

PRE-READING

ECONOMIC SCENARIOS FOR AI & THE FUTURE OF WORK

How can France prepare for the economic, fiscal,
and labour-market consequences of AI, and lead
the European conversation on the future of work?

17 SEP 2026

Collège de France — rue d'Ulm site, Paris

CONVENED BY

Philippe Aghion — Farhi Innovation Lab, Collège de France

Eric Hazan — Ardabelle Capital

In partnership with the Windfall Trust

Conducted in English · held under the Chatham House Rule

An anonymised synthesis will be circulated after the event

CONTENTS

Foreword

a day to establish where we actually disagree

Programme of the day

PART ONE — The general impact of AI on work

1. The fog

the three uncertainties, and what fixes each

2. Productivity: a promise beginning to materialise

the experiments, the macro signal, the funnel challenge, diffusion, and why organisations must change

3. The physical frontier

why robots may matter more to Europe than chatbots

4. Politics: the Luddite risk

opinion, energy, and what the labs are proposing

PART TWO — How work itself is changing

5. Jobs: from exposure to bundles

transition framework, weak and tight bundles, demand elasticity, agents

6. The divide inside the workforce

who actually uses AI — and whether the gap closes

7. Inside the firm: organisation, managers, meaning

absorption, middle management, three ways of working with AI, control

8. The broken first rung

apprenticeship, and what would have to be built

9. Where the pressure lands: a sectoral reading

software, professional services, finance, health, industry, the state

PART THREE — What to do about it

10. France in comparative perspective

Draghi on compute and adoption, structural differences, the specific vulnerability

11. Transforming the firm: what actually works

10-20-70, process redesign, the workforce, six tests

12. A preparedness agenda: five workstreams

with the objection to each

ANNEX

AI and productivity: rethinking workforce training and social cohesion

Philippe Tibi & Eric Hazan, Polytechnique Insights, April 2026

Exhibits to the annex

four exhibits

A day to establish where we actually disagree

In late August 2026, Bill Gates changed his mind in public. Three years earlier he had written that the coming disruption to work would be bumpy but manageable, and that AI would mostly help people work more efficiently. In an essay and a series of interviews he now argues the opposite: that without intervention there will be far fewer jobs than exist today, that the technology has crossed dangerous thresholds in biology, cybersecurity and white-collar work, and that “on the current course and speed, there's a very high chance of a net negative outcome”¹. Two elements of his argument matter for this seminar more than the headline. The first is why he thinks the historical analogies fail: earlier technological transitions unfolded over generations and ended by creating new work that required human cognition, whereas AI can substitute for that cognition directly. The second is his verdict on the response. He describes the plan from governments and technology leaders as essentially non-existent, criticizes both American parties for being neither knowledgeable nor thoughtful about it, and notes that the AI companies themselves are more worried in private than their public statements suggest². One does not have to accept his forecast to take the second point seriously. The evidence in this document does not support a prediction of mass job destruction, and several careful economists find no aggregate displacement at all. But the observation that there is no plan is not a forecast — it is a description of the present, and it is the reason forty people are being asked to spend a day on this. France will hold a presidential election in 2027. Between now and then, the country will have to form a view on a question it has so far discussed mostly in fragments: what artificial intelligence is doing to work, to the distribution of income, and to the institutions built to manage both. The fragments are not the problem. The problem is that they point in different directions. Frontier laboratories publish research showing that a fifth of occupations face near-term automation pressure. Central banks find half a million fewer software developers employed than pre-2023 trends implied. Other equally careful economists find no significant divergence at all between the most and least exposed occupations. Employers announce redundancies attributed to AI while their own data show something more ambiguous. And the public, largely uninterested in this methodological argument, has already concluded that something is happening to it. This is the situation the seminar of 17 September is designed to address. Not by resolving the empirical dispute — that will take years — but by putting the people who hold the different pieces of it in the same room for a day, and asking what a country would do if it took the uncertainty seriously.

Why this group, and why now

The room brings together four populations that rarely meet in this format. The economists whose work defines how AI's effect on productivity and employment is actually measured. The founders and operators building the systems, who know what is technically close and what is further away than the discourse suggests. The business and HR leaders translating all of this into decisions inside companies. And some political figures who have a view on the topic. Each of these groups holds information the others need. Economists know what the evidence supports and, more importantly, what it does not. Technologists know the trajectory of capability and cost. Employers know what is actually happening inside organisations, which is generally slower and messier than

¹B. Gates, interviews and essay of late August 2026, as reported by Semafor (25 August 2026), Axios (26 August 2026) and CNN.

²Ibid.

either the optimistic or the alarmed account. And political leaders know the constraint that determines whether any of this becomes policy. The timing is not incidental. The evidence of labour-market disruption will remain contested for several more years. Public opinion has already moved. Around seventy per cent of Americans expect AI to reduce employment; barely a quarter of Britons think it is positive for humanity. Policy will be made in response to perception well before it can be made in response to measurement. That gap — between what we can prove and what we will be asked to answer for — is where this seminar sits.

What we are trying to produce

The day is organised around a single question: how can France increase its preparedness for the economic, fiscal, labour-market and political consequences of transformative AI? It moves through three stages — the premise (what should we plausibly prepare for?), the scenarios (what would different trajectories mean concretely for France?), and the response (what institutional capacity, policy choices and political leadership would be required?). The ambition is deliberately bounded. A single day cannot produce a programme, and any document claiming otherwise would be a marketing exercise. What a day can produce is a shared map: agreement on what the evidence actually says, clarity on which disagreements are empirical and resolvable by better measurement, and identification of the two or three choices that cannot be delegated to the passage of time. The format follows from that ambition. Around forty participants rather than an audience. Two structured breakout discussions rather than a succession of speeches. Panels built around disagreement rather than consensus. And the Chatham House Rule throughout, so that people can say what they actually think about their own sector, their own firm and their own party. A short, anonymised synthesis will be circulated afterwards. It will record areas of convergence and, more usefully, the questions that remained open. No view will be attributed to any participant.

How to read this document

What follows is a background paper in three parts. The first examines the general impact of AI on work — what the productivity evidence shows, why it is so hard to measure, where the physical frontier may matter more than the cognitive one, and how public opinion is moving. The second looks at how work itself is changing: which jobs, which tasks, which people, inside which organizations, and across which sectors. The third turns to what can be done — France's specific position, what the evidence says about transforming a firm, and a preparedness agenda in five workstreams, each presented with the serious objection against it. The paper takes positions. It argues that organizational absorption, not model capability, is the binding constraint over the next three years; that the erosion of the entry-level rung is the most consequential and least visible risk; and that the distributional damage will arrive before the aggregate damage, through channels current statistics do not monitor. Participants who disagree are precisely the ones the day needs to hear from. A column published in Polytechnique Insights in April 2026, which sets out much of the underlying argument at greater length, is reproduced in the annex together with its four exhibits.

Eric Hazan

Ardabelle Capital

PROGRAMME

Programme of the day

Thursday 17 September 2026 · Collège de France, rue d'Ulm site, Paris. The session is conducted in English and held under the Chatham House Rule, except for the roundtable.

8:30 Welcome coffee

9:15 Opening remarks and framing

Eric Hazan, Ardabelle Capital and The Windfall Trust

9:30 Opening keynote

Yann LeCun — Founder, AMI Labs · Turing Award laureate

10:15 Scenario briefing

Presented by the Windfall Trust

10:30 Breakout discussion I — “If this scenario materialised, what would it imply?”

5 min introductions · 35 min discussion · 10 min plenary share

11:30 Break

11:50 Breakout discussion II — “What would we do about it?”

5 min introductions · 35 min discussion · 10 min plenary share

13:00 Lunch — buffet, and informal exchange

14:15 Address by Philippe Aghion

Professor, Collège de France & INSEAD · Nobel laureate in Economics

15:00 Panel 1 — AI, productivity and the future of work

With Antonin Bergeaud (HEC), Antoine Bozio (IPP), Fabien Curto Millet (Chief Economist, Google), Marine Boudot (Wonderful) and Charles Gorintin (Alan). Moderated by Eric Hazan.

16:00 Break

16:15 Panel 2 — AI and the transformation of work

Roles, organisations and ways of working. With Alex Richmond (OpenAI), Philippe Corrot (Mirakl), Julia Hobsbawm (author, The Nowhere Office), Gautier Cloix (H Company) and Xavier Jaravel (LSE). Moderated by David Djaïz.

17:15 Break

17:30 Closing political roundtable

A discussion on AI and the future of work between two to three candidates to the 2027 French presidential election.

18:30 Closing remarks

Preparing France for Transformative AI

What the evidence now shows on productivity, jobs, skills and the social contract — and what would have to be decided

Prepared for the working session of 17 September 2026. Sources are given in the footnotes throughout. The column published in Polytechnique Insights in April 2026 is reproduced in the annex.

KEY TAKEAWAYS

- The productivity gain is real at the level of the task, but it shrinks at every human bottleneck between the task and the finished product — and it has not yet appeared in aggregate statistics.
 - Exposure predicts almost nothing. What matters is whether a job is a loose bundle of separable tasks or a tightly enmeshed one — which is why AI erodes the junior version of a job while enhancing the senior version.
 - Adoption is itself unequal: the best-paid and most experienced workers are using AI several times more than everyone else, which turns a productivity tool into an inequality machine.
 - The most serious structural risk is not mass unemployment but the collapse of apprenticeship — the mechanism through which every profession has produced its own experts.
 - None of the gains are automatic. They require reorganizing firms at scale rather than adopting a tool, and a preparedness agenda — measurement, the entry-level rung, adoption, distribution, and the direction of the technology itself — whose hardest choices cannot be delegated to firms or to the passage of time.
-

The general impact of AI on work

What we know, what we cannot yet know, and why the answer keeps moving

1. The fog

Something unusual is happening to the debate about artificial intelligence and work: the more evidence we accumulate, the less certain the conclusion becomes. Three years into the generative era and roughly one into the agentic one, we have thousands of studies, dozens of dashboards, real-time trackers built on payroll and job-posting data — and no settled answer to the simplest question anyone asks, which is whether this technology is destroying jobs.

That is not a failure of economics. It is a description of where we are. Ben Casselman, who covers the economy for the New York Times, puts the problem bluntly: our statistical infrastructure was not designed to capture a change moving this fast. The monthly jobs report does not break out technology as an industry — the categories were fixed decades ago, and tech sits scattered across information, professional services and manufacturing. We do not reliably track what happens to recent graduates month by month. Private data has rushed to fill the gap, but ADP, LinkedIn, Indeed and the AI labs themselves are producing findings that point in opposite directions³.

Two things can nonetheless be said with confidence. The first is that the technology and its adoption are both moving very quickly, and that a statement signed by around two hundred economists in July 2026 warned this could be a transformation larger than the Industrial Revolution unfolding over a vastly shorter period. The second is that whatever is happening so far is subtle. If AI were already destroying employment at scale, imperfect data would still show it at the macro level.⁴

Whatever effect AI is having on employment today, it is subtle. That is not reassurance — it is the definition of the window in which preparation is still possible.

This document is an attempt to hold those two facts together. It is written for a room of economists, technologists, employers and political leaders who will spend a day at the Collège de France asking what France should do. It takes a position: that the current absence of visible displacement is not evidence of safety but the last moment in which choices can be made cheaply. And it tries to be honest about which claims rest on solid evidence, which rest on contested evidence, and which are simply assertions repeated until they sounded like findings.

³Ben Casselman, interviewed in “What A.I. Is Actually Doing to the Economy”, The Daily, The New York Times, 27 July 2026.

⁴The qualification matters. The claim concerns aggregate employment statistics. At sector and cohort level the data have already moved — see chapters 5 and 8 on entry-level hiring and on the composition of recruitment.

The case for taking a genuine take-off seriously

Because this document spends a good deal of time on why the alarmed accounts are not yet supported by aggregate data, it is worth being explicit about the opposite risk: that the caution becomes its own form of complacency. Three considerations argue for treating a real take-off, and a real disruption of work, as a live possibility rather than a tail scenario.

The first is the nature of what is being automated. Generative AI has captured the world's imagination precisely because it appears likely to automate tasks that previously required advanced cognitive skills — which raises the genuine prospect that many highly educated and experienced workers are replaced by algorithms. The question is no longer what happens when machines come for the jobs of handloom weavers and autoworkers, but what happens when they come for the jobs of scriptwriters, lawyers, middle managers and even senior executives⁵. No previous automation wave began at that end of the skill distribution.

The second is that the absence of a signal in aggregate data is exactly what the early phase of a genuine transformation looks like. The J-curve, the lag between capability and reorganisation, the measurement categories fixed decades ago: each is a reason why a large effect would remain invisible for some time. Consistency across micro studies, a steep adoption curve and a capability overhang of the size documented in chapter 5 are not the signature of a technology that fizzles.

The third is that the capability trajectory has not slowed. The length of task that an agent can complete reliably has moved from minutes to hours within a year; extrapolating that horizon forward, even conservatively, implies systems handling work that currently takes a person days. If that holds even approximately, the bundling logic that today protects senior and cross-functional roles becomes considerably less protective.

The absence of a signal in the aggregate data is exactly what the early phase of a genuine transformation looks like. Caution about the evidence is not the same as confidence that little will happen.

A framework published in September 2026 puts numbers on exactly this question, and it is worth describing because the morning of this seminar is built the same way. Korinek, Jones and their co-authors take a handful of plain questions — how much of today's work AI can do, how widely it is actually used, how much time it saves, and whether it works alongside people or instead of them — and translate the answers into paths for growth, wages and unemployment out to 2030⁶.

Three settings give three very different worlds. In the first, AI adds about 1.6% to GDP by 2030 and barely moves unemployment: a useful technology, absorbed without difficulty. In the second, it adds 8.3%, the share of national income paid out as wages falls from 60% to 56%, and unemployment among office and professional workers rises by half. In the third, it adds 32%, growth reaches 15% a year, the wage share falls to 45%, and nearly one in five of those workers is out of work⁷.

Two of their observations bear directly on the argument above. The first is that today's muted data are equally consistent with the first world and with the very beginning of the third — and that the two call for opposite

⁵S. Johnson and E. Hazan, "Human Capital for the Age of Generative AI", Project Syndicate, 30 May 2024.

⁶A. Korinek, C. I. Jones, S. Sacher, T. Cotter and P. McCrory, "Economic Scenarios for Transformative AI", The Anthropic Institute Working Paper No. 2026-02, September 2026. The accompanying scenario explorer is at anthropic.com/institute/econ-scenarios.

⁷Ibid., Table 3.

responses. The second is a matter of timing: the three paths barely separate before 2027. The evidence that would settle the question arrives after the decisions have to be taken⁸.

The three paths barely separate before 2027. The evidence that would settle the argument arrives after the decisions have to be taken.

They also asked nearly 11,000 Americans what they expect — when AI will be able to do each of eight tasks, how widely it will be used, how much time it will save. Views vary enormously, but the median answers land close to the middle scenario: a capable technology, adopted unevenly, with real but manageable disruption⁹. Their online explorer allows anyone to set the assumptions and see what follows, which makes it a useful instrument for the small-group discussions.

The conclusion we draw is not that disruption is certain, but that shared prosperity will not emerge from this technology by default. It can flow from it only if adoption is accompanied by upgraded skills and proactive worker redeployment — which means employers being candid about the skills gaps opening up, and governments focused on enabling every worker to upgrade in a timely and appropriate way¹⁰. That is a conditional statement, and the condition is the subject of Part Three.

Why the uncertainty is structural, not temporary

It is tempting to treat the current disagreement as a transitional problem that better data will resolve. It will not, entirely, and it is worth understanding why. Three distinct kinds of uncertainty are being conflated in most public discussion, and they have different remedies.

The first is measurement uncertainty: we do not know what is happening now. This is genuinely fixable. Statistical agencies could break out technology employment, track graduate outcomes monthly, and publish occupational data at a granularity that matches the speed of the change. Several institutions are already building the instruments — California's workforce-impact monitor, the OECD's real-time indicators, the Direction générale du Trésor's own assessment¹¹. This is the cheapest and most consensual item on any preparedness agenda.

The second is causal uncertainty: even where we see a change, we cannot cleanly attribute it. The hiring slowdown in exposed occupations began before ChatGPT was released, and the rate cycle, post-pandemic over-hiring and the end of cheap money are all plausible competing explanations. Casselman adds a reason for caution about corporate testimony: executives are currently rewarded by investors for attributing job cuts to AI, so announcement data overstates the contribution¹². This uncertainty narrows with time and better identification strategies, but it will never fully close.

The third is forecast uncertainty, and it is irreducible. Nobody knows how far capability will advance or how fast adoption will follow — a property of the object, not a gap in the literature. The right response is not to wait for it to resolve but to find choices that hold across several futures, which is what the morning of this seminar is for.

⁸Ibid.

⁹Ibid., Section 3.5. Survey of 10,980 US adults, August 2026.

¹⁰Ibid.

¹¹Office of the Governor of California, June 2026; OECD Employment Outlook 2026; Direction générale du Trésor, « L'intelligence artificielle : quels effets sur l'emploi ? », 30 June 2026.

¹²B. Casselman, The Daily, The New York Times, 27 July 2026.

2. Productivity: a promise beginning to materialize

Start with the promise, because everything else depends on it. If AI does not raise productivity materially, the distributional argument becomes a fight over a stagnant pie and the political stakes fall. If it does, then almost every institution built for the twentieth-century labour market comes into question.

