

RESEARCH ARTICLE

From AI Adoption to Institutional Preparedness: Governing Enterprises for AGI

Fang Sun[†] 

Department of Business Administration, Dongshin University, Jeollanam-do 58245, Republic of Korea

[†]Corresponding author. Email: sf2025@dsu.ac.kr

Abstract

Artificial general intelligence (AGI) is increasingly discussed not merely as a technical frontier but as a social and institutional challenge. Drawing on current generative and agentic AI as diagnostic evidence, this paper examines enterprise institutional preparedness for increasingly general and autonomous AI systems through three interrelated dimensions: organizational resilience, employee protection, and social responsibility. Using conceptual and normative analysis, the paper synthesizes AI governance, organizational resilience, labor studies, corporate social responsibility, and policy frameworks to develop a forward-looking governance architecture. The framework links three levels of institutional exposure: continuity of organizational functions, preservation of worker agency, and external social legitimacy. It identifies seven preparedness components: strategic foresight, governance architecture, risk and impact assessment, human oversight, employee protection, social responsibility, and continuous learning. The analysis further distinguishes scalable risks already visible in current AI from potentially discontinuous risks associated with greater autonomy and cross-domain integration. The study shifts the unit of analysis from tool-level AI risk management to enterprise-level institutional capacity to preserve organizational continuity, worker agency, and external social legitimacy.

Keywords: Artificial general intelligence; institutional preparedness; enterprise governance; organizational resilience; employee protection; corporate social responsibility

1. Introduction

Artificial general intelligence (AGI) has moved from a speculative technical concept to a central theme in debates about ethics, governance, labor, education, human-machine relations, and future society. Definitions differ, but most accounts emphasize broad task generality, adaptability, and autonomy rather than capability within a narrow application domain (Legg & Hutter, 2007; Morris et al., 2023). Fully autonomous AGI is not an established enterprise technology. Current generative and agentic AI systems are therefore treated in this paper as precursor technologies that reveal governance mechanisms that could become more consequential if AI systems become more general and autonomous.

For enterprises, increasingly general and autonomous AI raises questions that go beyond automation, productivity, and digital transformation. More capable systems could participate in strategic analysis, product development, customer interaction, legal review, risk management, supply-chain coordination, human resource decisions, and organizational learning. Such participation can create



opportunities for innovation and resilience while also increasing exposure to accountability gaps, overdependence, opacity, cybersecurity risks, workforce disruption, and the externalization of harm to workers, consumers, communities, and public institutions (Bostrom, 2014; Russell & Norvig, 2021).

Existing enterprise AI research often focuses on adoption, performance, managerial efficiency, or technical risk. The institutional gap addressed here is different: as AI systems become more general, agentic, and embedded across functions, enterprises must determine not only whether a system performs well, but whether the organization has rules, capabilities, and accountability arrangements for governing its consequences. Institutional preparedness refers to the organizational structures, cultural norms, employee protections, escalation rights, and social responsibility commitments that allow a firm to anticipate and govern disruptive AI-related change before failures become systemic. This shifts attention from isolated tool governance to enterprise-level institutional capacity.

This paper organizes institutional preparedness around three interdependent dimensions. Organizational resilience concerns continuity, adaptation, and recovery under technological disruption. Employee protection concerns worker agency, dignity, participation, skill transition, and fair treatment. Social responsibility concerns the effects of enterprise AI use on external stakeholders, public trust, and the distribution of benefits and harms (Floridi & Cowls, 2019; Jobin et al., 2019). Their integration is theoretically important because they represent three linked levels of governance failure: disruption to organizational continuity, erosion of worker agency, and loss of external social legitimacy. Preparedness demands will also vary across jurisdictions and firm sizes because enterprises face different regulatory expectations, labor institutions, and governance resources. The framework is therefore designed to be cross-regionally adaptable and proportionate to organizational capacity.

The study adopts a conceptual and normative analysis method. It uses *ex ante* conceptual analysis to integrate governance mechanisms that are currently dispersed across AI governance, organizational resilience, labor studies, and corporate responsibility. Evidence from generative AI, algorithmic management, and other advanced AI applications serves as diagnostic and analogical evidence for mechanisms that may persist or intensify as systems become more general and autonomous. This approach supports framework construction at a stage when AGI does not yet constitute a stable empirical category, while empirical calibration and validation remain tasks for subsequent research.

2. Conceptual Background and Analytical Approach

The concept of AGI is contested. Some scholars define it in relation to universal intelligence, others in terms of human-level performance across diverse tasks, and still others through levels of capability, autonomy, and generality (Legg & Hutter, 2007; Morris et al., 2023). For enterprise governance, this paper uses AGI as an analytical horizon rather than as a claim that a specific existing model is already AGI. The relevant transition is from AI used within bounded tasks and workflows toward systems capable of broader cross-domain reasoning, greater operational autonomy, and deeper participation in organizational decision processes.

2.1 Distinguishing Current Generative AI from Future AGI

Current generative AI and future AGI should not be treated as interchangeable categories. The distinction matters because the governance burden changes as task scope, autonomy, adaptation, accountability traceability, and enterprise integration increase. Table 1 states the conceptual boundary used throughout this paper.

Table 1. Conceptual Distinction Between Current Generative AI and Future AGI for Enterprise Governance

Dimension	Current Generative / Agentic AI	Future AGI as an Analytical Horizon	Governance Implication
Task scope	Broad capabilities but still bounded by products, tasks, and use cases	Cross-domain capability across diverse organizational functions	Governance expands from individual use cases to enterprise-wide coordination
Autonomy	Human-initiated use or bounded agentic workflows	Potentially sustained planning and action with wider discretion	Stronger authorization, escalation, pause, and override rights
Learning and adaptation	Predominantly externally trained, with limited in-use adaptation	Potentially more persistent and context-sensitive adaptation	Continuous monitoring, change control, and capability reassessment
Accountability traceability	Outputs usually linked to identifiable tools, owners, and workflows	Multi-step decisions may span systems, functions, and organizational levels	Greater need for traceable decision rights and responsibility allocation
Enterprise integration	Function-specific copilots, platforms, APIs, or decision aids	Cross-functional participation in organizational cognition and action	Institutional governance becomes as important as technical control

The evidentiary logic is therefore prospective but bounded. Some governance risks are continuous: bias, opacity, excessive monitoring, cybersecurity exposure, overreliance, and accountability diffusion are already observable in current AI applications and may scale as capability and organizational dependence increase. Other risks are potentially discontinuous, including cross-functional autonomous action, persistent goal pursuit, and cascading decisions that arise only if systems obtain substantially greater autonomy and generality.

Accordingly, current generative AI deployments and algorithmic-management cases serve as diagnostic windows into institutional mechanisms that can inform anticipatory governance design. Their relevance lies in revealing governance pathways that may persist or intensify as capability, autonomy, and organizational integration increase.

Artificial intelligence governance literature provides a foundation for analyzing AGI preparedness. Human-compatible AI, cooperative AI, and trustworthy AI approaches all emphasize that advanced AI systems should be aligned with human values, subject to oversight, and embedded in accountable institutions (Dafoe, 2018; Russell, 2019). Ethical AI scholarship further identifies recurring principles such as transparency, fairness, responsibility, privacy, safety, and human autonomy (Floridi & Cowls, 2019; Jobin et al., 2019). These principles are useful, but they are often stated at a high level. Enterprises need to translate them into internal governance practices.

Policy frameworks have increasingly moved from voluntary ethical principles toward operational risk governance. The OECD AI Principles and UNESCO's Recommendation emphasize human rights, fairness, and social welfare (Organisation for Economic Co-operation and Development, 2024; UNESCO, 2021). The EU AI Act provides a risk-based legal architecture (European Union, 2024). Importantly, the NIST AI Risk Management Framework already addresses organization-wide governance, accountability, leadership responsibilities, personnel training, human-AI oversight, and engagement with affected stakeholders (Tabassi, 2023). The present analysis builds on, rather than substitutes for, these provisions by examining how governance arrangements jointly sustain organizational resilience, employee agency, and external legitimacy as AI systems become more general and autonomous. Section 7 compares regulatory settings and their cross-border implications.

A second body of literature concerns organizational resilience. Resilience is not simply the ability to survive a crisis; it includes anticipation, coping, adaptation, and learning. Organizations become

resilient when they can detect weak signals, coordinate resources, preserve critical functions, and transform routines under uncertainty (Duchek, 2020; Lengnick-Hall et al., 2011). Increasingly general and autonomous AI could create disruption not only through direct automation but through rapid shifts in markets, knowledge production, risk exposure, and stakeholder expectations. Therefore, resilience becomes a central dimension of institutional preparedness.

