

Artificial Intelligence as a Productivity Shock: A Techno-Economic Analysis of India's AI Governance Framework and Inclusive-Growth Strategy

Author Name: **Naman Kamra**

Affiliation : Sri Guru Gobind Singh College of Commerce, Delhi University

Email: Namank1608@gmail.com

Corresponding Author: **Pari Rauthan**

Affiliation :Guru Gobind Singh Indraprastha University (GGSIPU)

Email: Parirauthan14@gmail.com

Corresponding Author: **Aditya Sharma**

Affiliation :Guru Gobind Singh Indraprastha University (GGSIPU)

Email: aditya.sharma1@timesinternet.in

Abstract

Artificial intelligence (AI) is increasingly framed by international development institutions not as a threat to employment in the Global South but as a productivity-enhancing general-purpose technology particularly suited to labour-abundant, capital-scarce economies. This paper undertakes a techno-economic analysis of India's emerging AI strategy, situating the IndiaAI Mission, the November 2025 India AI Governance Guidelines (the “Seven Sutras”), and associated sectoral deployments within the theoretical literature on general-purpose technologies, task-based automation-augmentation, and growth theory. Drawing on secondary evidence from the World Bank's *World Development Report 2026*, Government of India policy documents, and industry data, the paper argues that India's comparative advantage lies in factor augmentation rather than substitution: AI layered on existing digital public infrastructure is enhancing labour productivity in agriculture, education, and governance at population scale, while automation risk to Indian jobs remains comparatively contained. A Welch's independent-samples t-test on IMF AI Preparedness Index data for 174 economies confirms a large, statistically significant readiness gap between advanced and developing economies ($p < .001$, Cohen's $d = 2.62$), lending quantitative support to the paper's general-purpose-technology framing. The paper documents measured returns from case evidence in Odisha and Telangana, examines India's public investment in sovereign compute and data infrastructure, and identifies structural constraints -informal-sector opacity, energy and capital intensity, and enforcement gaps in a voluntary governance regime -that could limit inclusive diffusion. It concludes with policy implications for aligning India's “AI for All” vision with the Viksit Bharat 2047 growth trajectory.

Keywords: Artificial Intelligence; Digital Public Infrastructure; Factor Augmentation; AI Governance; IndiaAI Mission; Inclusive Growth

1. Introduction

1.1 Background and Motivation

Every general-purpose technology in economic history -the steam engine, electricity, the internal combustion engine, the semiconductor -has provoked the same anxious question: does it replace the worker, or does it extend what the worker can do (Bresnahan & Trajtenberg, 1995)? Artificial intelligence has reopened this question with unusual urgency, but the terms on which it is being answered differ sharply between rich and developing economies. Much of the global

commentary on AI and jobs has been written from the vantage point of high-income, capital-intensive labour markets, where the anxiety is that AI substitutes for cognitive, white-collar tasks that took a century of formal education to build.

India's policy establishment has, over the past two years, made a conscious argument that this framing does not travel well to a labour-abundant, capital-scarce economy, and that the more relevant economic question for India is not “how many jobs will AI destroy” but “how much can AI raise the output of the workers, farmers, teachers and clerks India already has, at a cost the state and citizen can actually afford.”

This argument has moved from rhetoric to institutional form with unusual speed. In March 2024, the Union Cabinet approved the IndiaAI Mission with an outlay of ₹10,371.92 crore over five years, structured around seven operational pillars spanning compute, data, indigenous foundation models, skilling, application development, safety and startup financing (Press Information Bureau, 2024). In November 2025, the Ministry of Electronics and Information Technology (MeitY) released the India AI Governance Guidelines, a principle-based, “techno-legal” framework anchored in seven guiding principles -the “Seven Sutras” -and recommending three new coordinating institutions (Ministry of Electronics and Information Technology, 2025).

In August 2026, the World Bank's *World Development Report 2026: The Promise of Artificial Intelligence* lent empirical weight to India's own framing, finding that jobs in high-income countries are roughly three times as likely to be at risk of automation from generative AI as those in low- and middle-income countries, while the productivity upside is nearly as large in developing economies as in rich ones (World Bank, 2026a). India featured prominently in that report as a proof-of-concept for low-cost, population-scale AI deployment in agriculture, health and governance.

1.2 Objectives of the Study

This paper has three objectives. First, it situates India's AI governance and mission architecture within the established economic literature on general-purpose technologies and task-based automation, so that the policy design can be read as an economic strategy and not merely an administrative one. Second, it assembles and interprets the available secondary evidence -international agency data, government progress reporting and independent industry data -on the measured economic returns to AI deployment in Indian agriculture, governance and education. Third, it identifies the structural and institutional constraints that could limit how evenly these gains are distributed across India's labour force, and draws out the policy implications for the country's stated *Viksit Bharat 2047* ambition.

