public regulation; purpose protection through PBC status, foundation control or purpose trusts; board safety committees; disclosure	OpenAI Foundation control; Anthropic Long-Term Benefit Trust; PBC status at OpenAI and Anthropic; EU AI Act; California SB 53
Exponential growth	Training compute doubling about every six months since 2010 (Sevilla et al., 2022)	Governance must scale from start-up to giant; ownership dilutes; benefits of founder control may decline with maturity	Governance that evolves over the life cycle; sunset clauses; professionalization of boards	Successive restructurings at OpenAI; new share classes in Anthropic's S-1; SpaceX listing
Intense competition	A few frontier firms; scale economies push towards concentration (Korinek & Vipra, 2025)	Race dynamics erode safety; unilateral restraint is costly	Precommitment devices that are hard to reverse; industry coordination; competition law; international coordination	Frontier Model Forum; safety frameworks revised under competitive pressure; CMA and FTC reviews of AI partnerships
Enormous capital requirements	Training costs growing about 2.4 times per year (Cottier et al., 2024)	Founders and nonprofits cannot finance frontier development; outside capital dilutes control	Dual-class and founder share classes; PBC or capped-profit subsidiaries; strategic partnerships; stock market listing	OpenAI's 2025 recapitalization; Microsoft, Amazon and Google stakes; proposed Class F share at Anthropic; SpaceX IPO
Rapid technological change	Regulation lags technology: the pacing problem (Collingridge, 1980; Marchant et al., 2011)	Fixed rules become obsolete; regulators lack information	Adaptive safety frameworks; principles-based regulation and codes of practice; technical expertise on boards	Anthropic Responsible Scaling Policy; OpenAI Preparedness Framework; Google DeepMind Frontier Safety Framework; EU General-Purpose AI Code of Practice

Note. Authors' compilation based on the sources discussed in Section 3.2.

3.3 Agency problems in AI firms

Each link in the hierarchy involves an agency relationship. Citizens delegate authority to legislators and regulators, owners to the board, the board to management, and managers to employees. At each link the agent knows more than the principal and may pursue interests of its own, and the mechanisms at the level above are in part responses to that problem (Fama & Jensen, 1983; Jensen & Meckling, 1976; Thomsen & Conyon, 2019).

These agency relationships are shaped by the industry characteristics discussed in Section 3.2, above all by knowledge intensity, which creates asymmetric information, and by externalities.

Founder control can be understood as an efficient response to asymmetric information because it enables entrepreneurs who possess unique knowledge to control AI firms and provides an economic incentive for them to do so efficiently since they have "skin in the game". Dual-class shares can be seen as an efficient adaptation to reconcile founder control with external financing (Section 4.2). The potential externalities of frontier AI on safety, security, employment and democratic processes invite purpose-protecting mechanisms such as public benefit corporation status, foundation control or purpose trusts. This creates an agency problem of its own, because those charged with protecting the purpose are themselves agents whose interest in safeguarding society cannot be taken for granted (Fried & Reiter, 2026). Proposals for self-regulation by the AI industry can be seen as a response to rapid technological change and intense competition, which imply that regulators may be fighting yesterday's wars.

Agency relationships in AI firms are also unusual in that stakeholders like employees, customers and suppliers can acquire considerable bargaining power over firm principals. Specialized researchers, whose knowledge is scarce and whose outside options are good, or strategic partners that supply capital, computing infrastructure and distribution can shape decisions taken by those with formal authority. Resource dependence theory describes the same phenomenon from the organization's side (Pfeffer & Salancik, 1978), which Pedersen and Thomsen (2026) apply to AI firms, distinguishing formal governance authority from effective purpose control. The OpenAI crisis of November 2023, in which the board's decision to dismiss the chief executive was reversed within days under pressure from employees and Microsoft, is a vivid illustration (Section 4.4).

3.4 The missing middle in AI governance

Mapped onto this hierarchy, conventional AI governance concentrates on the two ends. At the top are legislation such as the EU AI Act, AI regulators and enforcement. At the bottom are compliance functions, audits, model documentation, red-teaming and firm-level safety frameworks — the instruments summarized in Table 1. The middle of the hierarchy, ownership and the board, is largely absent from that literature. Yet the middle determines how the bottom works. Whether a safety framework is maintained under competitive pressure, whether a model release is delayed, and whether a compliance function has authority depend on what the board wants, and the board depends on who owns the firm. If external enforcement at the top is weak

and unstable (Section 2.6), the middle of the hierarchy carries more of the weight than it would in a well-regulated industry. This is the gap that the governance of AI firms fills more or less imperfectly.

3.5 Three principles

Three principles from corporate governance research help to analyse the mechanisms.

Logical priority and its limits. Formal authority flows downwards: owners can change the board, and boards can change management. Ownership is therefore the most powerful lever for anchoring a firm's purpose, which is why the debate on AI firms has focused on who owns them. But as Section 3.3 showed, formal authority can be constrained by agents who control critical resources. The hierarchy is a map of formal control rights, not a guarantee of effective control.

Substitution and complementarity. Governance mechanisms form a system in which some substitute for each other and others are complements (Milgrom & Roberts, 1995). Strong investor protection allows dispersed ownership, while large owners substitute for weak law (La Porta, Lopez-de-Silanes, Shleifer, & Vishny, 1998). Boards matter most when ownership is separated from control, and large owners who run the firm themselves rely less on board monitoring (Fama & Jensen, 1983). Two implications follow for AI. First, founder-controlled AI firms can be expected to have weaker boards in the monitoring sense, because founders substitute for board oversight (and can replace the board at will). Second, public regulation is the first-best response to externalities, but it may fail, for at least three reasons: political capture by vested interests; global competition, which makes unilateral regulation costly and international coordination difficult; and bureaucracy, which makes it hard to regulate a technology that changes faster than legislation. Where public regulation is weak, private mechanisms at the ownership and board levels can partly substitute for it. Arrow (1963) argued that when markets fail, non-market institutions tend to arise to bridge the gap, as medical ethics does for the market for medical care. The purpose-protecting devices adopted by some AI firms can be seen in this light. The substitution is imperfect, however, because private mechanisms bind only the firms that adopt them, not their competitors.

The golden mean. Every governance mechanism has costs as well as benefits, so good governance consists in finding the golden mean for each mechanism in each firm (Thomsen & Conyon, 2019). Up to a point more governance creates value; beyond that point it destroys value. The implication is that firms can be overgoverned as well as undergoverned.

Overgovernance destroys value, for example by taking up management attention, slowing decisions or discouraging initiative (Aghion & Tirole, 1997; Burkart, Gromb, & Panunzi, 1997). The optimum varies between firms and over time, and it may be a corner solution. Two features of AI pull in opposite directions. A fast-moving technology race raises the cost of governance that slows decisions, and overgovernance in one firm may simply shift activity to less governed competitors. But the possibility of severe and irreversible harms raises the cost of undergovernance. Undergovernance announces itself in scandals, whereas overgovernance shows up in decisions that were never taken; in AI, the stakes on both sides are unusually high.

3.6 What to expect: insider governance in high-technology firms

The framework explains what kind of governance to expect in AI firms. Ownership tends to be allocated to the parties for whom the costs of ownership — the costs of controlling managers and of bearing risk — are lowest (Hansmann, 1996). In knowledge-intensive firms where the value of strategic decisions is hard for outsiders to judge, these costs are high for outside investors and low for the founders who understand the technology. Founders may need control to pursue a vision that investors cannot easily verify, and investors may rationally accept this (Goshen & Hamdani, 2016). Founder control in AI is therefore what we would expect, not an anomaly (Section 4). By the substitution principle, founder control will tend to come with less board monitoring. The relevant normative question is therefore not whether AI firms should be founder-controlled but which complementary mechanisms should be combined with founder control as the firms grow, raise outside capital and come to affect society on a large scale. The leading firms give different answers: conventional checks and balances at Alphabet, a purpose trust at Anthropic, foundation control at OpenAI, and few checks at SpaceX.

3.7 A menu of governance options

Table 3 summarizes the governance options available at each level of the hierarchy, their strengths and limitations, and examples from the leading AI firms. It serves as the benchmark for the comparison in Sections 6 to 10.

Table 3. Governance options for AI firms

Level	Options	Strengths	Limitations	Examples in leading AI firms
Political governance and regulation	AI legislation; AI regulators and enforcement; competition law; state oversight of charitable assets; industry codes and soft law	Democratic legitimacy; binds all firms; addresses externalities directly	Political capture; slow relative to technology; jurisdiction-bound; regulatory competition; unstable	EU AI Act obligations for general-purpose AI; California SB 53; New York RAISE Act; CMA and FTC reviews of AI partnerships; Delaware and California Attorney General conditions on OpenAI's recapitalization; Frontier Model Forum
Ownership	Insider (founder) ownership and dual-class shares; nonprofit and foundation ownership; purpose trusts; steward ownership; public benefit corporation status	Strongest lever, since owners appoint the board; long horizon; can protect purpose against short-term investors; founders' knowledge	Entrenchment; key-person risk; capital constraints; arrangements can be amended by those they are meant to constrain; binds only adopters	Founder super-voting at SpaceX, Meta and Alphabet; proposed founder Class F share at Anthropic; OpenAI Foundation control; Anthropic Long-Term Benefit Trust; PBC status at OpenAI and Anthropic
Board	Independent directors and chairs; AI competencies on the board; safety or AI committees; purpose-holder appointments; directors' duties to balance stakeholder interests	Translates ownership into oversight; can intervene in strategy and releases; legal accountability	Information asymmetry; few directors with AI expertise; weak where founders control; dependent on owners for appointment	OpenAI Safety and Security Committee with power to halt releases; Trust-elected directors at Anthropic; independent chair and majority of independent directors at Alphabet
Management	Senior management accountability; frontier safety frameworks; compliance and AI governance functions; internal audit; management standards (ISO/IEC 42001, NIST AI RMF)	Technical expertise; fast; adaptable	Voluntary; revised under competitive pressure; decoupling of policy from practice; limited external verification	Anthropic Responsible Scaling Policy; OpenAI Preparedness Framework; Google DeepMind Frontier Safety Framework
Operations and culture	Culture and values; safety culture; collegial monitoring; employee voice; whistleblowing protection	Governs behaviour that formal rules cannot reach; employees have the best information	Hard to verify; can cut both ways; erodes with rapid growth and stock-based pay	Employee influence in the OpenAI 2023 crisis; employee Class E shares voting on changes to Anthropic's Trust; whistleblower protections in SB 53

Note. Authors' classification, building on Thomsen and Conyon (2019) and the sources discussed in Sections 2 and 4–9.

Read as a menu, the table makes two points. First, the options are wider than those discussed in the AI governance literature, which focuses on the first and the last two rows. Second, the levels interact. Purpose protection at the ownership level is of little value if the board lacks the competence to act on it, and a safety framework at the management level is only as durable as the board's and owners' commitment to it. The following sections examine the governance of leading AI firms level by level using this governance framework, starting with ownership.

4. Founder control: the context and the AI reality

4.1 The return of founder control

The classic description of the large corporation is one of dispersed ownership and managerial control (Berle & Means, 1932). Much of corporate governance theory takes this separation of ownership and control as its starting point (Fama & Jensen, 1983; Jensen & Meckling, 1976). Concentrated ownership has long been common in most of the world, where large firms are typically controlled by founders, families or the state (Faccio & Lang, 2002; La Porta, Lopez-de-Silanes, & Shleifer, 1999). What is new is that founder control has returned in the United States, the home of the widely held corporation, and that it has done so most strongly in technology. Increasing asymmetric information due to differences in technological competence and increasingly abundant sources of external finance can explain why.

Dual-class shares, which give founders and other insiders superior voting rights, illustrate the trend. They were rare among US technology IPOs in the 1980s and 1990s, but in every year since 2015, except 2016, between about 35% and 50% of technology IPOs have had dual-class shares, and in 2025 the share of all IPOs with dual-class shares reached 41%, the highest level in more than four decades of data (Ritter, 2026). Prominent examples include Alphabet (2004), Meta (2012), Snap (2017), which sold shares with no voting rights at all, and most recently SpaceX (2026).

