

CORPORATE ORDERING

How Corporations Navigate Social Conflict

Cambridge University Press (forthcoming October 2026)

Stavros Gadinis

George R. Johnson Professor of Law, UC Berkeley Law School
Faculty Co-Director, Berkeley Center for Law and Business

Corporate law is facing a challenge it was not designed to meet. Its foundational tools, such as fiduciary duties, the business judgment rule, and disclosure requirements, were built around the assumption that corporations can identify a clear objective and pursue it. That assumption presupposed a degree of alignment among shareholders and managers that modern corporations can no longer take for granted. Shareholder primacy assumes that shareholders share preferences that can be aggregated into a single mandate. Stakeholder theory introduces the interests of employees, communities, and other impacted parties, but assumes that these interests can be weighed against each other on a common scale. Both assumptions collapse in the face of the defining condition of modern corporate life: heterogeneity of preferences so deep that it resists reduction to any common denominator.

Modern corporations bring together shareholders with incommensurable time horizons, risk tolerances, and moral commitments. They employ workers whose values and identities are bound up in what the firm does and how it does it. They operate at a scale that makes their decisions consequential for communities, ecosystems, and democratic institutions that have no formal seat at the table. When a firm decides how to moderate speech, how to account for carbon across its supply chain, or how to set the parameters for an artificial intelligence system, it is not managing externalities. It is making choices that are simultaneously business decisions and acts of governance, and no existing framework in corporate law tells it how.

The book argues that leading corporations have not waited for law to catch up. Quietly and without a theoretical vocabulary to describe what they were doing, they have built a new form of governance infrastructure, which this book calls corporate ordering, to navigate decisions that carry genuine social conflict at their core. This infrastructure is not corporate social responsibility or ESG rebranded, nor is it political activism by another name. It is an institutional response to a structural problem, built around the core of the business model rather than at its edges, and organized around a precise architecture that mirrors the functional logic of legal orders themselves.

This argument is grounded in conversations with practitioners who built these governance structures from the inside. In ridesharing, those conversations included Tony West, General Counsel of Uber; Kristin Sverchek, former General Counsel and President of Lyft; and Kate Sampson, an insurance specialist who worked on the coverage frameworks that became the template for state regulation. In climate sustainability, the research drew on conversations with Ryan Hart at Salesforce, whose team developed the emissions accounting methodology that prefigured mandatory disclosure. In artificial

intelligence governance, the research included conversations with Che Chang at OpenAI, Tom Lue at Google DeepMind, and Jon Palmer and Daniel Klutz at Microsoft. These practitioners were not sources of quotable opinions but windows into how governance problems were actually experienced, and how the tools that now shape emerging public law came to be designed.

That architecture works through three interconnected functions. Rulemaking establishes standards before individual decisions arise, through broad consultative processes designed to surface the full range of affected interests and relevant expertise. Executive decision-making disciplines the operational choices that implement those standards, through disclosure requirements that force explicit engagement with social dimensions rather than allowing them to be absorbed silently into business necessity. Monitoring relies on independent oversight mechanisms to verify that practice aligns with stated principles and to generate the institutional memory that allows standards to evolve. Together, these three functions create a capacity for managing social complexity that is architecturally comparable to what director independence created for management oversight and what audit committee independence created for financial integrity.

The book's deeper claim is that corporate ordering is not merely a best practice for enlightened management but a structural feature of how modern societies govern themselves. Regulators and legislators frequently arrive after corporations have already created the templates they subsequently formalize. In platform governance, climate sustainability, and artificial intelligence, the first draft of regulatory practice is already being written inside firms. Recognizing this dynamic transforms what legal scholars, policymakers, investors, and business leaders should be asking. The question is no longer whether corporations should engage with social conflict. They already do, at scale, with consequences that reach millions of people. The question is whether the governance processes through which they do so are robust enough to deserve the authority they already exercise, and whether corporate law has the tools to assess them.

The argument proceeds in ten chapters organized around three movements: establishing the problem, developing the framework, and demonstrating it across three substantive domains before turning to legal reform.

Chapters 1 and 2 — The Problem

Chapter 1 identifies the structural forces that have dissolved the traditional boundary between business decisions and social governance: efficiency at scale, the forced transparency of the digital age, and the fragmentation of shareholder preferences into irreconcilable camps. Chapter 2 analyzes the resulting gap in corporate law doctrine. The business judgment rule offers discretion but no guidance when social dimensions cannot be separated from operational choices, and existing fiduciary duties and disclosure requirements do not supply the procedures needed to manage genuine value conflict.

Chapters 3 and 4 — The Framework

Chapter 3 provides the full theoretical articulation of corporate ordering, defining the concept precisely and developing the tripartite architecture in detail, with grounding in legal theory and political philosophy. Chapter 4 addresses the central objection: why not simply regulate? The chapter explains why corporations routinely act before regulation appears and how the tripartite framework can serve as a diagnostic instrument, identifying when corporate self-regulation carries genuine governance capacity and when external intervention is warranted.

Chapter 5 — Historical Foundations

This chapter traces corporate ordering to earlier episodes in which private firms pioneered regulatory frameworks that governments subsequently adopted as mandatory law, including early traffic regulation and the development of financial disclosure practices that culminated in the Securities Act of 1933. The chapter establishes that corporations have long functioned as incubators of legality, making the contemporary phenomenon a structural feature rather than a novelty.

Chapters 6, 7, and 8 — Three Case Studies

These chapters constitute the empirical core of the book, each built around conversations with the practitioners who designed the governance structures under study. Chapter 6 draws on those conversations to reconstruct how Lyft and Uber, working without regulatory guidance and under intense public pressure following a series of safety incidents, developed insurance and passenger safety frameworks that multiple states subsequently wrote into law. Chapter 7 traces how an internal team at BP, with no mandatory template to follow, took the first steps to establish the GHG Protocol that now underlies mandatory disclosure regimes taking shape on both sides of the Atlantic. Chapter 8 draws on conversations at Microsoft, OpenAI, and Google DeepMind to show how engineers, ethicists, and lawyers built AI governance structures in real time, without settled law or industry consensus, and how those internal choices became the foundation on which public oversight is now being constructed. Across all three chapters, the same analytical structure recurs: business and social dimensions become inseparable, preferences fragment, corporate ordering emerges, and regulators subsequently build outward from what firms have already created.

Chapter 9 — Legal Reform

Drawing on the case study evidence, this chapter proposes incorporating the best practices of corporate ordering into fiduciary duty. Corporate law has historically absorbed new decision-making tools once they became necessary for informed judgment, and the same logic now applies to social deliberation. The proposed standard requires boards to establish governance structures capable of genuine inquiry and independent oversight when credible evidence indicates significant social dimensions, without prescribing substantive outcomes or overriding business judgment.

Chapter 10 — Conclusion

The conclusion synthesizes the argument across three registers. Descriptively, it presents corporate ordering as a primary site of legal formation, where regulatory frameworks originate in corporate deliberation before governments formalize them. Normatively, it argues that deliberative quality provides the best available mechanism for corporate decision-making that affects public outcomes when electoral authorization is unavailable and shareholder preferences are irreconcilable. Methodologically, it calls for legal scholarship that treats corporate governance processes as a central rather than peripheral window into how contemporary societies actually govern themselves.

Contents of This Packet

Chapter 3, Corporate Ordering for Social Conflicts	pp. 95–128
Notes to Chapter 3	
Chapter 8, How to Govern AI? Opening and Sections 8.1–8.2	pp. 257–264
Chapter 8, Sections 8.5–8.6	pp. 276–297
Notes to Chapter 8	

Page numbers refer to the printed edition. Because the excerpts follow section boundaries, some pages include the end of a preceding section or the start of the next. Notes pages are reproduced in full and include notes to adjacent chapters. © Stavros Gadinis 2026. Published by Cambridge University Press & Assessment. DOI 10.1017/9781009873741.

3

Corporate Ordering for Social Conflicts

Chapters 1 and 2 revealed a striking gap in contemporary corporate law. Social considerations arise amidst significant uncertainty in many contexts, from developing sustainable product lines to entering emerging markets with complex governance challenges, and from disrupting traditional industries through new technologies to restructuring supply chains to address labor standards. In these situations, corporate directors and managers face decisions with enormous stakes, competing stakeholder preferences, and intense public scrutiny. Yet they must ultimately decide how to proceed with little or no guidance from corporate law. The uncertainty is compounded by the reality that different shareholders often have sharply divergent views about what constitutes responsible corporate behavior. These high-stakes choices unfold under the watchful eyes of activists, media, and regulators, amplifying the pressure on decision-makers.

Managers' choices will almost certainly be protected by the business judgment rule, so long as they can point to good faith and some evidence of deliberation. The ultimate outcome matters little; the process protection remains robust.¹ True, shareholder activists might stir up controversy at the annual meeting, and disclosure obligations may require awkward admissions in securities filings when an issue grows sufficiently to satisfy materiality thresholds. But ultimately, corporate law offers virtually no guidance on how these momentous decisions should be made. It does not prescribe what processes to follow, what factors to weigh, what deliberative structures to employ, or what standards of analysis to apply. Executives enjoy breathtaking discretion to chart whatever course they deem appropriate, armed with little more than their personal judgment and the comforting knowledge that courts will rarely second-guess their choices.²

Some would argue this is not a gap at all, but rather a defining feature of corporate law. On this view, the law deliberately leaves managers free to exercise their judgment as they see fit, positioned at the apex of a corporate hierarchy that is essential for efficient business operations. Chief executives, while not autocrats, serve as singular leaders and ultimate decision-makers, and this concentrated authority enables the swift, decisive action that competitive markets demand.³ Social considerations,

by contrast, embody the messy complexity of democratic society, an amalgamation of diverse viewpoints, competing interests, varied perspectives, and distinct backgrounds. Our political institutions manage this complexity through intricate networks of multiple decision-makers, elaborate checks and balances, and sophisticated institutional interactions. Such democratic complexity, the argument goes, would only come at the cost of the streamlined hierarchy that makes corporate governance effective.⁴

Yet this conception of managerial authority fundamentally mischaracterizes the decision-making processes that corporate law actually envisions and encourages. If we truly embraced managers as quasi-autocratic leaders limited only by external regulation, corporate governance mechanisms would be entirely superfluous. Instead, our governance framework reflects a more sophisticated understanding of managers as informed decision-makers who benefit from, and should be subject to, reasoned processes of information gathering and consultation, overseen by boards of directors and accountable to shareholders.⁵

Consider how corporations approach conventional business challenges. Managers do not make decisions in isolation, nor are they expected to rely solely on their personal judgment. Rather, they engage in elaborate processes of information collection and consultation that span both internal and external sources. They canvass markets to understand customer preferences and competitive dynamics. They commission product research to assess technical feasibility and consumer appeal. They conduct extensive market research to gauge demand and identify optimal positioning strategies. They establish specialized departments dedicated to generating the information necessary for sound decision-making, such as research and development, strategic planning, and government relations. They create board committees with specific expertise to oversee critical functions such as audit, compensation, and risk management.⁶

When major decisions loom, managers routinely seek guidance from networks of specialists: investment bankers for mergers and acquisitions, financial advisors for capital structure decisions, marketing experts for brand positioning, public relations specialists for communications strategy, and management consultants for operational improvements.⁷ The consulting industry, epitomized by firms such as McKinsey & Company, has never been perceived as undermining managerial prerogatives, but rather as enhancing executive capabilities. Even entrepreneurial founders in startup environments rely heavily on venture capital partners, who bring experience and market sophistication to navigate business uncertainties.⁸ The corporate world has developed sophisticated frameworks governing these information-gathering and advisory relationships, encompassing everything from financing arrangements to gatekeeping functions to credit rating processes.⁹

Crucially, the various steps corporations have developed for traditional business decisions have not remained merely matters of best practice. Over time, they have become embedded in corporate law itself. Courts now expect boards to demonstrate that they have followed appropriate deliberative processes, consulting relevant

experts, gathering pertinent information, and utilizing specialized committees when making major business decisions as appropriate.¹⁰ These tools might have begun as market-generated efforts to address a business need, but have since evolved into legally relevant steps boards must fulfill to claim business judgment rule protection, mitigate conflicts of interest, demonstrate their awareness of securities disclosures, or avoid compliance-related liability.¹¹ Yet as to social issues, no such procedural steps have been established or required. The law remains silent on what constitutes adequate deliberation in this domain.¹²

What this book documents is how companies are already developing and implementing innovative frameworks to enhance their decision-making processes regarding social issues. These developments reflect a growing recognition among corporate leaders that, rather than relying on their own intuitions, the prudent approach involves systematic information gathering and consultation, drawing on stakeholder input, expert analysis, specialized research, and structured deliberation.¹³ This represents precisely the same principle that guides sound decision-making in other areas of business operations.

This chapter synthesizes the methods, practices, and approaches that leading corporations have adopted to build effective decision-making processes for navigating complex social issues. From these observed practices, this book develops a new theoretical framework that connects the functions corporations seek to fulfill with the specific methodologies they select, illustrated through concrete examples that demonstrate these frameworks in action. Together, these functions form a framework this book terms “corporate ordering.” The following chapters then provide in-depth case studies that showcase how these principles apply across diverse substantive areas: rideshare regulation ([Chapter 6](#)), sustainability ([Chapter 7](#)), and artificial intelligence governance ([Chapter 8](#)). Together, these analyses reveal how systematic approaches to decision-making on social issues can enhance both corporate performance and societal outcomes.

The systematization and theoretical development of corporate ordering is not merely an academic exercise, but an essential tool for protecting shareholder interests. Just as corporations have recognized that informed decision-making in traditional business matters serves shareholders by reducing risks and improving outcomes, the same logic applies to social issues. When companies make social decisions without adequate information gathering and consultation processes, they expose shareholders to reputational risks, regulatory sanctions, consumer backlash, and operational disruptions that could have been anticipated and mitigated through better decision-making frameworks. Performing this function well, by establishing systematic processes for gathering relevant information and engaging appropriate stakeholders before making social decisions, ought to represent a core fiduciary service to shareholders.

This recognition demands that corporate law evolve to acknowledge and support these emerging practices. Yet normative debates about whether corporations

should address social issues at all have obscured the more fundamental question of how they should make such decisions when they inevitably do. [Chapter 9](#) will explore at length how these disputes have divided legal theorists and practitioners. Yet this divisiveness has prevented legal scholarship and doctrine from developing the analytical frameworks necessary to evaluate and improve corporate decision-making processes in this domain. The time has come to move beyond these normative impasses and focus on the practical imperative of ensuring that when corporations engage with social issues, as they unavoidably do, they employ processes that serve shareholder interests through informed, systematic decision-making.

3.1 ANALYZING CORPORATE FUNCTIONS UNDER THE DISTINCTION OF POWERS

Conventional wisdom thinks of corporate managers as exercising purely executive power. After all, the person at the top bears the title of Chief Executive Officer, and all corporate authority flows through this single individual under delegation from the same governing body, the board of directors.¹⁴ The organizational chart appears to show a straightforward hierarchy of executive power. When we observe corporate decision-making, it seems to follow familiar patterns of managerial authority: leaders set strategic direction, allocate resources, and respond to operational challenges through a unified chain of command. This intuitive understanding aligns with our conventional view of business management as focusing on executive decisions.

Yet this superficial simplicity obscures a remarkable reality. What appears to be a vertical line to executive authority is actually a complex system performing multiple functions under a single institutional umbrella.¹⁵ All power structures inherently perform three distinct yet interrelated functions: a legislative, an executive, and a judicial one. This analytical framework forms one of the most enduring insights of political theory, tracing its intellectual lineage back to Aristotle. It has of course found its most prominent modern application in constitutional design, where the separation of powers into distinct institutions serves as a cornerstone of democratic governance.¹⁶

However, the utility of this tripartite distinction extends far beyond the realm of constitutional and administrative law. Indeed, whether examining military commands, university rules, or private clubs, every hierarchy must establish norms (legislative function), implement decisions (executive function), and adjudicate conflicts (judicial function), revealing these three dimensions as fundamental pillars of organized authority itself.¹⁷ When we peer beneath the unified facade of corporate hierarchy, the full spectrum of the three functions readily emerges. Just as any hierarchical body, corporations have also evolved, consciously or not, to perform analogous functions through different mechanisms and with different constituencies in mind.

By applying the tripartite distinction to rulemaking, executive, and monitoring functions onto corporate decision-making, I do not mean to suggest that corporations are akin to states, mini-societies, or bearers of democratic obligations. Such analogies obscure rather than illuminate the unique characteristics of corporate influence. Treating corporations as if they were governments prevents us from understanding the true profile of corporate power, with its distinct strengths and limitations. Corporate influence operates through fundamentally different mechanisms than state power, as it relies on contractual relationships, market positioning, and economic leverage rather than sovereign authority, legal compulsion, and democratic legitimacy. Corporate decisions, while often far-reaching in their effects, lack the formal binding force of law and operate within constraints that differ markedly from those facing governmental actors.¹⁸

Rather, this approach rests on a fundamentally functional argument. Corporations, by their very nature as hierarchical organizations wielding significant economic and social influence, inevitably make choices that fulfill one or more of these essential governance functions. Thus, they are choosing decision-making tools that best serve the goals of each function. These are the choices that [Sections 3.1–3.4](#) will try to illuminate.

3.1.1 *The Tripartite Distinction as an Analytical Tool for Corporate Functions*

The functional differentiation becomes clear when we examine how corporate decisions actually unfold. Rulemaking functions occur when managers create frameworks of general and uniform applicability. These decisions typically face higher stakeholder heterogeneity because general rules apply across many different circumstances, potentially implicating problems managers may not fully understand or anticipate.¹⁹ The broad scope increases visibility and elevates the perceived importance of the solution, often seen as embodying management's values and predicting future moves. Such decisions benefit from wide consultation processes that engage large numbers of stakeholders, reach out to potential opponents, and provide opportunities for direct participation in decision-making.²⁰

Executive functions involve individual decisions with far-reaching implications, ranging from implementing long-term strategic visions, such as building new facilities or launching products, to responding quickly to emerging needs. While stakeholder engagement may be valuable, it can be difficult to organize effectively. These decisions are best served by robust disclosure practices that communicate choices and rationales, notify interested parties, and invite future engagement while building trust through transparency and accountability.²¹

Monitoring functions encompass specialized oversight and investigation of past failures and mistakes. These mechanisms both help managers maintain control over large organizations and constrain managerial behavior through statutory mandates,

regulatory obligations, and voluntary commitments. Given well-established incentives for managers to shirk responsibility or cover up misconduct, these functions raise critical concerns about independence, integrity, thoroughness, and access to resources.²²

The tripartite distinction provides analytical clarity by recognizing that different corporate functions require correspondingly different decision-making processes and governance approaches. This framework serves as an analytical tool rather than a rigid institutional blueprint, illuminating both the inherent bargaining imbalances between corporations and stakeholders and the distinct audiences each type of corporate action is intended to address. While ultimate authority formally resides with boards of directors and chief executives, these actors frequently choose to delegate specific functions to different institutional formations, both within the corporate hierarchy and external to it, that possess specialized capabilities or legitimacy to fulfill these distinct rulemaking, executive, and monitoring roles more effectively than centralized leadership by itself.²³

This tripartite analysis illuminates the true architecture of corporate power, revealing the complex functional edifice hidden within what appears to be simple business management. Understanding these distinctions is crucial because each function demands different governance approaches, accountability mechanisms, and stakeholder engagement processes. Corporate law's failure to recognize these functional differences leads to governance frameworks that treat fundamentally different types of corporate power as if they were identical, undermining both corporate effectiveness and shareholder accountability.

3.1.2 *Mapping Corporate Functions to Governance Solutions*

The practical application of this framework becomes evident when examining how different corporate functions demand correspondingly different governance approaches. Rather than adopting uniform oversight mechanisms across all business activities, effective corporate governance must recognize that each type of corporate decision-making presents unique challenges, stakeholder relationships, and performance requirements that necessitate tailored supervisory structures.²⁴ The accompanying table synthesizes these critical relationships, demonstrating that effective corporate governance cannot rely on one-size-fits-all solutions but must instead calibrate its mechanisms to the specific demands of each functional category.

Rulemaking functions, which involve creating frameworks of general and uniform applicability, require governance approaches that can handle high stakeholder heterogeneity and broad visibility. The governance goals center on wide outreach and notice-and-comment processes that engage diverse constituencies who may be affected by general rules applied across varying circumstances. Examples like ride-sharing platform policies or post-#MeToo workplace reforms illustrate how companies must reach beyond typical advisory channels to gather perspectives from stakeholders who may experience different impacts from uniform policies.²⁵

TABLE 3.1 *Mapping corporate functions to governance solutions*

Type of power	Application to corporations	Governance goals	Examples
<i>Rulemaking</i>	Creating frameworks of general and uniform applicability	Wide Outreach Notice-and-Comment	Rideshare; Post- #MeToo reforms
<i>Executive</i>	Individual decisions with far-reaching implications	Disclosure Accuracy	Net zero commitments
<i>Review</i>	Specialized oversight Averting conflicts	Independence Investigation	Ethics Committees; Civil Rights Audits

Executive functions, characterized by individual decisions with far-reaching implications, call for governance mechanisms focused on disclosure and accuracy. These decisions, whether strategic implementations like net zero commitments or urgent responses to emerging challenges, benefit from clear communication of choices and rationales rather than extensive prior consultation.²⁶ The governance priority shifts from broad participation to transparent explanation and reliable follow-through.²⁷

Review functions, encompassing specialized oversight and conflict avoidance, demand governance structures that prioritize independence and thorough investigation. Ethics committees, civil rights audits, and similar monitoring mechanisms must demonstrate credibility through structural independence from the decisions they evaluate and the resources necessary to conduct meaningful investigations.²⁸

This mapping reveals that corporations seeking to optimize their governance cannot simply adopt generic best practices, but must thoughtfully match their governance tools to the functional demands of different types of decisions. The framework thus serves not as a rigid template but as a diagnostic tool for identifying which governance mechanisms will be most effective for particular corporate choices. The following sections systematically examine how these principles apply across major corporate functional areas, with each section addressing a separate domain of business activity (Table 3.1).

3.2 CORPORATE RULEMAKING

The mere suggestion that corporations are themselves actively engaged rulemakers will raise eyebrows in some quarters. Lobbying by corporate interests has been a long-established concern in modern democracies, where politicians often find themselves dependent on large donations to launch successful electoral campaigns. These sums of money, the fear is, are only handed over in expectation that, once elected, politicians will pass laws that favor corporate perspectives, facilitate business goals, or otherwise offer preferential treatment to their supporters.²⁹ On the

regulatory side, corporations are widely considered well-placed to utilize administrative rulemaking to their benefit, by mustering resources to more effectively comment on policy proposals, submit complaints and requests, and represent themselves in exploratory hearings.³⁰ Moreover, corporations have organized into industry groups like the Business Roundtable that have successfully pressured governments and taken them to court when necessary.³¹ This corporate effectiveness in pushing for favorable legislation is in stark contrast to the nameless interests of the public, which unavoidably take a back seat. So, if we distrust corporate involvement in formal lawmaking by elected officials and their appointed executives, what can we expect when corporations themselves are the rulemakers? Why should we even entertain this idea?

In reality, I will argue here, corporate rulemaking is a routine activity in the market that corporations have undertaken since their inception, but it has become far more systematic and consequential as the drive for efficiency at scale has reshaped corporate operations. As outlined in [Chapter 1](#), modern corporations coordinate millions of employees, serve billions of customers, and orchestrate supply chains spanning dozens of countries. To achieve efficiency at this unprecedented scale, they rely on uniform standards and comprehensive monitoring systems that inevitably extend corporate influence far beyond traditional commercial relationships. The imperative of uniformity and standardization transforms individual business decisions into frameworks of general applicability that function as *de facto* social policy,³² whether in Apple's App Store policies affecting millions of apps worldwide, Amazon's fulfillment processes standardized across global operations, or Uber's platform rules governing interactions between drivers and passengers across dozens of countries.

Sometimes, corporations set out terms, conditions, norms, or practices that govern their interaction with their stakeholders and customers on a large scale. While some of these might be dictated or constrained by law or regulation, many of them operate in the private or voluntary domain. They represent management's preferred conditions for doing business or, more generally, for entering into a relationship with the party or parties involved. What makes these corporate choices distinctly rule-like, rather than mere contractual terms, is precisely this drive for standardization across vast operations. When Tesla establishes direct-to-consumer sales policies or when Facebook creates content moderation standards, these decisions must apply uniformly across millions of interactions because individual case-by-case determinations would be prohibitively costly and operationally unworkable.³³ The same mechanisms that enable extraordinary corporate efficiency also make corporate decisions into building blocks of social organization.

3.2.1 *Defining Corporate Rules: Generality and Uniformity*

What makes these corporate choices more like rules, and less like individual contractual terms, is that they apply in a general and uniform manner to a large set

of actors. First, these rules exhibit generality, as they are meant to encompass a class or category of people interacting with the company in the same way. Often, the corporation makes the choice without even knowing who the persons subject to it are going to be, because the company does not know the group of consumers or employees or local citizens with which it plans to engage.³⁴ Sometimes, the corporate choice is not even directed at particular persons, but simply chosen to define a term that is going to be useful to many individuals, such as what a company considers appropriate conduct for someone entering its premises. The generality of these rules is often embedded in the design of the product or service that the company is planning to offer, which is built to be used in specific ways, to accommodate or prioritize specific needs, and to appeal to a large audience.³⁵

General rules are important because they apply uniformly, in ways that are identical or under exemptions prescribed in the rule. Uniformity is a key building block of any hierarchy, because it allows managers to administer and coordinate their subordinates. Without uniformity, managers would have to engage in individual determinations, an often impractical task. Instead, managers can reduce transaction costs by reaching a determination once and avoid revisiting it until a problem occurs.³⁶ However, the decision to adopt a uniform rule also alters calculations for managers. By resolving all arising cases at once, management needs to “get it right,” making sure that their choice is, from a substantive perspective, the one that will better serve the corporation over a range of potential cases with slight variations.

But while rules clarify expectations and allow for greater predictability and better planning, they are also more likely to produce noticeable effects, since their results are going to be similar and more easily aggregated.³⁷ Thus, those impacted are going to understand the implications of the rule, as well as that there are other stakeholders in the same predicament as themselves. Moreover, general and uniform rules can be more easily communicated to third parties outside the corporate hierarchy, like consumers, investors, government authorities, and the media. When management adopts a rule, their choice will now be more visible, as it will affect greater numbers of cases all at once. It will be judged both by the group to which it may be directly addressed, as well as by other groups who may not be affected by it but who regularly interact with management.³⁸

Most rules and rulemaking frameworks discussed earlier take the guise of contractual commitments, and companies may seek stakeholders’ consent or consider them voluntarily incorporated in employment or other contracts. But these terms constitute contracts in formal nature only. They have not been the result of negotiations with impacted parties, which sometimes neither take part in nor are represented in their drafting. Parties rarely have the opportunity to voice consent or disagreement; at most, stakeholders may be presented with the option to either agree to these terms or reject any relationship with the corporation. Consent is often given in a formulaic manner, with stakeholders signing onto frameworks at moments when considering their implications is impractical.³⁹

However, the consequences of violating these rules are not only real, but often quite painful for stakeholders. Stakeholders directly subject to the corporate hierarchy, such as employees or suppliers, may be subject to disciplinary procedures, remuneration reductions, or negative reputational consequences.⁴⁰ To parties whose relationship with the corporation is more coincidental, managers can deny cooperation if they fail to conform with its terms, an option that may put smaller subcontractors out of business. When corporations have superior bargaining power compared to stakeholders, the line between consent and coercion begins to blur.⁴¹

3.2.2 *Corporate Rules in Practice: Application and Examples*

Once we understand corporate rules as key components of corporate hierarchies, it becomes clear that managers have relied on them routinely to organize production, coordinate subordinates, and establish efficient exchanges with third parties. Below are some examples.

3.2.2.1 Production and Quality Control

The cardinal mission of every firm is to organize production. In performing these tasks, managers often make determinations that take the shape of general and uniform rules that encapsulate the firm's approach to whole clusters of decisions.⁴² For example, managers may decide that they will only use packaging material that satisfies certain biodegradability requirements, occupy offices that obtain LEED certification, or hire employees with specific threshold qualifications. Each individual choice may have a social dimension, whether helping attain sustainability goals, setting terms that govern entire populations of workers, or impacting other stakeholders more generally.⁴³ These determinations are replicated in numerous choices around the firm, which aggregates their impact and renders them more noticeable to affected groups or external watchdogs. Particularly when a firm's business scales up, the effort to install a common approach to similar problems takes the shape of uniform and generally applicable rules.⁴⁴ By setting corporate policies that employees must follow, the company embraces a rulemaking function that helps coordinate production while solidifying the choices managers make vis-à-vis other stakeholders.