1.3 Scope and Method

The paper is primarily a conceptual and descriptive-analytical study built on secondary data, supplemented by one quantitative inferential test. It does not generate primary survey or field data; rather, it triangulates Government of India policy documents (Ministry of Electronics and Information Technology notifications, Press Information Bureau releases), the World Bank's *World Development Report 2026*, and industry-body data (NASSCOM, NASSCOM-Zinnov) against the canonical economics-of-technology literature. This method is appropriate to a fast-moving policy domain in which the primary data available to any researcher, at this stage, is what governments and multilateral institutions have themselves measured and published; the paper's contribution is to organise that evidence within a coherent theoretical frame rather than to generate new field measurements.

Figures reported from Government of India sources are presented as officially reported and are cross-referenced against independent reporting wherever such reporting exists. Where country-level microdata are publicly available -specifically, the IMF's AI Preparedness Index covering 174 economies (Cazzaniga et al., 2024) -the paper moves beyond description to test, using Welch's independent-samples t-test with a Mann-Whitney U robustness check, whether the advanced/developing-economy split that motivates the paper's argument is statistically as well as descriptively real (Section 5.3).

2. Review of Literature and Theoretical Framework

2.1 AI as a General-Purpose Technology

Bresnahan and Trajtenberg (1995) characterised a general-purpose technology (GPT) as one marked by pervasiveness across sectors, inherent potential for technical improvement, and “innovational complementarities” -the property that the technology raises the returns to innovation in the sectors that use it. Electricity and the steam engine both required

decades of complementary reorganisation of factories, supply chains and skills before their productivity effects became visible in aggregate statistics, a lag that Bresnahan and Trajtenberg used to explain why transformative technologies often coincide with a temporary productivity slowdown rather than an immediate boom. AI has been widely read as a contemporary GPT (Aghion, Jones, & Jones, 2017; Brynjolfsson & McAfee, 2014), and this framing matters for India specifically because it implies that the payoff to AI adoption depends less on the technology itself than on the complementary infrastructure -data, connectivity, institutions, and skills -into which it is deployed. India's strategy, as elaborated in Sections 4 and 7 below, can be read as a deliberate attempt to shorten this complementary-investment lag by layering AI onto digital public infrastructure that already reaches population scale.

2.2 The Automation-Augmentation Dichotomy

The task-based framework developed by Acemoglu and Restrepo (2019) offers the paper's central analytical lens. In their model, technological change affects labour demand through two counteracting channels: a displacement effect, in which capital takes over tasks previously performed by labour, and a reinstatement effect, in which new task categories are created in which labour retains a comparative advantage. Automation, in this framework, always reduces the labour share of value added, even when it raises aggregate productivity; new-task creation always raises the labour share. Applied to AI, this framework predicts starkly different outcomes depending on whether AI is deployed to substitute for existing tasks (automation) or to expand the range of tasks a given worker or institution can perform (augmentation) - for example, enabling a single agricultural extension system to serve millions of farmers who previously had no access to any advisory service at all, rather than replacing extension officers who already existed. India's own policy framing - visible in the IndiaAI Mission's language of "AI for All" and repeated official emphasis on "enhancing, not replacing" -is, in this reading, an explicit bet on the augmentation channel, and Section 5 below tests that bet against the available cross-country evidence.

2.3 AI, Growth Theory and the Developing-Economy Case

In the neoclassical growth-accounting tradition following Solow (1956), sustained increases in output per worker require growth in total factor productivity (TFP), since capital deepening alone runs into diminishing returns. Aghion, Jones and Jones (2017) extend this reasoning to AI specifically, modelling it as a form of automation capable of raising TFP growth, while noting that the aggregate growth effect depends on how quickly the "bottleneck" tasks that remain non-automatable come to dominate the value of output. For a capital-scarce economy such as India, this literature carries a distinctive implication: because the marginal productivity of additional capital is comparatively high, and because a large share of the workforce operates with minimal capital-intensive tooling to begin with, even a modest AI-enabled productivity uplift per worker can translate into a proportionally large aggregate effect -the same logic that underlies the World Bank's finding that developing economies have comparatively more to gain from AI-driven productivity enhancement (World Bank, 2026a).

2.4 Digital Public Infrastructure as Economic Infrastructure

A distinct strand of the literature treats digital identity, payment and data-exchange rails as public infrastructure in the classical economic sense -non-rival, interoperable platforms that lower transaction costs across the economy (North, 1990, on the role of institutions in lowering transaction costs more generally). India's Aadhaar digital identity system, the Unified Payments Interface (UPI), DigiLocker and the broader "India Stack" are frequently cited as an applied instance of this logic (Nilekani, 2025). The relevance for an AI-economics paper is that these systems constitute exactly the kind of complementary infrastructure that general-purpose-technology theory identifies as a precondition for realising productivity gains: India's argument, made explicitly by the architect of Aadhaar and UPI, is that "AI makes DPI better" -that an AI layer placed on top of an existing, population-scale digital rail is a lower-cost and faster path to productivity gains than building either AI capability or digital infrastructure from scratch (Nilekani, 2025).