A third body of literature concerns labor, algorithmic management, and employee protection. Research on digital platforms and algorithmic control shows that automated systems can intensify monitoring, obscure managerial responsibility, and reduce worker voice (De Stefano, 2019; Kellogg et al., 2020). Similar patterns already appear when advanced AI is used to evaluate employees, allocate tasks, recommend personnel actions, or mediate communication; future AGI could amplify these effects through greater autonomy and cross-functional reach. The key issue is not only job loss; it is whether employees retain agency, dignity, meaningful participation, and access to transition support.

The fourth body of literature is corporate social responsibility (CSR). CSR has evolved from philanthropic activities to a broader understanding of enterprise responsibilities toward stakeholders, social legitimacy, and shared value creation (Carroll, 1991; Freeman, 1984). As AI systems become more general and autonomous, CSR must include responsible innovation, AI impact assessment, fairness, transparency, sustainability, and protection of affected communities. Increasingly general AI deployment may produce benefits that are privately captured while risks are socially distributed. Institutional preparedness should therefore include mechanisms for social accountability (Elkington, 1997; Porter & Kramer, 2011).

Based on these literatures, this paper uses a three-dimensional analytical framework: organizational resilience, employee protection, and social responsibility. The dimensions are analytically distinct but institutionally coupled. Resilience addresses whether the enterprise can continue to function and adapt; employee protection addresses whether adaptation preserves worker agency and fair treatment; social responsibility addresses whether the firm prevents private technological gains from shifting harm to external stakeholders. A firm can therefore appear operationally resilient while weakening employee agency, or appear socially responsible while lacking internal accountability. Institutional preparedness integrates these levels so that continuity, worker protection, and legitimacy are governed as parts of the same enterprise-level problem. Table 2 summarizes these three domains, their principal risks, preparedness mechanisms, and expected outcomes.

Table 2. Three Dimensions of Enterprise Institutional Preparedness for AGI

Dimension	Main Risks	Preparedness Mechanisms	Expected Outcomes
Organizational Resilience	Operational disruption, overdependence on increasingly capable AI systems, cybersecurity risks	Scenario planning, fallback procedures, risk governance	Organizational adaptability and stability
Employee Protection	Job displacement, skill anxiety, surveillance, limited employee voice	Reskilling, human review, consultation, appeal channels	Worker dignity and fair transition
Social Responsibility	Bias, opacity, externalized harm, public distrust	Impact assessment, transparency, CSR reporting, stakeholder dialogue	Responsible innovation and social legitimacy

3. Why Enterprises Need Institutional Preparedness for AGI

Enterprise preparedness is necessary because increasingly general and agentic AI could alter the structure of organizational decision-making. Current systems are often deployed for bounded tasks such as content generation, classification, coding support, customer service, and decision assistance. If future systems acquire broader cross-domain capability and autonomy, they could synthesize market intelligence, simulate strategic scenarios, coordinate workflows, and influence decisions across functions. The governance issue is therefore not simply automation of work, but possible participation of AI in organizational cognition (Bostrom, 2014; Russell & Norvig, 2021).

As AI participates more deeply in organizational cognition, responsibility can become distributed across human instructions, model outputs, data pipelines, automated workflows, vendors, and managerial approvals. Existing AI accountability research already documents such responsibility gaps (Raji et al., 2020; Selbst et al., 2019). Greater cross-functional autonomy could deepen them by allowing one system's recommendations or actions to propagate across departments. Preparedness therefore requires traceable ownership and escalation rather than assuming that a final human click is sufficient accountability.

Preparedness is also needed because deeper reliance on increasingly capable AI may create organizational dependence. Routine dependence can weaken human expertise and institutional memory if employees cease practicing critical judgment because automated outputs appear faster or more authoritative. A firm may then become efficient in ordinary conditions but fragile during model failure, data drift, cyberattack, or unexpected social crisis. Human oversight should therefore be treated as an organizational capability that is maintained and tested, not as symbolic approval (Lee & See, 2004; Shneiderman, 2020).

A further reason for preparedness is labor transition. The labor impact of advanced AI is not limited to displacement. It also includes task redesign, skill polarization, work intensification, surveillance, and the redefinition of occupational identity. Studies of AI and work emphasize that automation can substitute for some tasks while complementing others, and that the final effects depend on institutions, training, bargaining power, and organizational choices (Gmyrek et al., 2023; World Economic Forum, 2025). Enterprises cannot assume that market adjustment alone will produce fair outcomes. They need rules for redeployment, reskilling, consultation, and protection.

Preparedness also matters because increasingly general AI can generate social externalities beyond the enterprise. Systems used in employment, customer interaction, credit, information services, logistics, or other consequential domains can affect access to opportunities and distribute harms unevenly. Studies of automated inequality and data-driven social power illustrate how technical systems can reshape social opportunities even when they are privately deployed (Eubanks, 2018; Zuboff, 2019). Future AGI could amplify these concerns through greater scale, speed, and autonomy.

Finally, preparedness is urgent because governance often lags behind technological diffusion. Enterprises tend to experiment with new technologies before legal rules become settled. This can create a gap between capability and accountability. The EU AI Act, NIST AI RMF, OECD Principles, and UNESCO Recommendation all indicate that AI governance is becoming more formalized, but enterprise-level implementation remains uneven (Organisation for Economic Co-operation and Development, 2024; UNESCO, 2021). Institutional preparedness helps firms move from reactive compliance to anticipatory governance. It enables organizations to ask not only what increasingly capable AI can do, but what should be delegated, what should remain human-controlled, and what responsibilities follow from deployment.

The necessity of preparedness also differs by firm size and resource base. Large enterprises may support dedicated legal, cybersecurity, risk, and responsible-AI teams, whereas SMEs often depend on external vendors and shared cloud services. A proportionate approach is therefore preferable to a uniform governance architecture: smaller firms still need an AI inventory, risk classification, named accountability, human review for high-impact decisions, vendor due diligence, and an

incident-escalation path, but they may obtain specialized assurance or legal expertise externally. This recognizes the real cost of governance without treating resource constraints as a reason to omit basic safeguards.

4. Organizational Resilience under AGI

Organizational resilience in the AGI era should be understood as the capacity to remain functional, adaptive, and ethically grounded under conditions of technological uncertainty. It is not enough for a firm to recover from disruption after it occurs. Enterprises need to anticipate possible disruptions, design flexible routines, maintain human competence, and create mechanisms for learning from near misses. Resilience literature emphasizes that organizations must combine prevention, adaptation, and transformation rather than treat resilience as a single crisis response capacity (Duchek, 2020; Lengnick-Hall et al., 2011).

The first element of resilience is strategic foresight. AGI development is uncertain in timing, capability, regulation, and social acceptance. Enterprises should therefore avoid making a single forecast. Instead, they can use scenario planning to consider multiple futures: gradual augmentation, rapid automation, regulatory restriction, public resistance, geopolitical fragmentation, or major safety incident. Each scenario has different implications for workforce planning, data governance, cybersecurity, compliance, and stakeholder trust. Foresight allows enterprises to build options rather than lock themselves into one technological pathway.

The second element is governance infrastructure. AI governance cannot be left entirely to IT departments or innovation teams. AGI preparedness requires a cross-functional governance body that includes senior management, legal experts, HR, cybersecurity, risk management, employee representatives, ethics advisors, and operational units. This body should define risk appetite, approve high-impact deployments, monitor compliance, and coordinate incident response. The NIST AI RMF is useful here because it frames risk management as an organizational practice involving governance, mapping, measurement, and management (European Union, 2024; Tabassi, 2023).

The third element is operational continuity. If increasingly general and autonomous AI systems become deeply embedded in workflows, their failure could disrupt business processes. Enterprises need fallback procedures, manual override capacity, data recovery mechanisms, vendor risk assessment, and human competence redundancy. Resilience research stresses that organizations can become fragile when efficiency eliminates slack, diversity, and local judgment (Linnenluecke, 2017; Williams et al., 2017). In more deeply AI-mediated contexts, the same danger appears when firms optimize around automated processes without preserving human ability to intervene.

The fourth element is knowledge resilience. Advanced AI can automate parts of knowledge work, and future AGI could extend this dependence across a wider range of organizational functions. Enterprises should not allow critical knowledge to become inaccessible to humans. When workers rely heavily on AI-generated analysis, they may lose the capacity to question assumptions, identify anomalies, and understand causal relations. Knowledge resilience requires documentation, peer review, training, communities of practice, and deliberate human learning. High-reliability organization research suggests that attention to weak signals, reluctance to simplify, and sensitivity to operations are crucial when systems become complex (Weick & Sutcliffe, 2015).

