

AI, Work and the Democratic Bargain

RUCHI GUPTA



India's AI risk is not mass technological unemployment but productivity without labour absorption: firms becoming more efficient while the economy loses the institutions through which young people enter work, acquire judgment and move into economic citizenship.

First published in **Strategic Priorities for India: Essays on AI Policy** (2026).

Executive Summary

Artificial intelligence is neither simply a threat to employment nor a guaranteed engine of prosperity; its consequences depend on how it is developed, deployed and governed. For India, the central risk is productivity without labour absorption: firms becoming more efficient while the economy becomes less capable of bringing young people into stable work, teaching them professional judgment and creating credible routes into the middle class.

The IMF estimates that roughly 60 per cent of employment in advanced economies is exposed to AI; among the major emerging markets it profiles, India's high-exposure share, at about 26 per cent, is the lowest.¹ But exposure is different from vulnerability. Advanced economies have older populations, greater fiscal capacity, deeper welfare systems and control over the firms driving the transition; India has a younger workforce in need of jobs, an overwhelmingly informal labour market, thin transition institutions and a growth path unusually dependent on services exports and entry-level professional employment. The argument is thus that a targeted shock could carry greater impact, because it strikes precisely the pathways on which India depends for labour absorption and social mobility.

The deeper risk is not to particular firms or workers but to the development model itself. India's path into the middle class has run through a narrow channel – services exports and entry-level cognitive work – and that is the channel AI compresses first. For a young, labour-surplus democracy, a shock to the pathway of mobility is not a sectoral problem but a systemic one.

The disruption operates through three channels and timelines. The contraction of entry-level cognitive work is underway: generative systems now produce the drafts, code and analysis previously assigned to juniors, thinning both hiring and the apprenticeship pipeline through which judgment is formed. The export-demand channel is gathering pace – the work need not move to a cheaper country when it can move in-house to the model. And the channel running through physical and informal work operates on two clocks: algorithmic management of platform work is already pervasive; physical replacement remains distant in India, but the harvesting of the data that will train it has begun on Indian factory floors.

The same capabilities offer gains India must capitalise on: wider access to scarce expertise and lower organisational costs for small firms and individual entrepreneurs. Upside and downside run through the same mechanisms; the task is to shape the terms of adoption.

India is alive to the potential of AI to reshape its economy and its governance, and a range of institutions and initiatives has grown up around it. This is welcome, but two things are missing: sustained attention to what AI is doing to work, and democratic representation in how the transition is governed, anchored in a dedicated parliamentary body.

This paper proposes three foundations for that response: seeing the transition clearly; rebuilding the pathway into work and protecting existing workers; and governing the transition democratically – supported by measures to spread adoption and to carry the argument outward through a Global South compact.

1. Work as Democratic Infrastructure

Debates about AI and employment often centre on the number of jobs that may disappear. That is a central concern, but not the only one. Work distributes income, teaches competence, establishes status, creates relationships and gives individuals independence from both family and state; for young people in particular, employment is the transition into full economic and civic citizenship.

Entry-level roles occupy a special place in this process. The first three to five years of working life in an occupation – the junior researcher and analyst, trainee nurse, legal assistant, fresher developer – are the period in which workers hold positions designed to be supervised. They do not begin with mature judgment; they acquire it by producing work, making errors, receiving corrections and observing how institutions operate. Routine work can be tedious, but it is where tacit knowledge is formed.

AI compresses the distance between production and review. Organisations once needed many junior workers to produce drafts, research, routine code and preliminary analysis, and fewer experienced people to assess and correct it. When the production layer becomes cheap and fast, the review layer can operate with fewer people beneath it. The ladder does not merely become shorter; its lower rungs may be removed while the upper rungs remain.

When AI makes it cheaper and faster to produce work without juniors altogether, every firm can rationally thin its junior layer, while the economy loses the institutional capacity through which inexperienced workers become experienced. This logic is well documented: firms under-invest in training because they cannot retain the benefits if the trained worker leaves, and so it can be rational to free-ride on workers trained elsewhere.² That problem gains new force from a technology that removes the very junior tasks through which such training was informally financed. No single firm behaves irrationally, but the overall effect is a market failure in the training of workers. This provides the economic, not merely democratic, justification for the interventions proposed here.

There are democratic stakes too. The democratic bargain is not a formal promise of equality but the expectation that citizenship comes with a credible route to economic agency: citizens tolerate differences in wealth and status when they believe mobility remains possible and effort can improve their circumstances. If large groups come to believe they are economically redundant, the legitimacy of markets and democratic institutions alike weakens.

AI could intensify India's existing weaknesses – inadequate job creation, informality, gender exclusion – by making it easier for firms to grow without employing proportionately more people. The relevant policy question is not only whether AI raises output, but whether it expands the number of people who can acquire skill, develop judgment and participate in economic life.

2. Less Exposed, More Vulnerable

The strongest objection to an India-centred argument is that advanced economies appear more exposed to AI, with their larger shares of professional and cognitive employment. The IMF's analysis bears this out: roughly 60 per cent of employment in advanced economies is exposed, against 40 per cent in emerging markets – and among the emerging markets it profiles, India's high-exposure share, about 26 per cent, is the lowest. But this framing sidesteps the developmental question. Exposure measures the proportion of tasks technology may affect; vulnerability concerns the consequences. A country can be less exposed in aggregate yet more vulnerable if the disruption concentrates in economically and politically central sectors.

India's vulnerability has four sources.

First, it is demographically young. A large young population becomes a dividend only when the economy can bring people into productive work; blocked entry converts the same structure into frustration, dependency and political instability.

Second, India's post-liberalisation mobility story has depended disproportionately on services. This was structural: late developers now tend to deindustrialise prematurely³, which is why India's escalator into the middle class ran through IT services and professional occupations rather than the factory. These sectors do not employ most Indians, but they have played [an outsized role in the creation of the Indian middle class](#), urban mobility, aspiration and export earnings.⁴ An AI shock concentrated here strikes the narrow channel through which structural transformation was occurring in India.

