



SAHILVERSE · FIELD GUIDE · OCTOBER 2026

The Task Shift Atlas

A field guide to AI and work: tasks before titles

AUTOMATE

codified, low stakes

AUGMENT

tool assists, you steer

PROTECT

presence, trust, care

RETHINK

codified but high stakes

This guide accompanies the interactive atlas at **aijobs.sahilverse.online**. It explains a simple way to think about how AI changes work, summarises what 2025–26 evidence does and does not show, walks through worked cases (including a South Asian small business), and gives exercises you can do alone, with a team, or with a family business.

What this is not: a forecast of job losses, career advice, or legal or financial guidance. The ratings in the atlas are editorial judgments, not measurements.

By Sahil Tayal · First edition 26 October 2026 · sahil@sahiltayal.com · sahiltayal.com



Contents

- 1. The core idea: occupations are bundles of tasks
- 2. Four bins: automate, augment, protect, rethink
- 3. The evidence, 2025–26, labelled
- 4. Capability is not adoption
- 5. Worked case: a Delhi kirana
- 6. Worked cases: an analyst and a care worker
- 7. The entry-level ladder
- 8. Field exercises
- 9. The 12 occupations at a glance
- 10. Read further

1. The core idea

“Will AI take my job?” is the wrong unit of analysis. A job title names a bundle of tasks held together by an employer, a contract, a customer relationship or a family arrangement. Technologies act on *tasks*. When some tasks in a bundle get cheaper, the bundle is redrawn: work is added, removed, merged, handed to someone else or turned into a new job.

That is why two people with the same title can face very different futures. A teacher who spends half the week on registers and worksheets faces a different shift from one who spends it running discussions and noticing which child is struggling at home.

Three distinctions run through this guide:

- **Occupation vs task.** Most research measures tasks and aggregates up; most headlines quote occupations.
- **Capability vs adoption.** That a model *can* do something says little about whether a given organisation will pay for it, trust it, or be allowed to use it.
- **Evidence vs scenario.** Observed employment data, employer expectations and editorial thought-experiments are different kinds of claim. This guide labels them.



The atlas method. Steps 2–3 are mechanical; steps 4–6 are where the real judgment lives.



2. Four bins

The atlas sorts each task with three ratings from 0 to 3 and four rules applied in order.

Rating	Question it asks	0 means	3 means
C — Codifiability	Can it be written down as rules, templates or learnable examples?	Tacit, improvised	Rule-following, many examples
R — Relational presence	Does it depend on trust, care, a body in the room or a live relationship?	Not at all	The task is the relationship
S — Stakes	How costly is an unnoticed error, and who bears it?	Cheap, visible	Costly, hidden, borne by the vulnerable

Order	Rule	Bin	Why this order
1	$R \geq 2$	Protect	Presence is treated as decisive — a value judgment stated openly.
2	$C = 3$ and $S \leq 1$	Automate	Fully codified and cheap to get wrong: the most automatable.
3	$C \geq 2$ and $S \geq 2$	Rethink	Codified but risky: the workflow, review and accountability must be redesigned.
4	otherwise	Augment	Partly codified or hard to codify, low presence: tools assist a person who steers.

Uncertainty. Ratings of 1 or 2 sit near thresholds. A task with three such ratings is marked *contested*; treat it as the weakest part of any reading.

A known weakness. A hard-to-codify, high-stakes judgment with no relational element (for instance, deciding whether a data result is causal) lands in “augment”. Many readers would prefer “protect”. The atlas records your disagreement rather than overriding it — disagreement is information.

Protect is not “safe”. Care work is mostly in the protect bin and is also among the lowest-paid work in many countries. Protection describes what a task depends on, not how it is valued.



3. The evidence, labelled

Each row states its source, date and the kind of claim. Figures are reported as the source frames them.