The fifth element is cybersecurity and model risk resilience. Advanced AI systems can be connected to internal databases, external APIs, customer interfaces, and operational systems, expanding the attack surface. Risks include prompt injection, data leakage, malicious automation, model manipulation, identity fraud, and cascading error. While technical controls are necessary, institutional controls are equally important: access rights, audit trails, procurement standards, vendor accountability, and emergency suspension rules. OpenAI's preparedness framework and frontier AI risk studies show that advanced AI systems require systematic evaluation of severe harms rather than ad hoc safety checks (Amodei et al., 2016; OpenAI, 2025).

The sixth element is cultural resilience. Enterprises need cultures that encourage responsible reporting, dissent, and learning. Employees should be able to challenge AI outputs, report unsafe use, and raise ethical concerns without fear of retaliation. A resilient organization does not treat skepticism as resistance to innovation. It treats it as a source of risk intelligence. In the AGI era, organizational culture must balance experimentation with caution, speed with reflection, and automation with human judgment.

Resilience design should also reflect sectoral exposure. Manufacturing firms may prioritize operational continuity, cyber-physical safety, and manual fallback capacity; financial firms may place greater weight on model risk, consumer impact, and traceable approvals; technology firms may face faster capability change, product-scale diffusion, and third-party ecosystem risks. Across sectors, implementation can be blocked by fragmented ownership, short budget cycles, unclear authority between technical and business units, and incentives that reward rapid deployment. Preparedness therefore requires not only formal rules but also decision rights, resources, and escalation mechanisms that can operate through ordinary organizational bureaucracy.

5. Employee Protection in the AGI Era

Employee protection is a central dimension of enterprise institutional preparedness because increasingly capable AI systems may transform the meaning, structure, and distribution of work. From the employee perspective, AI adoption can appear not only as productivity improvement but also as job insecurity, skill anxiety, monitoring pressure, reduced autonomy, and uncertainty about future career pathways. ILO analysis of generative AI indicates that many effects may involve augmentation rather than full automation, but exposure differs across occupations and countries (Gmyrek et al., 2023). This uneven exposure makes institutional protection necessary before more general and autonomous systems emerge.

The first employee protection issue is skill transition. Current generative AI is already changing how some employees perform cognitive tasks, while future AGI could deepen changes in how competence is learned, exercised, and evaluated. Routine cognitive tasks may be automated, while judgment, coordination, ethical reasoning, domain interpretation, and AI supervision become more important. WEF reports similarly emphasize that AI, big data, and technological literacy are among the fastest-growing skills, while human skills such as analytical thinking, resilience, leadership, and collaboration remain essential (World Economic Forum, 2025). Enterprises should therefore build reskilling systems rather than rely on individual adaptation alone.

The second issue is fair redeployment. If increasingly capable AI reduces demand for certain tasks, enterprises should develop internal mobility pathways before layoffs become the default response. Redeployment should include transparent criteria, training access, career counseling, and reasonable transition time. Labor economics research shows that technological change can both displace and create tasks, but outcomes depend on institutional arrangements and labor market adjustment mechanisms (Autor, 2015; Cazzaniga et al., 2024). Enterprise preparedness should therefore include a transition compact: a clear commitment to help workers move toward emerging roles when possible.

The third issue is worker voice. Employees are often the first to understand how AI systems affect real work, but they may be excluded from design and deployment decisions. Algorithmic management research shows that automated systems can shift power toward management while making rules less visible to workers (De Stefano, 2019; Kellogg et al., 2020). More autonomous and deeply integrated AI could intensify this problem if employees cannot understand why tasks are assigned, how performance is evaluated, or how AI-generated recommendations influence promotion and discipline. Institutional preparedness should include employee consultation, appeal channels, and participatory assessment.

The fourth issue is psychological safety. Rapid AI adoption can create fear of redundancy, pressure to constantly upgrade skills, and anxiety about being compared with machine performance;

increasingly general AI could deepen these concerns. Psychological safety does not mean avoiding change; it means creating conditions in which employees can learn, ask questions, admit uncertainty, and report risks. Enterprises should communicate openly about AI deployment plans, expected task changes, and support measures. Silence can produce rumors and resistance, while transparent communication can transform anxiety into collective problem-solving.

The fifth issue is protection against excessive monitoring. Increasingly capable AI systems may make it easier to collect, infer, and evaluate employee behavior in real time. This can create a workplace where employees feel constantly observed. Research on digital labor platforms shows that algorithmic evaluation can reduce autonomy and increase opacity (Rosenblat, 2018; Wood et al., 2019). Enterprises should define limits on employee data collection, prohibit invasive inference where unnecessary, and separate legitimate operational analytics from excessive surveillance. Employee protection requires both technical privacy safeguards and institutional rules.

The sixth issue is human oversight. The EU AI Act emphasizes human oversight for high-risk AI systems, and this principle is especially relevant for employment-related decisions (European Union, 2024). Enterprises should ensure that employees can request human review when advanced AI influences hiring, evaluation, promotion, task allocation, disciplinary action, or termination. Oversight should be meaningful: reviewers need authority, expertise, and access to relevant information. Human review should not be a symbolic rubber stamp over automated decisions.

The seventh issue is equality. AI systems trained on historical data may reproduce inequalities in hiring, evaluation, and opportunity distribution, and wider deployment could amplify these effects. Fairness research warns that abstract technical solutions may fail when social context is ignored (Raji et al., 2020; Selbst et al., 2019). Enterprises should conduct bias audits, monitor differential impacts, and involve affected groups in assessment. Employee protection is therefore not only about preserving jobs; it is also about ensuring that increasingly capable AI does not deepen existing workplace inequalities.

Vulnerability is not evenly distributed within the workforce. Temporary employees, contractors, low-skilled frontline staff, and workers with limited bargaining power may have less access to reskilling, less ability to contest automated decisions, and greater exposure to monitoring or rapid task substitution. Corporate incentives can also weaken protection when training, consultation, or human review is treated mainly as a cost that slows deployment. For this reason, preparedness should evaluate not only whether a policy exists but whether affected workers have practical access to training, appeal, representation, and an authorized human decision-maker. This moves employee protection from a general ethical commitment toward an institutional constraint on how technological gains are pursued.

6. Corporate Social Responsibility in the AGI Era

The prospect of increasingly general and autonomous AI broadens the meaning of corporate social responsibility because enterprise AI deployment may affect stakeholders far beyond shareholders and employees. Traditional CSR literature emphasizes responsibilities toward economic, legal, ethical, and philanthropic expectations, as well as stakeholder claims (Carroll, 1991; Freeman, 1984). As AI systems become more general and autonomous, these responsibilities extend to technological responsibility: the obligation to develop, procure, deploy, and monitor intelligent systems in ways that respect human dignity, fairness, safety, and social trust.

One CSR challenge is responsibility for indirect harm. Enterprises may use increasingly capable AI to personalize prices, screen customers, optimize logistics, generate marketing content, manage credit risk, or automate service interactions. These applications can affect consumer choices, access to opportunities, and exposure to manipulation. If harm occurs, firms may claim that the system acted unpredictably or that the vendor supplied the model. Such claims are insufficient. CSR requires enterprises to accept responsibility for systems they deploy, especially when they profit from their

use.

Another challenge is transparency. Stakeholders should know when advanced AI systems influence consequential interactions. Transparency does not require revealing trade secrets or overwhelming users with technical details. It requires meaningful disclosure, understandable explanations, and accessible channels for contestation. Ethical AI literature repeatedly identifies transparency as a central principle, but transparency must be translated into stakeholder-specific practices (Floridi & Cowls, 2019; Jobin et al., 2019). Customers, employees, regulators, and communities may need different forms of explanation.

A third CSR challenge is social impact assessment. Enterprises often evaluate AI projects through cost savings, efficiency, or market growth. As AI systems become more general and autonomous, they should also assess social risks: labor displacement, bias, misinformation, environmental cost, psychological dependence, community effects, and public trust. The UN Guiding Principles on Business and Human Rights emphasize that companies should identify, prevent, mitigate, and account for adverse human rights impacts (United Nations, 2011). Prospective AGI social impact assessment can be understood as an extension of this due diligence logic.

