



AI, Jobs, and the Social Contract

A Policy Paper by the Future of India Foundation

February 2026

FUTURE OF INDIA FOUNDATION

Making the World's Largest Democracy Work for the World's Largest Youth Population

futureofindia.in

C O N T E N T S

Executive Summary	3
I. Work, Democracy, and Power	4
II. Cross-Cutting Disruption Across Tiers and Types of Work	5
A. White-Collar Transformation.....	6
B. The AI-Washing Complication	6
C. Automation at the Physical Frontier	7
D. Gendered Impacts	8
E. Algorithmic Management.....	8
III. The Productivity Paradox and the Demand Risk.....	9
The J-Curve Counter-Argument.....	10
IV. Principles for Managing the Transition.....	11
Democratic Participation.....	11
Transparency and Accountability	11
Security in Transition	12
Shared Prosperity.....	13
Data and Knowledge Justice	13
Domestic Productive Capacity	13
Global Accountability	14
V. What Governments Should Do Now.....	14
Immediate: Within 12 Months	14
Medium-Term: 1–3 Years	15
Structural: 3–5 Years.....	16
VI. Conclusion: Work, Dignity, and Democratic Limits.....	17
Endnotes	19

Executive Summary

“AI will reorganise work at scale. How it reorganises opportunity, and in whose interest, is a matter of governance.”

Artificial intelligence is accelerating structural change across labour markets. But the current trajectory reveals something beyond a technology story: economic and decision-making power over AI deployment is concentrated in fewer than ten firms, while democratic oversight over how benefits and costs are distributed remains limited and uneven.

AI-driven disruption is not technologically predetermined but shaped by policy and institutional choices. Despite unprecedented investment, aggregate productivity gains remain modest, concentrated among a narrow set of frontier firms, and largely decoupled from wage growth or employment expansion. At the same time, disruption is already spreading across white-collar, blue-collar, and informal work simultaneously. Entry-level pathways are narrowing, organisational middles are thinning, and algorithmic management is reshaping work even where jobs persist. Women are disproportionately exposed, concentrated in early-automated roles, and their income loss carries macroeconomic consequences by suppressing household demand in economies that depend on it.

In labour-surplus democracies such as India, unmanaged automation risks compressing wages, weakening domestic demand, squandering the demographic dividend, and eroding the social contract. The question is not whether to adopt AI, but how to govern its adoption and in whose interest.

This brief advances seven principles for democratic management of the transition: democratic participation in deployment decisions; transparency and accountability for labour and data impacts; income security and structured pathways during transition; shared prosperity through fair distribution of productivity gains; data and knowledge justice; investment in domestic productive capacity; and international accountability for transnational AI firms.

These principles are translated into a sequenced action framework. Immediate steps include a national White Paper on AI and Work, early-warning labour observatories, and mandatory AI–Labour Impact Reports for large employers. Medium-term measures focus on parliamentary oversight, coordinated transition missions, and substantive reform of skilling systems. Structural reforms include building indigenous AI capability, embedding employment conditions in public procurement, and advancing global coordination on labour standards and data compensation.

For India and the broader Global South, the stakes are acute. With the world's largest youth populations and service-led growth models, these economies have the most to lose from unmanaged automation and the most to gain from governing it well. The benefits of AI will depend on whether democratic states can shape adoption to expand opportunity across the middle of the labour market, not merely concentrate returns at the top.

AI will reorganise work at scale. How it reorganises opportunity, and in whose interest, is a matter of governance.

I. Work, Democracy, and Power

In 2025, companies worldwide attributed nearly **55,000** job cuts directly to artificial intelligence — twelve times the number just two years earlier.¹² India's top six IT firms shed over **70,000 positions** in a single year.³ Yet a meta-analysis of 371 studies found no robust link between AI adoption and aggregate productivity gains.⁴ AI is displacing workers when it is — and even when it is not — creating value, and the gap between the two is where social contracts break.

What is clear is that AI will alter not only the number of jobs but also the structure, quality, and meaning of work. The effects are already visible in widespread layoffs and hiring freezes, but also in anticipatory decisions as companies reallocate capital toward AI capabilities, reshape hiring, and re-engineer workflows. The recent cycle is instructive: during the pandemic-era online boom, technology firms hired aggressively and wages spiked, only to be followed by cascading layoffs a few years later. The AI wave is repeating elements of that pattern — vast capital inflows and sharp rises in top-end pay alongside abrupt cuts elsewhere. Meanwhile, the ripple effects extend beyond the tech sector. In cities like San Francisco, Gurugram, and Bengaluru, elevated compensation and concentrated capital have priced other workers out of housing and services. Even now, large AI data-centre buildouts compete for energy and long-term capital that could otherwise fund housing, healthcare, or education.