3.2.2.2 Customer Interface

An aspect of production that has gained renewed attention in recent years concerns the customer interface.⁴⁵ Besides statutory or regulatory requirements, managers maintain significant leeway in structuring their firm's customer interface.⁴⁶ Companies have sought to take advantage of new technologies that allow more direct access to consumer data and preferences, more effective monitoring of consumer conduct, and greater use of social media to target specialized audiences.⁴⁷

These developments have brought new concerns regarding the monetization of consumer data, protection for privacy and free speech, and securing access to markets for disadvantaged groups.⁴⁸ Addressing these concerns often requires companies to shape frameworks that their employees can follow in a general and uniform manner, such as establishing protocols for managing the storage and use of consumer data.⁴⁹

The wide availability of product information has also changed the customer interface. In the past, consumer mobilization against corporate misconduct required significant effort, often by dominating national headlines or gaining regulatory attention.⁵⁰ Today, both monitoring and mobilization are readily collapsed into smartphone apps, while complex information becomes available at consumers' fingertips. This increase in disclosure requires companies to be more intentional about their choices and monitor public reactions more closely.⁵¹

3.2.2.3 Organizing the Workplace

In the centuries-old fight between capital and labor, the companies' power to regulate the workplace has traditionally featured as a critical battlefield. Managers have set out rules for structuring the internal line of command, how employees should behave toward each other, and how they should treat third parties. They determine dress codes, what customer requests employees can satisfy, and take possession of employee communications.⁵² Managers also play a critical role in defining employees' benefits and privileges beyond government minimums: parental leaves, health care coverage, and pension contributions involve important choices that are critical not only for the employees in question, but for society more generally.⁵³

A long literature has examined how these rule-based frameworks reproduce social conventions, structures, hierarchies, and values prevalent outside the firm's boundaries, turning the workplace into a microcosm of society.⁵⁴ Gender and race discrimination have been central areas of exploration, along with class origins, LGBTQ+ concerns, political and religious beliefs, immigration status, and other issues. Managers have the power to embrace, strengthen, or disregard these external factors.

In recent years, social movements have embraced the corporation not only as a reflection of society, but as a vehicle for change. Following #MeToo and Black Lives Matter, corporations became preferred terrain for social reformers seeking quick and meaningful progress. Popular accounts present this approach as a solution to political gridlock, but they tend to underestimate the significance of this strategic shift. By choosing to focus on changes within corporations, social movements prioritize the implementation phase, ensuring that reform goals are carried out where they are more likely to matter.⁵⁵

The initial corporate response was swift and comprehensive. Companies established chief diversity officers, implemented bias training programs, revised hiring practices to emphasize diverse candidate pools, and committed billions of dollars to racial equity initiatives.⁵⁶ Major corporations from Goldman Sachs to Starbucks

adopted policies requiring diverse interview panels, set numerical diversity targets for leadership positions, and tied executive compensation to diversity metrics.⁵⁷ These represented systematic changes to corporate rulemaking that affected millions of employees and shaped industry practices across entire sectors.

However, the pendulum has swung dramatically in the opposite direction. Beginning around 2023 and accelerating under the Trump administration's explicit anti-DEI positioning, corporations have systematically retreated from diversity initiatives. Major companies including Meta, Amazon, McDonald's, and Walmart have eliminated diversity, equity, and inclusion programs, removed diversity language from job postings, and disbanded DEI teams.⁵⁸ This retreat reflects multiple pressures: the Supreme Court's 2023 decision striking down affirmative action in college admissions created legal uncertainty about corporate diversity programs; Republican state attorneys general launched investigations into corporate DEI practices; and shareholder activists filed lawsuits challenging diversity initiatives as discriminatory against white employees.⁵⁹

The speed and scope of this reversal illustrate the central dynamic of corporate rulemaking. When social pressure mounted following Black Lives Matter protests, corporations rapidly implemented uniform policies affecting millions of workers because the reputational and business costs of inaction seemed prohibitive.⁶⁰ When the political and legal environment shifted, these same standardized policies were reversed with equal efficiency across entire corporate hierarchies. Individual employees had little voice in either direction. They experienced both the implementation and elimination of diversity programs as unilateral management decisions that reshaped their workplace environments without meaningful consultation.⁶¹

This cycle reveals how corporate rulemaking functions as a mechanism of social ordering that extends far beyond traditional commercial relationships. The adoption and subsequent abandonment of DEI policies affected not just hiring and promotion practices, but broader questions about workplace culture, speech norms, and the role of corporations in addressing societal inequities. The uniformity required for operational efficiency means that when corporations make these choices, they do so at scale, either creating or dismantling frameworks that govern the daily experiences of millions of workers and shape social expectations about corporate responsibility.

3.2.2.4 Rules of Disruption

Disruption has become a byword of business innovation in the last twenty years, as new technology has spurred development of new ideas. The gig economy has revolutionized services and upended employment relationships; short-term rentals have reshaped urban areas; smartphones have transformed daily lives; and social media has established new means of communication and news distribution. But when redesigning entire spheres of human activity, corporations come face to face with the societal implications of their success.⁶²

Companies typically adopt rules, norms, or practices that seek to bring order and predictability to their business activities, at once confirming the hierarchy they have established and governing the new terrain unfolding under them. Airbnb laid down rules on house parties when some became raucous and even resulted in deaths, and opened up rentals to refugees when housing them became problematic.⁶³ Gig economy companies like Uber sought to maintain the independent contractor status they claimed was preferred by workers while expanding the offerings available to them. These efforts represent legal ordering because they have general applicability and are impossible for individual workers to negotiate around.⁶⁴

Some companies may choose to stay largely inactive, deflecting the need for further intervention because they fear this may reduce the appeal of their product. Facebook's reticence in instituting frameworks for regulating posts illustrates this stance.⁶⁵ But this is no less rulemaking than active intervention: it illustrates the moral priorities of the company and represents its view on the norms that should govern the new interactions it has facilitated.

The legal ordering stakes are much higher for disruptors than for companies in established industries.⁶⁶ With the disruption fresh in the public's mind, the normative features of the company's conduct weigh significantly in society's assessment of the company's success. The new world that the company is ushering in has clear winners and losers, who are likely to be vocal about their grievances. Policymakers are lobbied on all sides and try to better reflect public attitude. Lacking precedent for addressing newly emerging challenges, companies and their stakeholders have no other recourse but going to first principles. But rekindling fundamental questions about social welfare benefits or the boundaries of free speech is bound to involve deeply held normative beliefs. When addressing these questions, the company is setting in motion practices that will engulf all its users and leave them with few options.⁶⁷

3.2.3 Governance Implications of Corporate Rulemaking

Corporate rulemaking instances unavoidably invoke the intense social impact challenge described in [Chapter 1](#), because they necessarily create highly visible standards that affect multiple constituencies simultaneously. Unlike individual business decisions that may fly under the radar, frameworks of general and uniform applicability, whether governing workplace conduct, supply chain standards, or customer data handling, become focal points for stakeholder attention precisely because of their broad reach and lasting footprint.⁶⁸ The stakes are inherently high: these rules will shape the experiences of employees, customers, suppliers, and communities, often for years to come, making them natural targets for organized advocacy and public scrutiny.

Moreover, corporate rulemaking directly confronts the challenge of heterogeneous shareholder preferences.⁶⁹ When companies establish frameworks that touch on social issues, from diversity policies to environmental standards to content moderation rules, they must navigate shareholders with fundamentally different

values and expectations about corporate strategy.⁷⁰ Some investors may prioritize aggressive social action while others resist what they view as mission creep, leaving management to make choices that will inevitably disappoint some constituencies.

Faced with these dynamics, corporations cannot simply pronounce rules through traditional internal processes and expect them to avoid contestation. Instead, they are better off turning to governance mechanisms that can help demonstrate the reasonableness and legitimacy of their choices while managing the inevitable tensions between competing stakeholder demands and shareholder preferences. [Section 3.2.3](#) explores the primary tools that companies are using to resolve these tensions and explains the distinctive contribution that each set of governance mechanisms makes to enhancing both the quality and acceptance of corporate rulemaking.

3.2.3.1 External Input and Validation Mechanisms

Companies increasingly recognize that effective rulemaking requires looking beyond internal expertise and business considerations to incorporate external perspectives and validation. One common approach involves engaging outside experts, such as academics, NGOs, or specialized consultants, who can provide independent perspectives and help evaluate the broader societal implications of particular frameworks.⁷¹ This strategy addresses the inherent limitations of internal decision-making by introducing external validation and subject-matter expertise that companies may lack internally. For instance, when developing environmental standards or human rights policies, companies frequently collaborate with environmental scientists or human rights organizations to ensure their frameworks reflect current best practices and address key concerns.

The value of expert engagement extends beyond mere technical accuracy. External experts can identify blind spots in corporate thinking, challenge assumptions about stakeholder impacts, and provide credibility that enhances the perceived legitimacy of resulting frameworks. However, this approach also presents challenges in expert selection, ensuring genuine independence, and managing potential conflicts between expert recommendations and business imperatives.⁷²

Equally important are extensive stakeholder consultation processes, where companies actively seek input from affected groups, organize listening sessions, and create feedback channels that allow stakeholders to shape framework development. These processes recognize that those who will be subject to corporate rules often possess crucial knowledge about potential impacts and implementation challenges that internal teams cannot anticipate.⁷³ For instance, Starbucks worked with Conservation International to create its Coffee and Farmer Equity (C.A.F.E.) Practices, a comprehensive set of guidelines designed to ensure ethical sourcing and sustainability in its coffee supply chain.⁷⁴ This collaboration enabled Starbucks to incorporate stakeholder perspectives and address key environmental and social concerns while developing frameworks that were both practically implementable and responsive to farmer needs.⁷⁵

A third external validation strategy involves working with other companies toward industry-wide or even market-wide standard setting or adopting standards set by well-respected independent bodies. This approach offers several advantages: it spreads the costs of framework development across multiple actors, reduces competitive disadvantages that might result from unilateral action, and enhances the credibility and reach of resulting standards.⁷⁶ For example, numerous global corporations, including General Motors, Shell, and BP, collaborated with the World Resources Institute and the World Business Council for Sustainable Development to develop the Greenhouse Gas Protocol, which has become the gold standard for corporate carbon accounting. This collaborative effort established a consistent and reliable methodology for measuring, managing, and reporting greenhouse gas emissions across industries worldwide.⁷⁷

3.2.3.2 Values-Based Internal Process and Implementation

Beyond external input, companies must develop robust internal processes for framework development and implementation. Some companies have found success in articulating foundational principles or value statements that serve as reference points for specific frameworks, providing consistency and clarity about the company's broader commitments and priorities.⁷⁸ This approach helps ensure that individual rules connect to broader corporate purposes while providing stakeholders with predictable criteria for evaluating corporate decisions. These foundational frameworks can take various forms, from formal corporate mission statements to detailed ethical codes to comprehensive sustainability commitments.⁷⁹

Equally important are procedural safeguards designed to enhance the quality and accountability of rulemaking processes. These include requirements for impact assessments before adopting new frameworks, mandatory comment periods that allow stakeholder input, regular review cycles that ensure frameworks remain current and effective, and clear documentation of decision-making rationales.⁸⁰ Transparency mechanisms play a crucial role in enabling external accountability. Companies increasingly publish detailed accounts of their framework development processes, including summaries of stakeholder input received, explanations of how feedback was incorporated or why it was not adopted, and clear statements of the business and social considerations that influenced final decisions.⁸¹

Effective governance must also address implementation and enforcement challenges. This includes developing clear metrics for evaluating framework effectiveness, establishing mechanisms for addressing stakeholder complaints or concerns about implementation, and creating processes for updating frameworks based on real world experience. Companies must also consider how to ensure consistent implementation across complex organizational structures, particularly for multinational corporations operating in diverse legal and cultural contexts. This often requires developing training programs, establishing clear accountability structures,

and creating feedback loops that can identify implementation problems before they become significant issues.

3.2.3.3 Building Legitimacy through Deliberative Governance

Perhaps the most challenging aspect of corporate rulemaking governance involves managing the inevitable tradeoffs between business considerations and stakeholder interests.⁸² No governance process can eliminate these tensions, but effective processes can help ensure that tradeoffs are made deliberately and transparently, with clear justification for the choices involved. This requires companies to develop frameworks for evaluating competing claims, establishing clear criteria for weighing different types of impacts, and creating mechanisms for explaining and justifying difficult decisions to stakeholders. Success often depends on early identification of potential conflicts and proactive engagement with affected parties to explore creative solutions that can address multiple concerns simultaneously.⁸³

Companies will need to present their choices as reasonable and justified, in an effort to convince audiences to adopt their products, work for them, or create new frameworks that take them into account. The governance mechanisms described earlier, while operating within the constraints of the corporate hierarchy, help ensure that framework development becomes a more deliberative and inclusive process, even as final authority remains with management. The goal is not just to make better decisions, but to create frameworks whose legitimacy stems from the careful consideration and incorporation of diverse perspectives and expertise.

Companies that invest in robust governance processes for their rulemaking activities often find that the resulting frameworks are more durable, face less resistance from stakeholders, and prove more effective in achieving their intended goals.⁸⁴ The additional time and resources required for comprehensive governance processes can be justified by reduced implementation costs, enhanced stakeholder relationships, and improved long-term business outcomes.⁸⁵ The challenge for companies is developing governance approaches that are appropriately tailored to the specific characteristics of their rulemaking activities – considering factors such as stakeholder diversity, potential impact magnitude, implementation complexity, and competitive dynamics – while maintaining the efficiency and decision-making authority necessary for effective business operations.

3.3 EXECUTIVE ACTION AND STRATEGIC DECISION-MAKING

While corporate rulemaking establishes general principles and creates systematic frameworks for ongoing operations, executive action encompasses the full range of specific decisions that address particular situations or launch discrete initiatives. These decisions may implement existing rulemaking commitments, but they also include countless business choices made independently of any established

framework. Thus, executive decisions represent the critical process of translating management's vision into concrete reality. This creates a fundamental dynamic: rulemaking establishes the architecture of corporate values and processes, while executive decisions represent the moment-by-moment construction work that can either reinforce or undermine that foundation.

This translation from principle to practice presents corporations with their most fundamental credibility challenge: how to convince shareholders and stakeholders that their specific actions genuinely advance their stated goals rather than merely paying lip service to them. The proliferation of terms such as “greenwashing,” “DEI washing,” “social washing,” and “sustainability washing” reflects not just growing skepticism about corporate sincerity, but a sophisticated understanding among stakeholders that the real test of corporate commitments lies not in their articulation but in their implementation through countless individual decisions.⁸⁶

The credibility gap has become a defining challenge of modern corporate governance. Companies routinely make sweeping commitments through their rule-making processes, adopting policies, publishing frameworks, and establishing principles on environmental sustainability, social justice, traditional values, or stakeholder capitalism. But stakeholders across the political spectrum increasingly judge these commitments not by their comprehensiveness or eloquence, but by whether they translate into consistent executive decisions that sometimes require short-term sacrifices for stated principles.⁸⁷ When a company commits to carbon neutrality but then makes acquisition decisions based solely on financial returns, or when a company pledges to avoid “woke politics” but then makes decisions that appear to contradict conservative values, or when a company establishes diversity principles but consistently promotes from networks that exclude underrepresented groups, the gap between stated commitments and executive action becomes a credibility crisis.⁸⁸

The governance challenge lies in creating systems that ensure executive decisions consistently implement and advance the commitments established through rule-making processes, while building stakeholder confidence that these decisions reflect authentic commitment rather than strategic positioning. This requires sophisticated mechanisms that go beyond traditional corporate governance focused on legal compliance and financial performance, extending into domains of credibility, authenticity, and consistent implementation of stated values through concrete action.

3.3.1 *The Nature and Scope of Executive Corporate Choices*

Executive choices represent decisions to move in specific ways on individualized matters, addressing particular situations through discrete actions rather than systematic policy development. These choices encompass the vast majority of daily corporate decision-making, from routine operational matters to transformative strategic initiatives that reshape entire business models.

The scope of executive choices spans an enormous range of corporate activity. At the operational level, executives make countless decisions about resource allocation, personnel matters, supplier selection, pricing strategies, and facility management. At the strategic level, they decide whether to launch new products, enter new markets, pursue acquisitions, divest business units, or respond to competitive threats. Crisis management requires executives to make rapid decisions about addressing emergencies, controversies, or unexpected market disruptions. What unites all these choices is their specificity to particular circumstances and their implementation of broader corporate direction through concrete action.

Most executive choices proceed without raising significant social concerns or attracting stakeholder attention beyond their immediate business implications. Decisions about inventory management, routine hiring, technology upgrades, or standard supplier agreements typically focus on operational efficiency, cost management, and performance optimization. These choices form the backbone of corporate operations, requiring managerial judgment about business trade-offs without invoking broader questions about corporate values or social responsibility.

Strategic executive choices are more likely to raise issues of social importance, both because of their larger impact on corporate direction and because they often bring underlying social issues to the surface. When companies launch new products, these choices inevitably raise questions about their social and environmental implications. When they enter new markets, particularly in different countries or regions, they confront issues about labor standards, environmental regulations, and community impact. When they pursue acquisitions, they must address how combining operations will affect various stakeholder groups and whether the combined entity will advance or undermine their stated commitments.

Society's attention sometimes turns to issues where corporations play major roles, creating moments when executive choices take on heightened social significance. Catalytic events such as George Floyd's murder in 2020 or awareness campaigns about environmental issues can transform routine corporate decisions into tests of corporate commitment to social values.⁸⁹ Similarly, backlash against ESG initiatives or concerns about corporate "woke politics" can create pressure for executives to demonstrate through specific actions whether their companies prioritize business fundamentals over political positioning.⁹⁰ These moments create pressure for executives to demonstrate through specific actions whether their companies' stated principles represent authentic commitments or merely strategic positioning.⁹¹

The corporate response to Russia's invasion of Ukraine illustrates how major events can transform executive choices into statements about corporate values. Many companies decided to withdraw from Russia, representing momentous strategic decisions with immediate implications for employees, business partners, and shareholders.⁹² These were discrete, one-time choices that implemented broader commitments to human rights or international law, requiring companies to sacrifice short-term financial interests for stated principles.⁹³

In contrast, when companies make statements supporting causes such as gender equality or Black Lives Matter, these proclamations represent attempts to establish ongoing commitments that require implementation through multiple executive decisions over time. Similarly, when companies pledge to avoid political activism or commit to religious freedom principles, these statements demand sustained implementation through operational choices. While withdrawal from Russia represents a discrete, easily observable strategic break, broader social commitments demand sustained substantiation through countless operational choices: hiring decisions, promotion practices, supplier selection, resource allocation, and community engagement initiatives. Each subsequent executive choice becomes an opportunity to either validate or contradict the original commitment, creating ongoing challenges for maintaining credibility and authentic implementation.⁹⁴

3.3.2 *Identifying and Assessing Social Significance: The Challenge of Recognition and Calibration*

Determining which executive choices carry social significance requires sophisticated judgment and contextual awareness that goes beyond the formal characteristics of the decision itself. Even seemingly routine executive choices can acquire social significance when they occur in contexts that reflect broader societal tensions, appear emblematic of corporate culture, or attract attention as examples of systematic corporate behavior. The dismissal of a senior officer following allegations of misconduct, and particularly the terms of their departure, represents an executive choice that might have profound social implications despite appearing routine from a legal perspective. Similarly, decisions about facility locations, supplier relationships, or employment practices can become focal points for stakeholder concerns about corporate values and priorities.⁹⁵

The challenge lies in recognizing that social significance exists along a sliding scale rather than as a binary classification. Some executive choices clearly invoke major social concerns, such as decisions about withdrawing from entire markets for political reasons, or high-profile hiring decisions that send clear signals about corporate values. Other choices operate in a gray area where social implications may be present but less pronounced, such as supplier selection decisions that involve modest differences in labor practices, or facility location choices that have some but not decisive environmental implications. Still other decisions appear entirely routine until external events or stakeholder attention transforms them into social issues.⁹⁶

This sliding scale of social significance has important implications for governance requirements. Executive choices with clear, major social implications demand the most rigorous disclosure, accuracy, and oversight mechanisms. Decisions in the gray area may require more limited governance responses, perhaps enhanced internal consideration of social factors and basic stakeholder communication, but not the full apparatus of third-party verification and comprehensive public reporting.

Routine decisions that seem unlikely to raise social concerns may require little special governance attention until circumstances change their significance. The key governance challenge lies in developing systems that can identify where particular decisions fall on this spectrum and calibrate appropriate responses accordingly.⁹⁷

3.3.3 *Governance Implications of Executive Actions: Disclosure and Accuracy*

The governance of executive decisions in the context of corporate social commitments requires fundamentally different approaches than traditional corporate governance focused primarily on financial performance and legal compliance. While conventional governance mechanisms aim to maximize shareholder value and prevent legal violations, governing executive decisions that implement social commitments demands systems designed to build and maintain stakeholder trust through demonstrated consistency between stated principles and concrete actions.⁹⁸

The central governance challenge stems from the credibility gap that opens between corporate commitments and their implementation. Unlike financial metrics, which provide relatively objective measures of performance, social commitments often involve complex, subjective judgments about whether specific decisions authentically advance stated principles or merely provide cover for business as usual approaches.⁹⁹ This complexity creates opportunities for companies to engage in various forms of “washing” – greenwashing, social washing, or diversity washing – where executive decisions are framed as advancing social commitments while actually serving primarily traditional business objectives.¹⁰⁰

3.3.3.1 Disclosure as the Foundation of Credible Implementation

Robust disclosure mechanisms represent the primary governance tool for addressing the credibility gap in executive decision-making.¹⁰¹ However, effective disclosure for social commitments requires approaches that go far beyond traditional corporate reporting. Rather than simply announcing decisions, companies must develop disclosure systems that explain decision-making rationales, demonstrate alignment with stated principles, acknowledge trade-offs and conflicts, and provide sufficient detail for shareholders and stakeholders to evaluate authentic commitment versus strategic positioning.¹⁰²

The most credible companies have learned that shareholder trust depends heavily on transparency about decision-making processes and outcomes. Successful companies typically explain not only what decisions they made, but why they made them and how these decisions advance specific stated commitments. When companies with human rights commitments make supplier selection decisions, they often provide detailed disclosure explaining how labor practice considerations factored

into the evaluation process alongside cost and quality criteria.¹⁰³ Similarly, companies with environmental commitments making supply chain decisions frequently detail how environmental factors influenced selection criteria and final choices. This level of specificity allows stakeholders to evaluate whether social commitments actually influenced decision-making or merely provided post hoc justification for choices made on other grounds.¹⁰⁴

Companies that have successfully managed shareholder expectations around social commitments typically acknowledge that implementing such commitments often requires difficult trade-offs with other business objectives.¹⁰⁵ Rather than claiming perfect alignment between social goals and business interests, credible companies acknowledge these tensions and explain how they balanced competing priorities. When companies choose suppliers with better labor practices despite higher costs, or reject profitable opportunities that conflict with environmental commitments, transparent discussion of these trade-offs tends to build stakeholder confidence by demonstrating that social commitments influenced decisions even when doing so imposed business costs. This approach helps distinguish companies that integrate social considerations into their core decision-making processes from those that consider them only when they align with other business interests.

Shareholders increasingly demand insight into corporate decision-making processes, not just outcomes, leading companies to provide disclosure that illuminates the internal mechanics of executive choice. Companies responding to this pressure typically explain how they structure decision-making to ensure consideration of social commitments, what internal controls prevent backsliding, and how accountability systems ensure consistent implementation across different business units and decisions.¹⁰⁶ This process-oriented disclosure addresses stakeholder concerns that social commitments might influence only high-profile decisions while leaving routine choices to operate under traditional business logic. Companies that reveal the systematic integration of social considerations into decision-making processes can demonstrate that their commitments represent more than aspirational statements or public relations strategies.¹⁰⁷

Companies have also learned to connect individual executive decisions to broader implementation strategies, helping stakeholders understand how specific choices fit into larger patterns of commitment implementation. Effective corporate communication typically explains how current decisions advance longer-term goals and what future decisions will be required to maintain progress. This forward-looking dimension of disclosure serves multiple purposes for companies: it demonstrates strategic thinking about commitment implementation, creates accountability for future consistency, and provides shareholders with frameworks for evaluating future decisions. When companies can articulate clear connections between individual executive choices and broader commitment implementation strategies, they provide shareholders with tools for holding management accountable for sustained consistency while building confidence in their authentic long-term commitment.

3.3.3.2 Accurate Measurement as the Cornerstone of Shareholder Demands

While disclosure provides the mechanism for stakeholder evaluation through voluntary sharing of information, accuracy represents the substantive requirement for building trust. Without accurate measurement and correct characterization, even the most comprehensive disclosure becomes merely an exercise in corporate communication that fails to serve stakeholders' evaluation needs. Similarly, extensive but vague disclosures that avoid specific commitments or measurable outcomes provide little value to stakeholders seeking to evaluate authentic implementation of social commitments. Companies that have successfully maintained credibility with shareholders who increasingly scrutinize corporate claims about social commitment implementation have learned that accuracy in this context involves several dimensions beyond simple factual correctness. The voluntary act of disclosure means little if what companies choose to share is mismeasured, mischaracterized, or incomplete in ways that mislead rather than inform stakeholder judgment.¹⁰⁸

The most successful companies have found that honestly characterizing the role social commitments played in their decision-making processes builds long-term shareholder confidence. When business considerations primarily drive decisions that happen to align with social commitments, credible companies acknowledge this reality rather than overstating the commitment motivation.¹⁰⁹ Conversely, when companies make decisions primarily to advance social goals despite business costs, accurate disclosure that highlights this authentic prioritization tends to reassure shareholders that management is making deliberate trade-offs rather than stumbling into costly decisions. This honest characterization helps shareholders evaluate whether management is strategically integrating social considerations or merely engaging in opportunistic positioning.¹¹⁰

Companies that maintain stakeholder trust over time typically provide realistic assessments of whether executive decisions actually advance stated commitments and by how much. Rather than overstating the social benefits of their decisions, successful companies tend to provide measured evaluations that give appropriate credit for genuine progress toward stated goals while avoiding exaggerated claims that might undermine future credibility. This balanced approach helps shareholders distinguish between companies making genuine progress on social commitments and those engaging in superficial gestures designed primarily for public relations value.¹¹¹

Particularly effective companies have learned that acknowledging limitations and setbacks, rather than undermining their credibility, actually strengthens stakeholder confidence in their long-term commitment authenticity.¹¹² Companies that honestly discuss where executive decisions fall short of stated commitments, where competing priorities prevented ideal implementation, and where unintended consequences undermined intended social benefits tend to build stronger relationships with shareholders who prefer transparency about challenges over misleading optimism.

This transparency signals to investors that management has realistic understanding of implementation difficulties and is likely to address problems systematically rather than ignore them.¹¹³

Finally, companies responding to shareholder demands for accountability have increasingly turned to quantitative precision wherever possible, recognizing that social commitments often resist easy quantification but that specific, measurable information enhances credibility. Rather than relying on vague claims about “improving diversity” or “reducing environmental impact,” companies that maintain strong stakeholder relationships typically provide concrete data about outcomes and progress toward specific targets. This quantitative approach allows shareholders to track progress over time and compare performance across companies, meeting investor demands for measurable accountability in areas that were traditionally considered too subjective for systematic evaluation.¹¹⁴

3.3.3.3 Additional Credibility-Enhancing Mechanisms

While disclosure and accuracy represent the primary tools for governing executive decisions that implement social commitments, companies responding to shareholder skepticism have developed several supporting mechanisms that enhance the credibility of their claims. These mechanisms address the fundamental challenge that shareholders face in evaluating corporate self-reporting about social commitment implementation.

Companies seeking to address shareholder concerns about the reliability of their social commitment claims have increasingly turned to third-party verification as a crucial external validation mechanism. Whether through environmental auditors, human rights monitors, or specialized certification bodies, independent verification helps companies address the inherent credibility challenges that corporate self-reporting faces with investors who remain skeptical about management’s incentives to overstate progress. However, shareholders have learned that the effectiveness of third-party assurance depends heavily on the independence, expertise, and methodological rigor of verifying organizations, leading to more sophisticated evaluation of verification quality rather than mere acceptance of any external validation.¹¹⁵

Building on external verification, companies have found that creating systematic mechanisms for stakeholder input and feedback enhances their credibility with shareholders who understand that stakeholder satisfaction often correlates with long-term business sustainability. These mechanisms might include regular stakeholder advisory panels, community consultation processes for decisions affecting local communities, or employee feedback systems for decisions affecting workplace culture and diversity. Shareholders increasingly view such stakeholder engagement not as corporate social responsibility theater, but as risk management tools that help companies identify potential problems before they become costly crises or regulatory interventions.¹¹⁶

Companies have also discovered that board-level oversight serves a dual function of ensuring management accountability while providing shareholders with confidence that social commitment implementation receives appropriate strategic attention. Effective boards regularly review executive decisions for consistency with stated commitments, hold management accountable for gaps between principles and practices, and ensure that social commitment implementation receives appropriate priority in strategic planning and resource allocation. This board involvement signals to shareholders that social commitments represent genuine strategic priorities rather than management pet projects or public relations initiatives that might divert resources from core business objectives.¹¹⁷

Finally, some companies have demonstrated their commitment to shareholders by linking executive compensation to successful implementation of social commitments through consistent decision-making. Such compensation linkages provide shareholders with assurance that management has personal financial incentives to follow through on stated commitments rather than abandoning them when implementation becomes costly or difficult. However, companies and their shareholders have learned that such linkages must be carefully designed to avoid creating perverse incentives that encourage form over substance, leading to increasingly sophisticated approaches that focus on meaningful outcomes rather than easily manipulated metrics.

