Disnesta Research

OPTIONAL WORK REPORT

AI and the Work It Replaces

The gap between what economies can automate and what they distribute

June 2026

A standing assessment of the distance between what artificial intelligence and robotics can automate and what economies distribute. This edition covers capability, displacement, productivity, the compute divide, policy, and skills.

COMPREHENSIVE EDITION

Disnesta Research · A monthly assessment from optionalwork.com · Published June, 2026

About this report

The Optional Work Report is a monthly assessment of the gap between what economies can automate and what they distribute. This Comprehensive Edition expands the June 2026 briefing into a full chapter treatment with sourcing detail, sector analysis, and supporting data tables. It is a Disnesta Research report prepared by Yaw Adom-Mensah.

How to read it

The report covers two subjects: the capability and cost of automation, and the distribution of the gains it produces. Chapters One and Two cover capability. Chapter Three covers labor market effects. Chapter Four covers the distribution of productivity gains. Chapter Five covers policy. Chapter Six ranks economies on two indices. Chapter Seven covers sectors. Chapter Eight covers skills and adjustment. The Outlook sets out the near term markers used to test the report's central claim.

Sources

Figures are drawn from named public sources, among them the World Economic Forum, the International Monetary Fund, the International Federation of Robotics, the United States Bureau of Labor Statistics, Goldman Sachs, Menlo Ventures, the European Commission, and GiveDirectly. Country standings are taken from the Optional Work Index and the National AI Compute Index, with methodology at optionalwork.com/research. Figures are current as of June 2026.

A note on numbers

Where a source reports a measured value, that value is used. Where a source reports a projection, it is labeled as such. Estimates from different institutions are presented separately rather than combined.

> **Editorial standard.** Measured results and projections are kept distinct, as are capability that has shipped and capability that is announced, and policy that is enacted and policy that is proposed.

Contents

FOREWORD

Capability has outpaced distribution

For three years the central question about artificial intelligence concerned capability: what the technology could do that it could not do before. That question remains open, but it no longer determines outcomes for working people. The determining question is now distributional.

For a growing class of routine cognitive and software mediated work, the capability question has moved from theory to deployment. Systems now plan and execute multi step work, operate software end to end, and perform a growing share of cognitive and, increasingly, physical tasks at a standard and a price that support substitution for paid labor. The binding constraint on optional work is therefore no longer the technical capacity to automate. It is the distribution of what automation produces.

The decisive variable is the widening distance between two rates of change. Capability advances on the cadence of product cycles, measured in months. Redistribution advances on the cadence of legislation, measured in years. While the two remain out of step, output accumulates faster than the institutions that would spread it, and the gains of automation concentrate where capital and compute are already located.

Two findings hold simultaneously. The productivity gains are real, and so is their concentration. A technology can raise output by a quarter in the tasks it touches while leaving most workers no better able to live on the result, because displaced workers do not receive the gains. Whether the period produces broad gains or concentrated ones is not determined by the technology. It is determined by policy, by wage bargaining, and by fiscal choices that remain largely unmade.

The analysis is structured to be testable. It sets out near term markers on both the capability side, which is likely to advance, and the distribution side, which may not. Those markers can be checked against the record over the year ahead.

The constraint on optional work is no longer what economies can automate. It is what they choose to share.

MAIN MESSAGES

The gap between capability and distribution is widening

92M JOBS PROJECTED DISPLACED BY 2030, AGAINST 170M CREATED (WORLD ECONOMIC FORUM)	**60%** OF ADVANCED ECONOMY JOBS EXPOSED TO AI, ABOUT HALF NEGATIVELY (INTERNATIONAL MONETARY FUND)
$37B ENTERPRISE GENERATIVE AI SPENDING IN 2025, UP FROM $11.5B IN 2024 (MENLO VENTURES)	**542k** INDUSTRIAL ROBOTS INSTALLED IN 2024, THE SECOND HIGHEST ANNUAL COUNT ON RECORD (IFR)

Three findings define the period. **Machine capability is advancing into autonomous execution** rather than faster assistance alone. **Displacement is beginning to appear in measured labor market data, not only in forecasts.** The resulting productivity is accruing to capital and to a narrow band of firms, while the principal mechanisms for spreading it, namely wages, taxes, and transfers, remain largely at the proposal stage.

On capability, the year's model releases extended AI from assistant to operator: systems that plan and execute multi step work and operate computers end to end. Cost fell as fast as capability rose, and frontier grade performance is now available at mid tier prices. That combination is what converts a benchmark result into a substitution decision.

On the physical side, automation extended beyond software. The International Federation of Robotics records more than half a million industrial robots installed in a single year and rising investment in humanoid developers, with the first general purpose machines entering commercial deployment. Manual and service work, which forms the base of the labor market, is no longer insulated from automation.

On distribution, the institutional response across governance, automation taxes, and transfers remains early and contested. The European Union has the first binding framework in force, and several governments are debating automation taxes. Evidence that broad distribution is feasible exists but remains experimental. The distance between capability and distribution is the subject of the analysis that follows.

The argument in six steps

1. Autonomy arrived in production

The capability that displaces rather than augments is autonomy: a system that removes the worker from a process rather than accelerating the worker within it. In 2025 and into 2026 that capability became available as a standard enterprise product rather than a research result.

2. The physical economy lost its insulation

For most of the AI era the displaced work was cognitive and the physical economy was safe. Reliable robotic performance of unstructured tasks, at deployable scale, is eroding that line and reaching the occupations that employ the largest numbers of people.

3. Displacement is beginning to appear in data

The earliest signals were projections. The record now holds measured reductions in labor demand and firms naming automation as the cause, with statistical agencies beginning to attribute displacement to AI directly.

4. The gains are concentrating

Enterprise AI spending is the fastest scaling category in software history, and measured productivity gains in knowledge work are large. Those gains are flowing to capital and to high skill labor, raising output without widening who can live on it.

5. Distribution is moving slowest

The institutions that would spread the gains move on the cadence of legislation, while capability moves on the cadence of product cycles. That mismatch, if sustained, is the central risk.

6. Geography decides who captures the gains

None of this reaches an economy that cannot run the compute. Effective access to AI compute is the most unequally distributed input, and it determines which economies produce automation's gains and which are only exposed to them.

KEY INDICATORS

KEY INDICATORS

The numbers that frame the period

Indicator	Value	Source
CURRENT AND MEASURED INDICATORS		
Enterprise generative AI spending, 2025 (surveyed)	$37 billion	Menlo Ventures
Enterprise generative AI spending, 2024 (surveyed)	$11.5 billion	Menlo Ventures
Industrial robots installed, 2024 (measured)	542,000	International Federation of Robotics
Operational robot stock worldwide (measured)	4.66 million	International Federation of Robotics
National AI Compute Index, global mean (current)	23	National AI Compute Index
Optional Work Index, global mean (current)	22.7	Optional Work Index
Optional Work Index, highest economy (current)	35.2 (Singapore)	Optional Work Index
PROJECTIONS AND ESTIMATES		
Projected new roles by 2030	170 million	World Economic Forum
Projected roles displaced by 2030	92 million	World Economic Forum
Net change, structural churn	+78 million (22% churn)	World Economic Forum
Projected robot installations, 2025	575,000	International Federation of Robotics
Advanced economy jobs exposed to AI	60%	International Monetary Fund
Emerging market jobs exposed	40%	International Monetary Fund
Low income country jobs exposed	26%	International Monetary Fund
US work hours exposed to automation	25%	Goldman Sachs
Productivity level lift, full adoption	about 15%	Goldman Sachs
Average task level productivity gain in studies	about 25%	Goldman Sachs, Penn Wharton

The upper block reports measured and surveyed values and current index readings. The lower block reports projections to 2030 and modeled estimates. The two are not combined. Figures are current as of June 2026.

Summary. Capability and cost have made automation a standard option; the labor market is registering the effect; the resulting productivity is concentrating; and the institutions that would spread it have moved only marginally.

CHAPTER ONE

What machines can now do

The capability that matters for labor is autonomy, and this is the period it arrived in production. AI systems have moved from completing tasks a worker assigns to executing multi step work unsupervised: planning, using software, and operating computers end to end.

Autonomy as the threshold that matters

Autonomy displaces rather than augments because it removes the worker from a process rather than accelerating the worker within it. An assistant raises a worker's output and leaves the worker in place. An operator performs the sequence of steps itself, and the worker becomes redundant for that sequence. A given model can serve either function depending on deployment, so the labor market effect of a release is determined not by the model itself but by how firms integrate it into their workflows.

The year's releases bear this out. Systems shipped in 2025 and into 2026 are distinguished less by a single benchmark result than by the length of the task they can complete without human intervention. The relevant measure is whether a model can take an objective, decompose it, act across several tools, recover from its own errors, and return a finished result. That is the functional profile required to substitute for a worker.