This volatility is not about AI but about power. A small number of firms and executives now make deployment and pricing choices that reshape labour markets across borders, with limited public accountability.

This concentration operates through four channels. First, *standard-setting power*: model capabilities, safety thresholds, and deployment norms are determined by a handful of labs — OpenAI, Anthropic, Google DeepMind, Meta — whose choices become de facto industry standards before regulators engage. Second, *labour-market power*: decisions to deploy agentic systems, eliminate entry-level roles, or restructure service delivery are made at firm level but ripple across sectors and borders within months. Third, *data power*: models are trained on the collective creative and informational output of societies, but the value extracted accrues overwhelmingly to model developers — an asymmetry especially stark for the Global South, whose languages, cultural production, and labour are inputs but whose populations are not shareholders. Fourth, *narrative power*: the framing of AI as inevitable, beneficial, and requiring only “responsible” deployment is itself produced by the firms that stand to gain most from rapid adoption.

Nobel laureates Acemoglu and Johnson have argued that technology has served broad prosperity only when countervailing institutions were strong enough to shape its trajectory. However, unlike earlier transitions shaped by thousands of competing firms and strong labour institutions, this one is directed by fewer than ten companies. Democratic oversight is therefore structurally necessary.

Innovation demands risk, and some disruption is inevitable. But in a democracy, power must be answerable to the people it affects, and innovation must improve lives — democratising rather than concentrating opportunity. That requires democratic institutions capable of anticipating and managing impact so that both risks and gains are shared fairly.

For India — and for much of the Global South — the stakes are even more urgent. These are youth-surplus economies and, for now, recipients of technologies developed elsewhere. Mishandled, AI-driven displacement will squander the demographic dividend; managed well, it can broaden opportunity. At the same time, Governments must ensure that global technology decisions adhere to democratic norms and shared responsibility.

This essay proceeds in four sections: first, the evidence from today's labour market; second, the emerging productivity paradox and demand-side risks; third, a principles-based framework for a just transition; and finally, a sequenced action plan — domestic and international — to keep people at the centre of AI-driven change.

II. Cross-Cutting Disruption Across Tiers and Types of Work

The impact of artificial intelligence on work is no longer speculative. Across sectors and skill levels, the first wave of disruption is already visible in layoffs, shrinking entry pathways, reorganised workflows, and new forms of workplace control. The pattern is not a simple reduction of jobs but a restructuring of how work is produced and supervised.

A. White-Collar Transformation

AI copilots and code-generation tools now automate large parts of routine coding, testing, and documentation in IT and tech services. In fact, Gartner predicts that **40% of enterprise apps will feature AI agents by end of 2026**, up from less than 5% in 2025.⁵ India's top IT firms have shed over 70,000 positions since 2024. Globally over 150,000 tech jobs were lost in 2025.⁶ As global clients build AI capabilities internally and the United States tightens visa and digital-services regimes, the export model that powered India's white-collar expansion faces structural pressure.

Entry-level jobs — tasks once performed by interns, analysts, and assistants — are now handled by generative tools.⁷ Indian IT services companies have reduced entry-level roles by 20–25 percent due to automation and AI;⁸ job platforms report a 35 percent decline in junior tech positions across major EU countries.⁹ KPMG reports that **64% of organizations have already scaled back entry-level hiring due to AI agents' influence, up from 18% last quarter**.¹⁰ These roles do more than provide income: they create socialisation, mentorship, and tacit learning. Their disappearance blocks upward mobility in youth-surplus economies and weakens firms' future talent pipelines.

Forrester finds that Gen Z workers have the highest AI readiness at 22 percent, compared to just 6 percent for Baby Boomers.¹¹ Yet companies are disproportionately shutting out the cohort most capable of working with AI by eliminating the entry-level positions through which they would enter the workforce.

Agentic systems like Claude Code are capable of planning, tool use, and executing multi-step tasks are also compressing managerial roles in law, finance, consulting, and operations. The organisational middle — professionals who oversee execution — is thinning. Senior specialists rise in value, while pathways for progression narrow, reducing the institutional learning that once sustained productivity growth.

Generative models trained on text, images, and music generated by human creatives now compete directly with their sources in the creative economy.¹² Journalism, design, translation, and other creative services have seen sharp declines in human commissions.¹³ As search and chat interfaces

answer queries directly, traffic to original creators falls.¹⁴ This breaks the link between creation and reward and undermines incentives for verified, high-quality information, particularly in countries where creators already face thin margins.