3.3.3.4 The Ongoing Challenge of Authentic Implementation

The governance of executive decisions that implement corporate social commitments represents an evolving challenge that requires continuous adaptation and refinement. As stakeholder expectations continue to rise and scrutiny of corporate behavior intensifies, companies must develop increasingly sophisticated approaches to demonstrating authentic commitment through consistent decision-making. The companies that succeed will be those that embrace disclosure and accuracy not as compliance burdens but as opportunities to build lasting stakeholder trust through demonstrated integrity in translating principles into practice.¹¹⁸

The effectiveness of these governance mechanisms depends ultimately on corporate leadership's genuine commitment to the principles they have articulated. No governance system can substitute for authentic commitment, but this reality makes the governance challenge more urgent, not less important. In an era where virtually all major corporations have learned to reproduce the same compelling message, the critical question for shareholders and stakeholders has become how to distinguish between companies that genuinely embrace their stated principles and those that treat them primarily as strategic positioning tools.¹¹⁹

The answer lies precisely in examining the concrete implementation of executive decisions and the rigor of the processes that govern them. While companies can craft similar-sounding policy statements and rulemaking frameworks, their executive decisions reveal authentic priorities through resource allocation, trade-off management,

and consistent implementation over time. Companies that have embedded social commitments into their actual decision-making processes demonstrate measurably different patterns of executive choice than those that treat such commitments as external constraints to be managed or marketing messages to be projected.¹²⁰

Moreover, the decision-making processes themselves serve as predictive indicators of future performance. Companies that have developed systematic approaches to integrating social considerations into executive decisions, created robust accountability mechanisms, and established transparent measurement systems are more likely to maintain consistent implementation over time than those relying on ad hoc responses or management rhetoric. In an era of increasing stakeholder activism and public scrutiny, this distinction between authentic implementation and strategic positioning will likely determine not just corporate reputation but fundamental business sustainability and long-term financial performance.¹²¹

3.4 THE MONITORING FUNCTION IN CORPORATE GOVERNANCE

While corporate monitoring is traditionally associated with ensuring employee compliance with legal mandates, modern corporations have developed complex internal bureaucracies that serve a far broader set of principles. The focus is on tracking progress and alignment across a wide spectrum of commitments, from environmental goals to workplace equity, and from ethical sourcing to community impact. This expanded monitoring function represents the crucial final step in corporate governance: ensuring that principles articulated in corporate rulemaking frameworks actually translate into practice through executive decisions, and that these decisions are faithfully implemented throughout the organization.¹²²

3.4.1 *The Expansive Reach of Modern Compliance*

What distinguishes contemporary corporate monitoring from its traditional counterpart is not merely its sophistication, but its remarkable scope. The modern compliance apparatus has evolved far beyond the conventional model of corporations as reluctant enforcers of externally imposed legal mandates. Today's monitoring systems encompass an extraordinary range of commitments that extend well into territories once considered purely strategic or aspirational. Companies now monitor adherence to self-imposed environmental targets, diversity and inclusion goals, supply chain ethical standards, data privacy commitments, community engagement promises, and broader stakeholder value propositions. Many of these goals do not result from direct legal mandates but emerge from business strategy, competitive positioning, investor expectations, or organizational values.¹²³

This expansion reflects a fundamental shift in how corporations conceive of their monitoring obligations. Rather than simply transposing regulatory requirements

into everyday corporate practice, companies now deploy sophisticated tracking mechanisms to oversee compliance with voluntary commitments that may be more stringent than legal minimums. A technology company might monitor not just its legal obligations under privacy law, but also its public commitments to data minimization that exceed regulatory requirements.¹²⁴ A manufacturing firm might track not only mandated emissions reporting but also progress toward self-imposed carbon neutrality goals.¹²⁵ A financial services company might monitor not just regulatory compliance with fair lending laws, but also voluntary commitments to community investment and financial inclusion.

This internal bureaucracy performs essential supervisory functions that extend well beyond conventional compliance, operating through an extraordinary level of managerial discretion exercised within sophisticated multilayered supervisory systems. At the operational level, companies deploy automated monitoring tools that track everything from employees' digital activities to factory emissions in real time. These systems feed into middle-tier review processes, where specialized compliance teams analyze patterns and flag potential issues. At the highest level, board committees and executive teams receive regular reports and metrics dashboards, creating a continuous flow of monitoring information up and down the corporate hierarchy.¹²⁶

This tiered structure is further reinforced by an intricate web of internal and external oversight mechanisms. Internally, companies typically maintain multiple channels for monitoring and reporting: routine supervisory reviews, anonymous reporting hotlines, regular audits by internal teams, and specialized oversight committees. These internal systems are complemented by external verification through third-party auditors, certification bodies, and independent monitors. Many companies have also established sophisticated appeals processes that allow stakeholders to challenge decisions, often involving multiple levels of review and independent arbitrators.

3.4.2 *Managerial Discretion in Designing Comprehensive Monitoring*

Within this detailed framework, managers must make countless decisions about how to structure monitoring processes, what metrics to track, and how to respond to violations. They must determine how different monitoring systems interact, when to escalate issues between tiers, and how to balance the need for oversight with operational efficiency. These choices, which include everything from the frequency of audits to the composition of review panels, shape the effectiveness and impact of corporate commitments, both those mandated by law and those voluntarily undertaken.¹²⁷

The scope of managerial discretion becomes apparent when examining specific monitoring challenges across key areas. In environmental monitoring, managers must decide how to account for carbon emissions: which calculation methodologies to use, what scope of emissions to include, and how to verify supplier data.

They face similar choices in water usage monitoring, particularly regarding supply chain impacts and regional variations in water stress.¹²⁸ In diversity initiatives, managers must determine how to measure progress across multiple dimensions: from basic metrics such as hiring and promotion rates to more nuanced measures of inclusion and intersectionality. Recent political developments, including executive actions and legal challenges, have further complicated these decisions by constraining the scope of permissible DEI programs in some jurisdictions. For ethical sourcing, they must choose how deep into the supply chain they should monitor, whether to rely on industry certifications or develop proprietary standards, and how to most effectively address violations. In privacy monitoring, they must balance thorough oversight of data handling with employee productivity and determine appropriate escalation protocols.¹²⁹

These choices are rarely straightforward and often involve balancing competing priorities with practical constraints. Each decision shapes not only how monitoring is implemented but also what behaviors and outcomes the organization prioritizes. Perhaps most crucially, managers must design the processes for handling violations and appeals. They must decide who has the authority to investigate complaints, what evidence standards to apply, and how to protect both the accusers and the accused. They must determine appropriate sanctions for different types of violations and establish procedures for appeal and review. These decisions require careful consideration of fairness, effectiveness, and organizational culture.

3.4.3 *How Companies Adjust Monitoring without Changing Policy Goals: From #MeToo to ESG Backlash*

The examples later illuminate the distinction between corporate rulemaking and monitoring. While rulemaking establishes or changes a company's fundamental values and principles, monitoring purports to implement those values more rigorously, comprehensively, or from multiple perspectives without altering them. In the cases that follow, companies maintained their stated commitments to political neutrality, balanced discourse, workplace diversity, and other values while developing more sophisticated systems to track, measure, and demonstrate adherence to those principles. Yet the very significance of these monitoring changes, with their capacity to reshape corporate behavior, employee experiences, and stakeholder relationships, naturally raises questions about whether such adjustments truly leave underlying values unchanged. These doubts, however, rather than undermining the monitoring function, actually underscore its fundamental importance in corporate governance. The power of monitoring systems to influence organizational culture, decision-making processes, and practical outcomes reveals why the design and implementation of these systems represents such a critical exercise of managerial discretion, one that can determine whether corporate principles remain aspirational statements or become operational realities.

The transformation of corporate monitoring in response to the #MeToo movement exemplifies how the modern compliance apparatus expands beyond legal mandates to encompass broader organizational values and cultural commitments. Prohibitions against sexual harassment in the workplace had been in place since at least the 1970s, and the legal framework appeared well established before 2017. With no new legislation required, #MeToo prompted companies to voluntarily expand their monitoring apparatus in ways that carried out the law's mission more effectively. Some companies created new reporting channels, offering immediacy through hotlines and guaranteeing anonymity.¹³⁰ Others provided enhanced support mechanisms, establishing rehabilitation groups and offering access to outside assistance for victims. Many companies developed efforts to strengthen bonds and relationships among women in the workplace through mentorship programs, regular caucuses and town hall meetings, and diversity groups. Most publicly traded companies increased the number of female directors, often under significant pressure from investors, recognizing that diverse governance structures might enhance monitoring effectiveness.¹³¹

The #MeToo response also fundamentally altered the calculation for individual managers responsible for sanctioning violators. Before 2017, typical resolution of sexual harassment allegations against high-profile executives involved allowing them to leave the company while maintaining stock options and golden parachutes. After 2017, this exit option became much harder to sustain, as managers recognized that appearing to protect abusers would undermine both organizational culture and external legitimacy.¹³²

These reforms illustrate the power and reach of voluntary compliance expansion. While entirely voluntary and arising from private initiative rather than legal mandate or fiduciary duty, these changes substantially strengthened enforcement frameworks and demonstrated the authority that corporations exercise in defining and implementing standards that exceed legal minimums. The cultural motivations driving #MeToo were channeled into institutional changes and individual choices that managers make, underscoring their role as key actors in expanding the scope and intensity of corporate monitoring beyond traditional legal compliance.¹³³

The responsiveness of corporate monitoring systems to cultural and political pressures extends beyond progressive movements like #MeToo to encompass conservative values and priorities as well. Recent years have witnessed significant corporate adjustments in response to conservative criticism and organized pressure campaigns, demonstrating that the expansion and contraction of voluntary monitoring systems reflects broader stakeholder dynamics rather than a unidirectional progressive trend.¹³⁴

The financial sector's response to "debanking" allegations illustrates how companies adjust their monitoring systems to better serve existing policy goals in response to conservative pressure. Following sustained criticism from Republican lawmakers and activists who alleged that banks were systematically discriminating against conservative customers and businesses, major financial institutions implemented

enhanced monitoring procedures designed to ensure political neutrality in their business decisions.¹³⁵ Bank of America, JPMorgan Chase, and Wells Fargo all established new internal review mechanisms for account closures, creating multitiered oversight processes that require detailed documentation of business justifications for ending customer relationships. These banks developed specialized training programs for relationship managers and created escalation procedures that flag potential politically sensitive decisions for senior review. Some institutions established dedicated ombudsman offices specifically to handle complaints about alleged political bias in banking decisions, complete with independent review processes and appeals mechanisms. Crucially, these changes did not represent a shift in policy goals. Banks maintained their commitment to serving all qualified customers regardless of political affiliation, but implemented more robust monitoring systems to demonstrate and ensure adherence to that principle.

Technology companies have similarly adjusted their content moderation monitoring systems in response to conservative criticism about alleged anti-conservative bias, while maintaining their stated commitment to balanced discourse. Following congressional hearings and public pressure campaigns alleging systematic suppression of conservative viewpoints, platforms such as Facebook (Meta) and Twitter (before its acquisition) implemented enhanced review processes specifically designed to identify potential bias in content moderation decisions. These companies created specialized oversight teams tasked with monitoring the treatment of political content across the ideological spectrum, developed new appeal mechanisms for users who believe their content was wrongly removed, and established external advisory boards that include conservative voices. Some platforms implemented “red team” exercises where conservative employees or consultants specifically review content moderation decisions for potential bias. These monitoring enhancements aimed to better serve the same underlying goal of fair and balanced content policies rather than changing those policies themselves.¹³⁶

Corporate hiring and workplace inclusion programs have also seen monitoring adjustments in response to conservative legal challenges, while maintaining commitments to diverse and inclusive workplaces. Following concerns about potential reverse discrimination and evolving legal interpretations of employment law, many companies have modified their diversity hiring monitoring systems to include enhanced documentation of candidate qualifications and selection criteria.¹³⁷ These companies developed new training programs for hiring managers that emphasize skills-based evaluation while maintaining diversity goals, created more detailed audit trails for hiring decisions, and established review mechanisms to ensure that diversity efforts do not inadvertently disadvantage any candidates based on protected characteristics.¹³⁸ Some organizations have expanded their diversity definitions to include veterans, first-generation college graduates, and employees from rural backgrounds, which may appeal more to conservative stakeholders while maintaining the fundamental commitment to workplace diversity.

3.4.4 *The Challenge of Performative Monitoring*

The expansion of corporate monitoring systems creates a fundamental tension between genuine oversight and performative compliance. Companies often respond to stakeholder pressure and regulatory expectations by layering on additional procedures, documentation requirements, and oversight roles without fundamentally altering decision-making processes or outcomes. This proliferation of monitoring apparatus can create an illusion of enhanced accountability while leaving underlying practices largely unchanged, a bureaucratic theater that generates extensive records but limited behavioral modification.

The distinction between effective and performative monitoring lies not in the formal structure of oversight systems, but in how companies design the practical mechanisms that translate monitoring into operational change. Organizations that successfully avoid performative compliance typically embed monitoring functions directly into business processes rather than treating them as parallel reporting exercises. They structure incentive systems that reward managers for substantive outcomes rather than procedural adherence, and they design escalation mechanisms that ensure identified problems trigger concrete remedial actions rather than additional documentation.¹³⁹

Critically, the effectiveness of these systems rarely emerges from management intuition or ad hoc decision-making. Instead, companies develop systematic approaches to structuring monitoring relationships, defining escalation thresholds, allocating investigative resources, and measuring outcomes.¹⁴⁰ The substantial discretion that corporations exercise in designing these systems, from determining which issues warrant board-level attention to establishing protocols for employee reporting, shapes whether monitoring functions serve as genuine accountability mechanisms or elaborate performances.¹⁴¹ The system Microsoft has put together to supervise responsible AI, described in [Chapter 8](#), illustrates how to translate genuine commitments into oversight.

Corporate monitoring is thus an essential governance function, operating alongside rulemaking and executive action as distinct but interconnected mechanisms of organizational control. While rulemaking establishes the formal parameters of acceptable behavior and executive action drives day-to-day operations, monitoring provides the feedback mechanisms that determine whether stated policies translate into actual practice. Designed under management discretion, monitoring represents an area where companies make consequential choices about institutional structure with limited external guidance or standardization. Without analyzing how organizations construct their monitoring systems, observers cannot fully comprehend how corporate rulemaking and executive action function in practice, since monitoring provides the institutional architecture that connects formal commitments to operational reality.

3.5 BEYOND PURPOSE: CORPORATE GOVERNANCE AS INSTITUTIONAL ARCHITECTURE

The framework developed in this chapter, distinguishing corporate rulemaking, executive action, and monitoring functions, represents more than an analytical exercise. It fundamentally reframes how we understand corporate power, institutional design, and the relationship between business organizations and broader social ordering. The implications span multiple domains, from practical governance design to theoretical questions about the origins of legal authority itself.

3.5.1 *Revealing the True Scope of Corporate Institutional Work*

The tripartite analysis exposes a critical blind spot in contemporary corporate law scholarship, which has remained fixated on the binary question of whether corporations should pursue shareholder profits, stakeholder interests or broader social welfare. This debate, while generating substantial academic literature, obscures the far more complex reality of what corporations actually do. By applying the rulemaking-executive-monitoring framework, we see that corporate activity extends well beyond either profit maximization or stakeholder capitalism into the fundamental work of social ordering itself.

Corporations do not simply choose between shareholder primacy and stakeholder interests; they actively construct the institutional frameworks through which millions of people organize their daily lives. When Apple establishes App Store policies, Amazon designs fulfillment protocols, or Uber creates platform rules, these companies are not merely pursuing business strategies, but are building the infrastructure of modern economic and social relationships. The tripartite framework reveals this institution-building work as distinct from both traditional profit-seeking and conventional social responsibility, operating through mechanisms that function as *de facto* governance systems affecting billions of individual interactions.

This recognition fundamentally alters how we evaluate corporate performance and accountability. Rather than asking whether companies adequately balance competing interests, we must examine whether they have developed appropriate institutional capabilities for the governance functions they inevitably perform. The framework shifts analytical focus from outcomes to processes, from goals to institutional design, recognizing that the path from corporate intentions to social results travels through complex organizational machinery that requires sophisticated stewardship to function effectively.

The emphasis on institutional processes rather than declared purposes addresses a persistent weakness in corporate law scholarship, which tends to assume that properly aligned incentives will automatically produce desired outcomes. The tripartite analysis demonstrates that this assumption overlooks the critical role of institutional design in translating intentions into results. Even corporations with

ideal incentive structures will fail to achieve their goals without governance systems appropriately calibrated to the distinct demands of rulemaking, executive action, and monitoring functions. The framework thus reveals corporate governance as fundamentally concerned with institutional architecture rather than simply incentive alignment.

3.5.2 *Practical Implications for Governance Design*

From a practical standpoint, the tripartite framework provides concrete guidance for designing more effective corporate governance systems. Current approaches to corporate oversight typically apply uniform mechanisms across fundamentally different types of corporate decisions, failing to recognize that rulemaking functions require different governance tools than executive actions, which in turn demand different approaches than monitoring activities.

The framework demonstrates why companies struggle when they apply identical governance approaches to fundamentally different corporate functions. Board committees designed for financial oversight may lack the stakeholder engagement capabilities necessary for effective rulemaking governance. Executive compensation systems optimized for strategic decision-making may poorly incentivize the independence and thoroughness required for effective monitoring. Disclosure requirements appropriate for routine business decisions may inadequately serve the transparency needs of executive actions implementing social commitments. The tripartite analysis provides a diagnostic tool for identifying these mismatches and developing more targeted governance solutions.

Legal practitioners and corporate advisors can use this framework to develop governance recommendations that match institutional tools to functional demands rather than relying on generic best practices. When companies face rulemaking decisions requiring broad stakeholder consultation, the framework suggests governance mechanisms focused on wide outreach and participatory processes. When they confront executive choices implementing stated commitments, it emphasizes disclosure and accuracy systems that build credibility through transparent communication. When they design monitoring systems for voluntary commitments, it highlights independence and investigative capacity as critical design elements.

The framework also helps explain why certain corporate governance reforms succeed while others fail. Reforms that align with functional requirements, such as stakeholder engagement processes for rulemaking decisions or independent oversight for monitoring functions, tend to enhance both corporate effectiveness and stakeholder satisfaction. On the other hand, reforms that ignore functional distinctions, such as applying uniform board oversight to all corporate decisions regardless of their different governance needs, often create compliance burdens without improving outcomes. [Chapter 9](#) includes a broader discussion of governance implications for boards and managers, as well as suggestions for reforming our understanding of fiduciary duties.

3.5.3 *Reconceptualizing Corporate Legal Authority*

The tripartite framework has profound implications for understanding the relationship between corporate power and legal authority. Traditional accounts position corporations as subjects of law, created by state authority and constrained by regulatory frameworks. This subject-centered view treats corporate rule-creation as either contracted-for private ordering or problematic regulatory capture that undermines legitimate governmental authority.

The institutional perspective developed here suggests a more complex relationship in which corporations function as significant sources of legal ordering rather than merely subjects of it. When corporations establish frameworks of general and uniform applicability that govern millions of interactions, they are not simply exercising delegated authority or filling regulatory gaps. They are participating in the fundamental work of legal system construction. This participation operates through different mechanisms than state lawmaking, relying on economic leverage and contractual relationships rather than sovereign authority, but it produces similar outcomes in terms of behavioral coordination and social ordering.

This recognition has important implications for how we understand regulatory relationships. Rather than viewing regulation as the beginning of the legal process and corporations as passive recipients of externally imposed rules, the framework suggests that corporate ordering often precedes and shapes subsequent regulatory responses. Companies first develop internal governance systems for addressing social challenges, establish industry practices through their operational choices, and create de facto standards through their rulemaking activities. Regulatory intervention typically represents attempts to formalize, modify, or supplement these existing corporate governance systems rather than creating entirely new legal frameworks from scratch.

Understanding this sequence has practical implications for both corporate strategy and regulatory design. Companies that recognize their role as legal system builders can develop more sophisticated approaches to governance challenges, anticipating regulatory attention and designing internal systems that can serve as foundations for broader industry standards. Regulators who understand corporate governance as institutional construction can design interventions that build on existing corporate capabilities rather than attempting to impose external systems that ignore established organizational capacities.

3.5.4 *Advancing Corporate Law beyond Normative Stalemates*

Perhaps most importantly, the tripartite framework offers a path beyond the normative debates that have paralyzed corporate law scholarship for decades. Disputes about whether corporations should focus on shareholders or stakeholders, whether they should address social issues or stick to business fundamentals, and whether they represent solutions to governmental failures or threats to democratic authority echo

through academic corridors without making much headway. These debates have prevented scholars from developing practical frameworks for evaluating and improving corporate governance where it actually operates.

The functional analysis developed here sidesteps these normative questions by focusing on institutional design rather than ultimate purposes. Rather than asking whether corporations should perform governance functions, it acknowledges that they inevitably do and examines how they can perform these functions more effectively. Rather than debating whether corporate social responsibility represents legitimate business activity, it analyzes what governance mechanisms best serve companies that have already decided to make social commitments. Rather than litigating whether corporate power threatens democratic values, it examines how corporate institutional design can better serve both business effectiveness and broader accountability.

This approach does not resolve normative questions about appropriate corporate purposes, but it makes progress on those questions by making them less necessary for developing better corporate governance. Companies with different values and priorities can benefit from the same institutional insights about matching governance mechanisms to functional demands. Progressive companies seeking to advance social justice goals and conservative companies focused on traditional business fundamentals both need effective governance systems for their rulemaking, executive, and monitoring functions, even though they may reach different conclusions about what those functions should accomplish.

The framework thus provides common ground for corporate governance improvement that transcends ideological divisions. By focusing on institutional effectiveness rather than ultimate goals, it offers practical guidance that can enhance corporate performance regardless of whether companies prioritize shareholder value, stakeholder interests, or other objectives. The tripartite framework reveals corporate governance as the fundamental institutional architecture that translates corporate intentions into social outcomes. This recognition opens new possibilities for corporate law scholarship and practice that move beyond tired debates about corporate purpose toward more sophisticated analysis of corporate capabilities and institutional design. In an era of increasing corporate influence and growing demands for corporate accountability, such analysis has never been more urgently needed.

179. Morgan Stanley, “Sustainable fund returns lag peers’ in second half of 2024,” March 4, 2025, www.morganstanley.com/insights/articles/sustainable-funds-performance-second-half-2024 (last accessed September 16, 2025).
180. Silla Brush, Saijel Kishan, and Alastair Marsh, “BlackRock exits major climate group amid Wall Street exodus,” Bloomberg Law, January 10, 2025, <https://news.bloomberglaw.com/ip-law/blackrock-exits-major-climate-group-amid-wall-street-exodus> (last accessed February 6, 2025).
181. Nathaniel Meyersohn, “Walmart rolls back DEI programs after right-wing backlash,” CNN, November 26, 2024, www.cnn.com/2024/11/26/business/walmart-diversity-dei-programs-backlash (last accessed September 16, 2025).
182. Pamela N. Danziger, “Disney shifts away from DEI to return to its entertainment business mission,” Forbes, February 13, 2025, www.forbes.com/sites/pamdanziger/2025/02/13/disney-shifts-away-from-dei-to-return-to-its-primary-business-mission/ (last accessed September 16, 2025).
183. Colin Hilson, “Climate change and the politicization of ESG in the US” (2024) 6 *Frontiers in Political Science* 1332399.
184. Virginia E. Harper Ho, “Corporate climate governance,” European Corporate Governance Institute Law Working Paper No. 804/2024, 2024, <https://ssrn.com/abstract=4934244> (last accessed September 16, 2025) (emphasizing governance structures and risk management as pragmatic approaches to climate issues).
185. Bainbridge, “Business judgment rule as abstention doctrine.”
186. Simeone at 980.
187. Zohar Goshen and Richard Squire, “Principal costs: A new theory for corporate law and governance” (2017) 117 *Columbia Law Review* 767–829 (arguing that shareholder heterogeneity produces divergent preferences, meaning that empowering shareholders directly may increase governance costs); Edward B. Rock, “The logic and (uncertain) significance of institutional shareholder activism” (1991) 79 *Georgetown Law Journal* 445–506 (arguing that institutional investors differ from retail shareholders in their time horizons and policy goals, complicating assumptions of unified shareholder preferences).
188. *Van Gorkom* at 872–73 (holding that directors breach their duty of care if they approve a merger without adequate information or deliberation).
189. Bebchuk, Kastiel, and Tallarita, “For whom corporate leaders bargain” (showing that corporate leaders primarily prioritize shareholder and managerial interests in dealmaking, while stakeholder perspectives are often excluded from the decision-making process).
190. *Caremark* at 971.
191. Frank H. Easterbrook and Daniel R. Fischel, *The Economic Structure of Corporate Law* (Cambridge, MA: Harvard University Press, 1991) (presenting corporate law as a set of standardized frameworks that enhance efficiency and adapt flexibly across organizational contexts); Reinier Kraakman et al., *The Anatomy of Corporate Law: A Comparative and Functional Approach*, 3rd ed. (Oxford: Oxford University Press, 2017).

3 CORPORATE ORDERING FOR SOCIAL CONFLICTS

1. *Aronson v. Lewis*, 473 A.2d 805 at 812 (Del. 1984).
2. Bainbridge, “Business judgment rule as abstention doctrine,” 90.
3. Bainbridge, “Director primacy,” 555.
4. Bainbridge, “Director primacy,” 556.
5. *Van Gorkom* at 872, 874–75.

6. Stephen Grace Jr., Lawrence Prendergast, and Susan Koski-Grafer, “Corporate governance and information gaps: Importance of internal reporting for board oversight,” *Business Law Today*, January 15, 2018, <https://businesslawtoday.org/2018/01/corporate-governance-and-information-gaps-importance-of-internal-reporting-for-board-oversight/> (last accessed November 26, 2025).
7. Stephen M. Bainbridge and M. Todd Henderson, “Boards-R-Us: Reconceptualizing corporate boards” (2014) 66(5) *Stanford Law Review* 1051–119 at 1092.
8. Marco Da Rin, Thomas Hellmann, and Manju Puri, “A survey of venture capital research,” NBER Working Paper No. 17523, 2011, www.nber.org/papers/w17523 (last accessed November 26, 2025).
9. Joshua D. Krebs, “The rating agencies: Where we have been and where do we go from here?” (2009) 3 *Journal of Business, Entrepreneurship & Law* 134–63 at 134.
10. Thomas A. Uebler, “Reinterpreting section 141(e) of Delaware’s General Corporation Law: Why interested directors should be ‘fully protected’ in relying on expert advice” (2020) 65 *The Business Lawyer* 1023–54.
11. Thomas W. France, “Complying with oversight fiduciary duty obligations in Delaware,” *Venable Insights*, June 4, 2020, www.venable.com/insights/publications/2020/06/complying-with-oversight-fiduciary-duty-obligat (last accessed November 15, 2025).
12. Matteo Gatti, “Corporate governing: Promises and risks of corporations as socio-economic reformers,” ECGI Law Working Paper No. 730/2023, 2023, www.ecgi.global/sites/default/files/working_papers/documents/corporategoverning_o.pdf (last accessed November 21, 2025).
13. Heiko Spitzack and Erik G. Hansen, “Stakeholder governance: How stakeholders influence corporate decision making” (2010) 10(4) *Corporate Governance* 378–91.
14. Business Roundtable, “Principles of corporate governance.”
15. Michael Mangum, “The power of corporate governance (part 1 of 2),” FMI Insights, September 2019, <https://fmicorp.com/insights/thought-leadership/the-power-of-corporate-governance-part-1-of-2/> (last accessed October 22, 2025).
16. Iain Stewart, “Aristotle, Montesquieu and Dicey on ‘Separation of Powers’” (2004) 4 *Macquarie Law Journal* 187–223 at 192.
17. Alex Zhang, “Separation of structures” (2024) 10 *Virginia Law Review* 599–689 at 630.
18. Rutger Claassen, “Political theories of the business corporation” (2023) 18 *Philosophy & Social Criticism* 1–13 at 2.
19. Bryan D. Jones, “Bounded rationality and public policy: Herbert A. Simon and the decisional foundation of collective choice” (2002) 35 *Policy Sciences* 269–84 at 275–76.
20. Bas Hillebrand, Paul H. Driessen, and Oliver Koll, “Stakeholder marketing: Theoretical foundations and required capabilities” (2015) 43 *Journal of the Academy of Marketing Science* 411–28 at 414.
21. Spitzack and Hansen, “Stakeholder governance,” 380–81.
22. Charl de Villiers and Ruth Dimes, “Determinants, mechanisms and consequences of corporate governance reporting: A research framework” (2020) 25 *Journal of Management and Governance* 7–26 at 13–14.
23. Stephen M. Bainbridge, “Why a board? Group decision-making in corporate governance” (2002) 55 *Vanderbilt Law Review* 1–34 at 7.
24. Marie-Estelle Rey, “The role of board-level committees in corporate governance,” OECD Corporate Governance Working Papers No. 24, 2022, www.oecd.org/content/dam/oecd/en/publications/reports/2022/09/the-role-of-board-level-committees-in-corporate-governance_907doa67/8a97a3f6-en.pdf at 6 (last accessed October 22, 2025).