Two releases illustrate the shift. OpenAI's GPT-5.4 shipped with a context window measured in the millions of tokens and a native interface for operating a computer directly, and it resolved roughly four in five tasks on the SWE-bench software engineering benchmark while reporting computer use performance at or above a human baseline. Anthropic's Claude Sonnet 4.6 brought capability previously reserved for a top tier model to mid tier pricing and was positioned as a standard enterprise model rather than a premium option. The significant point is not which model leads on a given test, but that benchmarks at the research frontier a year earlier are now resolved in routine production.

An assistant raises a worker's output and leaves the worker in place. An operator performs the sequence of steps itself.

The change of the period was the reliability of long horizon execution. A system that completes a ten step task correctly nine times in ten has different economic value from one that completes it correctly half the time, even where both can produce a correct result on occasion. Reliability is what converts capability into substitution, because a firm will not remove a worker from a process that fails unpredictably. The advance of 2025 and 2026 lay less in new capabilities than in the dependable performance of existing ones, which is the property a substitution decision requires.

This also accounts for the lag between the capability signal and the labor market signal. A model capable of the work exists before firms restructure their workflows to use it. That restructuring is slow and internal, and it is the principal bottleneck on displacement. Capability is necessary but not sufficient; the reorganization of work around it is a second, slower step, and it is the step that converts a model release into a change in headcount.

Cost as the second axis

Capability alone does not produce a labor market effect. Cost is the second factor, and during this period it fell as fast as capability rose. Once top tier performance reaches mid tier prices, the automation of cognitive work ceases to be confined to well capitalized firms and becomes available to most businesses. That diffusion is the mechanism by which a model release acquires labor market consequences, and it is why the cost trajectory matters as much as the capability trajectory.

Frontier model pricing illustrates the point. A leading model at or near the top of public benchmarks now lists at a few dollars per million input tokens and a few tens of dollars per million output tokens. At that level the marginal cost of an automated task falls well below the cost of the human equivalent across a wide class of routine cognitive work. Where the more capable tool is also the cheaper one, adoption is no longer constrained by capital but only by the willingness to reorganize.

The market structure beneath these releases reflects the shift. By the close of 2025 the enterprise market for model access had reorganized around coding and agentic workloads. By the measure compiled in Menlo Ventures' annual enterprise survey, Anthropic held roughly two fifths of enterprise large language model usage, up from around an eighth two years earlier; OpenAI's share had fallen to roughly a quarter from about half; and Google's had risen to about a fifth. Market share is not the focus here, but the speed of the reordering indicates how rapidly the underlying capability and pricing are moving.

FIGURE 1.1 ENTERPRISE LLM USAGE SHARE HAS REORDERED IN TWO YEARS

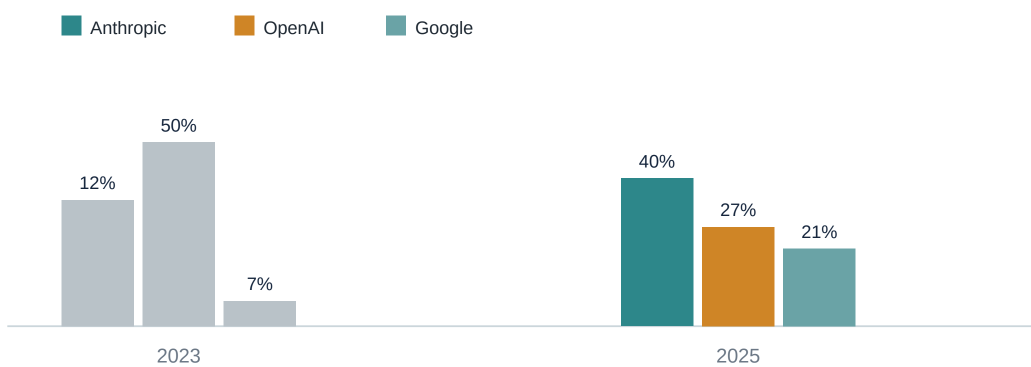

Source: Menlo Ventures, 2025 State of Generative AI in the Enterprise. Shares are of enterprise LLM usage as surveyed.

What the benchmarks do and do not show

Headline benchmark figures overstate real world reliability. The most cited coding benchmark, SWE-bench Verified, was introduced by Princeton researchers in 2023 and became the standard test of agentic software work. By early 2026 frontier scores had compressed near the top of the benchmark, with several models clustered around four in five tasks solved, partly as an artifact of saturation. In February 2026 OpenAI stepped away from SWE-bench Verified after a contamination audit found that leading models could reproduce known solutions verbatim on some tasks and that a share of the test cases were flawed. The field moved to a harder, contamination resistant successor, on which standardized scores fall to roughly a half to three fifths of tasks.

The capability is not illusory, but a single leaderboard figure is a poor guide to system performance in a working codebase, call center, or back office. A benchmark measures a narrow, curated slice of work under favorable conditions. The economically relevant question is reliability across the full distribution of actual tasks. On that question capability is high and rising but uneven, and the gap between a demonstration and a dependable deployment accounts for most of the remaining uncertainty.

Adoption data support this. Even after capability and pricing crossed the threshold, the share of firms using generative AI in regular production remained modest, and the dominant pattern shifted toward purchasing capability rather than building it: roughly three quarters of enterprise AI solutions were bought rather than developed in house by 2025. Capability is ready before the organization is, and that lag is why the labor market signal in Chapter Three is present but still early.

Reading the capability signal

Three claims from this chapter carry into the rest of the analysis, stated here in both their strongest and most cautious forms.

The strong form

Autonomous execution of multi step cognitive work is now available as a standard product at commodity pricing. For a wide class of routine knowledge tasks the cheaper option is the machine, and the remaining barrier is the willingness of firms to reorganize. On current trends that barrier falls steadily rather than abruptly, so displacement appears as a trend rather than a shock.

The cautious form

Benchmark scores overstate dependable performance, adoption remains uneven, and the reorganization of work is slow. The capability is real, but the path from capability to substitution runs through internal change measured in quarters and years rather than weeks. Forecasts of a sudden collapse in employment read the capability curve without the adoption curve.

What both forms share

Neither reading alters the central finding. Whether displacement arrives quickly or slowly, the resulting productivity must still be distributed, and distribution is determined by institutions that are not keeping pace with either reading. The capability debate concerns timing; the distribution debate concerns outcomes, and the latter is the decisive one.

> **Summary.** Capability is sufficient to make automation a standard option for routine cognitive work; cost has removed capital as the constraint; and adoption is the remaining bottleneck, which is now easing. These three conditions frame the labor data in Chapter Three.

CHAPTER TWO

Automation leaves the screen

For most of the AI era the physical economy was insulated and the displaced work was cognitive. That insulation is eroding. The threshold being crossed is reliable robotic performance of unstructured tasks at deployable scale, rather than any single machine.

542k INDUSTRIAL ROBOTS INSTALLED IN 2024, MORE THAN DOUBLE THE COUNT A DECADE EARLIER	**4.66M** OPERATIONAL ROBOT STOCK WORLDWIDE, UP 9% YEAR ON YEAR
$4.6B INVESTED IN HUMANOID DEVELOPERS IN 2025	**~1,000** HUMANOID ROBOTS IN VERIFIED ACTIVE COMMERCIAL DEPLOYMENT BY EARLY 2026

Sources: International Federation of Robotics, World Robotics 2025; humanoid investment and deployment figures as reported by IFR and TechCrunch.

Industrial robotics at scale

The International Federation of Robotics, in its World Robotics 2025 report, counted 542,000 industrial robots installed in a single year, more than double the number a decade earlier and the second highest annual total ever recorded, only marginally below the all time peak. Annual installations have now topped half a million units for four consecutive years, and the total operational stock of industrial robots in service worldwide reached 4.66 million units, a rise of roughly a tenth in a year. Installations are projected to keep climbing, to around 575,000 in 2025 and past 700,000 by 2028.

The installed base is geographically concentrated. Asia accounted for roughly three quarters of new deployments, and China alone took more than half of all industrial robots installed worldwide. The composition of Chinese demand is notable: domestic manufacturers now supply the majority of robots installed in China, a reversal from a decade earlier when foreign suppliers dominated. The capacity to build automation, and not only to buy it, is consolidating in the same economies best positioned to deploy it. The pattern recurs throughout the analysis and culminates in the compute divide of Chapter Six.

