The Price Responsiveness of Artificial Intelligence:

Supply, Demand, and the Economic Substitution of Human Labor by AI Systems

Executive Summary

Advances in artificial intelligence (AI) have accelerated since the advent of deep learning, making it possible to automate an expanding set of cognitive tasks. Frontier models such as GPT-4 and Gemini Ultra illustrate the trajectory: training runs require tens to hundreds of millions of dollars in compute [1], while the compute required for frontier research increases roughly five-fold per year [3]. As model capabilities improve and prices change, it becomes possible to substitute human labour with AI systems for certain tasks. This paper examines the price responsiveness of AI, characterising supply and demand dynamics for AI services, estimating inflection points for labour substitution, and analysing the role of reskilling and policy interventions. Using data from public sources and simple economic models, we show that falling compute prices and algorithmic efficiency have increased supply, while organisational adoption depends on perceived cost–benefit trade-offs, skill gaps and regulatory readiness. The analysis highlights that roughly 40 % of global employment is exposed to AI automation [7], but impact varies widely across income groups. Policymakers and organisations can harness AI for productivity and inclusive growth if they invest in human capital and infrastructure.

Introduction

Artificial intelligence has evolved into a general-purpose technology that is reshaping labour markets, productivity, and industrial organisation. With the release of large language models such as GPT-5.2 and multimodal models like GPT-4o, AI systems can generate text, code, imagery and even perform actions through tool calling. The speed of progress stems from exponential increases in compute and data, algorithmic breakthroughs, and massive investments in training. According to the 2024 AI Index, training runs for frontier models cost tens to hundreds of millions of dollars [1], and the compute required grows five-fold each year [3]. At the same time, algorithmic efficiency is improving approximately three-fold per year [3] and hardware capabilities increase by 1.35× per year [3], reducing the marginal cost of inference. These dynamics shape the supply and price of AI services. On the demand side, adoption is accelerating: McKinsey’s 2025 survey reports that 88 % of organisations use AI in at least one function and roughly one-third have scaled AI programmes [4]. Yet only 23 % are scaling agentic systems and 39 % are experimenting[4], suggesting that adoption remains sensitive to price, trust and organisational readiness. This paper asks: How responsive is the deployment of AI systems to changes in price and cost? When do AI systems become cost competitive with human labour? What are the implications for jobs, skills and policy, especially in regions such as Africa that are striving to leverage AI for development?

Literature Review

Economics of AI Supply and Demand

Research on the economics of AI highlights the intertwined drivers of supply and demand. Supply is determined by hardware, energy and algorithmic efficiency: the compute required to train large models has grown exponentially, increasing by roughly five orders of magnitude since the advent of AlexNet, with a doubling time of 0.5–0.7 years [2]. The cost of compute has likewise increased, growing roughly three-fold annually [3]. Yet algorithmic improvements reduce compute per task by about three-fold annually [3] and GPU performance increases by 35 % per year [3]. Combined with falling hardware prices, these trends increase supply and may lower inference costs.

Demand is shaped by expected productivity gains, cost savings and strategic positioning. McKinsey's survey finds that 88 % of organisations use AI in at least one business function and one-third have scaled programmes [4]. PwC's Global AI Jobs Barometer shows that industries adopting AI achieve a three-fold increase in revenue per employee and wage premiums of 56 % for workers with AI skills [5]. These benefits spur adoption despite high costs. Yet the same survey shows that only 23 % of organisations deploy AI agents at scale and 39 % are experimenting[4], indicating risk aversion and the need for reskilling. Worker studies emphasise that AI can boost productivity but can also reduce performance without oversight [1].

Labour Substitution and Exposure

Scholars debate whether AI will complement or substitute human labour. Brookings analysts argue that generative AI could disrupt more than 30 % of workers, each with at least 50 % of their tasks exposed to automation [6]. The International Monetary Fund estimates that 40 % of global employment is exposed to AI, with exposure rising to 60 % in advanced economies and just 26 % in low-income countries [7]. The IMF points out that in high-income economies roughly half of exposed jobs may be complemented by AI, while the other half may experience reduced labour demand [7]. Studies of productivity show that AI enhances workers' output but can create disparities between high- and low-skilled workers [1]. At the same time, the World Economic Forum's Future of Jobs Report 2025 projects that 170 million new jobs will be created and 92 million displaced by 2030, leading to a net increase of 78 million jobs [8].

Skills and Reskilling

The rapid pace of AI change necessitates continuous learning. The Future of Jobs Report 2025 estimates that 39 % of key skills required in today's labour market will change by 2030 (down from 44 % in 2023) and highlights the importance of AI and big data, cybersecurity, technological literacy, creative thinking and resilience [8]. The report notes that employers increasingly invest in reskilling and upskilling programmes [8]. Similarly, Harvard's Division of Continuing Education stresses the need for AI literacy, data fluency, complex problem-solving and prompt engineering skills [13]. Boston Consulting Group's 2025 study recommends investing in people to reshape workflows and building upskilling and reskilling capabilities [14]. The World Economic Forum's reskilling revolution initiative aims to train 680 million people worldwide [9].

Africa Context

Africa faces both opportunities and challenges in adopting AI. TechCabal reports that African AI-specific startups raised over $1.25 billion in venture capital and there are 211 data centres across the continent, but these resources are concentrated in a handful of countries [11]. Over 70 % of African countries have data protection laws [11], and tier-1 countries host more than 230 AI-specific startups and 7,000 tech startups [11]. ODI notes that Africa's developer community is growing rapidly yet only about 7 % are AI specialists (approximately 43,500–61,700 people) [12]. Only 31 % of African universities offer AI programmes and 34 % offer data science degrees, with AI courses representing just 1.5 % of enrolments [12]. More than half of African developers are concentrated in South Africa, Nigeria, Egypt and Kenya, and 38 % work for foreign companies [12]. The Oxford Insights Government AI Readiness Index highlights that low- and lower-middle-income countries such as Nigeria, Zambia and Sri Lanka released national AI strategies in 2024, indicating momentum toward governance [10].