A second trend provides further illustration. Firms stay private longer. The median age of US companies at IPO rose from about five years in 1999–2000 to 12–14 years in 2024–2025 (Ritter, 2026), and the number of US listed companies roughly halved between 1996 and 2012 (Doidge, Karolyi, & Stulz, 2017). Abundant private capital allows firms to grow very large without listing (Ewens & Farre-Mensa, 2020; Kahle & Stulz, 2017), so that ownership stays concentrated in founders and anchor investors for longer. At the end of 2025, there were about 1,290 privately held “unicorns” worth $1 billion or more, with a combined value of about $5.2 trillion; AI companies accounted for a large share of their number and value (WIPO, 2026).

4.2 Why founder control?

Founder control and particularly founder control through dual-class shares have been strongly criticized in the academic literature (Bebchuk & Kastiel, 2017; Gompers, Ishii, & Metrick, 2010; Masulis, Wang, & Xie, 2009), but the fact is that they have increasingly passed the market test in both tech and non-tech industries despite investor scepticism. When founders are substantially more competent and better informed than outsiders, investors may be better

off allowing founders to control firms and passively financing their expansion provided that they are reasonably protected against expropriation (see Adams & Ferreira, 2026, for a review). Governance risks of voting concentration must be weighed against the governance risks of free-rider problems under dispersed ownership. Getting thousands of investors to agree rapidly on a change of direction in a field they basically do not understand may be virtually impossible.

Empirically, founder-led firms have been found to invest more in innovation (Fahlenbrach, 2009), and founder chief executives appear to have a positive causal effect on performance (Adams, Almeida, & Ferreira, 2009). Goshen and Hamdani (2016) argue that founders need control to pursue an "idiosyncratic vision" that outside investors cannot easily verify, and that investors may rationally accept this. Aggarwal et al. (2022) show that the rise of dual-class IPOs is driven by founder-controlled firms and linked to founders' greater bargaining power when private capital is plentiful; founders with more bargaining power obtain larger wedges between votes and equity and are less likely to accept sunset clauses. Aran, Broughman and Pollman (2024) find no evidence that dual-class shares are used to entrench founder CEOs. Hopt and Kalss (2024) point to greater acceptance, perceived advantages for society, and increasing regulatory backing for dual-class shares in Europe.

Some studies have found a firm value discount associated with dual-class shares at least when voting rights substantially exceed cash flow rights (Gompers, Ishii, & Metrick, 2010; Masulis, Wang, & Xie, 2009; Anderson et al., 2025), but others have found a premium in the early years after the IPO albeit declining over time (Cremers, Lauterbach, & Pajuste, 2024; Kim & Michaely, 2019) or a premium rather than a discount in tech companies (Lel et al., 2025) or other firms with high corporate risks (Kim, 2023). Finally, Anderson et al. (2025) find that dual-class shares generate higher risk-adjusted shareholder returns, meaning that investors are amply compensated for governance risk.

Altogether there appears to be no conclusive evidence against founder control or dual-class shares. However, Hettler et al. (2021), Seaborn et al. (2020), and Zhang, Ye, Dong and Wang (2025) point to negative environmental and social externalities of dual-class firms, which could be particularly relevant in AI firms.

4.3 Regulatory and market responses

Regulators and markets have moved in two directions. Institutional investors and proxy advisors have opposed dual-class structures. The Council of Institutional Investors favours one share, one vote and sunset clauses (CII, n.d.), and ISS has recommended voting against directors of companies with unequal voting rights unless the structure has a sunset of seven years or less (ISS, 2025). Index providers initially excluded new multi-class companies. S&P Dow Jones Indices did so in 2017, and FTSE Russell required more than 5% of voting rights to be in the hands of free-float shareholders, though S&P reversed its exclusion in April 2023 (FTSE Russell, 2017; S&P Dow Jones Indices, 2023).

Competition for listings has pushed regulation in the opposite direction. Hong Kong and Singapore introduced weighted voting rights in 2018 (HKEX, 2018; SGX, 2018), the UK liberalized its listing rules for dual-class shares in 2021 and again in 2024 (FCA, 2024), and the EU's Listing Act of 2024 requires member states to permit, by December 2026, multiple-vote shares for companies seeking admission to trading on multilateral trading facilities, including SME growth markets (European Union, 2024b). In the United States, competition between states for incorporations has led to reforms favourable to controlling shareholders. Delaware's 2025 amendments (Senate Bill 21) created safe harbours for transactions with controlling shareholders (State of Delaware, 2025), partly to stem the departure of firms such as Tesla and SpaceX, which reincorporated in Texas in 2024. Texas, in turn, now allows companies to restrict shareholder litigation, for example by requiring a minimum ownership of up to 3% to bring derivative suits (State of Texas, 2025).

4.4 Founder control in AI firms

AI firms follow this trend of founder control, which in the AI industry takes several forms. At SpaceX, Elon Musk holds about 82% of the voting power through ten-vote Class B shares, including about 1.3 billion shares that are unvested but nonetheless give him voting rights (New York Times, 2026; SpaceX, 2026a); one billion of these vest at market capitalization milestones of up to $7.5 trillion combined with the establishment of a permanent human settlement, on Mars, of at least one million people (Fortune, 2026b; SpaceX, 2026a). At Tesla, shareholders approved a pay package worth up to about $1 trillion in November 2025, after Musk had said that he wanted about 25% of the voting power (Tesla, 2025). Musk's companies are thus linked by a controlling shareholder rather than by ownership, and in October 2026 the boards of Tesla and SpaceX formed special committees to review a possible merger (SpaceX, 2026b). At Meta, Mark Zuckerberg holds about 61% of the votes with an economic stake of

about 13% (Meta Platforms, 2026) and has personally directed the company's AI strategy, including the creation of Meta Superintelligence Labs in 2025. At Alphabet, the founders together hold about 53% of the votes (Alphabet, 2026a), and Sergey Brin has returned to hands-on work on the Gemini models (Fortune, 2026c).

The leading independent AI developers are also founder-led. Anthropic's seven co-founders each own only about 1.4% of the equity, but its registration statement provides that a "Founder LLC" of the co-founders will hold a single Class F share carrying 50.1% of the total voting power once the company lists (Anthropic PBC, 2026b; see also Mari, 2026; TechCrunch, 2026). At Mistral, founder shares with ten votes each reportedly give the three co-founders a large share of the votes relative to their capital (Mistral AI, 2025). DeepSeek is controlled by its founder, Liang Wenfeng, who retained a large majority of the equity and nearly all voting rights after the company's first outside funding round in 2026 (Caixin, 2026). Among newer laboratories, Mira Murati's Thinking Machines Lab has reportedly given its founder a board vote equal to all other board votes combined plus one (Fortune, 2025).

OpenAI is the exception that proves the rule. Its chief executive, Sam Altman, reportedly holds no direct equity in the company and is not a controlling shareholder. Yet the November 2023 crisis showed that Altman's de facto power, based on the support of employees and Microsoft, was greater than the formal authority of the board that dismissed him. In this sense, OpenAI illustrates founder control without ownership.

4.5 Implications

Founder control in AI firms is double-edged. On the one hand, it can deliver technology, entrepreneurship and innovation. It is probably no accident that the AI frontier has been pushed by founder-led firms rather than by widely held incumbents with a single share class such as IBM or Microsoft. Concentrated ownership can protect a long-term or safety-oriented purpose against pressure from short-term investors. If founders genuinely care about safety, their control may be the most effective safeguard available. Anthropic's registration statement warns investors that the directors' duty to balance stakeholder interests "may conflict with short-, medium-, or long-term financial interests" and may reduce the value of their shares (Anthropic PBC, 2026b).

On the other hand, founder control can leave one or a few individuals unchecked in firms whose activities may have very large external effects, which raises the question of whether founder control will remain appropriate as AI firms become some of the largest and most systemically

important companies in the world. The governance devices that some AI firms have adopted to protect purpose — public benefit corporation status, foundation control and purpose trusts — can be understood as attempts to retain the advantages of concentrated control while limiting its risks. The following sections examine these devices in four leading firms.

5. The firms

The analysis focuses on four firms: OpenAI (ChatGPT), Anthropic (Claude), SpaceX (Grok, through its subsidiary xAI) and Alphabet (Gemini, developed by Google DeepMind). They include the two leading independent frontier AI developers and two founder-controlled conglomerates in which AI is one of several businesses. Other important developers include Meta (Llama and Muse), Mistral (Le Chat), DeepSeek, Microsoft, Amazon, Cohere and several Chinese technology groups; Mistral, Meta and DeepSeek are discussed more briefly in Sections 4 and 10.

The unit of analysis is the parent company, because that is where control rights reside. This matters for two of the firms. Google DeepMind is a division of Google, which is wholly owned by Alphabet (Alphabet, 2026a). xAI, the developer of Grok, merged with X (formerly Twitter) in 2025 and was acquired by SpaceX in an all-stock transaction effective February 2026; SpaceX listed on Nasdaq in June 2026 (SpaceX, 2026a). For these firms, the governance of AI is the governance of the parent.

Facts concerning ownership, board composition and proposed governance arrangements were checked as of 26 October 2026. They are drawn primarily from company documents, regulatory filings and statements by public authorities, supplemented by reputable press reports where primary sources are not available. Where it matters, we distinguish between three categories of evidence: arrangements that are already legally effective; arrangements that have been formally announced or filed but are awaiting approval or completion; and proposals reported in the press on the basis of unpublished drafts. Anthropic's proposed Class F founder share is the main example of the second category. Ownership and board composition in this industry change rapidly. At the time of writing, OpenAI has confidentially submitted a draft registration statement for an initial public offering (OpenAI, 2026d), while Anthropic, after a confidential submission in June 2026 (Anthropic, 2026g), publicly filed its registration statement on Form S-1 on 26 October 2026 and has applied to list its Class A shares on the Nasdaq Global Select Market under the symbol "ANT" (Anthropic PBC, 2026b). SpaceX and

Tesla have formed special committees to review a possible merger (SpaceX, 2026b). The tables should therefore be read as a snapshot.

6. Purpose

Purpose is set at the top of the corporate governance hierarchy, in the charter that owners adopt and can change. All four firms state a broad purpose (Table 4). OpenAI's mission is "to ensure that artificial general intelligence benefits all of humanity" (OpenAI, n.d.). Anthropic's certificate of incorporation defines its public benefit as "the responsible development and maintenance of advanced AI for the long-term benefit of humanity" (Anthropic, 2023). SpaceX's mission, as stated in its 2026 prospectus, is "to build the systems and technologies necessary to make life multiplanetary", now extended, after the xAI merger, "to understand the true nature of the universe" (SpaceX, 2026a). Alphabet's mission, inherited from Google, is "to organize the world's information and make it universally accessible and useful" (Alphabet, 2026a); Google DeepMind separately describes its aim as building AI responsibly to benefit humanity (Google DeepMind, n.d.).

The legal significance of these statements differs. For OpenAI and Anthropic, both Delaware PBCs, the purpose is written into the certificate of incorporation, and directors are legally required to balance the stockholders' financial interests, the interests of those materially affected by the corporation's conduct, and the specified public benefit (Delaware General Corporation Law §365). In OpenAI's case, the conditions agreed with the Delaware Attorney General go further, and on safety and security matters, the directors must consider only the mission, not the financial interests of stockholders (Delaware Department of Justice, 2025). For SpaceX and Alphabet, ordinary for-profit corporations, the mission statements are not legally binding, and directors' fiduciary duties run to the corporation and its shareholders.

An important difference for the analysis is that the two AI-focused firms define their purpose in terms of the effects of AI on humanity, with explicit reference to safety and broad benefit, while the two conglomerates define their purpose in terms of a business vision. The governance arrangements at OpenAI and Anthropic are thus explicitly designed to protect a purpose that may conflict with profit, while those at SpaceX and Alphabet are designed to protect founder control.