25. Dinara Davlembayeva and Savvas Papagiannidis, “Platform-provider relationship dynamics in the sharing economy: Challenges and implications” (2023) 111 *Industrial Marketing Management* 242–56 at 252–54.
26. Harsh V. Singh and Olivia Zeydler, “Why the net-zero transition needs a transparency boost,” World Economic Forum, November 10, 2022, www.weforum.org/stories/2022/11/net-zero-transition-transparency-cop27-climate-change (last accessed October 25, 2025).
27. George S. Georgiev, “Disclosure as a corporate governance tool: Channels and challenges,” in Jeffrey N. Gordon and Wolf-Georg Ringe (eds.), *The Oxford Handbook of Corporate Law and Governance* (Oxford: Oxford University Press, 2025), pp. 2–39.
28. John M. Holcomb, “Ethics and compliance committees of corporate boards: An exploratory study” (2017) 14 *Corporate Ownership & Control* 114–21 at 115.
29. Tom Moore, “The corporate power reset that makes Citizens United irrelevant,” Center for American Progress, September 15, 2025, www.americanprogress.org/article/the-corporate-power-reset-that-makes-citizens-united-irrelevant/ (last accessed October 24, 2025).
30. Brian Libgober, “Meetings, comments, and the distributive politics of rulemaking” (2020) 15 *Quarterly Journal of Political Science* 449–81 at 451.
31. Lee Drutman, “How corporate lobbyists conquered American democracy,” The Atlantic, April 20, 2015, www.theatlantic.com/business/archive/2015/04/how-corporate-lobbyists-conquered-american-democracy/390822/ (last accessed October 24, 2025).
32. Eliska Drapalova and Kai Wegrich, “Platforms’ regulatory disruptiveness and local regulatory outcomes in Europe” (2024) 13 *Internet Policy Review* 2–28 at 11.
33. Alex Hern, “Facebook releases content moderation guidelines – rules for posts made by billions of users,” The Guardian, April 24, 2018, www.theguardian.com/technology/2018/apr/24/facebook-releases-content-moderation-guidelines-secret-rules (last accessed October 22, 2025); Derek Khanna, “Tesla and the car dealers’ lobby,” Cato Institute Regulation, Summer 2014, www.cato.org/regulation/summer-2014/tesla-car-dealers-lobby (last accessed October 24, 2025).
34. Alexander Benlian et al., “Algorithmic management” (2022) 64 *Business & Information Systems Engineering* 825–39 at 834.
35. Drapalova and Wegrich, “Platforms’ regulatory disruptiveness.”
36. Michael Klausner, “The contractarian theory of corporate law: A generation later” (2006) 31 *Journal of Corporation Law* 779–900 at 782.
37. Finn Janning, Wafa Khlif, and Coral Ingley, “Transparency is (full) disclosure in corporate governance,” in Finn Janning, Wafa Khlif, and Coral Ingley (eds.), *The Illusion of Transparency in Corporate Governance: Does Transparency Help or Hinder True Ethical Conduct?* (London: Palgrave Macmillan, 2020), pp. 57–82.
38. John G. Ruggie, “Multinationals, corporate social responsibility and public regulation: A perspective on corporate self-regulation in the global economy” (2017) 12 *Regulation & Governance* 317–33 at 320.
39. Brett H. McDonnell, “Stakeholder engagement” (2022) 51 *Florida State University Law Review* 305–72 at 316–8.
40. Economic Policy Institute, “Unequal power: How the assumption of equal bargaining power in the workplace undermines freedom, fairness, and democracy,” www.epi.org/unequalpower/home/ (last accessed October 26, 2025).
41. McDonnell, “Stakeholder engagement,” 318.
42. Coase, “The nature of the firm.”
43. Coline Serres, Marek Hudon, and François Maon, “Social corporations under the spotlight: A governance perspective” (2022) 37 *Journal of Business Venturing* 106192 at 2–3.

44. David Cray, “Control and coordination in multinational corporations” (1984) 15 *Journal of International Business Studies* 85–97 at 86.
45. Raphael Amit and Christoph Zott, “Value creation in e-business” (2001) 22 *Strategic Management Journal* 493–520 at 498.
46. Christian Homburg, Danijel Jozić, and Christina Kuehn, “Customer experience management: Toward implementing an evolving marketing concept” (2017) 45 *Journal of the Academy of Marketing Science* 377–401 at 379.
47. Michel Wedel and P.K. Kannan, “Marketing analytics for data-rich environments” (2016) 80 *Journal of Marketing* 97–121 at 99–100.
48. Lina M. Khan, “The separation of platforms and commerce” (2019) 119 *Columbia Law Review* 973–1098 at 1004.
49. Julie E. Cohen, “Turning privacy inside out” (2019) 20 *Theoretical Inquiries in Law* 1–37 at 26.
50. Michele Micheletti and Dietlind Stolle, “Mobilizing consumers to take responsibility for global social justice” (2007) 6 *Annals of the American Academy of Political and Social Science* 157–75 at 169.
51. Amitai Etzioni, “The limits of transparency,” in Emmanuel Alloa & Dieter Thomä (eds.), *Transparency, Society and Subjectivity: Critical Perspectives* (London: Palgrave Macmillan, 2018), pp. 176–201.
52. Pauline T. Kim, “Beyond principal–agent theories: Law and the judicial regulation of workers” (1997) 1997 *Law and Contemporary Problems* 535–75 at 544.
53. Robert Klonoski, “Defining employee benefits: A managerial perspective” (2016) 6 *International Journal of Human Resource Studies* 52–72 at 59.
54. Roger Friedland and Robert R. Alford, “Bringing society back in: Symbols, practices, and institutional contradictions,” in W. W. Powell and P. J. DiMaggio (eds.), *The New Institutionalism in Organizational Analysis* (Chicago: University of Chicago Press, 1991), pp. 232–67.
55. Panikos Georgallis, “The link between social movements and corporate social initiatives: Toward a multi-level theory” (2017) 142 *Journal of Business Ethics* 735–51 at 736.
56. Todd Corley, Vontrese (Voni) Pamphile, and Katina Sawyer, “What has (and hasn’t) changed about being a chief diversity officer,” *Harvard Business Review*, September 23, 2022, <https://hbr.org/2022/09/what-has-and-hasnt-changed-about-being-a-chief-diversity-officer> (last accessed October 24, 2025).
57. Umesh Chandra Tiwari, “DEI metrics in executive compensation,” *Harvard Law School Forum on Corporate Governance*, June 26, 2024, <https://corpgov.law.harvard.edu/2024/06/26/dei-metrics-in-executive-compensation/> (last accessed October 24, 2025).
58. Eric Lau and Taylor Telford, “‘DEI’ vanishing from corporate filings, mirroring business world’s retreat,” *The Washington Post*, April 30, 2025, www.washingtonpost.com/business/interactive/2025/dei-companies-sec-filings/ (last accessed October 24, 2025).
59. Lara A. Flath, David E. Schwartz, and Amy Van Gelder, “Corporate DEI policies face scrutiny following SCOTUS affirmative action decision,” *Skadden Insights*, September 2023, www.skadden.com/insights/publications/2023/09/quarterly-insights/corporate-dei-policies-face-scrutiny (last accessed October 24, 2025); Michael Hatcher et al., “Ten state attorneys general launch inquiry into major financial institutions’ DEI/ESG programs,” *Jackson Lewis P. C.*, January 29, 2025, www.jacksonlewis.com/insights/ten-state-attorneys-general-launch-inquiry-major-financial-institutions-dei-esg-programs (last accessed November 21, 2025); Taylor Telford, “The growing battle over corporate diversity practices, explained,” *The Washington Post*, October 2, 2023, www.washingtonpost.com/business/2023/10/02/corporate-diversity-inclusion-affirmative-action-ruling/ (last accessed October 24, 2025).

60. Geri Stengel, “Black Lives Matter protests move corporate D&I initiatives into the spotlight,” *Forbes*, June 17, 2020, www.forbes.com/sites/geristengel/2020/06/17/black-lives-matter-protests-moves-corporate-di-initiatives-into-the-spotlight/ (last accessed October 24, 2025).
61. Maria Aspan, “Corporate America’s retreat from DEI has eliminated thousands of jobs,” *NPR*, May 27, 2025, www.npr.org/2025/05/27/nx-s1-5307319/dei-jobs-trump (last accessed October 24, 2025).
62. Anurada Jain, “The gig economy: Disruption, innovation, and economic evolution” (2024) 30 *Educational Administration Theory and Practice Journal* 5133–40 at 5135.
63. Airbnb, “Official codification of party ban,” June 28, 2022, <https://news.airbnb.com/official-codification-of-party-ban/> (last accessed October 31, 2025); Airbnb.org, “Airbnb.org expands temporary housing support for refugees,” June 20, 2024, <https://news.airbnb.com/airbnb-org-expands-temporary-housing-support-for-refugees/> (last accessed October 31, 2025).
64. Audrea Chen, “Employee or independent contractor? A legal analysis of Uber’s worker misclassification,” *Columbia Law Review Online*, September 5, 2024, www.culawreview.org/current-events-2/employee-or-independent-contractor-a-legal-analysis-of-ubers-worker-misclassification (last accessed October 31, 2025).
65. Faiza Patel and Laura Hecht-Fellella, “Facebook’s content moderation rules are a mess,” *Brennan Center for Justice*, February 22, 2021, www.brennancenter.org/our-work/analysis-opinion/facebooks-content-moderation-rules-are-mess (last accessed October 31, 2025).
66. Jeroen K. G. Hopster and Matthijs M. Maas, “The technology triad: Disruptive AI, regulatory gaps and value change” (2024) 4 *AI and Ethics* 1051–65 at 1064.
67. Abdullah Alsaleh, “The impact of technological advancement on culture and society” (2024) 14 *Scientific Reports* 32140 at 4.
68. Kezia Farnham, “Corporate governance rules: What they are & why they matter,” *Diligent Corporation*, May 21, 2024, www.diligent.com/resources/blog/corporate-governance-rules (last accessed November 1, 2025).
69. Eduardo Schiehl, Henrique Castro Martins, and Krista B. Lewellyn, “Symbolic or substantive? How shareholder heterogeneity shapes the purpose of the corporation” (2025) 58 *Long Range Planning* 102582 at 19–20.
70. Robert P. Bartlett and Ryan Bubb, “Corporate social responsibility through shareholder governance,” *ECGI Working Paper Series in Law* No. 677, 2023, www.ecgi.global/sites/default/files/working_papers/documents/corporatesocialresponsibilitythroughshareholdergovernance_o.pdf (last accessed November 21, 2025).
71. Marisa McVey, “Untangling the authority of external experts in the corporate implementation of the UN Guiding Principles on Business and Human Rights” (2022) 21 *Journal of Human Rights* 620–38 at 627.
72. McVey, “Untangling the authority of external experts,” 627.
73. Spitzack and Hansen, “Stakeholder governance,” 378.
74. Starbucks, “C.A.F.E. Practices: Starbucks’ approach to ethically sourcing coffee,” February 28, 2024, <https://stories.starbucks.com/stories/2024/cafe-practices-starbucks-approach-to-ethically-sourcing-coffee> (last accessed November 1, 2025).
75. Starbucks, “Sustainability at the heart of Starbucks coffee sourcing,” <https://athome.starbucks.com/story/sustainability-heart-starbucks-coffee-sourcing/> (last accessed November 1, 2025).
76. Paul Wiegmann et al., “Competing standard-setting organizations: A choice architecture for innovators?” (2022) 51 *Research Policy* 1–15 at 11–12.
77. World Resources Institute and World Business Council for Sustainable Development, “About GHG Protocol,” <https://ghgprotocol.org/about-us> (last accessed November 1, 2025).

78. Olivier Serrat, “A primer on corporate values,” in *Knowledge Solutions: Tools, Methods, and Approaches to Drive Organizational Performance* (Singapore: Springer 2017), pp. 701–09.
79. Giorgia Profumo, Lara Penco, and Ilaria Tutore, “Mission statements and the sustainability communication” (2019) 2 *Symphonya: Emerging Issues in Management* 95–108 at 97–100.
80. Alperen Afşin Gözlügül, “Blinded by ‘fairness’: Why we need (strong) procedural safeguards in screening self-dealing and obtaining a fair price is not the answer” (2022) *European Business Organisation Law Review* 23(4) 634–67 at 634.
81. Ratna Sari and Muslim Muslim, “Corporate transparency and environmental reporting: Trends and benefits” (2024) 4 *Amkop Management Accounting Review* 1–18 at 14.
82. Ingo Pies et al., “Trade-offs in stakeholder theory: An ordonomic perspective” (2024) 20 *Society and Business Review* 975–95 at 976.
83. Simone de Colle, “A stakeholder management model for ethical decision making” (2005) 6 *International Journal of Management and Decision Making* 299–314 at 312–13.
84. Marwan Mansour et al., “The relationship between corporate governance quality and firm performance” (2022) 14 *Sustainability* 1–25 at 4.
85. Renata Wandroski Peris et al., “Does better corporate governance increase operational performance?” (2017) 17 *Corporate Governance: The International Journal of Business in Society* 524–37.
86. Timothy J. McClimon, “Bluwashing joins greenwashing as the new corporate whitewashing,” *Forbes*, October 3, 2022, www.forbes.com/sites/timothyjmcclimon/2022/10/03/bluwashing-joins-greenwashing-as-the-new-corporate-whitewashing/ (last accessed November 1, 2025).
87. Aaron K. Chatterji and Michael W. Toffel, “The new CEO activists,” *Harvard Business Review*, March–April 2018, <https://hbr.org/2018/01/the-new-ceo-activists> (last accessed November 1, 2025).
88. Lisa M. Fairfax, “Stakeholderism, corporate purpose, and credible commitment” (2022) 108 *Virginia Law Review* 1163–247 at 1166.
89. Bloomberg, “Corporate America promised to hire more people of color – it did,” September 27, 2023, www.bloomberg.com/graphics/2023-black-lives-matter-equal-opportunity-corporate-diversity/ (last accessed November 3, 2025).
90. Ioannis Ioannou, “What the ESG backlash reveals – and what comes next” (2025), *Forbes*, March 25, 2025, www.forbes.com/sites/lbsbusinessstrategyreview/2025/03/25/what-the-esg-backlash-reveals-and-what-comes-next/ (last accessed November 3, 2025).
91. David A. Cifrino, “The future of corporate sustainability and climate action in the anti-‘woke’ era,” *Harvard Advanced Leadership Initiative Social Impact Review*, February 8, 2025, www.sir.advancedleadership.harvard.edu/articles/future-corporate-sustainability-climate-action-in-anti-woke-era/ (last accessed November 4, 2025).
92. Yale Chief Executive Leadership Institute, “Over 1,000 companies have curtailed operations in Russia – but some remain,” January 28, 2024 (updated regularly since 2022), <https://som.yale.edu/story/2022/over-1000-companies-have-curtailed-operations-russia-some-remain> (last accessed November 4, 2025).
93. Michael A. Witt, “Corporate responses to the war in Ukraine: Moral principles, strategic rationales, and stakeholder pressures” (2022) 48 *Financial Research Letters* 1–4.
94. Cifrino, “The future of corporate sustainability and climate action.”
95. Kathleen Rehbein et al., “Corporate social responsibility and the issue environment: A longitudinal analysis” (2006) 111 *Business & Society Review* 175–91 at 182–85.
96. Ronald Mitchell et al., “Toward a theory of stakeholder identification and salience: Defining the principle of who and what really counts” (1997) 22 *Academy of Management Review* 853–86 at 869.

97. Organisation for Economic Co-operation and Development, “G20/OECD Principles of Corporate Governance,” 2023, www.oecd.org/content/dam/oecd/en/publications/reports/2023/09/g20-oecd-principles-of-corporate-governance-2023_60836fcb/ed750b30-en.pdf (last accessed November 15, 2025).
98. R. Edward Freeman, *Strategic Management: A Stakeholder Approach* (Cambridge: Cambridge University Press, 2015), pp. 25–26.
99. Fairfax, “Stakeholderism, corporate purpose, and credible commitment,” p. 1166.
100. McClimon, “Bluewashing joins greenwashing.”
101. OECD, “G20/OECD Principles of Corporate Governance,” p. 23.
102. Nicole Dando and Tracey Swift, “Transparency and assurance: Minding the credibility gap” (2003) 44 *Journal of Business Ethics* 195–200 at 196.
103. McVey, “Untangling the authority of external experts,” p. 627.
104. Thi Hong Van Hoang et al., “Environmental transparency and performance: Does the quality of ESG disclosure matter?” (2021) 10 *Corporate Governance: The International Journal of Business in Society* 1–11 at 6.
105. Alex Edmans, *Grow the Pie: How Great Companies Deliver Both Purpose and Profit* (Cambridge: Cambridge University Press, 2020).
106. Russ Banham, “Shareholders: The new rules of engagement,” Corporate BoardMember, <https://boardmember.com/shareholders-the-new-rules-of-engagement/> (last accessed 7 November 2025).
107. Zhiyang Liu et al., “Stakeholders and ESG disclosure strategies adoption: The roles of firm and stakeholder characteristics” (2022) 10 *Elementa: Science of the Anthropocene* 1–12 at 3–4.
108. Joanna Krasodomska et al., “The impact of companies’ trust-building efforts on sustainability reporting and assurance quality: Insights from Europe” (2024) 33 *Meditari Accountancy Research* 246–79 at 249.
109. Bentleys, “ESG reporting: Key insights for accurate and effective disclosure,” April 1, 2025, www.bentleys.com.au/resources/esg-reporting-key-insights-for-accurate-and-effective-disclosure/ (last accessed November 4, 2025).
110. Deloitte, “Using the COSO framework to establish internal controls over sustainability reporting (ICSR)” April 21, 2023, <https://dart.deloitte.com/USDART/home/publications/deloitte/heads-up/2023/coso-framework-for-internal-controls-over-sustainability> (last accessed November 4, 2025).
111. Reiner Quick et al., “Assurance on CSR reports impact on credibility perceptions of non-financial information by bank directors” 28 *Meditari Accountancy Research* 833–62 at 833.
112. Marta Isabel Garcia-Rivas et al., “Corporate social responsibility reports: A review of the evolution, approaches and prospects” (2023) 9 *Heliyon* 1–18 at 13–14.
113. Merle Ojasoo, “CSR reporting, stakeholder engagement and preventing hypocrisy in corporations” (2016) 6 *Journal of Global Entrepreneurship Research* 1–14 at 3–7.
114. Simone Altoe, Carmen Simion, and Martin Wennerström, “ESG-linked compensation: Driving action on sustainability,” Sustainalytics, April 3, 2024, www.sustainalytics.com/esg-research/resource/investors-esg-blog/incentivizing-change-how-esg-linked-compensation-can-advance-sustainability-initiatives (last accessed November 4, 2025).
115. Sarah Marschlich et al., “The effect of third-party certifications on corporate social responsibility communication authenticity and credibility” (2025) 9 *Corporate Communications: An International Journal* 1–20 at 5.
116. Ayes Saka-Helmhout et al., “Stakeholder engagement strategies for impactful corporate social innovation initiatives by multinational enterprises” (2024) 30 *Journal of International Management* 1–20 at 9–11.

117. OECD, “G20/OECD Principles of Corporate Governance,” p. 47.
118. Faiza Siddiqui et al., “The role of corporate governance and reputation in the disclosure of corporate social responsibility and firm performance” (2023) 9 *Heliyon* 1–14 at 4–5.
119. Erica Steckler et al., “Authenticity and corporate governance” 114 *Journal of Business Ethics* 951–63 at 955–6.
120. Tamanna Fatima and Said Elbanna, “Corporate social responsibility (CSR) implementation: A systematic review of 122 empirical studies” (2022) 31 *Business Ethics: A European Review* 678–95.
121. Fatima and Elbanna, “Corporate social responsibility (CSR) implementation,” p. 682.
122. Alan S. Gutterman, “CSR governance and organizational culture,” ResearchGate, June 14, 2023, www.researchgate.net/publication/371680183_CSR_Governance_and_Organizational_Culture (last accessed November 5, 2025).
123. Fatima and Elbanna, “Corporate social responsibility (CSR) implementation,” p. 679.
124. Associated Press, “Microsoft lets cloud users keep personal data within Europe to ease privacy fears,” January 11, 2024, <https://apnews.com/article/70e50c9051b57794f6074f15624a000f> (last accessed November 8, 2025).
125. Simon Jessop, “Companies setting climate transition plans up 44% in 2023, research shows,” Reuters, June 18, 2024, www.reuters.com/sustainability/climate-energy/companies-setting-climate-transition-plans-up-44-2023-research-shows-2024-06-18/ (last accessed November 8, 2025).
126. Oliver G. Kayas, “Workplace surveillance: A systematic review, integrative framework, and research agenda” (2023) 168 *Journal of Business Research* 1–16 at 6.
127. Bernd Manz, “Corporate governance reporting, disclosures, monitoring and analytics” (2025) 33(2) *Corporate Governance: An International Review* 145–62 at 150.
128. Kristin Burnham, “Scope 3 emissions top supply chain sustainability challenges,” MIT Sloan Ideas Made to Matter, December 18, 2024, <https://mitsloan.mit.edu/ideas-made-to-matter/scope-3-emissions-top-supply-chain-sustainability-challenges> (last accessed 14 November 2025).
129. Adwoa Bagalini, “5 ways intersectionality affects diversity and inclusion at work,” World Economic Forum, July 22, 2020, www.weforum.org/stories/2020/07/diversity-inclusion-equality-intersectionality/ (last accessed November 8, 2025).
130. Lauren Weber and Emily Glazer, “Sex scandals. Accounting fraud. It’s all showing up on the corporate hotline,” *The Wall Street Journal*, September 12, 2025, www.wsj.com/business/nestle-employee-complaint-hotline-74ac4fde (last accessed November 8, 2025).
131. Patrick Temple-West, “More women make it on to US company boards but fewer reach the C-suite,” *Financial Times*, May 20, 2024, www.ft.com/content/a586a0d6-d299-463c-a531-22539446ff87 (last accessed 14 November 2025).
132. Rachel Amow-Richman, James Hicks, and Steven Davidoff Solomon, “Do social movements spur corporate change? The rise of ‘MeToo termination rights’ in CEO contracts” (2022) 98 *Indiana Law Journal* 125–74 at 151.
133. Weber and Glazer, “Sex scandals.”
134. Patrick Temple-West, “Companies drop DEI targets from bonus plans on pressure from conservatives,” *Financial Times*, July 21, 2024, www.ft.com/content/63ac21a5-9929-4bc3-9089-03299f023bcd (last accessed November 10, 2025).
135. Ross Kerber, “Just 35 complaints of debanking cite political bias despite Trump order,” Reuters, August 19, 2025, www.reuters.com/legal/transaccional/just-35-complaints-debanking-cite-political-bias-despite-trump-order-2025-08-19/ (last accessed November 10, 2025).

136. Jay Peters, “Zuckerberg says he’s moving Meta moderators to Texas because California seems too ‘biased’,” *The Verge*, January 7, 2025, www.theverge.com/2025/1/7/24338305/meta-mark-zuckerberg-moving-meta-moderators-texas-california-bias (last accessed November 14, 2025); Lauren Goode, “Twitter’s former head of trust and safety finally breaks her silence,” *Wired*, November 21, 2023, www.wired.com/story/del-harvey-twitter-trust-and-safety-breaks-her-silence/ (last accessed November 14, 2025).
137. Ariel Zients, “Companies expand DEI programs to include first-generation professionals,” *Bloomberg Law*, April 5, 2024, <https://news.bloomberglaw.com/esg/companies-expand-dei-programs-to-include-first-generation-professionals> (last accessed November 14, 2025).
138. Society for Human Resource Management, “Regularly monitoring, using DE&I data is key to closing organizational gaps,” September 11, 2021, www.shrm.org/topics-tools/news/inclusion-diversity/regularly-monitoring-using-dei-data-key-to-closing-organizational-gaps (last accessed November 14, 2025).
139. Miriam H. Baer, “Governing corporate compliance” (2009) 50 *Boston College Law Review* 949–1019 at 1001.
140. Jean-Michel Ferat and Shelly Mady, “The significance of audit and monitoring in a compliance programme,” *Global Investigations Review*, August 7, 2024, <https://globalinvestigationsreview.com/guide/the-guide-compliance-archived/third-edition/article/the-significance-of-audit-and-monitoring-in-compliance-programme> (last accessed November 14, 2025).
141. Francesco Prencipe, “Accountability between compliance and legitimacy: A structured conceptual review of sustainability governance” (2025) 17 *Sustainability* 9305–24 at 9310.

4 WHY NOT REGULATE?

1. John C. Coffee, “No soul to damn, no body to kick’: An unsandalized inquiry into the problem of corporate punishment” (1981) 79 *Michigan Law Review* 386–459; Cary Coglianese and David Lazer, “Management-based regulation: Prescribing private management to achieve public goals” (2003) 37 *Law & Society Review* 691–730 (together explaining that regulation has long been the primary tool for addressing divergences between private incentives and public welfare, particularly when corporate behavior imposes externalities).
2. Jill E. Fisch and Jeff Schwartz, “How did corporations get stuck in politics and can they escape?” (2024) 3 *University of Chicago Business Law Review* 325–60.
3. Jonathan B. Wiener, “The regulation of technology, and the technology of regulation” (2004) 26 *Technology in Society* 483–500; Coglianese and Lazer, “Management-based regulation,” 700.
4. Ian Ayres and John Braithwaite, *Responsive Regulation: Transcending the Deregulation Debate* (Oxford: Oxford University Press, 1992), pp. 7–10.
5. John Braithwaite, “Enforced self-regulation: A new strategy for corporate crime control” (1982) 80 *Michigan Law Review* 1466–507.
6. Gunther Teubner, “Global Bukowina: Legal pluralism in the world-society,” in Gunther Teubner (ed.), *Global Law without a State* (Aldershot: Dartmouth, 1996), pp. 3–28; Brian Tamanaha, “A non-essentialist version of legal pluralism” (2000) 27 *Journal of Law and Society* 296–321.

How to Govern AI?

The general counsel of a major AI company got the telephone call that all in-house lawyers fret.¹ The U.S. Department of Defense wanted to invite the company to a meeting, in D.C., to talk about the implications of artificial intelligence for national security. It was a few months after OpenAI had released ChatGPT for free, taking the world by storm.² Suddenly, everyone had been left amazed by how articulate and knowledgeable it seemed. There was a lot of excitement about potential applications, from providing content and education to facilitating menial support tasks.³ There were also concerns about misinformation, ethics, and the reinforcing of harmful stereotypes, the revolutionization of industries, and the replacement of humans by machines.⁴ But national security was an issue with an entirely different set of considerations, because it affected the integrity of the country itself.⁵

It was also an issue about which any AI company would have thought only in passing at the time. Yet, management felt that they had to go to D.C. and answer the military's questions. To take the extraordinary step of inviting them, they must have done something to enter their radar.⁶ But what could that possibly be? They were determined to be prepared. When national security considerations enter the scene, the government has a lot of extraordinary, far-reaching powers that they could use against the company. They needed to be ready to put down as many fires as possible. So, the general counsel worked hard to be briefed about various aspects of the company's operations, the people who accessed its platform, and the sources of information it utilized. They traveled to D.C. and entered the meeting ready to defend the company and determined to be truthful. After the initial introductions and pleasantries, the senior military officers went to the heart of the topic. They had seen AI grab the world's attention and they were worried about the future. So, they thought, they should be proactive and ask AI companies: in their view, how could AI impact national security?⁷ Management immediately realized that they knew a lot more about the topic than the military officers across the table. They went on to have a very engaging meeting.

It sounds almost like a truism that government often finds itself one step behind the industry.⁸ Particularly in areas such as technology or finance, which bring together some of the best minds in the world with capital exploring investment opportunities,

regulators often find themselves struggling to follow developments.⁹ Still, the incident above is extraordinary because it shows not only how far behind governments can fall, but also how much they must rely on private companies to find their way. The U.S. military is, one would hope, far more sophisticated and well-resourced than the average regulatory agency.¹⁰ Still, when developments occur at breakneck speeds, governments have little option apart from turning to companies for information, analysis, and guidance.