Third, India's transition institutions are weak and its labour absorption problem predates AI. The "jobless growth" literature documents two decades in which output expanded far faster than decent employment;⁵ roughly nine in ten workers are informally employed.⁶ AI threatens to intensify India's existing absorption problem.

Fourth, India has limited control over the frontier: advanced compute, model development, cloud infrastructure and commercial decision-making remain concentrated elsewhere.

The question, then, is not how many tasks AI will automate but what it does to India's development model itself. India's mobility has depended on a narrow set of pathways; a technology that strikes those pathways threatens not a slice of the economy but the mechanism by which the country has turned growth into opportunity. That is the argument the rest of this paper develops.

3. Three Channels of Labour-Market Disruption

The three channels operate on different timelines. The compression of entry-level cognitive work is underway, with the first Indian evidence visible; the export-demand channel is a medium-term risk whose pace depends on how global clients reorganise; and disruption of

physical and informal work runs on two clocks – algorithmic management is already the governing structure of platform employment, while physical replacement remains the most distant, even as the collection of the data that will enable it has begun.

The entry-level ladder

Emerging US evidence suggests early-career employment has weakened in highly AI-exposed occupations, particularly where systems automate rather than complement: researchers have found a marked relative decline among workers aged 22 to 25 in the most exposed occupations,⁷ and the IMF has called reduced entry-level hiring a growing concern.⁸ These findings are contested – critics note the difficulty of separating AI effects from the post-2022 tech-hiring correction – but this paper’s argument rests not on estimates but on the mechanism. Junior work has traditionally consisted of tasks easy to specify and review: initial drafts, document checking, routine code, standard queries, summaries. These are among the tasks on which generative systems have improved most rapidly.

Nor is the evidence confined to the United States. US companies attributed about 55,000 announced job cuts to AI in 2025 – roughly fourteen times the 2023 figure.⁹ India’s top five IT firms shed nearly 70,000 positions in a single year.¹⁰ Tata Consultancy Services alone cut over 12,000 jobs in FY26 even while publicising large-scale AI deployment contracts.¹¹ Fresher hiring, which peaked at around 6 lakh in FY22, fell by 80 per cent to about 1.2 lakh by FY25.¹² This shows that displacement operates through both slower hiring and separations, so the damage concentrates on those trying to enter – which is why, as Section 7 argues, a transition system triggered by formal retrenchment will miss it. Whether these cuts are “caused by AI” in a strict sense is examined in Section 4; what is not in dispute is that the firms that built India’s largest professional apprenticeship system are shrinking its intake.

The optimistic outlook is that AI will enable less-experienced workers to perform at a higher level: it has been argued that well-designed AI could widen access to expertise and thus help rebuild the middle class.¹³ That possibility should be taken seriously, with policy implications adapted to the Indian context.

But individual capability and aggregate hiring are two separate issues. If one junior equipped with AI performs the work of three, the one who retains the position gains while two lose the opportunity to enter the profession. The distributional question is not whether AI makes individual workers more productive, but how many people retain access to the institutional pipeline in which capability is developed and honed. The answer is not to protect yesterday’s routine tasks but to reconstruct the training function they performed: supervised responsibility, real clients and real-world problems, room to make and correct mistakes, a route from assisted production to independent judgment.

The export-demand channel

India's labour-arbitrage model rests on foreign demand for Indian cognitive labour: a bank or retailer that once outsourced coding, testing, support or document processing to India may, as AI improves, perform that work internally with a smaller team. The work need not move to a cheaper country. It can move to the model.

What happens next depends on how global clients reorganise and whether Indian firms move from execution to higher-value capability. NITI Aayog's scenarios illustrate the stakes: under adverse conditions India's technology-services workforce could contract from roughly 75–80 lakh today to about 60 lakh by 2031; under a favourable scenario it could expand to about one crore.¹⁴

India's more than 1,700 Global Capability Centres, through which foreign firms employ Indian workers directly, may mitigate some of this impact.¹⁵ But there are two vulnerabilities here too. First, GCCs retain Indian employment largely as an execution layer while capability, intellectual property and decision rights concentrate in the foreign parent – and an execution layer is precisely what AI compresses first. Second, a political coalition is hardening in the United States around discouraging offshored labour and services, including models like the GCCs¹⁶.

The same logic extends to the export of workers themselves. India has relied on sending its knowledge workers abroad. AI-driven contraction in demand for imported cognitive labour, compounded by hardening visa politics, points to a fall in foreign demand for Indian labour.

The way forward thus requires redirecting expertise toward domestic developmental priorities through appropriate public investments and policy.

Physical and informal work

The third channel concerns the large majority outside formal professional occupations. Agriculture still employs roughly 43 per cent of India's workforce, and the question AI poses is whether the pathways out of agriculture will remain open and for how long.

Platform workers are already subject to algorithmic management: software assigns work, calculates pay, sets metrics, penalises and deactivates. In theory, such systems can reduce favouritism and improve allocation, but they have often made employment more arbitrary through opaque rules and the absence of workers' right to human review. The EU's Platform Work Directive offers a precedent India should incorporate: consequential employment decisions made through software must be reviewable.¹⁷

Manual workers are also entering the AI economy as producers of training data. Multiple news reports from Indian factories have documented workers wearing cameras while sewing, folding, sorting and assembling;¹⁸ the resulting first-person footage – egocentric data – teaches robotics systems how human hands perform physical tasks. In some cases workers were neither separately compensated nor fully informed.

Automation is already replacing physical labour across the developed world and in China, from automated checkouts and factory robots to self-driving cars.¹⁹ Low wages may postpone substitution in India for some years. However, workers can be simultaneously protected from replacement by their low wages and selected as low-cost sources of the data required to build future replacement systems. The worker receives wages for present labour; the firm acquires long-term option value. Consent here cannot be reduced to a signed form – a worker may agree to recording without understanding what the data will train or whether refusal would affect employment – so data governance must account for bargaining power, not merely formal permission. Whether, and on what terms, such replacement-training data may be collected at all is a question we must debate collectively.