A promise beginning to materialize

The optimistic case no longer rests on anticipation. Two independent efforts to measure what is actually happening inside firms now point the same way. Torsten Slok's team at Apollo Global Management documented concrete gains in DevOps, process automation and support functions in January 2026 — measured improvements at adopting firms rather than projections¹³. And the Early Signals of AI Impact dashboard, which tracks the question continuously across eighteen indicators and more than 460 sources, reports the same three things: adoption accelerating in every industry, measured productivity rising at the firms furthest along, and roles built on repetitive tasks being compressed¹⁴.

What the evidence does not yet show is movement in the aggregate. Headline output and economy-wide productivity have not visibly shifted. The column reproduced in the annex draws the conclusion this document adopts: work is changing faster than it is disappearing, and the interval between those two facts is the window in which preparation is still cheap.

That interval should not be mistaken for safety. A technology whose effects at task level are this consistent, and whose adoption curve is this steep, does not usually leave the aggregate untouched for long.

Adoption is accelerating in every industry, measured productivity is rising where firms have gone furthest, and the aggregate has not moved. All three are true at once — for now.

What the experiments actually show

The task-level evidence is now robust, consistent and in places remarkable — and it points in one direction. The most rigorous study to date is Brynjolfsson, Li and Raymond's *Generative AI at Work*, which followed the staggered introduction of a conversational assistant across 5,172 customer-support agents at a Fortune 500 software firm. Access to AI raised productivity, measured as issues resolved per hour, by 15% on average¹⁵.

The distribution of that gain matters more than its size. Less experienced and lower-skilled agents improved by roughly a third — around 34% in the working-paper estimates — while the most experienced saw small gains in speed and small declines in quality. The authors find suggestive evidence that the model disseminates the tacit knowledge of the best performers and moves newer workers down the experience curve, alongside improvements in customer sentiment and employee retention¹⁶.

¹³T. Slok, R. Shah and S. Galwankar, *Quantifying the Productivity Gains from AI Adoption*, Apollo Global Management, January 2026.

¹⁴Early Signals of AI Impact, real-time dashboard (jobsdata.ai), March 2026. Aggregates 460+ sources across 18 indicators.

¹⁵E. Brynjolfsson, D. Li and L. Raymond, “Generative AI at Work”, *The Quarterly Journal of Economics*, vol. 140, no. 2, May 2025, pp. 889–942. The published version reports 15% on a sample of 5,172 agents; the earlier NBER working paper (no. 31161, 2023) reported 14% on 5,179 agents.

¹⁶*Ibid.* The published paper notes small quality declines among the highest-skilled agents.

A second experiment, run with Boston Consulting Group by researchers from Harvard, MIT Sloan and Wharton, tested 758 consultants across eighteen realistic tasks. Inside AI's capability frontier, consultants using GPT-4 completed 12.2% more tasks, worked 25.1% faster and produced output rated more than 40% higher in quality. Outside it, on a complex managerial task chosen to sit beyond the model's reach, they were 19% less likely to reach the correct answer than colleagues working without AI — the finding that gave us the phrase “jagged technological frontier”¹⁷.

Noy and Zhang, working with professional writing tasks, found time-to-completion falling by around 40% and output quality rising by about 18%¹⁸. Taken together the literature converges on task-level gains in the range of 14 to 40%, concentrated among less-experienced workers, with real quality improvements inside the frontier and real degradation beyond it¹⁹.

AI compresses the skill distribution: it lifts the floor far more than the ceiling. That is the most consequential finding in the whole experimental literature — and it cuts both ways.

That compression is the single most policy-relevant result in the micro literature. A technology that makes novices substantially more productive and experts only marginally so is, on its face, an equalizing technology. It is also, as chapters 6 and 9 argue, a technology that erodes the wage premium attached to experience and the roles through which experience was acquired. The same finding supports the optimistic and the anxious reading; which one materializes depends on choices made inside firms and by governments.

The disagreement between these estimates is not arithmetic; it is about what is being counted. Acemoglu's figure captures the automation of existing goods and services. Aghion's captures something additional: the production of ideas. His argument is not that AI will automate more tasks, but that it will change the rate at which ideas are produced — intelligence as a recombination engine, generating new ideas from existing ones at a pace no research organization could previously sustain. Which of the two framings is right determines whether the appropriate policy posture is managing an adjustment or preparing for an acceleration.

¹⁷ F. Dell'Acqua, E. McFowland III, E. R. Mollick, H. Lifshitz-Assaf, K. Kellogg, S. Rajendran, L. Kraymer, F. Candelon and K. R. Lakhani, “Navigating the Jagged Technological Frontier”, *Organization Science*, 2026 (first circulated as Harvard Business School Working Paper 24-013, 2023).

¹⁸ S. Noy and W. Zhang, “Experimental evidence on the productivity effects of generative artificial intelligence”, *Science*, vol. 381, 2023.

¹⁹ For a consolidated review of this evidence and of the paradox it creates, see É. Hazan, *Does AI Really Boost Productivity? The Paradox, the Evidence, and What Policy-Makers Should Do About It*, Policy Brief No. 2026/02, IIAS · ENS, May 2026.

BOX · Philippe Aghion's work on AI, growth and employment

Because his estimate anchors one end of this debate, and because he co-convenes the day, it is worth setting out briefly where it comes from.

The frame is creative destruction. In the Schumpeterian growth models he built with Peter Howitt, growth comes from innovation and each innovation makes the previous one obsolete. Two things follow. Growth is not a smooth process but a process of replacement, which produces losers mechanically alongside aggregate gains. And the rate of innovation depends on institutions — competition, education, finance, social protection — at least as much as on the technology available.

His position on AI follows from a paper with Benjamin Jones and Charles Jones: AI does not only automate the production of goods and services, it automates the production of ideas. If research itself becomes more productive, the rate of innovation accelerates and the compounding dwarfs ordinary automation. The same paper contains its own counterweight, a Baumol-type argument that growth is limited not by what improves fastest but by what improves slowest — the remaining bottlenecks decide²⁰.

The empirical work, conducted on French data with Céline Antonin, Simon Bunel and Xavier Jaravel, has consistently found no net destruction of employment. Firms adopting AI between 2017 and 2020 increased both employment and sales: the productivity gains let them grow, and that growth outweighs displacement. A study of 120 million French job postings since 2019 finds diffusion highly uneven — concentrated in IT, finance and consulting, in large pioneer firms, and around research clusters — with hiring slowing in the most exposed occupations while total employment rises inside the firms that recruit AI skills²¹. The effect operates at the level of the firm rather than the occupation, and it redistributes more than it destroys. A further refinement in the 2025 work: what matters is not adoption but use — employment grows for some applications, such as cybersecurity, and falls for others, such as administrative processes.

The policy conclusion is where his research meets his public argument. The dividend is never automatic; it depends on institutions able to absorb the destruction. Hence his support for Danish-style flexicurity, for a bonus-malus on unemployment insurance so that firms which lay off easily contribute more, for competition policy against entrenched rents, and for investment in skills.

His is therefore an optimism conditional on institutional reform, which is not the same as technological optimism — and it is what makes his position useful to a seminar about preparedness. It also leaves an open question worth putting to him: between the 2017 framework, which rests part of the gain on accelerating discovery, and the recent scenario work finding that channel weak to 2030, which mechanism is actually doing the work?

And the macro signal may be starting to appear

Until recently the honest summary was that none of this was visible in aggregate statistics. That is now changing at the margin. The qualification matters, and it reconciles this section with the one above: the signal appears in a growth-accounting decomposition of where productivity growth has come from, not yet in headline GDP. Both statements can therefore stand together — the aggregate picture is no longer blank, but it is not yet obvious to the naked eye.

Growth-accounting work by Bontadini, Corrado, Haskel and Jona-Lasinio treats AI as both a general-purpose technology and an innovation in the method of innovation — a technology that raises the efficiency with which new knowledge itself is produced. Their finding is that AI can already be seen in US productivity statistics.

²⁰P. Aghion, B. Jones and C. Jones, “Artificial Intelligence and Economic Growth”, NBER Working Paper 23928, 2017.

²¹P. Aghion, S. Bunel, X. Jaravel, T. Mikaelson, A. Roulet and J. Sogaard, “How Different Uses of AI Shape Labor Demand: Evidence from France”, AEA Papers and Proceedings, vol. 115, May 2025, pp. 62–67; P. Aghion, A. Bergeaud, S. Bunel and P. Delboué, Diffusion de l’intelligence artificielle en France, July 2025; P. Aghion, C. Antonin, S. Bunel and X. Jaravel, “The Effects of Automation on Labor Demand: A Survey of the Recent Literature”, CEPR DP 16868, 2022.

Software products and software R&D alone contributed 50% of the 2% average growth rate in US nonfarm business labor productivity between 2017 and 2024, and 50% of its 1.2 percentage-point acceleration relative to 2012–2017²².

Their forward projection is equally relevant to this room: up to around 1 percentage point of additional annual labour productivity growth in the United States, but only about 0.3 points in Europe — a gap driven by a smaller ICT sector, less capital deepening and slower organizational adaptation rather than by any difference in the technology available²³.

Philippe Aghion's estimates sit in the same territory. With Simon Bunel, and subsequently with Xavier Jaravel, he puts the aggregate productivity gain at 0.8 to 1.3 percentage points a year over the coming decade — comparable to the acceleration delivered by electricity (about 1.3 points) or by information technology (about 0.8) in their respective eras; in more recent public remarks he has described at least 0.7 points as a floor once diffusion is real²⁴. Crucially, his framework makes the gain conditional on institutional adaptation: countries that reorganize work, invest in skills and reform labor markets capture the dividend; those that do not see AI reinforce existing inequalities instead.

Set against this, Daron Acemoglu's estimate remains the serious low-end anchor: roughly 0.53% of additional total factor productivity over a decade, which is the same claim as the frequently quoted 0.05 to 0.07% a year²⁵, and the Federal Reserve Bank of St Louis calculates that US labor productivity has run only about 1.9 percentage points above its 2015–2019 trend in the two and a half years since ChatGPT's release — an excess of roughly 0.7 points a year²⁶. The distance between 0.53% over a decade and 1.3 points a year is the single widest disagreement in this document, and no participant should leave the day without a view on it.

The challenge: and then the funnel narrows

None of what follows contradicts the evidence above. It qualifies it, and it should be read as the serious challenge to the optimistic case rather than as its refutation. The gains are real at the level of the task; the question is how much of them reaches the level of the product, and how quickly organizations learn to let them through.

The most uncomfortable counter-evidence comes from managers themselves. A large survey of roughly 6,000 chief executives, finance directors and senior executives across the United States, the United Kingdom, Germany and Australia found that more than 80% reported no change in productivity — measured as sales volume per employee — over the preceding three years, even as AI adoption among the firms surveyed rose from 61% to 71% in a single year²⁷. PwC's global CEO survey reached a similar conclusion: only around 30% of chief executives reported any revenue increase attributable to AI, while 22% said their costs had risen²⁸.

²²F. Bontadini, C. Corrado, J. Haskel and C. Jona-Lasinio, "AI as an Innovation in the Method of Innovation: Implications for Productivity Growth", AEA Papers and Proceedings, vol. 116, 2026, pp. 36–40. The contribution is attributed to software products and software R&D specifically, not to AI capital broadly defined.

²³Ibid.

²⁴P. Aghion and S. Bunel (2024); P. Aghion and X. Jaravel (2025), as summarised in É. Hazan, Does AI Really Boost Productivity? The Paradox, the Evidence, and What Policy-Makers Should Do About It, Policy Brief No. 2026/02, IIAS · ENS, May 2026. The 0.7-point figure is from « On ne voit pas encore d'effets négatifs de l'IA sur l'emploi », Les Échos, 22–23 May 2026.

²⁵D. Acemoglu, "The Simple Macroeconomics of AI", NBER Working Paper 32487, 2024.

²⁶Federal Reserve Bank of St Louis, as cited in É. Hazan, Does AI Really Boost Productivity? The Paradox, the Evidence, and What Policy-Makers Should Do About It, Policy Brief No. 2026/02, IIAS · ENS, May 2026.

²⁷NBER executive survey, February 2026, as reported in É. Hazan, Does AI Really Boost Productivity?, Policy Brief No. 2026/02, IIAS · ENS, May 2026. The brief's annex records a figure of 89% for the same survey; the lower figure is used here.

²⁸PwC Global CEO Survey, as cited in É. Hazan, Does AI Really Boost Productivity?, Policy Brief No. 2026/02, IIAS · ENS, May 2026.

Part of the explanation is measurable and mundane. Workday's 2026 study finds that between 37 and 40% of the time ostensibly saved by AI is currently consumed reviewing, correcting and verifying what it produces. AI does not yet automate tasks end to end; it displaces effort from production to supervision. Until organizations redesign workflows to distinguish output they can trust from output they must check, the net gain will remain well below the gross task-level improvement²⁹.

The most important recent research does not contradict those findings. It measures what happens to them downstream — and the answer is that they shrink at every human bottleneck.

Mert Demirer and co-authors at MIT measured software developers at several points in the workflow rather than one. At the top of the funnel the effect is explosive — almost 300% more files created or edited. By the review stage the boost has halved to around 150%. By the time it reaches actual software releases, the thing the company sells, it has shrunk fivefold again to roughly 30%³⁰.

A 30% improvement in a company's core product is a large number by any historical standard. But the shape of the funnel is the useful image: it explains why individual users feel transformed while their organizations report modest gains, and why the loudest claims come from the people furthest from the finished output.

The measurement problem runs deeper still. METR's earlier experiment produced a genuinely disconcerting result: software engineers believed AI was making them 20% faster, while precise measurement showed it was making them 20% slower. In a follow-up survey of 350 technical knowledge workers, METR asked about value rather than speed — how many copies of yourself would your team need to hire to replace you without AI; what fraction of your current value could you produce if the tools disappeared. Estimates fell from around 3x to around 2x, and to 1.6x on the strictest question. The researchers themselves treat these as upper bounds³¹.

Engineers believed AI made them 20% faster. Measured precisely, it made them 20% slower. Self-reported productivity is not a weak signal — it is sometimes an inverted one.

And there is a quality dimension that the enthusiasm has largely ignored. Telemetry from 22,000 developers on the Faros platform summarizes its findings in six words: more code, declining quality, accelerating incidents. Volume rises sharply at the start of the workflow; every subsequent handoff — review, testing, release — takes substantially longer, because each is a point where human attention and judgement determine what happens next. Incidents in production have tripled relative to a low-adoption baseline, and unlike earlier findings, high-performing engineering organizations show the same downstream deterioration as everyone else³².

Europe's constraint is diffusion, not technology

For a French audience the most actionable numbers in this chapter are not the productivity estimates but the adoption ones, because they identify a bottleneck that policy can act on directly. European adoption has accelerated sharply. The share of EU enterprises with ten or more employees using AI technologies rose from 8.7% in 2023 to 13.5% in 2024 and 20.0% in 2025 — a 130% increase in two years, with the OECD reporting a comparable 20.2% on a slightly broader sample. The ICT sector leads at 57.3%, professional and scientific services

²⁹Workday, 2026, as cited in É. Hazan, Does AI Really Boost Productivity?, Policy Brief No. 2026/02, IIAS · ENS, May 2026.

³⁰M. Demirer et al., cited in John Burn-Murdoch, "How much value is AI really creating?", Financial Times, 5 June 2026.

³¹METR, survey of 350 technical knowledge workers, discussed in J. Burn-Murdoch and S. O'Connor, "Are we thinking about AI and productivity all wrong?", Financial Times, 15 May 2026.

³²Faros, report on telemetry from 22,000 developers, cited in Financial Times, 15 May 2026.

at 36.8%³³. The aggregate conceals the problem. By firm size, 52% of large enterprises used AI in 2025 against 30% of medium-sized firms and 17.4% of small ones — a gap of nearly 35 points. By geography, Denmark leads at 42.0%, followed by Finland at 37.8% and Sweden at 35.0%, while Romania stands at 5.2%, Poland at 8.4% and Bulgaria at 8.5% — a tenfold spread inside a single market³⁴.

A 35-point adoption gap between large and small firms, and a tenfold gap between member states. Europe's productivity problem is not access to the technology; it is diffusion.

The consequences show up directly in the projections. The IMF puts medium-term productivity gains for Europe at roughly 1.1% cumulatively over five years — modest. A European Investment Bank study, using US peer adoption rates as an instrument, finds that AI adoption raises EU labour productivity by about 4% on average, driven primarily by capital deepening rather than by job displacement³⁵. The implication is uncomfortable and useful: much of the gap between European and American outcomes is a gap in adoption rather than in capability, which means it is addressable by the kind of measures set out in Part Three.

The Solow paradox, and why the analogy is precise

None of this means the gains are illusory. It means we are early. Demirer's own interpretation is that current organizational structures and marketplaces are not built to capture the underlying improvement — and the historical parallel he reaches for is exact. When factories first electrified, they replaced a giant steam engine with a giant electric motor and changed nothing else; productivity barely moved. The boom came decades later, when engineers fitted individual workstations with their own small motors and redesigned the factory floor around what electricity made possible³⁶. This is the J-curve that Casselman describes: a technology arrives, everyone fumbles with it, measured productivity may even dip, and then — once work is reorganized around it — the curve turns up sharply. It is also the reason the most striking productivity and revenue growth today is visible not at incumbents grafting AI onto existing workflows, but at the AI-native firms themselves.