7. Ownership and legal form

Ownership is logically prior to the other corporate governance mechanisms in Figure 1, since owners appoint the board. Table 4 summarizes ownership and legal form.

Table 4. Purpose, ownership and control of four leading AI firms (October 2026)

	OpenAI	Anthropic	SpaceX (incl. xAI)	Alphabet (incl. Google DeepMind)
Main AI product	ChatGPT	Claude	Grok	Gemini
Stated purpose	"To ensure that artificial general intelligence benefits all of humanity"	"The responsible development and maintenance of advanced AI for the long-term benefit of humanity"	"To build the systems and technologies necessary to make life multiplanetary … to understand the true nature of the universe"	"To organize the world's information and make it universally accessible and useful"
Legal form	OpenAI Group PBC: Delaware public benefit corporation, controlled by the OpenAI Foundation (Delaware nonprofit)	Delaware public benefit corporation (since 2021)	Texas corporation (reincorporated from Delaware, Feb. 2024)	Delaware corporation
Major owners	At recapitalization (Oct. 2025): OpenAI Foundation 26% (plus warrant); Microsoft c. 27%; employees and investors 47%. Diluted by $122bn round (Mar. 2026)	Amazon (est. mid-to-high teens %); Google c. 14% (capped at 15%); seven co-founders c. 10% combined (each pledged 80% to philanthropy); Amazon and Google each >5%, held as non-voting Class N shares; venture investors, employees	Elon Musk c. 82% of votes; c. 48% incl. unvested awards and options; venture investors; Alphabet 7.2% of Class A shares	Larry Page and Sergey Brin c. 53% of votes combined (much smaller economic stake); institutional investors
Control mechanisms	Foundation appoints and removes all PBC directors; Safety and Security Committee can halt model releases; Delaware and California AG conditions	Long-Term Benefit Trust (Class T shares) has the right to elect a board majority; Trust appointees held 4 of 7 seats (Apr. 2026), 3 of 6 after a May 2026 resignation (one Trust seat vacant per S-1); proposed Founder LLC with a Class F share carrying 50.1% of total voting power (S-1, Oct. 2026; effective only on listing)	Dual-class shares (10 votes per Class B share); Class B elects board majority; no sunset	Dual-class shares (Class A 1 vote, B 10 votes, C none)
Listing status	Private; confidential IPO filing (June 2026)	Private; public S-1 filed 26 Oct. 2026 (confidential submission June 2026); applied to list on Nasdaq as "ANT"	Listed on Nasdaq (June 2026); "controlled company"	Listed on Nasdaq (2004)

Note. Facts checked as of 26 October 2026 and subject to rapid change. Sources: OpenAI (n.d., 2025a); Delaware Department of Justice (2025); Anthropic (2023, 2026a); Forbes (2026); Fortune (2026a, 2026b); TechCrunch (2026); Anthropic PBC (2026b); SpaceX (2026a); Musk (2026); Alphabet (2026a, 2026c). Estimates from press reports are marked "est.".

OpenAI was founded in 2015 as a nonprofit research organization. In 2019, to attract capital and talent, it created a capped-profit subsidiary controlled by the nonprofit, in which Microsoft became the largest investor (Aprill et al., 2024; Pedersen & Thomsen, 2026). After a proposal to end nonprofit control was withdrawn following engagement with the California and Delaware Attorneys General, OpenAI completed a recapitalization in October 2025 (Taylor, 2025). The operating business is now OpenAI Group PBC, a Delaware public benefit corporation. The nonprofit, renamed the OpenAI Foundation, held 26% of the equity at closing plus a warrant for additional shares if the share price passes a specified threshold; Microsoft held about 27%; and employees and other investors held the remaining 47% (OpenAI, n.d., 2025a). These stakes have since been diluted by the March 2026 funding round, for which no updated ownership table has been published. The Foundation's stake was valued at about $130 billion at closing, which makes it, on paper, one of the largest charitable endowments in the world.

Anthropic was founded in 2021 by seven former OpenAI employees and incorporated from the outset as a Delaware PBC (Anthropic PBC, 2021; Anthropic, 2023; Pedersen & Thomsen, 2026). Contrary to some reports, its founders do not hold a majority of the equity. According to the registration statement, each co-founder beneficially owns between about 230 and 245 million Class A shares, about 1.7 billion combined, which by our calculation corresponds to roughly 10% of the company's fully diluted shares, or about 1.4% each; earlier press estimates were in the same range (Anthropic PBC, 2026b; Forbes, 2026; TechCrunch, 2026). Amazon and Google are large minority shareholders, with Google's stake reported at about 14%, contractually capped at 15%, and Amazon's at somewhat more (Fortune, 2026a). The registration statement confirms that each will hold more than 5% of the capital stock after conversion of preferred stock and convertible notes, in the form of non-voting Class N shares, and that Amazon's conversions are subject to an ownership cap, but the preliminary filing leaves the exact holdings blank (Anthropic PBC, 2026b). The remaining equity is held by venture and growth investors and employees. In September 2026, ahead of its planned listing, Anthropic was reported to be seeking to give its seven founders a combined 50.1% of voting power, without additional economic rights (Reuters, 2026b; TechCrunch, 2026). The public registration statement filed in October 2026 confirms the core of the proposal: immediately before the offering, a single Class F share without economic rights will be issued to a "Founder LLC" whose initial members are the seven co-founders, carrying 50.1% of the total voting power on matters submitted to stockholders generally, including the election of the directors

not elected by the Trust (Anthropic PBC, 2026b). The share is voted as decided by a majority of the co-founders. A successor can be admitted on a co-founder's death or incapacity until the tenth anniversary of the listing, and once fewer than three co-founders have remained for 90 days, a 450-day transition period begins, after which the Class F share is redeemed for no consideration (Anthropic PBC, 2026b; see also Mari, 2026). The arrangement has been formally filed but takes effect only with the listing, and it illustrates how strongly the founder-control model has spread even to the firm with the most elaborate purpose-protecting structure.

SpaceX is a Texas corporation. It was incorporated in Delaware in 2002 and reincorporated in Texas in February 2024, shortly after the Delaware Court of Chancery invalidated Elon Musk's pay package at Tesla. Since its June 2026 listing, SpaceX has two classes of common stock: Class A shares with one vote and Class B shares with ten votes. Musk holds about 82% of the voting power (SpaceX, 2026a) and about 48% including unvested awards and options (Musk, 2026). Holders of Class B shares elect a majority of the board for as long as any Class B shares are outstanding, and there is no time-based sunset clause. SpaceX is a "controlled company" under Nasdaq rules and uses the associated exemptions from board independence requirements.

Alphabet is a Delaware corporation with three classes of shares: Class A (one vote), Class B (ten votes) and Class C (no votes). The founders Larry Page and Sergey Brin together hold about 53% of the voting power, although their economic stake is much smaller (Alphabet, 2026a). A shareholder proposal to move to one share, one vote received 31% support at the June 2026 annual meeting (Alphabet, 2026b) — most of it, presumably, from outside shareholders.

The legal forms illustrate the regulatory competition among US states. OpenAI and Anthropic chose the Delaware PBC, which signals a commitment to purpose and gives directors legal room to pursue it, but provides limited enforcement. Only shareholders holding at least 2% of the shares (or, in listed companies, the lesser of 2% and shares worth $2 million) can sue directors for failing to balance the public benefit. SpaceX left Delaware for Texas, whose 2025 corporate law reforms allow companies to restrict shareholder litigation, for example by requiring a minimum ownership of up to 3% to bring derivative suits. SpaceX has adopted such a threshold, together with mandatory arbitration of shareholder claims under certain circumstances (SpaceX, 2026a).

8. How owners exert influence

Owners exert influence over the levels below them mainly through the right to appoint and remove directors. The four firms rely on different mechanisms to secure control and, in two cases, to protect purpose. At OpenAI, the OpenAI Foundation controls the operating company through a special class of shares that gives it the exclusive right to appoint and remove all directors of OpenAI Group PBC (OpenAI, n.d.; Pedersen & Thomsen, 2026). The arrangement agreed with the Delaware Attorney General adds several safeguards (Delaware Department of Justice, 2025). The PBC's mission must be identical to that of the Foundation. A Safety and Security Committee remains a committee of the Foundation, not of the PBC, and can require mitigation measures, including halting the release of models. Its chair, initially Zico Kolter, sits on the Foundation board but not on the PBC board, with observer rights at PBC board meetings. Within a year, at least two Foundation directors must not sit on the PBC board, and the Attorney General must receive advance notice of significant governance changes. The California Attorney General concluded a parallel memorandum of understanding (California Department of Justice, 2025). The Foundation also has a contractual right to the personnel, research, intellectual property and support it reasonably needs to pursue its mission (Pedersen & Thomsen, 2026). In short, OpenAI is controlled by a foundation with formal authority over the board and specific authority over safety decisions.

At Anthropic, control is shared between the founders, the investors and the Long-Term Benefit Trust (LTBT). The Trust is a Delaware common-law purpose trust whose financially disinterested trustees hold a special class of stock (Class T) with the right to elect and remove a growing number of directors (Anthropic, 2023; Morley, Berger, & Simmerman, 2023). The Trust's appointment rights were phased in with time and funding milestones, and in April 2026, with the appointment of Vas Narasimhan, Trust-appointed directors came to hold four of the seven board seats (Anthropic, 2026a). In May 2026, however, Jay Kreps, who had been appointed by the Trust in 2024, stepped down from the board (Kreps, 2026). The registration statement confirms that the seat has not been filled (Anthropic PBC, 2026b). Trust appointees therefore hold three of six occupied seats — half the board rather than a majority. The Trust's right to elect a majority of the directors is thus a right that is not always fully exercised. Under the registration statement, the board will have seven authorized seats, four elected by the Trust (Class T directors) and three by the holders of Class A and Class F shares voting together, among whom the Class F share carries a majority of the votes (Anthropic PBC, 2026b; see also Mari, 2026). The Trust's power is qualified in one respect. For one of its four seats (the

"Specified Class T Director"), the Trust can nominate, remove without cause or redesignate the director only with the concurrence of the nominating and governance committee — composed of one Class T director and two other directors — or, failing that, with the approval of holders of a new class of employee shares (Class E), which carry no economic rights and vote only on specified matters (Anthropic PBC, 2026b). Because the Trust can elect a board majority, Anthropic will be a "controlled company" under Nasdaq rules, although it states that it intends to use the associated exemption only for its nominating and corporate governance committee (Anthropic PBC, 2026b). Trustees serve annual terms without term limits, are appointed by the existing trustees in collaboration with the company, and are prohibited from holding Anthropic securities. Each receives annual cash compensation of $400,000 from the company (the chair a further $25,000), which also funds the Trust's operations — about $8.4 million in the first nine months of 2026 (Anthropic PBC, 2026b). Under the arrangements described in the registration statement, the Trust and the rights of the Class T share can be amended with the trustees' support by 80% of the board and 80% of the trustees. Without the trustees' consent, and as long as the Class F share is outstanding, they can be amended by the co-founders acting through the Class F share (unanimously during the first three years after listing, and later by two-thirds of the co-founders) together with a majority of the votes cast by Class A shares and a majority (later 75%) of the votes cast by the employees' Class E shares. Only after the Class F share has lapsed does the earlier stockholder "failsafe", an 85% supermajority of Class A shares, apply (Anthropic PBC, 2026b; Morley, Berger, & Simmerman, 2023). Trustee turnover has been high. Four of the five original trustees had left by early 2026, and one trustee appointed in January 2026 left in August to join the company as an executive. In October 2026, the experimental physicist and former Harvard Dean of Science Christopher Stubbs joined the Trust (Fortune, 2026d). At the time of writing, four of the Trust's five seats are filled, by Neil Buddy Shah (chair), Richard Fontaine, Ben Bernanke and Christopher Stubbs (Anthropic, 2026b; Anthropic PBC, 2026b; Fortune, 2026d). Anthropic is thus controlled by its founders and investors, but subject to a purpose trust that elects the board majority.