FIGURE 2.1 INDUSTRIAL ROBOT INSTALLATIONS, RECENT AND PROJECTED

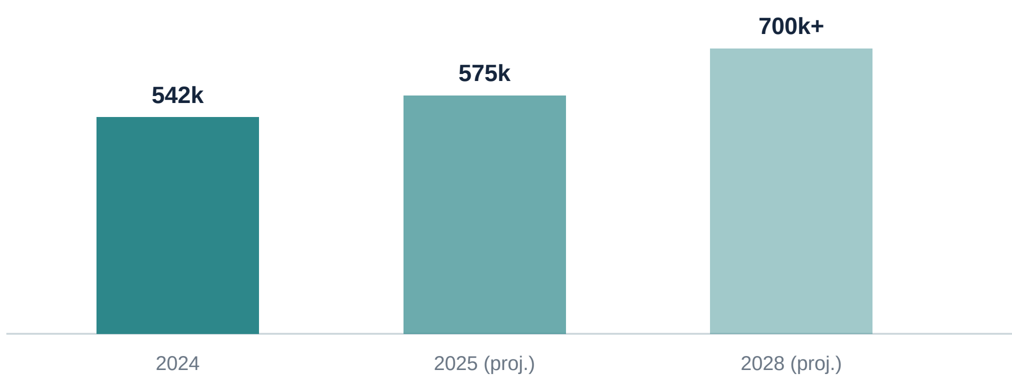

Source: International Federation of Robotics, World Robotics 2025. 2024 is a measured installation count; 2025 and 2028 are IFR projections.

A note of proportion applies. Industrial robots are specialized machines with few joints, tuned to repetitive, millimeter precise tasks at high speed, and the IFR holds that they will remain the backbone of high speed manufacturing for the foreseeable future. The growth above is substantial but incremental, a continuation of an established trend rather than a break from it. Any break lies in a different category of machine.

The humanoid threshold

The development that widens displacement is the general purpose robot, built to operate in environments designed for the human body rather than in a fixed cell. Investment in humanoid developers reached roughly $4.6 billion in 2025 by the IFR's accounting, and major technology firms, among them NVIDIA, Tesla, and Amazon, have announced significant programs alongside substantial funding from defense agencies. By early 2026 about a thousand humanoid robots were in verified active commercial deployment worldwide. The figure is small in absolute terms but significant, because it marks the transition from demonstration to deployment.

Two reference points define the frontier. Boston Dynamics demonstrated its Atlas robot performing autonomous factory tasks at a Hyundai plant, an instance of a general purpose machine carrying real work in an industrial setting rather than performing a choreographed routine. The Rivian backed developer Mind Robotics raised $615 million at a valuation of around $2 billion, a financing scale that indicates investor expectation that general purpose robotics is approaching commercial readiness. The significant point is the threshold itself: robots performing unstructured tasks reliably enough to deploy at scale.

A note on caution

The IFR counsels against extrapolating from staged demonstrations. The form follows function critique holds that the human body is not the optimal design for many tasks, and that humanoids will complement rather than replace the specialized robots already on factory floors. Public demonstrations of dancing or running humanoids are not evidence of production capability. Mass adoption in factories and homes is, by the IFR's assessment and by China's national planning, a development for the end of the decade rather than its middle.

That caution qualifies rather than negates the finding. Humanoids will not enter workplaces at scale within the year, but the category has crossed from research into early deployment, and the direction is established. Cognitive automation pressures the segment of the labor market occupied by routine knowledge work; physical automation reaches the base of the market, the manual and service occupations that employ the largest numbers of people and that were assumed safe longest because they required a body in an unstructured environment. Once a machine can reliably operate in that environment, the assumption no longer holds.

This is the structural basis for the labor market findings that follow. Cognitive displacement and physical displacement are two fronts of the same advance arriving at different times, and the second front reaches the larger population. A transition confined to knowledge workers reshapes a set of professions; one that also reaches manual and service work reshapes the labor market as a whole.

Where cognitive automation pressures a slice of the labor market, physical automation reaches its base.

National strategies and the embodied turn

The robotics buildout is increasingly a matter of industrial policy rather than private investment alone. China has placed AI powered robotics at the center of its national strategy, with its planning documents identifying humanoids as a potential disruptive technology on the order of the computer or the smartphone and setting targets for their eventual mass production. The country already operates an industrial robot stock several times larger than the next largest national base, and it is steadily substituting domestic suppliers for foreign ones across the value chain.

The United States approaches the same frontier from a different position. It is a leading market for industrial robots and a leading source of capital and research for humanoid development, yet it imports most of its industrial robots and has comparatively few domestic manufacturers, relying instead on a deep ecosystem of system integrators. Europe, for its part, emphasizes collaborative robots designed to work alongside people and frames the question in terms of safety and augmentation rather than replacement.

These differences are consequential. They determine which economies capture the gains of physical automation as producers and which experience it primarily as consumers and as exposed labor markets. The logic that governs compute, set out in Chapter Six, also governs robotics: the capacity to build the machines, and not only to use them, concentrates the gains. An economy that imports its automation gains the productivity but not the value added from production; an economy that builds it retains both.

Region	Posture toward general purpose robotics
China	National strategy, mass production targets, domestic supply substitution
United States	Capital and research leadership, limited domestic manufacturing, integrator depth
Europe	Collaborative robotics, safety and augmentation emphasis
Japan, Korea	Established manufacturing base, social and care robotics tradition

Source: International Federation of Robotics positioning papers and World Robotics 2025. Postures summarized from IFR analysis.

CHAPTER THREE

The labor market

Displacement is beginning to appear in data, not only in forecasts. The earliest signals were projections; the record now holds measured reductions in labor demand and firms naming automation as the cause.

From projection to data

The labor market is where the capability of the preceding chapters registers in measured outcomes, and it is where the evidence has shifted most. United States job openings have fallen by roughly two fifths from their 2022 peak. More than 260,000 technology workers were laid off in 2023, with many employers citing automation among the reasons. The 2023 Hollywood strikes made AI displacement a central labor demand for the first time in a major industrial action, securing contract language on the issue.

No single one of these facts establishes that AI caused the change. Labor markets move for many reasons, and a cooling from a pandemic era peak has causes unrelated to automation. The narrower and supportable finding is that the direction of labor demand has turned in the sectors most exposed to AI, that firms increasingly name automation as a reason, and that statistical agencies are beginning to attribute displacement to AI directly, with the United States Bureau of Labor Statistics expected to report AI attributable displacement in its annual data. The signal has moved from anecdote into the official record.

The most specific early evidence concerns younger workers in exposed fields. By one major bank's research, unemployment among workers in their twenties in technology exposed occupations rose by close to three percentage points from the start of 2025, faster than for their peers in other trades and for technology workers as a whole. This is consistent with automation pressuring entry level positions first, because entry level work concentrates the routine, well specified tasks that current systems handle best, and because firms can decline to hire more readily than they can decline to retain. The effect falls on the entry rung, which is where careers begin.

The shape of the churn

The headline numbers are large in both directions, and the relationship between them is the central point. The World Economic Forum projects 170 million new roles created by 2030 against 92 million displaced, a net increase of 78 million jobs, with gross churn amounting to 22 percent of the roughly 1.2 billion formal jobs in its dataset. The net figure understates the disruption: a net gain of 78 million conceals 92 million separations and the obsolescence of close to two fifths of current skills within five years. A worker displaced from a declining role is not necessarily the worker hired into a growing one, and the distance between the two in skill, geography, and time is where the hardship concentrates.

FIGURE 3.1 PROJECTED JOBS CREATED AND DISPLACED BY 2030

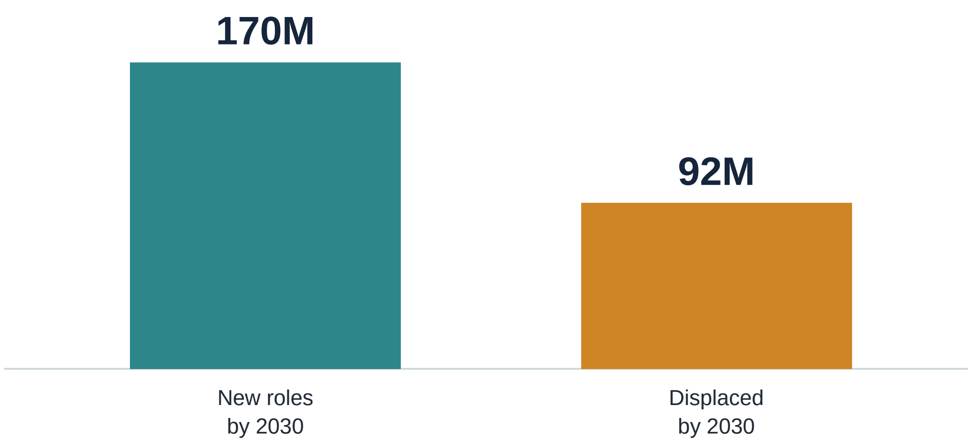

Source: World Economic Forum, Future of Jobs Report 2025. Both values are projections to 2030; net change is plus 78 million.

The churn takes the form of polarization rather than uniform loss. Roles are created at the top, in technology, data, and AI specialisms, and grow at the frontline base, where the Forum projects gains for farmworkers, delivery drivers, construction workers, and care roles driven by demographic and green transition trends rather than by automation. Roles contract in the middle and at the entry level, where routine cognitive tasks concentrate. The result is a hollowing of the middle of the distribution, the clerical and administrative roles that once offered a stable path, even as both ends grow.