This chapter explores two different pathways for governing AI. In the conventional regulatory model, AI companies become regulated entities subject to compliance and disclosure obligations and, possibly, licensing requirements.¹¹ Federal agencies supervise various AI products and aspects of the firm's operations, governance, and board personnel. In the corporate ordering model, AI companies establish their own internal governance regimes, disclosing their contours and inviting accountability and engagement. They may also come together, possibly in an industry-funded regulatory body, to set standards about their products and services.¹² Groups of experts in different fields collaborate to propose standards in their respective areas. The standard-setting process is transparent with input from governments, citizens' groups, NGOs, and other interested parties, informing the public about developments.

When faced with two different regulatory models, conventional legal analysis would contrast them by highlighting their strengths and weaknesses. My goal in this chapter is to argue that viewing this as a dilemma is misconstrued. Not only can the two models coexist, but in fact, the corporate ordering effort is a necessary precondition for any official regulatory regime. Corporations must do the legwork of preparing for regulation, for the simple reason that they understand better than anyone how their activity may impact other groups. This is particularly true in areas where technological developments disrupt existing industry structures at a very quick pace, like in AI. Moreover, corporations will also be required to implement any solutions to social concerns, so their input in designing them is crucial.

One might worry that, if the substantive core of a regulatory obligation originates from the regulated entity, that is, the company, then the likelihood of capture is extremely high. This concern is central to any self-regulatory effort. Addressing it requires corporations to implement measures that enhance their credibility, help them take into account other parties' views, and increase their legitimacy. The chapter offers some examples of how governance can help in this direction. But it also highlights that, if corporate ordering fails, then governments might remain in the dark and stakeholders may continue to suffer harm.

8.1 STAKES: INNOVATION AT FULL SPEED, SAFETY AND STEWARDSHIP

Few technological developments have elicited language as apocalyptic as artificial intelligence. The stakes could not have been higher, as AI is on track to revolutionize

the economy and throw the workforce into chaos. From a political economy perspective, lawmakers and tech moguls, progressives and conservatives, elite thinkers, and working class advocates are united in the belief that the risks posed by AI are too severe to be left unchecked.¹³ As governments worldwide race to establish regulatory frameworks, a fundamental tension lies at the heart of AI development: the conflict between rapid innovation and responsible stewardship. This tension crystallizes the competing imperatives of accelerated technological advancement and careful safety considerations in an era of unprecedented change. It also illustrates why AI is a hard case for the argument this book is advancing. Given the magnitude of the challenges, one would not expect private companies to be the ones taking the lead in setting up appropriate frameworks.

8.1.1 *The Stewardship Imperative: Safeguarding the Future*

The case for careful stewardship of AI development rests on an intricate web of concerns, ranging from existential threats to immediate societal risks. At the most fundamental level, stewardship advocates worry about AI systems that could become uncontrollable and act against human interests.¹⁴ There is a real possibility that highly developed systems may actively resist human oversight and pursue goals that conflict with human values. The challenge of ensuring these systems remain subordinate to human control becomes exponentially more complex as their capabilities increase.¹⁵ Industry leaders such as Bill Gates and Elon Musk have raised alarm about these existential threats, arguing that even though catastrophic scenarios might be decades away, the groundwork for prevention must be laid now.¹⁶ Their concern stems from a crucial insight: once superintelligent AI systems exist, implementing controls might be impossible.¹⁷ This creates an asymmetric risk scenario where the cost of moving too quickly far outweighs the benefits of rapid development.

The stewardship perspective gains additional urgency from more immediate threats. AI systems already demonstrate alarming capabilities for malicious deployment. Advanced AI can automate the discovery of software vulnerabilities, potentially compromising critical systems faster than human defenders can respond.¹⁸ The threat extends beyond traditional cybersecurity, as AI-powered systems could orchestrate sophisticated disinformation campaigns, manipulating public discourse and undermining democratic institutions at unprecedented scales.¹⁹

Critical infrastructure faces particular risks. AI-powered cyber intrusions could simultaneously target multiple systems, creating cascading failures across interconnected networks.²⁰ A coordinated attack could disable financial systems, disrupt energy grids, and paralyze transportation networks, causing widespread societal disruption. The automation and scale possible with AI attacks make traditional security measures increasingly inadequate.

Perhaps most insidiously, AI enables unprecedented surveillance capabilities. By aggregating data across multiple domains, including social media, financial

transactions, location data, and communication metadata, AI systems can create detailed profiles of individuals and groups.²¹ This mass surveillance capability threatens not just privacy but fundamental human liberties, potentially enabling social control mechanisms that would be impossible without AI.²² The combination of bulk data access and sophisticated analysis tools creates surveillance possibilities that even George Orwell could not have imagined.

8.1.2 *The Speed Imperative: Innovation as Economic and Strategic Necessity*

The argument for rapid AI development presents its own compelling logic, grounded in economic necessity and geopolitical reality. AI is not just another technological advance, but rather represents a fundamental transformation of human productive capabilities. The technology is already revolutionizing the workforce in ways that previous automation waves never approached.²³ Unlike earlier technological shifts that primarily affected routine manual labor, AI is disrupting professions that were thought to require uniquely human capabilities.²⁴

The scope of this disruption is staggering. Legal analysis, medical diagnosis, creative writing, and artistic production, all domains previously considered safe from automation, are being transformed by AI capabilities.²⁵ The recent Hollywood writers' strike highlights how AI is reshaping even creative industries.²⁶ With one in five U.S. workers facing potential job displacement or transformation, the economic imperative for rapid development becomes clear: those who fail to adapt quickly risk being left behind.²⁷

This transformation extends beyond domestic economic concerns into the realm of global power dynamics. Nations and organizations that successfully deploy AI gain unprecedented advantages in productivity and capability.²⁸ In critical sectors such as defense, transportation, and healthcare, AI advantages could become decisive. Military applications, from autonomous weapons to strategic planning systems, could shift the balance of global power.²⁹ Healthcare breakthroughs enabled by AI could provide dramatic advantages in life expectancy and quality of life.³⁰

The global race for AI supremacy creates a prisoner's dilemma in international relations. While all nations might benefit from careful, coordinated AI development, the potential advantages of pulling ahead create strong incentives for rapid deployment.³¹ The U.S. restriction on exporting advanced AI chips to China³² exemplifies this dynamic. While motivated by security concerns, U.S. policies also reflect the recognition that AI capabilities have become central to national power.

This tension between speed and stewardship reflects a singular challenge in human technological development. The same capabilities that promise extraordinary benefits also create equally novel risks. As public anxiety about AI grows, with Americans expressing more concern than excitement, rising from 37 percent in 2021 to 52 percent in 2023,³³ the challenge of balancing innovation and safety becomes

increasingly urgent. Yet, governments are in a significant disadvantage in grasping the full range of concerns AI rises, as well as coming up with rules that would address very different needs. Thus, it falls on companies at the forefront of AI development to determine how to best address these challenges. The [following section](#) explores how this abstract debate manifests in concrete organizational decisions, with different shareholders and stakeholders holding varying views.

8.2 HETEROGENEITY OF PREFERENCES

The differing perspectives on how to address AI concerns divide not only public opinion, but also some of the key actors spearheading the AI revolution. Throughout the business world, managers, investors, financiers, and employees take different positions regarding the necessity of stewardship and the urgency of adjusting to new realities. The November 2023 leadership crisis at OpenAI, one of the pioneers in LLMs, serves as a compelling illustration of how heterogeneous preferences among stakeholders can create significant organizational tensions in artificial intelligence development.³⁴ This section analyzes how divergent stakeholder priorities regarding AI safety, commercial development, and governance led to one of the most dramatic corporate conflicts in recent technology history.

The crisis began on November 17, 2023, when OpenAI's board abruptly fired CEO Sam Altman, citing concerns about his candor in communications.³⁵ This decision triggered a cascade of events that exposed deep-seated tensions between different stakeholder groups. The board's initial statement claimed that Altman "was not consistently candid in his communications with the board, hindering its ability to exercise its responsibilities."³⁶ However, this dramatic action revealed fundamental disagreements about OpenAI's direction and governance that had been brewing beneath the surface.³⁷

The immediate aftermath demonstrated the complex web of shareholder and stakeholder interests, with each group bringing distinct preferences and priorities to the forefront. The board, led by independent directors, emphasized their unique governance mandate, noting that unlike traditional companies, they weren't tasked with maximizing shareholder value. Instead, their stated mission focused on ensuring the creation of "broadly beneficial" artificial general intelligence.³⁸ But faced with opposition from key funders, who were ready to redirect their resources elsewhere, and the threat of an en masse resignation by the company's employees, the board had little option but to retreat.³⁹ With Altman reinstated, board members resigned and paved the way for a leadership reshuffle and direction realignment.⁴⁰

8.2.1 *Stewardship Advocates*

The perspective of stewardship advocates represents a fundamental commitment to ensuring artificial intelligence development serves humanity's long-term interests.

This stakeholder group, including Helen Toner, Ilya Sutskever, and Elon Musk, prioritizes the careful governance and responsible development of AI over commercial or technological imperatives.⁴¹

Helen Toner's position, as articulated through her academic work and board actions, emphasized the paramount importance of responsible AI development.⁴² Her paper "Decoding Intentions: AI and Costly Signals" had criticized OpenAI's rapid deployment of technologies like GPT-4, arguing that commercial pressures were compromising essential safety protocols.⁴³ The tension between Toner and Altman reached a critical point when Altman attempted to remove her from the board due to her paper's criticisms of OpenAI's approach.⁴⁴ Toner's later statement that "*even OpenAI's potential self-destruction might align with the mission*" of ensuring AGI benefits humanity demonstrates the depth of the safety-first perspective.⁴⁵ This view was so fundamental that when an OpenAI lawyer suggested the board's decision might lead to the company's downfall, Toner maintained her position, arguing that such an outcome could still serve the broader mission of ensuring beneficial AGI development.

Elon Musk, an original cofounder of OpenAI, represents perhaps the most forceful articulation of the stewardship position. His 2024 lawsuit against OpenAI alleged that the organization had "*strayed from its original mission of developing AGI for the benefit of humanity.*"⁴⁶ Musk's legal action argued that the shift to a for-profit model fundamentally compromised OpenAI's founding principles. His consistent warnings about AI's existential risks and his criticism that OpenAI's current focus on commercialization only exacerbates these dangers align closely with Toner's concerns about prioritizing market success over responsible development.⁴⁷

Ilya Sutskever's role in the crisis particularly illuminates the complexity of the stewardship position. As OpenAI's chief scientist, Sutskever initially supported Altman's removal, arguing during a company all-hands meeting that the action was "necessary" to protect OpenAI's mission of "making AI beneficial to humanity."⁴⁸ However, his subsequent reversal and public regret over participating in Altman's removal highlight the challenging balance between technological progress and responsible oversight.⁴⁹

8.2.2 Commercialization and Speed Supporters

The perspective favoring rapid development and market release of AI technologies represents a distinct philosophy about how AI development should proceed, most prominently championed by Sam Altman and strongly supported by OpenAI's technical workforce. This viewpoint emphasizes the importance of practical deployment and swift iteration while attempting to maintain alignment with the broader mission of beneficial AI development.⁵⁰

Sam Altman's leadership at OpenAI embodied a complex vision that sought to bridge rapid technological advancement with responsible development.⁵¹

His approach centered on the belief that practical deployment was essential for understanding and improving AI systems.⁵² When addressing OpenAI's mission, Altman consistently framed it in terms of both technological progress and societal benefit, stating that ensuring artificial general intelligence benefits all of humanity was not just a corporate objective but a moral imperative. In 2019, Altman defended the shift of OpenAI from a nonprofit to a capped-profit model.⁵³ He stated that OpenAI would need to “invest billions of dollars in upcoming years into large-scale cloud compute, attracting and retaining talented people, and building AI supercomputers,” making commercialization necessary to support this massive scale.

The perspective of senior technical leaders further reinforced this position. The resignation of key figures like research director Jakub Pachocki and head of preparedness Aleksander Madry suggested that those closest to the technical work saw Altman's leadership style as crucial for meaningful progress.⁵⁴ Research scientist Andrej Karpathy's public statement that “The board had a chance to explain their drastic actions and they did not take it” reflected the technical community's skepticism of the board's more cautious approach.⁵⁵

Microsoft's substantial investment in OpenAI aligned closely with this rapid deployment perspective. Their \$10 billion investment, structured significantly around cloud computing resources, demonstrated a belief in the importance of scaling AI development and deployment.⁵⁶ Satya Nadella's vision for the partnership emphasized both technological advancement and practical application, seeing rapid deployment as crucial for progress. His statement that “Microsoft's partnership with OpenAI is built on a foundation of ethical AI development” while simultaneously pushing for accelerated development and integration of OpenAI's technologies into Microsoft's products reveals the attempt to balance speed with responsibility.⁵⁷

This perspective gains particular support from market-oriented stakeholders, including venture capital firms such as Khosla Ventures. Vinod Khosla's description of Altman as a “once in a generation CEO” whose “positive mark on the world will be indelible” reflects the belief that rapid, market-oriented development under visionary leadership is crucial for realizing AI's potential benefits.⁵⁸

8.2.3 *Crisis and Overhaul*

The tension between rapid deployment and careful oversight came to a head during the leadership crisis in November 2023. The board's initial decision to remove Altman, citing concerns about his candor in communications, reflected deep-seated anxieties about the pace and direction of development. However, the swift and nearly universal backlash from employees, investors, and industry partners suggested strong support for the rapid development approach.

The employee response to Altman's removal provided crucial validation of his approach to AI development. The near-unanimous workforce reaction, with over 650 of approximately 770 employees threatening to resign, demonstrated remarkable

alignment with Altman's vision.⁵⁹ The employee letter to the board offered a detailed defense of this perspective on AI development: "OpenAI is the world's leading AI company. We, the employees of OpenAI, have developed the best models and pushed the field to new frontiers. Our work on AI safety and governance shapes global norms. The products we built are used by millions of people around the world. Until now, the company we work for and cherish has never been in a stronger position."⁶⁰ This statement reveals how the technical workforce viewed rapid development and deployment as complementary to, rather than in conflict with, responsible AI development. Their assertion that the board's actions had "jeopardized all of this work and undermined our mission and company" suggests they saw no contradiction between swift progress and ethical considerations.

The depth of employee commitment to this vision became particularly clear in their ultimatum to the board: "We are unable to work for or with people that lack competence, judgment and care for our mission and employees." The fact that Microsoft extended job offers to all OpenAI employees, and that many were prepared to follow Altman to a new venture, demonstrated how deeply the technical workforce aligned with his approach to AI development.⁶¹

The crisis resolution, including Altman's reinstatement and the reorganization of the board, represented an attempt to find a sustainable balance between speed and oversight. Microsoft's gaining of a nonvoting board observer position through Dee Templeton highlighted the formal recognition of the importance of maintaining development momentum while establishing new governance guardrails.⁶²

This speed and market release perspective fundamentally differs from both the stewardship view and pure commercial interests. It represents a belief that rapid development and deployment of AI technologies, when properly managed, serve both the advancement of the technology and the broader mission of ensuring AI benefits humanity. The perspective acknowledges risks but argues that these are best understood and mitigated through practical experience and iterative improvement rather than primarily through theoretical consideration and cautious deployment.

8.3 PRIVATE GOVERNANCE IN SPECIALIZED AI: ALPHAFOLD

AlphaFold is an artificial intelligence program developed by DeepMind, a subsidiary of Alphabet Inc. (Google's parent company).⁶³ It represents a significant advancement in the field of structural biology. AlphaFold's primary function is to predict the three-dimensional structures of proteins based on their amino acid sequences. This capability is a major breakthrough because understanding a protein's structure is crucial for comprehending its function and subsequently utilizing it in applications like drug discovery and vaccine development.

Before AlphaFold, determining protein structures was a time-consuming and complex process, often requiring years of experimental work.⁶⁴ AlphaFold, however, can predict structures with a level of accuracy comparable to laboratory methods,

This constitutional framework functions as a form of self-supervision during training, where the model learns to critique its own outputs according to the established principles. When generating responses, the AI system first produces a draft, then evaluates that draft against constitutional principles, identifies potential violations, and refines its output accordingly.¹¹⁹ This recursive self-improvement process creates an internal deliberative mechanism analogous to human ethical reasoning.

What makes constitutional AI particularly significant from a governance perspective is its transparency and modifiability. Unlike black box safety mechanisms, the constitutional principles can be publicly documented, debated, and revised based on stakeholder input and emerging ethical consensus. This approach received substantial praise for its emphasis on principled restraint rather than rushing to market, with observers such as Open AI's Helen Toner noting that Anthropic's decision to keep their product off the market was a costly signal of restraint, enhancing the credibility of its commitments to AI safety. By building ethical reasoning into the core of model development rather than treating it as an external constraint, constitutional AI represents a paradigm shift in how AI systems can be aligned with human values.

8.4.4.4 The Need for Governance Tools in Hybrid Corporate Structures

Hybrid governance structures like those pioneered by OpenAI and Anthropic represent meaningful attempts to address the unique challenges posed by advanced AI development. By creating formal mechanisms that prioritize broader societal benefit alongside commercial considerations, these structures establish an important foundation for responsible innovation. They signal to stakeholders, employees, and the public a serious commitment to developing AI that serves humanity's interests rather than merely maximizing shareholder returns.

Yet as the analysis earlier demonstrates, these structural innovations function best when complemented by robust operational governance practices. The translation of abstract principles into consistent organizational behavior requires more than well-crafted mission statements or innovative board compositions. It demands the patient, deliberate development of governance mechanisms that guide daily decision-making at all levels of the organization. Red teaming introduces essential external perspectives that counter institutional blindspots, system cards create transparency and accountability around model capabilities and limitations, and approaches like constitutional AI embed ethical reasoning into the technical development process itself.

8.5 CORPORATE ORDERING IN A LARGE CORPORATION: MICROSOFT

For a large corporation like Microsoft, implementing responsible AI governance presents distinct organizational challenges. AI technologies inevitably appear across

multiple products and software applications with widely varying commercial uses and contexts. Maintaining a consistent ethical approach across this diverse portfolio is a significant challenge in itself. Moreover, certain strategic decisions about acceptable uses and technological boundaries must be made centrally, while ensuring alignment among employees with vastly different skillsets and priorities, from engineers focused on innovation to legal teams concerned with compliance and risk.¹²⁰ As CEO Satya Nadella stated, “we learned that we needed to create a governance model that was inclusive and encouraged engineers, researchers, and policy practitioners to work shoulder to shoulder to uphold our AI principles. A single team or a single discipline tasked with responsible or ethical AI was not going to meet our objectives.”¹²¹ This insight drove the development of a comprehensive governance structure at Microsoft that integrates three essential functions: corporate rulemaking, executive decision-making, and continuous monitoring.

The success of Microsoft’s approach stems from its ability to combine a principles-based framework with operational reach across the organization. By establishing clear ethical guidelines through its Responsible AI Standard, Microsoft created a common language for evaluating AI systems while extending implementation through a hub-and-spoke model of AI champions embedded across engineering teams. This structure ensures consistent application of principles while maintaining flexibility for diverse technical contexts. When difficult tradeoffs arise, Microsoft’s Sensitive Uses review process functions as a quasi-judicial system, creating precedents that guide future development with clear escalation paths to the highest levels of the organization.¹²²

The paragraphs below discuss three core components of Microsoft’s governance structure. First, the corporate rulemaking efforts that establish standards, including the development of the Responsible AI Principles and Standard through the Aether Committee. Second, the executive decision-making functions that operationalize these principles through the Office of Responsible AI and its hub-and-spoke network of champions. Finally, the monitoring process is embodied by the Sensitive Uses review, which subjects potentially problematic AI applications to rigorous scrutiny and makes principled choices about tradeoffs between innovation and responsibility.

8.5.1 Microsoft’s AI Strategy: Values as Business Imperatives

In 2016, Microsoft CEO Satya Nadella published a pioneering article in *Slate* titled “The Partnership of the Future,” establishing the company’s early commitment to artificial intelligence development.¹²³ This piece reveals how, years before the current AI boom, Microsoft was already formulating a strategic approach to AI that placed values at its core. Nadella’s vision transcended the simplistic narrative of computers versus humans, instead proposing a future where “human and machine work together to solve society’s greatest challenges such as beating disease, ignorance, and poverty.”

What distinguishes Nadella's approach is his recognition that values are not peripheral considerations but fundamental to AI's business viability. He powerfully cites John Markoff's observation that "The best way to answer the hard questions about control in a world full of smart machines is by understanding the values of those who are actually building these systems." This insight forms the cornerstone of Microsoft's AI philosophy, which posits that technological development and values-based considerations are inseparable, with the latter directly determining market success. Without appropriate value systems guiding development, Nadella implicitly acknowledges that society might reject AI technologies, impose restrictive regulations, or ban applications altogether, effectively limiting AI's commercial potential.

Nadella frames this values-first approach through several interconnected principles: AI must augment rather than replace human capabilities; systems must be transparent and understandable; privacy and security protections must be built-in; and technology must preserve human dignity while maximizing efficiency.¹²⁴ These are not presented merely as ethical ideals but as practical business requirements that will determine whether AI technologies gain widespread acceptance. By emphasizing that AI must be designed with "algorithmic accountability" and against bias, Nadella positions responsible AI as a market imperative rather than a regulatory burden.

The business tradeoffs in AI development are addressed with similar nuance. Nadella acknowledges the tension between automation that displaces workers and technology that enhances human capabilities, suggesting that business leaders must "replace our labor-saving and automation mindset with a maker-and-creation mindset."¹²⁵ This reflects his understanding that AI's commercial potential depends on societal acceptance, which requires addressing legitimate economic concerns. For Nadella, establishing an "ethical and empathic framework" for AI design is not just morally appropriate, but it is first and foremost the essential foundation for realizing AI's business potential and securing Microsoft's position in this emerging market. His approach demonstrates remarkable foresight in recognizing that embedded values would become central to AI's commercial viability, a perspective that has only gained relevance in today's rapidly evolving AI landscape.

8.5.2 *Corporate Rulemaking: Responsible AI Principles and Standard*

Microsoft understood that to translate ethical aspirations into responsible practice, they first needed to identify key principles, and then develop these principles into actionable items. This corporate rulemaking effort began with defining core principles, which then evolved into the Responsible AI Standard, initially introduced in 2018 and later revised in 2022.¹²⁶

Microsoft recognized early that waiting for regulatory guidance was not an option. They needed to take the initiative themselves, not only because regulatory timelines would be impractically slow for their development cycles, but because they needed

to independently determine what the critical ethical issues were in the first place.¹²⁷ This proactive approach led them to establish a deliberate process of soliciting input from external stakeholders.¹²⁸ They actively consulted scholars, researchers, and academics to inform their emerging ethical framework.¹²⁹ They strategically expanded their knowledge base by drawing insights from adjacent communities focused on human rights, privacy, and related fields.

This external consultation phase was crucial in helping Microsoft develop standards that would be both robust and practical, and thus widely accepted and adopted. By channeling this diverse external input into their internal processes, they were able to create guidelines that had legitimacy both within and outside the company. Engineering teams, arguably the most influential group at Microsoft due to their product development capabilities, needed to be convinced of the framework's value and practicality.

As regulators eventually began showing interest in AI governance, Microsoft's approach evolved further. Their early investment in ethical frameworks positioned them to not only anticipate regulatory requirements but also to participate meaningfully in shaping the regulatory process itself.¹³⁰ While legal requirements would only establish a minimum threshold, Microsoft's proactive stance meant they were already operating beyond that floor. This forward-looking perspective has proven valuable, as the threat of regulation has served as a helpful motivator for organizational alignment around responsible AI practices. This methodical progression from conceptual frameworks to operational guidelines represents Microsoft's systematic approach to ensuring their AI technologies align with human values and societal needs.

8.5.2.1 The Microsoft Responsible AI Standard

The foundational premise of Microsoft's strategy was that AI technologies would gain market acceptance if they addressed user concerns and societal implications.¹³¹ Thus, the goal of the rulemaking framework would be to identify and categorize these concerns, so that they form a cohesive set of guidelines for the company's software engineers. The Microsoft Responsible AI Standard operationalizes six core AI principles: accountability, transparency, fairness, reliability and safety, privacy and security, and inclusiveness.¹³² Each principle, as the paragraphs later discuss, seeks to broaden the commercial appeal of the company's products by ensuring they address the concerns that might otherwise limit market acceptance and adoption.

At its heart, the standard recognizes that accountable AI systems require rigorous assessment of potential impacts on different populations.¹³³ Microsoft implements this through structured Impact Assessments during early development, helping teams identify and address potential negative effects before they manifest. This approach reduces legal and reputational risks while positioning Microsoft's products favorably for earning consumer confidence and navigating potential regulatory concerns.

Hand in hand with accountability comes transparency, which allows end users to understand what information forms the basis for their interactions with AI systems. Through Transparency Notes, users gain insight into an AI system's capabilities, limitations, and appropriate uses.¹³⁴ This principle reduces uncertainty about AI operations and enables informed decisions. Commercially, transparent AI systems address concerns about "black box" decision-making, potentially increasing adoption rates.¹³⁵ Businesses are more likely to invest in AI solutions when they understand how these tools make decisions that affect their operations.

The standard places equal emphasis on fairness, which ensures AI systems provide a consistent quality of service across different demographic groups, including marginalized communities. This principle enhances commercial value by expanding the market reach of AI products. Fair systems work effectively for diverse user bases, opening additional market segments.¹³⁶ Moreover, companies with demonstrably fair AI systems are better positioned to win contracts that require equitable treatment of all users, particularly in sectors such as finance, healthcare, and government services.

Reliability and safety form another cornerstone of the framework, focusing on predictable performance even under varied conditions. Microsoft understood that AI systems that fail unexpectedly or operate inconsistently quickly lose market share regardless of their technical sophistication. By emphasizing these qualities, Microsoft positions its AI offerings as dependable business tools rather than experimental technologies, allowing them to command premium pricing in mission-critical applications where reliability justifies higher investment.

Privacy and security concerns have emerged as significant barriers to AI adoption, which Microsoft addresses through strict data governance and management practices.¹³⁷ The standard mandates comprehensive protocols for data collection, processing, and storage that protect user information from misuse.¹³⁸ Beyond addressing user concerns, these practices ensure compliance with global regulations such as GDPR and CCPA, expanding Microsoft's market reach while avoiding costly penalties.¹³⁹

Rounding out the framework, Microsoft emphasizes inclusiveness by actively collaborating with diverse communities throughout the development process.¹⁴⁰ This approach ensures AI systems work effectively across global markets and varied user groups, opening new opportunities and use cases that might otherwise remain unexplored. Products designed with inclusivity from the ground up naturally align with the values of socially conscious consumers and investors, creating additional market advantages.¹⁴¹

The release of Microsoft's Responsible AI Standard in 2022 has garnered significant attention within the industry and press.¹⁴² Commentators have widely recognized it as a leading example of translating high-level ethical principles into enforceable, operational requirements.¹⁴³ The Standard has drawn particular praise for its depth, specificity, and rigorous oversight mechanisms, with many highlighting its emphasis on continuous evaluation, human oversight, and transparency as

setting a benchmark for the industry.¹⁴⁴ Microsoft's approach to fairness, especially in essential domains such as healthcare, finance, and employment, has been noted as a meaningful step toward mitigating algorithmic bias. However, critics have pointed to challenges in implementation, with some industry observers expressing concern about the complexity and resource demands of such a thorough standard, particularly for smaller organizations with limited resources.¹⁴⁵ There is also ongoing debate about enforcement mechanisms, with questions raised about whether external audits or regulatory oversight might ultimately be necessary to ensure compliance industry-wide.¹⁴⁶

8.5.2.2 The Aether Committee as Corporate Rulemaker: A Pioneering Multidisciplinary Brain Trust

Behind Microsoft's development of these AI principles and standards is the Aether Committee, a pioneering body essentially functioning as an internal think tank and policymaker.¹⁴⁷ AETHER stands for AI, Ethics, and Effects in Engineering and Research. This internal advisory body is dedicated to guiding the ethical development and deployment of artificial intelligence technologies.

The Committee consists of approximately forty top executives drawn from research, policy, and engineering teams, creating a multidisciplinary brain trust that combines technical knowledge with policy and ethical expertise. Co-led by Brad Smith, Microsoft's Board Chairman and former chief legal officer, and Kevin Scott, the company's Chief Technology Officer, the Committee meets monthly to submit to ORA progress and coordinate AI initiatives.¹⁴⁸ Their discussions range from evaluating specific high-risk applications to debating broader ethical frameworks, with division leaders sharing implementation updates while subject matter experts in fields ranging from machine learning to human rights law provide guidance. The Committee's influence was evident from its earliest days. It formulated Microsoft's first AI principles, which the company formally adopted in 2018, establishing the foundation for all subsequent governance efforts. The principles represented a quintessential example of corporate rulemaking, establishing uniform, generally applicable standards that would govern all of Microsoft's AI development.