At SpaceX and Alphabet, control is secured by dual-class shares. There is no specific mechanism to protect purpose. At SpaceX, Musk's voting control is near-absolute, as only Class B shareholders can remove him as chief executive, chair or director, and his holdings make him effectively irremovable. An award of up to one billion additional super-voting shares, vesting at market capitalization milestones of up to $7.5 trillion and the establishment

of a human settlement on Mars, would further entrench this control (Fortune, 2026b). At Alphabet, the founders' voting majority is held jointly by two people who have withdrawn from day-to-day management, which leaves more room for the board and the professional chief executive.

The contrast between OpenAI and Anthropic deserves emphasis. At OpenAI, the purpose holder — the Foundation — is also an owner with a large economic stake. It thus resembles a classic enterprise foundation, which owns a controlling stake in a business company and uses the dividends for charitable purposes (Hansmann & Thomsen, 2021). At Anthropic, the purpose holder — the Trust — has control rights but no economic stake. It thus resembles the purpose trusts used in steward-ownership arrangements, such as the Patagonia Purpose Trust, which holds the voting shares of Patagonia while a separate nonprofit receives the dividends (Chouinard, 2022; Ho & Harding, 2025). The first model ties the purpose holder's financial resources to the value of the firm while the second keeps the purpose holder financially independent but gives it no resources of its own.

9. Boards

The board is the link between ownership and management in the hierarchy. Table 5 summarizes the boards of directors of the four leading AI firms.

Table 5. Boards of directors (October 2026)

	OpenAI	Anthropic	SpaceX (incl. xAI)	Alphabet (incl. Google DeepMind)
Chair	Bret Taylor (former co-CEO, Salesforce)	Daniela Amodei (co-founder, president)	Elon Musk (also CEO and CTO)	John L. Hennessy (former president, Stanford)
CEO on board	Sam Altman	Dario Amodei (co-founder)	Elon Musk; also Gwynne Shotwell (president, COO)	Sundar Pichai; also founders Larry Page and Sergey Brin
Other directors	Adam D'Angelo (Quora); Sue Desmond-Hellmann (ex-Gates Foundation); Paul M. Nakasone (ret. US Army, ex-NSA); Adebayo Ogunlesi (Global Infrastructure Partners); Nicole Seligman (ex-Sony); David Vélez (Nubank); Robin Vince (BNY)	Yasmin Razavi (Spark Capital); Reed Hastings (co-founder, Netflix); Chris Liddell (ex-CFO, Microsoft); Vas Narasimhan (CEO, Novartis). Jay Kreps (Confluent) stepped down May 2026	Antonio Gracias (Valor Equity); Ira Ehrenpreis (DBL Partners); Randy Glein (DFJ Growth); Donald Harrison (Google); Steve Jurvetson (Future Ventures); Luke Nosek (Gigafund); Roelof Botha (ex-Sequoia)	Frances Arnold (Caltech); Marty Chávez (Sixth Street); John Doerr (Kleiner Perkins); Roger W. Ferguson Jr. (ex-Federal Reserve); K. Ram Shriram (Sherpalo Ventures); Robin L. Washington (Salesforce)
Purpose holder	OpenAI Foundation board: same members plus Zico Kolter (Carnegie Mellon; chair, Safety and Security Committee) and Paul Christiano, both non-voting observers on the PBC board	Long-Term Benefit Trust (five seats, four filled): Neil Buddy Shah (chair), Richard Fontaine, Ben Bernanke, Christopher Stubbs (Oct. 2026); high turnover (four of five original trustees replaced)	(Elon Musk)	(Larry Page and Sergey Brin)
Board–owner relation	Near-complete overlap between Foundation and PBC boards	Trust separate from board; Trust entitled to 4 of 7 seats, but its appointees (Hastings, Liddell, Narasimhan) hold 3 of 6 occupied seats; one Trust seat vacant (S-1)	Insiders and investors with long ties to Musk; Nasdaq controlled-company exemptions	Majority independent directors; founders hold voting majority

Note. Facts checked as of 26 October 2026. Sources: OpenAI (n.d., 2026b); Anthropic (n.d., 2026a, 2026b); Kreps (2026); Anthropic PBC (2026b); Fortune (2026d); Reuters (2026a); SpaceX (2026a); Alphabet (2026a).

Several patterns are worth noting. First, the chief executive sits on the board at all four firms. The chair is independent at OpenAI and Alphabet, but is the chief executive at SpaceX, and at Anthropic the chair is the co-founder and president, Daniela Amodei, sister of the chief executive, Dario Amodei (Anthropic, n.d.). At Anthropic, therefore, both the chair and the chief executive are founders, while the Trust-appointed directors hold half of the six occupied seats.

Second, the boards draw heavily on technology entrepreneurs, venture capitalists and former executives of large companies. At SpaceX, most of the outside directors are venture investors with long-standing ties to Musk, several of whom also sit on Tesla's board or invest in other Musk companies (Reuters, 2026a; SpaceX, 2026a). At Anthropic, the outside directors include an investor representative (Spark Capital), the co-founder of Netflix, a former chief financial

officer of Microsoft and the chief executive of Novartis. OpenAI's board has been rebuilt since 2023 and now includes a former Salesforce co-chief executive as chair, a retired general who led the US National Security Agency, an infrastructure investor, the former chief executive of the Gates Foundation, and the chief executives of Nubank and BNY. Few outside directors at any of the firms have a background in AI safety research; the most prominent exceptions are the two academics on OpenAI's Foundation board, Zico Kolter and Paul Christiano, neither of whom has a vote on the PBC board.

Third, the boards differ in their independence from the controlling owner. OpenAI's Foundation and PBC boards overlap almost completely; the Foundation board has two additional members (Kolter and Christiano). This overlap limits what research on enterprise foundations calls "managerial distance", understood as the separation between the foundation, which acts as owner, and the operating company (Hansmann & Thomsen, 2021; McNulty & Thomsen, 2026). At Anthropic, the Trust is entirely separate from the board and appoints, rather than overlaps with, its members, and with the founders as chief executive and chair and a seat for investors, the board is evenly balanced between Trust appointees and others. At SpaceX, the board is dominated by insiders and directors with close ties to the controlling shareholder (SpaceX, 2026a). At Alphabet, a majority of directors are independent under stock exchange rules, but the founders retain voting control.

10. Benchmarking AI firm governance

Measured against the menu of governance options in Table 3, the four firms differ mainly at the ownership and board levels. Table 6 summarizes their governance in a simple typology, based on two dimensions: who controls the firm, and whether purpose is protected by governance institutions.

Table 6. A typology of AI firm governance

	OpenAI	Anthropic	SpaceX (incl. xAI)	Alphabet (incl. Google DeepMind)
Summary	Nonprofit (foundation) control with purpose protection	Founder leadership with trust-based purpose protection	Largely unchecked founder control	Founder control with checks and balances
Who controls	OpenAI Foundation (board appointment rights); CEO is not a controlling shareholder	Founders and investors, under a purpose trust entitled to elect a board majority (its appointees currently hold 3 of 6 seats)	Elon Musk (c. 82% of votes)	Page and Brin (c. 53% of votes)
Purpose protection	Foundation control; PBC; Safety and Security Committee; AG conditions	PBC; Long-Term Benefit Trust	None (mission not legally binding)	None (mission not legally binding)
Checks on management	Independent chair; formally strong board	Trust-appointed directors (half the board); founder (president) as chair	Few: Texas law, arbitration and 3% threshold limit shareholder suits	Independent chair; professional CEO; independent board majority
Main weakness	Board overlap; resource dependence (2023 crisis); IPO pressure	Trust amendable without trustees' consent by founders, employees and shareholders; founder super-voting proposal; IPO pressure	Entrenchment; conflicts across Musk companies	No AI-specific purpose protection

Note. Authors' classification based on Tables 4 and 5.

OpenAI combines nonprofit control with substantial outside ownership. The chief executive is not a controlling shareholder, the chair is independent, and the formal purpose is protected by foundation control, PBC status, the Safety and Security Committee and conditions enforced by two state Attorneys General. Yet the 2023 crisis showed that formal authority does not guarantee effective control. When the board dismissed the chief executive, the combined response of employees and Microsoft forced it to reverse the decision (CMA, 2025; Harvard Law Review, 2025; Pedersen & Thomsen, 2026).

Anthropic combines founder leadership with a purpose trust that has the right to elect a majority of the board and currently holds half the seats. Its arrangements are the most elaborate attempt to protect purpose in the industry, but they are also the most dependent on the continued goodwill of the founders and other stockholders, who can amend them without the trustees' consent, and the founders' proposed super-voting shares would shift the balance further towards founder control (Anthropic PBC, 2026a, 2026b; Mari, 2026; Morley, Berger, & Simmerman, 2023).

SpaceX represents largely unchecked founder control, protected by dual-class shares and Texas law. Alphabet represents founder control with conventional checks and balances: dual-class

shares, but also an independent chair, a professional chief executive and a board with a majority of independent directors.

At the lower levels of the hierarchy, the evidence points in the same direction. The firm-level safety frameworks discussed in Section 2.5 have been weakened under competitive pressure, which suggests that management-level commitments are only as durable as the commitment of the board and owners above them. At the operations level, employees have proved to be a powerful but unpredictable governance mechanism. At OpenAI in 2023 they helped reverse a decision of the board, and at Anthropic a new class of employee shares (Class E) will vote on changes to the Trust arrangements (Section 8). Employee influence can support purpose, but it can also work against it, for example when employees' interests in stock-based pay diverge from the purpose.

As shown in Section 4.4, the three firms not included in the tables fit the same pattern of founder control. Meta has experimented with stakeholder governance in a different domain. Its Oversight Board, which reviews content moderation decisions, is funded through an irrevocable trust (Oversight Board, 2025a, 2025b). There is, however, no comparable body for AI. Mistral, founded in Paris in 2023, is a privately held French company (SAS) led by its three founders, with venture investors and strategic partners as shareholders; the Dutch semiconductor equipment maker ASML has been its largest shareholder, with a stake of about 11% (ASML, 2025). There is no public evidence that Mistral has adopted the French société à mission form. DeepSeek, based in Hangzhou, was funded entirely by its founder and his quantitative hedge fund High-Flyer until 2026. It is the most closely held of all the firms considered here, although its governance also has to be understood in the context of Chinese state influence.

In short, every leading AI firm is either closely held or founder-controlled through dual-class shares. The variation lies in whether founder control is combined with purpose-protecting devices.

11. Enterprise foundations and AI

The arrangements at OpenAI and Anthropic resemble enterprise foundations: foundations that own and control business companies (Thomsen, 2017; Thomsen & Kavadis, 2022). Enterprise foundations are self-owning entities with no shareholders or members, governed by a charter that typically combines a commercial purpose — the long-term continuation of the company

— with a philanthropic purpose. They are relatively common in Northern Europe, where companies such as Bosch, Carlsberg, IKEA, Maersk and Novo Nordisk are controlled by foundations (Sanders & Thomsen, 2023). Similar arrangements have attracted attention in the United States as a way of anchoring corporate purpose (Eldar & Ørberg, 2026; Hwang & Lund, 2025; Mayer, 2021) and in the broader movement for steward ownership (Manelli et al., 2026).

Research on enterprise foundations provides some grounds for optimism about foundation ownership of AI firms. Despite the absence of owners with a personal profit motive, foundation-owned companies perform at least as well as comparable investor-owned companies in terms of profitability and stock returns (Thomsen & Rose, 2004), survive longer (Thomsen, Poulsen, Børsting, & Kuhn, 2018), pay their employees better and offer more stable employment (Børsting & Thomsen, 2017), and display better social and environmental performance (Schröder & Thomsen, 2025). The proposed explanation is that foundations are long-term owners, less exposed to short-term market pressure and better able to sustain stakeholder trust (Thomsen, 2017; Thomsen, Poulsen, Børsting, & Kuhn, 2018). These are exactly the qualities that proponents of purpose-protecting governance in AI hope for.