A fourth challenge is benefit distribution. Advanced AI may produce substantial productivity gains, and future AGI could extend them across a wider range of activities. If these gains accrue mainly to capital owners while workers face insecurity, the technology may increase inequality and reduce social legitimacy. Shared value approaches argue that business success and social progress can be mutually reinforcing when firms redesign value creation around broader stakeholder needs (Porter & Kramer, 2011). In a future AGI context, this could include profit-sharing, training investment, reduced work intensity, support for affected communities, and responsible supply-chain practices.

A fifth challenge is environmental responsibility. Contemporary evidence links the expansion of AI-related computing to rising electricity demand from data centers. The International Energy Agency (2025) estimates that data centers consumed approximately 415 TWh globally in 2024 and projects around 945 TWh in 2030 under its base-case scenario, with AI an important driver of growth. These totals cover data centers as a whole, however, and must not be misrepresented as AI-only electricity consumption. A U.S. energy-use study by Shehabi et al. (2024) similarly estimates historical data-center electricity consumption and develops future demand scenarios. These sources provide sector-level estimates and projections, not measurements of the enterprises considered in this conceptual paper. As a proposed, not yet validated, assessment practice, enterprises could document AI model or API usage, computational workloads, electricity data supplied by vendors, and the carbon intensity of relevant power sources, while disclosing allocation boundaries and uncertainty. Such indicators should not be interpreted as precise firm-level AI energy or emissions estimates without reliable measurement. Elkington (1997) provides the broader sustainability rationale rather than direct empirical evidence of contemporary AI electricity demand.

A sixth challenge is public trust. Enterprises do not operate outside society. If major firms deploy increasingly general AI irresponsibly, public trust in AI, markets, and institutions may decline. Conversely, responsible enterprise behavior can support public legitimacy by showing that innovation can be governed. International governance initiatives such as the Bletchley Declaration and the Frontier AI Safety Commitments agreed at the 2024 Seoul AI Summit reflect growing recognition that advanced AI safety requires cooperation among governments, firms, and civil society (Department for Science, Innovation and Technology, 2023, 2024). Enterprises should see themselves as participants in this broader governance ecosystem.

Corporate social responsibility in the AGI era should therefore be embedded in enterprise governance. It should not be limited to sustainability reports or public statements. It should include board oversight, risk committees, stakeholder consultation, impact reporting, grievance mechanisms, procurement standards, and independent audits. The boundary of responsibility expands when enterprise technologies acquire greater autonomy and social reach.

Cross-border deployment further complicates responsibility. A multinational enterprise may procure a model in one jurisdiction, process data in another, and affect workers or customers elsewhere. Vendor contracts do not eliminate the deploying firm's obligations to assess foreseeable impacts, preserve traceable accountability, and comply with applicable employment, data, consumer, and AI rules. CSR therefore operates alongside legal and regulatory accountability rather than as a purely voluntary mechanism; where high-impact AI causes discrimination, misinformation, privacy violations, or other harms, enterprise exposure may include remediation duties, regulatory sanctions, litigation, and reputational loss.

7. Policy Contexts and Governance Requirements

Enterprise institutional preparedness is shaped by a changing policy environment. The OECD AI Principles, updated in 2024, provide values-based guidance for trustworthy AI, including human-centered values, transparency, robustness, accountability, and inclusive growth (Organisation for Economic Co-operation and Development, 2024). These principles matter for enterprises because they frame AI actors as responsible participants in social governance rather than purely technical innovators. Enterprises preparing for AGI should translate these principles into internal standards for deployment approval, documentation, oversight, and stakeholder protection.

UNESCO's Recommendation on the Ethics of Artificial Intelligence provides a broader human rights and social justice framework. It emphasizes dignity, diversity, environmental sustainability, fairness, and inclusive governance (UNESCO, 2021). For enterprises, this implies that AGI preparedness should include equity and social responsibility, not just safety and compliance. A firm could comply with narrow technical rules while still producing social harm through exclusionary design, unequal access, or irresponsible labor practices. UNESCO's framework helps expand the moral horizon of enterprise governance.

The EU AI Act is particularly important because it shifts AI governance toward binding risk-based regulation. It regulates prohibited practices, high-risk systems, transparency obligations, and general-purpose AI models (European Union, 2024). Even enterprises outside the European Union may be affected through supply chains, market access, and emerging global standards. The Act also signals that employment, education, essential services, and fundamental rights are high-sensitivity domains. Enterprise AGI governance should assume that future regulatory expectations will increasingly demand documentation, risk assessment, human oversight, and accountability.

NIST's AI Risk Management Framework already offers a substantial organizational governance architecture, not merely technical risk controls. Through Govern, Map, Measure, and Manage, it addresses leadership accountability and roles, staff training, human-AI configuration and oversight, risk documentation, and engagement with internal and external stakeholders (Tabassi, 2023). These established practices provide the foundation for enterprise AGI preparedness. The present framework integrates them across three substantive domains—organizational resilience, employee protection, and social responsibility—to examine the institutional demands of greater cross-functional autonomy, organizational dependence, and potentially cascading consequences.

NIST's Generative Artificial Intelligence Profile and ISO/IEC 23894 further operationalize risk management (Autio et al., 2024; International Organization for Standardization & International Electrotechnical Commission, 2023). These resources support procurement review, monitoring, impact identification, incident response, and documentation. Building on these standards, the proposed framework examines how firms can maintain operational continuity, meaningful worker agency, and social legitimacy together when increasingly general AI affects multiple business functions. Its prospective relationships provide a basis for subsequent empirical validation.

The policy landscape is also more diverse than a solely EU-U.S. comparison suggests. China combines innovation support with classified governance, data and content obligations, security assessment, and provider responsibilities for generative AI services (Cyberspace Administration of

China, 2023). Japan's AI Guidelines for Business emphasize a risk-based, human-centered approach that translates AI principles into responsibilities for developers, providers, and business users (Ministry of Economy, Trade and Industry & Ministry of Internal Affairs and Communications, 2024).

The Republic of Korea's AI framework legislation combines industrial development with transparency, safety, trust, and accountability requirements (Republic of Korea, 2025). Singapore has developed organization-facing governance frameworks and, in 2026, issued a framework for agentic AI that combines technical and non-technical risk controls while emphasizing continuing human accountability (Infocomm Media Development Authority, 2026). These approaches differ in legal force and institutional design, but together they show that enterprises operating across borders face heterogeneous documentation, transparency, oversight, and assurance expectations. This heterogeneity can raise legal-mapping, documentation, assurance, localization, and staff-training costs for multinational firms. Cross-border preparedness therefore includes the capacity to map overlapping requirements and avoid treating compliance with one jurisdiction as globally sufficient.

Labor policy is also relevant. ILO research on generative AI and jobs emphasizes that technology can augment or automate work depending on task structure and institutional context (Gmyrek et al., 2023). IMF and WEF reports similarly highlight labor-market exposure, skill change, and the need for reskilling (Cazzaniga et al., 2024; World Economic Forum, 2025). These policy discussions suggest that enterprises should not wait for governments to manage workforce transitions. Firms are direct sites of technological change and should therefore build internal labor protection mechanisms.

Business and human rights frameworks provide another governance layer. The UN Guiding Principles on Business and Human Rights require companies to conduct due diligence, address adverse impacts, and provide remediation where appropriate (United Nations, 2011). As AI systems become more general and autonomous, human rights due diligence may include assessment of algorithmic discrimination, labor displacement, privacy invasion, manipulation, and downstream social harm. Enterprise preparedness should therefore connect AI governance with existing compliance, ethics, and sustainability systems rather than create an isolated AI committee.

Overall, the policy context suggests three governance requirements. First, enterprises must institutionalize risk management. Second, they must preserve human oversight and employee protection in high-impact contexts. Third, they must accept broader social responsibility for increasingly general AI deployment. These requirements form the foundation for the integrated framework proposed below.

8. An Integrated Framework for Enterprise Institutional Preparedness

This section proposes an integrated framework for enterprise institutional preparedness in the AGI era. Its seven mutually reinforcing components are strategic foresight, governance architecture, risk and impact assessment, human oversight, employee protection, social responsibility, and continuous learning. Building on NIST AI RMF's organizational governance, human oversight, training, and stakeholder involvement (Tabassi, 2023), this paper advances an enterprise-level, anticipatory synthesis of three mutually dependent preparedness domains: the continuity and adaptability of organizational functions, the protection of worker agency and fair transition, and the maintenance of external social legitimacy. It further considers their interaction under prospective cross-domain AI autonomy and dependence, distinguishing what current governance evidence supports from what remains conjectural. These are proposed conceptual relationships, not empirically validated effects. The three dimensions describe substantive preparedness domains, whereas the seven components specify complementary governance arrangements through which firms can build and maintain that preparedness. Figure 1 visualizes the relationship between the three domains and seven preparedness components.