2.5 Research Gap

The literature reviewed in Sections 2.1–2.4 supplies the theoretical tools this paper uses, but each strand was developed with a different empirical setting in mind. The task-based automation-augmentation framework (Acemoglu & Restrepo, 2019; Autor, 2015) is calibrated almost entirely to United States and OECD labour-market data, where occupational classifications are granular and formal-sector employment is the norm. The growth-theoretic treatment of AI (Aghion,

Jones, & Jones, 2017) is likewise a general-equilibrium exercise pitched at the level of the world's largest economies rather than any single developing country's policy architecture. The IMF's cross-country preparedness data (Cazzaniga et al., 2024) and the World Bank's *World Development Report 2026* (World Bank, 2026a) go further by including developing economies, but both stop at the level of aggregate national indices and cross-country regressions; neither examines how a specific developing-country government has structured its compute, data, skilling and safety institutions in response to that aggregate picture. Studies of India's own digital public infrastructure (Nilekani, 2025) are, conversely, institutional and journalistic in character, describing what Aadhaar, UPI and DigiLocker do without connecting that description to the general-purpose-technology or task-based literatures reviewed above.

Three specific omissions follow from this. First, no study identified for this paper reads the IndiaAI Mission architecture and the November 2025 India AI Governance Guidelines through the automation-augmentation lens developed by Acemoglu and Restrepo (2019): the policy documents themselves are not written in economic language, and the economic literature has not yet been applied to them. Second, the cross-country automation-risk statistics that anchor the developing-economy optimism of the *World Development Report 2026* are, as Section 8.1 argues in detail, built substantially from formal-sector employment data, and no study identified for this paper has set out what this omission means for a majority-informal labour force such as India's specifically. Third, while the IMF's AI Preparedness Index has been used descriptively in policy commentary, it has not, to this paper's knowledge, been tested statistically as a two-group comparison and then read back against a single developing country's institutional strategy in the way Section 5.3 does below.

This paper's contribution is correspondingly threefold. First, it applies the task-based automation-augmentation framework and general-purpose-technology theory to a specific developing-country governance and mission architecture -India's Seven Sutras and the IndiaAI Mission's seven pillars -treating policy design itself as economic evidence rather than administrative background. Second, it supplies a country-level statistical test of the advanced/developing-economy AI-readiness gap that the existing literature has so far only described in aggregate terms, and connects that test back to the paper's institutional argument (Section 5.3). Third, it foregrounds the informal-sector measurement problem that is largely absent from a literature calibrated to formal-sector labour markets, offering an explicit qualification to the otherwise optimistic developing-economy reading of AI that this paper's own argument largely endorses (Section 8.1).

3. India's AI Governance Architecture: A Techno-Legal Framework

3.1 The Seven Sutras

The India AI Governance Guidelines, unveiled by MeitY on 5 November 2025 under the chairmanship of Professor Balaraman Ravindran of IIT Madras, are organised around seven guiding principles, described in the guidelines as “sutras”: Trust, People First, Innovation over Restraint, Fairness and Equity, Accountability, Understandable by Design, and Safety, Resilience and Sustainability (Ministry of Electronics and Information Technology, 2025). Several of these principles are adapted, with cross-sectoral generalisation, from the Reserve Bank of India's earlier Framework for Responsible and Ethical Enablement of Artificial Intelligence (FREE-AI), published in August 2025 for the financial sector (Saikrishna & Associates, 2025). The guidelines' own stated governing philosophy is captured in the phrase “Do No Harm,” and the framework explicitly favours innovation over restraint as an ordering principle -a sequencing choice with direct economic content, discussed further in Section 3.3.

Table 1. India AI Governance Guidelines -The Seven Sutras

No.	Sutra (Principle)	Core Emphasis
1	Trust	Foundational precondition for adoption and public confidence in AI systems
2	People First	Human-centric design; citizen welfare as the reference point for deployment
3	Innovation over Restraint	Ex-ante enablement of innovation prioritised over prescriptive restriction

4	Fairness & Equity	Non-discrimination and distributive justice in AI-enabled decisions
5	Accountability	Clear allocation of responsibility across developers, deployers and users
6	Understandable by Design	Transparency and explainability built into system design
7	Safety, Resilience & Sustainability	Risk mitigation, robustness and long-run sustainability of AI systems

Source: Compiled by the author from Ministry of Electronics and Information Technology (2025).

3.2 Institutional Architecture

The guidelines recommend three new coordinating institutions: an AI Governance Group (AIGG) to coordinate policy across ministries without itself acting as a sectoral regulator; a Technology and Policy Expert Committee (TPEC) to provide technical advice; and an AI Safety Institute (AIS), operating on a hub-and-spoke model, to conduct evaluation, testing, conformity assessment and standards-setting, and to maintain India's AI-incident database in line with OECD definitions (Ministry of Electronics and Information Technology, 2025). Sectoral regulators -the Reserve Bank of India, the Securities and Exchange Board of India, and the Insurance Regulatory and Development Authority of India among them -retain independent statutory authority within their domains; the AIGG's role is coordinative rather than supervisory.