B. The AI-Washing Complication

Whether workers are displaced by actual AI or by the expectation of AI, the effect on entry-level pathways and worker bargaining power is the same. Capital markets reward headcount reductions even when efficiency gains remain unproven — the narrative itself reshapes labour markets. The evidence confirms this pattern.

Forrester found that **55 percent of employers regret laying off workers for AI capabilities that do not yet exist.**¹⁵ Oxford Economics has argued the role of AI in recent cuts is “overstated,” noting that productivity growth has not accelerated in ways consistent with widespread labour replacement.¹⁶ Companies may be using AI as investor-friendly cover for corrections after pandemic-era overhiring.¹⁷

When firms that cut prematurely do rehire, Forrester finds they often do so at significantly lower salaries — producing not job recovery but a downward restructuring of compensation.¹⁸

The Klarna case illustrates the pattern. The company replaced 700 customer-service employees with AI, quality declined, customers revolted, and it had to rehire humans — but on different terms.¹⁹ This cycle of premature displacement, quality failure, and restructured re-employment is likely to recur across sectors.

C. Automation at the Physical Frontier

Retail, logistics, and warehousing are rapidly integrating robotics and computer vision in blue-collar work. Amazon’s systems set the global benchmark,²⁰ and Indian firms such as Reliance Retail,²¹ Flipkart,²² and Delhivery²³ are piloting automated checkout, robotic sorting, and vision-based inspection.

While low wages in India will slow substitution, leasing models and automation-as-a-service accelerate adoption. Construction and agriculture are shifting toward modular, factory-style processes that reduce on-site labour. Drivers and delivery workers are under new pressure from self-driving vehicles, drones, and route-optimisation algorithms that squeeze earning potential even before full automation arrives.

In India and much of the Global South, between **80 and 90 percent of employment lies outside formal wage relationships** in informal work.²⁴ AI’s influence here is less visible but equally consequential. Gig-platform workers — drivers, delivery personnel, domestic service providers — are now managed by algorithms that assign tasks, track performance, and trigger automated penalties or deactivations without human recourse. Small retailers and street vendors compete with AI-driven inventory systems and direct-to-consumer platforms that bypass

traditional distribution. Artisans see their designs replicated by generative models without attribution or payment.²⁵ Because informal workers lack contracts, unions, or social-security coverage, these disruptions directly translate into income volatility. For economies where informal work is the norm rather than the exception, AI governance cannot focus only on formal employment. It must address how algorithmic systems reshape power in unregulated markets, how traditional livelihoods adapt or disappear, and how millions of workers outside the tax and social security net can be included in transition support.

D. Gendered Impacts

Automation is not gender-neutral.²⁶ A large share of early-exposed tasks — reception, transcription, billing, customer support — are concentrated among women.²⁷ Their loss reduces income, mobility, and long-term agency. Algorithmic scheduling further penalises workers with caregiving interruptions, pushing many out of the labour force.

Because women have a high marginal propensity to consume, these income shocks carry macroeconomic consequences: they suppress household demand and local growth precisely when economies need them most. This dynamic is especially dangerous for India, where household consumption is a majority contributor to GDP²⁸ and where female labour-force participation is already low by international standards. When the segments most exposed to automation are also the segments whose spending sustains domestic demand, displacement becomes a macro-critical feedback loop.

E. Algorithmic Management

Even where employment persists, AI is reshaping how performance is monitored and rewarded through algorithmic management. “Bossware” — systems that track keystrokes, sentiment, or time online — promise efficiency but often erode autonomy and increase psychological pressure.²⁹ Workers are held accountable for outputs they do not fully control. Employee concerns about job loss due to AI have risen from 28 percent in 2024 to 40 percent in 2026, and 62 percent of employees feel leaders underestimate AI’s emotional and psychological impact. Yet, under appropriate oversight, the same technologies could reduce bias, improve feedback, and enhance safety — demonstrating that governance choices determine whether AI strengthens or weakens dignity at work.

Taken together, these disruptions reveal a systemic reordering: the bottom of the labour market narrows, the middle hollows, and the top concentrates — across white-collar, blue-collar, and informal tiers simultaneously.

III. The Productivity Paradox and the Demand Risk

The standard counterargument is that productivity gains ultimately raise living standards and create new work. For AI so far, the evidence is mixed: productivity growth across major economies remains modest, and its benefits are concentrated among a small group of frontier firms. The much-publicised rise of “prompt engineering” illustrates how quickly apparent new opportunities can collapse once models learn to automate the very tasks they created.