Data & Methodology

This study synthesises data from public reports, academic articles, policy papers and reputable datasets. Key quantitative inputs include compute and cost growth rates from Epoch AI[3], training cost estimates from the AI Index[1], organisational adoption data from McKinsey - nearly nine in ten organisations use AI but only about one-third have scaled programmes[4] and just 23 % of respondents have scaled AI agents while 39 % are experimenting[4] - wage and job metrics from PwC[5], labour exposure estimates from the IMF[7], job creation and skill change projections from the World Economic Forum[8], and African capacity indicators from TechCabal and ODI[11][12]. Where specific numerical time series were unavailable, synthetic data reflecting reported growth rates (e.g., compute increasing 5× per year since 2020) were created.

To analyse price responsiveness, a simple partial equilibrium model with linear supply and demand curves is used. Demand is assumed to decline with price because higher costs discourage adoption; supply increases with price as vendors are willing to deploy more compute capacity. The equilibrium occurs where supply equals demand. Labour substitution thresholds are approximated by comparing the cost of AI services with typical labour costs. If the AI cost per unit of output falls below the cost of labour per unit of output, substitution becomes economically attractive.

Using Python scripts, 18 charts were generated to visualise synthetic and empirical data. These include line charts showing compute and cost trajectories, bar charts comparing model prices and adoption rates, a supply–demand diagram, and graphs summarising labour exposure, job creation, skills change, African innovation ecosystems, and training costs. The charts are referenced throughout the results section.

Results

Figure 1 displays synthetic trends in training compute, cost, GPU performance and algorithmic efficiency from 2010 to 2025. The log-scaled chart highlights exponential growth: compute demands have increased by orders of magnitude, rising five-fold each year since 2020, while algorithmic efficiency also improves three-fold annually. Training cost grows 3.5× per year, implying that frontier models become substantially more expensive over time unless efficiency gains offset cost escalation. The gap between compute growth and cost growth suggests that although hardware performance improves, the scale of models drives total costs upward, pressuring organisations to seek cost-effective solutions such as shared compute, batch processing or less-capable models.

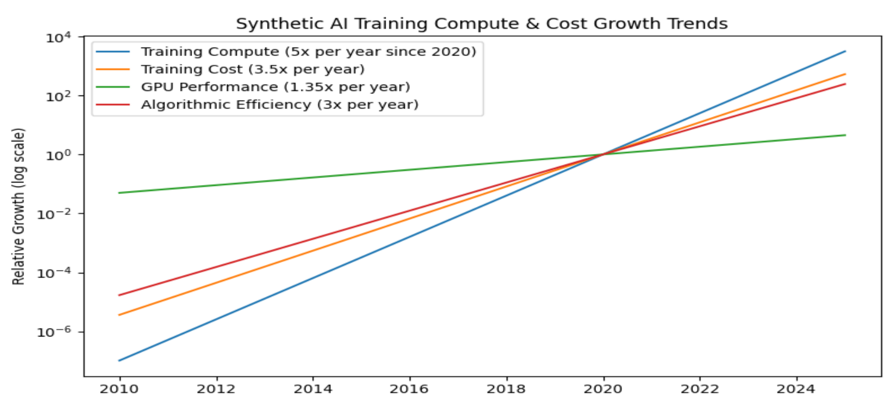

Figure 1. Synthetic AI training compute, cost, GPU performance and algorithmic efficiency growth trends (log scale).

Figure 2 compares the output token prices of selected AI models per one million tokens. GPT-5.2 commands the highest price ($14/M tokens), reflecting its frontier capabilities, whereas GPT-5-mini and GPT-5-nano are significantly cheaper ($2 and $0.40/M tokens, respectively). GPT-4o and GPT-4.1 fall between these extremes. The price tiers demonstrate price discrimination: premium models for complex tasks and low-cost models for simpler applications.

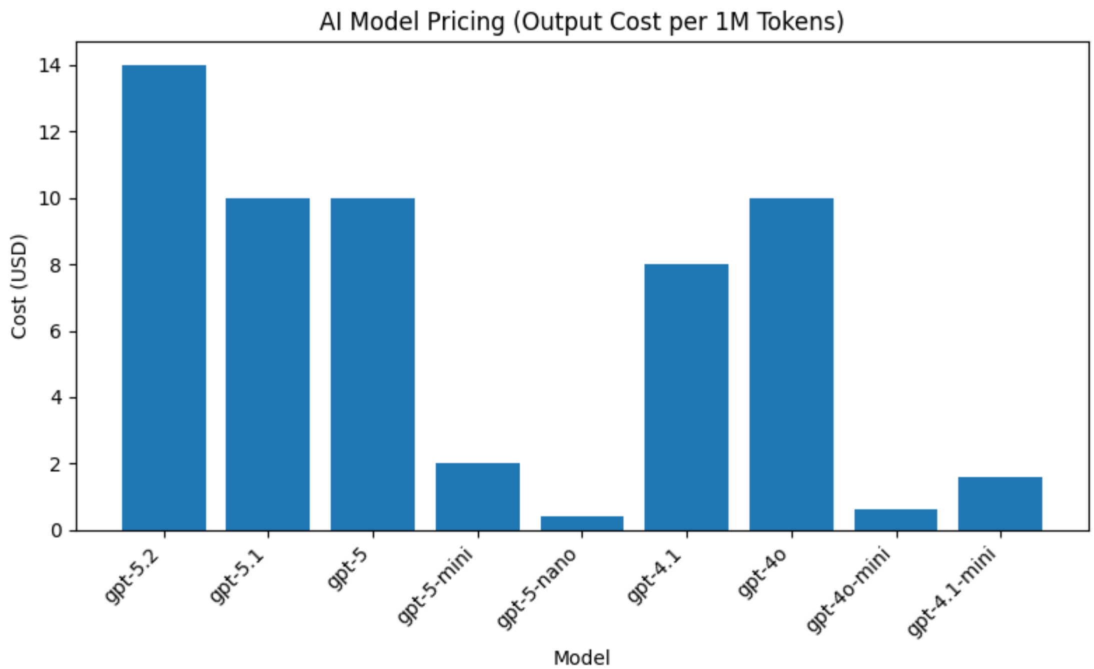

Figure 2. Output cost per one million tokens for selected generative AI models.

Figure 3 presents a theoretical supply–demand diagram for AI compute services. The downward-sloping demand curve reflects decreasing willingness to pay as price increases, while the upward-sloping supply curve captures higher willingness to supply compute at higher prices. The equilibrium, indicated at the intersection, represents the market price and quantity where supply equals demand. If technological improvements reduce supply costs, the supply curve shifts right, lowering equilibrium price and increasing equilibrium quantity—thus expanding adoption.