The International Monetary Fund describes the scale of exposure and its uneven geography. By its analysis, AI will affect about 60 percent of jobs in advanced economies, roughly 40 percent in emerging markets, and about 26 percent in low income countries, with the higher exposure of rich economies reflecting their concentration of cognitive, office based work. Of the affected jobs in advanced economies, the Fund estimates that roughly half may be harmed, with AI taking over key tasks and lowering the demand for labor, while the other half may benefit from higher productivity. The same analysis warns that, in most scenarios, AI is likely to worsen overall inequality, a conclusion that anticipates the distribution argument of Chapter Four.

FIGURE 3.2 SHARE OF JOBS EXPOSED TO AI, BY ECONOMY GROUP

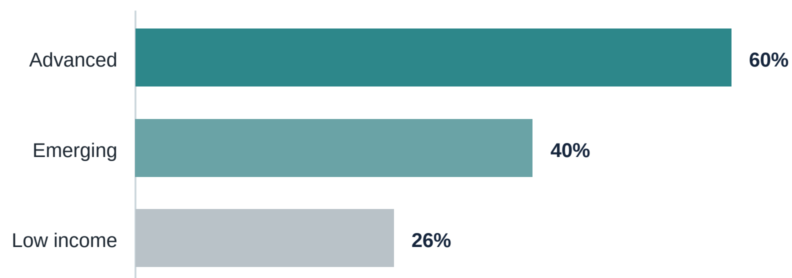

Source: International Monetary Fund, Gen-AI: Artificial Intelligence and the Future of Work. About half of advanced economy exposure is estimated to be negative.

A leading bank's research adds a sober counterweight to the largest figures. While roughly 300 million jobs across the United States and Europe are exposed to some degree of AI automation, and while AI could in principle automate tasks accounting for around a quarter of all United States work hours, the same research finds that under current use cases only about 2.5 percent of United States employment would be at risk of displacement. Exposure is not the same as substitution. Most jobs are partially exposed, which historically means complemented rather than replaced, and the conversion of exposure into job loss depends on adoption, which remains the bottleneck identified in Chapter One.

Reconciling these figures yields the chapter's conclusion. The largest numbers describe potential, the breadth of work that touches tasks AI can perform. The smallest numbers describe the near term, the share of jobs at risk given current usage. The outcome for the period sits between them and shifts over time, as adoption converts potential into realized displacement at a pace set by the speed of firm reorganization. The labor market is registering the beginning of that conversion rather than its end.

What the polarization means for a working life

Aggregate churn figures obscure the effect of polarization on individual careers. The growing roles at the top require skills that take years to build and that are unevenly distributed across the population. The growing roles at the base are frontline and care roles that pay modestly and are physically demanding. The contracting roles in the middle are the clerical, administrative, and routine analytical jobs that historically allowed a worker without a specialized degree to reach a stable, middle income position.

As the middle contracts, the path from the base to the top loses its intermediate steps. A worker can no longer progress from an entry level administrative role into a more skilled one within the same firm, because the intermediate roles are those being automated. The Forum's finding that close to two fifths of current skills will be transformed or obsolete within five years indicates how frequently a worker must reskill to remain in place, and how much of the cost of that reskilling falls on the worker rather than the firm or the state.

The institutional response is reskilling, and employers report an intention to provide it: a large majority report plans to prioritize upskilling. Whether intention translates into provision at the necessary scale is an open question, addressed directly in the distribution chapters. Reskilling is a redistribution of opportunity, and like the other distribution mechanisms examined here, it is advancing more slowly than the capability that makes it necessary.

> **The entry level signal.** The clearest early evidence of AI displacement is concentrated among younger workers in exposed fields. If automation continues to pressure entry level positions first, the long run effect bears not only on who holds jobs today but on whether the careers of the next decade can begin.

Measured or projected signal	Reading
US job openings down about 40% from 2022 peak	Labor demand cooling, exposed sectors leading
260,000+ tech layoffs in 2023, automation often cited	Firms naming the cause
Tech worker unemployment in their 20s up ~3pp from early 2025	Entry rungs pressured first
WEF: 170M created, 92M displaced, 22% churn	Net gain over substantial gross churn
IMF: 60% advanced economy exposure, half negative	Scale of exposure, uneven geography
Goldman: 2.5% of US employment at near term risk	Exposure is not yet substitution

CHAPTER FOUR

Where the productivity goes

The productivity gains are real, which makes their distribution the decisive question. The issue is not whether AI produces gains but who receives them.

The financial signature of the shift

Enterprise spending on generative AI reached 37 billion dollars in 2025, up from 11.5 billion in 2024, a 3.2 fold rise in a single year and, on the report's longer baseline, roughly a sixteen fold increase over 2023. Menlo Ventures, which compiles the figure from a survey of around 500 enterprise decision makers and a bottom up market model, calls it the fastest scaling category in the history of software, reaching a scale in three years that earlier software categories took more than fifteen years to achieve. That spending is the financial signature of work shifting from people to machines.

The composition of the spending tells a sharper story than the total. More than half of it, roughly 19 billion dollars, went to the application layer, the products that sit on top of the models and do specific jobs, while the remainder, around 18 billion, went to the infrastructure beneath. Within the application layer, coding was the single largest category at about 4 billion dollars, the clearest case of a model capability translating directly into a category of automated professional work. Vertical applications aimed at specific industries drew about 3.5 billion, with healthcare the leading buyer.

CHAPTER FOUR · WHERE THE PRODUCTIVITY GOES

FIGURE 4.1 ENTERPRISE GENERATIVE AI SPENDING, 2023 TO 2025

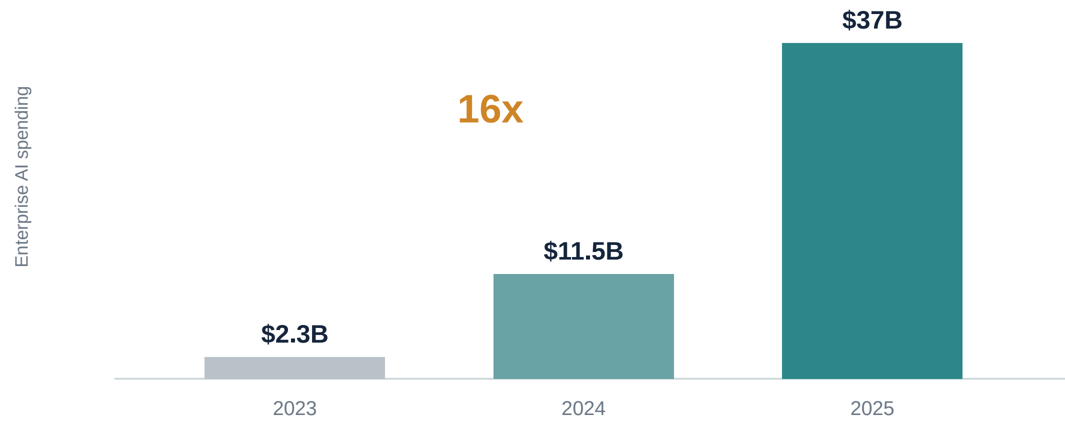

Source: Menlo Ventures, State of Enterprise AI. The 2024 and 2025 figures are surveyed; the multiplier is measured against the 2023 baseline.

The return that justifies this spending is a measurable lift in productivity. Goldman Sachs estimates that, once fully adopted and incorporated into regular production, generative AI will raise the level of labor productivity in the United States and other developed markets by around 15 percent, a lift its economists value at roughly 4.5 trillion dollars in annual output for the United States economy. Studies of task level performance find larger gains still: a review of the literature finds productivity improvements ranging from about 10 to 55 percent, averaging around 25 percent, with projected cost savings expected to grow toward 40 percent over the coming decades. United States labor productivity has posted its strongest growth since 2010.

The same qualification from Chapter One applies. The large productivity figures describe potential at full adoption, and adoption remains partial. Aggregate productivity statistics have only recently begun to move, because the share of firms using AI in regular production is still modest. The gains are present where AI is deployed, but the economy wide effect is still emerging. Spending of 37 billion dollars runs ahead of the measured macroeconomic result, the pattern expected early in a technology's diffusion rather than at maturity.

Aggregate abundance is not shared abundance

Aggregate gains are not the same as shared gains. Productivity gains that accrue to capital and to a narrow band of high skill labor raise output without broadening who can live on it, because displaced workers do not receive them. A 15 percent rise in productivity raises what the economy produces by 15 percent; it does not, in itself, raise what a displaced worker can spend. The mechanisms that would connect the two, namely rising wages, taxes on the gains, and transfers to those who lose work, are the mechanisms that Chapter Five finds advancing slowest.

The distributional concern is supported by analysis. The International Monetary Fund concludes that, in most scenarios, AI is likely to worsen inequality within countries, because the technology tends to complement high income, cognitively intensive work while substituting for the middle. Where a technology raises the return to capital and to scarce high skill labor faster than it raises wages at the middle and bottom, the gap widens even as the total grows. Output and inequality can rise together in this transition, and on present evidence they are doing so.