Across advanced and emerging economies, labour-productivity growth has hovered below 1.5 percent despite unprecedented investment in AI infrastructure.³⁰ Medium-term macro-level productivity gains in the United States³¹ and Europe³² have been estimated to be “modest”. Firm-level data show sharp gains for a narrow set of leaders but little diffusion to the median enterprise. Capital markets have amplified the distortion: technology stocks rise on announced head-count reductions even when efficiency gains remain unproven, rewarding labour cuts over real innovation.

This narrow pattern of productivity deepens inequality. Wage shares in national income continue to fall while capital returns rise,³³ producing economies that grow on paper but generate stagnant or declining real earnings for workers. For labour-surplus democracies, such imbalance is macro-critical: weaker wage growth constrains consumption, narrows the tax base, and dampens private investment — creating a feedback loop of low demand and high precarity. Weak domestic demand also increases dependence on exports at a moment of rising global protectionism, adding vulnerability to external shocks.

The composition of job loss compounds the risk. As detailed in Section II, entry-level roles, youth employment, and women’s work — segments with the highest marginal propensity to consume — are among the most exposed. In globally linked labour markets, automation abroad similarly reduces service-export employment in India and other developing economies.

Evidence of augmentation is real but limited. GitHub Copilot users complete coding tasks faster;³⁴ radiologists using AI achieve higher diagnostic accuracy;³⁵ legal associates and customer-service agents resolve cases more efficiently.³⁶ Yet a study of over 10,000 developers found that AI-augmented teams completed 21 percent more tasks but saw code-review bottlenecks increase by 91 percent — individual acceleration without organisational productivity gain.³⁷ A randomised controlled trial found that experienced open-source developers using AI tools completed tasks *19 percent slower* than without them — while believing they had been 24 percent faster.³⁸ These gains rarely expand employment. Instead, firms expect the same teams to deliver more output, compressing career ladders and training opportunities. The result is “efficiency without inclusion” — productivity that lifts profits but not wages.

The J-Curve Counter-Argument

Proponents argue this pattern reflects a J-curve: productivity dips during adoption before eventually rising, as happened with electricity and computing. A February 2026 study of 809 US banks found precisely this pattern — AI-adopting banks showed a 428-basis-point decline in return on equity during implementation, consistent with an “innovation tax” that may eventually yield gains.³⁹ This argument deserves engagement, but it carries a critical assumption: that the trough will be managed. Historical J-curves unfolded within national economies with functioning labour institutions, welfare states, and decades of adjustment time. The AI transition is faster, more global, and governed by fewer actors. Without deliberate intervention, the trough may not be a trough at all — it may become a permanent reallocation of returns from labour to capital.

A further challenge lies in timing. Historical transitions unfolded over decades; AI adoption is compressing disruption into years. New roles will emerge, but the gap between displacement and re-employment can stretch long enough to erode social trust and demand. Managing that interval is as important as fostering the next wave of innovation.

The core warning remains clear. When technological progress is decoupled from labour inclusion, it creates macro-economic fragility. Taiwan is a cautionary example: while the economy grew by an impressive 7.64% in Q3 2025 driven by semiconductor exports, private consumption rose by less than 1%, indicating that wealth failed to trickle down to labour.

NITI Aayog’s own projections for India are sobering: in a business-as-usual scenario, IT services headcount could fall from 7.5–8 million to 6 million by 2031.⁴⁰ Without broad-based productivity growth that translates into incomes, AI adoption risks producing an economy with strong capital returns but weak demand. This is the dynamic that labour-surplus democracies must avoid.

Yet the same technology that threatens to narrow opportunity can, under different conditions, broaden it. If AI is powerful enough to reorganise workflows in the world’s largest companies, it can also raise the effectiveness of workers, smaller firms, and public systems. True augmentation succeeds where AI handles pattern-recognition and routine generation while humans provide judgment, context, and relational skill. The diffusion of such augmentation, however, depends on access to infrastructure, data, and institutional capacity. Unless deliberate policy spreads these capabilities beyond the frontier, productivity will remain concentrated and social returns limited.

These patterns make one conclusion unavoidable: productivity alone will not guarantee shared prosperity. Deliberate public policy is required to ensure that technological efficiency translates into broad-based income growth.

IV. Principles for Managing the Transition

Technological disruption is not destiny. Its consequences depend on the choices societies make about fairness, participation, and responsibility. The following principles outline a framework for governing AI-driven change so that innovation strengthens, rather than erodes, social stability and democratic legitimacy.