Figure 1 — Weighted average cost (USD) per Intelligence-Index task; lower is better. The spread across current models is nearly two orders of magnitude. Source: Artificial Analysis.

³³ Eurostat, use of artificial intelligence technologies in enterprises, 2025; OECD, 2026. Figures as compiled in É. Hazan, Does AI Really Boost Productivity?, Policy Brief No. 2026/02, IIAS · ENS, May 2026.

³⁴ Ibid.

³⁵ IMF Working Paper WP/25/067; European Investment Bank Working Paper 2026/02, as compiled in É. Hazan, Does AI Really Boost Productivity?, Policy Brief No. 2026/02, IIAS · ENS, May 2026.

³⁶ J. Burn-Murdoch, Financial Times, 5 June 2026.

Cost is the other half of the productivity equation, and it moves in two directions at once. The unit cost of model output has collapsed across generations — the “LLMflation” effect³⁷ — while agentic and reasoning workloads consume far more tokens per task. The practical consequence is visible in corporate budgets: Uber's chief executive revealed the company had exhausted its entire 2026 AI budget in a single quarter and was moving much of its usage to cheaper models, reserving frontier tools for special cases³⁸. Substitution economics depend on this number, and it is not yet stable.




Figure 2 — The productivity funnel. The explosive gain at the level of the task survives only partially to the level of the product.

What the funnel means for a chief executive

The Demirer finding has a direct managerial reading that is worth stating without hedging. If your organisation is measuring AI's impact at the top of the funnel — lines of code, drafts produced, tickets handled, documents generated — you are measuring the number that will flatter you most and predict least. The number that matters is the one at the bottom: releases shipped, cases closed, products launched, decisions actually made. The ratio between the two is a measure of your organization's absorptive capacity, and it is probably the single most useful diagnostic any firm in this room could run.

The Faros telemetry sharpens the point. The bottleneck is not the model; it is human attention at each handoff — review, testing, release. Volume rises at the top, and each downstream stage takes substantially longer because more work is arriving at it. Quality falls because reviewers are overwhelmed rather than because generated code is inherently worse. Incidents in production tripled relative to a low-adoption baseline. And crucially, high-performing engineering organizations show the same deterioration, which rules out the comfortable conclusion that this is a maturity problem afflicting laggards³⁹.

The practical implication is uncomfortable for the way most firms are currently deploying capacity at the review and integration stages, not generation capacity, is the binding constraint. Investing further in generation without investing in absorption makes the problem worse, not better. This may be the most immediately actionable finding in this entire document.

³⁷“LLMflation: LLM inference cost”, a16z.

³⁸Cited in J. Burn-Murdoch, Financial Times, 5 June 2026; see also “Microsoft's AI cost problem”, Fortune, 22 May 2026.

³⁹Faros, telemetry from 22,000 developers, cited in Financial Times, 15 May 2026.

The value question nobody is asking

There is a further layer that the METR work exposes and that deserves more attention than it has had. When you ask people how much faster they work with AI, you invite them to count time saved on tasks — including tasks they would never have performed at all before. Building an interactive dashboard to accompany a report, running a more elaborate analysis, producing three variants of a design: these would have taken a long time manually, so the notional time saved is large. But they are additions, not improvements. The hypothetical saving is real; the value added is marginal⁴⁰.

This is why METR reframed the question around value: how many copies of yourself would your team need to hire to replace you without AI; what fraction of your current value you could still produce if the tools vanished. The answers came in materially lower — around 2x on the first framing, 1.6x on the strictest — and the researchers treat even these as upper bounds, noting that when they examined the most striking self-reported gains closely, the work produced was unlikely to be as valuable as claimed⁴¹.

Sarah O'Connor's response to this line of research contains the deeper point. Individual productivity is not a particularly meaningful metric in the first place, not if you work inside a company rather than alone, and certainly not at the macroeconomic level. Organizations and economies are interdependent systems. The value each person generates is not visible in isolation and does not depend only on them⁴². A technology that boosts the productivity of individuals in one part of a company does not necessarily produce more value overall and, as the software evidence shows, can produce less. It needs an ecosystemic change.

Why the gains require changing the organization, not just adopting the tool

If there is one conclusion this chapter should leave, it is that the productivity of AI is not a property of the technology. It is a property of the system the technology is placed in — and the evidence for that proposition is now several decades old.

Brynjolfsson, Rock and Syverson give it its clearest formulation. General-purpose technologies, they argue, both enable and require substantial complementary investment: the co-invention of new processes, new products, new business models and new human capital. Those investments are largely intangible and are poorly captured in national accounts, which produces a predictable sequence — measured productivity is understated in the early years, while firms are paying for reorganization, and overstated later, when the benefits are harvested. They call it the productivity J-curve, and they find it empirically: adjusting US data for intangibles tied to software and computer hardware puts total factor productivity 15.9% above official measures by the end of 2017. Their assessment of AI-related intangible capital is that its effect on measured productivity is, so far, small but growing⁴³.

The firm-level mechanism behind that aggregate pattern was documented long before generative AI. Bresnahan, Brynjolfsson and Hitt showed that the returns to information technology came not from the technology alone but from its combination with workplace reorganization and higher skills — three complements that had to move

⁴⁰ METR study, discussed in J. Burn-Murdoch and S. O'Connor, *Financial Times*, 15 May 2026.

⁴¹ Ibid.

⁴² Ibid.

⁴³ E. Brynjolfsson, D. Rock and C. Syverson, "The Productivity J-Curve: How Intangibles Complement General Purpose Technologies", *American Economic Journal: Macroeconomics*, vol. 13, no. 1, January 2021, pp. 333–372.

together, and whose returns accrued only to firms that moved all three⁴⁴. Paul David's account of electrification makes the same point historically: the dynamo delivered little until factories were rebuilt around it⁴⁵.

The dynamo delivered almost nothing until factories were rebuilt around it. The pattern has not changed: the technology is the cheap part, and the reorganization is where the gains are.

Two further conditions follow, and both are addressed later in this document. The first is scale. Agrawal, Gans and Goldfarb distinguish point solutions, which insert AI into an existing process and yield modest returns, from system solutions, which redesign the surrounding architecture and yield large ones — and note that the second kind requires changes no single team can make, which is why they arrive late and unevenly⁴⁶. Beyond the firm, the same logic applies to the economy: new entrants built around the technology, reallocation of capital and labor towards them, and adjustments in regulation, professional licensing and public services. That is the deeper reason the macro signal lags the micro evidence — and, as chapter 10 argues, the reason Europe's binding constraint is diffusion rather than capability.

The second condition is that the complementary investment has to reach individuals, not only processes. The experimental evidence in this chapter shows AI raising the productivity of novices by roughly a third while adding little for experts — which means the return depends on whether people are equipped to use it well and to judge what it produces. That argues for building systems that raise human capability rather than route around it: what Brynjolfsson calls avoiding the Turing trap, the temptation to aim at human imitation rather than human augmentation because imitation is easier to measure and to sell⁴⁷. Productivity, organization and skills are not three separate agendas. On this evidence they are one, and a country that funds only the first will measure very little.

⁴⁴T. Bresnahan, E. Brynjolfsson and L. Hitt, "Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence", *The Quarterly Journal of Economics*, vol. 117, no. 1, February 2002, pp. 339–376.

⁴⁵P. A. David, "The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox", *American Economic Review*, vol. 80, no. 2, 1990, pp. 355–361.

⁴⁶A. Agrawal, J. Gans and A. Goldfarb, *Power and Prediction: The Disruptive Economics of Artificial Intelligence*, Harvard Business Review Press, 2022; see also their "Artificial Intelligence: The Ambiguous Labor Market Impact of Automating Prediction", *Journal of Economic Perspectives*, vol. 33, no. 2, 2019.

⁴⁷E. Brynjolfsson, "The Turing Trap: The Promise and Peril of Human-Like Artificial Intelligence", *Daedalus*, vol. 151, no. 2, Spring 2022, pp. 272–287; see also D. Autor, "Applying AI to Rebuild Middle Class Jobs", NBER Working Paper 32140, 2024.

3. The physical frontier

Almost everything in the preceding section concerns cognitive work, which is where the attention and the anxiety are concentrated. That may be a mistake of perspective.

Machines fitted with AI can now carry out physical actions using cameras and sensors, guided by world models rather than by text prediction. Daniela Rus, who directs MIT's Computer Science and Artificial Intelligence Laboratory, frames the economics simply: large language models may deliver faster near-term gains, but physical AI has the larger long-run upside because it targets physical bottlenecks — and most working hours are still spent moving atoms rather than bits⁴⁸.

Early deployments are concrete rather than speculative. Foxconn found that vision-guided, self-adjusting robotic arms improved assembly cycle times by up to 30% while cutting error rates by 25%. Amazon's AI-equipped robot fleets learned to navigate around moving humans and obstacles, cutting travel time by 10%. In a 2026 Capgemini survey, more than half of industrial executives cited labor shortages, costs and regulation among their top five reasons for adopting physical AI⁴⁹.

Most working hours are still spent moving atoms, not bits. The largest productivity prize may sit on the factory floor, not in the inbox.

For France and for Europe, this observation carries a strategic implication that deserves airing in the room. The cognitive frontier is dominated by American firms and American capital. The physical frontier — industrial robotics, embodied AI, the redesign of production — plays to European industrial strengths, addresses a labor shortage rather than creating a surplus, and is politically far less explosive than automating white-collar work. China's latest five-year plan already includes a strategic pivot toward embodied AI⁵⁰. There is a plausible argument that Europe's comparative advantage lies here, and that the current fixation on chatbots is an imported framing.

Why the physical frontier may matter more for France than for the United States

The argument of the preceding section has a specific national implication that deserves developing, because it cuts against the grain of most French AI discussion.

The cognitive frontier — foundation models, agents, the productivity of knowledge work — is dominated by American firms, American capital and American compute. France participates as a customer and, through Mistral and a handful of others, as a small competitor. The returns from that frontier accrue overwhelmingly to model developers and infrastructure providers, which the sectoral mapping in Figure 1 shows capturing value across essentially every industry.

The physical frontier looks different. Industrial robotics, embodied AI and the redesign of production play to an existing European industrial base. They address a labor shortage rather than creating a surplus — more than half of industrial executives in a 2026 Capgemini survey cited labor shortages, costs and regulation among their top

⁴⁸ Cited in Tej Parikh, "Robots, not chatbots, will realise AI's potential", Financial Times, 28 June 2026.

⁴⁹ Ibid.

⁵⁰ Ibid.

reasons for adopting physical AI. And they are politically far less explosive: automating tasks nobody wants to do, in sectors that cannot find workers, is a fundamentally different proposition from automating the white-collar careers of graduates⁵¹.

There is also a sovereignty argument. Multitasking machines reduce the need for very large assembly lines, which means production can take place domestically at a lower marginal cost — directly complementing European efforts on supply-chain resilience in critical industries. China has already made this pivot explicit, with embodied AI written into its latest five-year plan⁵².

Automating jobs nobody can fill is a different political proposition from automating the careers of graduates. France's industrial base may be an advantage the current debate is ignoring.

None of this is an argument for neglecting the cognitive frontier, where the productivity of French services and administration is at stake. It is an argument that the balance of French attention is currently inherited from an American debate about knowledge work, and that a country with France's industrial history, energy mix and demographic profile might reasonably weight the two differently.

4. Politics: the Luddite risk

Public opinion is running well ahead of the labour-market data, and that gap is itself a political fact that will shape what is possible.

Around 70% of Americans believe AI will lead to fewer jobs⁵³. In the United Kingdom, a King's College London study finds only 24% of citizens think AI is positive for humanity, and a third of university students believe AI will eliminate jobs fast enough to provoke civil unrest. YouGov finds half of Americans expecting AI to reduce their ability to form meaningful relationships and almost three quarters saying development is moving too fast⁵⁴. A Gallup poll finds a third of Gen Z Americans describing their feelings toward AI as anger — and these are, as Clara Shih observes, precisely the people whose participation the transition requires⁵⁵.

Corporate rhetoric is actively feeding this. Standard Chartered announced the end of 8,000 roles and described it as the replacement of “lower-value human capital” with AI. Block laid off close to half its workforce, its chief executive citing coding agents, and investors responded with a 25% after-hours share-price surge⁵⁶. Casselman's caution is essential here: executives are currently rewarded for attributing cuts to AI, which makes announcement data an unreliable measure of displacement — but announcements shape hiring behavior and political mood long before any measurable displacement occurs⁵⁷.

⁵¹T. Parikh, “Robots, not chatbots, will realise AI's potential”, Financial Times, 28 June 2026.

⁵²Ibid.

⁵³B. Casselman, The Daily, 27 July 2026.

⁵⁴Cited in Camilla Cavendish, “The new Luddite movement”, Financial Times, 23 May 2026.

⁵⁵C. Shih, The New York Times Magazine, 14 June 2026.

⁵⁶Financial Times, 23 May 2026; The New York Times, 30 April 2026.

⁵⁷The Daily, 27 July 2026.

Energy has become the visible surface of the discontent. Data centers now consume around 22% of all electricity in Ireland, up from 5% in 2015, and since August at least twelve US states have debated temporary bans on data-centre construction⁵⁸. Cavendish's warning is that if governments do not moderate the pace, voters may do it themselves — and that the twenty-first-century Luddites will be the middle classes as well as the working classes.

What the polling actually rewards

The Democratic strategist David Shor's polling, presented to an audience of AI employees in San Francisco in February 2026, is worth studying by anyone preparing a political argument in France. It found 79% of voters worried about government not having a plan to protect workers, and 72% concerned that AI drives down wages for people like them. Universal basic income polls badly; a federal jobs guarantee polls well. Voters are indifferent to beating China and enthusiastic about AI curing diseases. And populist framing outperforms: the best-testing advertisement declared that those who profit from AI should pay their fair share, closing with the line “They work for the bots. We work for you.”⁵⁹

79% of voters worry that government has no plan to protect workers. That number is a political opportunity for someone. The question is who, and with what content.

The window this opens is unusually wide. AI appears to expand the Overton window for economic policy that would normally be unsellable — which is precisely why the ideas now circulating deserve serious scrutiny rather than reflexive dismissal or reflexive enthusiasm.

The energy question, and why it is now political

One further front deserves attention because it is where public discontent has become physical rather than rhetorical. Data centers now account for around 22% of all electricity consumed in Ireland, up from 5% in 2015. In the United States, the Center for Biological Diversity estimates that on current trends data centers could account for close to half of all power-sector emissions compatible with national climate targets. Since August, at least twelve US states have debated temporary moratoria on data-center construction⁶⁰.

This matters for a French audience for a specific reason. France's low-carbon electricity is a genuine comparative advantage in an AI build-out — arguably its single clearest one. But the political economy of data centers is turning hostile faster than the industrial opportunity is being captured, and the two are on a collision course. A country that wants to host compute will have to make an explicit public argument for it, including on price and grid access for households, rather than assuming the industrial logic speaks for itself.

What the labs themselves are proposing

It is a striking feature of this moment that the most radical redistributive proposals in circulation are coming from the companies building the technology, and this deserves to be examined rather than either dismissed or celebrated. OpenAI's white paper, “Industrial Policy for the Intelligence Age”, proposes a 32-hour working week, higher taxes on corporations and capital gains, and a public wealth fund giving citizens an equity stake in AI companies. Sam Altman had argued as early as 2021 that sufficiently capable systems would shift power from

⁵⁸Financial Times, 23 May 2026.

⁵⁹D. Shor, presentation of 25 February 2026, reported in The New York Times, 30 April 2026.

⁶⁰Cited in C. Cavendish, “The new Luddite movement”, Financial Times, 23 May 2026.

labor to capital and that the answer was aggressive taxation of assets — land and AI-company shares⁶¹. Anthropic's leadership frames the outcome as a societal choice rather than a technological destiny, with Jack Clark arguing that society could choose to expand the share of human labor in relational roles — teaching, nursing, care — and to pay those roles substantially more⁶². On the other hand, economists like David Autor and Neil Thompson propose universal basic capital: every child receives a portfolio of productive assets at birth, granting a permanent claim on capital income rather than a politically fragile transfer⁶³.

The skepticism this deserves is not about sincerity but about mechanism. As one former OpenAI employee puts it, the industry is engaged in cutthroat competition over world-changing technology; unless the incentives change, we should not be surprised when companies cut corners even when they have said the right things⁶⁴. Neither major lab has yet endorsed specific legislation implementing the redistribution it describes. Proposals without an enforcement mechanism are a description of a preference, not a policy.

The most radical redistribution proposals now come from the firms building the technology. None has yet been attached to a mechanism that would compel it.

⁶¹Cited in J. Sun, The New York Times, 30 April 2026, and in Financial Times, 9 April 2026.

⁶²Ibid.

⁶³D. Autor and N. Thompson, Digitalist Papers, vol. 2.

⁶⁴S. Adler, cited in The New York Times, 30 April 2026.

PART TWO

How work itself is changing

Which jobs, which tasks, which people — and inside which organisations

5. Jobs: from exposure to bundles

The public debate remains organised around a question — will AI destroy jobs? — that the best research has abandoned. Three frameworks have replaced it, and together they explain most of what we observe.