Gender across all three channels

AI's effects are likely to hit women workers more. Clerical and administrative occupations have the highest generative-AI exposure of any category, and women's employment is concentrated in exactly these roles.²⁰ Women hold a substantial share of India's business-process and support employment – among the few jobs offering educated women in smaller cities formal wages, transport and social legitimacy for working; a demand collapse removes one of the few sanctioned routes into the labour force.

Consequently, any transition policy must disaggregate by gender at every step.

4. Productivity Without Absorption

The central economic danger is that productivity gains may be captured without increasing employment, wages or domestic demand. For a firm, producing more with fewer workers can be rational. For a labour-surplus economy, the same outcome may be destabilising: if output rises while labour income stagnates, gains accrue to capital owners and specialised workers, and since richer households spend a smaller share of additional income, aggregate demand may weaken even as measured productivity rises. This is productivity without absorption – an economy becoming technically more efficient while becoming less effective at incorporating people into productive work.

The dynamic has a developed literature. Acemoglu and Restrepo distinguish automation that creates new tasks for labour from “so-so automation” – good enough to displace workers, not good enough to generate the productivity that would create new demand for them²¹ – and Acemoglu and Johnson argue that technology raises broad living standards only when institutions force a direction of innovation that complements rather than replaces labour.²² The stakes of that directional choice are highest in a labour-surplus democracy, where the surplus is young and the legitimacy of the political order rests on absorbing it.

Evidence on AI productivity remains mixed – a 2025 METR trial even found experienced developers slower with then-current tools²³ – and benchmark performance cannot be treated as proof of organisational value. In early 2026 the same researchers revisited the question and concluded that AI tools probably accelerate developers more than they did a year earlier, though the size of the gain remains uncertain.²⁴

This evidentiary gap produces an apparent contradiction: firms announce AI-driven workforce reductions faster than rigorous studies can find the productivity gains that would justify them. A meta-analysis pooling 371 estimates found no robust link between AI adoption and aggregate labour-market outcomes,²⁵ and Oxford Economics has called AI's role in recent cuts overstated.²⁶ Part of the resolution is AI-washing: "AI" has become a face-saving label under which companies execute corrections for pandemic-era over-hiring and ordinary margin defence – a narrative that reads to investors as transformation rather than distress, and one capital markets reward, so the narrative itself reshapes labour markets whether or not the capability exists. The implication cuts both ways: AI-attributed displacement statistics overstate the technology's current power, but firms' willingness to invoke AI when shedding workers reveals what they expect and intend, and normalises labour-light growth before the technology has earned it.

The strongest rejoinder is the J-curve: general-purpose technologies from electricity to computing depressed measured productivity during adoption before yielding large gains.²⁷ The argument deserves engagement, but it smuggles in an assumption: that the trough will be managed. Historical J-curves unfolded within nation-states with labour institutions, welfare systems and decades of adjustment time; this transition is faster, transnational and governed by fewer actors. In fact, concerns about an economic transformation "larger than the Industrial Revolution, but unfolding over a vastly shorter time frame [leading to] large-scale job displacement" have prompted more than 200 economists, AI researchers and Nobel laureates to sign a joint statement to "act now"²⁸. The worry is that without deliberate intervention the trough may not be a trough at all but a permanent reallocation of returns from labour to capital – and for a labour-surplus democracy, a decade-long trough may mean a wasted generation, not just lower growth.

Individual acceleration and social benefit are distinct: policy should favour technological uses that raise worker capability or expand access, and be wary of subsidising adoption whose principal effect is to transfer income from workers to capital.

None of this is an argument for preserving inefficient employment. It is an argument for measuring efficiency correctly. A technology that lowers one firm's wage bill while weakening skills formation and transferring transition costs to households and the state may be privately efficient and socially costly.

Where the unabsorbed go

An absorption failure does not appear as mass unemployment – India's workers cannot afford to be unemployed – but as displacement into lower-quality work. The destinations are already visible.

Over the past eight years the shift out of agriculture has essentially stalled: the farm's share of employment sits almost exactly where it was in 2017-18, around 44 per cent, having risen during the pandemic rather than continuing its long decline.²⁹ And the work the non-farm economy added was not, for the most part, regular salaried employment: the salaried share

has, if anything, slipped, while a growing share of workers has been absorbed into self-employment – much of it low-earning own-account work and unpaid labour in family enterprises.³⁰ Growth, in other words, has not been translating into better jobs.

The second destination is the gig economy, the de facto shock absorber for (semi-)educated young men in Indian cities: work available immediately, algorithmically managed, disconnected from any ladder into skill or career progression. Official estimates put the gig workforce at 77 lakh in 2020-21 and project it to reach 2.35 crore by 2029-30 – from about 1.5 per cent of the workforce to roughly 4 per cent.³¹

The third is the government-examination economy, where blocked absorption becomes politically legible. The aspiration for a government job is linked to the absence of adequate well-paying private jobs: lakhs of applicants for a few thousand posts, sustained by a coaching industry that absorbs years of young lives and family savings in a lottery with dwindling odds. The recurring protests over examination leaks and recruitment processes are a reaction to rationed entry-level opportunity.³² This is Section 1's democratic-bargain argument operating in real time; if AI narrows private pathways further, pressure on this valve will grow.

These destinations are the counterfactual policy must weigh. The question is never "AI adoption versus the status quo", but AI adoption governed well versus a labour surplus draining into disguised underemployment, algorithmic precarity and examination queues.