However, the research also suggests three reasons for caution. The first concerns the life cycle. Enterprise foundations have typically been established in mature companies that have reached a certain size and prosperity (Thomsen, 2017). Foundation ownership is a poor fit for young, fast-growing firms because it limits access to external equity and strategic flexibility. It fits best in the middle of the life cycle, when the company is established and the main governance challenge is long-term commitment. It may fit poorly again if a company needs to be sold or merged. Frontier AI firms are at the early, capital-hungry growth stage, which helps explain why OpenAI moved from a pure nonprofit to a capped-profit model and then to conventional equity, and why its stated rationale was that investors "need conventional equity and less structural bespokeness" (OpenAI, 2024; Pedersen & Thomsen, 2026). A related prediction is that the planned stock market listings will put the foundation and trust arrangements under further pressure, as public investors demand clarity about who controls the company.

The second concerns governance competence. Research on foundation governance emphasizes that formal control rights are not enough; effective foundation ownership requires managerial distance between the foundation and the company, credible ownership commitments to the purpose, and a competent foundation board (McNulty & Thomsen, 2026). Good practice in Danish industrial foundations is to separate the foundation and company boards with limited

overlap (Hansmann & Thomsen, 2021). By this standard, OpenAI's near-complete board overlap is a weakness, while Anthropic's separation between Trust and board is closer to best practice, although the Trust's small size, high turnover and dependence on the company for its funding raise questions about its capacity.

The third concerns resource dependence. At both OpenAI and Anthropic, the effectiveness of purpose control has been dependent on critical resources — capital, computing infrastructure, specialized employees and key customers — controlled by actors whose interests may diverge from the purpose. OpenAI's board had the formal authority to dismiss its chief executive but could not sustain the decision when employees threatened to leave and Microsoft offered them an alternative (CMA, 2025). Anthropic maintained contractual restrictions on military uses of its models against pressure from the US government, but this was led by management (Amodei, 2026), and the Trust's contribution is unclear. Foundation-like control therefore establishes a formal basis for purpose-oriented intervention, but its practical effect depends on the governance competence of the purpose holder and on the firm's resource position.

Compared with classic enterprise foundations, the AI arrangements are also less complete. In a Danish enterprise foundation, the foundation typically holds a controlling majority of votes and the arrangement is protected by foundation law and supervision by a public foundation authority (Sanders & Thomsen, 2023). The OpenAI Foundation holds a minority of the equity, but all the board appointment rights, with state Attorneys General acting as a partial substitute for a foundation authority. Anthropic's Trust holds no economic stake, and its rights can be amended without its consent. Both arrangements are therefore better described as "foundation-like" than as enterprise foundations proper. In October 2026 Anthropic announced that 5% of its shares would be committed to a wholly owned Anthropic Public Benefit Fund, which may donate to nonprofits or make direct distributions, for example to people affected by AI's effects on the labour market (Anthropic, 2026d). According to the registration statement, the company has contributed non-voting shares representing 5% of its fully diluted shares as of 30 June 2026 to the fund, and the board may contribute up to a further 2% of outstanding shares each time the share price rises 50% above the previous milestone; until the board approves a governance structure for the fund, the board decides its allocations, and the designated initial managers are Dario Amodei, Daniela Amodei, Jack Clark and the Trust chair, Neil Buddy Shah (Anthropic PBC, 2026b). Separately, each co-founder has pledged to dedicate 80% of his or her personal equity to philanthropic causes (Anthropic PBC, 2026b). If realized, this would move Anthropic a small step closer to the enterprise foundation model by giving a purpose-bound entity an

economic stake, although an entity owned by the company itself is not an independent owner in the way an enterprise foundation is.

12. Conclusion

This paper has analysed the governance of the leading AI firms using a framework that views governance as a hierarchy of connected mechanisms, from regulation through ownership and boards to management and operations. All the leading firms are closely held or founder-controlled through dual-class shares. In this respect, the AI industry reflects a broader revival of founder control in corporate governance, which is what one would expect in knowledge-intensive, high-technology firms. What distinguishes the firms is whether founder control is combined with institutional protection of purpose. OpenAI is controlled by a foundation, Anthropic by founders under a purpose trust, Alphabet by founders with conventional checks and balances, and SpaceX by a single founder with few checks.

The comparison has implications for how AI firms should be governed. Frontier AI may create significant externalities, including risks to safety, security, employment and democratic processes that are not fully borne by the firms' owners. Public regulation is the first-best response, but for the reasons discussed in Section 3 it is likely to remain weak and unstable. Private-law arrangements at the ownership and board levels are therefore second-best options that will realistically not solve the problem, but may contribute to a solution. They are typically opt-in, as firms choose them voluntarily, perhaps encouraged by investors, employees or regulators. They include corporate purpose clauses, possibly supported by comply-or-explain recommendations; director or corporate liability for breach of purpose; purposeful ownership through enterprise foundations, perpetual purpose trusts or steward ownership; stakeholder or expert boards of the kind adopted by Anthropic and, in a different domain, Meta; legal forms such as the French société à mission and the German proposals for companies in steward ownership; and certification, of which the PBC form is a weak version.

The experience of the leading AI firms suggests that the value of these devices depends on their credibility. A purpose clause that can be amended by the controlling owner, or a trust that can be overridden by investors, offers limited protection when the stakes are highest. Conversely, purpose guardians who are protected from owners face the familiar problem that leaders with broad stakeholder mandates may lack the incentives to deliver on them (Fried & Reiter, 2026). Foundation ownership is the strongest commitment device, because it removes the possibility of private appropriation and is supervised by public authorities, but it fits less well with the

agility and capital needs of startup firms. Imperfect as they are, these devices are superior, from a governance perspective, to unchecked founder control (autarchy or non-governance), because they delegate some power to purpose guardians and subject founders to checks through disclosure and commitment. The most realistic role for private law is therefore to complement, not replace, public oversight.

Several questions remain open and will be answered by events in the coming months. The planned listing of OpenAI and the imminent listing of Anthropic, which publicly filed its registration statement in October 2026, will test whether foundation-like control and purpose trusts can survive the transition to public markets. Anthropic's plan, set out in its registration statement, to give its founders majority voting power will show how far purpose-protecting governance can coexist with founder control. And the possible merger of SpaceX and Tesla will reveal how far a single controlling shareholder can reshape an AI firm's governance without effective checks. The facts reported here will change quickly, but the underlying governance questions — who controls the firms that build AI, what they are obliged to pursue, and who can change these arrangements — will not. Research on corporate governance and enterprise foundations has much to contribute to answering them.

References

Acemoglu, D. 2021. Harms of AI. *NBER Working Paper* No. 29247. Cambridge, MA: National Bureau of Economic Research.

Adams, R., Almeida, H., & Ferreira, D. 2009. Understanding the relationship between founder–CEOs and firm performance. *Journal of Empirical Finance*, 16(1): 136–150.

Adams, R. B., & Ferreira, D. 2026. The governance of dual-class firms. ECGI Finance Working Paper No. 1152/2026, European Corporate Governance Institute. http://dx.doi.org/10.2139/ssrn.7360478

Adams, R. B., Hermalin, B. E., & Weisbach, M. S. 2010. The role of boards of directors in corporate governance: A conceptual framework and survey. *Journal of Economic Literature*, 48(1): 58–107.

Aggarwal, D., Eldar, O., Hochberg, Y. V., & Litov, L. P. 2022. The rise of dual-class stock IPOs. *Journal of Financial Economics*, 144(1): 122–153.

Aghion, P., & Tirole, J. 1997. Formal and real authority in organizations. *Journal of Political Economy*, 105(1): 1–29.

Agrawal, A., & Knoeber, C. R. 1996. Firm performance and mechanisms to control agency problems between managers and shareholders. *Journal of Financial and Quantitative Analysis*, 31(3): 377–397.

Aguilera, R. V., Desender, K., Bednar, M. K., & Lee, J. H. 2015. Connecting the dots: Bringing external corporate governance into the corporate governance puzzle. *Academy of Management Annals*, 9(1): 483–573.

Aguilera, R. V., Filatotchev, I., Gospel, H., & Jackson, G. 2008. An organizational approach to comparative corporate governance: Costs, contingencies, and complementarities. *Organization Science*, 19(3): 475–492.

Ahmed, N., Wahed, M., & Thompson, N. C. 2023. The growing influence of industry in AI research. *Science*, 379(6635): 884–886.

Ali, S. J., Christin, A., Smart, A., & Katila, R. 2023. Walking the walk of AI ethics: Organizational challenges and the individualization of risk among ethics entrepreneurs. *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*: 217–226. https://doi.org/10.1145/3593013.3593990

Alphabet. 2026a. *Annual report on Form 10-K for fiscal year 2025* and *Definitive proxy statement (DEF 14A), April 24, 2026*. Mountain View, CA: Alphabet Inc.

Alphabet. 2026b. Current report (Form 8-K), results of the June 5 annual meeting. June 11. https://www.sec.gov/Archives/edgar/data/1652044/000119312526267578/d57679d8k.htm

Alphabet. 2026c. Schedule 13G, Space Exploration Technologies Corp., Class A common stock. Filed with the US Securities and Exchange Commission. August 14. https://www.sec.gov/Archives/edgar/data/0001181412/000119312526352289/xslSCHEDULE_13G_X02/primary_doc.xml

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

Anderljung, M., Barnhart, J., Korinek, A., Leung, J., O'Keefe, C., Whittlestone, J., et al. 2023. Frontier AI regulation: Managing emerging risks to public safety. *arXiv*: 2307.03718.

Anderson, R., Ottolenghi, E., Reeb, D., & Savor, P. 2025. Control risk premium: Dual-class shares, family ownership, and minority investor returns. *Financial Management*, 54(2): 199–236.

Andhov, A., & Murray, I. 2026. OpenAI: Governance for public good or private gain? SSRN Working Paper No. 6673598.

Anthropic. n.d. Leadership at Anthropic. https://www.anthropic.com/company/leadership (accessed October 21, 2026).

Anthropic. 2023. The Long-Term Benefit Trust. September 19. https://www.anthropic.com/news/the-long-term-benefit-trust

Anthropic. 2026a. Anthropic's Long-Term Benefit Trust appoints Vas Narasimhan to Board of Directors. April 14. https://www.anthropic.com/news/narasimhan-board

Anthropic. 2026b. Ben Bernanke appointed to Anthropic's Long-Term Benefit Trust. July 9. https://www.anthropic.com/news/ben-bernanke

Anthropic. 2026d. Investing in the public benefit. October 18. https://www.anthropic.com/news/public-benefit-fund

Anthropic. 2026e. Anthropic raises $65B in Series H funding at $965B post-money valuation. May 28. https://www.anthropic.com/news/series-h

Anthropic. 2026f. Anthropic's Responsible Scaling Policy version 3.0. February 24. https://anthropic.com/responsible-scaling-policy/rsp-v3-0

Anthropic. 2026g. Anthropic confidentially submits draft S-1 to the SEC. June 1. https://www.anthropic.com/news/confidential-draft-s1-sec

Anthropic PBC. 2021. Certificate of incorporation of Anthropic, PBC. Filed January 26 with the Delaware Secretary of State, Division of Corporations, entity file no. 4860621.

Anthropic PBC. 2026a. Amended and restated certificate of incorporation of Anthropic, PBC. Filed May 27 with the Delaware Secretary of State, Division of Corporations, entity file no. 4860621; amended September 17.

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

Aprill, E. P., Chan Loui, R., & Horwitz, J. R. 2024. Board control of a charity's subsidiaries: The saga of OpenAI. *Tax Notes Federal*, 182: 289–296.