3.3 A Regulatory-Economics Reading: Principle-Based versus Prescriptive Design

From a regulatory-economics standpoint, the choice of a voluntary, principle-based framework -as opposed to the European Union's prescriptive, tiered-risk model -is best understood as a deliberate trade-off between compliance cost and regulatory certainty. A prescriptive, ex-ante licensing regime of the kind embodied in the EU AI Act imposes fixed compliance costs that fall disproportionately on smaller firms and are more easily absorbed by large, capital-rich developers; a voluntary, principle-based regime lowers the fixed cost of entry for India's very large population of AI-enabled startups (Section 7.2) at the price of weaker ex-ante assurance and a greater reliance on existing sectoral law for enforcement.

The guidelines are explicit that they are advisory rather than legally binding, with no statutory penalties attached to non-compliance as of the time of writing (India AI Rulebook, 2026). This is not a design flaw so much as a revealed policy preference consistent with the "Innovation over Restraint" sutra: it prioritises the diffusion speed of a general-purpose technology -which, per Bresnahan and Trajtenberg (1995), is itself a first-order determinant of the technology's eventual productivity payoff -over the tighter ex-ante risk containment that a prescriptive regime would offer. The trade-off is discussed critically in Section 8.3.

4. The IndiaAI Mission: Public Investment as Infrastructure for Innovation

4.1 Outlay and the Seven Pillars

The IndiaAI Mission, approved by the Union Cabinet in March 2024 with a total outlay of ₹10,371.92 crore over five years, is organised around seven pillars: IndiaAI Compute, AIKosh (Data), Foundation Models, FutureSkills, Application Development, Safe and Trusted AI, and Startup Financing (Press Information Bureau, 2024). The mission's own stated objective -"Making AI in India, and Making AI Work for India" -captures a twin ambition of building sovereign technological capability while directing the bulk of near-term public spending toward deployment and diffusion rather than frontier research alone.


Table 2. The IndiaAI Mission -Seven Pillars

Pillar	Core Objective	Illustrative Progress Indicator
1. Compute	Affordable, subsidised access to GPU infrastructure	38,000+ GPUs empanelled by Dec. 2025; subsidy up to 40%
2. Data (AIKosh)	Sovereign, non-rival data and model repository	Sectoral datasets and pre-trained models on a single national platform
3. Foundation Models	Indigenous large and small language models	~20 companies supported in building indigenous LLMs/SLMs
4. FutureSkills	AI-ready human capital pipeline	Targets: 500 PhDs, 5,000 postgraduates, 8,000 undergraduates
5. Application Development	AI solutions for socially relevant problem statements	Deployments across agriculture, health and education
6. Safe & Trusted AI	Responsible AI R&D and risk mitigation	IndiaAI Safety Institute established; projects on deepfakes and bias
7. Startup Financing	Risk capital for AI-focused ventures	Dedicated financing window within the ₹10,371.92 crore outlay

Source: Compiled by the author from Press Information Bureau (2024) and IMPRI (2025).

4.2 Compute as a Quasi-Public Good

The Compute pillar is the clearest instance of the state acting as a provider of quasi-public infrastructure. Rather than each firm or research group independently importing and financing graphics-processing-unit (GPU) capacity, MeitY has empanelled cloud-service providers offering GPUs -including NVIDIA's H100, H200 and A100 chips, AMD's MI300X, Intel's Gaudi 2 and AWS Trainium -at subsidised rates of roughly ₹65 per GPU-hour, with the mission underwriting subsidies of up to 40 percent of the market rate. Empanelled capacity has scaled rapidly: from an initial target of 10,000 GPUs at the mission's approval in March 2024, to roughly 18,417 GPUs in early rollout, to 34,333 GPUs by May 2025, crossing 38,000 GPUs by December 2025 (IMPRI, 2025). Figure 1 traces this trajectory.


Figure 1. IndiaAI Mission -Empanelled GPU Compute Capacity, March 2024–December 2025

Source: Compiled by the author from IMPRI (2025).

The economic logic mirrors the classical case for public provision of infrastructure with high fixed costs and non-rival characteristics up to a point of congestion: individual startups and academic groups in India, particularly outside the largest metropolitan hubs, face capital-market frictions that make it difficult to finance GPU capacity independently,

and a shared, subsidised compute pool lowers the fixed cost of entry into AI application development in a manner analogous to the classical infant-industry argument for public investment in shared infrastructure.

4.3 AIKosh and Data as a Non-Rival Input

The AIKosh platform, launched in March 2025, functions as a sovereign data and pre-trained-model repository, aggregating datasets that were previously fragmented across ministries, state governments, hospitals and universities into a single access point, with built-in tools for fine-tuning (IMPRI, 2025). Because data, unlike physical capital, is non-rival in use -one model's access to a dataset does not diminish another's -centralising and standardising access to public datasets is a comparatively low-cost way for the state to raise the effective capital stock available to every downstream AI developer simultaneously, addressing what would otherwise be a significant market failure arising from data fragmentation and duplicated data-cleaning costs across the ecosystem.