5. The Upside India Cannot Afford to Miss: Labour-Complementing AI

A purely defensive response to AI would be a mistake. India has too many shortages to ignore AI's productive potential. In a country short of teachers, doctors, nurses, lawyers, extension workers and local administrators, AI can augment scarce human capacity: help a nurse triage, a teacher personalise, a farmer access advice, a panchayat official navigate government schemes, a citizen understand a government form in her own language. AI's first democratic promise for India lies in expanding access to expertise that has long been scarce, expensive, urban or English-mediated.

AI also lowers the cost of organisational infrastructure – a small team can now do what once required separate support functions. This matters where most Indians actually work: MSMEs, civil-society organisations, local governments, first-generation entrepreneurs. New forms of work will emerge around deployment, evaluation and data stewardship; but new work does not automatically reach those displaced from old work, and the test is whether it arrives before existing ladders collapse.

Access to easy generation, however, is not access to value. AI lowers the cost of producing output; it does not comparably lower the cost of identifying worthwhile problems, earning trust, reaching users or obtaining finance. A 2026 study of more than 160,000 start-up launches found that generative AI produced a sharp rise in solo entrepreneurship – but solo founders did not become more prevalent among the highest-performing ventures.³³ AI is so

far increasing the number of people who can attempt a business more clearly than the number who build exceptional ones. The missing public infrastructure is not just compute but judgment: policy should democratise the full chain from capability to value – tools, knowledge, validation, mentoring, finance, users and procurement.

6. Who Governs the Transition: Power, Sovereignty and India's Institutional Landscape

Three concentrations of power

The labour-market problem is intensified by the concentration of power over AI, at three levels.

The first is corporate: a handful of companies determine the capabilities, prices, access conditions and safety practices of the most advanced systems, and their product decisions can restructure work across many countries before national institutions have accepted or even debated the consequences.

The second is jurisdictional. In June 2026 the US government directed Anthropic to suspend access to advanced models for foreign nationals; Anthropic temporarily disabled the models for all customers, restoring access to a substantially attenuated model subsequently.³⁴ Similarly, OpenAI, at the US government's request, released its most capable models only to a small group of "trusted partners" vetted by Washington, staggering their wider release under a new executive framework that lets the government review frontier models before they reach the public.³⁵ This demonstrated that access outside the United States can be interrupted by an American political decision unrelated to the priorities of countries relying on the systems. Alongside continuing chip export controls,³⁶ it establishes a broader reality: compute and frontier models are an instrument of national-security and competition policy.

The third is domestic: workers, students and small firms have little influence over decisions to introduce monitoring, automate services or collect data. Those most exposed to the consequences are least likely to participate in deployment decisions.

Sovereignty and India's institutional landscape

The diffusion of open-weight models and falling inference costs reduce some dependence, and India should welcome it – but it comes with its own caveats: the open weights Indian firms increasingly adopt are, to a substantial degree, Chinese – the DeepSeek and Qwen families and their successors.³⁷ Reliance on American frontier models carries export-control and pricing risk; reliance on Chinese open weights carries different questions of security review, data practice and geopolitical exposure.

Sovereignty requires the practical capacity to choose among suppliers, evaluate claims independently, maintain critical public systems during external disruptions and set domestic conditions for use.

India is alive to the potential of AI to reshape its economy and its governance, and a range of institutions and initiatives has grown up around it. The IndiaAI Mission, approved in 2024 with an outlay of over ₹10,000 crore, supports compute, foundation models, datasets, applications and skilling.³⁸ The AI Governance Guidelines of November 2025 set out a principles-based framework for its use.³⁹ An inter-ministerial AI Governance and Economic Group, constituted in 2026, coordinates policy across departments and is tasked, among other things, with examining AI's effect on jobs.⁴⁰ A parliamentary standing committee has reported on the impact of AI. Alongside these sit the Digital Personal Data Protection Act, 2023, whose rules were notified in November 2025,⁴¹ and the labour codes, in force since November 2025, which recognise gig and platform workers for the first time and create a worker re-skilling fund.⁴² Behind them all stands the Digital Public Infrastructure model – the government's central narrative of technology-led inclusion, and a set of rails on which several of the proposals below could run.⁴³ The proposals in this paper are designed to work through these institutions rather than around them; each has a natural home in one or more of them.

At the same time, two central gaps remain: legislative and democratic.

First, the data protection framework has a workplace lacuna. Whether consent within an employment relationship – where refusal carries consequences – can be meaningful, whether the egocentric video of a worker's hands recorded to train the very systems that may one day replace her is her personal data, and whether workers hold any claim on the value of the training data derived from them, all fall between its provisions. Similarly, the labour codes are silent on the mechanism by which platform work is actually controlled – algorithmic allocation, rating-based discipline, automated deactivation. Whether copyrighted work may be used to train a model at all is contested before the courts and the subject of a proposed compulsory-royalty licence – which, notably, accepts that those whose work trains a model have a claim on the value it creates but does not give them the option to opt out and determine that value independently.⁴⁴

The second gap is democratic. Because AI strikes the development model itself, the questions it raises are not technical settings for ministries and firms to adjust but choices about the country's economic direction, and those belong in Parliament. The deliberation so far has been framed through industry and technology rather than youth and labour: the parliamentary report on AI drew heavily on the sector's own optimistic employment projections⁴⁵.

7. Recommendations: Three Foundations

India must judge its AI strategy by how many people the economy absorbs into increasingly productive work – not by the productivity gains of its largest firms.

The objective of democratic policy should be to ensure that productivity gains from AI are widely shared, to preserve credible routes into capability and responsibility for young workers, to protect existing workers against opaque decisions and abrupt losses, and to give Indian institutions the capacity to evaluate and govern the systems reshaping the economy.

The scale of the transformation demands a coherent response rather than fragmented interventions. And given India's status as the world's largest democracy with the world's largest youth population, the response must show that innovation and equitable development can proceed together.

The frontier firms' own answers to this problem do not translate well in India. In the advanced economies the distributional question is starting to be asked as one of sharing the gains directly: proposals to hand the public a slice of the leading AI firms – a few per cent of the equity – or to fund a universal basic income out of the wealth they generate.⁴⁶ Neither is available to India, since the companies whose value those schemes would tap are American and Chinese, and a basic income at any meaningful level is beyond India's fiscal capacity.