Exposure is the wrong unit

OpenAI's jobs-transition framework, applied to more than 900 occupations covering some 152 million jobs, begins by rejecting exposure as a predictor. It adds three dimensions: human necessity (is a person required for regulatory, relational or physical reasons?), demand elasticity (does cheaper output expand the market?), and realized usage. The resulting map: roughly 18% of jobs at higher short-term automation risk, 24% likely to reorganize as their task mix shifts, 12% that could grow because lower cost expands access, and 46% with no clear near-term direction⁶⁵.

Two of its findings should unsettle both camps. The capability overhang is enormous, realized usage lags theoretical exposure by around 68 percentage points in the most at-risk categories. And the unemployment data do not yet track the risk categories: between early 2024 and early 2026, the largest rise in unemployment was among workers in jobs the framework classifies as facing the least immediate change⁶⁶.

Jobs are bundles — and some bundles are loose

The most illuminating recent framework comes from Luis Garicano and co-authors at the LSE. Jobs are bundles of tasks, but bundles differ in how tightly they are bound. A contractor or a junior hire generally holds a weak bundle: daily work consists largely of writing code to specification, tasks that can be handed to someone else — or to a machine — without disrupting the workflow. AI breaks off a large chunk and leaves a role with much diminished scope. A senior developer, or a coder whose programming sits inside domain expertise, holds a tight bundle: extracting the coding is much harder, so the bundle survives and AI becomes an assistant rather than a competitor⁶⁷.

The same technological capability shrinks one job and expands another. It erodes the junior version of a role even as it enhances the senior version.

This framework does real explanatory work. It accounts for why senior hiring holds up while junior hiring does not. It explains why the effect appears profession-agnostic — reported by software engineers, journalists,

⁶⁵ A. M. Richmond, The AI Jobs Transition Framework: Mapping AI's Near-Term Impact on Jobs, OpenAI Economic Research, April 2026.

⁶⁶ Ibid.

⁶⁷ L. Garicano et al., cited in J. Burn-Murdoch and M. Murgia, "The AI job loss story is all about bundles", Financial Times, 9 April 2026; see also L. Freund and L. Mann; J. Gans and A. Goldfarb.

musicians, financial professionals and lawyers alike. And it identifies the mechanism: when coding is automated, the senior developer gains time for the many other valuable parts of the job; take coding away from a junior developer and very little remains⁶⁸.

Official statistics have now caught up with the payroll data. Leland Crane and Paul Soto at the Federal Reserve, using an expanded definition of coders that includes contractors and programmers outside the tech industry, find through the flagship US labor force survey what Brynjolfsson found through payrolls: roughly half a million fewer coders working today than pre-LLM trends would have implied. This is not an absolute decline but a marked slowdown in growth⁶⁹.

Demand is the forgotten variable

Against the displacement story stands the most useful historical corrective in this literature. John Burn-Murdoch's walk through the software era shows how badly the “can AI do this task?” framing predicts employment. Large productivity increases in software since the 1990s were accompanied by rising, not falling, employment in web development, because pent-up demand outstripped the labor saved. The same held for accountants, architects and advertising creatives. In manufacturing, where demand for most goods is largely satiated, productivity gains translated into falling employment instead. The digital boom hollowed out retail — not through robots but by moving commerce online — while creating employment in warehousing and logistics. Spreadsheets devastated bookkeeping clerks while multiplying accountants and analysts. And radiologists have been insulated less by capability than by regulation and insurance, which make fully automated radiology de facto prohibited even where systems outperform humans⁷⁰.

The conclusion is uncomfortable for forecasters in both directions: so far, the occupations most exposed to AI have been about as likely to have grown as to have shrunk.

Where the demand goes: the missing half of every forecast

Almost every exposure study answers the question “can AI perform these tasks?” and stops. The historical record says that question determines relatively little. What determines employment is what happens to demand when the price of output falls — and the answer varies enormously by sector, which is why the same technology produces opposite outcomes in adjacent industries.

Consider the software era, which is the closest available analogue and the one that is usually misremembered. Productivity in producing software rose enormously from the 1990s, and employment in web development rose with it, because pent-up demand for software vastly outstripped the labor saved per unit of code. The same pattern held, less dramatically, across professional services: accountants, architects and advertising creatives all became more productive, and all became more numerous, because appetite for their services grew faster than their productivity⁷¹.

⁶⁸ Ibid.

⁶⁹ L. Crane and P. Soto, cited in Financial Times, 9 April 2026.

⁷⁰ J. Burn-Murdoch, “What the AI ‘jobpocalypse’ narrative misses”, Financial Times, 25 April 2026; on satiated manufacturing demand, J. Bessen.

⁷¹ J. Burn-Murdoch, “What the AI ‘jobpocalypse’ narrative misses”, Financial Times, 25 April 2026.

In manufacturing the same productivity growth produced the opposite result, for a simple reason identified by James Bessen: demand for most manufactured goods is largely satiated. When people do not want more of something, making it more cheaply reduces employment⁷².

Second-order effects can dwarf the direct ones. The digital boom did not replace shop assistants with robots; it moved commerce from bricks-and-mortar stores to websites, hollowing out retail employment — while simultaneously creating employment in warehousing, distribution and delivery. Within a single sector, the same tool can move two occupations in opposite directions: spreadsheets devastated bookkeeping clerks in the 1980s while multiplying accountants, financial analysts and data analysts⁷³.

Spreadsheets destroyed bookkeeping clerks and multiplied accountants. The same tool, the same decade, the same industry, opposite outcomes.

And regulation can matter more than capability. Radiologists have been substantially insulated from displacement not because AI cannot read scans — in many settings it now outperforms humans — but because regulation and insurance make fully automated radiology de facto prohibited. That example is the single most useful precedent for thinking about lawyers, auditors, notaries and the substantial share of French employment that sits behind a license⁷⁴.

The unpredictability cuts in both directions, and the historical examples are humbling. Bank tellers survived the ATM comfortably. What eventually destroyed the role was the smartphone, which removed the reason to enter a branch at all. Newspapers were not undone by a technology that produced journalism more cheaply; they were undone by online advertising and search engines dismantling the business model that funded it⁷⁵.

The agentic shift, and why it changes the unit of analysis

Everything above concerns AI as an assistant operating at the level of the task. The move to agents — systems that plan and execute multi-step work autonomously — changes the unit from the task to the workflow, and that is a different economic object.

The relevant metric here is the METR benchmark, which measures the length of task an agent can complete reliably. Mark Chen, OpenAI's head of research, describes the trajectory bluntly: a year ago the models were handling tasks a human could do in minutes; now the horizon is measured in hours; extrapolating forward, tasks that would take a human day⁷⁶. If that extrapolation holds even approximately, the bundling logic changes. Tight bundles are protected because the tasks within them are interdependent and hard to separate — but an agent that can hold a multi-day workflow in view is precisely a machine for handling interdependence.

Madhumita Murgia draws the implication that should most concern the senior professionals in this room: displacement falls on junior employees today, but that may not last. As agents take on higher-level tasks, it is a matter of time before experienced workers become project managers of AI rather than practitioners⁷⁷. Whether that is a promotion or a demotion is exactly what Panel 2 exists to discuss.

⁷²Ibid.

⁷³Ibid.

⁷⁴Ibid.

⁷⁵Ibid.

⁷⁶M. Chen, cited in J. Burn-Murdoch and M. Murgia, Financial Times, 9 April 2026.

⁷⁷Ibid.

One finding should be held alongside it: Anthropic's own research found that junior engineers relying on coding agents did not work much faster and understood their work less afterwards⁷⁸. The agentic transition may thus be accelerating automation while degrading the human capacity to supervise it — not a stable configuration.

6. The divide inside the workforce

If there is one finding in the last year that deserves more attention than it has received in France, it is this one — because it converts a productivity story into an inequality story without requiring a single job to disappear.

A poll of 4,000 workers in the United States and United Kingdom, conducted for the Financial Times by Focalddata, found adoption heavily skewed toward the best paid: more than 60% of the highest earners use AI daily, against just 16% of lower earners. The divide runs between occupations rather than within them — lawyers, accountants and software developers use these tools at similar rates whether junior or senior, but far more than lower-paid colleagues in the same industries. A persistent gender gap runs alongside it: women are roughly 20% less likely to use AI than men, on Google's own data⁷⁹.

Daron Acemoglu's reading is characteristically direct: the rhetoric says these tools are democratizing, but in reality they require education, abstract and quantitative skills, and familiarity with computers and code. His conclusion — that AI will increase inequality between labor and capital — is stated as near-certainty⁸⁰. Christopher Pissarides offers the same observation in gentler form: the more intelligent the technology we invent, the more your own intelligence matters. When technology was simpler, IQ mattered less⁸¹.

More than 60% of the highest earners use AI daily. Among lower earners, 16%. The tool that was supposed to level the field is being picked up fastest by those already ahead.

Two findings complicate the pessimism productively. The first is Carl Benedikt Frey's historical point: the same pattern appeared during the personal-computing revolution and evened out as use became widespread. His caveat is the operative one — it depends how long the gap takes to close; a decade or two would be genuinely worrying⁸². The second is that the gap is tractable. Google's chief economist Fabien Curto Millet points to research showing that a single AI training session for UK workers tripled daily usage among women over 55. The FT's own data identify corporate training as the single biggest driver of AI use at work⁸³.

There is also a surprise buried in the survey that bears directly on the apprenticeship question below. The heaviest users of AI at work are not the youngest but those in their thirties with longer tenures — consistent with OpenAI's own observation that AI complements proficiency, allowing established experts to be more productive. AI appears to be a multiplier on expertise, not a substitute for acquiring it⁸⁴.

⁷⁸Anthropic research, cited in J. Sun, The New York Times, 30 April 2026.

⁷⁹M. Murgia and J. Burn-Murdoch, "High earners race ahead on AI as workplace divide widens", Financial Times, 23 April 2026. Fieldwork: UK 26 February–2 March 2026 (n=2,365); US 6–9 March 2026 (n=1,754).

⁸⁰Ibid.

⁸¹Ibid.

⁸²Ibid.

⁸³Ibid.

⁸⁴Ibid.


Is the divide temporary or structural?

Everything depends on this question, and the honest answer is that we do not know — but the two candidate answers have very different policy implications, and both are represented in the evidence.

The optimistic reading belongs to Carl Benedikt Frey, and it is historically grounded. The same pattern appeared during the personal-computing revolution: early adoption skewed heavily toward the educated and the well-paid, wage premia for computer use were substantial, and the gap closed as use became universal and the skill ceased to be scarce. On this reading, today's 46-point gap is a transitional artefact of an early diffusion curve. Frey's own caveat is the operative sentence: it depends how long the gap takes to close; if it is a decade or two, that is more worrying⁸⁵.

The pessimistic reading belongs to Acemoglu and, in gentler form, to Pissarides. The claim is that this technology is different in kind: using it well requires education, abstract and quantitative reasoning, and comfort with computational tools — which are precisely the attributes already unequally distributed. Pissarides's formulation is that the more intelligent the technology, the more your own intelligence matters; when tools were simpler, IQ mattered less⁸⁶. If that is right, diffusion will not close the gap, because the constraint is not familiarity but capability. Two pieces of evidence bear on the question and point in opposite directions. In favor of the optimistic reading: the gap is demonstrably tractable. A single AI training session for UK workers tripled daily usage among women over 55, and the FT's own data identify corporate training as the single largest driver of workplace AI use. That is not what an insurmountable cognitive barrier looks like⁸⁷.

In favor of the pessimistic reading: the heaviest users are not the youngest, who are presumably the most comfortable with new tools, but people in their thirties with longer tenures — consistent with the observation that AI complements proficiency. If the tool multiplies existing expertise rather than substituting for it, then diffusion alone cannot equalize outcomes, because what is unequally distributed is not access but the expertise the tool multiplies⁸⁸. The synthesis that emerges, and that this document adopts, is that the gap is tractable but not self-correcting. Training closes it; time alone does not. That is a policy conclusion, not a forecast, and it is the strongest argument in this entire document for treating workforce training as urgent rather than eventual.



Source: chart by the authors, from the FT-Focaldata AI workforce tracker (n=4,119), Financial Times, 23 April 2026. Women are ~20% less likely to use AI than men

Figure 3 — The adoption divide. Daily AI use at work, by earnings level.

⁸⁵C. B. Frey, cited in M. Murgia and J. Burn-Murdoch, Financial Times, 23 April 2026.

⁸⁶Ibid.

⁸⁷Ibid.

⁸⁸Ibid.

7. Inside the firm: organisation, managers, meaning

The macro picture may be contested, but the organizational picture is clear, and it is where most of the people in this room actually have agency.

The technology is not the variable

The most consistent finding across firm-level research is that outcomes depend far less on which model a company deploys than on whether it redesigns processes, retrains teams and redefines roles. Ethan Mollick's experiment with 776 employees at Procter & Gamble found that individuals using AI performed as well as two-person teams not using it — and, more interestingly, that it blurred the line between roles: business people began producing technical ideas, technical people business ones⁸⁹.

The management literature has converged on a corollary: AI transformation requires redesigning work rather than cutting roles, and firms that begin from a headcount target rather than a workflow analysis systematically underperform⁹⁰. The paradox worth stating plainly to any executive team is that AI does not reduce workload; it intensifies it. Employees who adopt these tools handle more cases, answer more requests, absorb new tasks. The gain is real and it is reinvested in additional work rather than returned as time⁹¹.

Middle management under pressure

Middle managers sit at the point of maximum strain, for a structural reason: much of the historical value of the layer — routing information, coordinating, reporting — is precisely what automated systems now handle. But the evidence does not show them being replaced; it shows them overloaded, absorbing the coordination cost of running parallel human and machine workflows⁹². Acemoglu's assessment is that managers currently feel behind the curve, unclear what is expected of them, and under materially higher stress; the tools and expectations have not been standardized⁹³.

Three ways of working with AI — and only two of them build skill

The finding that outcomes depend on organizational choices has a counterpart at the level of the individual, and it is unusually concrete. Field work with management consultants using the same tools, in the same firm, identified three distinct patterns of use. Centaurs delegate specific tasks to AI while keeping authority over the shape of the work: there is a clear division of labor, and the person decides what to hand over. Cyborgs weave the machine into a continuous back-and-forth, so that human and model iterate on the same piece of work. Self-automators hand over whole workflows and step back⁹⁴.

The skill consequences diverge sharply. Centaurs deepen their own domain expertise, because they keep making the judgement about what can safely be delegated. Cyborgs acquire new capabilities of a different kind — largely

⁸⁹E. Mollick, in “Who Will Actually Thrive in the Hybrid A.I.-Human Work Force”, The New York Times Magazine, 14 June 2026.

⁹⁰“AI transformation requires redesigning work, not cutting roles”, Harvard Business Review, August 2026.

⁹¹“AI doesn't reduce work — it intensifies it”, Harvard Business Review, February 2026.

⁹²“AI adoption is overloading your middle managers”, Harvard Business Review, June 2026.

⁹³The New York Times Magazine, 14 June 2026.

⁹⁴F. Dell'Acqua and co-authors, field work with Boston Consulting Group consultants; see also S. Randazzo, “Cyborgs, Centaurs and Self-Automators”, Harvard Business School Working Paper 26-036, 2025, and M. Sigelman et al., Beyond the Binary: How Automation and Augmentation Are Combining to Reshape Work, January 2026.

the skill of working with the machine. Self-automators develop nothing at all: the work gets done and the person does not advance⁹⁵.

Same firm, same tools, three patterns of use — and three completely different outcomes for the person's skill. The technology did not decide which; the individual and the organization did.

Two conclusions follow for anyone deploying these tools at scale. The first is that “has the team adopted AI?” is the wrong question; the useful one is which of these three patterns adoption has produced, because only two of them leave the organization with more capability than it started with. The second is that all three modes presuppose expertise the person already has — the judgement to know what to delegate, and the ability to evaluate what comes back. That is the hinge between this chapter and the next: the modes that build skill require skill to begin with, which is precisely what the erosion of entry-level work is putting at risk.

The part that gets automated may be the part people valued

There is a dimension of this transition that economic data cannot capture, and Sarah O'Connor's reporting for her book on the subject surfaces it better than any dataset. Interviewing translators, software engineers, Amazon warehouse workers and miners, she found a story that fits neither the utopian nor the dystopian narrative: people were not freed for more creative work, nor were they thrown out of their jobs. They were losing the parts of their jobs they found meaningful, and working faster to keep pace with machines. Warehouse work became more productive and simultaneously lonelier — standing in one place for ten hours while robots bring the shelves. Translators moved from translating to editing machine output, faster and for less⁹⁶.

Some workers were not displaced and not liberated. They were left doing the residue — the parts of the job the machine could not do, at the machine's pace.

Her finding on what separates good from bad deployments is the most actionable sentence in this entire document: it depends on whether workers themselves have control over how and when to use the tools. Where employers trusted staff to experiment, outcomes were better — people know their own jobs in detail and can judge what is safe to offload. Where AI was imposed, and particularly where management redesigned jobs around the strengths and weaknesses of the machine rather than of the humans, workers were crunched into roles that plugged the inadequacies of the AI⁹⁷.