5. The Economic Lens: AI as a Productivity Shock in a Labour-Abundant, Capital-Scarce Economy

5.1 Factor Augmentation versus Substitution

Section 2.2 introduced the Acemoglu-Restrepo distinction between automation (a displacement effect that reduces labour's share of value added) and new-task creation or augmentation (a reinstatement effect that raises it). India's official framing of its AI strategy leans explicitly and consistently toward the augmentation channel: successive government communications describe AI as “enhancing, not replacing” work, and frame the central Indian use case as closing service-delivery gaps in health, education and agricultural advisory at a population scale that human staffing alone could never reach, rather than displacing an existing, adequately staffed service. This is a testable economic claim, not merely a rhetorical one, and Section 5.2 sets out the cross-country evidence that bears on it.

5.2 Evidence from the World Development Report 2026

The World Bank's *World Development Report 2026: The Promise of Artificial Intelligence* provides the most direct empirical test available of the augmentation hypothesis for developing economies as a group. The report finds that 4.5 percent of existing jobs in low- and middle-income economies are at risk of automation by generative AI, compared with 14.2 percent in high-income economies -meaning jobs in high-income countries are roughly three times as exposed to automation risk as those in developing economies. At the same time, 16.2 percent of jobs in developing economies could see productivity meaningfully enhanced by AI, a figure close to the 18.7 percent estimated for high-income economies (World Bank, 2026a). Figure 2 presents this comparison.




Figure 2. Automation Risk and Productivity-Enhancement Potential: Developing versus High-Income Economies

Source: Compiled by the author from World Bank (2026a, 2026b).

The combination is analytically significant: developing economies face markedly lower automation-related displacement risk while capturing an almost equivalent share of the augmentation upside. In the Acemoglu-Restrepo terms of Section 2.2, this is precisely the profile one would expect where AI is deployed predominantly to extend the

reach of scarce professional and administrative capacity -doctors, teachers, extension officers, judicial staff -into populations that were previously unserved, rather than to displace an already-adequate stock of formal-sector labour. The World Bank's own report language urges developing economies to “tune out much of the noise” surrounding white-collar displacement fears that dominate discourse in richer countries, on the grounds that these fears are calibrated to labour markets structurally unlike their own (World Bank, 2026a).

5.3 A Statistical Test of the Global AI-Preparedness Gap

The general-purpose-technology argument developed in Section 2.4 implies a testable prediction: if complementary infrastructure is the binding constraint on realising AI's productivity potential, then countries classified as advanced economies should show measurably higher AI readiness than emerging-market and developing economies, independent of the automation/productivity split reported in Section 5.2. Unlike the World Development Report's two aggregate percentages, the IMF's AI Preparedness Index (APII) -a composite 0–1 measure spanning digital infrastructure, human capital and labour-market policy, innovation and economic integration, and regulation and ethics -is published at the level of individual economies, covering 174 countries for 2023 (Cazzaniga et al., 2024). This makes it possible to move from a two-point comparison to a proper inferential test. Countries were classified into Advanced Economies (AE) and Emerging Market and Developing Economies (EMDE) using the standard IMF World Economic Outlook grouping.

The null hypothesis (H_0) is that mean APII does not differ between AE and EMDE; the alternative (H_1) is that it does. Because Levene's test rejected the assumption of equal variances between the two groups ($F = 7.53$, $p = .007$), Welch's independent-samples t-test -which does not assume equal variances -was used as the primary test, with a Mann-Whitney U test run alongside it as a distribution-free robustness check on the same data.

Table 3. Independent-Samples t-Test: AI Preparedness Index, Advanced Economies versus Emerging Market and Developing Economies (2023)

Statistic	Value
Advanced Economies (AE)	$n = 38$; mean = 0.686; SD = 0.089
Emerging Market & Developing Economies (EMDE)	$n = 136$; mean = 0.407; SD = 0.111
Levene's test for equality of variances	$F = 7.53$, $p = .007$
Welch's independent-samples t-test	$t(72.71) = 16.17$, $p < .001$
Mean difference (AE – EMDE), 95% CI	0.279 [0.244, 0.313]
Effect size (Cohen's d)	2.62 (very large)
Mann-Whitney U (robustness check)	$U = 5,001$, $p < .001$

Source: Author's calculation from IMF AI Preparedness Index country-level data (Cazzaniga et al., 2024; data table as compiled in Zhu, 2024).




Figure 3. Distribution of AI Preparedness Index Scores, Advanced versus Emerging Market and Developing Economies
Source: Author's calculation from Cazzaniga et al. (2024); data table as compiled in Zhu (2024).

The test finds a large and statistically significant gap: advanced economies score, on average, 0.28 points higher on the 0–1 AIPI scale than emerging-market and developing economies ($t = 16.17$, $df = 72.71$, $p < .001$), and the effect size (Cohen's $d = 2.62$) is very large by conventional benchmarks. The Mann-Whitney U test, which does not depend on distributional assumptions, confirms the result ($p < .001$), so the finding is not an artefact of the parametric test's assumptions. This lends direct statistical support to the general-purpose-technology reading advanced in Section 2.4: the AE-EMDE gap is not confined to automation-exposure statistics but extends to the underlying complementary infrastructure -digital, human-capital and regulatory -that theory identifies as the actual precondition for converting AI into realised productivity gains.