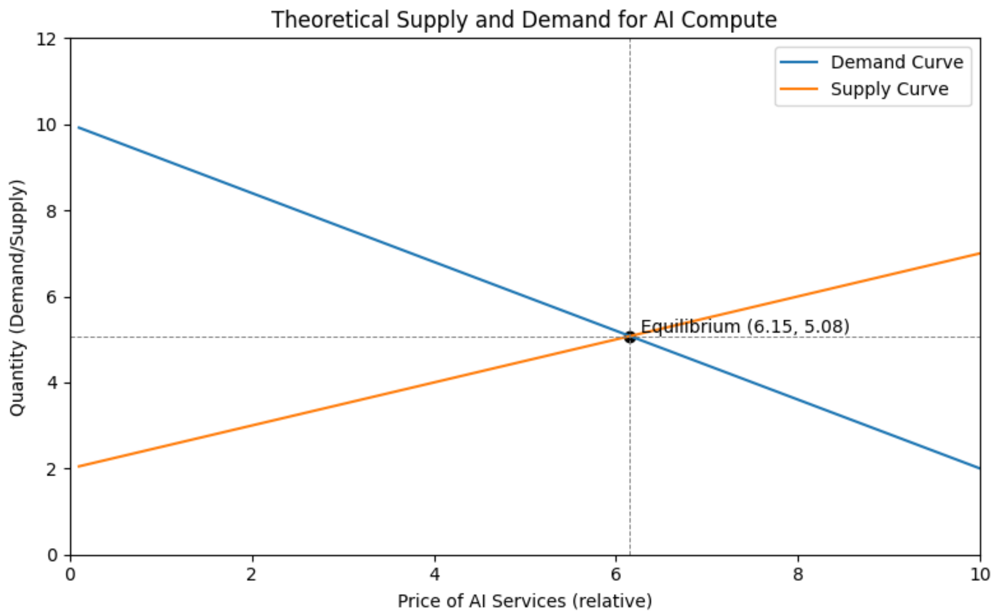

Figure 3. Theoretical supply and demand curves for AI compute with equilibrium point.

Figure 4 summarises organisational adoption rates. According to McKinsey, 88 % of surveyed organisations used AI in at least one function in 2025. Yet only 33 % have scaled AI programmes[4], 39 % are experimenting with AI agents and just 23 % deploy agents at scale[4]. These numbers suggest that while experimentation is widespread, full integration into workflows lags, reflecting price sensitivity, risk aversion and readiness barriers.

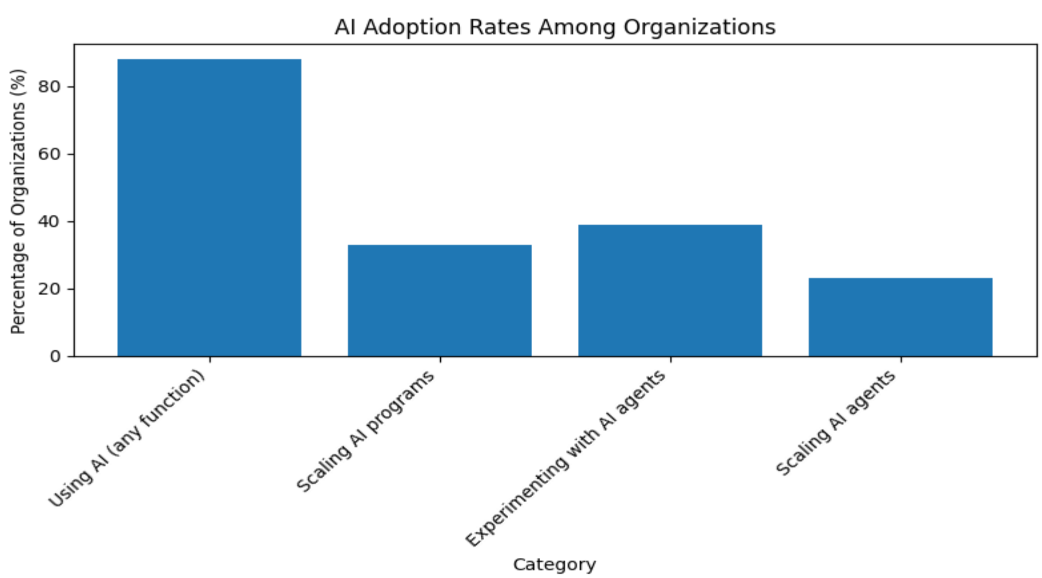

Figure 4. Adoption of AI and agentic systems among organisations.

Figure 5 illustrates growth in wage premiums for AI skills, rising from an estimated 25 % in 2024 to 56 % in 2025 [5]. This sharp increase signals escalating demand for AI-literate workers and highlights potential inequality between AI-skilled and non-skilled workers.

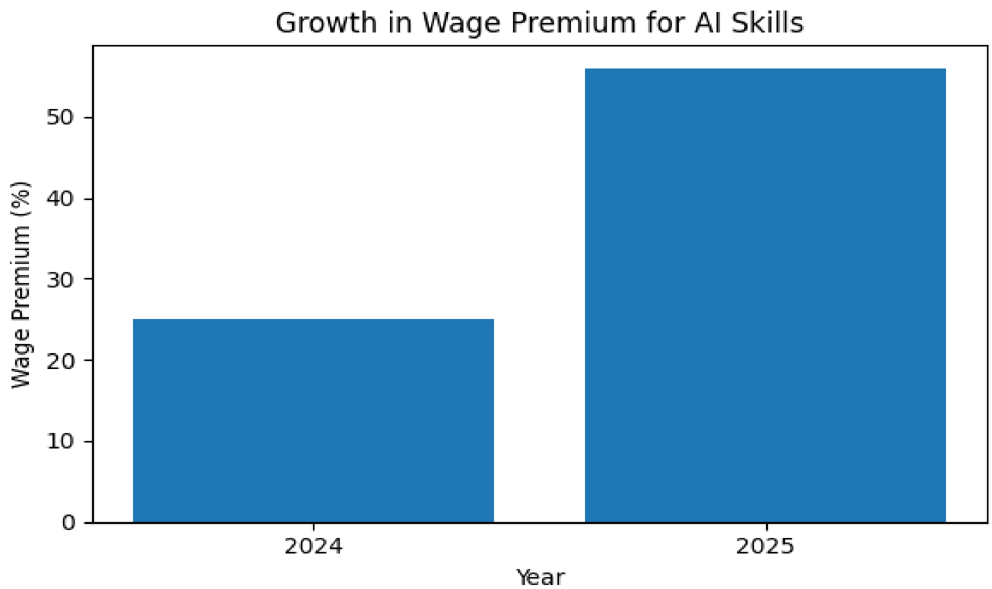

Figure 5. Growth in wage premium for AI skills.