Figure 1. Integrated Governance Framework for Enterprise Institutional Preparedness in the AGI Era

The first component is strategic foresight. Enterprises should establish processes for monitoring advanced-AI capability trends, regulatory changes, stakeholder expectations, and sector-specific risks. Foresight should not be limited to technology departments. It should involve executives, employees, legal teams, risk officers, customers, and external experts. Scenario analysis can identify potential futures in which increasingly general AI accelerates productivity, disrupts employment, creates regulatory pressure, or produces social backlash. A prepared enterprise uses these scenarios to build flexible strategies.

The second component is governance architecture. Enterprises should create an AGI preparedness committee or integrate advanced-AI oversight into existing board-level risk and sustainability structures. The committee should define decision rights, risk appetite, escalation procedures, deployment thresholds, and accountability lines. It should also review major high-impact AI use cases before implementation. Governance architecture matters because the relevant risks are cross-functional. A system introduced in one department may affect legal exposure, employee relations, cybersecurity, brand reputation, and social trust.

The third component is risk and impact assessment. Enterprises should assess not only technical performance but also social consequences. A complete assessment should include safety, security, privacy, bias, labor impact, environmental cost, stakeholder trust, and long-term dependence. It should map who is affected, how harms may occur, and what mitigation measures are available. AI

accountability research emphasizes that audits and impact assessments can help close governance gaps when designed with organizational responsibility in mind (Raji et al., 2020; Selbst et al., 2019).

The fourth component is human oversight. Enterprises should define which decisions must remain human-led, which can be AI-assisted, and which may be automated only under strict conditions. Oversight should include authority to pause, reverse, or modify AI-driven actions. It should also include training for managers and employees so that oversight is informed rather than ceremonial. Human factors research on automation trust shows that users may overtrust or undertrust automated systems depending on design, experience, and institutional context (Lee & See, 2004; Shneiderman, 2020).

The fifth component is employee protection. Enterprises should build a worker transition system before increasingly general AI causes more extensive task and role disruption. This system should include skill mapping, reskilling pathways, internal mobility, consultation, psychological support, and appeal mechanisms for AI-influenced decisions. It should also define limits on employee monitoring and protect the right to human review. Employee protection is not only a moral issue; it is also a resilience issue because alienated and anxious employees are less likely to support responsible transformation.

The sixth component is social responsibility. Enterprises should integrate advanced-AI governance into CSR and sustainability reporting. They should disclose high-impact use cases, explain risk management measures, report labor transition policies, and provide stakeholder grievance mechanisms. CSR theory suggests that legitimacy depends on responsiveness to stakeholder expectations and social values (Carroll, 1991; Freeman, 1984). As AI becomes more general and autonomous, legitimacy will increasingly depend on whether firms can demonstrate that their technological power is governed responsibly.

The seventh component is continuous learning. AGI preparedness cannot be completed once. Capabilities, risks, and regulations will evolve. Enterprises should conduct post-deployment reviews, incident learning, external audits, red-team exercises, employee feedback sessions, and periodic policy updates. Continuous learning helps firms avoid both complacency and panic. It allows institutional rules to evolve with technological and social change.

The integrated framework can be summarized as follows: anticipate future scenarios, institutionalize governance, assess risks and impacts, preserve human oversight, protect employees, assume social responsibility, and learn continuously. Each component addresses a specific weakness in common AI adoption practices. Foresight counters short-termism. Governance counters fragmentation. Impact assessment counters narrow efficiency metrics. Oversight counters excessive automation. Employee protection counters social insecurity. CSR counters externalization of harm. Continuous learning counters governance obsolescence.

The framework can also be translated into observable readiness indicators. Illustrative readiness indicators include a scheduled scenario-review process for strategic foresight; a named accountable owner and escalation route for governance architecture; documented high-impact classification for risk and impact assessment; tested pause, override, and appeal rights for human oversight; training, consultation, and monitoring limits for employee protection; material-impact disclosure and grievance mechanisms for social responsibility; and documented incident review and policy updating for continuous learning. These indicators provide auditable signs of preparedness that future research can refine and validate.

Implementation can begin with a readiness audit. The audit should review existing AI policies, data governance rules, procurement contracts, employee consultation procedures, cybersecurity controls, and CSR reporting systems. The purpose is to identify gaps between current digital governance and the stronger institutional demands implied by AGI. For example, a firm may have a data privacy policy but no rule for human review of AI-mediated employment decisions, or it may have an ethics statement but no impact assessment procedure.

A second implementation stage is prioritization of high-impact use cases. Enterprises do not need to apply the same level of scrutiny to every AI tool. Systems that affect employment, safety, finance, healthcare, consumer rights, or public information should receive stronger review. This proportional logic reflects the risk-based direction of contemporary AI governance and helps firms allocate limited governance resources where they matter most (European Union, 2024; Tabassi, 2023).

A third implementation stage is institutional embedding. AGI preparedness should be embedded into existing routines such as procurement, internal audit, training, compliance, product approval, and stakeholder reporting. If preparedness remains a separate document, it will have little practical effect. Institutional embedding means that responsible AI checks become part of ordinary decision processes rather than extraordinary interventions.

A fourth implementation stage is stakeholder feedback. Employees, customers, suppliers, and affected communities can reveal risks that internal technical teams may overlook. Stakeholder feedback should therefore be treated as a source of governance intelligence. This is especially important for AGI because the effects of a general system may emerge in unexpected social contexts. Feedback mechanisms can include surveys, grievance channels, worker councils, customer disclosures, and independent advisory panels.

A fifth implementation stage is external assurance. Internal governance may be insufficient when incentives favor rapid deployment. Enterprises should consider third-party audits, certification, public reporting, and expert review for high-impact advanced-AI applications. External assurance cannot eliminate all risks, but it can increase credibility, identify blind spots, and strengthen accountability. This is consistent with broader trends in AI audit research and social responsibility reporting (Raji et al., 2020; Selbst et al., 2019).

Implementation should be proportionate to organizational capacity. For SMEs, the minimum core should be narrower but not optional: inventory high-impact AI uses, classify material risks, name an accountable senior owner, preserve human review for consequential decisions, conduct vendor due diligence, and maintain an incident-escalation procedure. Larger organizations can add dedicated committees, internal assurance functions, red-team programs, and more extensive stakeholder reporting. This tiered approach avoids assuming that every firm can maintain the same governance infrastructure while preserving essential controls.

This framework also clarifies the difference between enterprise AGI preparedness and ordinary digital transformation. Digital transformation often emphasizes data, platforms, process optimization, and customer experience. AGI preparedness emphasizes institutional capacity to govern highly general, autonomous, and socially consequential systems. It is not only a management project but also a social responsibility project.

9. Discussion: From Technological Adoption to Institutional Readiness

The analysis suggests that enterprises should move beyond a narrow adoption mindset. Many organizations approach AI as a tool to be acquired, integrated, and scaled. This mindset may be appropriate for limited applications, but it is insufficient for AGI. A general and autonomous system can influence multiple functions and stakeholder relationships at once. The enterprise must therefore ask institutional questions: Who authorizes deployment? Who monitors consequences? Who protects employees? Who communicates with stakeholders? Who is accountable when harm occurs?

The concept of institutional preparedness highlights the importance of preconditions. Enterprises often attempt to govern technology after implementation. In the AGI era, this sequence is risky. Governance should precede deployment, not follow crisis. Preparedness includes rules, roles, competencies, values, and feedback systems that enable responsible action. It also includes a willingness to refuse or delay deployment when risks cannot be managed. Responsible innovation sometimes requires restraint.

The three dimensions proposed in this paper are mutually dependent. Organizational resilience

without employee protection can become a managerial strategy for absorbing shocks at workers' expense. Employee protection without organizational resilience may fail if the firm lacks resources and adaptive capacity. CSR without internal governance may become symbolic. An integrated approach is therefore necessary. Enterprises should not treat resilience, labor protection, and social responsibility as separate departments. They are parts of the same institutional readiness problem.

This distinction also separates symbolic preparedness from substantive preparedness. Symbolic preparedness exists when a firm can point to an AI policy, ethics statement, committee, or nominal human-in-the-loop process but the responsible actors lack authority, resources, access to relevant information, or the ability to stop and reverse a decision. Substantive preparedness requires operational decision rights, documented escalation paths, protected worker and stakeholder voice, incident learning, and accountability that remains traceable when AI is supplied by an external vendor. The difference is therefore not whether governance artifacts exist, but whether they change organizational decisions under pressure.