Meaning, control and the limits of the productivity frame

There is a dimension of this transition that no productivity statistic will capture, and it may determine whether AI deployment succeeds or fails inside firms more reliably than any technical variable.

Sarah O'Connor spent a year interviewing people at the frontline of the change — translators, software engineers, Amazon warehouse workers, miners in Sweden — and found a story that fits neither prevailing narrative. People were not liberated for more creative work. Nor were they, mostly, thrown out of their jobs.

⁹⁵Ibid.

⁹⁶S. O'Connor, We Are Not Machines: The Fight for the Future of Work, discussed in “Why (happy) humans still matter at work”, Financial Times, 17 June 2026.

⁹⁷Ibid.

What many described was losing the parts of their work they had found meaningful and working faster to keep pace with machines. Some felt they were becoming more machine-like themselves⁹⁸.

The Amazon example is precise and instructive. Robots made warehouse workers considerably more productive by bringing shelves to them. The same change intensified the pace and made the job much lonelier, standing in one place for ten hours rather than moving through a building among colleagues. Translators moved from translating to post-editing machine output, faster and for less money, with the creative element removed and the tedium retained⁹⁹.

One reader's account, published in the same series, generalizes the pattern beyond translation: software engineers reporting that the most creative activities of the job are being replaced by the tedium of checking machine output, that skills atrophy as they write less code, and that there is less scope to demonstrate style and originality¹⁰⁰.

The question is not only how much work AI removes, but which parts. A job stripped of its craft and left with its supervision is not obviously a better job.

O'Connor's finding about what separates good deployments from bad ones is the most directly actionable sentence available to the employers in this room, and it is worth quoting in substance. What makes the difference is whether workers themselves have control over how and when to use the tools. Where employers trusted staff to experiment, outcomes were better — most people want to do a good job, know their own work in detail, and can judge which parts are safe to offload and where a new capability opens something up. Where deployment was imposed, and particularly where management redesigned roles around the strengths and weaknesses of the machine rather than of the humans, workers were crunched into jobs that existed to plug the inadequacies of the AI. That is neither satisfying nor, as she notes, something humans are well suited to¹⁰¹.

The two-speed organization

One structural observation deserves to be on the table for any French corporate leader present. The most striking productivity and revenue growth is not visible at incumbents adding AI to existing workflows; it is visible at firms built around the technology from the beginning, whose products are written and reviewed by it. Demirer's interpretation of his own results is that incumbent structures and marketplaces are not organized to capture the underlying gains, and that history suggests the real jumps in productivity come from new companies and new processes rather than from established firms grafting technology onto old workflows¹⁰².

If that is right, the competitive threat to a large French company is not that its own AI program underperforms. It is that a firm designed around these tools reaches the same market with a fundamentally different cost structure. That reframes the internal debate from “how do we adopt?” to “what would we build if we were starting now?” — a question most transformation programs are structurally incapable of asking.

⁹⁸S. O'Connor, We Are Not Machines: The Fight for the Future of Work, discussed in Financial Times, 17 June 2026.

⁹⁹Ibid.

¹⁰⁰Reader correspondence, “The AI Shift: Readers tell us how their jobs are changing”, Financial Times, 11 June 2026.

¹⁰¹Financial Times, 17 June 2026.

¹⁰²J. Burn-Murdoch, Financial Times, 5 June 2026.

8. The broken first rung

If the reader retains one risk from this document, it should be this one. It is the least visible, the slowest to materialize, and the hardest to reverse.

Every profession has produced its own experts through the same mechanism for centuries: junior people do the grunt work, are assessed on it, and learn by doing. Ethan Mollick's formulation is that apprenticeship has worked for four thousand years, and that it has now collapsed — you cannot simply say we should hire junior workers if you have not thought about how to train them when the training tasks have been automated¹⁰³.

The empirical support has strengthened considerably. Brynjolfsson's canaries analysis found employment for 22-to-25-year-olds in highly exposed roles falling around 6% between late 2022 and mid-2025, with declines near 20% for young software developers¹⁰⁴. The Crane and Soto work corroborates the direction using official statistics. And Anthropic's own research adds a mechanism that should worry anyone building a training program: junior engineers who relied on AI coding agents not only failed to complete tasks much faster, they also understood their work less when quizzed afterwards¹⁰⁵.

Put those findings beside the Focaldata result that AI complements proficiency, and the shape of the trap becomes clear. AI is most useful to those who already have expertise. It is simultaneously eroding the roles in which expertise was acquired, and degrading the learning of those who use it before they have any. Mollick's observation is that field experience is what allows you to evaluate work you did not produce — the ability to glance at an article, a paper, a piece of legislation or a block of code and judge its quality. That is precisely the skill required to supervise AI agents, and precisely the skill the new entry-level path no longer builds¹⁰⁶.

AI is most valuable to those who already have expertise — and it is dismantling the roles through which expertise was acquired. That is not a hiring problem. It is a reproduction problem.

Carl Benedikt Frey supplies the sentence that should govern the policy response: most economists acknowledge technological progress causes adjustment problems in the short run; what is rarely noted is that the short run can be a lifetime¹⁰⁷.

¹⁰³E. Mollick, The New York Times Magazine, 14 June 2026.

¹⁰⁴E. Brynjolfsson, B. Chandar, R. Chen, Canaries in the Coal Mine?, Stanford Digital Economy Lab, 2025.

¹⁰⁵Anthropic research, cited in Jasmine Sun, "Silicon Valley Is Bracing for a Permanent Underclass", The New York Times, 30 April 2026.

¹⁰⁶The New York Times Magazine, 14 June 2026.

¹⁰⁷Cited in The New York Times, 30 April 2026.

Five types of technology and their labour-market consequences

Only new-task-creating technologies are unambiguously favourable to workers

Type of technology	Example	Labour productivity	Value of human expertise	Labour's share of income	Pro-worker?
1. Labour-augmenting	Electric cable stripper	+	+/-	≈ 0	Ambiguous
2. Capital-augmenting	Faster cable stripper	+	+/-	≈ 0	Ambiguous
3. Automation	Cable-installing robot	+	- expertise made obsolete	- fewer tasks done by labour	Not pro-worker
4. New task-creating	Ethernet, fibre optics	+	+ new expertise needed	+ more tasks done by labour	Unambiguously pro-worker
5. Expertise-levelling	Blood oximeter	+	+/- entrants gain, incumbents lose	+/- offsetting effects	Ambiguous

Source: chart by the authors, redrawn from D. Acemoglu, D. Autor and S. Johnson, "Building Pro-Worker Artificial Intelligence", Brookings, 2026.

Figure 4 — Types of technologies and their labor-market consequences. Only new-task-creating technologies are unambiguously pro-worker. Source: Acemoglu, Autor & Johnson (Brookings, 2026), via MIT Sloan.

This is where the taxonomy above stops being academic. Expertise-levelling technologies let novices perform expert tasks — good for access, ambiguous for the incumbent expert, and corrosive of the incentive to acquire depth. Only new-task-creating technologies are unambiguously pro-worker, and they are the most demanding in training terms. Acemoglu, Autor and Johnson argue that current incentives push systematically toward automation rather than augmentation, and that this is a market failure rather than a series of bad corporate decisions: path dependence in business models, an AGI-oriented research culture largely indifferent to labor effects, and concentration that suppresses alternative designs¹⁰⁸.

What would actually have to be built

The diagnosis of a broken first rung is now widely shared. The response is not, and the proposals in circulation are of very uneven seriousness. Four are worth examining.

The first is to reform curricula, from primary school through the grandes écoles, toward the capacities models cannot replicate: critical judgement, reasoning under uncertainty, contextual creativity, relational intelligence. This is necessary and it is slow. It also carries a risk that the debate has not confronted: research on unstructured AI use in education has found materially negative learning outcomes, while structured AI tutoring shows large positive effects. The difference between those two findings is the entire design question, and getting it wrong at national scale would damage a generation¹⁰⁹.

The second is mass retraining targeted at exposed occupations, delivered through short certified pathways accessible to people in work, under tripartite governance. France has more institutional machinery for this than most countries. The obstacle is not capacity but trigger: retraining systems are designed to activate on redundancy, and the scenario evidence says the interventions that matter must precede measurable displacement.

¹⁰⁸ D. Acemoglu, D. Autor and S. Johnson, Building Pro-Worker Artificial Intelligence, Brookings / NBER WP 34854, 2026.

¹⁰⁹ Discussion in "Who Will Actually Thrive in the Hybrid A.I.-Human Work Force", The New York Times Magazine, 14 June 2026; D. Acemoglu references large-scale Chinese evidence on adverse outcomes from student AI use.

The third is the reconstruction of apprenticeship itself, and it is the one most specific to the diagnosis. If junior roles no longer generate learning because their content has been automated, then the learning has to be produced deliberately somewhere else. Ethan Mollick's suggestion is that universities are, in fact, rather good at training and assessing, and could extend professional education into the gap AI is creating¹¹⁰. Clara Shih's is that every young person should own a project end to end, in order to develop the judgement required to supervise machine output — while conceding that this does not work for everyone and correlates uncomfortably with background and socioeconomic status¹¹¹. Neither is a system. Both are gestures toward one.

The fourth is financing, and it is where the argument becomes concrete. Firms are investing hundreds of billions in compute, data centers and models. A comparable order of investment in human capital is not a moral supplement to that program; on the scenario evidence it is the condition under which the program produces broadly shared returns rather than concentrated ones.

Retraining systems activate on redundancy. The evidence says the interventions that matter must come first — which is exactly when they are hardest to fund.

The direction of the technology, not only the response to it

All four proposals above accept the technology as given and organize the human response around it. Acemoglu, Autor and Johnson argue that this is a category error, because the direction of technological development is itself a choice — and one currently being made badly.

Their diagnosis of the market failure has three components. Path dependence: large firms have built business models around automation tools from which they derive their profitability, and those models shape what gets built next. Ideology: the dominant research culture in frontier laboratories is oriented toward general capability and is largely indifferent to labor-market effects. And concentration: a small number of firms determine the research agenda, which suppresses alternative designs that might augment rather than replace¹¹².

Their policy levers follow from that diagnosis rather than from the usual training agenda. Competition enforcement, to open the market to entrants with different business models. Legal protection of professional expertise, so that workers whose know-how is absorbed into AI systems hold rights over that contribution — a principle already operationalized in the 2023 SAG-AFTRA agreements, which obliged studios to obtain consent and provide compensation before using an actor's image or voice for model training. And the institutionalization of workers' voice in deployment decisions, at firm and sectoral level¹¹³.

Simon Johnson's summary — that we are not currently on the pro-worker AI path — is a claim about incentives, not about intentions. It is also the point at which this debate stops being about labor economics and becomes about competition policy, intellectual property and corporate governance. That is uncomfortable territory for a discussion framed around the future of work, and it is probably where the most consequential decisions actually sit.

¹¹⁰E. Mollick, The New York Times Magazine, 14 June 2026.

¹¹¹C. Shih, *ibid*.

¹¹²D. Acemoglu, D. Autor and S. Johnson, Building Pro-Worker Artificial Intelligence, Brookings / NBER WP 34854, 2026.

¹¹³*Ibid*.

9. Where the pressure lands: the sectoral picture




Figure 5 — The AI dividend flows differently by sector and by stakeholder. Consumers gain almost everywhere; employees lose in a majority of sectors; compute and AI providers capture value across the board.

Most participants will want to know what this means for their own industry. The honest starting point is that the sector is the wrong unit. The consultancy work that has looked hardest at this finds the value concentrating by function rather than by sector — which is why two companies in different industries can face more similar pressure than two companies in the same one.

Three quarters of the value sits in four functions

McKinsey's estimate is that generative AI could add the equivalent of 2.6 to 4.4 trillion dollars a year across the sixty-three use cases it examined — roughly 6.1 to 7.9 trillion once the effect of embedding the technology in software used for other purposes is included. The number worth remembering is not the total but its distribution: about 75% of it falls in four functions. Customer operations, marketing and sales, software engineering, and research and development¹¹⁴.

That concentration explains a great deal of what the rest of this document describes. Every sector has a customer-service function, a marketing function, a software function and, in most cases, an R&D or product function. The exposure of a bank, a retailer and an industrial group therefore looks similar at the level of these functions and quite different at the level of the business itself. It also explains why the occupations under most pressure — contact-centre work, junior software development, content production, first-line analysis — recur across industries that otherwise have little in common.

The sector is the wrong unit. Three quarters of the value sits in four functions that exist in every sector — which is why exposure travels by job family rather than by industry.

¹¹⁴ McKinsey Global Institute, "The Economic Potential of Generative AI: The Next Productivity Frontier", June 2023, and subsequent sector work.

Which sectors gain most, and in what sense

Two different questions are often conflated. Measured as a share of industry revenue, the largest gains are expected in banking, high technology and life sciences. Measured in absolute terms, the largest pools sit elsewhere: retail and consumer packaged goods together could see 400 to 660 billion dollars a year. Banking's own potential is put at 200 to 340 billion a year, equivalent to 9 to 15% of operating profits, with the largest absolute gains in corporate and retail banking rather than in the trading businesses that attract most attention¹¹⁵.

One sectoral estimate deserves to be read carefully by anyone running a service operation: generative AI could reduce human-serviced customer contacts by up to half in banking, telecommunications and utilities¹¹⁶. That is not a productivity improvement of the kind described in chapter 2. It is the removal of a category of work — and it lands on exactly the roles that have historically offered entry into stable employment without a degree.

The levers, sector by sector

Below, the principal levers identified for each sector with the orders of magnitude attached to them. These are potentials rather than forecasts — what could be captured if the use cases were implemented — and they are useful less as predictions than as a checklist of where to look first.

Banking and insurance. The largest identified value pools are risk and legal work, corporate banking and retail banking, ahead of the trading and capital-markets activities that attract most of the attention. The sector total is put at 200 to 340 billion dollars a year, equivalent to 9 to 15% of operating profits — a lower profit uplift than other sectors precisely because margins are already high. Service deflection is the most immediate lever: up to half of human-serviced customer contacts¹¹⁷.

Life sciences and health. The lever here is research rather than administration, which makes it the sector where the technology reaches the core of the business rather than its overheads. AI applied to accelerating R&D is credited with 360 to 560 billion dollars of annual value across all industries, with molecule screening, trial design and evidence synthesis as the specific applications in life sciences¹¹⁸. Alongside it sits diagnostic support, where the constraint is regulatory rather than technical — the radiology case discussed in chapter 5.

Retail and consumer goods. The largest absolute pool of any sector: 400 to 660 billion dollars a year, which at the upper bound would represent an uplift of around 44% to profits, reflecting thinner margins¹¹⁹. The levers are concentrated in marketing and merchandising — the productivity of the marketing function alone is estimated at 5 to 15% of total marketing spend — together with assortment, pricing and customer service.

Industry and manufacturing. Two distinct levers, on different timescales. In the near term, R&D acceleration: for industries producing complex manufactured products, development processes could be compressed by 20 to 80% depending on the sector¹²⁰. In the medium term, the physical frontier described in chapter 3 — vision-guided robotics, predictive maintenance, quality control — which addresses labor shortages rather than creating surpluses, and which plays to European industrial strengths.

¹¹⁵McKinsey Global Institute, “The Economic Potential of Generative AI: The Next Productivity Frontier”, June 2023, and subsequent sector work.

¹¹⁶Ibid.

¹¹⁷McKinsey Global Institute, “The Economic Potential of Generative AI”, June 2023; 63 use cases across 16 business functions, 850+ occupations and 2,100 work activities in 47 countries.

¹¹⁸McKinsey, “The next innovation revolution — powered by AI”, 2025.

¹¹⁹McKinsey Global Institute, “The Economic Potential of Generative AI”, June 2023; 63 use cases across 16 business functions, 850+ occupations and 2,100 work activities in 47 countries.

¹²⁰McKinsey, “The next innovation revolution — powered by AI”, 2025.

Energy and utilities. The levers are operational rather than cognitive: generation scheduling, grid balancing, demand forecasting and asset performance. Utilities deploying optimization engines of this kind report reductions in operating costs and in carbon intensity in the range of 5 to 15%, depending on asset mix and starting point¹²¹. For France, with a low-carbon generation base and an ageing grid, this is the sector where the technology and the energy transition point in the same direction.

Technology and software. The sector where all the levers apply directly, because the product consists of intellectual property: with no physical prototyping or manufacturing to constrain them, R&D acceleration, code generation and testing can be applied end to end¹²². This is also why the sector is the laboratory for everything in this document — and why its experience generalizes less well than it appears to.

Professional and business services. The levers are knowledge synthesis, drafting and first-pass analysis. What distinguishes this sector is not the size of the opportunity but the elasticity of demand: unmet need for legal, accounting and advisory work is large and price-rationed, which is the mechanism that historically turned productivity gains into more employment rather than less, as discussed in chapter 5.

Public services. The levers are the same as in any large administration — case handling, documentation, first-line advice, procurement — but the decision rule is different, because no market forces it. Health and education are the two sectors where the potential is greatest and the record to date weakest, which makes them the place where the design of deployment matters most.

Where the value is, and where the money is going

One finding from the same body of work deserves particular attention from the investors and executives present, because it identifies an opportunity rather than a risk: investment is poorly aligned with potential. Consumer goods, retail and financial services represent hundreds of billions of identified value while remaining comparatively underinvested; healthcare and technology show much better alignment, with capital flowing to drug discovery, diagnostics and software development¹²³.

