



The Age of Autonomous Intelligence

Part III: AI — Solution or Multiplier of Long-Term Sustainability Risks?



Priti Shokeen
PhD, Managing Director,
Sustainable Investment,
TD Asset Management Inc.



Jackie Cheung
Vice President,
Sustainable Investment
TD Asset Management Inc.

Part III of the series examines the value creation associated with the Artificial Intelligence (AI) revolution—who benefits, at what cost, and what this implies for governance, fairness, ethics, and the long-term sustainability of capital markets. Building on Parts I and II, the focus is on **the stakeholder group likely to be most impacted by widespread AI adoption: human employees.**

A Technological Shift with System-Level Consequences

We are witnessing a technological transformation widely viewed as one of the most consequential shifts in modern history. Rapid advances in AI project exponential gains in productivity, with implications reaching far beyond operational efficiency.

Increasingly, AI is also being framed as a potential solution to some of the world's most complex sustainability challenges, particularly climate change. Where the limits of human cognition and traditional computing have constrained progress, AI-enabled systems could unlock new efficiencies in power, transportation, and food production. Academic research suggests these advancements may accelerate both mitigation and adaptation efforts faster than previously anticipated.

Beyond environmental benefits, AI may also support social progress. Studies¹ highlight its potential to advance outcomes aligned with the United Nations Sustainable Development Goals, while optimism

continues to grow around its application in areas such as disease detection and treatment. Yet, even as enthusiasm builds, critical questions remain. As the global race toward artificial general intelligence accelerates, the long-term implications for investment returns, social stability, and sustainability are far from settled.

One immediate challenge stems from the rapid expansion of data centers and compute infrastructure. Global capacity is expected to double to roughly 200 gigawatts by 2030², bringing with it rising energy demand, potential upward pressure on electricity prices, and renewed geopolitical competition. These developments could also complicate progress toward emissions reduction targets. In our view, these challenges are not insurmountable. However, they highlight the trade-offs that must be carefully managed if short-term innovation is to translate into durable, long-term value creation.

The Central Question: Human Capital in an AI-Driven Economy

The more profound existential concern lies in the future of human employment and the deployment of human capital. While this issue has been widely debated, recent studies and scenarios have forced markets, policymakers, and investors to confront its implications more directly. We believe that the evolving role of human contribution—and the distribution of labour income—represents the most pressing issue for long-term consumption models and the effective functioning of capital markets.

If large-scale human replacement becomes central to the AI success narrative, it risks evolving into a systemic challenge. Historically, technological revolutions have disrupted labour markets while ultimately creating more jobs than they destroyed.

However, as discussed in [Part I of this series >](#), the prevailing AI narrative today increasingly centers on replacement rather than augmentation. Skills transferability is limited, and while humans continue to excel in areas such as judgement, creativity, and higher-order reasoning, even these advantages may be challenged as AI capabilities advance.

Although it is impossible to forecast unemployment levels, sectoral impacts, or regional outcomes with precision, the directional risks are increasingly evident. This suggests a need to shift the dominant narrative from AI-driven productivity gains toward a framework that explicitly prioritizes pro-worker AI outcomes.

Human

The Existential Problem: Speed and Timing

The primary risk we identify is one of timing. If AI adoption outpaces workers' ability to reskill or redeploy, the resulting disruption could have adverse consequences for economic systems in capitalist societies. While humans are typically adaptive and solution-oriented, the current environment is characterized by more questions than answers.

Debates surrounding responsible AI development, covering issues such as hallucinations, accountability, and governance, have persisted for more than a decade (see [Challenges and Risks in Generative AI: Considerations for Investors >](#)). Yet uncertainty remains regarding whether AI will encounter cognitive or practical limits. Given the scale of capital investment, market optimism, and geopolitical importance attached to AI leadership, current assumptions appear to favour continued, largely unconstrained development. This raises a fundamental issue of value distribution: who ultimately captures productivity gains, and at what social and economic cost?

As of March 2026, there is no consensus on AI's long-term labour market impact. Some market participants argue that job displacement will be gradual³ and that many professions are unlikely to be replaced in the near term. At the same time, other commentators acknowledge that AI is likely to exacerbate inequality, often citing K shaped outcomes^{4,5}. Reconciling these positions is difficult without explicitly acknowledging job losses or income deterioration as central to the discussion.

The declining share of labour in GDP⁶—driven by automation and weakened bargaining power—raises important questions about fiscal sustainability, government transfers, and social security systems, which already represent a growing share of public expenditure.

The Secondary Problem: Ethics and Ownership of Intelligence

A related ethical question concerns ownership: who does intelligence, and the value it creates, belong to? AI systems are trained on vast amounts of human-generated knowledge, yet the ownership of resulting intellectual property remains contested. Several technology firms have begun formalizing contractual arrangements to address this issue, including Anthropic AI's proposed \$1.5B copyright settlement in September 2025, the largest copyright recovery to date⁷. The debate, however, remains unresolved.