India's own AIPI score of 0.49 is informative in this context. It sits above the EMDE group mean of 0.41, placing India at approximately the 72nd percentile within its income peer group despite ranking 72nd of 174 economies globally (Figure 3). Read alongside the large AE-EMDE gap documented above, this suggests India begins from a comparatively favourable position within the developing-economy cohort -consistent with the digital-public-infrastructure, demographic and startup-depth advantages catalogued in Section 7 -even though it remains well short of the advanced-economy benchmark on every AIPI sub-component.

5.4 Implications for India's Factor Endowments

India's economy remains, in relative terms, labour-abundant and capital-scarce: a very large share of the workforce is engaged in agriculture and informal services with limited capital intensity, while formal capital markets remain comparatively shallow relative to the size of the workforce. In a standard two-factor production framework, a technology that behaves as labour-augmenting -raising the effective productivity of each unit of labour without requiring proportional capital deepening -has a disproportionately large effect on output precisely in an economy of this kind, because the marginal product of the abundant factor (labour) is being raised directly. This is the theoretical basis for treating AI, when deployed through low-cost, shared digital infrastructure rather than expensive proprietary enterprise systems, as a productivity shock with an unusually favourable benefit-to-cost ratio for an economy structured like India's -a claim examined empirically through sectoral evidence in Section 6.

6. Sectoral Evidence: AI Deployment and Measured Economic Returns

6.1 Agriculture: Krushi Samruddhi and AI-Based Weather Advisory

Agriculture remains the largest single source of employment in India, engaging an estimated 45–50 percent of the workforce, which means that even small per-farmer productivity gains scale into a large aggregate economic effect. The Ama Krushi programme, launched in Odisha in 2018 and now operating under the name Krushi Samruddhi, provides AI-powered agricultural advisory services and has scaled from roughly 3 million farmers reached in 2022 to more than

7 million today (World Bank, 2026a). Figure 4 illustrates this growth alongside the programme's reported benefit-cost ratio of between 9:1 and 15:1.




Figure 4. Ama Krushi / Krushi Samruddhi -Farmer Reach in Odisha, 2022 versus Current

Source: Compiled by the author from World Bank (2026a).

A separate, independently studied case from Telangana offers a rare instance of AI's economic effect being measured through a controlled empirical design rather than programme self-reporting: an AI-based weather-forecasting system, evaluated by the University of Chicago's Climate and Energy Policy Institute, enabled smallholder farmers to adjust planting and investment decisions in response to monsoon forecasts issued roughly a month in advance. Some farmers raised pre-harvest expenditure by close to one-third in response to favourable forecasts, while others achieved net savings of up to USD 560 per farmer; the World Bank reports that affected households also reduced debt and increased the survival rate of newly established non-farm businesses (World Bank, 2026a). Because the farmers reached had historically limited access to formal government extension advisories, the intervention is best read as an instance of augmentation -creating an advisory capacity that had not previously existed for this population -rather than automation of an existing extension service.

6.2 Governance: From Digital India to "GovAI"

India's public-administration strategy explicitly positions AI as a layer placed on top of digital public infrastructure that the state already owns and operates -Aadhaar, UPI, DigiLocker, the CoWIN vaccination platform and the National Data and Analytics Platform -rather than as a parallel system built from scratch. Official framing describes this as a shift "from Digital India to Intelligent India," with the AI layer intended to move public administration from reactive to predictive, people-centric service delivery. The BHASHINI multilingual translation initiative, deployed across more than 500 government portals, is presented as the connective layer that allows India's linguistic diversity to be served without a proportional increase in multilingual human staffing -itself a clear example of augmentation, since it extends existing administrative capacity to language communities it previously could not efficiently serve. The stated economic rationale is a reduction in the transaction costs of accessing the state: faster grievance redress, more precisely targeted welfare delivery, reduced leakage through fraud detection, and administrative decision-making informed by real-time data rather than periodic surveys.

6.3 Education and Skilling: FutureSkills and YuvaAI for All

The government's education strategy rests on a similar augmentation logic, aimed explicitly at "personalisation at a scale no teacher-pupil ratio can match" -adaptive, local-language learning and automated tutoring intended to extend individualised instruction to student populations that India's teacher-to-student ratios could not otherwise support. On the supply side, the FutureSkills Prime platform has set targets of training 500 AI PhDs, 5,000 postgraduates and 8,000 undergraduates, while the broader YuvaAI for All initiative reports enrolment above 8.6 million with more than 2.8

million completions, supported by a reported network exceeding 500 AI and Data Labs established across states and Union Territories, and a newly announced Centre of Excellence for AI in Education with a budget outlay of ₹500 crore.