Consequently, India's answer has to be structural. Where the frontier can contemplate sharing the proceeds after the fact, India has to shape how growth is produced and spread in the first place – dispersed across regions and firms rather than concentrated, and complementary to labour rather than replacing it – and it has to expand and remake the fragmented schemes it already runs into a social security floor strong enough to hold the economy steady as it changes. For India, distribution cannot be a dividend added at the end; it has to be built into how growth is generated.

This is what Jacob Hacker called predistribution,⁴⁷ and what Acemoglu and Johnson describe as steering the direction of technology⁴⁸ – with the difference that for India it is not a preference but an imperative: the redistributive instruments available to the frontier are structurally unavailable to it.

One objection to this agenda is state capacity. But Indian state capacity is uneven rather than uniformly weak. The question is not whether the Indian state can execute, but whether it will point that capacity at the labour-market transition. This is why India's response to AI is a political and democratic question, not only a developmental one.

What follows rests on three foundations, in order of priority. Without the first – the capacity to see what AI is doing to work – everything else is guesswork. The second rebuilds the pathway into work that AI most directly threatens and protects those already on it. The third gives the transition a democratic home, so that these become choices a polity makes rather than outcomes a market imposes. The remaining measures – spreading adoption and carrying the argument to the Global South – support these three.

Foundation 1: Observe the transition and make adoption answerable

India cannot govern a labour-market transition that it does not track. Existing statistics are poorly suited to detecting the forms of disruption discussed in this paper: fewer entry-level vacancies, slower fresher recruitment, declining hours or wages, movement into lower-quality work, and the growing use of algorithmic systems to allocate, monitor and terminate work.

The first requirement is therefore a national AI and Work Observatory. It should integrate official labour-force and payroll data with sectoral hiring information, employer disclosures, apprenticeship outcomes and carefully validated indicators from private recruitment platforms. Its first task should be a baseline of entry-level hiring, wages, retention and progression across sectors and states; thereafter, regular indicators disaggregated by age, sex, region and employment status, including informal and platform work. Its purpose is not to declare whether AI caused every observed change, but to identify where conditions are deteriorating enough to require investigation or response.

A White Paper on AI and Work should be completed within the first six months and tabled in Parliament. It should map exposure and vulnerability, identify legislative gaps, cost the principal interventions and establish the thresholds that would trigger enhanced monitoring or transition support. It should present alternative scenarios rather than a single forecast. Developed through public consultation, it would provide the shared evidence base and language needed for coordinated national action.

Firms above a certain employee threshold should be subject to mandatory disclosures regarding tasks changed, the expected and actual effects on recruitment and existing workers, and the evidence behind claimed productivity gains. The purpose is not a veto over reorganisation; it is to make large employment claims answerable to workers, investors and the public.

Foundation 2: Rebuild pathways into work and protect existing workers

Rebuild and expand pipelines into work

India should establish sectoral AI-era apprenticeship compacts, beginning where entry-level work is already being reorganised and employers can jointly define credible routes to competence: IT services, financial and business operations, legal services, health administration and selected public services are plausible starting points.

Apprentices should work with AI rather than be trained for a workplace in which it does not exist – but competence should be measured by domain knowledge, verification, judgment and the ability to recognise when automated output is unreliable, through real work and progressive responsibility rather than simulated exercises or generic certification. Vocational training must be redesigned in partnership with employers to focus on durable capabilities rather than narrow technical tasks likely to be automated.

The collective-action problem identified earlier in this paper requires pooled rather than purely firm-level solutions. Sectoral financing arrangements, employer contributions and public support can spread the cost. Major public contracts and incentives should reward firms that maintain credible training and progression pipelines.

Safeguards are needed to prevent apprenticeship from becoming a label for cheap temporary labour. Standards should specify supervision, compensation, learning content and progression, and public funding should follow verified outcomes – placement, wages, retention – not enrolment or certification.

Education reform and skilling should proceed on two timelines. Assessment can change quickly: greater weight on oral defence, source verification, problem definition and the ability to explain and challenge AI-generated work. Foundational literacy, numeracy, reasoning and domain knowledge are the longer-term priority – people cannot verify output in subjects they do not understand – and they are not achievable through short AI courses. Entry pathways should also be designed to include women, smaller cities and people outside elite institutions.

And protect existing workers

The most immediate regulatory gap concerns workplace decisions insulated from scrutiny, whether through automation or unilateral management. Workers should know when systems are being used to make consequential decisions about recruitment, allocation of work, pay, performance, discipline, suspension or termination. They should receive understandable reasons and have access to meaningful human review.

No worker should be dismissed, penalised or deactivated solely through an unreviewable automated process. These protections are particularly urgent for platform workers, for whom algorithmic management is already the governing structure of employment.

Firms should also be required to consult workers or their representatives before high-impact workplace deployments, particularly where systems affect monitoring, pay, allocation, recruitment or termination.

Protection also means income security, and here India's task is less to invent than to deliver. The Code on Social Security already promises gig and platform workers accident, health, maternity and old-age cover, financed by contributions from the aggregators that engage them.⁴⁹ The promise needs delivery on the ground: a contribution set high enough to fund real benefits, the cover the law already provides actually activated, and protection made portable across platforms and states so that a worker does not lose it by changing how she earns. Portable social protection should be developed in parallel, treated as structural reform rather than a near-term deliverable.

Where AI produces sustained and measurable gains, workers should have credible routes to share in them. The relevant test is not whether employment rises in exact proportion to output, but whether productivity gains accrue exclusively to owners and a narrow group of specialists while transition costs are transferred to households and the state.