These principles may face resistance — not because they are impractical but because they redistribute power. The firms subject to disclosure and procurement conditions are the same firms that shape technology policy through industry associations and direct engagement with ministries. AI investment attracts capital, headline GDP growth, and high-tax-bracket employment, while displacement costs are diffuse and lagged, creating a structural bias toward AI-friendly policy. This is precisely why the first principle must be democratic governance itself: without institutional capacity to countervail concentrated private power, every other protection becomes precarious.

Democratic Participation

Decisions on how AI is deployed must be subject to public deliberation and legislative oversight rather than confined to executive, technical, or corporate elites. This is the anchoring commitment: if societies cannot collectively shape the terms on which technology enters their labour markets, every other principle in this framework remains aspirational. Regulatory institutions must be structurally independent from the sectors they oversee.

Policy instruments include annual parliamentary hearings on the labour-market impact of AI; mandatory public reporting by large firms on employment and automation outcomes; regional consultations with workers, women's organisations, youth groups, and MSMEs; and legally binding requirements for firms and governments to respond to the recommendations arising from these deliberations.

Transparency and Accountability

Power that cannot be seen cannot be governed. Before societies can manage AI-driven change, they must be able to measure it. This principle brings together three interconnected requirements: making market concentration visible through competition policy, making firm-level decisions traceable through mandatory disclosure, and making workplace AI systems auditable.

- **Competition policy** must evolve to address AI-specific concentration. Current antitrust frameworks assess harm primarily through consumer prices; AI concentration manifests through labour-market power, data control, and standard-setting. Merger reviews and major automation approvals in AI-adjacent sectors should assess employment effects alongside price effects. Structural remedies — including mandatory interoperability of foundation models, data-portability requirements, and limits on

vertical integration between model developers and deployment platforms — should be available to regulators.

- **Mandatory AI–Labour Impact Reports** for enterprises with more than 1,000 employees should disclose projected job losses, redeployment plans, and mitigation measures. These reports should be publicly accessible and collation tabled in Parliament. An annual public dashboard should track wage share, employment levels, and productivity trends at the sectoral level.
- **Dignity and accountability at work** requires that workers know when decisions affecting them are automated, understand the criteria applied, and have the right to human review. Mandatory algorithmic impact assessments should precede the deployment of workplace AI. Systems that penalise factors beyond workers' control should be prohibited.

Security in Transition

People displaced by technological change must have predictable income support, paid learning time, and redeployment guarantees. Entry-level pathways — which provide not only income but socialisation, mentorship, and tacit learning — must be protected or re-created as stepping stones for the young. The speed of AI-driven displacement makes this especially urgent: historical transitions unfolded over generations; this one compresses change into years, and the gap between displacement and re-employment can stretch long enough to erode social trust and demand.

Policy instruments include advance notice requirements before major automation; income-linked transition support replacing 70–80 percent of prior earnings during retraining; paid learning time protected by law and subsidised through a skills and transition fund; regional transition hubs linking employers, training providers, and placement services; expanded early career programs and portable benefits for informal and platform workers.

Gender-responsive design is essential. A large share of early-exposed tasks — reception, transcription, billing, customer support — are concentrated among women. Because women have a high marginal propensity to consume, these income shocks carry macroeconomic consequences, suppressing household demand precisely when economies need it most. Transition support must account for these asymmetries in design, eligibility, and delivery.

Shared Prosperity

AI-driven productivity must translate into higher wages, wider ownership, and fairer profit-sharing — rather than concentrating returns among a few firms or investors. The goal is not only

to redistribute gains but to enable diffusion: ensuring that AI-driven efficiency strengthens the middle of the labour market by raising the effectiveness of workers and smaller firms. When technological progress is decoupled from labour inclusion, it creates macro-economic fragility — economies with strong capital returns but weak demand, narrowing tax bases, and rising precarity.

Policy instruments include profit-sharing, employee-ownership, and sectoral collective-bargaining mechanisms; employment-linked tax and public-procurement incentives; and public-procurement rules that favour solutions which maintain or expand employment. Societies must decide democratically how much automation is acceptable and under what conditions. Some countries may choose to slow adoption in labour-intensive sectors, accepting modestly higher costs in exchange for employment stability. Others may accelerate automation but pair it with strong redistribution and retraining. The point is not to prescribe a single answer but to insist that the choice be made democratically.

Data and Knowledge Justice

AI systems trained on the collective creative and informational output of societies must ensure transparency, consent, and fair compensation for the data and content they use. This asymmetry is especially stark for the Global South, whose languages, cultural production, and labour are inputs but whose populations are not shareholders.