Aran, Y., Broughman, B. J., & Pollman, E. 2024. CEO turnover at dual-class firms. ECGI Law Working Paper No. 839/2025. https://ssrn.com/abstract=5063277

Armstrong, S., Bostrom, N., & Shulman, C. 2016. Racing to the precipice: A model of artificial intelligence development. *AI & Society*, 31(2): 201–206.

Arrow, K. J. 1963. Uncertainty and the welfare economics of medical care. *American Economic Review*, 53(5): 941–973.

ASML. 2025. ASML and Mistral AI enter into strategic partnership. September 9. https://www.asml.com/en/news/press-releases/2025/asml-mistral-ai-enter-strategic-partnership

Bartz v. Anthropic PBC. 2025. Order on fair use. No. 3:24-cv-05417-WHA, Dkt. 231, June 23. United States District Court for the Northern District of California.

Bartz v. Anthropic PBC. 2026. Order granting final approval of class action settlement. No. 3:24-cv-05417-AMO, Dkt. 680, July 20. United States District Court for the Northern District of California.

Bebchuk, L. A., & Kastiel, K. 2017. The untenable case for perpetual dual-class stock. *Virginia Law Review*, 103(4): 585–631.

Bebchuk, L. A., & Tallarita, R. 2020. The illusory promise of stakeholder governance. *Cornell Law Review*, 106: 91–178.

Bengio, Y., et al. 2026. *International AI safety report 2026*. London: UK Department for Science, Innovation and Technology / AI Security Institute. https://internationalaisafetyreport.org/publication/international-ai-safety-report-2026

Bengio, Y., Hinton, G., Yao, A., Song, D., Abbeel, P., et al. 2024. Managing extreme AI risks amid rapid progress. *Science*, 384(6698): 842–845.

Berle, A. A., & Means, G. C. 1932. *The modern corporation and private property*. New York: Macmillan.

Birhane, A., Steed, R., Ojewale, V., Vecchione, B., & Raji, I. D. 2024. AI auditing: The broken bus on the road to AI accountability. *2024 IEEE Conference on Secure and Trustworthy Machine Learning*: 612–643.

Birkstedt, T., Minkkinen, M., Tandon, A., & Mäntymäki, M. 2023. AI governance: Themes, knowledge gaps and future agendas. *Internet Research*, 33(7): 133–167.

Bommasani, R., Klyman, K., Longpre, S., Kapoor, S., Maslej, N., Xiong, B., Zhang, D., & Liang, P. 2023. The foundation model transparency index. *arXiv*: 2310.12941.

Bresnahan, T. F., & Trajtenberg, M. 1995. General purpose technologies 'Engines of growth'? *Journal of Econometrics*, 65(1): 83–108.

Burkart, M., Gromb, D., & Panunzi, F. 1997. Large shareholders, monitoring, and the value of the firm. *Quarterly Journal of Economics*, 112(3): 693–728.

Børsting, C., & Thomsen, S. 2017. Foundation ownership, reputation, and labour. *Oxford Review of Economic Policy*, 33(2): 317–338.

Caixin. 2026. DeepSeek reaches $52 billion valuation in round backed by Tencent, CATL. July 17. https://www.caixinglobal.com/2026-07-17/deepseek-reaches-52-billion-valuation-in-round-backed-by-tencent-catl-102465358.html

California Department of Justice. 2025. Attorney General Bonta issues statement on OpenAI's recapitalization plan. October 28. https://oag.ca.gov/news/press-releases/attorney-general-bonta-issues-statement-openai%E2%80%99s-recapitalization-plan

Chouinard, Y. 2022. Earth is now our only shareholder. Patagonia, September 14. https://www.patagonia.com/ownership/

Cihon, P., Schuett, J., & Baum, S. D. 2021. Corporate governance of artificial intelligence in the public interest. *Information*, 12(7): 275.

CMA (Competition and Markets Authority). 2025. Microsoft Corporation's partnership with OpenAI, Inc. Decision on relevant merger situation. March 5. https://assets.publishing.service.gov.uk/media/67fe26ef712bf73dea135449/____Full_text_decision__.pdf

Coles, J. L., Daniel, N. D., & Naveen, L. 2008. Boards: Does one size fit all? *Journal of Financial Economics*, 87(2): 329–356.

Collingridge, D. 1980. *The social control of technology*. London: Frances Pinter.

Conference Board. 2026. *Report: In the S&P 500, disclosure of AI risks surges from 12% to 83%*. New York: The Conference Board, April 22. https://www.conference-board.org/press/governing-AI-2026

Costanza-Chock, S., Raji, I. D., & Buolamwini, J. 2022. Who audits the auditors? Recommendations from a field scan of the algorithmic auditing ecosystem. *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*: 1571–1583.

CII (Council of Institutional Investors). n.d. Dual-class stock. https://www.cii.org/dualclass_stock (accessed October 22, 2026).

Cottier, B., Rahman, R., Fattorini, L., Maslej, N., Besiroglu, T., & Owen, D. 2024. The rising costs of training frontier AI models. arXiv: 2405.21015.

Cremers, M., Lauterbach, B., & Pajuste, A. 2024. The life cycle of dual-class firm valuation. *Review of Corporate Finance Studies*, 13(2): 459–493.

Dafoe, A. 2018. *AI governance: A research agenda*. Oxford: Centre for the Governance of AI, Future of Humanity Institute.

Demsetz, H., & Lehn, K. 1985. The structure of corporate ownership: Causes and consequences. *Journal of Political Economy*, 93(6): 1155–1177.

DSIT (Department for Science, Innovation and Technology). 2025. Tackling AI security risks to unleash growth and deliver Plan for Change. February 14. https://www.gov.uk/government/news/tackling-ai-security-risks-to-unleash-growth-and-deliver-plan-for-change

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

Doidge, C., Karolyi, G. A., & Stulz, R. M. 2017. The U.S. listing gap. *Journal of Financial Economics*, 123(3): 464–487.

EDPB (European Data Protection Board). 2024. *Opinion 28/2024 on certain data protection aspects related to the processing of personal data in the context of AI models*. Adopted December 17, published December 18.

Eldar, O., & Ørberg, M. 2026. The anatomy of nonprofit control of business enterprise. *Yale Journal on Regulation*, 43(2): 335–414.

European Commission. n.d. The General-Purpose AI Code of Practice. https://digital-strategy.ec.europa.eu/en/policies/contents-code-gpai (accessed October 2026).

European Union. 2024a. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). *Official Journal of the European Union*, L series. https://eur-lex.europa.eu/eli/reg/2024/1689/oj

European Union. 2024b. Directive (EU) 2024/2810 of 23 October 2024 on multiple-vote share structures in companies that seek admission to trading of their shares on a multilateral trading facility. https://eur-lex.europa.eu/eli/dir/2024/2810/oj/eng

European Union. 2026. Regulation (EU) 2026/1744 of 8 July 2026 amending Regulations (EU) 2024/1689, (EU) 2018/1139 and (EU) 2023/1230 (Digital omnibus on AI). *Official Journal of the European Union*, July 24. https://eur-lex.europa.eu/eli/reg/2026/1744/oj

Ewens, M., & Farre-Mensa, J. 2020. The deregulation of the private equity markets and the decline in IPOs. *Review of Financial Studies*, 33(12): 5463–5509.

Executive Order 14365. 2025. Ensuring a national policy framework for artificial intelligence. December 11. Washington, DC: The White House.

Faccio, M., & Lang, L. H. P. 2002. The ultimate ownership of Western European corporations. *Journal of Financial Economics*, 65(3): 365–395.

Fahlenbrach, R. 2009. Founder-CEOs, investment decisions, and stock market performance. *Journal of Financial and Quantitative Analysis*, 44(2): 439–466.

Fama, E. F., & Jensen, M. C. 1983. Separation of ownership and control. *Journal of Law and Economics*, 26(2): 301–325.

Fan, J. S., & Nguyen, X.-T. 2025. Novel corporate governance structures. *Harvard Journal of Law & Technology*, 38(4): 1027–1102.

Feffer, M., Sinha, A., Deng, W. H., Lipton, Z. C., & Heidari, H. 2024. Red-teaming for generative AI: Silver bullet or security theater? *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 7(1): 421–437.

FCA (Financial Conduct Authority). 2024. FCA overhauls listing rules to boost growth and innovation on UK stock markets. July 11. https://www.fca.org.uk/news/press-releases/fca-overhauls-listing-rules-boost-growth-and-innovation-uk-stock-markets

Floridi, L. 2019. Translating principles into practices of digital ethics: Five risks of being unethical. *Philosophy & Technology*, 32(2): 185–193.

Forbes. 2026. Fortunes of Anthropic's seven cofounders more than double to $16.6 billion each. June 1. https://www.forbes.com.au/news/billionaires/fortunes-of-anthropics-seven-cofounders-more-than-double-to-16-6-billion-each/

Fortune. 2025. The new extreme in founder-friendly deals: Complete board control. May 2. https://fortune.com/2025/05/02/the-new-extreme-in-founder-friendly-deals-complete-board-control/

Fortune. 2026a. Amazon and Google have billions riding on Anthropic. The IPO will finally reveal how much. June 4. https://fortune.com/2026/06/04/amazon-google-billions-anthropic-ipo/

Fortune. 2026b. 'Tap-dancing through a minefield': Elon Musk's $3.5 trillion Tesla-SpaceX merger talks expose his power—and scant accountability. October 21. https://fortune.com/2026/10/21/elon-musk-trillionaire-spacex-tesla-merger-corporate-accountability/

Fortune. 2026c. Google Gemini boss describes working with founders Larry Page and Sergey Brin to win the AI future. Fortune, February 12. https://fortune.com/2026/02/12/google-gemini-boss-founders-larry-page-sergey-brin-to-win-ai-future/

Fortune. 2026d. Anthropic adds Harvard astrophysicist to the trust that controls its board. October 21. https://fortune.com/2026/10/21/anthropic-adds-harvard-astrophysicist-to-the-trust-that-controls-its-board/

Fried, J. M., & Reiter, I. 2026. AI corporate governance and Ben & Jerry's risk. ECGI Law Working Paper No. 927/2026.

FTSE Russell. 2017. FTSE Russell voting rights consultation: Next steps. https://www.lseg.com/content/dam/ftse-russell/en_us/documents/other/ftse-russell-voting-rights-consultation-next-steps.pdf

Future of Life Institute. 2026. *AI safety index: Summer 2026*. July 14. https://futureoflife.org/ai-safety-index-summer-2026/

Giroud, X., & Mueller, H. M. 2010. Does corporate governance matter in competitive industries? *Journal of Financial Economics*, 95(3): 312–331.

Glaeser, E. L., & Shleifer, A. 2001. Not-for-profit entrepreneurs. *Journal of Public Economics*, 81(1): 99–115.

Gofman, M., & Jin, Z. 2024. Artificial intelligence, education, and entrepreneurship. *Journal of Finance*, 79(1): 631–667.

GPDP (Garante per la protezione dei dati personali). 2023. Intelligenza artificiale: il Garante blocca ChatGPT. Raccolta illecita di dati personali. Assenza di sistemi per la verifica dell'età dei minori. https://www.garanteprivacy.it/web/guest/home/docweb/-/docweb-display/docweb/9870847

GPDP (Garante per la protezione dei dati personali). 2024. ChatGPT, il Garante privacy chiude l'istruttoria. OpenAI dovrà realizzare una campagna informativa di sei mesi e pagare una sanzione di 15 milioni di euro. https://www.garanteprivacy.it/home/docweb/-/docweb-display/docweb/10085432

Gebru, T., Morgenstern, J., Vecchione, B., Vaughan, J. W., Wallach, H., Daumé III, H., & Crawford, K. 2021. Datasheets for datasets. *Communications of the ACM*, 64(12): 86–92.

Goldblatt, A. M. 2026. The inverted benefit corporation: OpenAI and the limits of mission governance. SSRN Working Paper No. 6613658.