Worker-generated training data require a more specific framework than ordinary consumer consent. Employers should disclose what is recorded, what systems it will train, how long it is retained and with whom it is shared; refusal should not attract retaliation where collection is unnecessary to the present role; and where data production is a distinct part of the job, additional compensation or collective negotiation is appropriate. Whether training data whose sole purpose is to replace the worker altogether should be permitted, at least for the time being, is a separate question.

Foundation 3: Govern the transition democratically

The third foundation is democratic governance for the transition. This requires two things: a federal architecture and a permanent parliamentary home.

The architecture should be federal but coordinated. The Centre should establish a national floor of workplace rights, reporting standards, portability and apprenticeship quality; states should experiment above it according to their economic structure.

A permanent Parliamentary Committee on AI, Work and Democratic Transition should review the observatory's findings, examine major public deployments and hold regular public hearings in different regions – with workers, platform workers, women's groups, employers, states, young people and MSMEs instead of a single national consultation. Its purpose is not to remove technological choices from political contest but to give the contest a constitutional forum, a shared factual base and a durable minimum consensus that survives changes in government.

Supporting measures: spread productive adoption and create demand for labour

The foundations govern the transition; they do not by themselves create the work. That is the growth-side task.

India should support uses of AI that expand access and raise the capability of workers and smaller organisations. Public-interest missions in health, education, agriculture and local administration should begin as evaluated pilots. Success should be measured through service reach, quality, error rates, worker capability and distributional effects, not staff reduction alone.

MSMEs need practical infrastructure rather than appeals to adopt AI: shared tools, Indian-language interfaces, custom workflows, sector-specific resources, independent evaluation, mentors, buyers and public procurement. India also needs public and university research, language resources, shared compute, technical expertise within government and institutions capable of auditing systems used in workplaces and public services. Critical systems should not depend entirely on one supplier or jurisdiction.

Public procurement should reward systems that expand access, permit meaningful human review and meet interoperability and data-protection standards – not require contractors to preserve fixed headcounts, which would reward artificial labour intensity; where a deployment substantially reorganises work, the contract should instead require consultation, training and transparent reporting.

But AI governance cannot substitute for an employment strategy. Apprenticeships improve entry into work that exists; they do not create enough demand for labour. A serious absorption strategy must expand the sectors in which India has unmet needs and where technology can complement labour: care, health, education, housing, construction, urban services, tourism, logistics, climate adaptation, labour-intensive manufacturing and the formalisation of small enterprises. The direction of public investment matters as much as the design of AI regulation.

Supporting measure: a Global Compact for AI, Work and Shared Prosperity

AI systems are developed in one jurisdiction, trained on data and labour drawn from many others, and capable of altering employment across borders with no corresponding system of accountability.

India should begin with a focused Global South coalition on labour-impact reporting, algorithmic management, independent evaluation and consultation where a global firm's technological or sourcing decisions affect employment elsewhere. Countries that contribute labour, languages, data and markets should not remain rule-takers in a system built partly from their resources.

India's argument should not be defensive: the Global South is a potential producer, adaptor and governor of AI, not merely a site of displacement or a source of training data.

8. Conclusion

India's central AI risk is that firms become more productive while the economy becomes less capable of absorbing people, developing their judgment and connecting greater capability to greater security and voice. That outcome would look efficient at the level of the firm and corrosive at the level of society.

India stands at a crossroads. AI can widen access to expertise, lower the cost of organisation, strengthen small firms and improve public services. But upside and downside run through the same mechanisms. A tool that helps one junior perform at a higher level may also reduce the number of juniors hired; an automated system that may improve efficiency may also become arbitrary and make management unanswerable; data that raises the capability of a robot may also be extracted from workers who receive none of its future value.

The democratic bargain requires a credible route into productive work, a reasonable floor of security and a voice over consequential decisions. Rebuilding that bargain means steering technology for mass benefit, treating entry-level work as an apprenticeship institution, sharing productivity gains, protecting workers against opaque systems and directing public investment towards labour-absorbing demand.

India can obtain cheaper production while allowing its entry ladders to decay. Or it can use AI to widen access to expertise, distribute growth across regions and classes and proactively build the pipelines and institutions that turn inexperienced citizens into capable workers. The question is not whether India will adopt AI. It is whether productivity will be allowed to sever its link with capability, labour income and economic citizenship.