Policy instruments include mandatory disclosure of training-data sources and affirmative consent mechanisms for creators; collective licensing or statutory royalty frameworks to compensate creators; data-contribution agreements between governments and model developers ensuring technology transfer or revenue sharing; and protections for underrepresented languages and cultural traditions.

Domestic Productive Capacity

Countries must build the capacity to develop, adapt, and deploy AI in ways that serve national priorities — particularly expanding employment, raising the productivity of smaller firms, and strengthening public systems. Without indigenous capability, governance will remain reactive and dependent. India's public R&D stands at 0.65 percent of GDP. Its 1,600 Global Capability Centres, expanding startup ecosystem, and the IndiaAI Mission's compute infrastructure provide a foundation — but without deliberate policy to convert these assets into indigenous innovation capacity, they will remain platforms for implementing technologies designed and owned abroad.

This requires sustained public investment in AI research and universities generally; national compute infrastructure accessible to universities, startups, MSMEs, and public agencies; open public datasets; technology-transfer requirements in government procurement; support for domestic firms to move from execution to innovation; and competitive compensation to retain research talent. India's digital public infrastructure demonstrates the state can operate at

population scale; the question is whether that capacity will be directed toward sovereign AI capability.

Global Accountability

The handful of frontier firms whose decisions shape employment across borders must be subject to international accountability mechanisms. No single country — not even a large one — can govern transnational technology firms alone.

Policy instruments include cross-border disclosure of labour and data impacts; a Global Data and Labour Compensation Pool funded through a levy on AI company revenues; G20/ILO/WTO coordination on minimum labour standards for AI-enabled services; collective negotiation by Global South governments for benefit-sharing; and AI Transition Apprenticeship and Skills Mobility Visas. The infrastructure for such coordination is emerging: the UN-backed Global Dialogue on AI Governance, South Korea's AI Basic Act (January 2026), and India's AI Governance Guidelines (November 2025) all signal momentum.

V. What Governments Should Do Now: A Sequenced Democratic Action Framework

The principles outlined above must translate into policy design. What follows is organised into three tiers: immediate (within twelve months), medium-term (one to three years), and structural (three to five years). Earlier actions create the evidentiary and institutional foundations on which later ones depend.

Immediate: Within 12 Months

Publish a White Paper on AI and Work

The purpose is to build a shared, evidence-based understanding of how AI is transforming employment by sector, region, gender, and skill level — so that industrial, labour, and social-protection policies start from a common empirical foundation. The purpose is not perfect prediction but a transparent baseline that enables coherent industrial policy, budgeting, and social protection. A time-bound White Paper should build a shared evidence base by mapping how AI affects employment, modelling scenarios, and assessing readiness. Specifically, it should map how AI affects employment across sectors, regions, skills, and gender; model short-, medium-, and long-term scenarios under different rates of adoption; and assess whether existing education, skilling, and social-protection systems can absorb disruption; and audit India's domestic AI productive capacity — its research talent pipeline, R&D investment relative to competitors, the readiness of GCCs and startups to move from execution to innovation, and the compute and data infrastructure available to Indian institutions. Developed through public consultation and tabled in Parliament, it would provide the shared evidence base needed for coordinated national action. India's Union Budget 2026–27 announced a high-level committee to assess AI's impact on the services sector — the White Paper should be this committee's first deliverable.

Establish Early-Warning and Monitoring Systems

Governments should create labour-market observatories that integrate administrative data, firm disclosures, and periodic surveys to monitor employment trends, wage shares, entry-level hiring, and algorithmic management across sectors. Public dashboards, updated quarterly, can make these changes visible to citizens, researchers, and legislators. Early-warning indicators should trigger automatic policy reviews, ensuring that responses remain adaptive rather than reactive.

Require AI–Labour Impact Reports

Large employers introducing automation that affects more than 1,000 workers should be required to submit impact assessments outlining projected job losses, redeployment plans, and mitigation measures. These reports should be publicly accessible and collation tabled in Parliament.

Medium-Term: 1–3 Years

Anchor AI and Work in Parliament and Public Deliberation

Parliaments should establish standing committees on AI and employment to hold annual hearings requiring ministers to report on employment trends, retraining and redeployment progress, and the use of data in model training. Regional and national consultations — including workers, youth groups, women’s organisations, and MSMEs — should directly inform legislative priorities.