Gompers, P. A., Ishii, J., & Metrick, A. 2010. Extreme governance: An analysis of dual-class firms in the United States. *Review of Financial Studies*, 23(3): 1051–1088.

Google DeepMind. n.d. About Google DeepMind. https://deepmind.google/about/ (accessed October 22, 2026).

Goshen, Z., & Hamdani, A. 2016. Corporate control and idiosyncratic vision. *Yale Law Journal*, 125: 560–617.

Groves, L., Metcalf, J., Kennedy, A., Vecchione, B., & Strait, A. 2024. Auditing work: Exploring the New York City algorithmic bias audit regime. *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*: 1107–1120.

Guiso, L., Sapienza, P., & Zingales, L. 2015. The value of corporate culture. *Journal of Financial Economics*, 117(1): 60–76.

Hagendorff, T. 2020. The ethics of AI ethics: An evaluation of guidelines. *Minds and Machines*, 30(1): 99–120.

Hansmann, H. 1980. The role of nonprofit enterprise. *Yale Law Journal*, 89(5): 835–901.

Hansmann, H. 1996. *The ownership of enterprise*. Cambridge, MA: Harvard University Press.

Hansmann, H., & Thomsen, S. 2021. The governance of foundation-owned firms. *Journal of Legal Analysis*, 13(1): 172–230.

Hart, O., & Zingales, L. 2017. Companies should maximize shareholder welfare not market value. *Journal of Law, Finance, and Accounting*, 2(2): 247–275.

Harvard Law Review. 2025. Developments in the law—Artificial intelligence, chapter four: Amoral drift in AI corporate governance. *Harvard Law Review*, 138(6): 1633–1656.

Hendrycks, D., Mazeika, M., & Woodside, T. 2023. An overview of catastrophic AI risks. arXiv: 2306.12001.

Hettler, B., Forst, A., Cordeiro, J., & Chavez, S. 2021. Excess insider control and corporate social responsibility: Evidence from dual-class firms. *Journal of Accounting and Public Policy*, 40(6): 106877.

Ho, L., & Harding, M. 2025. A commitment theory of express trusts. *Law Quarterly Review*, 141: 247–271.

Hopt, K. J., & Kalss, S. 2024. Multiple-voting shares in Europe: A comparative law and economic analysis. ECGI Law Working Paper No. 786/2024.

HKEX (Hong Kong Exchanges and Clearing). 2018. Hong Kong's listing regime enters new era, featuring emerging and innovative firms. April 24. https://www.hkex.com.hk/News/Regulatory-Announcements/2018/180424news?sc_lang=en

Hwang, C., & Lund, D. S. 2025. Purpose and nonprofit enterprise. *Columbia Law Review*, 125: 2267–2330.

In re OpenAI, Inc., Copyright Infringement Litigation, No. 1:25-md-03143-SHS-OTW (S.D.N.Y., 2026) (stipulated omnibus sealing order for summary judgment and Daubert briefing). September 3. https://dockets.justia.com/docket/new-york/nysdce/1:2026cv06966/670598

ISS (Institutional Shareholder Services). 2025. United States proxy voting guidelines: Benchmark policy recommendations. Published December 9, effective for meetings on or after February 1, 2026. https://www.issgovernance.com/file/policy/active/americas/US-Voting-Guidelines.pdf

ISO/IEC. 2023. *ISO/IEC 42001:2023 Information technology — Artificial intelligence — Management system*. Geneva: International Organization for Standardization.

ISS-Corporate. 2025. Roughly one-third of large U.S. companies now disclose board oversight of AI, ISS-Corporate finds. March 19. https://www.iss-corporate.com/press/roughly-one-third-of-large-u-s-companies-now-disclose-board-oversight-of-ai-iss-corporate-finds/

Jensen, M. C. 1993. The modern industrial revolution, exit, and the failure of internal control systems. *Journal of Finance*, 48(3): 831–880.

Jensen, M. C., & Meckling, W. H. 1976. Theory of the firm: Managerial behavior, agency costs and ownership structure. *Journal of Financial Economics*, 3(4): 305–360.

Jensen, M. C., & Meckling, W. H. 1992. Specific and general knowledge, and organizational structure. In L. Werin & H. Wijkander (Eds.), *Contract economics*: 251–274. Oxford: Blackwell.

Jobin, A., Ienca, M., & Vayena, E. 2019. The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1(9): 389–399.

Kahle, K. M., & Stulz, R. M. 2017. Is the US public corporation in trouble? *Journal of Economic Perspectives*, 31(3): 67–88.

Kim, S. 2023. Dual-class share structure and firm risks. *Pacific-Basin Finance Journal*, 80:102084.

Kim, H., & Michaely, R. 2019. Sticking around too long? Dynamics of the benefits of dual-class voting. ECGI Finance Working Paper No. 590/2019.

Korinek, A., & Stiglitz, J. E. 2019. Artificial intelligence and its implications for income distribution and unemployment. In A. Agrawal, J. Gans, & A. Goldfarb (Eds.), *The economics of artificial intelligence: An agenda*: 349–390. Chicago: University of Chicago Press.

Korinek, A., & Vipra, J. 2025. Concentrating intelligence: Scaling and market structure in artificial intelligence. *Economic Policy*, 40(121): 225–256.

Kreps, J. 2026. Post announcing departure from the Anthropic board of directors. LinkedIn, May 22. https://www.linkedin.com/in/jaykreps/

La Porta, R., Lopez-de-Silanes, F., & Shleifer, A. 1999. Corporate ownership around the world. *Journal of Finance*, 54(2): 471–517.

La Porta, R., Lopez-de-Silanes, F., Shleifer, A., & Vishny, R. W. 1998. Law and finance. *Journal of Political Economy*, 106(6): 1113–1155.

Lel, U., Netter, J., Poulsen, A., & Qin, Z. 2025. Dual-class shares and firm valuation: Market-wide evidence from regulatory events. *Financial Review*, 60(4): 1251–1275.

Levine-Schnur, R., & Ofir, M. 2025. GenAI models and the hybrid governance trap. *Cornell Journal of Law and Public Policy*, 34(3): 339–363.

Liang, W., Rajani, N., Yang, X., Ozoani, E., Wu, E., Chen, Y., Smith, D. S., & Zou, J. 2024. Systematic analysis of 32,111 AI model cards characterizes documentation practice in AI. *Nature Machine Intelligence*, 6(7): 744–753.

Manelli, L., Pek, S., Waldkirch, M., Hachigian, H., Jamal, A., Thomsen, S., Besharov, M., Hand, M., Segrestin, B., Levillain, K., & Hatchuel, A. 2026. Beyond ownership as usual: The implications of steward-ownership for management research. *Journal of Management Inquiry*, 35(3): 227–252.

Marchant, G. E., Allenby, B. R., & Herkert, J. R. (Eds.). 2011. *The growing gap between emerging technologies and legal-ethical oversight: The pacing problem*. Dordrecht: Springer.

Milgrom, P., & Roberts, J. 1995. Complementarities and fit: Strategy, structure, and organizational change in manufacturing. *Journal of Accounting and Economics*, 19(2–3): 179–208.

Mäntymäki, M., Minkkinen, M., Birkstedt, T., & Viljanen, M. 2022. Defining organizational AI governance. *AI and Ethics*, 2(4): 603–609.

Mari, M. 2026. Public purpose and founder control at Anthropic. *CLS Blue Sky Blog*, October 2. https://clsbluesky.law.columbia.edu/2026/10/02/public-purpose-and-founder-control-at-anthropic/

Masulis, R. W., Wang, C., & Xie, F. 2009. Agency problems at dual-class companies. *Journal of Finance*, 64(4): 1697–1727.

Mayer, C. 2021. The future of the corporation and the economics of purpose. *Journal of Management Studies*, 58(3): 887–901.

McNulty, T., & Thomsen, S. 2026. Foundation governance for the purposeful ownership of enterprise. *European Management Review*. https://doi.org/10.1111/emre.70044

Meta Platforms. 2026. Preliminary information statement (PRE 14C). October 21. https://www.sec.gov/Archives/edgar/data/1326801/000162828026066710/meta-20261021.htm

Meyer, J. W., & Rowan, B. 1977. Institutionalized organizations: Formal structure as myth and ceremony. *American Journal of Sociology, 83*(2): 340–363.

Mitchell, M., Wu, S., Zaldivar, A., Barnes, P., Vasserman, L., Hutchinson, B., Spitzer, E., Raji, I. D., & Gebru, T. 2019. Model cards for model reporting. *Proceedings of the Conference on Fairness, Accountability, and Transparency*: 220–229.

Mistral AI. 2025. Statuts mis à jour, 6 August 2025, filed copy made available 12 October 2026, RCS Paris no. 952 418 325. https://www.pappers.fr/entreprise/mistral-ai-952418325/documents/Mistral%20AI%20-%20Copie%20des%20statuts%20mis%20%C3%A0%20jour%2012-10-2026.pdf

Mittelstadt, B. 2019. Principles alone cannot guarantee ethical AI. *Nature Machine Intelligence*, 1(11): 501–507.

Mökander, J., Schuett, J., Kirk, H. R., & Floridi, L. 2024. Auditing large language models: A three-layered approach. *AI and Ethics*, 4(4): 1085–1115.

Morley, J., Berger, D. J., & Simmerman, A. L. 2023. Anthropic Long-Term Benefit Trust. *Harvard Law School Forum on Corporate Governance*, October 28. https://corpgov.law.harvard.edu/2023/10/28/anthropic-long-term-benefit-trust/

Munn, L. 2023. The uselessness of AI ethics. *AI and Ethics*, 3(3): 869–877.

Musk, E. R. 2026. Schedule 13G, Space Exploration Technologies Corp., Class A common stock. Filed with the US Securities and Exchange Commission. August 13. https://www.sec.gov/Archives/edgar/data/1181412/000110465926095936/xslSCHEDULE_13G_X02/primary_doc.xml

New York State. 2026. Responsible AI Safety and Education Act. General Business Law, Article 44-B, as amended. https://www.nysenate.gov/legislation/laws/GBS/A44-B

New York Times. 2026. How SpaceX Is Structured to Favor Elon Musk. May 26. https://www.nytimes.com/2026/05/26/technology/spacex-elon-musk-pay-board-governance.html

NIST (National Institute of Standards and Technology). 2023. *Artificial intelligence risk management framework (AI RMF 1.0)*. NIST AI 100-1. Gaithersburg, MD: NIST.

Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. 2019. Dissecting racial bias in an algorithm used to manage the health of populations. *Science*, 366(6464): 447–453.

OECD. 2024. *Recommendation of the Council on Artificial Intelligence* (OECD/LEGAL/0449), adopted 2019, amended 2023 and 2024. Paris: OECD.

Office of the New York State Comptroller. 2025. *Enforcement of Local Law 144 – Automated employment decision tools* (Report 2024-N-6). December 2.

OpenAI. n.d. Our structure. https://openai.com/our-structure/ (accessed October 2026).

OpenAI. 2024. Why OpenAI's structure must evolve to advance our mission. December 27. https://openai.com/index/why-our-structure-must-evolve-to-advance-our-mission/

OpenAI. 2025a. The next chapter of the Microsoft–OpenAI partnership. October 28. https://openai.com/index/next-chapter-of-microsoft-openai-partnership/

OpenAI. 2025b. Preparedness framework, version 2. April 15. https://cdn.openai.com/pdf/18a02b5d-6b67-4cec-ab64-68cdfbddebcd/preparedness-framework-v2.pdf

OpenAI. 2026b. David Vélez and Robin Vince join the boards of the OpenAI Foundation and OpenAI Group PBC. July 21. https://openai.com/index/david-velez-robin-vince-join-openai-boards/

OpenAI. 2026c. OpenAI raises $122 billion to accelerate the next phase of AI. March 31. https://openai.com/index/accelerating-the-next-phase-ai/

OpenAI. 2026d. Confidential submission of draft S-1 to the SEC. June 8. https://openai.com/index/openai-submits-confidential-s-1/

Oudin, P., & Groza, T. 2024. The governance of AI companies: Reconciling purpose with profits. SSRN Working Paper No. 4972751.