Retail could capture 400 to 660 billion dollars a year, yet few companies in the sector invest at top-quartile levels. The gap between where the value sits and where the money goes is itself the opportunity.

A second corrective is worth stating plainly, because the public conversation has drifted away from it. The same analysis estimates that extensive use of non-generative AI — the machine learning and analytics that have been maturing for fifteen years — would deliver 11 to 17.7 trillion dollars of economic value, roughly four times the generative figure¹²⁴. For most organizations in most sectors, the largest available gains are still in forecasting, optimization and prediction rather than in text generation. That is an unfashionable conclusion, and it is probably the most immediately actionable one in this chapter.

¹²¹McKinsey / QuantumBlack energy practice, reported client outcomes, 2025–2026.

¹²²McKinsey, “The next innovation revolution — powered by AI”, 2025.

¹²³McKinsey Global Institute, “The Economic Potential of Generative AI”, June 2023; 63 use cases across 16 business functions, 850+ occupations and 2,100 work activities in 47 countries.

¹²⁴Ibid.

Five patterns of change, and why the label matters more than the industry

The most useful recent taxonomy comes from BCG, which modelled roughly 165 million US jobs across about 1,500 roles and sorted them by what actually happens to the work. It is worth restating in full, because the categories are more informative than any sector ranking¹²⁵.

Amplified roles are those where AI expands what a person can do and demand grows to absorb it. Rebalanced roles keep their headcount but see their task mix shift substantially. Enabled roles — about 23% of jobs — have AI embedded in daily activity, changing how tasks are performed without altering how the work is structured; workers are simply expected to use it. Divergent roles, about 12%, are the ones this document has argued matter most: junior and entry-level positions face the earliest automation of structured tasks while senior responsibilities persist or grow. Substituted roles, another 12%, are those where AI replaces human input in core tasks and demand does not expand to compensate — contact-centre representatives being the clearest case.

The headline arithmetic: 50 to 55% of jobs meaningfully reshaped within two to three years, and 10 to 15% — some 16 to 25 million US positions — vulnerable to elimination over four to five years¹²⁶. Read against chapter 5, the convergence with the OpenAI framework is notable: two independent exercises, different methods, and a similar conclusion that reorganisation dominates elimination by a factor of roughly four.

The two questions that decide a sector's outcome

Underneath the taxonomy sit two tests that are worth applying to your own industry, because they determine which category your roles fall into.

The first is demand headroom: if output becomes cheaper, is there unmet demand to absorb it? This is the elasticity question from chapter 5, and it is why software employment grew through decades of productivity gains while manufacturing employment did not. The second is structural scalability: even where demand exists, can supply actually expand? Capital intensity, physical capacity and capex requirements can prevent a sector from growing into its own productivity gain¹²⁷. A professional-services firm can take on more clients next quarter; a hospital, a utility or a manufacturer cannot.

The combination is what separates a sector where AI grows employment from one where it compresses it. It is also a more useful diagnostic than exposure, because both questions can be answered by people who know their industry.

¹²⁵ Boston Consulting Group, “AI Will Reshape More Jobs Than It Replaces”, April 2026. The model covers roughly 165 million US jobs across about 1,500 roles.

¹²⁶ Ibid.

¹²⁷ Boston Consulting Group, “AI Will Reshape More Jobs Than It Replaces”, April 2026. The model covers roughly 165 million US jobs across about 1,500 roles.

What workers across sectors are actually reporting

Survey evidence across 14 markets and nearly 12,000 workers gives the cross-sector view from inside. Some 72% say AI has already considerably changed the skills expected of them. Nearly half — 47% — now spend more time managing and directing AI than doing the work themselves. Two thirds of regular users say it has improved their job satisfaction, while four in ten report higher cognitive load: a paradox in which the work becomes both better and harder¹²⁸.

Two findings from the same survey bear directly on chapter 11. Organisations that have moved past deploying tools toward redesigning end-to-end workflows capture materially more value and report better employee experience. And strategic clarity, rather than better tooling, is what lifts measurable business impact — by around 25 percentage points¹²⁹.

Companies that cut beyond what AI can actually deliver will see productivity fall, institutional knowledge disappear and talent leave. Some of the current reductions are running ahead of the technology.

The sector that no consultancy report covers

One sector goes almost unmentioned in this literature and is disproportionately important in France. A very large share of exposed cognitive work — administrative processing, case handling, documentation, first-line advice — sits inside the state, the hospitals, the collectivités and the social-security system.

This has three consequences. Adoption will be slower, because public procurement and employment rules are slower. The distributional effect will be muted, because employment protection is stronger. And the decision is genuinely political, in the sense that no market will force it. That combination makes the French public sector the clearest case in this whole document where the direction of deployment — augmenting the work of public servants or reducing their number — is an explicit choice rather than an emergent outcome. It is therefore also the place where a government could demonstrate, at scale, what pro-worker AI actually looks like.

¹²⁸ Boston Consulting Group, “AI at Work 2026”, fourth annual survey, 11,749 workers across 14 markets, June 2026.

¹²⁹ Ibid.

What to do about it

Inside the firm, and at the level of the country

10. France in comparative perspective

Draghi: compute as the condition of adoption

Mario Draghi set out his own answer to this question in the Financial Times on 11 September 2026 — six days before this seminar — and it is worth opening the chapter with it, because it locates the European constraint one step further upstream than the preceding pages.

His figures are stark. The productivity gap between the euro area and the United States widened from 9 dollars an hour in 2018 to 21 dollars in 2025. European total factor productivity growth has been approximately zero since 2022, while ECB scenarios suggest rapid AI adoption could add 0.3 to 0.4 percentage points a year. And the constraint he identifies is physical: the European Union hosts under 5% of the world's AI compute capacity against 75% for the United States. The gap between European demand and installed supply is around 3 gigawatts today, roughly a quarter of existing capacity, and on current trends reaches 14 gigawatts by 2030¹³⁰.

His proposed remedy is concrete and unusually practical for a piece of macroeconomic advocacy: European companies should pool their purchasing commitments into contracts large enough to finance new data-centre construction. He points to the group — ASML, Capgemini and Amadeus among them — that has committed to multiyear purchases of Mistral's European Compute Units, an arrangement intended to underwrite a gigawatt of capacity by 2030¹³¹. It is demand aggregation rather than subsidy, and it is the kind of measure that does not require a treaty change.

Under 5% of the world's AI compute capacity, against 75% for the United States. If compute is scarce and expensive in Europe, the firms rationed first are precisely the ones already least likely to adopt.

It would be a misreading to treat this as a rival diagnosis to the one in chapter 2. Draghi's argument is about adoption too: the ECB scenarios he cites are explicitly conditional on rapid adoption, and compute is the condition that makes it possible. His remedy runs through adoption rather than around it — pooled purchase commitments are demand for AI services converted into financing for the capacity to supply them. Infrastructure and diffusion are two links in one chain, not two candidate explanations.

What the two readings together suggest is more uncomfortable than either alone. A compute constraint does not bind evenly: scarce and costly capacity is rationed by ability to pay and by proximity to existing infrastructure, which means it bites hardest on smaller firms and on member states with no capacity of their own. Those are

¹³⁰ M. Draghi, "Europe must build more AI data centres", Financial Times, 11 September 2026.

¹³¹ Ibid.

exactly the groups on the wrong side of the adoption gap described in chapter 2 — 52% of large European firms against 17.4% of small ones, and a tenfold spread between member states. The infrastructure deficit and the diffusion deficit therefore compound rather than compete, and a policy that addresses only one will leave the other in place.

For France the two links meet on favourable ground. A low-carbon generation base makes hosting compute more defensible here than almost anywhere in Europe, and a tradition of tripartite governance offers a route to the diffusion problem that a purely market-led system does not. The obstacle on the infrastructure side is not electricity but the political economy of data centres described in chapter 4; the obstacle on the diffusion side is the one set out in chapter 11.

Most of the evidence in this document is American. That is a genuine limitation. The United States has a more fluid labor market, weaker employment protection, a larger frontier-AI sector and a different wage structure; findings on entry-level displacement do not transfer mechanically.

Three structural features shape French exposure. Employment protection and collective bargaining mean adjustment will appear as reduced hiring rather than visible layoffs — slower, less explosive, and considerably harder to measure. A large public and para-public sector means much exposed cognitive work sits inside institutions whose adoption decisions are political rather than commercial. And a comparatively thin domestic frontier-AI industry means France will import capability rather than capture the returns from producing it, which shifts the distributional question toward trade and sovereignty¹³².

French projections for 2026–2030 identify intermediate administrative professions, certain support-function managers and skilled office workers as most exposed. The social model does not confer immunity; it changes the shape of the adjustment¹³³.

The nearer-term risk may not be displacement at all but under-adoption. European firms are adopting more slowly, and the productivity evidence is unambiguous that gains accrue to organisations that redesign work rather than bolt the technology on. A France that avoids disruption by failing to adopt has not solved the problem — it has chosen the concentrated-gains scenario without the compensating growth.

Against that, three assets are worth naming. France has unusual academic depth in the economics of innovation and labor. It has a tradition of tripartite governance that could deliver sectoral retraining pathways faster than a purely market-led system. And it has an industrial-policy tradition, which matters if the answer to concentration involves more than competition enforcement.

The specific French vulnerability is the interaction between two things: a labour-market entry problem that long predates AI, and the erosion of the entry-level rung that is the best-supported finding in this whole literature. Youth unemployment and a difficult school-to-work transition are chronic French weaknesses. AI is now attacking the mechanism through which young people enter, in a country whose entry mechanism was already the weak point. That interaction, rather than aggregate displacement, is the risk this seminar is best placed to address.

¹³² Direction générale du Trésor, « L'intelligence artificielle : quels effets sur l'emploi ? », 30 June 2026; OECD Employment Outlook 2026.

¹³³ Exposition à l'IA générative et emploi : application à la classification socio-professionnelle française, 2024.

Four questions for the day

1. The premise. What developments in capability and cost should France plausibly prepare for over the next five years — and which observable indicators would tell us early which trajectory we are on? Candidates: the closing of the capability overhang; the composition rather than volume of hiring; the cost per completed agentic task; the labor share in AI-intensive sectors.
2. The distribution. If the aggregate gain is real but flows to consumers, shareholders and compute providers rather than employees, what mechanism could redirect part of it — and who would have to decide? Sovereign funds, universal basic capital, taxation of automation, and the objections to each.
3. The first rung. If entry-level roles are automated and apprenticeship breaks, what is the credible French answer — in curricula, in apprenticeship contracts, in firms' own hiring practice, and in who pays for the training that used to happen implicitly on the job?
4. The political ownership. Which of these choices can be delegated to firms and social partners, and which require explicit political commitment ahead of 2027 — before the polling that shows 79% of voters expecting no plan hardens into something less manageable?

The purpose of the day is not to settle these questions. It is to establish, among people who rarely sit in the same room, where the genuine disagreements lie — and which of them are empirical and resolvable by better measurement, and which are ultimately political and will have to be decided by someone with a mandate.

11. Transforming the firm: what the evidence says actually works

The preceding chapters describe a transition. This one is about the part of it that the people in this room control directly. Most of what has been written about AI in companies is either vendor material or anxiety; a smaller body of work has measured what separates the firms capturing value from those spending money without result. That evidence converges on a conclusion that is both encouraging and inconvenient: the constraint is almost never the technology.


The 10-20-70 rule, and why most programmes get it backwards

The most useful single heuristic comes from BCG's cross-industry benchmarking, and it is a statement about where effort should go. Top-performing organisations devote roughly 10% of their effort to algorithms, 20% to technology and data, and 70% to people, processes and cultural transformation. The failure mode is the inverse: programmes that allocate most of their attention and budget to models and infrastructure, and treat adoption as a communications exercise at the end¹³⁴.

¹³⁴BCG, “Closing the AI Impact Gap”, 2025, and “How Agents Are Accelerating the Next Wave of AI Value Creation”, December 2025.

The 10-20-70 principle

Allocation of effort in AI transformations that deliver measurable value



Source: chart by the authors, from BCG, "Closing the AI Impact Gap" (2025) and "How Agents Are Accelerating the Next Wave of AI Value Creation" (2025).

Figure 6 — The 10-20-70 principle. The majority of the value comes from the part most transformation programs treat as an afterthought.

This is not a soft finding dressed up as a hard one. It is the organizational counterpart of the productivity funnel described in chapter 2. If the gain generated at the level of the task dissipates at every human handoff, then the returns to investing in more generation capacity fall sharply while the returns to investing in absorption — process design, role definition, review capacity, decision rights — remain high. The 70% is where the funnel gets unblocked.

McKinsey's version of the same insight is expressed as a set of capabilities rather than a ratio: a transformation roadmap tied to identified value rather than to technology; a talent bench with genuine expertise; an operating model capable of moving at pace; a flexible technology environment; data embedded across the organization rather than in a central team; and — the one most often missing — an explicit capability for adoption and scaling, which is what converts pilots into results¹³⁵.

Ten per cent algorithms, twenty per cent technology and data, seventy per cent people and processes. Most programs are funded in exactly the opposite proportion.

Redesign the process, not the org chart

The second consistent finding is that value comes from end-to-end process redesign rather than from inserting AI into existing workflows. BCG's own agentic deployments show this starkly: end-to-end redesign is what separates organizations achieving cost reductions of 60% or more from those capturing less than 20%. The near-term objective, as they put it, is not to deploy more agents¹³⁶.

A concrete example makes the logic legible. One industrial company redesigning its request-for-quotation process targeted a distribution in which roughly 70% of RFQs are resolved by agents without human intervention — the smaller transactions with established specifications; around 20% require human intervention working alongside agents; and the remaining 10%, the most complex and highest-value transactions, require intensive human judgement supported by AI. The company first assessed whether its systems and data could give the agents adequate context, and is implementing the change over fifteen to eighteen months in two releases¹³⁷.

¹³⁵ McKinsey (QuantumBlack / McKinsey Digital), Rewired methodology; see also E. Hazan, A. Madgavkar, M. Chui et al., A New Future of Work, McKinsey Global Institute, 2024.

¹³⁶ BCG, "Reinventing the Operating System of Work with AI", July 2026.

¹³⁷ BCG, "Scaling AI Requires New Processes, Not Just New Tools", March 2026.

Three features of that example are worth extracting. The target is a distribution of work, not a headcount. The sequencing runs from data readiness to deployment rather than the reverse. And the timeline is measured in quarters, not weeks — which is the honest answer to a board asking when the returns arrive.

The contrast with the historical record is instructive and returns us to the electrification analogy. Traditional process optimization was built around a scarce resource: human capacity. It reduced waste through standardization and flow control. When the scarce resource changes, the optimization problem changes with it — and an organization that keeps optimizing for the old scarcity while paying for the new capability gets the worst of both.

The workforce is the transformation

BCG's more recent framing is blunter still: AI transformation is a workforce transformation. Within the 70%, the firms that outperform are roughly four times more likely to run structured AI-learning programs and to protect dedicated time for employees to learn¹³⁸. That finding sits directly alongside the FT–Focaldata result that corporate training is the single biggest driver of workplace AI use — and alongside the evidence that a single training session tripled daily usage among women over 55.

There is also a widely misread datapoint worth putting on the record for the executives present: roughly two thirds of companies expect to maintain the size of their workforce¹³⁹. The dominant corporate expectation is not mass reduction. It is reconfiguration — different roles, different skill mixes, different spans of control. That is a harder management problem than a headcount target, which is precisely why so many firms reach for the headcount target instead.

Where agentic deployment does compress structure, the pattern BCG anticipates is specific: significantly fewer frontline employees in transactional processes, with a corresponding reduction in management layers, and managers running smaller teams focused on higher-value work¹⁴⁰. Set that beside the finding in chapter 7 that middle managers are currently overloaded rather than removed, and the sequencing becomes clear — the overload comes first, the restructuring later, and the firms that manage the first well are better placed to make the second deliberate rather than reactive.

What good looks like: six tests

Drawing the strands together, six questions distinguish a serious AI transformation from an expensive one. They are offered as a diagnostic that any organization represented here could run this quarter.

1. Where does your measurement sit in the funnel? If you are reporting drafts generated, tickets handled or code written, you are measuring the top. The number that matters is at the bottom: releases shipped, cases closed, decisions made. The ratio between the two is your absorptive capacity.
2. What proportion of the program budget is going to the 70%? If people, process redesign and role definition are receiving less than half the effort, the benchmark says the returns will disappoint regardless of model choice.
3. Have you redesigned a process end to end, or inserted AI into an existing one? The gap between 60% cost reduction and under 20% appears to run through exactly this distinction.

¹³⁸BCG, “AI Transformation Is a Workforce Transformation”, March 2026.

¹³⁹BCG, “Closing the AI Impact Gap”, 2025.

¹⁴⁰BCG, cited in Forbes, 26 January 2026.

4. Do workers have control over how and when they use the tools? The evidence on deployment quality is unambiguous: where staff were trusted to experiment, outcomes were better; where jobs were redesigned around the strengths and weaknesses of the machine, workers were left plugging its inadequacies.
5. Who is being trained, and who is not? If adoption in your organization mirrors the 62%-versus-16% pattern of the wider labor market, you are importing an inequality problem into your own payroll, and you are leaving most of the available productivity on the table.
6. What happens to your junior intake? If AI is absorbing the tasks through which your firm has always produced its own senior people, you have a reproduction problem that will become visible in five to ten years and cannot be fixed quickly at that point.