The effort to combine multiple disciplines reflects the understanding that input from all perspectives is necessary for responsible AI development. Aether serves as a cross-disciplinary forum, bringing together experts from technology, ethics, law, and policy to evaluate the societal implications of AI technologies. The committee provides recommendations on policies, processes, and best practices to ensure that AI systems are developed and deployed responsibly.¹⁴⁹ It functions in an advisory capacity to Microsoft's senior leadership and the Office of Responsible AI, influencing decisions on AI-related initiatives and customer engagements.

To address the multifaceted challenges of AI ethics, Aether comprises several specialized working groups, each focusing on specific areas such as fairness and

inclusiveness, reliability and safety, privacy and security, transparency, and accountability.¹⁵⁰ These working groups develop tools, best practices, and implementation guidance to operationalize responsible AI across Microsoft's products and services.

But the institutional experimentation does not stop there. Once the company established executive teams dedicated to implementing the standard, it also designed channels for providing feedback from these teams into the corporate rulemaking effort. The Aether Committee participates in the Responsible AI Council, which convenes representatives from research, policy, engineering, and business. These forums are described as both “challenging and refreshing,” fostering collective wisdom and generating new ideas through the combination of diverse viewpoints.¹⁵¹ The committee's processes are iterative, with ongoing dialogue and feedback loops, allowing for the identification of blind spots, the surfacing of new risks, and the adaptation of policies as technology and societal expectations evolve.

8.5.3 *Executive Decision-Making: The Office of Responsible AI*

With a set of AI principles established by 2019, Microsoft was ready to take the next step in responsible AI governance by creating a dedicated team designed to bring these principles into practice. This team took shape as the Office of Responsible AI (ORA), established in 2019 as the heart of Microsoft's governance structure for responsible AI.¹⁵² The ORA is structured under Corporate External Legal Affairs (CELA), though not all members are practicing lawyers. It has evolved into a hybrid organization combining legal expertise, technical knowledge, and policy experience.¹⁵³ Its staff includes lawyers and compliance professionals alongside engineers, ethicists, and policy experts who can engage meaningfully with development teams on technical questions while maintaining focus on ethical and societal impacts. As of early 2025, Microsoft has invested significantly in this area, with over 400 employees working on responsible AI initiatives, of which approximately half are dedicated full-time to ORA, and the other half dispersed around the company on various teams working on AI.¹⁵⁴

While this might sound like a traditional compliance department, the ORA represents a significant departure in both structure and function. Unlike conventional legal or compliance teams that typically react to problems after they arise, the ORA works proactively with engineering teams during the development process. This required a fundamental shift in how oversight functions operate within the company. A crucial part of the ORA's operation is its daily engagement with engineers to build a compliance platform for all teams to submit their AI uses, partially automating the oversight process. Through this hybrid structure, the ORA effectively combines executive decision-making functions with continuous monitoring to ensure those decisions align with Microsoft's ethical principles and standards.

This institutional hybrid was Microsoft's answer to a key challenge: how to gain buy-in from the company's famed engineering teams, which wield exceptional influence due to their central role in product development. In a company that prizes

innovation above all, presenting a set of principles that not only appears consistent but also evolves alongside products without stifling future developments was crucial. This approach was essential to establishing legitimacy and building the close collaboration necessary to ensure that ethical standards became embedded in technical processes rather than remaining abstract principles.

8.5.3.1 Hub and Spoke: Embedding Responsibility throughout the Company

Microsoft employs a “hub-and-spoke” governance model for responsible AI that seeks to integrate privacy, security, and accessibility into Microsoft’s products and services. This model allows Microsoft to centralize key governance functions while distributing implementation responsibilities throughout the organization.

At the center of this model, the Office of Responsible AI serves as the primary hub, connecting to two other essential components: the Aether Committee, which leverages top scientific and engineering talent for subject-matter expertise, and the Responsible AI Strategy in Engineering (RAISE) group, which enables engineering groups to implement responsible AI processes through systems and tools.¹⁵⁵ This interconnected hub ensures that issues identified by the ORA provide direct input to the company’s policymaking efforts, creating a continuous feedback loop between practical implementation challenges and policy development. Together, these three groups set a consistent bar for responsible AI across the company.

From this central hub, the “spokes” extend outward through a network of “Responsible AI Champions,” that is, senior employees embedded across different engineering and sales teams. These Champions, appointed by company leadership, serve as the eyes and ears of the responsible AI program on the ground. They actively raise awareness about Microsoft’s approach to responsible AI, help teams assess ethical and societal considerations in their specific contexts, and cultivate a culture of responsible innovation within their units.

When engineering teams encounter potential ethical challenges or need to evaluate the implications of new AI applications, Champions serve as their first point of contact, helping to navigate the company’s responsible AI requirements and connecting them with appropriate resources from the ORA. This structure reflects Microsoft’s recognition that effective AI governance requires both centralized expertise and distributed implementation capabilities.

8.5.3.2 Operationalizing “Fit for Purpose” through the Champion Network

The “Fit for Purpose” standard serves as the key mechanism through which engineers interact with the Office of Responsible AI, ensuring their AI systems remain responsible throughout development and deployment.¹⁵⁶ Microsoft defines this

standard as the requirement that AI systems must provide valid solutions to the problems they are designed to solve. This means each system should effectively address its intended use cases, with evidence supporting that it is an appropriate and effective solution for those specific problems.

Microsoft operationalizes “Fit for Purpose” through its Responsible AI Standard and associated Impact Assessment process.¹⁵⁷ Engineering teams must engage with the ORA through this process, explicitly assessing and documenting how their AI systems will solve the problems posed by their intended uses. The documentation process becomes the primary interface between technical teams and the governance structure, creating accountability at every stage of development.

The documentation process flows through a comprehensive and structured Impact Assessment for each AI system. Teams must articulate a clear definition of the problem being solved for each intended use, provide detailed documentation of the model’s inputs and how well they represent the intended concepts and limitations, and document the model’s outputs and their validity in representing the target phenomena. They must also assess the model’s generalizability based on its training and testing data, define performance metrics and thresholds for responsible release, and identify error types and acceptable error rates.

Evidence requirements create another layer of interaction between engineers and the governance framework. Teams engage with multiple stakeholders and research methodologies to gather support for their system’s fitness.¹⁵⁸ This may involve soliciting input from domain experts, gathering results from user studies, focus groups, or surveys, and conducting experiments or systematic research showing the system’s effectiveness compared to alternatives.

To illustrate how this works in practice, consider an AI system designed to allocate hospital beds. The engineering team would work with the ORA through the Champions network to document how the system’s predictions help solve the bed allocation problem. They would gather evidence through pilot studies or expert feedback demonstrating that the system improves efficiency compared to existing methods. Throughout this process, they would acknowledge any limitations in the model’s predictions or generalizability, creating a transparent record of the system’s capabilities and constraints.

This ongoing dialogue between engineers and the ORA is facilitated by the Champions, who guide teams through probing questions that encourage critical thinking about AI systems. Questions like “What is the problem to be solved for this intended use?” and “What evidence will establish that the system is fit for this intended use?” prompt engineers to consider the ethical and practical implications of their work. The Champions ensure that teams justify their design decisions with evidence and stakeholder input, legitimizing AI development through a rigorous review process that the ORA ultimately approves. This continuous interaction between technical development and ethical oversight is what distinguishes Microsoft’s approach to responsible AI governance.

8.5.3.3 From Ground-Level Implementation to Board Oversight

Microsoft's responsible AI governance structure creates a clear line of accountability from day-to-day engineering decisions all the way to board-level oversight. The ORA has a governance and enablement team that connects to the board through two key leaders who embody the cross-disciplinary approach necessary for effective AI governance. Brad Smith, the board chairman and former chief legal officer, represents the legal and strategy perspective, while the Chief Technology Officer, Kevin Scott, represents the engineering perspective. These two executives serve as the primary conduit to the board, ensuring that both ethical and technical considerations inform the highest level of decision-making. Importantly, the Environmental, Social, and Public Policy Committee of the Microsoft Board provides oversight of the responsible AI program, establishing a direct line to board-level governance.¹⁵⁹

Below these executives sits the Responsible AI Council, consisting of approximately forty corporate vice presidents who meet monthly to set the agenda, coordinate efforts, share progress, and serve as division leads. They function as the ORA's eyes and ears on the ground, helping to adopt policies in practice and providing feedback. They also connect with internal audit and other oversight functions to ensure comprehensive coverage of AI governance issues. The Aether Committee also participates in the Responsible AI Council. This governance structure has evolved organically over time in response to the growing importance of AI in Microsoft's business. Each addition to the governance framework has built upon earlier structures, creating a progressively more robust system for ensuring responsible AI development. The Responsible AI Council now serves as a forum for leadership alignment and accountability in implementing Microsoft's responsible AI program across the entire organization.¹⁶⁰

This multilayered approach creates a governance structure that balances centralized expertise with distributed implementation, connecting board-level oversight to day-to-day engineering decisions through multiple layers of engagement. By taking feedback from governments, public policy teams, civil society, and international stakeholders, Microsoft's responsible AI governance demonstrates how corporate oversight can be both internally effective and externally responsive. The result is a sophisticated response to the unique challenges of overseeing AI development and deployment, a governance structure that can adapt to the evolving landscape of AI technology while maintaining a focus on responsible innovation.

8.5.4 Monitoring and Review: Sensitive Uses

Microsoft's governance structure faces its true test in handling sensitive use cases, where competing priorities and complex tradeoffs demand careful navigation.¹⁶¹ The sensitive uses process exemplifies how Microsoft's three governance functions work together in practice: rulemaking through the standards that define sensitive uses,

executive decision-making in determinations about whether to launch specific products, and monitoring through the multilayered review and appeal process for sensitive uses. In 2019, Microsoft founded the program to subject its most sensitive and novel AI use cases to rigorous, specialized review that results in tailored guidance. Since its establishment, the Sensitive Uses team has processed over 900 submissions, with the pace accelerating dramatically to 300 cases in 2023 alone, nearly 70 percent of them related to generative AI. This surge reflects both the rapid advancement of AI technology and Microsoft's expanding efforts to scrutinize its applications.

The paragraphs below describe the review process for identifying and assessing sensitive uses, highlighting that the effort for careful calibration of diverse considerations and thorough documentation lend it a quasi-judicial character. Some landmark examples of that review process in practice illustrate how this process serves the company's business strategy.

8.5.4.1 The Sensitive Uses Review Process: A Quasi-Judicial Approach

The review process itself operates like a sophisticated filtering system, with multiple levels of scrutiny that reach all the way to the top of the organization. It begins when an employee flags a potential sensitive use through a central submission tool. The system captures cases involving AI applications that could deny individuals access to consequential services such as healthcare or financial support, create risks of physical or psychological harm (particularly for vulnerable populations), or infringe on fundamental human rights such as freedom of expression or privacy. This scope, aligned with Microsoft's Responsible AI Standard, ensures that Microsoft identifies not just obvious risks but also subtle ways AI might impact human welfare.

Once submitted, the case is first routed to a local Responsible AI Champion, who serves as an initial assessor. The Champion, working closely with the Office of Responsible AI and the project team, investigates the case details and checks whether similar scenarios have previously been addressed. For familiar situations with established precedent, they can apply existing guidance. But when they encounter novel challenges, as often happens with rapidly evolving technologies such as generative AI, the case moves to the Sensitive Uses Working Group for deeper evaluation.¹⁶²

This working group represents Microsoft's brain trust for ethical AI development, bringing together diverse experts to deliberate on complex cases. Besides technical feasibility, they consider broader societal implications and potential unintended consequences. For particularly thorny issues, or issues where there are disagreements between the engineers and the Office of Responsible AI, cases can be escalated all the way to the Aether Committee itself, ensuring the company's most senior leaders weigh in on precedent-setting decisions. If disagreements persist, the case will go to the CEO, in a clear indication of just how seriously Microsoft takes these issues and how the very top of the organization is involved in responsible AI

governance.¹⁶³ This escalation path from a local champion to working groups to committees and ultimately to the CEO demonstrates Microsoft's commitment to ethical AI at every level of the organization.

The outcome is not simply a yes or no verdict. The Sensitive Uses team provides detailed, practical guidance tailored to each project's specific challenges. These recommendations might include technical safeguards, additional employee training, or modifications to the project's scope. Sometimes, when the risks appear unmanageable, teams are directed to abandon projects entirely. This is a clear indication that Microsoft's commitment to responsible AI development can trump commercial opportunities.

Microsoft's Sensitive Uses Review Program bears a striking resemblance to formal legal proceedings in both structure and function. Like a judicial system, it applies broad principles such as privacy, safety, or human rights to specific technological applications, creating a body of precedent that guides future decisions. The program's multilayered review process mimics the hierarchical structure of courts, with initial assessments by Responsible AI Champions serving as a "first instance" review, the Sensitive Uses Working Group functioning as an appellate body, and the Aether Committee representing a supreme court of sorts, with the CEO as the final authority. This structure ensures a degree of independence from the engineering teams developing the products, similar to how judicial independence separates legal judgment from the interested parties in a case. The rigorous documentation requirements parallel legal briefs and opinions, forcing teams to articulate their reasoning and consider counterarguments. Like court proceedings, the process allows all stakeholders to present their perspectives and concerns, with dissenting voices having a path to escalate their objections. Even the outcome, which consists in detailed guidance with specific requirements rather than simple approval or rejection, resembles judicial remedies that often include specific conditions and limitations. This quasi-judicial approach lends legitimacy to Microsoft's decisions, establishing a governance system that, while internal to the company, mimics the procedural fairness we typically associate with formal legal institutions.

8.5.4.2 Example: Facial Recognition in Law Enforcement

A landmark example of the Sensitive Uses review process in action emerged through Microsoft's handling of facial recognition technology in law enforcement. When faced with requests to deploy real time facial recognition on police body cameras, the case went through this review process and ultimately reached the highest levels of consideration.¹⁶⁴

The outcome might have seemed unthinkable in an earlier era of tech development: Microsoft said no. The company knew this decision would spark intense controversy. While facial recognition technology promised to enhance public safety and aid law enforcement, it also raised serious ethical concerns.¹⁶⁵ The technology

had shown significant accuracy problems, with particularly high misidentification rates for women, young people, and people of color, potentially subjecting these groups to unwarranted police scrutiny. Beyond accuracy issues, the technology's deployment would enable unprecedented levels of public surveillance, allowing for pervasive monitoring of citizens without their consent. There were also valid concerns about potential abuse, as the technology could be used for discriminatory targeting or unauthorized tracking of individuals' movements and activities.

Microsoft determined such applications were premature given both technical limitations and societal implications. The decision extended beyond considerations of technical readiness, reflecting Microsoft's recognition that the deployment of facial recognition technology in law enforcement applications necessitated comprehensive societal dialogue and robust regulatory frameworks.

In the wake of Microsoft's decision, other companies moved to fill the market gap. Most notably, Axon, which had evolved from its origins as a taser manufacturer into a comprehensive law enforcement technology provider, expanded aggressively into the police technology space. By 2024, Axon had further extended its reach by incorporating advanced AI capabilities into its body camera systems, including automated audio summarization powered by LLMs. The contrast between Microsoft's cautious approach and Axon's expansion illustrates how the Sensitive Uses process helps Microsoft navigate complex tradeoffs between commercial opportunities and potential societal harm, even when competitors choose different paths.

Microsoft's stance on facial recognition technology has continued to evolve. As recently as May 2024, the company expanded its prohibition to include "any real-time facial recognition on mobile cameras used by any law enforcement globally," strengthening its previous restrictions.¹⁶⁶ This policy adjustment came shortly after Axon announced a new AI technology utilizing OpenAI's GPT-4 to translate audio captured by body cameras into police reports, demonstrating Microsoft's ongoing vigilance regarding how its AI technologies might be deployed.¹⁶⁷

Microsoft's decision to decline the opportunity to provide facial recognition technology for police body cameras represented a calculated strategic choice. While the company recognized the significant revenue potential in equipping police forces with this technology, Microsoft's leadership made a deliberate assessment that the reputational risks outweighed the short-term financial gains. By saying no to what could have been a lucrative market, Microsoft prioritized its long-term brand positioning as a leader in responsible AI, a differentiator they viewed as increasingly valuable in a competitive tech landscape. This difficult decision reflected Microsoft's understanding that maintaining trust with consumers, enterprise clients, and regulators would ultimately yield greater returns than pursuing every possible AI application regardless of societal impact. Rather than seeing responsible AI governance as merely a compliance function, Microsoft positioned ethical considerations as a core competitive advantage, betting that

their restraint would enhance their credibility when offering AI solutions in more sensitive sectors such as healthcare, finance, and government services where trust is paramount. The decision demonstrated how responsible AI practices had become integrated not just into Microsoft's ethical framework, but into its fundamental business strategy.

8.5.4.3 Azure Custom Neural Voice: Balancing Innovation and Safeguards

Azure AI's Custom Neural Voice represents another important case where Microsoft's Sensitive Uses review process demonstrated its value.¹⁶⁸ This innovative speech technology enables the creation of synthetic voices that sound nearly identical to the original source. This technology is already used by enterprise customers like AT&T and Progressive, with promising applications in education, accessibility, and entertainment. However, the potential for abuse through inappropriate impersonation and listener deception required careful consideration.

Consistent with Microsoft's measured approach for higher-risk AI systems, Custom Neural Voice underwent several Sensitive Uses reviews as it evolved toward broader availability. These reviews led to the adoption of a layered control framework for Custom Neural Voice. Microsoft limited customer access to the service, ensured acceptable use cases were defined and communicated through an application form, implemented speaker consent mechanisms, created specific terms of use, published transparency documentation detailing risks and limitations, and established technical guardrails to help ensure the speaker's active participation when creating a synthetic voice.¹⁶⁹ Through these and other controls, Microsoft is helping protect against misuse while maintaining beneficial uses of the technology. One such beneficial use includes what is known as voice banking, allowing people who are at risk of losing their voice to "bank" their voice for later use by training a synthetic voice model through Custom Neural Voice.¹⁷⁰

8.5.4.4 Fairness in Speech-to-Text Technology

Microsoft's Responsible AI Standard has also been instrumental in addressing fairness issues in speech-to-text technology. This technology can improve individuals' lives, from sending hands-free texts to helping people with hearing loss communicate. Yet an academic study revealed that this technology produced nearly double the error rates for members of some Black and African American communities than for white users.¹⁷¹

These results spurred Microsoft to take a closer look at the AI systems that power its speech-to-text technologies, to ensure they had been sufficiently trained on the rich diversity of speech patterns.¹⁷² The company turned to an expert sociolinguist and expanded its data collection efforts to narrow the performance gap in its solutions. During the process, Microsoft involved stakeholders from outside the

company, such as experts and people from diverse communities. These lessons were invaluable, and have since been incorporated into the Responsible AI Standard to further articulate specific steps teams must take to ensure Microsoft AI systems are designed to provide a similar quality of service for identified demographic groups, including marginalized groups, to help the company and other organizations harness the benefits of these technologies and avoid potential harms in the future.

8.6 SPEED AND STEWARDSHIP: WHAT CAN REGULATION CONTRIBUTE?

The tension between speed and stewardship that defines corporate governance in AI also shapes governmental approaches. States, like corporations, face the same foundational dilemma: whether to privilege innovation and competitiveness or to institutionalize deliberation and restraint. The underlying logic is not uniquely public, as it originates within corporate environments that first grapple with the challenges of deploying high-risk technologies. Governments enter later, adopting the conceptual vocabulary and operational routines already developed by companies to manage AI risk. What distinguishes the public sphere is not its orientation but its instrument. States possess superior tools for coordination, enforcement, and subsidization that enable them to amplify or contain private action.

This section traces how different governmental entities have responded to the dilemma. It begins with the Trump Administration, whose approach to AI regulation exemplifies the speed imperative in its most explicit form. The federal government there uses its distinctive tools to support rapid technological development through deregulation and subsidization, constructing what might be called a state-managed accelerationism. It then turns to California, where lawmakers have built a countermodel organized around transparency, auditing, and public reporting, a model that converges in design and ambition with the European Union's Artificial Intelligence Act (AIA). The juxtaposition highlights how both state and supranational institutions extend, rather than replace, corporate practices of governance: whether accelerating or deliberating, they institutionalize frameworks that first emerged within private enterprise.

8.6.1 *Tools for Boosting Speed: The Trump Administration's Approach*

The Trump Administration exemplifies a government that governs artificial intelligence through the speed imperative: the conviction that technological leadership depends on accelerating development faster than any rival can respond.¹⁷³ In this view, AI policy becomes a competitive instrument of national power rather than a site of public deliberation. The administration operationalizes this imperative through two tools: deregulation and subsidization. The first removes procedural friction, while the second channels public resources toward private acceleration.

Together, they define a model of the government as financier and deregulatory catalyst of technological expansion.

8.6.1.1 Deregulation: Removing Friction

Federal AI policy has evolved rapidly across administrations, reflecting different judgments about how government should relate to technological development. During the Biden Administration, the prevailing emphasis was on stewardship: ensuring that innovation unfolded alongside ethical, labor, and environmental considerations. Executive actions and agency initiatives, particularly from the Office of Science and Technology Policy (OSTP) and the National Institute of Standards and Technology (NIST), sought to create a coordinated framework for responsible AI, centered on transparency, accountability, and risk management.¹⁷⁴ These measures aimed to guide rather than constrain private innovation, using standards and interagency coordination to shape early norms in an emerging field even in the absence of legislation.

The Trump Administration redirected this approach. Through a series of executive orders, it shifted the language of AI policy from stewardship to speed. The most emblematic, *Removing Barriers to American Leadership in Artificial Intelligence*, rescinded many of the prior executive measures and instructed agencies to identify and remove regulatory obstacles to deployment. OSTP and NIST were encouraged to promote adoption and competitiveness. In effect, the Administration redefined the federal role: from framing rules for responsible innovation to advancing the speed of technological progress.

The deregulatory pivot also introduced a cultural dimension. Some executive orders discouraged or eliminated references to diversity, equity, inclusion, and climate change in federal funding and procurement programs. For supporters, this signaled a commitment to ideological neutrality and market focus; for critics, it blurred the boundary between administrative simplification and cultural intervention. What remained consistent across both administrations was the assumption that government could shape AI development indirectly, either by coordinating standards or by accelerating implementation, but not by building an elaborate regulatory apparatus through legislative intervention.¹⁷⁵

8.6.1.2 Subsidization: Financing Speed

While deregulation reduced procedural friction, subsidization sought to build the underlying momentum. The Trump Administration's emphasis on speed did not simply involve removing constraints, but it also reoriented the use of federal resources toward expansion. The idea was that state investment could amplify innovation where private markets alone might move too slowly, using fiscal capacity as a competitive instrument.

This orientation took its most explicit form in *Winning the Race*, a policy document that recast AI strategy as an investment agenda rather than a regulatory framework.¹⁷⁶ Drafted under the supervision of Michael Kratsios and David Sacks, the report proposed that the United States secure technological leadership not primarily by tightening oversight, but by scaling infrastructure and funding. It called for increased federal grants, public–private partnerships, and procurement programs to accelerate AI adoption across industry, defense, and research. Agencies were encouraged to prioritize projects consistent with national competitiveness goals, and to avoid duplicative or “burdensome” oversight efforts. In this way, fiscal coordination replaced administrative regulation as the government’s preferred mode of influence.

The report also framed AI as a domain of strategic competition. It argued that the United States could not lead in the global race without a robust domestic foundation in computing power, semiconductor production, and data infrastructure. Public funding, therefore, was not merely economic support but a tool of industrial resilience and geopolitical positioning. Yet this approach also introduced a degree of centralization. While the document promoted deregulation, it simultaneously concentrated discretion within the Executive Office to decide where and how federal resources should be allocated. The result was a hybrid model, market-driven in rhetoric, state-directed in practice.

Elements of this strategy were codified in the *Big Beautiful Bill*.¹⁷⁷ Although Congress removed its most ambitious deregulatory proposal, a decade-long moratorium on state AI regulation, the final statute institutionalized large-scale federal investment in AI infrastructure. It provided funding for domestic data centers, semiconductor manufacturing, and supply chain expansion, and included restrictions on participation by foreign entities.¹⁷⁸ These provisions linked acceleration to industrial policy, embedding speed within a framework of economic nationalism. By tying eligibility for subsidies to compliance with sourcing and governance requirements, the Big Beautiful Bill blurred the boundary between market autonomy and state direction. Public funds were deployed not to constrain or command but to incentivize. The government selected which sectors and partnerships would advance most rapidly, using capital allocation as its organizing mechanism.

Subsidization also carried a cultural dimension. In documents such as *Winning the Race*, funding was linked not only to technical capacity but also to a broader notion of “ideological neutrality.”¹⁷⁹ Federal support was conditioned on projects and jurisdictions perceived as free from what the administration described as politicized or value-laden approaches to technology, such as references to diversity, equity, or environmental impact. This framing reflected an effort to align technological development with a particular vision of neutrality: innovation unencumbered by social programming. While its practical implications were uneven, the linkage illustrates how fiscal instruments can be used not only to accelerate production but also to signal the cultural and political orientation that public investment is meant to promote.

Subsidization also carried a cultural dimension. In documents such as *Winning the Race*, funding was linked not only to technical capacity but also to a broader notion of “ideological neutrality.” Federal support was conditioned on projects and jurisdictions perceived as free from what the administration described as politicized or value-laden approaches to technology, such as references to diversity, equity, or environmental impact. This framing reflected an effort to align technological development with a particular vision of neutrality: innovation unencumbered by social programming. While its practical implications were uneven, the linkage illustrates how fiscal instruments can be used not only to accelerate production but also to signal the cultural and political orientation that public investment is meant to promote.

8.6.2 *California’s Countermodel and the European Parallel*

California has taken the opposite approach, constructing a deliberative model of AI governance that parallels the European Union’s AIA, as discussed later. The state’s legislative initiatives, most notably SB 53, the Transparency in Frontier AI Act, and the Privacy Protection AI package, reflect an effort to define processes and responsibilities rather than prescribe outcomes. SB 53 requires developers of frontier or foundation models to prepare and publicly disclose safety and security protocols, conduct catastrophic-risk assessments, and report serious incidents to state authorities. It also calls for independent audits and annual updates to risk-management frameworks beginning later in the decade. The accompanying Privacy Protection AI regulations extend this procedural logic to automated decision-making systems, introducing risk assessments, opt-out rights, and audit requirements for significant uses of AI in employment, finance, housing, and healthcare.¹⁸⁰

California’s framework treats review and documentation as the infrastructure of trust. Where the federal model defines efficiency as the removal of delay, California defines safety as the product of testing and revision. Its architecture parallels the EU’s administrative approach to AI governance: both institutionalize foresight through risk assessment, conformity testing, and documentation. The result is an emerging transatlantic convergence around procedural regulation, in contrast to the federal accelerationism that prioritizes immediate deployment. California’s deliberative model thus acts as the domestic counterpart to the EU’s rule-based system: each constructs legitimacy through procedure rather than velocity.

8.6.3 *The EU’s Effort to Regulate AI: Institutionalizing Best Practices*

For many in the EU, a mandatory regulatory regime for AI is a no brainer. According to Dragoș Tudorache, the co-leading negotiator for the European parliament, AI “is a pioneering technology and we need to harness it, which means putting rules in place. Self-due diligence by companies is not enough.”¹⁸¹ First proposed by the

European Commission in April 2021, the AIA introduced a graded system of oversight, classifying AI systems according to the level of risk they pose, from minimal to prohibited. Policymakers initially conceived of it as a form of consumer-safety regulation, requiring developers to test and certify the quality of the data their systems use, and to inform users about the nature and limits of the technology.¹⁸² But this early approach, which resembled product safety law, was quickly overtaken by the pace of innovation. The unveiling of ChatGPT and similar systems in late 2022 transformed the debate, revealing the broad reach of foundation models capable of generating complex text, code, and imagery from vast and largely unfiltered datasets. Concerns that such systems could be used to manipulate human behavior, evade cybersecurity controls, or even design biological weapons led the European Parliament to expand the scope of the Act substantially. Following two years of negotiation, it was approved in late 2023, ushering in the EU's first comprehensive framework for regulating artificial intelligence.