A 15 percent lift in what the economy produces is not, by itself, a 15 percent lift in what a displaced worker can spend.

The evidence that broad distribution is possible exists, but it remains experimental rather than institutional. The clearest demonstration is the long running basic income study that GiveDirectly has conducted in rural Kenya since 2017, the largest and longest study of its kind, a roughly 30 million dollar program reaching around 23,000 people across nearly 200 villages, with some recipients guaranteed a basic income for twelve years. Independent researchers, including economists at the Massachusetts Institute of Technology, found that recipients did not work less or drink more, the outcomes critics predict. They invested, became more entrepreneurial, and earned more, and the guarantee of a long horizon income in particular encouraged saving and risk taking.

The study establishes feasibility rather than policy. It shows that unconditional income, at meaningful scale and over a long horizon, can be delivered and that it produces investment rather than idleness. It does not show that any wealthy economy facing AI driven displacement has decided to build such a mechanism. Feasibility is demonstrated; the corresponding policy is not, and Chapter Five measures the distance between them.

Where the application spending lands

Because the application layer is where AI does specific jobs, its composition is a map of which kinds of work are being automated first and most profitably. The table below sets out the segments of enterprise application spending, and each line is also a line about labor: the larger the spend, the more directly a category of human work is being substituted or compressed.

Segment, 2025	Spend	What it automates
Application layer, total	$19B	Products doing specific jobs on top of models
Horizontal AI	$8.4B	Productivity across all functions
Departmental AI	$7.3B	Specific roles such as engineering and sales
Coding, within departmental	$4.0B	Software development, the single largest category
Vertical AI	$3.5B	Industry specific, led by healthcare
Infrastructure	$18B	Foundation models, training, and tooling

Source: Menlo Ventures, 2025 State of Generative AI in the Enterprise. Figures are surveyed enterprise spend for 2025.

Two facts in this table carry the chapter's argument. First, coding alone draws 4 billion dollars, more than any other application category, which is why software recurs as the leading edge of automation throughout the analysis. Second, the application and infrastructure layers together form a market in which the returns accrue principally to the firms that build and sell the capability and to the firms productive enough to deploy it, rather than to the workers whose tasks it absorbs. The composition of spending indicates where the productivity is accruing, and it is not accruing to the displaced.

> **The central question.** Whether the period's gains are broadly or narrowly shared is the most important question in the transition, and it is not determined by the technology. It is determined by the wage bargain, the tax system, and transfers, examined next.

The productivity and employment paradox

The same evidence that establishes the productivity gains also explains why the labor disruption to date is smaller than headline figures suggest; the two readings are consistent. Goldman Sachs estimates that some 300 million jobs across the United States and Europe are exposed to AI automation, and that roughly a quarter of all work hours in the United States could be performed by the technology. These figures are the source of the displacement pressure documented in Chapter Three.

The same analysis finds that only about 2.5 percent of United States employment is at near term risk of outright displacement under current use cases. The gap between the two figures reflects the difference between exposure and substitution. Most exposed work is partly automated rather than fully replaced, which raises the output of the worker who retains the job before it removes the job. Productivity can therefore rise sharply while measured unemployment moves slowly, and the labor signal of the period is compression and slowed hiring rather than mass layoffs.

The current calm should not be read as a settled outcome. Exposure that appears as augmentation today becomes substitution as capability advances and cost falls, the supply trajectory treated here as nearly certain. The 2.5 percent figure is a snapshot under current use cases rather than a ceiling. The displacement is real but early, the productivity is real and already present, and the interval between the two is the period in which the distribution question of the next chapter has to be addressed.

> **Exposure and substitution.** A job exposed to AI is one whose tasks the technology can perform, not one it has taken. Near term risk is narrow and long term exposure is wide, and the difference between them is time. Current displacement is the beginning of the trajectory rather than its conclusion.

CHAPTER FIVE

The policy response

The institutions that would distribute the gains are moving slowest, and that lag is the central risk. The developments here are institutional rather than technical, and they remain early and contested.

What has actually become law

The most concrete institutional development is the European Union's Artificial Intelligence Act, the first binding, horizontal AI governance framework anywhere in the world. It entered into force on the first of August 2024 and applies in phases through 2027. Its bans on prohibited practices took effect in February 2025; its obligations for general purpose AI models took effect in August 2025; and most of its high risk obligations apply from August 2026. The Act classifies systems by risk and carries penalties of up to 35 million euros or 7 percent of global annual turnover, whichever is higher, which makes it the first AI rulebook with teeth proportionate to the largest developers.

The Act's implementation shows both the reach and the limits of governance. Its Code of Practice for general purpose models, finalized in July 2025, has been signed by around two dozen providers, among them Anthropic, OpenAI, Google, Microsoft, Amazon, and Mistral, an indication that the major developers are, for now, choosing to demonstrate compliance rather than contest the framework. Yet at least one large developer, Meta, declined to sign, and a subsequent simplification package proposed in late 2025 could push some high risk deadlines later still. Governance is real, it is binding, and it is already being negotiated down at the margins even before it fully applies.

The Act governs how AI is built and deployed; it does not distribute the gains. It sets rules for safety, transparency, and accountability, which are necessary, but it contains no automation tax, no automation dividend, and no transfer mechanism. The most developed instrument of AI policy is a framework for managing risk rather than for sharing surplus. The governance question and the distribution question are distinct, and far more progress has been made on the first than on the second.

EU AI Act milestone	Date
Entry into force	1 August 2024
Prohibited practices and AI literacy apply	2 February 2025
General purpose AI model obligations apply	2 August 2025
Most high risk obligations and penalties apply	2 August 2026
High risk AI in regulated products	2 August 2027

Source: European Commission, Regulation (EU) 2024/1689. Subsequent simplification proposals may adjust later deadlines.

What remains proposal

On the distribution side, the developments are a list of debates rather than laws. Several member states have opened discussion of automation taxes, levies on the substitution of capital for labor intended to fund the social response. Capable models have been open sourced, beginning with Meta's release of its LLaMA family, which widens access to capability and is, in a limited sense, a form of distribution: it spreads the means of production rather than the proceeds. Against these sit countervailing moves, including OpenAI's transition toward a for profit structure, which sharpened the question of how the gains of firms built on publicly funded research are to be shared.

The asymmetry between the two sides is the chapter's finding, expressible as a difference in cadence. Capability advances on the cadence of product cycles, measured in months. Redistribution advances on the cadence of legislation, measured in years. If that mismatch persists, output accumulates faster than the means to spread it, and the central gap widens by arithmetic rather than by intention.

Capability advances on the cadence of product cycles. Redistribution advances on the cadence of legislation.

The proposals on the horizon

The instruments that would close the distribution gap exist as concrete proposals, set out below because the Outlook treats their enactment or non enactment as the decisive test of the coming year. None is law at scale in any major economy. Each is specified, debated, and unproven.

- **A first national automation tax.** A levy on automation or on the capital that displaces labor, designed to fund the social response. Debated in several jurisdictions, enacted in none at national scale.
- **Automation dividend legislation.** A mechanism to route a share of automation's gains directly to citizens, treating the productivity surplus as partly common property.
- **Median wages rising to capture a share of productivity.** Not a single law but a labor market outcome, in which the wage bargain shifts enough that workers capture part of the productivity gain rather than seeing it accrue entirely to capital.
- **A first government funded basic income pilot at scale.** A publicly financed counterpart to the privately funded GiveDirectly study, which would move unconditional income from philanthropy and experiment into state policy.

These proposals are tracked because they are testable. Within a year each will either have advanced or not, and the record can register which. They serve as the distribution side's benchmark, measuring institutional commitment rather than capability. The supply side is likely to meet its targets; whether any of these four advances is the open question of the period.

> **Summary of policy.** The first binding rules for how AI is made and used are in force, while the mechanisms for sharing AI's gains are largely unbuilt. The first is governance, the second is distribution. The transition's outcome depends more on the second, and the second is where almost nothing has been enacted.

Developments on the distribution side over the next year will matter more for the outcome than developments on the capability side, because the capability side is nearly certain to advance while the distribution side may not. The most consequential developments of the coming year are therefore the least visible, because legislation proceeds slowly and quietly while model releases are rapid and prominent.

CHAPTER SIX

The compute divide and where countries stand

None of this reaches an economy that cannot run it, and access to AI compute is distributed more unequally than any other input.

Effective access, not raw hardware

The decisive resource of the transition is the compute on which models run, which cannot be copied, rather than the models themselves, which can. Compute access is not only a question of chip ownership but of effective access: whether an economy can afford to run and power the hardware at scale. The National AI Compute Index measures this, conditioning raw hardware on affordability and energy, and finds an average of 23 across 150 economies, with access concentrated in a small group of strong performers and a long tail near zero.

A mean of 23 on a 0 to 100 scale, with a long tail near zero, is among the most concentrated distributions of any economic input. The capacity to produce automation's gains, as distinct from importing its effects, is held by a small number of economies. The distribution of effective compute access closely tracks which economies will capture the gains of the transition.