Source & date	Kind of claim	What it says	What it cannot say
ILO & NASK, Working Paper 140, May 2025	Potential exposure (task scoring + survey + experts)	About one in four workers globally are in occupations with some generative-AI exposure; 3.3% in the highest gradient (4.7% of women's employment, 2.4% of men's). Clerical work most exposed.	Job losses. The authors stress exposure is not displacement.
WEF Future of Jobs Report, Jan 2025	Employer expectations (1,000+ firms, 55 economies)	170 million jobs created and 92 million displaced by 2030: net 78 million.	What will actually happen; which workers move into the new jobs.
McKinsey State of AI, Nov 2025; chart 3 Sep 2026	Organisational self-report	88% use AI in at least one function; nearly two-thirds not yet scaling. 32% expected headcount cuts in 2025; 14% reported them a year later; 39% expect cuts next year.	Worker-level outcomes; causality.
Brynjolfsson, Chandar & Chen, Stanford, rev. Aug 2026	Descriptive payroll data (US)	No economy-wide displacement; employment of 22–25-year-olds in highly exposed occupations ~19% below less-exposed peers, mainly via reduced hiring; declines where AI use substitutes.	Causal effect (authors' own caveat); transfer to India.
NITI Aayog roadmap, Oct 2025	Policy scenario (with NASSCOM, BCG)	For India's tech-services sector: losses up to 1.5 million or creation up to 4 million, depending on choices.	An independent forecast; informal-sector outcomes.
Anthropic Economic Index, 2025–26	Usage logs of one AI product	How people actually use one assistant across tasks, including automation-like vs augmentation-like use.	Use of other tools; non-users. Published by an AI developer.

Interpretation (the author's reading, not a finding): across these sources, the clearest current signal is slower *entry* into exposed work, not mass layoffs. Expectations of cuts have tended to run ahead of reported cuts.

Open questions: do augmentation gains reach workers as pay or time, or as higher targets? What happens in informal work, where most South Asian workers are, and where data are thinnest?



4. Capability is not adoption

Between “a model can do this” and “this shop now does it differently” sit many filters. The atlas asks six questions:

Condition	Why it matters in South Asian small business
Reliable device and connection	Phones are often shared; data packs are budgeted.
Tools in the working language	Hindi, Marathi, Tamil, Bengali and code-mixed Hinglish are the real working languages.
Affordable tool and upkeep	A monthly subscription is weighed against stock and rent on thin margins.
Digital, clean data	Khata books, paper bills and memory still hold much of the record.
Legal and contractual room	Confidentiality, client terms and data rules vary.
Worker voice	Family helpers and assistants rarely choose the tools or metrics.

5. Worked case: a Delhi kirana (fictional composite)

Meena runs a neighbourhood shop. A distributor offers a free AI catalogue bot for WhatsApp: it answers price questions and takes orders. Her nephew handles the shop phone after school.

Task	C R S	Model	Meena's call and reason
Answer price questions on WhatsApp	3 1 0	Automate	Automate after hours only — regulars message to be recognised.
Reorder from distributors	2 1 2	Rethink	Suggestions help, but the bot does not know who gives credit.
Decide udhaar (credit)	1 3 3	Protect	Trust, family circumstance, neighbourhood ties.
Record cash, UPI and khata	3 0 1	Automate	Augment instead — she checks against the drawer daily.
Handle a short-weight complaint	1 3 2	Protect	Goodwill is the shop's edge.

Mistake avoided: letting the bot reply to every regular in the same template. **Binding constraint:** one shared phone and an English-first interface — access, not capability. **Outcome:** partial adoption, deliberately bounded.



6. Two more worked cases (fictional composites)

Arjun, second-year analyst, Bengaluru

His team adopts a coding assistant; the SQL and decks that filled his week now take an afternoon. Sorting puts SQL and decks in **automate**, data cleaning in **rethink** (silent errors now flow straight into decisions), and framing questions and explaining findings in **protect**.

The instructive mistake: treating saved time as safe time. The routine work was also his training. He asks for a deliberate learning ladder — owning one causal question each quarter, reviewed by a senior. This is the organisational side of the Stanford finding on entry-level hiring.