Create Coordination Mechanisms

Inter-ministerial and state transition councils, chaired at a senior level, should align policy and budgets across education, industry, labour, finance, and technology portfolios. These councils should design just-transition missions that provide paid learning time and redeployment support, expanded early-career programs, issue early alerts of automation in high-risk sectors, and pilot regional transition models — testing approaches in select states or cities before national scaling.

Take Skilling Seriously

Skilling must not be treated as a rhetorical response to immediate job loss. It must build a genuine continuum from foundational learning to advanced capabilities. At the base, digital literacy, numeracy, and critical thinking must become universal, starting in primary education. Vocational training must be redesigned in partnership with employers to focus on durable capabilities rather than narrow technical tasks likely to be automated. Entry-level hiring must also be restored and expanded, since it provides the critical bridge between learning and work. Governments should establish a Skills and Transition Fund — financed through a modest levy on payroll or automation gains — to finance paid learning time, lifelong training credits, and redeployment programs.

Structural: 3–5 Years

Build a Global Compact for Fairness and Inclusion

The infrastructure for global AI coordination is emerging: the UN-backed Global Dialogue on AI Governance, South Korea’s AI Basic Act (January 2026), and India’s AI Governance Guidelines (November 2025) all signal momentum.

Governments should work collectively through the G20, ILO, and WTO to establish international labour and transparency standards for AI-enabled services. Such coordination should also address fair compensation when data or creative content from one country is used to train models commercialised globally, and promote technology transfer and skills development in emerging economies.

Hold Frontier Firms Accountable

Frontier firms should be required to publish annual, independently audited AI–Labour Impact Reports, disclosing direct and downstream employment effects, data geography, and compensation to creators and affected workers. Mechanisms must ensure that they return value to the countries and communities whose data, talent, and content underpin their models. Profits accruing from automation should contribute to a collective Global Data and Labour Compensation Fund that supports worker transitions and public innovation.

Build Domestic AI Capability

India must move from being an execution layer for AI technologies designed elsewhere to building its own. This requires a coordinated national effort across three fronts: expanding the research talent pipeline by investing in faculty development, doctoral programmes, and competitive compensation to retain India’s best AI researchers; raising public R&D expenditure; and using government procurement to mandate that AI solutions be co-developed with Indian research institutions or use domestically built models where capability exists.

India’s 1,600 Global Capability Centres, its expanding startup ecosystem, and the IndiaAI Mission’s compute infrastructure provide a foundation. But without deliberate policy to convert these assets into indigenous innovation capacity, they will remain platforms for implementing technologies designed and owned abroad. Procurement conditions should be specific and enforceable: co-development with Indian research institutions; full disclosure of training-data sources and data geography; employment-impact assessments prior to deployment; and contractual commitments to maintain or expand human employment in the relevant public service. Contracts should include clawback provisions triggered by failure to meet employment commitments. These conditions are not protectionism — they are the standard due diligence that any sophisticated buyer applies to transformative technology.

VI. Conclusion: Work, Dignity, and Democratic Limits

Societies have managed large-scale labour transitions before — from agriculture to industry, from manufacturing to services. What distinguishes successful transitions from failed ones is the presence of institutions capable of managing displacement and distributing gains: strong unions that gave workers bargaining power, welfare states that provided income security, public education systems that enabled skill acquisition, and democratic processes that forced capital to accommodate labour's interests. Where these institutions were weak or absent, transitions produced prolonged social instability and the erosion of democratic legitimacy.

What makes the AI transition different is its speed and scope. Earlier shifts unfolded over generations and affected sectors sequentially; AI compresses change into years and reaches across manual and cognitive work simultaneously. And where earlier transitions occurred within regulable nation-states, AI is governed by transnational firms whose decisions outpace policy.

India stands at a critical moment. With the world's largest youth population and a service-led economy, it has the most to lose from unmanaged automation and the most to gain from governing it well. By combining foresight with participation, and national policy with global fairness, it can shape an AI transition that expands opportunity rather than constricts it.

“Democracies cannot remain stable if people become economically disposable. They survive because they insist that every individual counts. If AI is to serve humanity, it must be governed by the same principle that defines democracy itself: power must be answerable to the people it affects.”

ABOUT THE FUTURE OF INDIA FOUNDATION

The Future of India Foundation is an independent, Section 8 non-profit organisation working to strengthen the responsiveness and effectiveness of India's democratic institutions in an era of rapid demographic, economic, and technological change. With nearly two-thirds of India's population under the age of 35, the Foundation focuses on how youth opportunity is created, constrained, and governed across India's districts.

The Foundation's work is anchored in the belief that youth opportunity is not merely a by-product of economic growth, but an institutional outcome shaped by the interaction of public services, labour markets, regulation, and democratic accountability at the local level.