Read through a human-capital lens -Becker's foundational insight that education raises the stock of productive skill embodied in a worker -this pillar can be understood as the mission's attempt to ensure that the labour-supply side of the economy is upgraded in step with the demand-side deployment of AI tools described in Sections 6.1 and 6.2, without which the productivity gains documented elsewhere in this paper would be constrained by a shortage of AI-literate operators, teachers and administrators capable of using the tools effectively.

7. India's Structural Advantages for AI Diffusion

7.1 Digital Public Infrastructure as an AI Platform

As Section 2.4 argued theoretically, India's "India Stack" -the interoperable, open-API combination of Aadhaar (identity), UPI (payments), DigiLocker (documents), CoWIN and the Open Network for Digital Commerce (ONDC) - constitutes exactly the complementary infrastructure that general-purpose-technology theory identifies as the binding constraint on realising a new technology's productivity potential. Nandan Nilekani, widely credited as the architect of Aadhaar and UPI, has summarised the logic directly: AI does not need to be built on a green field in India, because it can instead be layered on the digital public rails India already owns -a framing that recasts India's AI strategy as an infrastructure-reuse problem rather than a from-scratch technology race (Nilekani, 2025).

7.2 Demographic Dividend, Startups and Global Capability Centres

India's working-age population remains, in relative terms, the youngest large-population cohort among major economies, and this demographic profile intersects with a startup ecosystem of roughly 1.8 million registered startups, of which some 89 percent of new startups are reported to have incorporated AI into their products or services. India also hosts one of the largest concentrations of Global Capability Centres (GCCs) in the world -captive innovation and delivery hubs operated by multinational enterprises -with independent industry data placing the ecosystem at more than 2,100 centres, employing over 2.3 million professionals and generating close to USD 98.4 billion in FY2026 revenue, up from roughly USD 19 billion in 2015 (PwC India, 2026; Business Standard, 2026). Figure 5 traces this growth trajectory, which functions in this paper as a proxy for the pace at which AI-enabled, higher-value digital work is scaling within the Indian economy: GCCs have moved, on the industry's own account, from cost-driven "delivery engines" toward AI-native "enterprise nerve centres" with growing product-ownership responsibility. India's broader technology sector, of which GCCs form a growing share, is projected by NASSCOM to cross USD 315 billion in revenue in FY2026 (NASSCOM, 2026).




Figure 5. India's Global Capability Centre (GCC) Sector Revenue Growth, 2015–FY2026

Source: Compiled by the author from PwC India (2026) and Business Standard (2026).

7.3 Towards a Sovereign AI Stack

Beyond deployment, the IndiaAI Mission's Foundation Models pillar is supporting an estimated 20 companies in building indigenous large and small language models, aimed at reducing long-run dependence on foreign frontier models for applications with strategic or linguistic specificity. Government briefing materials describe a six-layer “sovereign AI stack” running from the energy and chip layers at the base, through data centres and compute, to a data layer of Indic-language corpora, an indigenous-models layer, and an application layer where socio-economic deployment occurs -with a parallel assurance track covering evaluation, testing, conformity assessment and standards. This layered framing is a useful organising device for the state's own investment priorities: it makes explicit that “sovereignty” in AI, for a capital-scarce economy, is being pursued unevenly across layers, with the heaviest near-term public investment concentrated in the compute-access and application layers examined in Section 4, rather than in chip fabrication or frontier-model training, where capital requirements are largest and near-term comparative advantage is weakest.

8. Discussion: Constraints and the Political Economy of Adoption

8.1 The Informal Sector and Measurement Gaps

The World Bank's cross-country automation-risk estimates, like most labour-market statistics used in the AI-and-jobs literature, are built substantially from formal-sector employment data. India's own workforce is majority informal, and the task content of informal work -highly variable, often physically situated, and poorly captured in standard occupational classifications -is not well represented in the datasets that generate figures such as the 4.5 percent automation-risk estimate cited in Section 5.2.

This is not a reason to dismiss the finding, but it is a reason for caution in extrapolating a reassuring aggregate statistic to every segment of the Indian labour force: certain formal or semi-formal task categories that AI can automate comparatively easily -data entry, basic customer support, first-line call-centre work, routine bookkeeping -are exactly the entry-level, formal-sector jobs that have historically served as a bridge out of informality for India's young workforce, and the World Bank report itself flags this category as being at meaningfully greater risk even within a generally favourable developing-economy picture.

8.2 Compute, Energy and Capital Intensity

The Compute pillar's rapid scale-up, documented in Section 4.2 and Figure 1, is a genuine infrastructure achievement, but it also represents a real and growing claim on scarce public capital and, increasingly, on electricity supply. Data-centre and AI-compute capacity is energy-intensive, and the same capital-scarcity that makes AI's labour-augmenting effect proportionally valuable (Section 5.4) also means that diverting public and private capital toward compute infrastructure carries a meaningful opportunity cost against other productive investment.

The sixth layer of the “sovereign AI stack” described in Section 7.3 -energy -is, on the government's own framing, the least developed of the six, and the pace at which India can expand affordable, reliable power for AI infrastructure is likely to be a binding constraint on how far the compute-subsidy model documented in Section 4.2 can scale before subsidy costs become fiscally difficult to sustain at current levels.