Critical AI scholarship also draws attention to concentrated control over data, models, computing infrastructure, and deployment channels. When firms that benefit commercially from AI also define the internal rules by which it is evaluated, governance may become vulnerable to conflicts of interest and weak contestability (Crawford, 2021; Zuboff, 2019). Institutional preparedness therefore requires counterweights such as independent assurance, protected worker voice, documented escalation rights, and board-level accountability rather than reliance on corporate ethics statements alone.

Current AI incidents provide diagnostic illustrations of this gap. The iTutorGroup settlement showed how automated recruitment rules could produce discriminatory outcomes and that organizational responsibility remained with the employer (U.S. Equal Employment Opportunity Commission, 2023). In *Moffatt v. Air Canada*, the tribunal held the airline responsible for inaccurate information provided by its website chatbot, illustrating that delegation to an automated interface does not remove enterprise accountability (*Moffatt v. Air Canada*, 2024). These incidents illustrate governance mechanisms that could become more consequential as AI systems gain broader authority and organizational integration.

There is also a tension between competitive pressure and responsible preparedness. Firms may fear that cautious governance will slow innovation. However, the absence of preparedness may create larger costs through accidents, employee resistance, regulatory penalties, reputational harm, and loss of public trust. In high-impact technologies, governance can be a source of competitiveness because it supports reliability, legitimacy, and stakeholder confidence. The AI risk management literature increasingly shows that governance is not opposed to innovation; it is a condition for sustainable innovation (European Union, 2024; Tabassi, 2023).

Another implication concerns human agency. More autonomous AI may tempt enterprises to delegate not only tasks but judgment. Yet organizations are moral and legal actors because human beings remain responsible for their actions. If future AGI-like systems become decision partners, enterprises must preserve meaningful human participation. This means investing in employee capability, maintaining human review, encouraging dissent, and ensuring that automation does not silently restructure power relations. Human agency is not automatically preserved by keeping a human in the loop; it requires institutional design.

The paper also suggests that enterprise preparedness for increasingly general AI is not limited to large technology firms. Manufacturing firms, financial institutions, healthcare providers, logistics companies, education providers, and public contractors may all adopt increasingly general AI systems. Smaller enterprises may lack resources, but they still need basic governance tools such as vendor due diligence, employee consultation, data safeguards, and clear use policies. Policymakers, industry associations, and standards bodies can support SMEs by providing practical templates and shared resources.

10. Managerial and Policy Implications

The framework developed in this paper has several managerial implications. First, enterprise leaders should recognize that AGI preparedness is a board-level issue rather than a purely technical matter. Boards and senior executives should ask whether the organization has a clear inventory of high-impact AI use cases, whether responsibility is assigned for AI-related harm, and whether employees have meaningful mechanisms to challenge or appeal AI-influenced decisions. Without senior-level ownership, AGI governance may remain fragmented across departments.

Second, enterprises should connect preparedness for increasingly general AI with enterprise risk management. Many firms already have risk systems for cybersecurity, finance, compliance, supply chains, and business continuity. Risks from increasingly general AI should not be managed in isolation from these systems. A more general AI system may create financial, legal, operational, reputational, labor, and social risks at the same time. Integrating such risks into enterprise risk management helps avoid the common problem of treating AI as an innovation tool until a crisis reveals its systemic nature.

Third, human resource departments should play a central role. Increasingly capable AI is likely to affect job design, skills, performance evaluation, employee monitoring, training, and career development. HR should not merely implement training after technology has been deployed. It should participate in deployment decisions from the beginning. This includes evaluating the labor impact of high-impact AI use cases, negotiating transition plans, protecting worker voice, and designing reskilling pathways that are accessible to different employee groups (Brynjolfsson et al., 2023; Tambe et al., 2019).

Fourth, enterprises should build a formal AI literacy system. AI literacy should not be limited to technical employees. Managers, frontline workers, legal teams, HR specialists, and customer-facing employees all need role-specific understanding of current advanced-AI capabilities and limits, together with the prospective governance implications of more general and autonomous systems. Such literacy should include prompt use, output verification, bias awareness, data protection, escalation procedures, and ethical responsibilities. A workforce that understands the technology is more capable of using it responsibly and identifying risks early.

Fifth, enterprises should avoid using productivity as the only evaluation criterion. If advanced AI systems increase output while intensifying stress, reducing autonomy, or undermining trust, the organization may gain short-term efficiency but lose long-term resilience. Evaluation metrics should include employee well-being, quality of decisions, error correction, stakeholder satisfaction, security incidents, fairness indicators, and learning outcomes. This broader evaluation logic is consistent with the idea that technological capability must be aligned with institutional responsibility.

Implementation pathways should also reflect sector-specific exposure. Manufacturing firms should prioritize cyber-physical continuity, manual fallback, and supplier-system dependencies; financial firms should emphasize model-risk governance, traceable approval, customer redress, and regulatory documentation; technology firms should emphasize release gates, rapid capability reassessment, and third-party ecosystem controls. These are different operational paths to the same institutional preparedness principles rather than separate frameworks.

Sixth, policymakers should support enterprise preparedness through proportionate guidance, shared infrastructure, and targeted incentives. SMEs may lack the resources to design governance systems independently, so public agencies and industry associations can provide model policies, common risk-assessment templates, subsidized training, shared testing or assurance services, and sector-specific certification pathways. Support should prioritize high-impact uses and help smaller firms meet a minimum governance baseline rather than requiring them to reproduce the structures of large technology companies.

Seventh, professional capacity is itself a governance constraint. Many firms lack personnel who combine technical AI knowledge with legal, HR, risk, audit, and stakeholder expertise. Professional

associations, universities, and certification bodies can help develop role-based training for AI risk owners, auditors, HR specialists, procurement teams, and human reviewers. Industry-level pools of qualified experts may also allow SMEs to obtain external assurance without maintaining large internal teams. This intermediary layer can translate abstract principles into repeatable enterprise routines for oversight, worker impact assessment, incident response, and responsible procurement.

Finally, enterprises should communicate preparedness transparently. Responsible communication does not mean exaggerating technical capacity or promising risk-free systems. It means explaining where increasingly general AI is used or planned, what safeguards are in place, how employees and stakeholders can raise concerns, and how the organization will respond to harm. Transparent communication can build trust and reduce the social uncertainty associated with advanced AI deployment.

11. Limitations and Future Research Directions

This paper has several limitations. Most importantly, its empirical grounding comes from current generative AI, algorithmic management, and other advanced AI applications, which are not equivalent to AGI. The framework therefore combines empirically informed mechanisms with forward-looking elements. Risks such as bias, monitoring, opacity, overreliance, and accountability diffusion are grounded in current evidence; risks associated with sustained cross-domain autonomy and cascading organizational action remain prospective. This temporal gap limits generalizability and means that the framework should be treated as an analytical starting point rather than a verified model.

A second limitation concerns sectoral variation. Banks, hospitals, manufacturing firms, logistics platforms, technology companies, and media organizations do not face identical exposure. Future research could use matched comparative cases or sector-specific longitudinal studies to examine which preparedness components are activated by different combinations of safety sensitivity, regulatory intensity, labor exposure, and operational dependence. Such designs would move beyond generic sector labels and identify where the framework requires adaptation.

A third limitation concerns firm size and feasibility. Large firms may have dedicated compliance, cybersecurity, sustainability, and legal departments, whereas SMEs often depend on vendors and external expertise. Future studies should test the proposed minimum-core model for SMEs and compare the cost, implementation burden, and effectiveness of internal versus shared or third-party governance arrangements. This would help determine which components can be simplified without weakening essential accountability.

A fourth limitation concerns global diversity. Cross-regional research should compare not only formal regulation but also enforcement intensity, documentation requirements, worker protections, data and localization rules, liability allocation, certification practices, and organizational capacity. Comparing these dimensions across major regulatory environments would clarify how multinational enterprises reconcile inconsistent governance expectations and where cross-border compliance generates duplication or gaps.

A fifth limitation concerns employee experience. Future empirical work could combine employee surveys with interviews, governance-document analysis, or comparative case studies to examine whether formal safeguards are accessible in practice, especially for temporary, frontline, and lower-bargaining-power workers. Longitudinal designs could also test whether reskilling, appeal rights, and human review remain substantive after initial AI deployment rather than becoming symbolic over time.

Future research should also develop and validate metrics for institutional preparedness. The main challenge is that visible governance artifacts, such as committees or policy statements, may not measure effective authority or implementation. Scale development should therefore distinguish formal structures from operational capability and combine observable indicators, incident evidence, worker access to safeguards, and outcome-based measures. The illustrative indicators proposed in

this paper can serve as candidate items, but weighting, reliability, cross-sector comparability, and predictive validity require empirical testing.