Sunita, home care worker, Pune

An agency introduces an app with voice care-logs, medication reminders and GPS check-ins. Logs → **augment**, if the app understands Marathi. Reminders → **rethink**: who is accountable when a dose is missed? Bathing, observation and coordinating with family → **protect**.

The edge case: GPS check-ins add monitoring of the worker rather than support for the client. Technology can intensify work as readily as lighten it. **Her question for the agency:** will time saved on logs become time with clients, or more visits per day?

What the three cases share

- The decisive factor was rarely what the model could do; it was access, accountability and who controls the change.
- Each person kept the relationship-heavy tasks and redesigned the high-stakes codified ones.
- Each case ended in a conversation, not a prediction.



7. The entry-level ladder



If the bottom rung is automated, organisations must build a new one — or seniors stop being made.

Junior roles have long been where people learn by doing the codified, lower-stakes parts of the work: first drafts, routine tickets, reconciliations, research memos. These are exactly the tasks the atlas most often puts in *automate*.

If they disappear without replacement, two things happen. Fewer people are hired at the bottom, and those who are hired lose the practice that made earlier generations competent. The 2026 Stanford revision reports the adjustment for young US workers is happening mainly through hiring rather than separations — consistent with this mechanism, though the authors do not claim causality.

Design responses an organisation can try

- **Reviewed real work.** Juniors handle real but reversible cases with senior review, instead of synthetic busywork.
- **Shadowing the judgment.** Make seniors narrate why they accept or reject an AI output.
- **Rotations into protected tasks.** Client conversations, investigations and care are learned alongside people.
- **Visible learning goals.** Track what a junior can now do unaided, not only output volume.

These are proposals, not tested interventions. The Human Capital Commons (humancapitalcommons.sahiltayal.com) has teaching units on deliberate practice, feedback and mentorship that go deeper.



8. Field exercises

Exercise A — Your own task inventory (30 minutes)

- Write the 10 tasks that filled your last two working weeks. Use verbs: “reconcile...”, “calm...”, “decide...”.
- Rate each C, R, S from 0 to 3. Do not peek at the rules yet.
- Apply the four rules. Circle every task where the result feels wrong and write one sentence on why.
- Estimate roughly how much time each task takes. The atlas counts tasks; your life is measured in hours.

Exercise B — The adoption audit (team, 45 minutes)

- Pick three “automate” tasks. For each, walk through the six access conditions.
- Name who would decide to adopt a tool, who would pay, and who would carry errors.
- Agree one condition that would make you pause or reverse the change.

Exercise C — Family business conversation (kirana, workshop, stall)

- Ask each family member and helper to list what they do in a normal day. You will find invisible work.
- Sort together. Disagreement about the protect bin is usually disagreement about what the business is for.
- Decide one tool experiment with a fixed end date, and what you will look at to judge it.

Exercise D — Skills transfer map

- For each protected or augmented task, name the underlying ability (e.g. “reading customers' unspoken needs”).
- List two other places that ability is used. Ask someone who works there whether it travels as you think.

Conversation prompt by bin	Ask
Automate	If this is automated, what happens to the time — and who decides?
Augment	Who trains me on the tool, and who checks its output?
Protect	How will this part of the work be recognised and resourced?
Rethink	Who is accountable for errors, and how will newcomers still learn this?



9. The 12 occupations at a glance (model suggestions)

Editorial ratings; your sort may differ. The interactive atlas lets you disagree and export your version.