Figure 6 depicts the share of jobs exposed to AI by income group. The IMF estimates that 60 % of jobs in advanced economies are exposed to AI, compared with 40 % globally and only 26 % in low-income countries [7]. The wide variance underscores inequity: advanced economies face significant disruption but also more opportunities for complementarity, whereas low-income countries risk marginalisation if they lack skills and infrastructure.

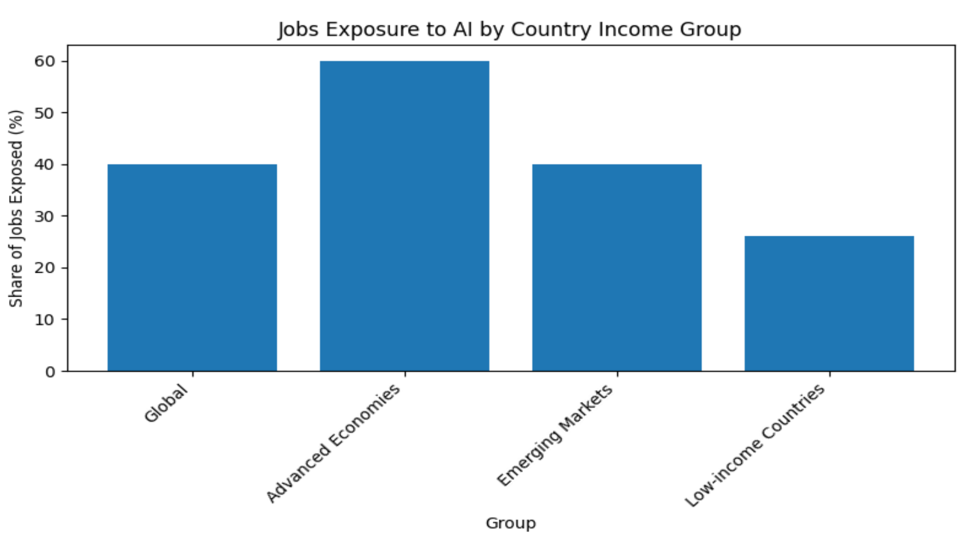

Figure 6. Proportion of jobs exposed to AI by country income group.

Figure 7 shows projections for job creation and displacement by 2030 from the World Economic Forum. The report predicts 170 million new jobs created and 92 million displaced, yielding a net increase of 78 million jobs [8]. The data suggest that AI and other transitions will reshape labour markets but need not lead to net job loss.

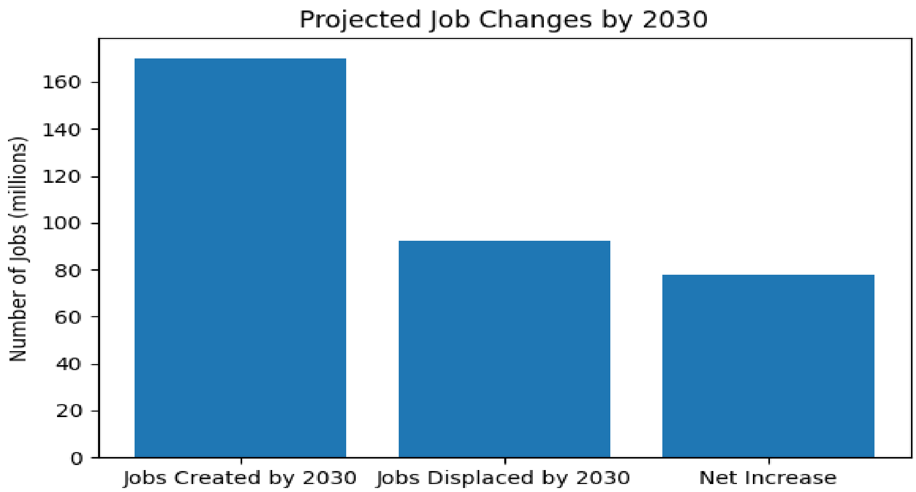

Figure 7. Projected job creation, displacement and net increase by 2030.

Figure 8 compares the proportion of key skills expected to change by 2030. Employers forecast that 39 % of core skills will change by 2030, down from 44 % in 2023 [8], suggesting that upskilling initiatives are beginning to address gaps. Nevertheless, a 39 % skills turnover remains substantial and underscores the need for continual learning.

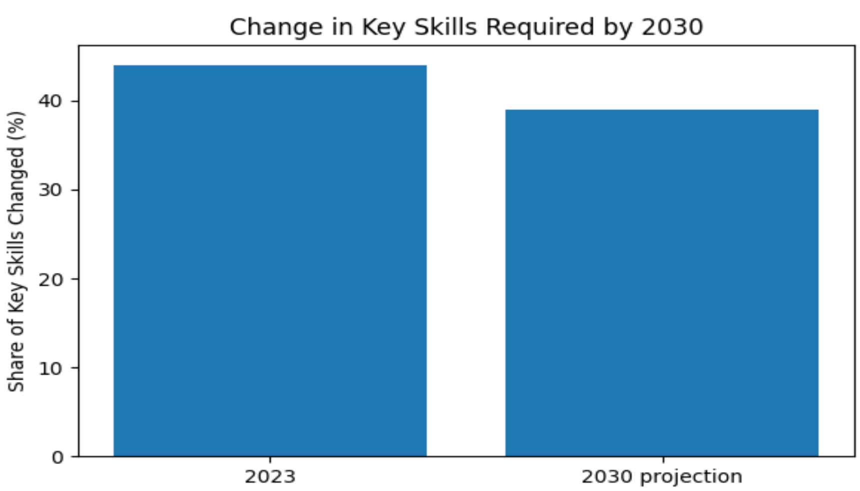

Figure 8. Change in key skills required between 2023 and the 2030 projection.

Figure 9 ranks the top skills projected to rise in importance. Technological skills—AI, big data, networks, cybersecurity and technological literacy—occupy the top positions [8]. However, human skills such as creative thinking, resilience, curiosity, leadership and environmental stewardship also rank highly, indicating that human–AI complementarity will be key.