Another promising direction is the study of failure cases. Much AI governance literature focuses on principles and best practices, but organizational learning often depends on understanding breakdowns. Future research could examine incidents where AI systems produced harmful decisions, how enterprises responded, whether employees had voice, and what institutional reforms followed. Such studies would clarify the difference between symbolic preparedness and effective preparedness.

Finally, future research should explore the relationship between AGI preparedness and competitive advantage. Responsible governance is often framed as a constraint on innovation, but it may become a source of trust, resilience, and legitimacy. Firms that can demonstrate robust preparedness may be better positioned to attract talent, satisfy regulators, retain customers, and participate in high-trust ecosystems. This possibility deserves more theoretical and empirical attention.

12. A Practical Enterprise Preparedness Checklist

For practical use, enterprises can translate the framework into a proportionate preparedness checklist. The checklist begins with an inventory of current and planned advanced-AI uses and then applies greater scrutiny as consequence, autonomy, and stakeholder exposure increase. Table 3 adds broad risk triggers, review cycles, and assurance expectations so that the framework can be used as a practical governance aid rather than as a purely descriptive list.

Table 3. Practical Checklist for Enterprise AGI Preparedness

Checklist Area	Key Question	Risk Grade /Trigger	Review Cycle	Assurance /Accountability
AI Inventory	Where are advanced AI systems used, piloted, or procured?	All uses; flag consequential or autonomous use	Quarterly for fast-changing use; at least annually	Named business owner
Impact Classification	Which uses affect employment, safety, rights, finance, health, or public information?	Low / medium / high based on consequence and autonomy	Before deployment and after material change	Risk or legal review for high-impact use
Accountability	Who owns the business decision, technical system, risk, and escalation path?	Medium-high when responsibility is distributed	At approval and at least annually	Traceable named owners; board/risk oversight for high-impact use
Employee Protection	Are workers informed, trained, consulted, and able to request human review?	High when AI affects hiring, evaluation, discipline, or termination	Before deployment and semi-annually for high-impact use	Human appeal route; worker/HR review
Failure Testing	Have error, bias, security, leakage, and downtime scenarios been tested?	High for safety-critical, rights-sensitive, or autonomous use	Pre-deployment; annual or after major updates	Independent testing or external assurance where material
Stakeholder Communication	Who communicates with affected stakeholders when material harm occurs?	Medium-high when external stakeholders may be affected	Annual readiness check and after incidents	Executive owner and documented incident protocol
Social Responsibility Reporting	Are material impacts, mitigation actions, and grievance outcomes documented?	High where impacts extend beyond the firm	At least annually	Internal audit; external assurance where warranted
Periodic Review	Are policies, training, controls, and vendor assumptions still current?	All uses; priority to rapidly changing/high-impact systems	At least annually; more often after capability or legal change	Risk committee or designated senior owner

A three-level grading logic provides a practical starting point. Low-risk uses are bounded productivity tools with limited consequence; medium-risk uses influence material business or stakeholder decisions but retain clear human control; high-risk uses affect rights, safety, employment, essential services, or operate with substantial autonomy. This three-level classification is designed as a practical governance heuristic and can be adapted to applicable law and sector-specific requirements.

Review intensity should increase with risk and capability change. High-impact applications warrant pre-deployment review, documented human oversight, failure testing, and independent or third-party assurance where internal independence is insufficient. Periodic review should occur at least annually and sooner after major model updates, incidents, changes in autonomy, or material regulatory change. This creates a repeatable governance cycle while allowing SMEs to use external assurance instead of maintaining large internal specialist teams.

13. Conclusion

This paper has examined enterprise institutional preparedness for a future in which AI systems become increasingly general and autonomous. It distinguishes current generative and agentic AI from AGI and uses present evidence as a diagnostic basis for anticipatory governance rather than as direct proof of AGI effects. The analysis focuses on three connected dimensions: organizational resilience, employee protection, and social responsibility.

The paper contributes to AGI social impact research by translating broad concerns into enterprise-level institutional questions. Its central argument is that preparedness links three levels of governance: continuity of organizational functions, preservation of worker agency, and external social legitimacy. These dimensions are interdependent. Operational resilience can be achieved at workers' expense, formal employee safeguards can remain symbolic without authority, and CSR commitments can lack credibility without internal accountability.

The proposed framework emphasizes seven components: strategic foresight, governance architecture, risk and impact assessment, human oversight, employee protection, social responsibility, and continuous learning. It also proposes illustrative indicators, a proportionate minimum core for SMEs, and a risk-graded enterprise checklist. Together, these elements translate the conceptual framework into proportionate governance practices across firms of different sizes and sectoral contexts.

If increasingly general and autonomous AI systems become embedded in enterprise decision processes, technological capacity alone will not determine whether deployment is responsible or resilient. Enterprises will also need institutions that preserve accountability, maintain human capability, protect workers, and address external social effects. Institutional preparedness is therefore best understood as a continuing governance capability rather than a one-time compliance exercise.

Data Availability Statement

No new data were generated or analyzed in this conceptual study.

Funding

This research received no external funding.

Conflicts of Interest

The author declares no conflict of interest.

Ethical Approval and Consent to Participate

Not Required

References

1. Amodei, D., Olah, C., Steinhardt, J., Christiano, P., Schulman, J., & Mané, D. (2016). *Concrete problems in AI safety* (arXiv:1606.06565) [Preprint]. arXiv. <https://doi.org/10.48550/arXiv.1606.06565>
2. Autio, C., Schwartz, R., Dunietz, J., Jain, S., Stanley, M., Tabassi, E., Hall, P., & Roberts, K. (2024). *Artificial intelligence risk management framework: Generative artificial intelligence profile* (NIST AI 600-1). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.AI.600-1>
3. Autor, D. H. (2015). Why are there still so many jobs? The history and future of workplace automation. *Journal of Economic Perspectives*, 29(3), 3–30. <https://doi.org/10.1257/jep.29.3.3>
4. Bostrom, N. (2014). *Superintelligence: Paths, dangers, strategies*. Oxford University Press. <https://global.oup.com/academic/product/superintelligence-9780199678112>
5. Brynjolfsson, E., Li, D., & Raymond, L. R. (2023). *Generative AI at work* (Working Paper No. 31161). National Bureau of Economic Research. <https://doi.org/10.3386/w31161>
6. Carroll, A. B. (1991). The pyramid of corporate social responsibility: Toward the moral management of organizational stakeholders. *Business Horizons*, 34(4), 39–48. [https://doi.org/10.1016/0007-6813\(91\)90005-G](https://doi.org/10.1016/0007-6813(91)90005-G)
7. Cazzaniga, M., Jaumotte, F., Li, L., Melina, G., Panton, A. J., Pizzinelli, C., Rockall, E. J., & Tavares, M. M. (2024). *Gen-AI: Artificial intelligence and the future of work* (Staff Discussion Note 2024/001). International Monetary Fund. <https://doi.org/10.5089/9798400262548.006>
8. Crawford, K. (2021). *Atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press. <https://doi.org/10.2307/j.ctv1ghv45t>
9. Cyberspace Administration of China. (2023). *Interim measures for the management of generative artificial intelligence services* [Regulation]. https://www.cac.gov.cn/2023-07/13/c_1690898327029107.htm
10. Dafoe, A. (2018). *AI governance: A research agenda*. Governance of AI Program, Future of Humanity Institute, University of Oxford. <https://www.governance.ai/research-paper/agenda>
11. De Stefano, V. (2019). “Negotiating the algorithm”: Automation, artificial intelligence, and labor protection. *Comparative Labor Law & Policy Journal*, 41(1), 15–46. https://cllpj.law.illinois.edu/archive/vol_41/
12. Department for Science, Innovation and Technology. (2023). *The Bletchley Declaration by countries attending the AI Safety Summit, 1–2 November 2023*. <https://www.gov.uk/government/publications/ai-safety-summit-2023-the-bletchley-declaration>
13. Department for Science, Innovation and Technology. (2024). *Frontier AI safety commitments, AI Seoul Summit 2024*. <https://www.gov.uk/government/publications/frontier-ai-safety-commitments-ai-seoul-summit-2024>
14. Ducheck, S. (2020). Organizational resilience: A capability-based conceptualization. *Business Research*, 13, 215–246. <https://doi.org/10.1007/s40685-019-0085-7>