Fairness concerns of this nature are not new. Previous industrial revolutions often required workers to train replacements or offshore counterparts, contributing to downward pressure on wages and rising inequality. The distinction today lies in speed and scale. The “worker” being trained—an AI system—is faster, more focused, and operates continuously. The resulting margin expansion disproportionately benefits shareholders, senior executives, and, in some cases, private investors.



The Tertiary Problem: Corporate Governance and Culture

While **AI governance** > has been widely discussed, less attention has been paid to how corporate approaches to AI shape organizational culture, talent development, and human capital formation. If AI displaces entry-level and repetitive roles, hiring freezes and reduced early-career opportunities may disproportionately affect new graduates.

This raises questions about how future leaders will be developed without traditional cycles of mentoring, experiential learning, and progressive responsibility. Leadership styles, collaborative learning, workplace norms, and professional networks are all shaped through sustained human interaction. Removing early-career roles risks eliminating a critical rung on the career ladder, with second-order effects on creativity, engagement, and societal development.

Advocating for Action: Stewardship in the World of AI

Over the past two decades, rapid technological advancement has repeatedly introduced complex social and governance challenges, from data privacy and cybersecurity to ethical accountability and intellectual property theft. Historically, corporate and policy responses have tended to lag innovation, often emerging only after harm has occurred.

We argue that this cycle must be broken. In the case of AI, corporate and policy frameworks should precede, rather than follow, large-scale deployment to mitigate the most severe risks to capital markets and social stability. A pro-worker AI framework may offer a pathway to preserving resilient income and consumption patterns, which remain foundational to long-term market sustainability.

In our early engagements on AI, discussions were broad and fragmented, spanning adoption, governance, education, and controls. To sharpen our stewardship approach, we have focused board-level engagement on three core areas:

1. Whether boards possess the appropriate skills and knowledge to oversee AI-related risks, and whether effective oversight structures are in place
2. How AI adoption has reshaped organizational structure and culture
3. Whether AI-driven productivity gains can be reliably quantified



Board Skills and Oversight

The quality, composition and effectiveness of boards remain one of the core pillars in which we evaluate the corporate governance practices of portfolio companies. Naturally, evaluating whether boards possess relevant AI expertise—through training or experience—is a foundational step in understanding oversight quality. We seek to

understand how boards intend to evolve, including changes through appointments or retirements, and what skills are being prioritized. With the right skills in place, well-structured, largely independent and high-performing boards can help support the oversight of material risks, including financially material environmental and social risks.

Organizational Structure and Corporate Culture

The second line of inquiry related how AI has shaped investee companies gave varied responses in our experience. Some focused on discrete tools and efficiency gains, while others speak at a cultural level. Understanding how AI reshapes corporate culture has proven particularly insightful.

In some of our engagements with portfolio companies, we have seen a higher degree of emphasis on people-centric philosophies but supported or empowered by technology – and more specifically, AI. In some cases, such axioms are embedded into corporate descriptions and communications. These statements represent what MIT Sloan Professor Edgar H. Schein describes as cultural “artifacts”—visible expressions of organizational values⁸. However, looking at artifacts alone are insufficient, and management’s espoused values may also not align with employees’ underlying assumptions. To assess culture meaningfully, stewardship requires engagement across organizational levels, boards, management, and non-management employees, to compare perspectives and identify inconsistencies.

Schein’s model also emphasizes that culture is transmitted through shared learning and the passing on of solutions to new members. That is, this passing on of solutions is a testament to whether such shared solutions or values are valid. In an agentic AI environment, where AI increasingly substitutes for entry level positions, opportunities for this transmission diminish. The implications are twofold.

First: Accountability remains with human decision-makers, even as AI performs more analytical work

Second: The developmental feedback loop, where poor work becomes a teachable moment, is weakened.

While re-querying AI systems may be efficient, it removes the process of developing human capital, diminishing long-term organizational capability. More importantly, without structured opportunities to test, reinforce, and evolve shared assumptions, with human participants that can receive feedback and be held accountable, organizational culture may be at risk of erosion over time.

Oversight

Value Creation for Whom? Revisiting Stakeholder Theory

Although shareholder primacy has long dominated corporate strategy, market outcomes are shaped by broader societal assumptions. As discussed in [Part I, the “Polanyi Moment”](#) > highlights periods when market systems are forced to re-embed within social structures. While thinkers such as Adam Smith and Milton Friedman shaped prevailing views on profit maximization, these ideas have been challenged.

Stakeholder theory, advanced by Edward Freeman, argues that long-term value creation requires balancing the interests of employees, communities, and the environment alongside those of shareholders⁹. This perspective has gained traction in parts of the developed world, grounded in the belief that unchecked profit maximization may ultimately undermine shareholder value itself—particularly when shareholders are also workers through pension and retirement systems.

As markets continue to reward short-term AI-driven gains, the question arises whether companies are prioritizing rapid wins at the expense of labour, potentially increasing reliance on government transfers or universal income mechanisms.