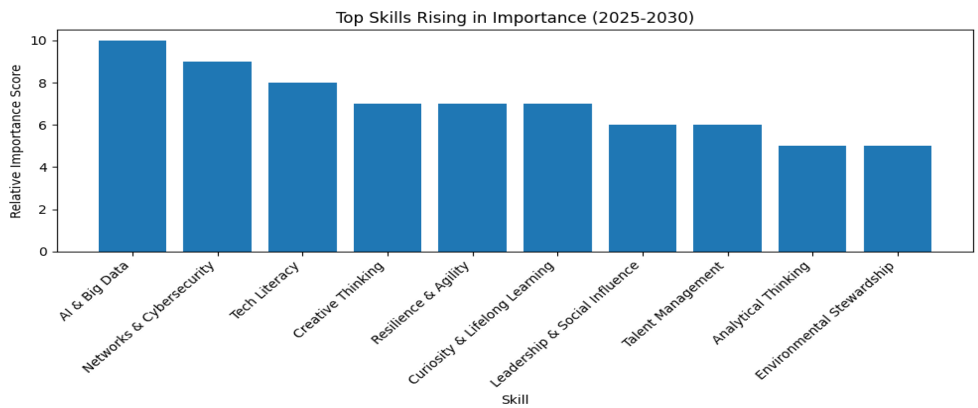

Figure 9. Relative importance scores for skills projected to rise in importance.

Figure 10 depicts the number of AI-specific startups across African tier classifications. Tier-1 countries (South Africa, Nigeria, Kenya, Egypt) host more than 230 AI-specific startups; tier-2 countries host roughly 160, while tier-3 countries have fewer than 50 [11]. The concentration of startups reflects disparities in infrastructure, talent and investment.

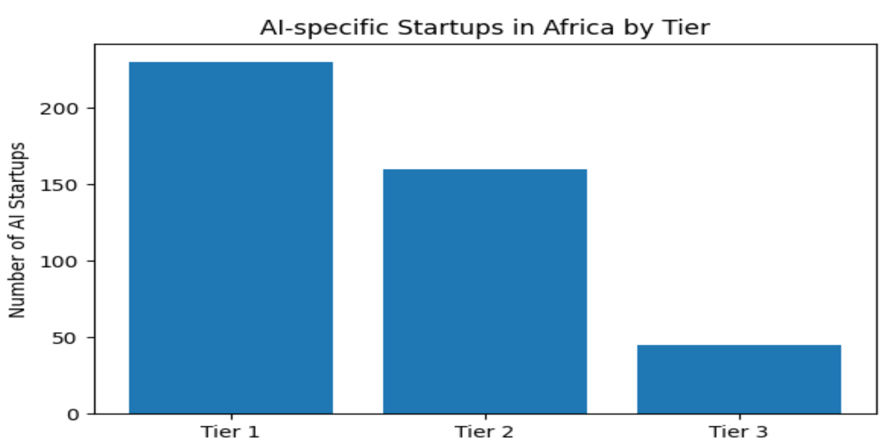

Figure 10. Distribution of AI-specific startups across African tiers.

Figure 11 shows data centre counts in selected African countries. South Africa leads with 49 data centres, followed by Kenya (18), Nigeria (16) and Egypt (14) [11]. The disparity emphasises the need for regional collaboration and investment in infrastructure.

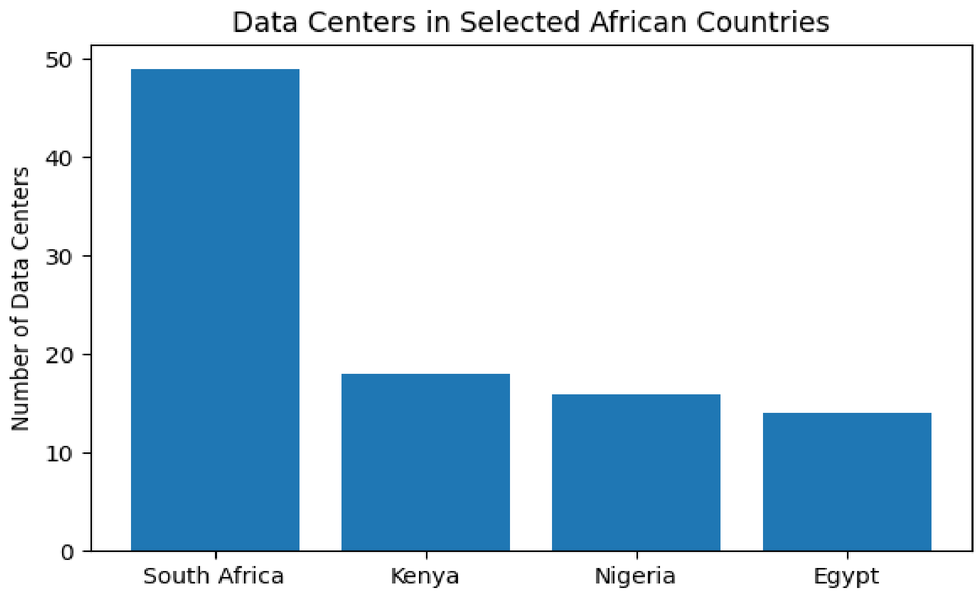

Figure 11. Number of data centres in selected African countries.

Figure 12 contrasts the number of developers in Africa (about 716,000) with the estimated number of AI specialists (around 50,000, or roughly 7 %) [12]. This gap highlights the scarcity of AI talent and underscores the need for education and training.