8.3 A Voluntary Governance Regime: Enforcement Gaps

Section 3.3 argued that India's voluntary, principle-based governance design is a deliberate trade-off favouring diffusion speed over ex-ante assurance. That trade-off carries identifiable costs. Because the guidelines are advisory rather than legally binding, and because the AI Safety Institute's testing and standards infrastructure was, as of the guidelines' release, still being operationalised, there is currently a gap between the stated ambition of the “Accountability” and “Understandable by Design” sutras and the state's practical capacity to enforce them uniformly, particularly for AI systems deployed by smaller firms or state-government agencies outside the direct purview of sectoral regulators such as the Reserve Bank of India.

The first significant enforcement action taken under the guidelines is likely to be closely watched, both domestically and by other developing economies weighing similar principle-based frameworks, as a signal of whether voluntary compliance can be sustained without periodic recourse to binding sectoral law.

9. Conclusion and Policy Implications

This paper has argued that India's AI strategy is best read as a coherent, if still incomplete, economic wager: that a labour-abundant, capital-scarce economy stands to gain disproportionately from AI deployed as a labour-augmenting general-purpose technology layered onto pre-existing digital public infrastructure, rather than treating the automation anxieties calibrated to high-income labour markets as directly applicable to India's own. The evidence assembled here - the World Bank's cross-country automation and productivity-enhancement estimates, the statistically confirmed AI-preparedness gap documented in Section 5.3, the measured returns from Odisha's agricultural-advisory programme and Telangana's weather-forecasting intervention, and the scale-up of subsidised public compute - lends empirical support to this wager without settling it conclusively.

The Seven Sutras governance framework and the IndiaAI Mission's seven pillars together constitute the institutional scaffolding for this bet, but three areas identified in Section 8 will determine whether the aggregate, population-scale gains documented in this paper translate into genuinely inclusive growth: closing the measurement gap around informal-sector task exposure, resolving the energy and capital constraints on compute expansion, and building enforcement capacity commensurate with the guidelines' stated principles before the voluntary governance model is tested by a significant AI-related harm. For policymakers, the priority implied by this analysis is not to slow AI deployment but to accelerate investment in exactly the complementary factors -human capital, energy infrastructure, and regulatory enforcement capacity -that general-purpose-technology theory identifies as the actual determinants of whether a productivity shock of this kind is broadly shared or narrowly captured, as India works toward the Viksit Bharat 2047 horizon.

References

- Acemoglu, D., & Restrepo, P. (2019). Automation and new tasks: How technology displaces and reinstates labor. *Journal of Economic Perspectives*, 33(2), 3–30.
- Aghion, P., Jones, B. F., & Jones, C. I. (2017). *Artificial intelligence and economic growth* (NBER Working Paper No. 23928). National Bureau of Economic Research.
- 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.
- Bresnahan, T. F., & Trajtenberg, M. (1995). General purpose technologies: 'Engines of growth'? *Journal of Econometrics*, 65(1), 83–108.
- Brynjolfsson, E., & McAfee, A. (2014). *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. W. W. Norton & Company.
- Business Standard. (2026, May 6). India's GCC revenue nears \$100 billion as AI drives sector growth.
- Cazzaniga, M., Jaumotte, F., Li, L., Melina, G., Panton, A. J., Pizzinelli, C., Rockall, E., & Tavares, M. M. (2024). *Gen-AI: Artificial intelligence and the future of work* (IMF Staff Discussion Note SDN/2024/001). International Monetary Fund.
- IMPRI. (2025). *Developing IndiaAI's dream: Understanding the IndiaAI Compute pillar's GPU infrastructure growth*. Impact and Policy Research Institute.
- India AI Rulebook. (2026). *AI governance guidelines 2025*.
- Ministry of Electronics and Information Technology. (2025, November 5). *India AI governance guidelines: Enabling safe and trusted AI innovation*. Government of India.
- NASSCOM. (2026). *Technology sector in India: Strategic review 2026*. National Association of Software and Service Companies.
- Nilekani, N. (2025, April 11). Remarks at Carnegie India Global Technology Summit 'Sambhavna'. As reported in *YourStory*.
- North, D. C. (1990). *Institutions, institutional change and economic performance*. Cambridge University Press.

Press Information Bureau. (2024, March). *Cabinet approves over Rs 10,300 crore for IndiaAI Mission*. Government of India.

PwC India. (2026). *Navigating the skills imperative for India's GCCs in the AI era*.

Saikrishna & Associates. (2025, November 12). Decoding the India AI governance guidelines.

Solow, R. M. (1956). A contribution to the theory of economic growth. *The Quarterly Journal of Economics*, 70(1), 65–94.

World Bank. (2026a). *World Development Report 2026: The promise of artificial intelligence*. World Bank Group.

World Bank. (2026b, August 4). *AI offers lifeline to developing economies in an era of weak growth* [Press release]. World Bank Group.

Zhu, K. (2024, October 31). Mapped: Which countries are most prepared for AI? *Visual Capitalist*.