Occupation	Automate	Augment	Protect	Rethink
Frontline retail / kirana counter	Answer stock and price questions on WhatsApp Record daily cash, UPI and khata entries	Arrange shelves for festival season	Decide who gets goods on udhaar (credit) Handle a complaint about a short-weight packet	Reorder fast-moving goods from distributors
HR generalist	Draft job descriptions and offer letters Answer routine policy questions	—	Investigate a workplace grievance Run onboarding for a new cohort	Screen CVs against a shortlist Prepare headcount and attrition reports
School teacher	Prepare worksheets and question banks Fill attendance and admin registers	—	Give feedback on a student essay Notice a child who is struggling at home Run a classroom discussion	Grade objective tests
Graphic designer	Produce size variants of one approved design Retouch product photographs	—	Interpret a vague brief with a client Decide the visual direction for a brand	Generate early moodboards Check rights and licensing of assets
Marketing executive	Write first drafts of social posts and ads Segment the customer list Report weekly campaign numbers	Decide positioning against a competitor	Interview customers about why they churned	Approve claims in an ad for legal accuracy
Data / business analyst	Write routine SQL and spreadsheet formulas Build the monthly reporting deck	Decide whether a result is causal or coincidental	Frame the business question with a stakeholder Explain a counter-intuitive finding to leadership	Clean and reconcile messy data
Home care worker	Write the daily care log	—	Help a client bathe, dress and move safely Notice changes in mood or health Coordinate with family and doctors	Remind about medication schedules Plan weekly shopping and meals
Customer support agent	Answer common FAQ-type queries Summarise a ticket history before a call Tag and route incoming tickets	Decide whether to grant an exception refund	Calm an angry customer whose order failed twice	Spot a pattern pointing to a product bug
Bookkeeper / accounts assistant	Enter invoices and receipts	—	Explain cash-flow trouble to a shop owner Chase clients for missing documents	Reconcile bank statements Prepare GST returns Spot signs of fraud or leakage



Occupation	Automate	Augment	Protect	Rethink
Junior software developer	Write boilerplate and CRUD code Write unit tests	Review AI-generated code for security flaws	Clarify requirements with a product manager	Debug a production incident Learn the codebase by doing small tickets
Logistics dispatcher	Plan daily delivery routes Track and update shipment status	—	Assign drivers to shifts Handle a breakdown in the middle of a route Settle disputes between drivers and shopkeepers	Forecast weekly volume
Paralegal / legal assistant	—	Check a contract for jurisdiction-specific risk	Interview a client about what happened	Summarise case documents Search precedents Draft standard agreements Manage filing deadlines



10. Read further

- ILO & NASK (May 2025). *Generative AI and Jobs: A Refined Global Index of Occupational Exposure*, ILO Working Paper 140. ilo.org/node/706766
- World Economic Forum (Jan 2025). *The Future of Jobs Report 2025*. weforum.org/publications/the-future-of-jobs-report-2025/
- McKinsey & Company (5 Nov 2025). *The state of AI in 2025: Agents, innovation, and transformation*; and (25 Aug 2026) *The state of AI in 2026: On the road to ROI*. mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai
- McKinsey & Company (3 Sep 2026). *AI job losses fall short of forecasts* (chart). mckinsey.com/featured-insights/charts/ai-job-losses-fall-short-of-forecasts
- Brynjolfsson, E., Chandar, B. & Chen, R. (2025, rev. 12 Aug 2026). *Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence*. Stanford Digital Economy Lab. digitaleconomy.stanford.edu/publications/canaries-in-the-coal-mine/
- NITI Aayog Frontier Tech Hub (Oct 2025). *Roadmap for Job Creation in the AI Economy*. niti.gov.in
- Anthropic (2025–26). *Anthropic Economic Index* reports. anthropic.com/economic-index — published by an AI developer; this guide was drafted with help from that company's model.

Related Sahilverse worlds

- The Human Capital Commons — humancapitalcommons.sahiltayal.com
- The Kirana Observatory — thekiranaobservatory.sahiltayal.com
- Future of Human Work — work.sahiltayal.online
- The full atlas — sahiltayal.com

Limits of this guide

Sources were selected for relevance and variety, not by systematic review, and were checked on 26 October 2026. Cases are fictional composites. Ratings are the author's judgments. Nothing here predicts employment for any person, firm or country. Corrections are welcome at sahil@sahiltayal.com.

© 2026 Sahil Tayal. Free to share non-commercially with attribution and a link to aijobs.sahilverse.online.