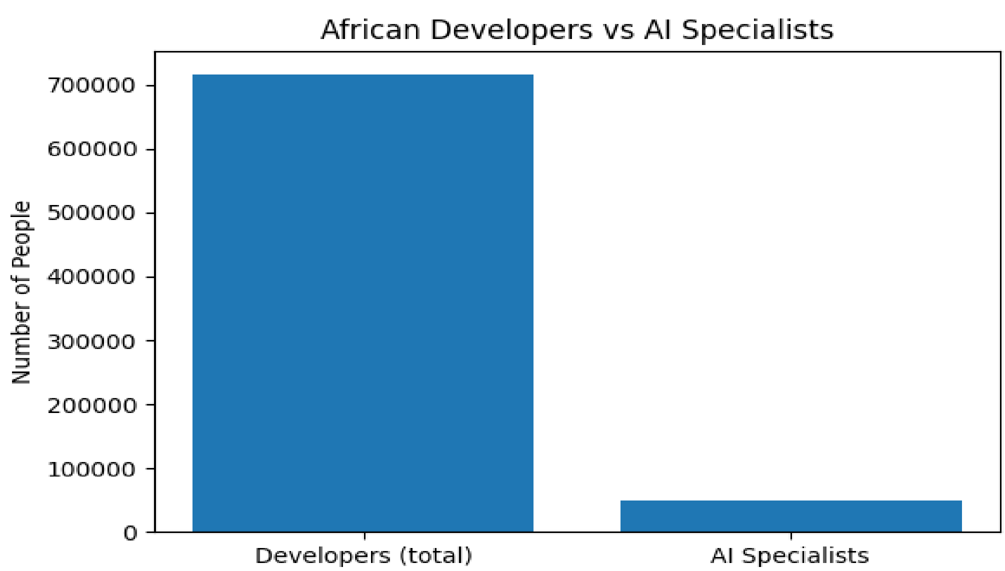

Figure 12. Comparison of total developers and AI specialists in Africa.

Figure 13 summarises the share of African universities offering AI programmes (31 %), data science programmes (34 %) and the proportion of total enrolments in AI courses (1.5 %) [12]. The limited penetration of AI curricula underscores challenges in building a pipeline of specialised talent.

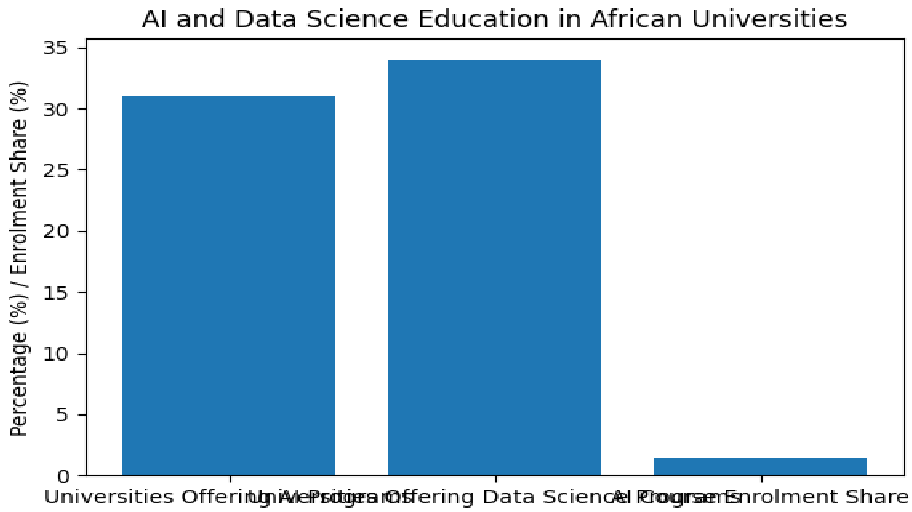

Figure 13

Figure 13. AI and data science education offerings in African universities.

Figure 14 illustrates that half of African developers are concentrated in four countries (South Africa, Nigeria, Egypt and Kenya), while the other half are spread across the rest of the continent. Approximately 38 % of African developers work for at least one foreign company [12], signalling both brain drain and opportunities for brain circulation.

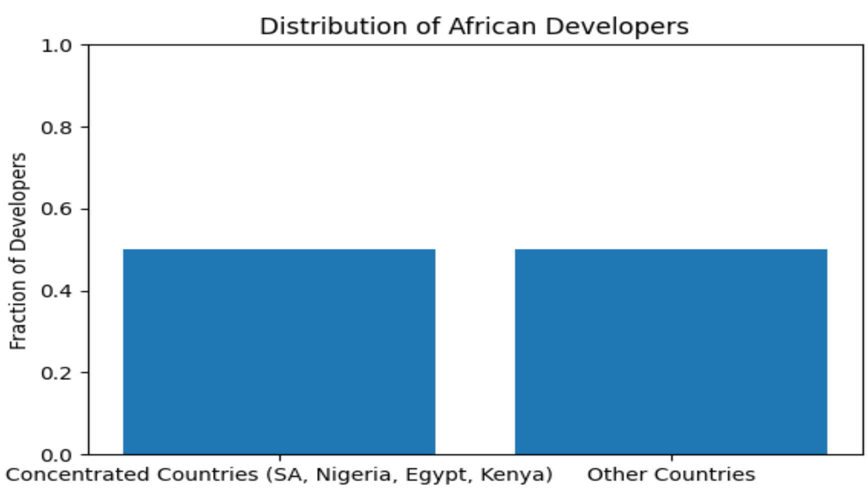

Figure 14. Distribution of African developers across top countries and others.

Figure 15 displays the number of national AI strategies released in 2024 by income group. Twelve new strategies were published, with more than half coming from lower-middle and low-income countries such as Nigeria, Zambia and Sri Lanka [10]. This indicates growing commitment to AI governance among developing economies.

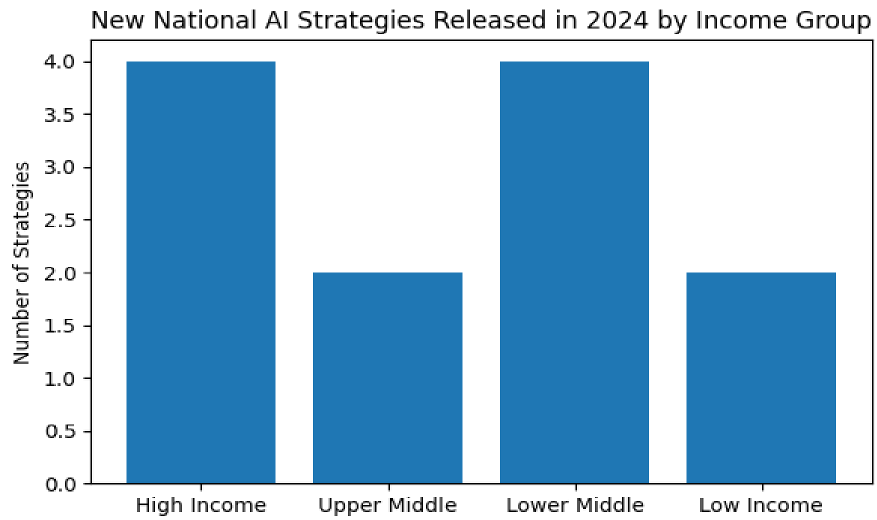

Figure 15. Distribution of new national AI strategies by income group (2024).