A key feature of the AIA is its explicit prohibition of certain AI practices, alongside its careful distinction between high-risk and lower-risk applications. The prohibited category reflects issues of fundamental rights and human dignity, encompassing activities such as social scoring that could determine access to services, emotion recognition in workplaces or schools, and remote biometric identification in public spaces, except under narrow law-enforcement exceptions. These prohibitions signal a broader European concern that algorithmic systems, when deployed without constraint, can intrude upon privacy, amplify discrimination, or erode individual autonomy.¹⁸³

Lower-risk systems, by contrast, are primarily subject to transparency obligations. Consumers must be able to tell when they are interacting with a machine, when images or voices are synthetically generated, or when outputs draw upon copyrighted materials. The Act's provisions on transparency extend to the developers of foundation models themselves, who must provide summaries of the data used in training and disclose whether it includes protected content. Regulators are expected to develop further guidance for such systems, encouraging companies to adopt voluntary codes of conduct covering human oversight, data accuracy, and fairness in algorithmic design. In this sense, the AIA not only prescribes mandatory standards but also invites a broader process of self-governance aligned with its objectives.¹⁸⁴

The central thrust of the regulatory regime for high-risk systems, however, goes beyond disclosure. It establishes an extensive compliance architecture that binds companies to conduct a detailed assessment of the risks their systems pose, evaluate their likelihood and severity, and adopt proportionate mitigation measures. These requirements operate alongside preexisting legal duties in areas such as data protection, consumer protection, and nondiscrimination, but they add new dimensions that are specifically attuned to the epistemic and moral risks of machine learning.¹⁸⁵ Providers must document every stage of this process, maintain traceability logs, and continuously monitor and update their systems once they are deployed. For larger organizations, compliance will require building internal audit and ethics functions,

designating personnel responsible for AI safety, and sometimes undergoing external conformity assessments administered by independent bodies accredited by national regulators.¹⁸⁶ In cases involving critical uses of biometric or safety-related data, external certification will be mandatory. The European Commission retains the power to update these conformity mechanisms through delegated legislation, ensuring that regulatory standards evolve in parallel with technological capabilities.¹⁸⁷

Enforcement will follow the typical EU model, relying on national supervisory authorities to implement and monitor compliance while coordination and standard-setting occur at the European level. A newly created European Artificial Intelligence Office, supported by a board of representatives from member states and an expert panel, will guide implementation, issue interpretive guidance, and facilitate uniform application across the Union. The office will also oversee the development of shared practices for high-risk and foundation models, balancing enforcement with capacity-building and cooperation among regulators.¹⁸⁸

Criticism of the AIA has focused largely on its costs and on the potential chilling effect on innovation. Some estimates, such as those from the Center for Data Innovation, have placed the cost of compliance for the European economy in the tens of billions, though the European Commission itself expects the burden to be considerably lower. The concern remains, however, that the effort to document, justify, and monitor AI systems could exhaust the limited resources of smaller firms or startups.¹⁸⁹ OpenAI's early statements that the obligations might make it impossible to operate in Europe, captured a broader anxiety among technology entrepreneurs, especially in Germany and France, where fledgling AI sectors face steep competition from American platforms.¹⁹⁰ Yet the Act's supporters argue that such fears overlook the long-term benefits of predictability, safety, and public trust. By imposing clear boundaries on what is acceptable, and by formalizing the obligations of those who design and deploy high-impact systems, the AIA seeks not to suppress innovation but to stabilize it within a framework of accountability.

8.6.4 *From Corporate Ordering to Public Law: Corporate Templates and the EU AIA Act*

The AIA Act's risk-based structure, described earlier, reflects a broader EU aspiration to make technological progress compatible with fundamental rights, as Anu Bradford has argued.¹⁹¹ Yet the Act's method for achieving this balance, through procedures of consultation, documentation, and review, did not emerge in a vacuum. It evolved in dialogue with the practices of leading firms that had already begun to develop internal governance mechanisms for artificial intelligence. The Act can thus be read not only as a legislative innovation but also as an act of institutional borrowing. It translates corporate self-regulation into a coordinated public framework, retaining many of its underlying features while adding enforcement and coherence across jurisdictions.

The governance architectures at Microsoft and DeepMind exemplify this convergence. Both companies had begun to construct procedural systems that mirrored public administrative functions. DeepMind's work on AlphaFold relied on engagement with domain experts, publication of confidence metrics, and staged data release. These steps, while voluntary, resembled the stages of public rulemaking: consultation with stakeholders, deliberation over risks, and publication of clear parameters for responsible use.¹⁹² Microsoft's internal system, anchored in the Office of Responsible AI, the Aether Committee, and a network of responsible AI "champions," performed a similar role on the decision-making and oversight side.¹⁹³ It established structured review, escalation protocols, and documentation duties, producing a framework that internalized many of the features of regulatory administration.

The AIA Act reproduces this tripartite structure – rulemaking, executive decision-making, and oversight – almost intact. Its risk management provisions require consultation with relevant expertise and documentation of design choices; its conformity assessment mechanisms parallel corporate escalation pathways; and its emphasis on independent review finds a counterpart in the requirement for external certification of certain high-risk systems. The procedural grammar is remarkably similar. What changes is the locus of authority. The Act takes the internal governance practices of a few pioneering companies and turns them into a shared baseline for all market participants.

This alignment does not mean the Act lacks significance. By embedding these practices in law, it replaces discretion with obligation and ensures a level of coordination that private governance could not readily achieve on its own. The European AI Office and national supervisory authorities mirror the internal committees and oversight boards of large firms, but they also link them across organizational boundaries, allowing for exchange of information, consistent interpretation, and collective learning. The Act thereby turns an individualized model of responsibility into a public system of mutual accountability.

Still, the continuity between private and public governance is striking. The AIA Act does not create a new model so much as it stabilizes one that already existed in the private sphere. Its contribution lies in extending these practices to firms that might not have adopted them voluntarily and in formalizing coordination among those that have. The distance between self-regulation and state regulation is, in this respect, smaller than often assumed, less a gap than a gradient. Both rely on similar mechanisms of consultation, documentation, and oversight; both seek legitimacy through transparency and reason-giving. The Act's primary intervention is to ensure that what was once a matter of corporate commitment becomes an expectation of all.

In this sense, the AIA Act represents a process of translation rather than rupture. It takes the procedural norms that firms such as Microsoft and DeepMind had already internalized and embeds them in a framework of public enforcement and institutional coordination. The result is a system that preserves the logic of private governance while distributing it more widely and binding it to collective authority.

The line between corporate self-regulation and public regulation, though formally present, becomes thin enough that in practice it often appears continuous.

This alignment does not mean the Act lacks significance. By embedding these practices in law, it replaces discretion with obligation and ensures a level of coordination that private governance could not readily achieve on its own. The European AI Office and national supervisory authorities mirror the internal committees and oversight boards of large firms, but they also link them across organizational boundaries, allowing for exchange of information, consistent interpretation, and collective learning. The Act thereby turns an individualized model of responsibility into a public system of mutual accountability.

Still, the continuity between private and public governance is striking. The AIA Act does not create a new model so much as it stabilizes one that already existed in the private sphere. Its contribution lies in extending these practices to firms that might not have adopted them voluntarily and in formalizing coordination among those that have. The distance between self-regulation and state regulation is, in this respect, smaller than often assumed, less a gap than a gradient. Both rely on similar mechanisms of consultation, documentation, and oversight; both seek legitimacy through transparency and reason-giving. The Act's primary intervention is to ensure that what was once a matter of corporate commitment becomes an expectation of all.

In this sense, the AIA Act represents a process of translation rather than rupture. It takes the procedural norms that firms such as Microsoft and DeepMind had already internalized and embeds them in a framework of public enforcement and institutional coordination. The result is a system that preserves the logic of private governance while distributing it more widely and binding it to collective authority. The line between corporate self-regulation and public regulation, though formally present, becomes thin enough that in practice it often appears continuous.

8.7 CONCLUSION: BUILDING AI GOVERNANCE FROM WITHIN

The story of AI governance to date is not one of regulatory foresight but of private initiative, compelled less by idealism than by necessity. It was the magnitude of what AI might become, its potential to transform industries, reshape societies, and pose existential risks, that pushed private firms to act. What began as a pragmatic response to technological power is evolving into something far more ambitious: an attempt to build governance from the inside out. These companies came to understand that building trustworthy AI required building the procedural and institutional infrastructure to oversee it, treating governance not as a compliance task but as a core function. They recognized that without internal rulemaking, formalized decision rights, and transparent mechanisms for dissent and review, their principled aspirations would remain rhetorical, valuable, but ultimately unenforceable.

These efforts emerged in response to real tensions within organizations, especially the growing divergence of views among investors, boards, and technical leaders.

145. Tovin Lapan, “How Salesforce wants to make its supply chain more sustainable,” *Fortune*, June 22, 2021, <https://fortune.com/2021/06/22/salesforce-supply-chain-sustainability/> (last accessed October 9, 2025).
146. Salesforce, “Reducing value chain emissions by enabling supplier action,” 2021, <https://relayto.com/explore/reducing-value-chain-emissions-by-enabling-supplier-action-vquyovrv5nx334/qwf7urGh1> (last accessed October 21, 2025).
147. Science Based Targets initiative, “Supplier engagement case study – Salesforce,” <https://sciencebasedtargets.org/companies-taking-action/case-studies/supplier-engagement-case-study-salesforce> (last accessed January 8, 2025).
148. Salesforce, “Supplier sustainability,” www.salesforce.com/company/supplier-sustainability/ (last accessed October 21, 2025).
149. Salesforce, “Reducing value chain emissions by enabling supplier action,” 2021, www.salesforce.com/content/dam/web/en_us/www/documents/white-papers/salesforce-emissions-supplier-enablement.pdf (last accessed October 9, 2025).
150. Salesforce, “FY24 stakeholder impact report,” 2024, www.salesforce.com/en-us/wp-content/uploads/sites/4/documents/white-papers/salesforce-fy24-stakeholder-impact-report.pdf (last accessed October 21, 2025).
151. Salesforce, “Sustainability Exhibit,” 4.
152. Salesforce, “FY24 stakeholder impact report,” 12.
153. Salesforce, “Schedules of selected environmental and social metrics,” 2024, www.salesforce.com/en-us/wp-content/uploads/sites/4/documents/white-papers/salesforce-fy24-stakeholder-impact-report.pdf (last accessed October 21, 2025).

8 HOW TO GOVERN AI?

1. Interview with the author, October 10, 2023, notes on file with the author.
2. Bernard Marr, “A short history of ChatGPT: How we got to where we are today,” *Forbes*, 19 May 2023, www.forbes.com/sites/bernardmarr/2023/05/19/a-short-history-of-chatgpt-how-we-got-to-where-we-are-today/ (last accessed September 22, 2025).
3. Sarah Kessler and Ephrat Livni, “The optimist’s guide to artificial intelligence and work,” *The New York Times*, May 20, 2023, www.nytimes.com/2023/05/20/business/dealbook/the-optimists-guide-to-artificial-intelligence-and-work.html (last accessed September 22, 2025).
4. Cade Metz, “What exactly are the dangers posed by A.I.?” *The New York Times*, May 1, 2023, www.nytimes.com/2023/05/01/technology/ai-problems-danger-chatgpt.html (last accessed September 22, 2025).
5. Emmet Lyons, “DeepSeek AI raises national security concerns, U.S. officials say,” *CBS News*, January 29, 2025, www.cbsnews.com/news/deepseek-ai-raises-national-security-concerns-trump/ (last accessed September 22, 2025).
6. Interview with the author.
7. Interview with the author.
8. Jonathan B. Wiener, “The regulation of technology, and the technology of regulation” (2004) 26 *Technology in Society* 483–500.
9. Kevin Wei et al., “How do AI companies ‘fine-tune’ policy? Examining regulatory capture in AI governance” (2024) *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society* 1539–55 (observing that in rapidly evolving fields like AI, regulators struggle to keep pace and often rely on industry input, heightening risks of regulatory capture).

10. US Department of Defense, “Department of Defense releases the President’s Fiscal Year 2025 defense budget,” March 11, 2024, www.defense.gov/News/Releases/Release/Article/3703410/departement-of-defense-releases-the-presidents-fiscal-year-2025-defense-budget/ (last accessed September 22, 2025).
11. Claudio Novelli et al., “A robust governance for the AI Act: AI Office, AI Board, Scientific Panel, and national authorities” (2025) 16 *European Journal of Risk Regulation* 566–90 (reporting an example of the conventional regulatory model, establishing compliance obligations, disclosure duties, and supervisory bodies for regulated AI systems).
12. Daniel E. Walters and Hannah J. Wiseman, “Self-regulation in emerging and innovative industries” (2025) 62 *Houston Law Review* 543–610 (documenting an example of the self-regulatory model, showing how companies can form industry-funded bodies to set standards and incorporate expert input).
13. Dario Amodei, “Anthropic CEO: Don’t let AI companies off the hook,” *The New York Times*, June 5, 2025, www.nytimes.com/2025/06/05/opinion/anthropic-ceo-regulate-transparency.html (last accessed September 22, 2025).
14. Judd Rosenblatt, “AI is learning to escape human control,” *The Wall Street Journal*, June 1, 2025, www.wsj.com/opinion/ai-is-learning-to-escape-human-control-technology-model-code-programming-o66b3ec5 (last accessed September 22, 2025).
15. Stuart Russell, “Provably beneficial artificial intelligence,” in *The Next Step: Exponential Life* (Madrid: BBVA OpenMind, 2017), pp. 1–15.
16. Bill Gates, “The risks of AI are real but manageable,” *Gates Notes*, July 11, 2023, www.gatesnotes.com/meet-bill/tech-thinking/reader/the-risks-of-ai-are-real-but-manageable (last accessed September 29, 2025); Marco Quiroz-Gutierrez, “Elon Musk says there’s a 10%–20% chance AI ‘goes bad,’” *Fortune*, October 30, 2024, <https://fortune.com/2024/10/30/elon-musk-ai-could-go-bad-existential-threat-xai-fundraising/> (last accessed September 29, 2025).
17. Damita Menezes, “AI will be smarter than humans in 20 years, ‘godfather of AI’ warns,” *NewsNation*, August 4, 2025, www.newsnationnow.com/cuomo-show/ai-humans-20-years-godfather-of-ai-warns/ (last accessed September 29, 2025).
18. Maanak Gupta et al., “From ChatGPT to ThreatGPT: Impact of generative AI in cybersecurity and privacy” (2023) 11 *IEEE Access* 80218–45.
19. Daniel Thilo Schroeder et al., “How malicious AI swarms can threaten democracy,” *OSF Preprints*, 2025, https://doi.org/10.31219/osf.io/qm9yk_v2 (last accessed October 10, 2025) (arguing that coordinated AI systems can fabricate messages, infiltrate communities, and manipulate public discourse at scale, posing systemic risks to democracy).
20. U.S. Department of Homeland Security, “DHS publishes guidelines and report to secure critical infrastructure and weapons of mass destruction from AI-related threats,” April 29, 2024, www.dhs.gov/news/2024/04/29/dhs-publishes-guidelines-and-report-secure-critical-infrastructure-and-weapons-mass (last accessed September 29, 2025).
21. Darrell M. West, “How AI can enable public surveillance,” *Brookings*, April 15, 2025, www.brookings.edu/articles/how-ai-can-enable-public-surveillance/ (last accessed September 30, 2025).
22. Charlie Savage, “Spy agency memo sets rules for artificial intelligence and Americans’ private data,” *The New York Times*, November 14, 2024, www.nytimes.com/2024/11/14/us/ai-privacy-guidelines-intelligence.html (last accessed September 30, 2025).
23. Jordyn Holman and Jeanna Smialek, “Will A.I. boost productivity? Companies sure hope so,” *The New York Times*, April 1, 2024, www.nytimes.com/2024/04/01/business/economy/artificial-intelligence-productivity.html (last accessed September 30, 2025).
24. Noam Scheiber, “Which workers will A.I. hurt most: The young or the experienced?” *The New York Times*, July 7, 2025, www.nytimes.com/2025/07/07/business/ai-job-cuts.html (last accessed September 30, 2025).

25. Alexa Mikhail, “Microsoft claims its AI tool can diagnose complex medical cases four times more accurately than doctors,” *Fortune*, July 3, 2025, <https://fortune.com/2025/07/03/microsoft-ai-diagnostic-orchestrator-diagnosing-complex-cases-medical-doctors/> (last accessed September 30, 2025); Noam Scheiber, “Will A.I. upend white-collar work? Consider the Hollywood editor,” *The New York Times*, July 30, 2024, www.nytimes.com/2024/07/30/business/economy/artificial-intelligence-hollywood-unions.html (last accessed September 30, 2025).
26. Jake Coyle, “In Hollywood writers’ battle against AI, humans win (for now),” *Associated Press*, September 27, 2023, <https://apnews.com/article/hollywood-ai-strike-wga-artificial-intelligence-39ab72582c3a15f77510c9c30a45ffc8> (last accessed September 30, 2025).
27. Molly Kinder et al., “Generative AI, the American worker, and the future of work,” *Brookings*, October 10, 2024, www.brookings.edu/articles/generative-ai-the-american-worker-and-the-future-of-work/ (last accessed September 30, 2025).
28. Adam Satariano and Paul Mozur, “The global A.I. divide,” *The New York Times*, June 21, 2025, www.nytimes.com/interactive/2025/06/23/technology/ai-computing-global-divide.html (last accessed September 30, 2025).
29. Sheera Frenkel, “The militarization of Silicon Valley,” *The New York Times*, August 4, 2025, www.nytimes.com/2025/08/04/technology/google-meta-openai-military-war.html (last accessed September 30, 2025).
30. Lisa D. Ellis, “The benefits of the latest AI technologies for patients and clinicians,” *Harvard Medical School*, August 30, 2024, <https://learn.hms.harvard.edu/insights/all-insights/benefits-latest-ai-technologies-patients-and-clinicians> (last accessed September 30, 2025).
31. Kevin Lynch, “The race for AI supremacy: A new frontier of international competition,” *Forbes*, July 28, 2025, www.forbes.com/councils/forbestechcouncil/2025/07/28/the-race-for-ai-supremacy-a-new-frontier-of-international-competition/ (last accessed September 30, 2025).
32. Tripp Mickle, “Nvidia says U.S. will restrict sales of more of its A.I. chips to China,” *The New York Times*, April 15, 2025, www.nytimes.com/2025/04/15/technology/nvidia-h2o-chip-china-restrictions.html (last accessed September 30, 2025).
33. Alec Tyson and Emma Kikuchi, “Growing public concern about the role of artificial intelligence in daily life,” *Pew Research Center*, August 28, 2023, www.pewresearch.org/short-reads/2023/08/28/growing-public-concern-about-the-role-of-artificial-intelligence-in-daily-life/ (last accessed September 30, 2025).
34. David Goldman, “How OpenAI so royally screwed up the Sam Altman firing,” *CNN*, November 20, 2023, www.cnn.com/2023/11/19/tech/sam-altman-open-ai-firing-board (last accessed September 30, 2025).
35. Matt O’Brien, “ChatGPT-maker OpenAI fires CEO Sam Altman, the face of the AI boom, for lack of candor with company,” *Associated Press*, November 18, 2023, <https://apnews.com/article/altman-ai-chatgpt-murati-893e4a460c10eb3a8fiafe6156eca3> (last accessed September 30, 2025).
36. OpenAI, “OpenAI announces leadership transition,” November 17, 2023, <https://openai.com/index/openai-announces-leadership-transition/> (last accessed September 30, 2025).
37. Matt O’Brien and Haleluya Hadero, “OpenAI brings back Sam Altman as CEO just days after his firing unleashed chaos,” *Associated Press*, November 22, 2023, <https://apnews.com/article/altman-openai-chatgpt-31187f7f6eca8ff9d0eef7585aac6ace> (last accessed September 30, 2025).
38. OpenAI, “Our structure,” 2023, <https://openai.com/our-structure> (last accessed September 30, 2025); OpenAI, “Introducing OpenAI,” December 11, 2015, <https://openai.com/index/introducing-openai/> (last accessed September 30, 2025).

39. Ty Roush, “More than 700 OpenAI employees threaten to quit – and join Microsoft – unless board resigns,” *Forbes*, November 20, 2023, www.forbes.com/sites/tylerroush/2023/11/20/more-than-500-openai-employees-threaten-to-quit-over-sam-altmans-removal/ (last accessed September 30, 2025).
40. Deepa Seetharaman and Tom Dotan, “Sam Altman rejoins OpenAI board along with three new directors,” *The Wall Street Journal*, March 8, 2024, www.wsj.com/tech/ai/openai-board-brings-sam-altman-back-adds-three-new-members-bbdb4807 (last accessed September 30, 2025).
41. Medium, “OpenAI in the context of steward-ownership – blog post 1 on the ‘OpenAI saga’,” December 20, 2023, https://medium.com/@purpose_network/openai-in-the-context-of-steward-ownership-blog-post-1-on-the-openai-saga-dd07feageecd (last accessed September 30, 2025).
42. OpenAI, “Helen Toner joins OpenAI’s board of directors,” September 8, 2021, <https://openai.com/index/helen-toner-joins/> (last accessed September 30, 2025).
43. Andrew Imbrie, Owen J. Daniels, and Helen Toner, “Decoding intentions: artificial intelligence and costly signals,” Center for Security and Emerging Technology, 2023, <https://cset.georgetown.edu/wp-content/uploads/CSET-Decoding-Intentions.pdf> at 27 (last accessed September 30, 2025).
44. Cade Metz, Tripp Mickle, and Mike Isaac, “Before Altman’s ouster, OpenAI’s board was divided and feuding,” *The New York Times*, November 21, 2023, www.nytimes.com/2023/11/21/technology/openai-altman-board-fight.html (last accessed September 30, 2025).
45. Meghan Bobrowsky and Deepa Seetharaman, “The OpenAI board member who clashed with Sam Altman shares her side,” *The Wall Street Journal*, December 7, 2023, www.wsj.com/tech/ai/helen-toner-openai-board-2e4031ef (last accessed September 30, 2025).
46. Jemma Green, “Elon Musk’s lawsuit against OpenAI sparks debate over AI transparency,” *Forbes*, March 7, 2024, www.forbes.com/sites/jemmagreen/2024/03/07/elon-musks-lawsuit-against-openai-sparks-debate-over-ai-transparency/ (last accessed September 30, 2025).
47. Adi Robertson, “Elon Musk sues OpenAI for abandoning its mission to benefit humanity,” *The Verge*, March 1, 2024, www.theverge.com/2024/3/1/24087473/elon-musk-openai-lawsuit-nonprofit-mission (last accessed September 30, 2025); *Elon Musk v. OpenAI et al.*, *Complaint* (Superior Court of California, County of San Francisco, February 29, 2024), www.courthousenews.com/wp-content/uploads/2024/02/musk-v-altman-openai-complaint-sf.pdf (last accessed September 30, 2025).
48. Kyle Wiggers, “Most of OpenAI’s employees threaten to quit if Sam Altman isn’t reappointed CEO,” *TechCrunch*, November 20, 2023, <https://techcrunch.com/2023/11/20/most-of-openais-employees-threaten-to-quit-if-sam-altman-isnt-reappointed-ceo/> (last accessed September 30, 2025).
49. Cade Metz, “OpenAI’s chief scientist and co-founder is leaving the company,” *The New York Times*, May 14, 2024, www.nytimes.com/2024/05/14/technology/ilya-sutskever-leaving-openai.html (last accessed September 30, 2025).
50. Gerrit De Vynck and Nitasha Tiku, “AI execs used to beg for regulation. Not anymore,” *The Washington Post*, May 8, 2025, www.washingtonpost.com/technology/2025/05/08/altman-congress-openai-regulation/ (last accessed September 30, 2025).
51. Sam Altman, “Sam Altman on the future of AI and humanity,” TED, January 7, 2025, www.ted.com/pages/sam-altman-on-the-future-of-ai-and-humanity-transcript (last accessed September 30, 2025).
52. Sam Altman, “The intelligence age,” September 23, 2024, <https://ia.samaltman.com/> (last accessed September 30, 2025).

53. OpenAI, “OpenAI LP,” March 11, 2019, <https://openai.com/index/openai-lp/> (last accessed September 30, 2025).
54. Livemint, “Three senior OpenAI researchers resign after CEO Sam Altman ousted,” November 18, 2023, www.livemint.com/ai/artificial-intelligence/sam-altman-ousted-three-senior-openai-researchers-resign-11700299676106.html (last accessed September 30, 2025).
55. Andrej Karpathy, X, November 19, 2023, <https://x.com/karpathy/status/1726289070345855126> (last accessed September 30, 2025).
56. Cade Metz and Karen Weise, “Microsoft to invest \$10 billion in OpenAI, the creator of ChatGPT,” The New York Times, January 23, 2023, www.nytimes.com/2023/01/23/business/microsoft-chatgpt-artificial-intelligence.html (last accessed September 30, 2025).
57. Microsoft, “OpenAI forms exclusive computing partnership with Microsoft to build new Azure AI supercomputing technologies,” July 22, 2019, <https://news.microsoft.com/source/2019/07/22/openai-forms-exclusive-computing-partnership-with-microsoft-to-build-new-azure-ai-supercomputing-technologies/> (last accessed September 30, 2025).
58. Shanthi Rexaline, “Microsoft, investors stand solidly behind ‘once in a generation’ CEO Sam Altman, putting OpenAI board under pressure,” Benzinga, November 19, 2023, www.benzinga.com/analyst-ratings/analyst-color/23/11/35865361/microsoft-investors-stand-solidly-behind-once-in-a-generation-ceo-sam-altman-puttin (last accessed September 30, 2025).
59. Jacob Kastrenakes and Tom Warren, “Hundreds of OpenAI employees threaten to resign and join Microsoft,” The Verge, November 20, 2023, www.theverge.com/2023/11/20/23968988/openai-employees-resignation-letter-microsoft-sam-altman (last accessed September 30, 2025).
60. OpenAI Employees, “Letter to the OpenAI board,” DocumentCloud, November 20, 2023, www.documentcloud.org/documents/24172377-letter-to-the-openai-board (last accessed September 30, 2025).
61. Tom Gerken, “Microsoft offers to match pay of all OpenAI staff,” BBC, November 21, 2023, www.bbc.com/news/technology-67484455 (last accessed September 30, 2025).
62. Dina Bass and Rachel Metz, “Microsoft picks Dee Templeton as OpenAI board observer,” Bloomberg Law, January 5, 2024, <https://news.bloomberglaw.com/artificial-intelligence/microsoft-picks-dee-templeton-as-openai-board-observer> (last accessed September 30, 2025).
63. Jeremy Khan, “DeepMind spins out new Alphabet company focused on drug discovery,” Fortune, November 4, 2021, <https://fortune.com/2021/11/04/deepmind-spins-out-alphabet-company-isomorphic-drug-discovery-company/> (last accessed September 30, 2025).
64. John Jumper et al., “Highly accurate protein structure prediction with AlphaFold” (2021) 596 *Nature* 583–89.
65. Protein Structure Prediction Center, “CASP14 press release,” December 1, 2020, https://predictioncenter.org/casp14/doc/CASP14_press_release.html (last accessed September 30, 2025).
66. DeepMind, “AlphaFold: A solution to a 50-year-old grand challenge in biology,” Google DeepMind Blog, November 30, 2020, <https://deepmind.google/discover/blog/alphafold-a-solution-to-a-50-year-old-grand-challenge-in-biology/> (last accessed September 30, 2025).
67. Madhumita Murgia, “DeepMind research cracks structure of almost every known protein,” Financial Times, July 28, 2022, www.ft.com/content/6a088953-66d7-48db-b61c-79005a0a351a (last accessed September 30, 2025).
68. Anastassis Perrakis and Titia K. Sixma, “AI revolutions in biology: The joys and perils of AlphaFold” (2021) 28 *Nature Structural & Molecular Biology* 649–50.