None of these six requires a view on whether artificial general intelligence is imminent. They are robust across the range of plausible trajectories — which is the property a decision should have when the underlying forecast is genuinely uncertain.

12. A preparedness agenda: five workstreams

This closing section is deliberately more prescriptive than what precedes it. It is offered not as a set of recommendations from the conveners but as a structured target for disagreement: five workstreams, each with a defensible rationale drawn from the evidence above, each with an identified objection, and each with a question that only the room can answer.

Workstream 1 — Build the instrument panel

The rationale is the least contested proposition in this entire document: governments cannot respond to a transition they cannot see, and French statistical infrastructure — like everyone else's — was designed for a slower economy. Casselman's illustration is that the monthly US jobs report contains no line for technology employment, because the industry categories were fixed decades ago. France faces the same structural problem with the same causes¹⁴¹.

What a serious instrument panel would contain follows directly from the mechanisms identified earlier. Hiring composition rather than volume — seniority mix, credentials required, entry-level vacancy rates — because the evidence says composition moves first¹⁴². The capability overhang, tracked as the gap between realised and theoretical exposure by occupation, because a rapid closing is the clearest early warning of accelerating substitution. Graduate outcomes at monthly rather than annual frequency. Adoption by firm size and by wage decile, since the FT–Focald data evidence shows the divide running between occupations and earnings levels. And the labour share in AI-intensive sectors, which is the ultimate distributional test.

Precedents exist and can be borrowed rather than invented: California's workforce-impact monitor, the OECD's real-time indicators, the Mayor of London's evidence base, and the AI labs' own economic indices — the last to be used critically, since their producers are not disinterested¹⁴³.

¹⁴¹B. Casselman, *The Daily*, The New York Times, 27 July 2026.

¹⁴²"AI is not reducing employment, but rather who gets hired", ProMarket, June 2026.

¹⁴³Office of the Governor of California, June 2026; OECD Employment Outlook 2026; Greater London Authority, Mayor's AI and Jobs Taskforce, July 2026; Anthropic Economic Index; OpenAI economic research.

The objection: measurement is what governments do instead of acting, and an instrument panel can become an alibi for four years of study. The question for the room: what threshold, on which indicator, would trigger which action — decided in advance?

Workstream 2 — Rebuild the first rung

The rationale is the best-supported finding in this document. Entry-level roles are being automated first; expertise has always been built in those roles; AI multiplies existing expertise but does not create it; and junior workers using AI heavily understand their work less¹⁴⁴. The mechanism by which every French profession reproduces itself is being dismantled, quietly, with a lag of a decade before the consequence becomes visible.

What this would require goes beyond training policy as currently constituted. Apprenticeship contracts redesigned around learning content rather than task content, since the tasks that generated learning are exactly those being automated. Explicit obligations or incentives for firms to maintain junior intake — the private return on which has fallen while the social return has not. Extension of professional education into the gap, using the institutions that are demonstrably good at training and assessment. And, uncomfortably, a national position on AI use in education, given evidence that unstructured use harms learning while structured tutoring helps substantially.

The objection: mandating junior hiring against economic logic is exactly the labor-market rigidity that European firms are already criticized for, and it may reduce hiring overall rather than redistribute it¹⁴⁵. The question for the room: who pays for training that used to happen implicitly on the job — the firm that no longer captures its return, the state, or the individual?

Workstream 3 — Close the adoption divide

The rationale is that the divide is measurable, large and tractable. A 46-point gap in daily use between the highest and lowest earners is not a technology problem; a single training session tripling usage among women over 55 is not the signature of an insurmountable barrier¹⁴⁶. Corporate training is the single biggest driver of workplace AI use, which means the lever is known and already in the hands of the employers in this room.

This is the workstream requiring least public money and offering the most immediate return, because it operates on adoption rather than on displacement. It also has a productivity rationale independent of equity: European under-adoption is a live risk, and the fastest route to closing it runs through the two-thirds of the workforce not currently using these tools daily.

The objection: training people to use tools that may be obsolete in eighteen months is a poor investment, and the skill is a moving target with no destination¹⁴⁷. The question for the room: is the object of training the tool, or the judgement required to evaluate what the tool produces?

¹⁴⁴Anthropic research and E. Brynjolfsson et al., cited above.

¹⁴⁵See D. Ball's warning, The New York Times Magazine, 14 June 2026.

¹⁴⁶M. Murgia and J. Burn-Murdoch, Financial Times, 23 April 2026.

¹⁴⁷C. Shih, The New York Times Magazine, 14 June 2026.

Workstream 4 — Decide the distributional question before it decides itself

The rationale is that the sectoral evidence shows employees losing in a majority of industries while consumers, shareholders and compute providers gain; that the labor share is the mechanism through which this becomes a political crisis; and that public opinion has already moved well ahead of the data.

One recent estimate makes the arithmetic concrete, and it is more encouraging than the rhetoric on either side suggests. In the most disruptive of the three scenarios described in chapter 1, the economy is a third larger by 2030, but office and professional workers earn less and more of them are out of work. Their combined losses come to roughly a third of what the economy gains. A transfer of about 9% of GDP — comparable in size to Social Security and Medicare combined — would restore their incomes in full and still leave everyone else more than 20% better off¹⁴⁸.

The authors then add the sentence that matters most for this room: transfers on that scale in response to technological change have no precedent, and the experience of past disruptions is that they do not happen on their own. The resources exist in every scenario. Whether they reach the people who bear the cost is not something growth delivers by itself.

The options on the table are more radical than French debate has yet absorbed, and several originate with the AI firms themselves: sovereign or citizen funds capitalized by contributions from the firms that benefit most; universal basic capital, granting every child a permanent claim on capital income; taxation of automation, of capital, or of data rather than of labor alone; and a shorter working week. Each has a serious objection — capital mobility, innovation deterrence, regulatory capture where public finances become bound to AI valuations¹⁴⁹.

What the polling suggests about feasibility should inform the choice. Universal basic income polls badly; a jobs guarantee polls well; populist framing on who profits outperforms technocratic framing substantially¹⁵⁰. A proposal that is economically elegant and politically unsellable is not a proposal.

The objection: any of these mechanisms requires international coordination that does not exist, given the mobility of capital and of firms. The question for the room: which of these could France do alone, which requires Europe, and which is simply not available?

Workstream 5 — Choose a direction for the technology, not only a response to it

The rationale is Acemoglu, Autor and Johnson's: the direction of technological development is a choice, currently made by a small number of firms whose incentives point toward automation rather than augmentation, and the resulting under-investment in pro-worker AI is a market failure rather than a preference¹⁵¹.

The levers are unfamiliar in a labor-policy context because they are not labor policy: competition enforcement to open the market to firms with different business models; legal protection of professional expertise absorbed into training data, on the SAG-AFTRA precedent; institutionalized worker voice in deployment decisions; and public procurement used deliberately, since the state is among the largest buyers of these systems in France and could specify augmentation rather than substitution.

¹⁴⁸ A. Korinek, C. I. Jones, S. Sacher, T. Cotter and P. McCrory, "Economic Scenarios for Transformative AI", The Anthropic Institute Working Paper No. 2026-02, September 2026. The accompanying scenario explorer is at anthropic.com/institute/econ-scenarios. See Section 4.3, "Who gains and who loses".

¹⁴⁹ "State-owned AI", Unhedged, Financial Times, 11 June 2026; OpenAI, "Industrial Policy for the Intelligence Age"; D. Autor and N. Thompson, Digitalist Papers, vol. 2.

¹⁵⁰ D. Shor polling, reported in The New York Times, 30 April 2026.

¹⁵¹ D. Acemoglu, D. Autor and S. Johnson, Building Pro-Worker Artificial Intelligence, Brookings / NBER WP 34854, 2026.

This is also where the physical frontier argument belongs. If Europe's advantage lies in embodied AI and industrial applications — addressing labor shortages rather than creating surpluses — then industrial and labor policy point the same way for once, and the case for public investment is unusually strong¹⁵².

The objection: steering technology invites regulatory capture, and the record of states picking technological directions is not encouraging. The question for the room: is public procurement a sufficient lever, or does anything meaningful require regulation?

Each of these five workstreams has a serious objection. A day that produced agreement on all five would have failed; a day that established which objections are decisive would have succeeded.

What would count as success

A realistic ambition for one day is not a program but a shared map: agreement on what the evidence says, clarity on which disagreements better measurement could settle, and identification of the two or three choices that cannot be delegated to firms, to social partners or to time.

The synthesis circulated afterwards will record convergence where it exists and, more usefully, the questions left open. If it records only convergence, the day will have been too comfortable to be worth the time of the people in the room.

¹⁵²T. Parikh, Financial Times, 28 June 2026.

AI and productivity: rethinking workforce training and social cohesion

Philippe Tibi · Professor of Strategy and Finance at École Polytechnique (IP Paris), Founder of Pergamon Campus

Eric Hazan · Professor of Digital Strategy at Sciences Po, Co-Founder of Ardabelle Capital, Senior Partner Emeritus at McKinsey & Company

First published in Polytechnique Insights, 14 April 2026.

Key takeaways

- AI has the potential to improve productivity without reducing workload: the actual gains are reinvested in increased activity rather than in freed-up time.
- The same technology can either enhance human capabilities or replace them, depending on how it is deployed.
- Young graduates are paying a heavy price as junior roles are the first to disappear, severing the traditional chain of transmission of professional expertise.
- A Marshall Plan for training is urgently needed. Curriculum reform, mass retraining and investment in human capital on a par with AI infrastructure.
- Like the “Engels pause”, AI could concentrate its benefits among the few — unless we invent new forms of solidarity and deliberately direct progress towards workers.

Rarely has a technology simultaneously sparked so much enthusiasm and so much anxiety at the same time. On the one hand, promise of productivity gains unseen since the 1970s; on the other, the specter of dwindling job prospects for an entire generation of young graduates. But these two realities are not contradictory. Together, they contribute to a fundamental contemporary equation: AI is an amplifier—of value for those who prepare for it, of inequality for those left behind. To fully grasp this challenge, we must follow the thread that runs from productivity to employment, from employment to business organization, and then from organization to a collective ambition for training that matches the scale of the challenge — the journey this article sets out to take.

Productivity: a promise beginning to materialize

Let’s start with some good news: signs of a real productivity surge are beginning to mount. In January 2026, economists at Apollo Global Management, led by Torsten Slok, identified tangible increases in DevOps, process automation and support functions¹⁵³. The most comprehensive summary available to date is the Early Signals of AI Impact dashboard, which aggregates more than 460 sources in real time — academic studies, sector reports, field data — across 18 indicators of AI-driven labor transformation: individual productivity, recruitment trends, shifts in required skills, team restructuring. Its consensus assessment is clear: the adoption of AI tools is accelerating across all industries, productivity measured in pioneering companies is rising, and roles involving repetitive tasks are being reduced. This crucial nuance can be summed up in a single word: jobs are changing faster than they are disappearing, and no aggregate macroeconomic shift is yet measurable. It is precisely this ‘yet’ that lends the issue of preparedness its sense of urgency¹⁵⁴. But AI does more than just replace tasks: it

¹⁵³Torsten Slok, Rajvi Shah and Shruti Galwankar, Quantifying the Productivity Gains from AI Adoption, Apollo Global Management, January 2026.

¹⁵⁴Early Signals of AI Impact, real-time dashboard (jobsdata.ai), March 2026. Aggregates 460+ sources across 18 indicators.

transforms the very pace of innovation. Bontadini, Haskel and their co-authors describe it as a ‘meta-innovation’ — an innovation in the method of innovation — with multiplier effects on total factor productivity that the US economy is only just beginning to feel, whilst Europe is lagging behind in adoption to a worrying extent¹⁵⁵.

AI does more than just replace tasks: it transforms the very pace of innovation.

Nevertheless, two caveats are in order. The first relates to Moravec’s paradox: AI excels at cognitively complex tasks — writing, summarizing, coding — but remains clumsy at the contextual manual tasks that seem easy to us. This discrepancy partly explains the gap between the touted potential and actual adoption in the economy¹⁵⁶. The second caveat is even more counter-intuitive: according to a Harvard Business Review study published in February 2026, AI does not lighten the workload — it intensifies it. Employees who adopt AI tools handle more cases, respond to more requests and take on new tasks. The productivity gain is real, but it does not automatically translate into freed-up time: it is reinvested in additional work¹⁵⁷.

Another key finding from the research by the NBER and MIT is that AI can boost workers’ productivity without necessarily automating entire tasks. This distinction between augmentation and automation is not merely semantic — it determines the full range of the technology’s social impacts. In augmentation mode, AI becomes a cognitive prosthesis: it extends the individual’s capabilities, enabling them to handle greater complexity, produce faster, and better document their decisions. Jobs are preserved and often enhanced. In automation mode, it replaces human activity itself: jobs may disappear or undergo radical transformation. Both modes coexist in most organizations — sometimes within the same role, depending on the tasks and the circumstances¹⁵⁸.

Acemoglu, Autor and Johnson refine this framework further by distinguishing five categories of technology, two of which deserve particular attention. ‘Expertise-levelling’ technologies enable non-experts to perform tasks previously reserved for specialists — medical AI for diagnostic support is one example. ‘New task creation’ technologies, on the other hand, are the only ones that are ‘unambiguously favorable to workers’: they generate a net demand for human expertise by opening up activities that did not previously exist. It is in this latter category that the most solid promise of AI lies — though it is also the most demanding in terms of training¹⁵⁹.

The impact on jobs: who are the canaries in the coal mine?

The metaphor chosen by Erik Brynjolfsson for his summer 2025 study is striking: in coal mines, canaries were sent in to detect toxic gases before the miners entered. Young workers in sectors exposed to AI are, today, those canaries¹⁶⁰.

The ADP data — the most granular ever used for this type of analysis — speaks for itself: in the United States, between late 2022 and July 2025, employment among 22- to 25-year-olds in roles highly exposed to AI fell by 6%. For young software developers, the decline reached 20% compared to the peak in November 2022. Entry-

¹⁵⁵ Filippo Bontadini, Carol Corrado, Jonathan Haskel and Cecilia Jona-Lasinio, “AI as an Innovation in the Method of Innovation: Implications for Productivity Growth”, AEA Papers and Proceedings, vol. 116, 2026, pp. 36–40 (first circulated SPRU/Imperial, October 2025).

¹⁵⁶ Anjana Susarla, A Gap in AI Adoption? Moravec and the AI Productivity Paradox, Forbes, January 2026.

¹⁵⁷ Harvard Business Review, AI Doesn’t Reduce Work — It Intensifies It, February 2026.

¹⁵⁸ Acemoglu, Autor, Johnson, op. cit.; MIT Sloan, Pro-Worker AI Doesn’t Just Happen, 2026.

¹⁵⁹ Acemoglu, Autor, Johnson, op. cit. The authors distinguish five categories of technology (labour-augmenting, capital-augmenting, automating, expertise-levelling, new-task-creating). Only the last is “unambiguously pro-worker”.

¹⁶⁰ Erik Brynjolfsson, Bharat Chandar, Ruyu Chen, Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of AI, Stanford Digital Economy Lab, August 2025.

level customer support roles also lost nearly 11% of their junior staff. It is the entry-level positions that are disappearing first¹⁶¹.

This assessment echoes the findings of the Financial Times' investigation into the 'great graduate jobs drought': entire cohorts are graduating from top universities and struggling to find roles commensurate with their qualifications, not because AI has replaced them, but because it has rendered the foundational learning experience provided by junior roles redundant¹⁶². The consequence is perverse and deserves our attention: if young people no longer acquire their expertise on the job in these first roles, the chain of transmission of professional expertise is broken. AI can amplify expertise — but it cannot create it from scratch. If the generation entering the labour market today does not build its foundations, who will train the experts of 2035?

However, we must not confuse exposure with net destruction — and indeed question the causality. The author himself points out that the slowdown in hiring in jobs exposed to AI began in spring 2022, i.e. before the release of ChatGPT in November 2022; the Fed's rise in key interest rates is at least as plausible an explanation¹⁶³. Even more surprising, Yale's Budget Lab observes no significant difference between employment levels in the jobs most exposed to AI and those least exposed — "the lines look the same", in the words of its chair. What the exposure metrics do capture with certainty is a decline in job vacancies in the 40% of occupations most exposed since the rise of ChatGPT — which is not equivalent to net job losses. The economy as a whole continues to create jobs; it is their structure that is changing, and with it the nature of the skills required¹⁶⁴.

In France, the picture is similar. Projections for 2026–2030 identify several socio-professional categories that are particularly vulnerable: intermediate administrative professions, certain support function managers, and skilled office workers. France is not immune due to its social model — it is exposed, like its partners, to an accelerated restructuring of its labor market¹⁶⁵.

David Autor offers a more nuanced and, ultimately, more optimistic perspective: AI could restore value to skilled labor by amplifying what humans do best — judgement, relationships and contextual creativity. It could even be the catalyst for a revival of skilled middle-class jobs, provided that the technology is geared towards augmenting human capabilities rather than simply replacing them. But this prospect requires active policy measures. It will not materialize spontaneously¹⁶⁶.