Notes

1. International Monetary Fund, Staff Discussion Note SDN/2024/001, M. Cazzaniga et al., "Gen-AI: Artificial Intelligence and the Future of Work," 14 January 2024. <https://www.imf.org/en/Publications/Staff-Discussion-Notes/Issues/2024/01/14/Gen-AI-Artificial-Intelligence-and-the-Future-of-Work-542379>
2. The under-provision of general training is a standard result in labour economics: firms cannot fully appropriate the returns to training a worker who may leave, and so under-invest. Gary S. Becker, *Human Capital: A Theoretical and Empirical Analysis*, University of Chicago Press, 1964.
3. Dani Rodrik, "Premature Deindustrialization," *Journal of Economic Growth* 21(1), 2016. https://drodrik.scholar.harvard.edu/files/dani-rodrik/files/premature_deindustrialization_revised2.pdf
4. Future of India Foundation, "IT Sector Helped Build the Indian Middle Class. For Its Workers, It Is Starting to Sputter," YouthPOWER, 2 October 2026. <https://youthpower.in/insights/it-and-the-middle-class>
5. Santosh Mehrotra & Jajati K. Parida, "India's Employment Crisis: Rising Education Levels and Falling Non-agricultural Job Growth," Centre for Sustainable Employment Working Paper #23, Azim Premji University, 2019, https://cse.azimpremjiuniversity.edu.in/wp-content/uploads/2019/10/Mehrotra_Parida_India_Employment_Crisis.pdf; and the State of Working India reports, Centre for Sustainable Employment, Azim Premji University, <https://cse.azimpremjiuniversity.edu.in/state-of-working-india/>
6. International Labour Organization & Institute for Human Development, *India Employment Report 2024*, 2024. https://www.ilo.org/sites/default/files/2024-08/India%20Employment%20-%20web_8%20April.pdf
7. Erik Brynjolfsson, Bharat Chandar & Ruyu Chen, "Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence," Stanford Digital Economy Lab, 2025. <https://digitaleconomy.stanford.edu/publication/canaries-in-the-coal-mine-six-facts-about-the-recent-employment-effects-of-artificial-intelligence/>
8. International Monetary Fund (Kristalina Georgieva), "New Skills and AI Are Reshaping the Future of Work," IMF Blog, 14 January 2026 (drawing on IMF Staff Discussion Note SDN/2026/001). <https://www.imf.org/en/blogs/articles/2026/01/14/new-skills-and-ai-are-reshaping-the-future-of-work>
9. Challenger, Gray & Christmas, "2025 Year-End Report": AI was cited in 54,836 announced US job cuts in 2025, against 3,900 in 2023. These are announced, US-based, employer-attributed cuts. <https://www.challengergray.com/blog/2025-year-end-challenger-report-highest-q4-layoffs-since-2008-lowest-ytd-hiring-since-2010/>
10. "Top 5 Indian IT companies lost 69,167 employees in FY24," *Business Today*, 27 April 2024. <https://www.businesstoday.in/technology/news/story/top-5-indian-it-companies-lost-69167-employees-in-fy24-report-427193-2024-04-27>
11. TCS announced reductions of about 12,000 roles (~2 per cent of its workforce) over FY26, July 2025. <https://www.peoplesmatters.in/news/performance-management/tcs-announces-12000-layoffs-by-fy26-in-strategic-workforce-realignment-42242>
12. Fresher-intake estimate (peak of roughly 6 lakh in FY22 to about 1.2 lakh by FY25) from talent-analytics firm Xpheno, reported in "Indian IT's fresher hiring slump signals structural shift, not just slowdown," *Business Today*, 21 March 2026. <https://www.businesstoday.in/technology/story/indian-its-fresher-hiring-slump-signals-structural-shift-not-just-slowdown-521374-2026-03-21>
13. David Autor, "Applying AI to Rebuild Middle Class Jobs," NBER Working Paper 32140, 2024. <https://www.nber.org/papers/w32140>

14. NITI Aayog, with NASSCOM and the Boston Consulting Group, "Roadmap for Job Creation in the AI Economy," October 2025. https://niti.gov.in/sites/default/files/2025-10/Roadmap_for_Job_Creation_in_the_AI_Economy.pdf
15. NASSCOM data on India's Global Capability Centres (more than 1,700 as of 2025). <https://nasscom.in/gcc2025/>
16. "U.S.'s HIRE Bill may choke India's IT, BPM, GCC sectors," The Hindu, 27 September 2025. <https://www.thehindu.com/business/uss-hire-bill-may-choke-indias-it-bpm-gcc-sectors/article70101354.ece>
17. Directive (EU) 2024/2831 of the European Parliament and of the Council of 23 October 2024 on improving working conditions in platform work. <https://eur-lex.europa.eu/eli/dir/2024/2831/oj/eng>
18. "'Who is going to pay us when we're replaced by robots?' The Indian factory workers told to film themselves for AI," The Guardian, Global development, 24 June 2026. <https://www.theguardian.com/global-development/2026/jun/24/indian-factory-workers-told-film-themselves-for-ai-robots>
19. International Federation of Robotics, World Robotics 2025: China accounted for about 54 per cent of global industrial-robot installations in 2024, with an operational stock exceeding two million units. <https://ifr.org/ifr-press-releases/news/global-robot-demand-in-factories-doubles-over-10-years>
20. Pawel Gmyrek, Janine Berg & David Bescond, "Generative AI and Jobs: A Global Analysis of Potential Effects on Job Quantity and Quality," ILO Working Paper 96, August 2023 (clerical work the most exposed occupational category). https://www.ilo.org/sites/default/files/2024-07/WP96_web.pdf
21. Daron Acemoglu & Pascual Restrepo, "Automation and New Tasks: How Technology Displaces and Reinstates Labor," Journal of Economic Perspectives 33(2), 2019. <https://pubs.aeaweb.org/doi/10.1257/jep.33.2.3>
22. Daron Acemoglu & Simon Johnson, Power and Progress: Our Thousand-Year Struggle Over Technology and Prosperity, PublicAffairs, 2023. <https://www.hachettebookgroup.com/titles/daron-acemoglu/power-and-progress/9781541702554/>
23. METR, "Measuring the Impact of Early-2025 AI on Experienced Open-Source Developer Productivity," July 2025 (experienced developers were, on average, slower with the tools available in early 2025). <https://metr.org/blog/2025-07-10-early-2025-ai-experienced-os-dev-study/>
24. METR, "We are Changing our Developer Productivity Experiment Design," 24 February 2026 (developers are likely more sped up than a year earlier, though the evidence on magnitude is weak). <https://metr.org/blog/2026-02-24-uplift-update/>
25. Enrico Santarelli, Emanuela Carbonara & Ishita Tripathi, "Assessing the Impact of AI on Labor Market Outcomes: A Meta-Analysis," SSRN working paper, 2025. The pooled unit is 371 estimates and the outcome studied is aggregate labour-market outcomes; the paper is an unrefereed preprint. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=5126345
26. Oxford Economics research briefing, January 2026, reported in "AI is a 'convenient corporate fiction' for layoffs," Fortune, 7 January 2026. <https://fortune.com/2026/01/07/ai-layoffs-convenient-corporate-fiction-true-false-oxford-economics-productivity/>
27. Erik Brynjolfsson, Daniel Rock & Chad Syverson, "The Productivity J-Curve: How Intangibles Complement General Purpose Technologies," American Economic Journal: Macroeconomics 13(1), 2021 (NBER Working Paper 25148), <https://www.nber.org/papers/w25148>; on the electrification lag, Paul A. David, "The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox," American Economic Review 80(2), 1990.
28. "We Must Act Now: A Statement on AI's Transformation of the Economy," Stanford Digital Economy Lab, 13 July 2026. <https://www.wemustactnow.ai/>
29. Ministry of Statistics and Programme Implementation, Periodic Labour Force Survey. Agriculture's share of employment (usual status) was 44.1 per cent in 2017-18, rose to about 46.5 per cent by 2020-21, and was 43.0 per cent in the 2025 (calendar-year) report. <https://www.mospi.gov.in/>
30. Ministry of Statistics and Programme Implementation, Periodic Labour Force Survey, distribution of workers by employment status: regular wage/salaried employment fell from about 22.8 per cent (2017-18) to 21.7 per cent (2023-24); self-employment rose from 52.2 to 58.4 per cent; unpaid helpers in household enterprises rose from 13.6 to 19.4 per cent. <https://www.mospi.gov.in/>
31. NITI Aayog, "India's Booming Gig and Platform Economy: Perspectives and Recommendations on the Future of Work," 2022 (77 lakh gig workers in 2020-21, projected to 2.35 crore by 2029-30); https://www.niti.gov.in/sites/default/files/2023-06/Policy_Brief_India%27s_Booming_Gig_and_Platform_Economy_27062022.pdf
32. See, for example, the 2022 Railway Recruitment Board (RRB-NTPC) protests: "RRB NTPC Exam Row: Railways Reaches Out to 2 Lakh Aspirants Over Protests," India.com, 2022, <https://www.india.com/news/india/rrb-ntpc-exam-row-railways-reaches-out-to-2-lakh-aspirants-over-protests-5222033/>; alongside the 2022 Agnipath protests and recurring examination paper-leak agitations.
33. "Generative AI Fuels Solo Entrepreneurship, but Teams Still Lead at the Top," arXiv:2605.10291, 2026 – an analysis of more than 160,000 Product Hunt launches. <https://arxiv.org/abs/2605.10291>