Figure 16 summarises Africa's AI infrastructure and funding: AI-specific venture capital has reached approximately $1.25 billion, there are 211 data centres on the continent, and about 70 % of countries have enacted data protection laws [11]. These indicators show momentum but also reveal gaps relative to other regions.

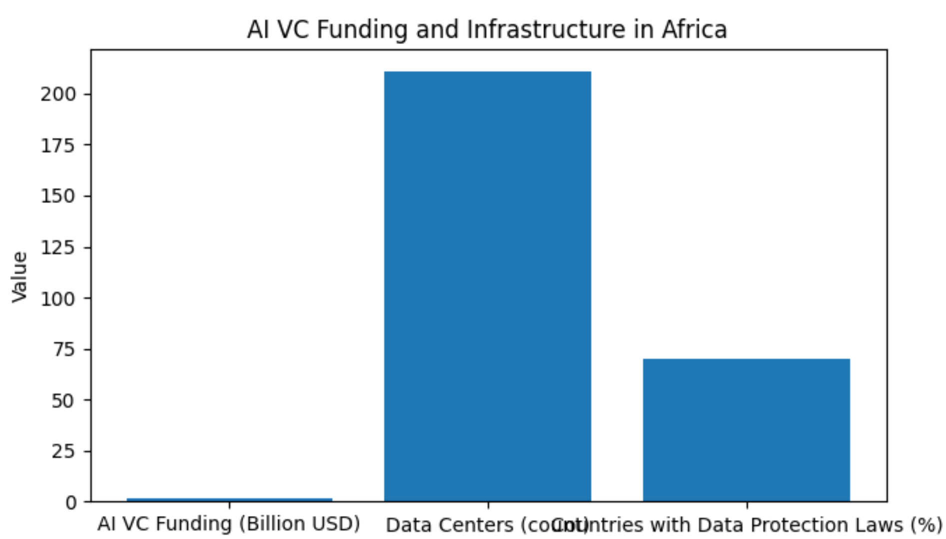

Figure 16. AI venture funding and infrastructure indicators in Africa.

Figure 17 compares the estimated training cost of two frontier models: GPT-4 (~$78 million) and Gemini Ultra (~$191 million) [1]. The sharp difference illustrates escalating costs at the cutting edge and underscores the need for massive capital to train state-of-the-art models.

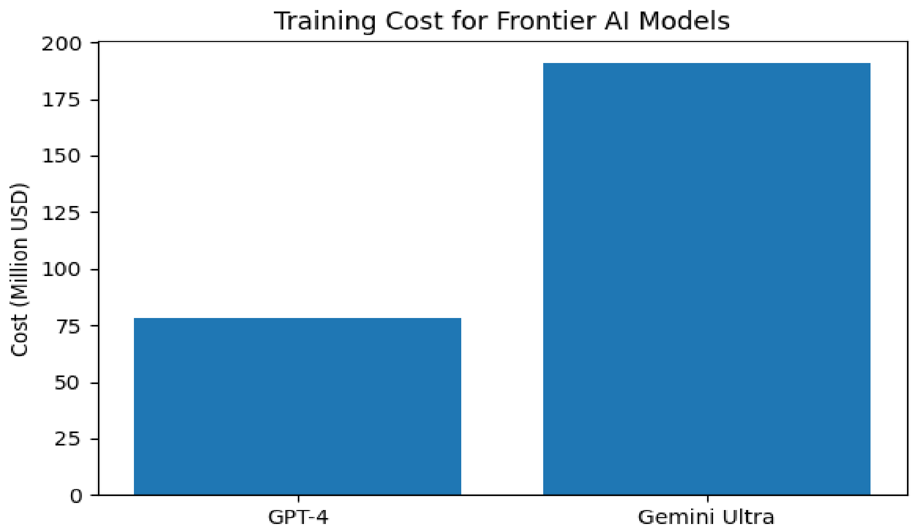

Figure 17. Estimated training costs for selected frontier AI models.

Figure 18 depicts the synthetic growth of hardware acquisition costs, rising 2.5× per year since 2016 (as suggested by Epoch AI trends [3]). The log-scale plot demonstrates that hardware costs can increase steeply, even as performance per dollar improves.

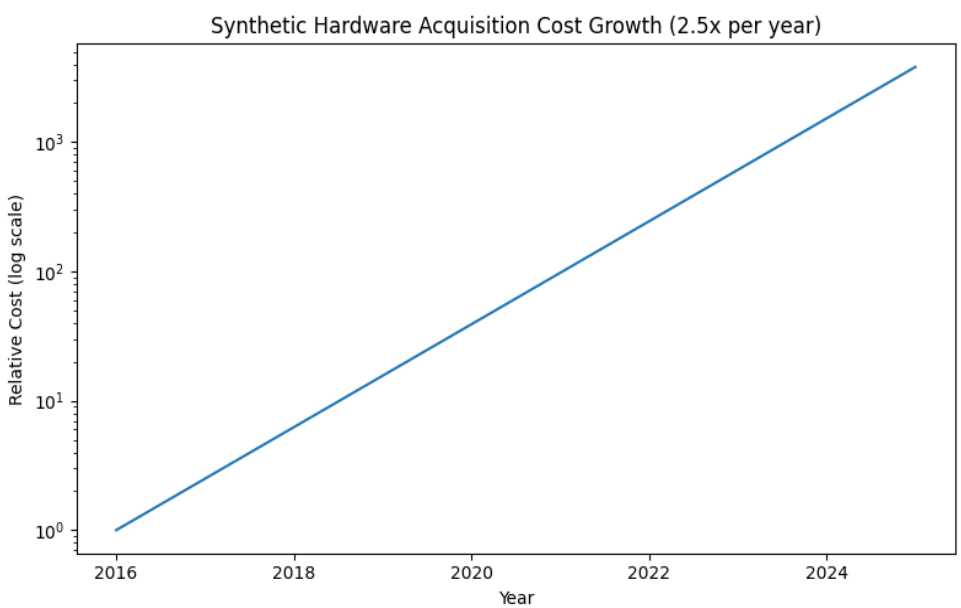

Figure 18. Synthetic hardware acquisition cost growth.