69. Thomas C. Terwilliger et al., “AlphaFold predictions are valuable hypotheses and accelerate but do not replace experimental structure determination” (2024) 21 *Nature Methods* 13–24.
70. Priscila S. F. C. Gomes, Diego E. B. Gomes, and Rafael C. Bernardi, “Protein structure prediction in the era of AI: Challenges and limitations when applying to *in silico* force spectroscopy” (2022) 2 *Frontiers in Bioinformatics* 983306.
71. Derek Lowe, “AlphaFold’s place in the world,” *Science*, December 11, 2023, www.science.org/content/blog-post/alphafold-s-place-world (last accessed September 30, 2025).
72. Oleg Kovalevskiy, Juan Mateos-Garcia, and Kathryn Tunyasuvunakool, “AlphaFold two years on: Validation and impact” (2024) 121 *Proceedings of the National Academy of Sciences* e2315002121.
73. Pankaj Malik et al., “Utilizing AlphaFold predictions to understand protein function and dynamics,” ResearchGate, January 2024, www.researchgate.net/publication/377895875_utilizing_alphafold_predictions_to_understand_protein_function_and_dynamics (last accessed September 30, 2025).
74. Ali Burak Ünal, Nico Pfeifer, and Mete Akgün, “A privacy-preserving approach for cloud-based protein fold recognition” (2024) 5 *Patterns* 101023.
75. Ewen Callaway, “Could AI-designed proteins be weaponized? Scientists lay out safety guidelines,” *Nature*, March 8, 2024, www.nature.com/articles/d41586-024-00699-0 (last accessed September 30, 2025).
76. Brendan Bordelon, “The fight over AI biosecurity risk takes a twist,” *Politico*, February 6, 2024, www.politico.com/newsletters/digital-future-daily/2024/02/06/the-fight-over-ai-biosecurity-risk-takes-a-twist-00139945 (last accessed September 30, 2025).
77. Andriy Kryshchak et al., “Critical Assessment of Methods of Protein Structure Prediction (CASP) – round XIII” (2019) 27 *Proteins: Structure, Function, and Bioinformatics* 323–38.
78. Koray Kavukcuoglu et al., “How our principles helped define AlphaFold’s release,” Google DeepMind Blog, September 14, 2022, <https://deepmind.google/discover/blog/how-our-principles-helped-define-alphafolds-release/> (last accessed September 30, 2025).
79. Andrew W. Senior et al., “Improved protein structure prediction using potentials from deep learning” (2020) 577 *Nature* 706–10.
80. EMBL-EBI, “About us,” www.ebi.ac.uk/about/ (last accessed September 30, 2025).
81. Mihaly Varadi et al., “AlphaFold protein structure database: massively expanding the structural coverage of protein-sequence space with high-accuracy models” (2022) 50 *Nucleic Acids Research* D439–44.
82. Anne Trafton, “Analyzing the potential of AlphaFold in drug discovery,” MIT News, September 6, 2022, <https://news.mit.edu/2022/alphafold-potential-protein-drug-0906> (last accessed September 30, 2025).
83. Derek Lowe, “Why AlphaFold won’t revolutionize drug discovery,” *Chemistry World*, August 5, 2022, www.chemistryworld.com/opinion/why-alphafold-wont-revolutionise-drug-discovery/4016051.article (last accessed September 30, 2025).
84. Perrakis and Sixma, “AI revolutions in biology.”
85. Olivia Shone, “5 key features and benefits of large language models,” Microsoft Blog, October 9, 2024, www.microsoft.com/en-us/microsoft-cloud/blog/2024/10/09/5-key-features-and-benefits-of-large-language-models/ (last accessed September 30, 2025).
86. Cade Metz and Karen Weise, “A.I. is getting more powerful, but its hallucinations are getting worse,” *The New York Times*, May 5, 2025, www.nytimes.com/2025/05/05/technology/ai-hallucinations-chatgpt-google.html (last accessed September 30, 2025).

87. Muhammad Muzammil Rawjan, “The ethical AI imperative: How OpenAI is leading the way in responsible development,” RT Insights, February 21, 2024, www.rtinights.com/the-ethical-ai-imperative-how-openai-is-leading-the-way-in-responsible-development/ (last accessed September 30, 2025).
88. “Amoral drift in AI corporate governance” (2025) 138 *Harvard Law Review* 1633–56 at 1634.
89. Natalia Garrido-Skurkowicz, Rafael Wittek, and Chaïm la Roi, “Performance of hybrid organizations: Challenges and opportunities for social and commercial enterprises” (2022) 15 *Journal of Social Entrepreneurship* 1058–87 (drawing on empirical evidence to examine the governance challenges and tensions hybrid organizations face in balancing social and commercial logics).
90. Devin Coldewey, “OpenAI shifts from nonprofit to ‘capped-profit’ to attract capital,” TechCrunch, March 11, 2019, <https://techcrunch.com/2019/03/11/openai-shifts-from-nonprofit-to-capped-profit-to-attract-capital/> (last accessed September 30, 2025).
91. Microsoft, “Microsoft and OpenAI extend partnership,” January 23, 2023, <https://blogs.microsoft.com/blog/2023/01/23/microsoftandopenaiextendpartnership/> (last accessed September 30, 2025).
92. Connie Loizos, “Sam Altman’s leap of faith,” TechCrunch, May 18, 2019, <https://techcrunch.com/2019/05/18/sam-altmans-leap-of-faith/> (last accessed September 30, 2025).
93. Sigal Samuel, “OpenAI as we knew it is dead,” Vox, September 26, 2024, www.vox.com/future-perfect/374275/openai-just-sold-you-out (last accessed September 30, 2025).
94. Gary Marcus, “OpenAI’s Sam Altman is becoming one of the most powerful people on earth. We should be very afraid,” The Guardian, August 3, 2024, www.theguardian.com/technology/article/2024/aug/03/open-ai-sam-altman-chatgpt-gary-marcus-taming-silicon-valley (last accessed September 30, 2025).
95. Robert Hart, “Sam Altman cut from OpenAI’s startup fund – here’s what the ChatGPT maker invests in,” Forbes, April 2, 2024, www.forbes.com/sites/roberthart/2024/04/02/sam-altman-cut-from-openais-startup-fund-heres-what-the-chatgpt-maker-invests-in/ (last accessed September 30, 2025).
96. Keach Hagey, “The secrets and misdirection behind Sam Altman’s firing from OpenAI,” *The Wall Street Journal*, March 28, 2025, www.wsj.com/tech/ai/the-real-story-behind-sam-altman-firing-from-openai-efd51a5d (last accessed September 30, 2025).
97. Hagey, “The secrets and misdirection behind Sam Altman’s firing from OpenAI.”
98. Wired, “More than 730 OpenAI staff threaten to quit after Sam Altman ouster,” November 20, 2023, www.wired.com/story/openai-staff-walk-protest-sam-altman (last accessed September 30, 2025).
99. Max Zahn, “Sam Altman hired by Microsoft, 600 OpenAI employees threaten to quit in protest of his ouster,” ABC News, November 20, 2023, <https://abcnews.go.com/Business/sam-altman-hired-microsoft-600-openai-employees-threaten/story?id=104942311> (last accessed September 30, 2025).
100. Ellen Glover, “Anthropic: What we know about the company behind Claude AI,” Built In, November 12, 2025, <https://builtin.com/artificial-intelligence/anthropic-ai> (last accessed November 15, 2025).
101. Anthropic, “The long-term benefit trust,” May 29, 2023, www.anthropic.com/news/the-long-term-benefit-trust (last accessed September 30, 2025).
102. Billy Perrigo, “How Anthropic designed itself to avoid OpenAI’s mistakes,” Time, May 30, 2024, <https://time.com/6983420/anthropic-structure-openai-incentives> (last accessed September 30, 2025).

103. Alex Kantrowitz, “After OpenAI’s chaos, Anthropic has an opportunity (and its own untraditional board),” *Big Technology*, December 1, 2023, www.bigtechnology.com/pl/after-openais-chaos-anthropic-has (last accessed September 30, 2025).
104. Bechuk and Tallarita, “The illusory promise of stakeholder governance,” 100–05 (arguing that broad stakeholder commitments lack operational content and fail to provide managers with concrete guidance for resolving trade-offs).
105. Krystal Hu, “OpenAI removes Sam Altman’s ownership of its startup fund,” *Reuters*, April 1, 2024, www.reuters.com/technology/openai-removes-sam-altmans-ownership-its-startup-fund-2024-04-01 (last accessed October 1, 2025).
106. Anca Dragan et al., “Taking a responsible path to AGI,” *DeepMind Blog*, April 2, 2025, <https://deepmind.google/blog/taking-a-responsible-path-to-agi/> (last accessed September 30, 2025).
107. Jonas Schuett et al., “Towards best practices in AGI safety and governance: A survey of expert opinion,” *arXiv*, May 11, 2023, <https://doi.org/10.48550/arXiv.2305.07153> (last accessed September 30, 2025).
108. Anne-Claire Pache, Julie Battilana, and Channing Spencer, “An integrative model of hybrid governance: The role of boards in helping sustain organizational hybridity” (2024) 67(2) *Academy of Management Journal* 437–67 (arguing that sustaining hybrid organizations requires governance mechanisms beyond managerial discretion, as boards must actively manage trade-offs between competing logics).
109. Fernando Delgado et al., “Stakeholder participation in AI: Beyond ‘add diverse stakeholders and stir’,” *arXiv*, November 1, 2021, <https://arxiv.org/abs/2111.01122> (last accessed September 30, 2025) (arguing that effective governance requires not just including diverse stakeholders but also structuring deliberation, negotiation, and communication about trade-offs).
110. *Reuters*, “OpenAI responds to warnings of self-governance by former board members, the Economist Reports,” May 30, 2024, www.reuters.com/technology/openai-responds-warnings-self-governance-by-former-board-members-economist-2024-05-30 (last accessed October 1, 2025).
111. Anup Banerjee, Martin Carlsson-Wall, and Mattias Nordqvist, “Hybrid board governance: Exploring the challenges in implementing social impact measurements” (2024) 56(5) *The British Accounting Review* 101359 (arguing that hybrid governance structures can provide space for social mission but struggle with implementation, measurement, and accountability); Liliana Ávila and Marlene Amorim, “Operational challenges in hybrid organizations: Insights for future research,” in João Reis, Marlene Amorim, Nuno Melão, and Paulo Matos (eds.), *Operations and Service Management* (Hershey, PA: IGI Global, 2017) 955–73 (discussing how hybrid organizations face inherent tensions in aligning social and commercial goals, creating complexity in translating missions into practice).
112. Qinghua Lu et al., “Responsible AI pattern catalogue: A collection of best practices for AI governance and engineering” (2024) 56(7) *ACM Computing Surveys* Art 173, 1–35.
113. Jacob Metcalf and Ranjit Singh, “Scaling up mischief: Red-teaming AI and distributing governance” (2024) 6(2) *Harvard Data Science Review* 1–7.
114. OpenAI, “GPT-4o system card,” August 8, 2024, <https://openai.com/index/gpt-4o-system-card> (last accessed October 1, 2025).
115. Anthropic, “Anthropic’s Responsible Scaling Policy,” September 19, 2023, www.anthropic.com/news/anthropics-responsible-scaling-policy (last accessed October 1, 2025).
116. Margaret Mitchell et al., “Model cards for model reporting” (2019) *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAT*)* 220–29.

117. OpenAI, “GPT-4 system card,” 2023, <https://cdn.openai.com/papers/gpt-4-system-card.pdf> (last accessed October 1, 2025); Anthropic, “System card: Claude Opus 4 & Claude Sonnet 4,” <https://anthropic.com/model-card> (last accessed October 1, 2025); Meta, “System cards: A new resource for understanding how AI systems work,” November 16, 2022, <https://ai.meta.com/blog/system-cards-a-new-resource-for-understanding-how-ai-systems-work> (last accessed October 1, 2025).
118. Yuntao Bai et al., “Constitutional AI: Harmlessness from AI feedback,” arXiv, 2022, <https://arxiv.org/abs/2212.08073> (last accessed September 30, 2025).
119. Anthropic, “Collective constitutional AI: Aligning a language model with public input,” October 17, 2023, www.anthropic.com/research/collective-constitutional-ai-aligning-a-language-model-with-public-input (last accessed October 1, 2025).
120. Ian W. Eisenberg, Lucía Gamboa, and Eli Sherman, “The unified control framework: Establishing a common foundation for enterprise AI governance, risk management and regulatory compliance,” arXiv, 2025, <https://arxiv.org/abs/2503.05937> (last accessed September 30, 2025).
121. Natasha Crampton, “Reflecting on our responsible AI program: Three critical elements for progress,” Microsoft, May 1, 2023, <https://blogs.microsoft.com/on-the-issues/2023/05/01/responsible-ai-standards-principles-governance-progress> (last accessed October 1, 2025).
122. Natasha Crampton, “Microsoft’s framework for building AI systems responsibly,” Microsoft, June 21, 2022, <https://blogs.microsoft.com/on-the-issues/2022/06/21/microsofts-framework-for-building-ai-systems-responsibly> (last accessed October 1, 2025).
123. Satya Nadella, “The partnership of the future,” Slate, June 18, 2016, <https://slate.com/technology/2016/06/microsoft-ceo-satya-nadella-humans-and-a-i-can-work-together-to-solve-societys-challenges.html> (last accessed October 1, 2025).
124. Mike Murphy, “Microsoft CEO Satya Nadella has 10 commandments for how AI and humans should act,” Quartz, July 20, 2022, <https://qz.com/microsoft-ceo-satya-nadella-has-10-commandments-for-how-ai-and-humans-should-act-1849195111> (last accessed October 1, 2025).
125. Nadella, “The partnership of the future.”
126. Microsoft, “Responsible AI Standard: Reference guide, version 2,” June 2022, <https://iti.com/wp-content/uploads/2024/03/Microsoft-Responsible-AI-Standard-Reference-Guide.pdf> (last accessed October 1, 2025).
127. Brad Smith, “Meeting the AI moment: Advancing the future through responsible AI,” Microsoft, February 2, 2023, <https://blogs.microsoft.com/on-the-issues/2023/02/02/responsible-ai-chatgpt-artificial-intelligence> (last accessed October 1, 2025).
128. Microsoft, “Responsible AI transparency report 2025,” 2025, <https://cdn-dynmedia-1.microsoft.com/is/content/microsoftcorp/microsoft/msc/documents/presentations/CSR/Responsible-AI-Transparency-Report-2025-vertical.pdf#page=31> at 31 (last accessed October 1, 2025).
129. Natasha Crampton, “Advancing AI responsibly,” Microsoft, 2023, <https://unlocked.microsoft.com/responsible-ai> (last accessed October 1, 2025).
130. Teresa Hudson and Natasha Crampton, “Our 2025 responsible AI transparency report: How we build, support our customers, and grow,” Microsoft, June 20, 2025, <https://blogs.microsoft.com/on-the-issues/2025/06/20/our-2025-responsible-ai-transparency-report> (last accessed October 1, 2025).
131. Microsoft Research Asia, “Societal AI: Research challenges and opportunities,” March 2025, www.microsoft.com/en-us/research/wp-content/uploads/2025/03/Societal-AI-Research-Challenges-and-Opportunities.pdf at 12 (last accessed October 1, 2025).

132. Microsoft, “AI principles and approach,” www.microsoft.com/en-us/ai/principles-and-approach (last accessed August 12, 2025).
133. Microsoft, “Responsible AI standard v2 – general requirements,” June 2022, <https://msblogs.thesourcemediaassets.com/sites/5/2022/06/Microsoft-Responsible-AI-Standard-v2-General-Requirements-3.pdf> at 4 (last accessed October 1, 2025).
134. Microsoft, “Transparency note for Microsoft 365 Copilot,” August 6, 2025, <https://learn.microsoft.com/en-us/copilot/microsoft-365/microsoft-365-copilot-transparency-note> (last accessed October 1, 2025).
135. Carlo Giovine et al., “Building AI trust: The key role of explainability,” QuantumBlack, November 26, 2024, www.mckinsey.com/capabilities/quantumblack/our-insights/building-ai-trust-the-key-role-of-explainability (last accessed October 1, 2025).
136. Greg Satell and Yassmin Abdel-Magied, “AI fairness isn’t just an ethical issue,” Harvard Business Review, October 20, 2020, <https://hbr.org/2020/10/ai-fairness-isnt-just-an-ethical-issue> (last accessed October 6, 2025).
137. Daria Korobenko, Anastasija Nikiforova, and Rajesh Sharma, “Towards a privacy and security-aware framework for ethical AI: Guiding the development and assessment of AI systems,” *Proceedings of the 25th Annual International Conference on Digital Government Research* (ACM 2024).
138. Microsoft Azure, “Chapter 1: Understanding security and responsible AI,” Security and Responsible AI Guide, 2025, https://azure.github.io/Security-and-Responsible-AI-Guide/chapters/chapter_01_understanding_security_and_responsible_ai (last accessed October 6, 2025).
139. Alex Fleck, “Responsible AI: Why it matters and how we’re infusing it into our internal AI projects at Microsoft,” Microsoft Blog, June 5, 2025, www.microsoft.com/insidetrack/blog/responsible-ai-why-it-matters-and-how-were-infusing-it-into-our-internal-ai-projects-at-microsoft/ (last accessed October 6, 2025).
140. Joyce Chou, Roger Ibars, and Oscar Murillo, “In pursuit of inclusive AI,” Microsoft Inclusive Design, 2023, <https://inclusive.microsoft.design/tools-and-activities/InPursuitofInclusiveAI.pdf> (last accessed October 6, 2025).
141. Lama Lteif et al., “A conceptual framework for marketplace inclusion” (2025) 44(2) *Journal of Public Policy & Marketing* 214–31 (arguing that inclusive design and marketing practices expand participation by addressing exclusionary barriers and creating access for diverse consumer groups).
142. Dean Carignan, “Lessons from Microsoft’s responsible AI journey,” Emerj, February 15, 2023, <https://emerj.com/lessons-from-microsofts-responsible-ai-journey-dean-carignan-microsoft/> (last accessed October 6, 2025).
143. James Darley, “Inside Microsoft’s 2025 responsible AI transparency report,” Almagazine, June 24, 2025, <https://aimagazine.com/articles/inside-microsofts-2025-responsible-ai-transparency-report> (last accessed October 6, 2025) (highlighting Microsoft’s steps toward openness and accountability as regulators and customers scrutinize AI governance more intensely).
144. Casey Ross and Brittany Trang, “Microsoft is selling AI in health care, and helping to set its standards: is that a problem?” STAT+, April 18, 2024, www.statnews.com/2024/04/18/microsoft-health-care-artificial-intelligence-regulatory-capture/ (last accessed October 6, 2025).
145. Malak Sadek et al., “Challenges of responsible AI in practice: scoping review and recommended actions” (2025) 40 *AI & Society* 199–215.
146. Edwin A. Farley and Christian R. Lansang, “AI auditing: First steps towards the effective regulation of artificial intelligence systems” (2025) 38 *Harvard Journal of Law & Technology Digest* (Spring) 1–41 at 38 (proposing government-mandated AI auditing as a

necessary foundation for enforceable regulatory oversight); Brian Judge, Mark Nitzberg, and Stuart Russell, “When code isn’t law: Rethinking regulation for artificial intelligence” (2025) 44 *Policy & Society* 85–97 (arguing that AI’s scale and opacity demand direct governmental regulation and oversight rather than reliance on voluntary or industry-led self-regulation).

147. Eric Horvitz, “The Aether Committee at Microsoft,” [erichorvitz.com](https://erichorvitz.com/Aether_Committee_Microsoft.htm), 2019, https://erichorvitz.com/Aether_Committee_Microsoft.htm (last accessed October 6, 2025).
148. Crampton, “Reflecting on our responsible AI program.”
149. Todd Bishop, “Microsoft is turning down some sales over AI ethics, top researcher Eric Horvitz says,” *GeekWire*, October 30, 2018, www.geekwire.com/2018/microsoft-cutting-off-sales-ai-ethics-top-researcher-eric-horvitz-says/ (last accessed October 6, 2025).
150. Natasha Crampton, “The building blocks of Microsoft’s responsible AI program,” *Microsoft*, January 19, 2021, <https://blogs.microsoft.com/on-the-issues/2021/01/19/microsoft-responsible-ai-program/> (last accessed December 17, 2025).
151. Crampton, “Reflecting on our responsible AI program.”
152. Ernie Fernandez, “7 top questions government leaders are asking about AI,” *Microsoft*, 2023, <https://wwps.microsoft.com/blog/top-questions-ai> (last accessed October 6, 2025).
153. Microsoft, “Artificial intelligence overview,” 2025, <https://learn.microsoft.com/en-us/compliance/assurance/assurance-artificial-intelligence> (last accessed October 6, 2025).
154. Microsoft, “Responsible AI transparency report 2024,” May 2024, <https://cdn-dynmedia-1.microsoft.com/is/content/microsoftcorp/microsoft/msc/documents/presentations/CSR/Responsible-AI-Transparency-Report-2024.pdf> (last accessed November 18, 2025).
155. Microsoft, “AI RMF RFI comments,” 2022, www.nist.gov/document/ai-rmf-rfi-comments-microsoft (last accessed October 6, 2025).
156. Microsoft, “Responsible AI standard v2 – general requirements,” 6.
157. Microsoft, “Microsoft responsible AI impact assessment template,” 2022, <https://msblogs.thesourcemediassets.com/sites/5/2022/06/Microsoft-RAI-Impact-Assessment-Template.pdf> (last accessed October 6, 2025).
158. Microsoft, “Responsible AI standard v2 – general requirements,” A1.1.
159. Microsoft, “Artificial intelligence overview.”
160. Microsoft, “Reflecting on our responsible AI program.”
161. Microsoft, “Responsible AI transparency report 2025.”
162. Microsoft, “Governing AI: A blueprint for the future,” May 25, 2023, <https://cdn-dynmedia-1.microsoft.com/is/content/microsoftcorp/microsoft/msc/documents/presentations/CSR/Governing-AI-A-Blueprint-for-the-Future.pdf> at 10 (last accessed October 6, 2025).
163. Microsoft, “Putting principles into practice: How we approach responsible AI at Microsoft,” 2020, www.microsoft.com/cms/api/am/binary/RE4pKH5 at 12 (last accessed October 6, 2025).
164. Microsoft, “Governing AI: A blueprint for the future,” 34.
165. Michael Smith and Seumas Miller, “The ethical application of biometric facial recognition technology” (2022) 37 *AI & Society* 167–75; Robin A. Waelen, “The struggle for recognition in the age of facial recognition technology” (2023) 3 *AI and Ethics* 215–22.
166. Microsoft, “Code of conduct for Microsoft AI services,” 2025, <https://learn.microsoft.com/en-us/legal/ai-code-of-conduct> at § 19 (last accessed October 13, 2025); Masha Borak, “Microsoft further limits law enforcement use of its facial recognition,” *Biometric Update*, May 3, 2024, www.biometricupdate.com/202405/microsoft-further-limits-law-enforcement-use-of-its-facial-recognition (last accessed October 13, 2025).
167. Axon, “How Axon is using AI responsibly,” 2024, www.axon.com/resources/how-axon-is-using-ai-responsibly (last accessed October 13, 2025); Axon, “A closer look at Draft One,” 2024, www.axon.com/resources/closer-look-draft-one (last accessed October 13, 2025).

168. Leah Culler, “Are you talking to me? Azure AI brings iconic characters to life with custom neural voice,” Microsoft, February 3, 2021, <https://news.microsoft.com/source/features/ai/custom-neural-voice-ga/> (last accessed October 7, 2025).
169. Eric Hal Schwartz, “Microsoft restricts custom voice and facial recognition AI access,” Voicebot.ai (Synthedia), June 21, 2022, <https://voicebot.ai/2022/06/21/microsoft-restricts-custom-voice-and-facial-recognition-ai-access/> (last accessed October 13, 2025).
170. Betsy Weber, “Tech for food: Giving a voice to the voiceless,” Microsoft Tech Community, August 20, 2024, <https://techcommunity.microsoft.com/blog/mvpawardprogramblog/tech-for-good-giving-a-voice-to-the-voiceless/4217303> (last accessed October 7, 2025).
171. Allison Koenecke et al., “Racial disparities in automated speech recognition” (2020) 117 *Proceedings of the National Academy of Sciences* 7684–89.
172. Shaun Sutner, “Microsoft reins in AI facial and voice recognition tech,” TechTarget, June 23, 2022, www.techtarget.com/searchenterpriseai/news/252521952/Microsoft-reins-in-AI-facial-and-voice-recognition-tech (last accessed October 13, 2025).
173. The White House, “Winning the race: America’s AI action plan,” July 2025, www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf (last accessed November 23, 2025).
174. United States, Preparing Americans for high-paying skilled-trade jobs of the future, 90 Fed. Reg. 26148 (28 April 2025); The White House, “Removing barriers to American leadership in artificial intelligence,” January 23, 2025, www.whitehouse.gov/presidential-actions/2025/01/removing-barriers-to-american-leadership-in-artificial-intelligence/ (last accessed November 23, 2025).
175. Bobby Allyn, “Trump’s new AI policies keep culture war focus on tech companies,” NPR, 23 July 2025, www.npr.org/2025/07/23/nx-si-5476771/trump-artificial-intelligence-woke-eo (last accessed November 23, 2025).
176. The White House, “Winning the race: America’s AI action plan.”
177. One Big Beautiful Bill Act, Pub. L. No. 119–21, 139 Stat. 72 (2025).
178. Will Oremus, “In dramatic reversal, Senate votes to kill AI-law moratorium,” *The Washington Post*, July 1, 2025 www.washingtonpost.com/politics/2025/07/01/ai-moratorium-big-beautiful-bill/ (last accessed November 23, 2025); Regina Sam Penti and Ama A. Adams, “AI and tech under the One Big Beautiful Bill Act: Key restrictions, risks and opportunities,” *Ropes & Gray*, July 14, 2025, www.ropesgray.com/en/insights/alerts/2025/07/ai-and-tech-under-the-one-big-beautiful-bill-act-key-restrictions-risks-and-opportunities (last accessed November 25, 2025).
179. The White House, “Winning the race: America’s AI action plan,” 2.
180. Office of the Governor of California, “Governor Newsom signs SB 53, advancing California’s world-leading artificial intelligence industry,” September 29, 2025, www.gov.ca.gov/2025/09/29/governor-newsom-signs-sb-53-advancing-californias-world-leading-artificial-intelligence-industry/ (last accessed November 25, 2025); Erin M. Bosman et al., “At the frontier – California enacts AI safety and transparency regulation TFAIA (SB 53),” Morrison & Foerster, October 1, 2025, www.mofo.com/resources/insights/251001-california-enacts-ai-safety-transparency-regulation-tfaia-sb-53 (last accessed November 22, 2025).
181. Javier Espinoza, “EU agrees landmark rules on artificial intelligence,” *Financial Times*, December 9, 2023, www.ft.com/content/cc8a1bod-e676-44b2-a4ab-69a8b28232a5 (last accessed November 25, 2025).
182. European Commission, “Regulatory framework proposal on artificial intelligence: Questions & answers,” April 21, 2021, https://ec.europa.eu/commission/presscorner/detail/en/qanda_21_1683 (last accessed November 22, 2025).

183. European Union, “Article 5: Prohibited AI practices,” AI Act Service Desk, <https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-5> (last accessed November 23, 2025); European Parliament, “Texts adopted – Artificial Intelligence Act,” June 14, 2023, www.europarl.europa.eu/doceo/document/TA-9-2023-0236_EN.html (last accessed November 25, 2025).
184. Regulation (EU) 2024/1689 of the European Parliament and of the Council of June 13, 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act), OJ L, July 12, 2024.
185. AIA, Arts 9 and 17.
186. AIA, Art 12.
187. AIA, Arts 19–21.
188. European Union, “The AI Office: What is it, and how does it work?” March 21, 2024, <https://artificialintelligenceact.eu/the-ai-office-summary/> (last accessed November 25, 2025).
189. Benjamin Mueller, “How much will the (EU) AI Act cost?” Center for Data Innovation, July 2021, www2.datainnovation.org/2021-ai-a-costs.pdf (last accessed November 25, 2025).
190. Pascale Davies, “Why OpenAI’s voice mode, Meta’s Llama and Apple’s AI won’t be coming to Europe yet,” Euronews, October 8, 2024, www.euronews.com/next/2024/10/08/why-openais-voice-mode-metas-llama-and-apples-ai-wont-be-coming-to-europe-yet (last accessed November 25, 2025).
191. Anu Bradford, *Digital Empires: The Global Battle to Regulate Technology* (Oxford: Oxford University Press, 2023).
192. DeepMind, “How our principles helped define AlphaFold’s release,” September 14, 2022, <https://deepmind.google/en/blog/how-our-principles-helped-define-alphafolds-release/> (last accessed November 25 2025).
193. Microsoft, “Our approach to responsible AI,” 2023, www.microsoft.com/en-us/ai/responsible-ai (last accessed November 25 2025).

9 A NEW FIDUCIARY STANDARD

1. Melvin A. Eisenberg, “The divergence of standards of conduct and standards of review in corporate law” (1993) 62 *Fordham Law Review* 437–503.
2. Tony West, “Uber delivers U.S. Safety Report,” Uber, December 6, 2019, www.uber.com/en-IL/newsroom/2019-us-safety-report/ (last accessed November 5, 2025).
3. Ryan Lawler, “On-demand ridesharing startup Lyft adds insurance between rides,” TechCrunch, March 13, 2014, <https://techcrunch.com/2014/03/13/lyft-insurance-between-rides/> (last accessed November 15, 2025).
4. World Resources Institute, “Safe climate, sound business” (describing a collaborative effort by GM, BP, the World Resources Institute, and Monsanto that initiated the partnerships later leading to the Greenhouse Gas Protocol).
5. Esmat Zaidan and Imad Antoine Ibrahim, “AI governance in a complex and rapidly changing regulatory landscape: A global perspective” (2024) 11 *Humanities and Social Sciences Communications* 1–18.
6. Crampton, “Microsoft’s framework for building AI systems responsibly.”
7. Elhauge, “Sacrificing corporate profits in the public interest.”
8. Charles Hansen, “The duty of care, the business judgment rule, and the American Law Institute corporate governance project” (1993) 48 *The Business Lawyer* 1355–76; Bernard Black, Brian Cheffins, and Michael Klausner, “Outside director liability” (2006) 58 *Stanford Law Review* 1055–160.