CHAPTER SIX · THE COMPUTE DIVIDE

FIGURE 6.1 EFFECTIVE AI COMPUTE ACCESS, LEADING ECONOMIES

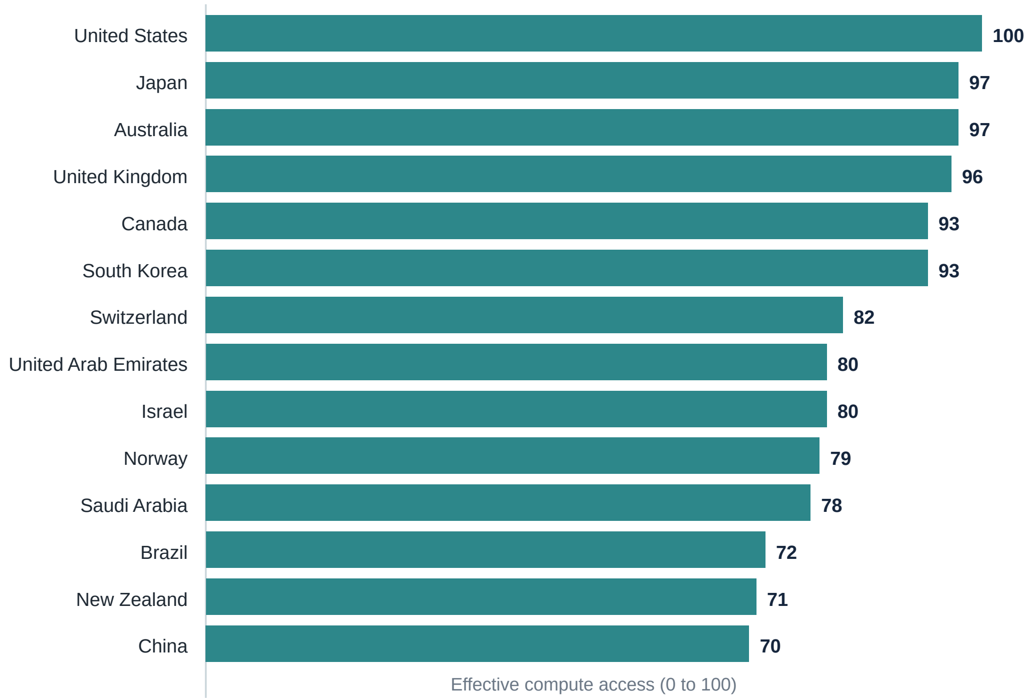

Source: National AI Compute Index, June 2026. Values shown for the leading economies; the global mean across 150 economies is 23.

This divide sets the geography of the transition. The economies positioned to capture automation's gains are those with the compute to produce them, while those most exposed to displacement through global labor markets are often least able to build offsetting capacity at home. An economy whose workers compete in global services markets can experience the displacement of automation without owning the means to produce its gains. The compute divide therefore compounds the labor exposure of Chapter Three rather than offsetting it.

The Optional Work Index

The Optional Work Index scores how optional paid work has become across 150 economies on a 0 to 100 scale, combining capability and distribution into a global mean of 22.7, with Singapore highest at 35.2. The index is constructed so that a high score requires both a large supply of automation and the means to distribute its proceeds. Leaders hold their position by sharing the gains rather than by automating the most: an economy that automates heavily but shares little does not lead, whereas one that automates moderately and distributes well can.

FIGURE 6.2 THE OPTIONAL WORK INDEX, SCALE AND REFERENCE POINTS

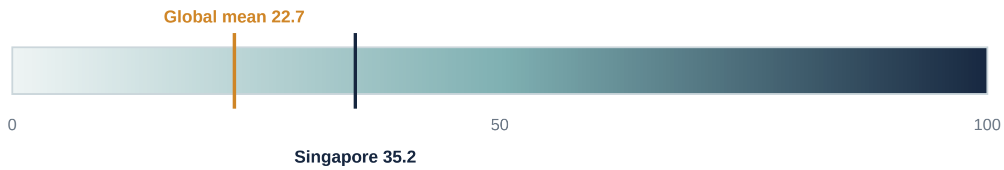

Optional Work Index (0 to 100). Even the leader sits near the bottom third of the scale.

Source: Optional Work Index, June 2026. Country level detail is published in the live index; figures shown are the global mean and the highest scoring economy.

The notable feature of the index is the absolute level rather than the spread between economies. A global mean of 22.7 and a leader at 35.2 indicate that even the most advanced economy has made paid work only modestly optional. The transition is real and accelerating but early. No economy has reached a state in which work is broadly optional for its population. Capability is arriving faster than distribution, so the practical experience of optional work remains limited everywhere.

Read together, the two indices describe the same divide from two angles. The National AI Compute Index measures which economies can produce the gains. The Optional Work Index measures which have converted those gains into reduced dependence on paid work. The first is concentrated; the second is uniformly low and led by economies that combine capacity with a commitment to distribution. The distance between the two indices, between producing gains and distributing them, is the subject of the analysis expressed numerically.

Why geography is destiny in this transition

The compute divide is unlikely to be eliminated by markets, because compute combines three properties that favor concentration: high capital intensity, dependence on cheap and abundant energy, and increasing returns to scale in the surrounding ecosystem of talent, data, and infrastructure. Each rewards economies that already lead and penalizes those that lag. An economy without affordable power cannot run the hardware even where it can purchase it, and one without the ecosystem cannot retain talent even where it trains it.

The compute divide therefore underlies the labor and distribution chapters rather than standing apart from them. An economy's position on the compute index largely determines whether it experiences AI as an owner or as a worker. Owners capture the productivity, set the terms, and build offsetting capacity. Workers experience the displacement through global labor markets with limited domestic capacity to absorb it. The same logic governs robotics, as Chapter Two showed: the capacity to build the machines, and not only to use them, determines who retains the value.

For the most exposed economies, the policy implication matches the Fund's: the priority is to build the foundation, digital infrastructure and a digitally capable workforce, before the divide hardens. For the leading economies, the priority is the distribution question of Chapter Five, because they will capture the gains regardless and the open question is how those gains are shared. The geography sorts economies into two distinct policy problems, and a single global response will not serve both.

> **Reading the two indices.** Compute access is concentrated near the top of the scale for a small number of economies and near zero for a long tail. Optional work is low everywhere and led by economies that pair capacity with distribution. The first index identifies which economies can produce the gains; the second, which are sharing them. The gap between them characterizes the transition.

CHAPTER SEVEN

Four sectors, four speeds

The transition does not arrive everywhere at once. It moves fastest where work is digital, well specified, and expensive, and slowest where it is physical, ambiguous, and low paid. Four sectors span the range.

Software: the leading edge

Coding is the first established category of automated professional work, and the evidence is clear. It is the single largest application of enterprise AI spending, at roughly 4 billion dollars, more than any other category. Around half of developers now use AI coding tools daily, and teams report velocity gains of 15 percent and more across the software development lifecycle, from prototyping to deployment. Capability rose sharply once models could interpret entire codebases and execute multi step tasks rather than completing isolated functions.

Software leads for reasons that generalize. Its work is text, its correctness is testable, and its practitioners are highly paid, so a tool that automates part of it pays for itself immediately and can be verified automatically. The same three properties, a digital medium, testable output, and high wages, predict which other kinds of work will follow. They also account for the early labor signal of Chapter Three: pressure on entry level technology hiring is most visible where automation is most advanced.

CHAPTER SEVEN · FOUR SECTORS, FOUR SPEEDS

Healthcare: the surprising fast adopter

Long regarded as a digital laggard, healthcare has become one of the fastest adopters of AI, deploying it at roughly twice the rate of the broader economy and leading vertical AI spending. The value is concentrated in administrative automation, above all in medical scribing, where AI reduces the non clinical workload that consumes a large share of clinicians' time. Here automation augments rather than replaces: it removes paperwork from a profession short of practitioners, raising the output of scarce labor rather than displacing abundant labor.

Healthcare illustrates the beneficial case, the half of exposure the Fund describes as potentially positive. Where a sector faces a shortage of skilled workers and a surplus of routine administrative burden, automation relieves the burden and allows the scarce worker to perform more of the valuable work. The distributional concern of Chapter Four is weaker here, because the gains accrue partly to overstretched workers and to patients rather than to capital alone. Not every sector resembles healthcare, but the case shows that the outcome depends on the structure of the work rather than on the technology alone.

Creative and knowledge work: the contested middle

The creative and knowledge professions occupy the contested middle, where augmentation and displacement are hardest to separate. Employment growth in marketing, graphic design, office administration, and call centers has run below trend amid reports of reduced labor demand attributed to AI efficiency. These are the routine cognitive roles concentrated in the middle of the distribution, the roles that polarization removes in Chapter Three. The 2023 Hollywood strikes made the conflict explicit, turning AI displacement into a formal bargaining demand for the first time in a major industrial action.