34. Anthropic, "Statement on the US government directive to suspend access to Fable 5 and Mythos 5," 12 June 2026, <https://www.anthropic.com/news/fable-mythos-access>; and "Redeploying Fable 5," 30 June 2026, <https://www.anthropic.com/news/redeploying-fable-5>
35. "OpenAI limits GPT-5.6 rollout after government request, says restrictions shouldn't be the norm," TechCrunch, 26 June 2026. <https://techcrunch.com/2026/06/26/openai-limits-gpt-5-6-rollout-after-government-request-says-restrictions-shouldnt-be-the-norm/>
36. U.S. Bureau of Industry and Security, advanced-computing and semiconductor manufacturing export controls: interim final rule of 7 October 2022 and the update of 17 October 2023. <https://www.bis.gov/press-release/commerce-strengthens-export-controls-restrict-chinas-capability-produce-advanced-semiconductors-military>
37. DeepSeek (DeepSeek-AI) and Qwen (Alibaba) are Chinese open-weight model families that have been widely adopted internationally. The characterisation of their adoption specifically by Indian firms is an inference from that broader trend rather than a separately sourced figure.
38. Union Cabinet approval of the IndiaAI Mission, with an outlay of ₹10,371.92 crore, 7 March 2024, Press Information Bureau. <https://www.pib.gov.in/PressReleasePage.aspx?PRID=2012355>
39. Ministry of Electronics and Information Technology, "India AI Governance Guidelines," 5 November 2025. <https://static.pib.gov.in/WriteReadData/specificdocs/documents/2025/nov/doc2025115685601.pdf>
40. "Government Constitutes AI Governance and Economic Group (AIGEG) to Lead India's National AI Governance Strategy," Press Information Bureau, April 2026. <https://www.pib.gov.in/PressReleasePage.aspx?PRID=2252739>
41. Digital Personal Data Protection Act, 2023; Digital Personal Data Protection Rules notified 13 November 2025, Ministry of Electronics and Information Technology. <https://static.pib.gov.in/WriteReadData/specificdocs/documents/2025/nov/doc20251117695301.pdf>
42. Ministry of Labour and Employment, the four Labour Codes brought into force on 21 November 2025 (recognising gig and platform workers; the Industrial Relations Code establishes a worker re-skilling fund), Press Information Bureau. <https://www.pib.gov.in/PressReleasePage.aspx?PRID=2192463>
43. On India's Digital Public Infrastructure – the "India Stack" of Aadhaar, the Unified Payments Interface and related layers. https://en.wikipedia.org/wiki/Digital_public_infrastructure
44. The lawfulness of using copyrighted work to train models is contested in ANI Media Pvt Ltd v OpenAI before the Delhi High Court; the compulsory-licence proposal appears in the Department for Promotion of Industry and Internal Trade expert committee's "Working Paper on Generative AI and Copyright (Part I)," December 2025. <https://www.dpiit.gov.in/static/uploads/2025/12/ff266bbeed10c48e3479c941484f3525.pdf>
45. Standing Committee on Communications and Information Technology (Lok Sabha), Twenty-Seventh Report, "Impact of Emergence of Artificial Intelligence and Related Issues," presented 30 March 2026.
46. On proposals for public equity stakes in leading AI firms, see "OpenAI proposes US government own 5% stake to address political blowback," CNBC, 2 July 2026. <https://www.cnbc.com/2026/07/02/openai-proposes-us-government-own-5percent-stake-to-address-political-blowback.html>
47. The term "predistribution" was coined by Jacob S. Hacker, "The Institutional Foundations of Middle-Class Democracy," Policy Network, 6 May 2011. <https://en.wikipedia.org/wiki/Predistribution>
48. Daron Acemoglu & Simon Johnson, Power and Progress: Our Thousand-Year Struggle Over Technology and Prosperity, PublicAffairs, 2023.
49. Code on Social Security, 2020, which defines gig and platform workers and provides for a Social Security Fund financed by aggregator contributions, Press Information Bureau. <https://www.pib.gov.in/PressReleasePage.aspx?PRID=2200767>