It builds long-term public policy infrastructure through four complementary streams: a national civic-tech platform that maps youth opportunity at the district level; institutional research that tracks legislative and governance performance; documentation of youth lived experience across states; and strategic analysis of structural forces such as artificial intelligence, platformisation, and labour market transformation.

Together, these efforts aim to support evidence-based policymaking, strengthen institutional accountability, and contribute to the development of durable democratic capacity in a young, rapidly changing society.

futureofindia.in

Endnotes

1. <https://www.challengergray.com/wp-content/uploads/2026/02/CR126007123.pdf>
2. <https://www.challengergray.com/blog/december-2023-challenger-report-job-cuts-fall-to-second-lowest-monthly-total-of-the-year-job-cuts-steady-as-hiring-plummets>
3. <https://economictimes.indiatimes.com/tech/information-tech/top-it-companies-cut-over-72000-jobs-in-fy24/articleshow/109653893.cms>
4. <https://cmr.berkeley.edu/2025/10/seven-myths-about-ai-and-productivity-what-the-evidence-really-says/>
5. <https://www.gartner.com/en/newsroom/press-releases/2025-08-26-gartner-predicts-40-percent-of-enterprise-apps-will-feature-task-specific-ai-agents-by-2026-up-from-less-than-5-percent-in-2025>
6. <https://techcrunch.com/2025/12/22/tech-layoffs-2025-list/>
7. AI is 'breaking' entry-level jobs that Gen Z workers need to launch careers, LinkedIn exec warns — Fortune, May 2025
8. <https://restofworld.org/2025/engineering-graduates-ai-job-losses/>
9. <https://talentup.io/blog/why-entry-level-jobs-in-europe-are-becoming-harder-to-find-in-2025/>
10. <https://kpmg.com/us/en/media/news/q4-ai-pulse.html>
11. <https://www.forrester.com/predictions/>
12. Understanding the AI Class Action Lawsuits — Authors Guild
13. A survey by the Society of Authors in April 2024 found that translators and illustrators have already lost work due to generative AI
14. How will generative AI impact website rankings and traffic? — Search Engine Land
15. <https://www.forrester.com/predictions/>
16. <https://www.oxfordeconomics.com/wp-content/uploads/2026/01/Evidence-of-an-AI-driven-shakeup-of-job-markets-is-patchy.pdf>
17. <https://hbr.org/2026/01/companies-are-laying-off-workers-because-of-ais-potential-not-its-performance>
18. https://www.theregister.com/2025/10/29/forrester_ai_rehiring/
19. <https://www.entrepreneur.com/business-news/klarna-ceo-reverses-course-by-hiring-more-humans-not-ai/491396>
20. Amazon Plans to Replace More Than Half a Million Jobs With Robots — New York Times, October 2025
21. How Reliance Retail's Freshpik and Yousta are implementing self-checkouts — India Retailing
22. Flipkart Deploys 100 Automated Robots In Its Bengaluru Delivery Hub — Inc42
23. Robots take over ecommerce-driven logistics from humans — Economic Times
24. https://www.wiego.org/wp-content/uploads/2020/10/WIEGO_Statistical_Brief_N24_India.pdf
25. Is A.I. Art Stealing from Artists? — The New Yorker
26. <https://www.ilo.org/publications/generative-ai-and-jobs-2025-update>
27. Institute for Women's Policy Research
28. <https://www.pib.gov.in/PressReleaseDetailm.aspx?PRID=2219912®=3&lang=2>
29. How AI is used to surveil workers — MIT Technology Review, February 2025
30. OECD Compendium of Productivity Indicators 2025
31. Acemoglu, D. (2024), The Simple Macroeconomics of AI — MIT
32. AI and Productivity in Europe — IMF Working Paper, April 2025
33. Wages still below pre-pandemic level, while corporate profits soared to 15-year high in FY24 — Indian Express
34. Research: quantifying GitHub Copilot's impact on developer productivity and happiness — GitHub Blog
35. AI in diagnostic imaging: Revolutionising accuracy and efficiency — ScienceDirect
36. Thomson Reuters: How AI is Transforming the Legal Profession
37. <https://www.faros.ai/ai-productivity-paradox>
38. <https://metr.org/blog/2025-07-10-early-2025-ai-experienced-os-dev-study/>
39. <https://arxiv.org/pdf/2602.02607>
40. https://niti.gov.in/sites/default/files/2025-10/Roadmap_for_Job_Creation_in_the_AI_Economy.pdf