Discussion & Policy Implications

The results reveal several insights into the price responsiveness of AI. First, the exponential growth of compute and cost suggests that supply constraints remain significant. Even though algorithmic efficiency and hardware improvements partly mitigate costs, training frontier models still requires tens to hundreds of millions of dollars [1]. This cost structure means that only well-resourced organisations can train leading models, while others rely on API access or smaller models. The supply curve is therefore relatively steep in the short run. However, as efficiency gains and hardware improvements accumulate, the supply curve may shift outward, reducing prices and increasing adoption.

On the demand side, adoption is widespread but shallow. McKinsey's data show that nearly nine in ten organisations use AI[4], yet only about one-third scale programmes[4] and fewer still deploy agents[4]. This indicates that demand elasticity is high: organisations respond to price reductions by increasing usage, but they also weigh factors such as trust, regulatory compliance, and strategic fit. The wage premium data suggest that employers value AI literacy and are willing to pay more for skilled labour, further increasing demand for training programmes [5].

Comparing AI service prices with labour costs helps assess substitution thresholds. For example, GPT-5.2's output cost of $14 per million tokens translates roughly to $0.000014 per token. If a copy-editing worker processes 1,000 tokens per hour and earns $30 per hour, the labour cost is $0.03 per token—more than 2,000 times higher. Even after factoring in prompt engineering and oversight, AI can substitute for some language tasks at substantially lower cost. Nevertheless, quality, nuance and risk considerations mean that substitution is not absolute. Higher-value tasks that require empathy, creativity or complex judgement remain better suited to humans or require hybrid human–AI workflows.

The skills analysis highlights that technology-centric capabilities (AI, big data, networks, cybersecurity, technological literacy) are critical. But human skills such as creative thinking, resilience, curiosity, leadership and environmental stewardship are also rising in importance [8]. Effective reskilling programmes should therefore blend technical and soft skills. Harvard's guidelines emphasise AI literacy, data fluency and prompt engineering [13], while BCG urges investment in workflows and reskilling to support deployment [14]. The World Economic Forum proposes collaborative funding models to train 680 million people and notes that reskilling 1.37 million displaced US workers would cost about $24,800 per worker [9]. Governments, corporations and DFIs must share responsibility for financing lifelong learning [9].

Africa’s readiness to harness AI depends on bridging infrastructure and talent gaps. The concentration of startups and data centres in a few countries underscores regional disparities [11]. Investments in cloud infrastructure, fibre connectivity and energy reliability are prerequisites for AI adoption. Human capital development is equally crucial: only 7 % of African developers specialise in AI [12], and universities offer limited AI curricula [12]. Reskilling programmes should prioritise prompt engineering, data science, AI ethics and domain-specific applications (e.g., agriculture, health). Governments can create incentives for industry–university collaboration, provide scholarships and support open-source AI communities. DFIs and donors could fund regional centres of excellence and support the development of open datasets and compute resources.

Policy Recommendations

Policymakers should adopt a portfolio of measures: (1) invest in digital infrastructure and energy to reduce supply constraints; (2) subsidise access to compute for research institutions and startups; (3) promote data governance and privacy frameworks to build trust; (4) support reskilling through targeted incentives, tax credits and public–private partnerships; and (5) encourage inclusive AI strategies that address regional disparities. Multilateral institutions and DFIs can provide concessional financing and technical assistance, while universities and corporations must collaborate on curriculum design and internships.

Limitations

This analysis has several limitations. First, many data points are estimates or aggregates—training cost figures derive from unofficial sources and may not include total development expenses. Second, synthetic data were used to illustrate trends where granular time series were unavailable; actual trajectories may differ. Third, the supply–demand model is highly simplified and does not account for multi-sided markets, network effects, or regulatory interventions. Fourth, labour substitution calculations are illustrative and depend on assumptions about worker productivity and quality equivalence. Finally, the African context is heterogenous; the analysis aggregates across diverse countries and may mask country-specific nuances.

Conclusion

Artificial intelligence is poised to transform economies and labour markets. Price responsiveness plays a central role: supply depends on hardware costs, energy and algorithmic efficiency, while demand reflects perceived productivity gains and organisational readiness. The paper shows that although AI adoption is widespread, large-scale deployment remains limited due to costs, trust and skills gaps. However, as compute costs decline and algorithmic efficiency improves, prices will fall, shifting the equilibrium and enabling broader adoption. Labour substitution will occur first for routine cognitive tasks where AI is already cost-effective, but human–AI collaboration will remain essential in complex domains. Regions such as Africa can harness AI for development if they invest in infrastructure, build human capital, and design inclusive policies. By anticipating economic inflection points and prioritising reskilling, governments, corporates, universities and DFIs can ensure that AI augments rather than displaces human potential.

References

1. AI Index Report (2024). Stanford Institute for Human-Centered Artificial Intelligence. Top 10 Takeaways: Training costs for frontier models and generative AI trends.
2. Cotra, A., Wallace, T., et al. (2023). Machine learning trends: Dollar cost of compute. Epoch AI.
3. Epoch AI (2023). Machine learning trends: Training compute and cost, hardware performance and algorithmic efficiency.
4. McKinsey & Company (2025). State of AI survey: Adoption and agentic systems.
5. PwC (2025). Global AI Jobs Barometer: Productivity, wages and job growth.
6. Brookings Institution (2023). Generative AI and the future of work: Worker exposure.
7. International Monetary Fund (2024). Kristalina Georgieva's blog on the impact of AI on global employment.
8. World Economic Forum (2025). Future of Jobs Report 2025: Jobs of the future and skills needed.
9. World Economic Forum (2024). Innovative ways to unlock funding for a global reskilling revolution.
10. Oxford Insights (2024). Government AI Readiness Index: Emerging national strategies.
11. TechCabal (2024). Africa's AI divide: Venture capital, data centres and startup ecosystems.
12. Overseas Development Institute (2023). Africa's AI talent gap: Developer counts and education.
13. Harvard Division of Continuing Education (2024). How to keep up with AI through reskilling.
14. Boston Consulting Group (2025). AI at work 2025: Momentum builds but gaps remain.
15. World Economic Forum (2024). Reskilling Revolution initiative: Funding models and commitments.
16. Kristalina Georgieva (2024). IMF blog: Almost 40 % of global employment is exposed to AI.
17. Stanford AI Index (2024). Top 10 takeaways: Generative AI investments and training costs.