For investors, understanding AI-driven productivity gains first requires reliable, standardized disclosure.

Standardized disclosures surrounding crystallized productivity gains would be helpful for investors to separate efficiencies already had, against hopeful anecdotes of potential use cases which are less certain. Without quantification, it is difficult to distinguish realized efficiencies from aspirational narratives.

Where gains are realized, shareholders are likely to capture a significant share through margin expansion and equity appreciation. Many management executives also participate through equity-based compensation plans, which are often explicitly tied to performance thresholds. In more extreme examples, some manager-owners have even asked shareholders to vote and approve significant share ownership grants tied to extraordinary performance hurdles, locking in at the very least the earning potential tied to productivity gains from AI if astronomical performance-hurdles are met.

This raises the question of whether employees—whose labour supports both production and consumption—should also participate in AI-driven value creation. Employees may face structural disadvantages: continued participation in equity plans is often contingent on employment, while many employees are excluded entirely based on seniority or eligibility.

AI Sustainability Trade-Offs Investors Can’t Ignore

Potential Upside



Productivity gains

Margin expansion

Faster innovation cycles

Competitive advantage

Potential Risks



Energy-intensive data centers

Rising electricity demand

Labour displacement

Governance and disclosure gaps

What Investors Should Look For



Quantified AI productivity disclosures

Evidence of worker augmentation

Board-level AI oversight

Human capital strategy

Reimagining Employee Participation

One potential solution is to reimagine equity participation through 'employee benefit trust'-type structures, allocating a portion of AI-driven productivity gains to current and former employees, irrespective of ongoing employment. Such structures would be distinct from traditional compensation plans and could be managed similarly to pension arrangements.

While not presented as a definitive solution, this approach illustrates how private-sector innovation could mitigate labour displacement risks more effectively than eventual government intervention.

Governance of such frameworks could also benefit from meaningful employee representation

at the board level, balancing shareholder and management interests. International examples, such as employee representation at the board level in German markets¹⁰, demonstrate the viability of such governing structures. More thinking would have to go into the actual level representation of 'employee representative' directors, but the idea would be to fairly represent the employee-class at the board, balancing the interests of shareholders and management. Having such representation at the board level would add appropriate perspectives to discussions and company strategy, which helps enhance the ability of the board to think long-term and manage risks.

Technological change has a history of provoking policy responses. As AI accelerates economic divergence, the implications may increasingly show up in fiscal decisions and monetary trade-offs. Our concluding article shifts the focus from innovation to policy—and what that shift could mean for markets. ■



Employee participation

Connect with TD Asset Management



¹ Vinues, R., Azizpour, H., Leite, I., Balaamm M., Dignum, V., Domisch, S., et al. (2020). The role of artificial intelligence in achieving the Sustainable Development Goals. Nature Communications, 11, 233.

² JII, 2026 Global Data Center Outlook Navigating AI demand, power constraints and global opportunities in 2026. January 2026.

³ Board of Governors of the Federal reserve System, Artificial Intelligence and the Labor Market: A Scenario-Based Approach, May 2025.

⁴ Macrosphere: Asian economics in a global context. Leif Eskesen, Chief Economist, CLSA. 23 March 2026.

⁵ BlackRock, Larry Fink's 2026 Annual Chairman's Letter to Investors, March 23, 2026.

⁶ Bureau of Labour Statistics, U.S. Department of Labor (March 5, 2026). Productivity and Costs, Fourth Quarter and Annual Averages 2026, Preliminary. (<https://www.bls.gov/news.release/pdf/prod2.pdf>)

⁷ NPR, Anthropic settles with authors in first-of-its-kind AI copyright infringement lawsuit. September 5, 2025.

⁸ MIT Sloan Management Review, January 1984.

⁹ R. Edward Freeman. Strategic Management: A Stakeholder Approach (1984).

¹⁰ Harvard law School Forum on Corporate Governance, How Law Firms Can Lead the Agentic AI Era — And What Clients Now Expect, March 2026.

Certain statements in this document may contain forward-looking statements ("FLS") that are predictive in nature and may include words such as "expects", "anticipates", "intends", "believes", "estimates" and similar forward-looking expressions or negative versions thereof. FLS are based on current expectations and projections about future general economic, political and relevant market factors, such as interest and foreign exchange rates, equity and capital markets, the general business environment, assuming no changes to tax or other laws or government regulation or catastrophic events. Expectations and projections about future events are inherently subject to risks and uncertainties, which may be unforeseeable. Such expectations and projections may be incorrect in the future. FLS are not guarantees of future performance. Actual events could differ materially from those expressed or implied in any FLS. A number of important factors including those factors set out above can contribute to these digressions. You should avoid placing any reliance on FLS.

TD Asset Management Inc. is a wholly-owned subsidiary of The Toronto-Dominion Bank.

®The TD logo and other TD trademarks are the property of The Toronto-Dominion Bank or its subsidiaries.