This is the sector where the central tension is most visible in individual careers. The same tool that makes a skilled designer or analyst more productive also reduces the number of designers or analysts a firm needs. Which effect dominates depends on whether demand for the output grows fast enough to absorb the productivity. Where demand grows sufficiently, the tool augments; where it does not, the tool displaces. The balance in the creative middle is being determined now, and the early evidence indicates compression.

Manufacturing and logistics: the slow front opening

Physical work has been the slowest front, protected by the difficulty of operating reliably in unstructured environments, and it is now opening. Industrial robotics continues its incremental climb, with installations above half a million units a year and an operational stock past 4.6 million, concentrated in Asia and above all in China. The newer development is the general purpose robot, still small in number but past the threshold from demonstration to deployment, extending toward the manual and service occupations that employ the largest share of the workforce.

Manufacturing and logistics matter most for the breadth of the transition. The three preceding sectors, software, healthcare administration, and creative work, are large but bounded, whereas physical and service work forms the base of the labor market. As Chapter Two set out, the significance of the humanoid threshold is that it extends automation from a segment of the workforce to its foundation. This is the slowest of the four fronts, and the one whose eventual arrival reshapes the labor market most.

Sector	Speed	Dominant effect so far
Software	Fastest	Substitution at the entry level, augmentation above it
Healthcare admin	Fast	Augmentation of scarce skilled labor
Creative and knowledge	Moderate	Contested, tilting toward compression
Manufacturing and logistics	Slowest, now opening	Incremental robotics, early general purpose deployment

Synthesis of Menlo Ventures, Goldman Sachs, and International Federation of Robotics data cited in this report. Effects are early stage readings.

> **The pattern across sectors.** Automation arrives first where work is digital, testable, and well paid. It augments where labor is scarce and displaces where labor is abundant and routine. The same technology produces different outcomes depending on the structure of the work, so general statements about whether AI helps or harms workers are unreliable. The outcome depends on the sector and on the distribution of the gains.

CHAPTER EIGHT

The skills and adjustment response

Displacement and creation do not involve the same workers. The net figure of Chapter Three conceals a churn that only retraining can bridge, and the scale of that retraining is a direct measure of the difficulty of the transition.

The size of the retraining problem

The net gain of 78 million roles, 170 million created against 92 million displaced, rests on an assumption that should be made explicit: that workers losing roles can move into the roles being created. The World Economic Forum quantifies the difficulty. About 22 percent of all jobs are expected to be structurally transformed by 2030, churning beneath the net total, and roughly 39 percent of the core skills a worker needs are expected to become outdated over the same period. A workforce whose required skills turn over by nearly two fifths in five years does not adjust without intervention.

This is the adjustment problem. Job creation at the top of the distribution does not assist a displaced worker in the middle unless that worker can reach the new role, which requires acquiring skills absent from the job left behind. The faster capability advances, the faster the required skills turn over, and the wider the gap between the work contracting and the work expanding. Skill turnover converts a favorable net figure into individual hardship, and it is rising.

CHAPTER EIGHT · THE SKILLS AND ADJUSTMENT RESPONSE

Who is expected to do the retraining

Employers report an awareness of the problem. In the Forum's survey, 85 percent name upskilling their existing workforce as a primary strategy for the period to 2030, 77 percent plan to retrain workers to operate alongside automated systems, and 70 percent expect to hire for newly created skills. Against these intentions, 63 percent of employers name skills gaps as the single largest barrier to adopting the technology. The shortage that constrains adoption is the same one that retraining is intended to close, so the skills response is at once the obstacle and the remedy.

The allocation of the cost of retraining remains unsettled. Where the cost falls to the firm, it reaches the workers the firm chooses to retain, who are rarely the workers most exposed to displacement. Where it falls to the worker, it falls hardest on those with the least time and money to absorb it. Where it falls to the state, it joins the distribution mechanisms from Chapter Five that remain largely proposals. Retraining is widely endorsed in principle and unevenly funded in practice, the same pattern that characterizes the rest of the distribution side.

What employers report for 2025 to 2030	Share
Name upskilling as a primary workforce strategy	85%
Plan to reskill workers to work alongside automation	77%
Expect to hire for newly created skill sets	70%
Name skills gaps as the top barrier to adoption	63%
Core worker skills expected to be outdated by 2030	39%

Source: World Economic Forum, Future of Jobs Report 2025.

> **The adjustment gap.** The supply side trains models; the distribution side must retrain people, who retrain far more slowly than models improve. The skills response is where the two cadences meet most directly: capability advancing on product cycles against a workforce rebuilt only on the cadence of education and careers. Where retraining is left to fund itself, the gap it is meant to close instead widens.

OUTLOOK · WHAT THE NEXT YEAR WILL DECIDE

What the next year will decide

The developments ahead resolve the central question in one direction or the other. The supply targets are ambitious and nearly certain. The distribution targets are fewer, later, and conditional.

The supply side will almost certainly advance

On the supply of automation the near term targets are specific and the direction is established. General reasoning benchmarks are converging toward human level performance and, where they saturate, are being replaced by harder successors rather than abandoned, indicating that capability has outrun the tests. Humanoid deployments are positioned to rise from the order of one thousand units toward ten thousand. Thresholds of this kind, once crossed, are not reversed; capability does not regress and cost does not rise. The supply side can be forecast with confidence.

For that reason the uncertainty does not lie on the supply side. To state that capability will advance over the next year is nearly tautological: the model that exists does not disappear, the price that fell does not return, and the robot that works does not cease working. The uncertain variable is on the distribution side, where little is settled and little has been built.

The distribution side is the open question

On distribution the targets are fewer, later, and conditional: a first national automation tax, automation dividend legislation, median wages rising to capture a share of productivity, and a first government funded basic income pilot at scale. Each is a proposal rather than a law, and each could advance or stall in the coming year. The central claim is arithmetic: if the supply targets are met and these are not, the distance between what economies can automate and what they share widens by arithmetic, without a deliberate decision.

FIGURE O.1 CAPABILITY AND DISTRIBUTION ADVANCE AT DIFFERENT CADENCES

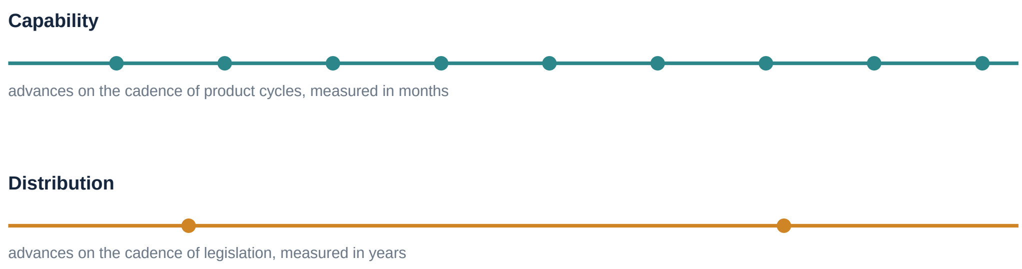

Schematic. The difference in cadence is the mechanism by which output can outpace its distribution.

The development to watch through the rest of 2026 is therefore not whether capability advances, which is nearly certain, but whether the distribution side moves to meet it. That is the test the analysis sets. The supply targets are likely to be met. The question that determines the outcome is whether even one of the four distribution proposals becomes law at scale, because that, rather than the next model, determines whether the gains of the transition are broadly or narrowly shared.

The conclusion returns to the opening question. The binding constraint on optional work is no longer the technical capacity to automate. It is the distribution of what automation produces, and that distribution has barely begun.

What would change the verdict

The central claim is conditional and therefore testable. The widening it describes follows only if the supply side advances while the distribution side does not. Two kinds of evidence would alter the conclusion, and both can be monitored over the year ahead.

The first would narrow the gap from the supply side: a plateau in capability or a halt in the fall of cost, either of which would slow the substitution pressure treated here as nearly certain. The current trajectory does not point that way, but a model generation that fails to improve, or compute that ceases to fall in cost, would weaken the case directly. Neither is expected, but either would constitute the clearest refutation.

The second, and more likely, would close the gap from the distribution side: a first automation levy enacted at scale, a funded basic income program that outlasts a pilot, or a wage settlement that lets median pay capture a measurable share of the productivity gain. Any one of these, enacted rather than proposed, would indicate that distribution has begun to move. The purpose of the analysis is not to predict which way the year resolves but to specify the question precisely enough that the outcome, when it arrives, is unambiguous.

> **The test.** Whether a single distribution mechanism becomes law at scale before the next model release. If it does, the gap can narrow. If the model release comes first and alone, as in every quarter to date, the gap widens in the absence of any decision to the contrary.

METHOD AND SOURCES

How this report is built

Figures are drawn from named public sources, among them the World Economic Forum, the International Monetary Fund, the International Federation of Robotics, the United States Bureau of Labor Statistics, Goldman Sachs, Menlo Ventures, the European Commission, and GiveDirectly. Country standings are taken from the Optional Work Index and the National AI Compute Index, whose methodology, including the scoring functions and thresholds, is set out at optionalwork.com/research. Figures are current as of June 2026.