Oversight Board. 2025a. Oversight Board charter. June. https://www.oversightboard.com/wp-content/uploads/2025/07/Oversight-Board-Charter-June-2025.pdf

Oversight Board. 2025b. Oversight Board Trust agreement. October 16, 2019, amended and restated June 18, 2025, as amended December 22, 2025. https://www.oversightboard.com/wp-content/uploads/2026/04/1.-Oversight-Board-Trust-Agreement-December-2025.pdf

Papagiannidis, E., Mikalef, P., & Conboy, K. 2025. Responsible artificial intelligence governance: A review and research framework. *Journal of Strategic Information Systems*, 34(2): 101885.

Pedersen, F. H., & Thomsen, S. 2026. Enterprise foundations in AI. Working paper, Copenhagen Business School.

Pfeffer, J., & Salancik, G. R. 1978. *The external control of organizations: A resource dependence perspective*. New York: Harper & Row.

Rajan, R. G., & Zingales, L. 1998. Power in a theory of the firm. *Quarterly Journal of Economics*, 113(2): 387–432.

Raji, I. D., & Buolamwini, J. 2019. Actionable auditing: Investigating the impact of publicly naming biased performance results of commercial AI products. *Proceedings of the 2019 AAAI/ACM Conference on AI, Ethics, and Society*: 429–435.

Raji, I. D., Smart, A., White, R. N., Mitchell, M., Gebru, T., Hutchinson, B., Smith-Loud, J., Theron, D., & Barnes, P. 2020. Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*: 33–44.

Rakova, B., Yang, J., Cramer, H., & Chowdhury, R. 2021. Where responsible AI meets reality: Practitioner perspectives on enablers for shifting organizational practices. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW1), Article 7: 1–23.

Rediker, K. J., & Seth, A. 1995. Boards of directors and substitution effects of alternative governance mechanisms. *Strategic Management Journal*, 16(2): 85–99.

Reuters. 2026a. IPO-bound SpaceX's board is stacked with Musk-empire loyalists. May 21. https://www.reuters.com/legal/transactional/ipo-bound-spacexs-board-is-stacked-with-musk-empire-loyalists-2026-05-21/

Reuters. 2026b. Anthropic seeks 50.1% voting control for co-founders ahead of IPO, The Information reports. September 24. https://www.reuters.com/legal/legalindustry/anthropic-seeks-palantir-style-voting-control-seven-co-founders-ahead-ipo-2026-09-24/

Ritter, J. R. 2026. *IPO statistics* (Tables 4, 4a and 23, updated December 2025). University of Florida. https://site.warrington.ufl.edu/ritter/ipo-data/

Roe, M. J. 2003. *Political determinants of corporate governance: Political context, corporate impact*. Oxford: Oxford University Press.

S&P Dow Jones Indices. 2023. S&P Dow Jones Indices announces results of S&P Composite 1500 index consultation on share class eligibility rules. April 17. https://press.spglobal.com/2023-04-17-S-P-Dow-Jones-Indices-Announces-Results-of-S-P-Composite-1500-Index-Consultation-on-Share-Class-Eligibility-Rules

Sanders, A., & Thomsen, S. (Eds.). 2023. *Enterprise foundation law in a comparative perspective*. Cambridge: Intersentia.

Schröder, D., & Thomsen, S. 2025. Foundation ownership and sustainability. *Journal of Corporate Finance*, 91: 102740.

Schuett, J. 2025a. Three lines of defense against risks from AI. *AI & Society*, 40: 493–507.

Schuett, J. 2025b. Frontier AI developers need an internal audit function. *Risk Analysis*, 45(6): 1332–1352.

Schuett, J., Reuel, A.-K., & Carlier, A. 2025. How to design an AI ethics board. *AI and Ethics*, 5(2): 863–881.

Seaborn, P., Olsen, T. D., & Howell, J. 2020. Is insider control good for environmental performance? Evidence from dual-class firms. *Business & Society*, 59(4): 716–748.

Sevilla, J., Heim, L., Ho, A., Besiroglu, T., Hobbhahn, M., & Villalobos, P. 2022. Compute trends across three eras of machine learning. *Proceedings of the 2022 International Joint Conference on Neural Networks (IJCNN)*: 1–8. IEEE.

Shleifer, A. 2004. Does competition destroy ethical behavior? *American Economic Review*, 94(2): 414–418.

Shleifer, A., & Vishny, R. W. 1986. Large shareholders and corporate control. *Journal of Political Economy*, 94(3): 461–488.

Shleifer, A., & Vishny, R. W. 1997. A survey of corporate governance. *Journal of Finance*, 52(2): 737–783.

SGX (Singapore Exchange). 2018. SGX launches rules for listing of dual class shares companies. June 26. https://links.sgx.com/1.0.0/corporate-announcements/AKWHJIGEQZEIDPKX/20180626_SGX_launches_rules_for_listing_of_DCS_companies.pdf

Smith, C. W., Jr., & Watts, R. L. 1992. The investment opportunity set and corporate financing, dividend, and compensation policies. *Journal of Financial Economics*, 32(3): 263–292.

SpaceX. 2026a. *Prospectus (Form 424(b)(4))*, Space Exploration Technologies Corp., June 11. https://www.sec.gov/Archives/edgar/data/1181412/000162828026042639/spaceexplorationtechnologi.htm

SpaceX. 2026b. Quarterly report (Form 10-Q) for the quarterly period ended September 30, 2026. October 22. https://www.sec.gov/Archives/edgar/data/1181412/000162828026066880/spcx-20260930.htm

State of California. 2025. Transparency in Frontier Artificial Intelligence Act (SB 53). Chapter 138, Statutes of 2025. https://leginfo.legislature.ca.gov/faces/billNavClient.xhtml?bill_id=202520260SB53

State of Delaware. 2025. Delaware General Corporation Law, section 144, as amended by Senate Substitute 1 for Senate Bill 21. https://delcode.delaware.gov/title8/c001/sc04/

State of Texas. 2025. Business Organizations Code, chapter 21, section 21.552, as amended by Senate Bill 29. https://tcss.legis.texas.gov/resources/BO/htm/BO.21.htm

Taeihagh, A. 2021. Governance of artificial intelligence. *Policy and Society*, 40(2): 137–157.

Tallarita, R. 2023. AI is testing the limits of corporate governance. *Harvard Business Review*, December 5.

Taylor, B. 2025. Built to benefit everyone. OpenAI, October 28. https://openai.com/index/built-to-benefit-everyone

TechCrunch. 2026. Anthropic's founders seek voting control ahead of IPO. September 25. https://techcrunch.com/2026/09/25/anthropics-founders-seek-voting-control-ahead-of-ipo/

Tesla. 2025. Current report (Form 8-K), shareholder approval of the 2025 CEO Performance Award. November 7. https://www.sec.gov/Archives/edgar/data/1318605/000110465925108507/tm2530590d1_8k.htm

Thomsen, S. 2017. *The Danish industrial foundations*. Copenhagen: Djøf Publishing.

Thomsen, S., & Conyon, M. 2019. *Corporate governance and board decisions*. Copenhagen: Djøf Forlag.

Thomsen, S., & Kavadis, N. 2022. Enterprise foundations: Law, taxation, governance, and performance. *Annals of Corporate Governance*, 6(4): 227–333.

Thomsen, S., & Rose, C. 2004. Foundation ownership and financial performance: Do companies need owners? *European Journal of Law and Economics*, 18(3): 343–364.

Thomsen, S., Poulsen, T., Børsting, C., & Kuhn, J. 2018. Industrial foundations as long-term owners. *Corporate Governance: An International Review*, 26(3): 180–196.

Toner, H., & McCauley, T. 2024. AI firms mustn't govern themselves, say ex-members of OpenAI's board. *The Economist*, May 26. https://www.economist.com/by-invitation/2024/05/26/ai-firms-mustnt-govern-themselves-say-ex-members-of-openais-board

Trautman, L. J., & Foster, L. D., II. 2025. Sam Altman, OpenAI, and the importance of corporate governance. *Case Western Reserve Journal of Law, Technology & the Internet*, 16(1): 133.

Tribunale di Roma. 2026. Judgment in OpenAI v. Garante per la protezione dei dati personali. R.G. 4785/2025, March 18. https://dei.web.uniroma1.it/sites/default/files/allegati/2026-05/Trib_Roma_OpenAI_Garante_2026.pdf

U.S. Department of Commerce. 2025. Statement from U.S. Secretary of Commerce Howard Lutnick on transforming the U.S. AI Safety Institute into the pro-innovation, pro-science U.S. Center for AI Standards and Innovation. June 3. https://www.commerce.gov/news/press-releases/2025/06/statement-us-secretary-commerce-howard-lutnick-transforming-us-ai

Villalonga, B., & Amit, R. 2006. How do family ownership, control and management affect firm value? *Journal of Financial Economics*, 80(2): 385–417.

Walsh, J. P., & Seward, J. K. 1990. On the efficiency of internal and external corporate control mechanisms. *Academy of Management Review*, 15(3): 421–458.

Wan, A., Klyman, K., Kapoor, S., Maslej, N., Longpre, S., Xiong, B., Liang, P., & Bommasani, R. 2025. The 2025 foundation model transparency index. *arXiv*: 2512.10169.

Ward, A. J., Brown, J. A., & Rodriguez, D. 2009. Governance bundles, firm performance, and the substitutability and complementarity of governance mechanisms. *Corporate Governance: An International Review*, 17(5): 646–660.

Wasserman, N. 2023. OpenAI's failed experiment in governance. *Harvard Business Review*, November 30.

Weiss, G. 2026. Aligned structuring of AI startups. *Arizona State Law Journal*, 58(1): 429–495.

Westphal, J. D. 2023. Systemic symbolic management, corporate social responsibility, and corporate purpose: A cautionary tale. *Strategy Science, 8*(2): 221–232.

Williamson, O. E. 2000. The new institutional economics: Taking stock, looking ahead. *Journal of Economic Literature*, 38(3): 595–613.

WIPO (World Intellectual Property Organization). 2026. Global unicorn valuation 2025. February 2. https://www.wipo.int/en/web/global-innovation-index/w/blogs/2026/global-unicorn-valuation-2025

Wright, L., Muenster, R. M., Vecchione, B., Qu, T., Cai, P., Smith, A., Metcalf, J., Matias, J. N., et al. 2024. Null compliance: NYC Local Law 144 and the challenges of algorithm accountability. *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*: 1701–1713.

Zhang, H., Ye, H., Dong, L., & Wang, J. 2025. CSR performance of dual-class firms in different life cycle stages. *Applied Economics*, 57(22): 2839–2852.

Zingales, L. 2000. In search of new foundations. *Journal of Finance*, 55(4): 1623–1653.

european corporate governance institute

About ECGI

ECGI is an international non-profit membership organisation that provides a platform for debate and dialogue among academics, policymakers, and business leaders, with a focus on major corporate governance, ESG, and stewardship issues.

The views expressed in this working paper are those of the authors, not those of the ECGI or its members.

www.ecgi.global

european corporate governance institute

ECGI Working Paper Series in Law

Editorial Board

Editor	**Amir Licht,** Professor of Law, Harry Radzyner Law School, Reichman University
Consulting Editors	**Hse-Yu Iris Chiu,** Professor of Corporate Law and Financial Regulation, University College London **Martin Gelter,** Professor of Law, Fordham University School of Law **Geneviève Helleringer,** Professor of Law, ESCP and Oxford Law Faculty **Kathryn Judge,** Professor of Law, Columbia Law School **Wolf-Georg Ringe,** Professor of Law & Finance, University of Hamburg
Series Manager	**Asif Malik,** ECGI Working Paper Series Manager

https://www.ecgi.global/publications/working-papers