15. Elkington, J. (1997). *Cannibals with forks: The triple bottom line of 21st century business*. Capstone. <https://books.google.com/books?id=QLg9PgAACAAJ>
16. Eubanks, V. (2018). *Automating inequality: How high-tech tools profile, police, and punish the poor*. St. Martin's Press. <https://books.google.com/books?id=dQjVzAEACAAJ>
17. European Union. (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence. *Official Journal of the European Union*. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>
18. Floridi, L., & Cows, J. (2019). A unified framework of five principles for AI in society. *Harvard Data Science Review*, 1(1). <https://doi.org/10.1162/99608f92.8cd550d1>
19. Freeman, R. E. (1984). *Strategic management: A stakeholder approach*. Pitman. <https://books.google.com/books?id=4PUJAQAAMAAJ>
20. Gmyrek, P., Berg, J., & Bescond, D. (2023). *Generative AI and jobs: A global analysis of potential effects on job quantity and quality* (ILO Working Paper No. 96). International Labour Organization. <https://doi.org/10.54394/FHEM8239>
21. Infocomm Media Development Authority. (2026). *Model AI governance framework for agentic AI*. <https://www.imda.gov.sg/-/media/imda/files/about/emerging-tech-and-research/artificial-intelligence/mgf-for-agentic-ai.pdf>
22. International Energy Agency. (2025). *Energy and AI*. <https://www.iea.org/reports/energy-and-ai>
23. International Organization for Standardization & International Electrotechnical Commission. (2023). *Information technology—Artificial intelligence—Guidance on risk management* (ISO/IEC 23894:2023). International Organization for Standardization. <https://www.iso.org/standard/77304.html>
24. Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1(9), 389–399. <https://doi.org/10.1038/s42256-019-0088-2>
25. Kellogg, K. C., Valentine, M. A., & Christin, A. (2020). Algorithms at work: The new contested terrain of control. *Academy of Management Annals*, 14(1), 366–410. <https://doi.org/10.5465/annals.2018.0174>
26. Lee, J. D., & See, K. A. (2004). Trust in automation: Designing for appropriate reliance. *Human Factors*, 46(1), 50–80. https://doi.org/10.1518/hfes.46.1.50_30392
27. Legg, S., & Hutter, M. (2007). Universal intelligence: A definition of machine intelligence. *Minds and Machines*, 17(4), 391–444. <https://doi.org/10.1007/s11023-007-9079-x>
28. Lengnick-Hall, C. A., Beck, T. E., & Lengnick-Hall, M. L. (2011). Developing a capacity for organizational resilience through strategic human resource management. *Human Resource Management Review*, 21(3), 243–255. <https://doi.org/10.1016/j.hrmr.2010.07.001>
29. Linnenluecke, M. K. (2017). Resilience in business and management research: A review of influential publications and a research agenda. *International Journal of Management Reviews*, 19(1), 4–30. <https://doi.org/10.1111/ijmr.12076>
30. Ministry of Economy, Trade and Industry, & Ministry of Internal Affairs and Communications. (2024). *AI guidelines for business (Version 1.0)*. https://www.meti.go.jp/english/press/2024/0419_002.html
31. *Moffatt v. Air Canada* (2024). 2024 BCCRT 149. British Columbia Civil Resolution Tribunal. <https://www.canlii.org/en/bc/bccrt/doc/2024/2024bccrt149/2024bccrt149.html>
32. Morris, M. R., Sohl-Dickstein, J., Fiedel, N., Warkentin, T., Dafoe, A., Faust, A., Farabet, C., & Legg, S. (2023). *Levels of AGI for operationalizing progress on the path to AGI* (arXiv:2311.02462)

- [Preprint]. arXiv. <https://doi.org/10.48550/arXiv.2311.02462>
33. OpenAI. (2025). *Preparedness framework (Version 2)*. <https://cdn.openai.com/pdf/18a02b5d-6b67-4cec-ab64-68cdfbddebcd/preparedness-framework-v2.pdf>
 34. Organisation for Economic Co-operation and Development. (2024). *OECD AI principles*. <https://www.oecd.org/en/topics/ai-principles.html>
 35. Porter, M. E., & Kramer, M. R. (2011). Creating shared value. *Harvard Business Review*, 89(1–2), 62–77. <https://hbr.org/2011/01/the-big-idea-creating-shared-value>
 36. Raji, I. D., Smart, A., White, R. N., Mitchell, M., Gebru, T., Hutchinson, B., Smith-Loud, J., Theron, D., & Barnes, P. (2020). Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency* (pp. 33–44). Association for Computing Machinery. <https://doi.org/10.1145/3351095.3372873>
 37. Republic of Korea. (2025). *Framework Act on the Development of Artificial Intelligence and the Creation of a Foundation for Trust* (Act No. 20676). National Law Information Center. <https://www.law.go.kr/lsInfoP.do?lsiSeq=268543>
 38. Rosenblat, A. (2018). *Uberland: How algorithms are rewriting the rules of work*. University of California Press. <https://www.ucpress.edu/books/uberland/hardcover>
 39. Russell, S. (2019). *Human compatible: Artificial intelligence and the problem of control*. Viking. <https://www.penguinrandomhouse.com/books/566677/human-compatible-by-stuart-russell/>
 40. Russell, S., & Norvig, P. (2021). *Artificial intelligence: A modern approach* (4th ed.). Pearson. <https://www.pearson.com/en-us/subject-catalog/p/artificial-intelligence-a-modern-approach/P200000003500/9780137505135>
 41. Selbst, A. D., Boyd, D., Friedler, S. A., Venkatasubramanian, S., & Vertesi, J. (2019). Fairness and abstraction in sociotechnical systems. In *Proceedings of the Conference on Fairness, Accountability, and Transparency* (pp. 59–68). Association for Computing Machinery. <https://doi.org/10.1145/3287560.3287598>
 42. Shehabi, A., Smith, S. J., Hubbard, A., Newkirk, A., Lei, N., Siddik, M. A. B., Holecck, B., Koomey, J. G., Masanet, E. R., & Sartor, D. A. (2024). *2024 United States data center energy usage report* (LBNL-2001637). Lawrence Berkeley National Laboratory. <https://doi.org/10.71468/P1WC7Q>
 43. Shneiderman, B. (2020). Human-centered artificial intelligence: Reliable, safe and trustworthy. *International Journal of Human–Computer Interaction*, 36(6), 495–504. <https://doi.org/10.1080/10447318.2020.1741118>
 44. Tabassi, E. (2023). *Artificial intelligence risk management framework (AI RMF 1.0)* (NIST AI 100-1). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.AI.100-1>
 45. Tambe, P., Cappelli, P., & Yakubovich, V. (2019). Artificial intelligence in human resources management: Challenges and a path forward. *California Management Review*, 61(4), 15–42. <https://doi.org/10.1177/0008125619867910>
 46. UNESCO. (2021). *Recommendation on the ethics of artificial intelligence*. <https://www.unesco.org/en/legal-affairs/recommendation-ethics-artificial-intelligence>
 47. United Nations. (2011). *Guiding principles on business and human rights: Implementing the United Nations “Protect, Respect and Remedy” framework*. https://www.ohchr.org/documents/publications/guidingprinciplesbusinesshr_en.pdf

48. U.S. Equal Employment Opportunity Commission. (2023). *iTutorGroup to pay \$365,000 to settle EEOC discriminatory hiring suit*. <https://www.eeoc.gov/newsroom/itutorgroup-pay-365000-settle-eeoc-discriminatory-hiring-suit>
49. Weick, K. E., & Sutcliffe, K. M. (2015). *Managing the unexpected: Sustained performance in a complex world* (3rd ed.). Wiley. <https://doi.org/10.1002/9781119175834>
50. Williams, T. A., Gruber, D. A., Sutcliffe, K. M., Shepherd, D. A., & Zhao, E. Y. (2017). Organizational response to adversity: Fusing crisis management and resilience research streams. *Academy of Management Annals*, 11(2), 733–769. <https://doi.org/10.5465/annals.2015.0134>
51. Wood, A. J., Graham, M., Lehdonvirta, V., & Hjorth, I. (2019). Good gig, bad gig: Autonomy and algorithmic control in the global gig economy. *Work, Employment and Society*, 33(1), 56–75. <https://doi.org/10.1177/0950017018785616>
52. World Economic Forum. (2025). *The future of jobs report 2025*. <https://www.weforum.org/publications/the-future-of-jobs-report-2025/>
53. Zuboff, S. (2019). *The age of surveillance capitalism: The fight for a human future at the new frontier of power*. PublicAffairs. <https://www.hachettebookgroup.com/titles/shoshana-zuboff/the-age-of-surveillance-capitalism/9781610395700/>