The distinction between measured and projected

Measured values and projections are kept distinct. Installation counts, spending totals, market shares, and study outcomes are measured. Job creation and displacement to 2030, productivity at full adoption, and deployment trajectories are projections, labeled as such wherever they appear. Measured and projected figures are not averaged, and estimates from different institutions are not combined into a single number.

The standard for inclusion

A development is included only where it can be attributed to a named, public, and checkable source. Anecdote, private briefing, and unsourced claim are excluded. Where a source reports uncertainty, that uncertainty is carried forward. Where two credible sources conflict, both are presented.

What the report does not do

The report does not forecast a single figure for displacement, recommend a particular policy, or claim to know the outcome of the transition. It sets out the evidence on each side and the near term markers whose arrival or non arrival will resolve the central claim.

APPENDIX A

The two indices, defined

The Optional Work Index

The Optional Work Index scores how optional paid work has become in an economy, on a 0 to 100 scale, across 150 economies. Its defining property is that it combines two things and rewards only their combination. One component measures the supply of automation: the capability and diffusion of systems that can perform paid work. The other measures the distribution of its proceeds: the means by which the gains reach the population rather than pooling at the top. The index reports the geometric mean of the two, so that a high score requires both. An economy that automates heavily but shares little, or shares well but automates little, cannot lead. Only an economy strong on both axes can.

This construction is the reason the index reads the way Chapter Six describes. The leaders hold their place by pairing capacity with distribution, not by automating the most. The full set of scoring functions, weights, and thresholds is published in the technical white paper. The values used in this report are the global mean of 22.7 and the highest scoring economy, Singapore, at 35.2.

The National AI Compute Index

The National AI Compute Index measures effective access to AI compute, on a 0 to 100 scale, across 150 economies. Its central idea is that raw hardware is not the right unit, because owning chips is not the same as being able to run and power them at scale. The index therefore conditions hardware on affordability and energy, producing a measure of compute a country can actually use rather than compute it nominally possesses. The result is a highly concentrated distribution, with a global mean of 23, a small group of strong performers near the top of the scale, and a long tail reading near zero.

The two indices are designed to be read together. The compute index measures the capacity to produce automation's gains. The optional work index measures the extent to which those gains have been turned into a population's freedom from paid work. The first is concentrated; the second is low everywhere and led by economies that distribute. The distance between them is the report's subject. Full methodology for both indices, including every scoring function and threshold, is set out at optionalwork.com/research.

APPENDIX B

Data behind the figures

Figure 3.1 and the labor projections

Series	Value	Type
New roles by 2030	170 million	Projection
Displaced by 2030	92 million	Projection
Net change	+78 million	Projection
Structural churn share	22%	Projection

Figure 4.1, enterprise generative AI spending

Year	Spending	Type
2023	$2.3 billion	Baseline
2024	$11.5 billion	Surveyed
2025	$37 billion	Surveyed

Figure 6.1, effective compute access, leading economies

Economy	Score	Economy	Score
United States	100	United Arab Emirates	80
Japan	97	Israel	80
Australia	97	Norway	79
United Kingdom	96	Saudi Arabia	78
Canada	93	Brazil	72
South Korea	93	New Zealand	71
Switzerland	82	China	70

Global mean across 150 economies is 23. Source: National AI Compute Index, June 2026.

Indices summary

Index	Global mean	Leader
National AI Compute Index	23	United States, 100
Optional Work Index	22.7	Singapore, 35.2

APPENDIX C

Glossary of terms

The terms used throughout this report, defined as they are used here.

Optional work

The condition in which paid work becomes a choice rather than a necessity, because the output an economy produces can support its people whether or not they sell their labor for it. The premise of this report is that the constraint on optional work is no longer what economies can automate but what they choose to share.

Optional Work Index (OWI)

A composite index scoring how optional paid work has become across 150 economies on a 0 to 100 scale. It is built as the geometric mean of a supply measure, the capability and cost of automation, and a distribution measure, the means to spread its gains, so that a high score requires both. The global mean reads 22.7, with Singapore highest at 35.2.

National AI Compute Index (NCI)

A standalone index measuring effective access to AI compute across 150 economies on a 0 to 100 scale, conditioning raw hardware on whether a country can afford to run and power it. The global mean reads 23, with the United States set at the reference value of 100.

Augmentation

The use of a tool to raise the output of a worker who stays in the loop. Augmentation accelerates the worker within a task rather than removing the worker from it, and it tends to dominate where labor is scarce.

Substitution and displacement

The replacement of a worker's task by an automated system that performs it without the worker. Displacement removes the worker from the loop rather than speeding the worker within it, and it tends to dominate where labor is abundant and routine.

Autonomy

The capacity of an AI system to plan and execute multi step work unsupervised, using software and operating computers end to end, rather than completing single tasks a worker assigns. Autonomy is the capability threshold that this report treats as the line between augmentation and displacement.

Churn

The gross movement of jobs created and destroyed beneath a net total. A net gain of 78 million roles conceals a churn of roughly 22 percent of all jobs structurally transformed by 2030, which is why the net figure understates the disruption.

Effective compute access

A measure of usable AI compute that conditions raw hardware on affordability and power, on the principle that capacity a country cannot afford to run is not capacity it can use. It is the basis of the National AI Compute Index.

Frontier model

The most capable general purpose AI systems available at a given time. The report's labor relevant claim is that frontier grade performance now sells at mid tier prices, which is what turns a benchmark result into a hiring decision.

General purpose robot

A robot, often humanoid in form, built to perform a range of unstructured physical tasks rather than a single fixed operation. The report treats the humanoid threshold as the point at which physical automation reaches the base of the labor market rather than a slice of it.

SWE-bench

A software engineering benchmark, introduced in 2023, that tests whether AI systems can resolve real code issues. Its verified subset was effectively saturated near 80 percent and has since been succeeded by a harder, contamination resistant successor, itself a sign that capability has outrun the tests.

Skills obsolescence

The rate at which the skills a role requires become outdated. Roughly 39 percent of core worker skills are expected to be outdated by 2030, the mechanism that converts a healthy net employment figure into individual hardship absent retraining.

Vertical and horizontal AI

Vertical AI is built for a specific industry such as healthcare or legal work; horizontal AI serves general functions across industries. Enterprise spending divides across both, with vertical applications led by healthcare.

Automation tax

A proposed levy on automated production or on the firms that deploy it, intended to fund the distribution of automation's gains. As of this report it remains proposal rather than law in every economy.

Automation dividend

A proposed transfer that returns a share of automation's productivity gains to the public directly. Like the automation tax, it is concrete but unproven, with no example yet enacted at scale.

Basic income

A regular cash transfer paid to individuals regardless of work status. The longest running evidence cited here is the GiveDirectly study in Kenya, where recipients invested more and worked no less, though no government funded pilot yet exists at national scale.

REFERENCES

Sources

European Commission. Regulation (EU) 2024/1689 of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). In force 1 August 2024, applying in phases through 2027. eur-lex.europa.eu/eli/reg/2024/1689/oj

GiveDirectly. Universal Basic Income: the world's largest and longest study, Kenya. Program ongoing since 2017; twelve year transfers of about $0.75 per day to roughly 23,000 people across 195 villages, with an independent evaluation by Banerjee, Niehaus, Suri, and colleagues. givedirectly.org/ubi

Goldman Sachs. Briggs, Joseph, and Devesh Kodnani. The Potentially Large Effects of Artificial Intelligence on Economic Growth. Global Economics Analyst, March 2023. goldmansachs.com/insights

International Federation of Robotics. World Robotics 2025: Industrial Robots. Frankfurt, September 2025. ifr.org/worldrobotics/report-2025

International Monetary Fund. Cazzaniga, Mauro, and others. Gen-AI: Artificial Intelligence and the Future of Work. Staff Discussion Note SDN/2024/001, January 2024. imf.org/en/publications/staff-discussion-notes/issues/2024/01/14/gen-ai-artificial-intelligence-and-the-future-of-work-542379

Menlo Ventures. 2025: The State of Generative AI in the Enterprise. December 2025. menlovc.com/perspective/2025-the-state-of-generative-ai-in-the-enterprise

Princeton University and collaborators. Jimenez, Carlos E., and others. SWE-bench: Can Language Models Resolve Real-World GitHub Issues? 2023 (arXiv:2310.06770); SWE-bench Verified and a contamination resistant successor, 2026. swebench.com

United States Bureau of Labor Statistics. Job Openings and Labor Turnover Survey (JOLTS). bls.gov/jlt

World Economic Forum. The Future of Jobs Report 2025. Geneva, January 2025. weforum.org/publications/the-future-of-jobs-report-2025

Optional Work Index and National AI Compute Index. June 2026. Methodology in the technical white paper at optionalwork.com/research

Optional Work Report. A monthly assessment from optionalwork.com. June 2026, Comprehensive Edition. Published June, 2026.