Businesses and AI: between hesitant adoption and profound transformation

While the macroeconomic effects remain to be consolidated, the effects of AI on the internal organization of businesses are already noticeable. The first lesson from the most rigorous empirical studies is that results depend far less on the technology itself than on the organizational choices that accompany it. Two companies equipped with the same tools can achieve radically different results depending on whether they have redesigned their processes, trained their teams and redefined roles — or not¹⁶⁷.

¹⁶¹ Ibid. The ADP data cover several tens of millions of US employees — the most granular empirical base available on AI's impact on young workers' employment.

¹⁶² Financial Times, The Great Graduate Job Drought, 2026.

¹⁶³ Budget Lab at Yale, Labor Market AI Exposure: What Do We Know?, February 2026; Gad Levanon, AI Exposure and Labor Market Change, LinkedIn Pulse, 2026.

¹⁶⁴ David Autor, Anton Korinek, Natasha Sarin, interview by David Leonhardt, "What Economists Really Think About A.I. and Jobs", The New York Times, 4 February 2026.

¹⁶⁵ Exposition à l'IA générative et emploi : application à la classification socio-professionnelle française, 2024; Zety, Les métiers en France face au risque de l'IA (2026–30).

¹⁶⁶ David Autor, AI and the Future of Work, Issues.org; NBER working paper w32140, Applying AI to Rebuild Middle Class Jobs.

¹⁶⁷ MIT Initiative on the Digital Economy, Effective AI in the Workplace: What the Research Shows, February 2026; SYPartners, Still Loading: Humans at Work in the AI Age, 2026.

This tension between automation and augmentation lies at the heart of the Beyond the Binary report: the two dynamics coexist and combine. The same technology can automate routine tasks whilst enhancing the capacity to focus on high-value-added tasks. Empirically, three adoption profiles have been identified: ‘cyborgs’, who seamlessly integrate AI into their daily workflow; ‘centaurs’, who alternate between working autonomously, and delegating to AI depending on the nature of the tasks; and ‘self-automators’, who program AI to take over entire portions of their work. These three modes do not have the same implications for the skills required, the risks of obsolescence or training needs. The question is therefore not “does AI replace?” but “how is it deployed, by whom, for what purposes, and with what safeguards?”¹⁶⁸

Companies which have successfully deployed AI that enhances their employees’ skills have done so deliberately, refusing to let the logic of automation alone dictate architectural choices.

One category under particular scrutiny is that of middle managers. AI brings to light hierarchical layers whose added value—information sharing, coordination, reporting—is precisely what automated systems can now handle¹⁶⁹. But the managerial role is not limited to the flow of information: it includes talent development, conflict management, contextual interpretation, and trust in teams—skills that the McKinsey Global Institute identifies as the big winners of the AI era, for which demand is growing precisely because LLMs cannot replicate them. The manager of tomorrow is not the one who filters information — it is the one who creates meaning¹⁷⁰.

Finally — and this is perhaps the most important lesson — AI that benefits workers is not built by default. Research from MIT Sloan shows that companies which have successfully deployed AI that enhances their employees’ skills have done so deliberately, refusing to let the logic of automation alone dictate architectural choices. This is not an abstract moral imperative: it is a prerequisite for sustainable performance, in a context where talent retention and internal trust are becoming decisive competitive advantages¹⁷¹. Acemoglu, Autor and Johnson go further: it is not merely a matter of corporate choice, but a documented market failure. Three cumulative factors are structurally driving the industry towards automation: path dependence (large firms have built their business models on automation tools from which they derive their profitability); the dominant ideology in research laboratories, which is oriented towards AGI and largely insensitive to labor effects; and the growing concentration of the sector, which stifles alternative models. “We are not currently on the pro-worker AI path,” states Simon Johnson. This failure calls for responses that go beyond the goodwill of companies¹⁷².

For a Marshall Plan for training

If AI is reducing entry-level roles, it is undermining the main mechanism through which companies have always passed on their human capital: on-the-job learning. All types of human-AI interaction — from ‘cyborgs’ to ‘centaurs’ and ‘self-automators’ — require prior expertise that AI serves to amplify. But this expertise is built up in the early years of a professional career, in junior roles which are precisely the ones that are disappearing. The

¹⁶⁸ Matt Sigelman et al., Beyond the Binary: How Automation and Augmentation Are Combining to Reshape Work, January 2026.

¹⁶⁹ Le Figaro Décideurs, IA, recherche d’efficacité : les managers intermédiaires sont-ils une espèce en voie de disparition ?, February 2026.

¹⁷⁰ McKinsey Global Institute, Human Skills Will Matter More Than Ever in the Age of AI, 2025; Eric Hazan, Anu Madgavkar, Michael Chui et al., A New Future of Work: The Race to Deploy AI and Raise Skills in Europe and Beyond, MGI, 2024.

¹⁷¹ MIT Sloan, Pro-Worker AI Doesn’t Just Happen: Companies Need to Act, 2026; Daron Acemoglu, David Autor, Simon Johnson, Building Pro-Worker Artificial Intelligence, The Hamilton Project / Brookings, NBER Working Paper 34854, February 2026.

¹⁷² Acemoglu, Autor, Johnson, op. cit.

systemic risk is real: by short-circuiting early-career learning, AI risks creating a generation of professionals without solid foundations — buildings without a ground floor¹⁷³.

This assessment calls for a response on the scale of the Marshall Plan. The analogy is not merely rhetorical: the 1948 plan was a collective and structured response to the massive destruction of physical capital. AI is causing the accelerated destruction of human capital — not through violence, but through obsolescence. The Oxford research on technological unemployment serves as a useful reminder that every major wave of automation has brought about unexpected adaptations; yet no previous wave has had this pace or this cross-cutting impact across the entire spectrum of cognitive functions, from the bottom of the ladder to the most highly skilled professions¹⁷⁴.

Skills have an increasingly short lifespan and continuing professional development systems were designed for a world where obsolescence was measured in decades, not years.

Three areas of work are required in parallel. The first is curriculum reform: from primary education to the grandes écoles, we must integrate not only skills for interacting with AI, but above all the abilities that LLMs cannot replicate — critical judgement, reasoning, contextual creativity, and relational intelligence. Research on K-12 skills in the AI era shows that this reform is feasible on a large scale if we commit the necessary resources to prioritizing it¹⁷⁵.

The second priority is a massive retraining program, targeting the 40% of occupations most at risk. Granular data from Yale's Budget Lab enables these interventions to be prioritized with unprecedented precision. The challenge is to develop short, certified pathways accessible to those currently in employment, with tripartite governance involving the state, professional bodies and businesses. Speed is crucial here: skills have an increasingly short lifespan and continuing professional development systems were designed for a world where obsolescence was measured in decades, not years.

The third area of focus is financial. Companies are investing hundreds of billions in AI infrastructure — data centers, chips, foundational models. A comparable investment in human capital is not merely a moral obligation: it is the prerequisite for the model's social sustainability. The economic scenarios for transformative AI developed in February 2026 in London, as part of a forward-looking exercise bringing together economists and policymakers around the Windfall Trust and Gresham College, illustrate this precisely: in all trajectories where productivity gains are widely distributed, investment in training has consistently preceded — or accompanied — technological deployment. The opposite is equally true: where training has been deferred, the benefits of AI have been concentrated among a few players, widening inequalities rather than reducing them¹⁷⁶.

These three areas of focus — curricula, reskilling, and funding — are necessary. But they are not enough unless the root cause is addressed: the direction of AI itself. Acemoglu, Autor and Johnson identify three complementary policy levers. The first is the application of competition law to break up concentration in the sector and open the market to new entrants with business models less focused on pure automation. The second

¹⁷³Steven Randazzo (Warwick/Harvard Business School WP 26–036), *Cyborgs, Centaurs and Self-Automators*, 2025.

¹⁷⁴Anselm Küsters and Benjamin Schneider, *What is Technological Unemployment?*, Oxford Economic and Social History Working Paper no. 218, March 2025.

¹⁷⁵*Which Skills Matter Now? A Data-Driven Framework for K-12 in the Age of AI*, V5.1, 2025.

¹⁷⁶Danny Buerkli (Windfall Trust) and Daniel Susskind (Gresham College, Oxford), *Securing humanity's AI future — Economic Scenarios for Transformative AI*, London, 5 February 2026.

is the legal protection of professional expertise: workers whose know-how is absorbed by AI systems should have rights over this contribution, as a safeguard against ‘theft of expertise’¹⁷⁷. The third is the institutionalization of workers’ voices in deployment decisions — at both company level and in sectoral regulation. Training humans to adapt to AI is essential; orienting AI towards humans is just as urgent¹⁷⁸.

In the face of social upheaval, creating new forms of solidarity

This training plan is therefore essential for those whose skills are threatened by artificial intelligence. However, history teaches us that industrial revolutions disrupt traditional social frameworks. When they are massive, productivity gains are not merely a statistical fact; they herald a profound structural transformation.

In the 19th century, the harnessing of mechanical and thermal power brutally drove peasants from their fields and displaced artisans in favor of factories operated by low-skilled workers. This process created an urban proletariat plunged into poverty by the glut of labor supply. Despite a spectacular rise in GDP, workers’ living conditions stagnated for nearly half a century. This is the ‘Engels’ pause’¹⁷⁹, the period during which the industrial elite captured the bulk of the wealth generated by productivity gains. This upheaval was also characterized by a devastating loss of social bearings: forced displacement, the breakdown of family solidarity, the erosion of traditional workplace hierarchies, and subsistence wages. The economy was thus ‘disembedded’ from society, to use Polanyi’s term¹⁸⁰. The intellectual and political response was vigorous: from Karl Marx’s *Capital* (1867) to the formation of revolutionary and reformist movements, a balance of power developed between capital and labor, gradually leading to modern social security, without the ‘great revolution’ that had been predicted.

This historical overview is, of course, simplified. The slow construction of the social contract was neither natural nor peaceful. It took place under the brutal pressure of political and geopolitical upheavals: colonial expansion, the clash of imperialisms, the Soviet revolution and the world wars. The social contract also long relied on the emergence of alternative jobs in the service sector, capable of absorbing the workforce rendered redundant in agriculture and industry. Douglass North emphasizes that the survival of such a system depends on the evolution of its institutions to manage these transitions¹⁸¹.

AI could bring about changes comparable to those of previous industrial revolutions.

Because it is cognitive, the AI revolution will spread more rapidly than that of the engine — which makes intervention all the more urgent: a 60-year ‘Engels pause’ would be unacceptable in the 21st century, even if shortened. In its physical applications, AI relies on existing robot architectures or those that can be reprogrammed within a relatively short timeframe. Like electricity, AI is a cross-cutting technology capable of massively increasing productivity across all sectors. It is all a question of the pace of deployment and real-world impact, but the risk exists and must be analyzed and anticipated. AI could bring about changes comparable to those of previous industrial revolutions. It is more a question of time, particularly in a Europe that has underinvested in technology and R&D for over 10 years. Furthermore, for the past 20 years, the internet revolution has concentrated profits among a small number of players capable of investing heavily in industrial

¹⁷⁷Collective compensation agreements already exist in this area: SAG-AFTRA obtained in 2023, after a 118-day strike, contractual commitments from studios to obtain actors’ consent and compensate them before any use of their image or voice for AI training.

¹⁷⁸Acemoglu, Autor, Johnson, op. cit. Nine policy directions analysed: antitrust, protection of workers’ expertise, institutionalisation of their voice, sectoral focus on health and education.

¹⁷⁹Robert C. Allen (2009), “Engels’ Pause: A Pessimist’s Guide to the British Industrial Revolution”, *Explorations in Economic History*, vol. 46.

¹⁸⁰Karl Polanyi, *The Great Transformation* (1944).

¹⁸¹North, D. C. (1990). *Institutions, Institutional Change and Economic Performance*. Cambridge University Press.

and human capital. This accumulation of capital enables American Big Tech to position itself at the forefront of AI, by acquiring rival start-ups, recruiting the best talent and securing the electrical resources needed for deployment. These strategies are reminiscent of those of the oil oligopolies of the second industrial revolution, which were ultimately dismantled at the start of the 20th century.

Faced with the risk of being left behind, we struggle to envisage future sectors to fall back on. Although an ageing population, education, leisure and healthcare are powerful drivers of employment, they are currently poorly paid and unattractive. At the same time, technology is reigniting geopolitical tensions: it is, for example, one of the root causes of Sino-American friction. We must therefore prepare for a dual revolution — the internal impact of technology on our production systems and the external effects of a new distribution of wealth. As Acemoglu and Johnson point out¹⁸², following on from Korinek and Stiglitz¹⁸³, adjustment is unlikely to occur spontaneously and requires us to ‘steer’ technical progress. This position is contested by those who fear the risks of regulatory capture; nevertheless, it raises an interesting question that we probably need to anticipate.

It is therefore prudent to work on innovative paradigms, particularly in terms of institutions and financing. The options frequently proposed all represent a break with existing systems: universal basic income, free access to public services, and a deliberate reduction in productivity through shorter working hours. It is too early to take a definitive stance, but these programs still need to be financed. The issue appears straightforward, since rising productivity increases the pool of wealth to be distributed. There is talk of taxes on robots¹⁸⁴ and productive capital, or even on data, since taxing labor alone will be insufficient¹⁸⁵. Or the creation of a very large social protection fund, managed like a pension fund and funded by contributions from new companies, in return for the provision of the social infrastructure to which they owe their success — modelled on the Alaska Permanent Fund, which redistributes a fraction of the state’s oil revenues to citizens, or the academic proposals for a ‘citizen fund’ developed by John Roemer and taken up by Daron Acemoglu. In the same vein, David Autor and Neil Thompson (MIT) propose a ‘Universal Basic Capital’: every child would receive a portfolio of productive assets at birth, granting them a permanent right to capital income — including income from automation — without relying on politically vulnerable budgetary transfers. Further ideas are needed, as the sums mentioned here do not measure up to the challenges at stake. Moreover, these proposals require a global consensus due to the mobility of capital and entrepreneurs. This debate must take place before undesirable solutions are imposed on us by developments in the labor market where we would be ‘acted upon’ rather than ‘acting’.

The economist David Autor puts it this way: AI could rebuild the middle class by amplifying what humans do best. Brynjolfsson reminds us, with data to back it up, that for the time being, it is the youngest who are paying the price of the transition. Between these two truths, there is no contradiction — there is an agenda.

AI is an amplifier. It amplifies the productivity of organizations that prepare for it; it amplifies inequalities in those that do not. It can amplify the value of human judgement; it can also amplify the vulnerability of those who lack access to training. The real question is not a technological one. It is a political one: do we have the collective will to make training as urgent a priority as investment in digital infrastructure? And to find creative and realistic solutions to fund the transformation of the system. The answer to this question will determine whether the AI dividend benefits a few — or everyone.

¹⁸² Acemoglu, D., & Johnson, S. (2023). *Power and Progress*. PublicAffairs.

¹⁸³ Korinek, A., & Stiglitz, J. E. (2021). “Steering Technological Progress”, AEA Papers and Proceedings, vol. 111.

¹⁸⁴ Robert Shiller, “Robotization of the Workforce and the Need for a Robot Tax”, 2017; Bill Gates, Quartz interview, February 2017.

¹⁸⁵ Korinek & Stiglitz (2021), op. cit.

Four exhibits accompanying the column




Exhibit 1 — Canaries in the coal mine: junior employment trends, 2022–2025. Source: Brynjolfsson, Chandar and Chen, Stanford Digital Economy Lab (2025), ADP data; chart prepared for Polytechnique Insights, March 2026.

2

Five Types of Technologies — Impacts on Work (Acemoglu, Autor, Johnson 2026)

Theoretical Framework — NBER / United States

Category	Definition	AI Example	Employment Impact	Skills Impact
Labor-Augmenting	Makes the worker more efficient at current tasks	Copilot, AI writing assistants	✓ Preserved	↑ Enhanced
Capital-Augmenting	Improves existing machines and algorithms	Automated process optimization	= Neutral	= Neutral
Automating	Replaces humans with machines for entire tasks	RPA, doc processing, automated support	✗ Reduced	↓ Devalued
Expertise-Leveling	Enables non-experts to perform expert tasks	AI-assisted medical, legal diagnosis	↔ Mixed	↔ Redistributed
Creating New Tasks ★	Opens unprecedented activities that did not exist before	AI governance, prompt engineering, algorithmic ethics	✓✓ Creates	↑↑ New

★ Only category “unambiguously favorable to workers” — the only one that generates net demand for human expertise.


Source: Acemoglu, Autor, Johnson — Building Pro-Worker Artificial Intelligence, Hamilton Project / Brookings, NBER WP 34854, Feb. 2026.

Exhibit 2 — Five types of technologies and their impacts on work. Source: Acemoglu, Autor and Johnson, Hamilton Project / Brookings (2026); chart prepared for Polytechnique Insights, March 2026.

3

AI Exposure by Socio-Professional Category in France (Composite Score, 2025-2030)

France



Source: France Stratégie, Les Métiers en 2030 (2021); OECD, Employment Outlook 2023; INSEE, AI, Growth and Employment, Aghion et al. (2019); WEF, Future of Jobs Report 2025. Scores synthesized by the authors.


Exhibit 3 — AI exposure by socio-professional category in France, composite score 2025–2030. Source: France Stratégie, OECD, INSEE and WEF; scores synthesised by the authors for Polytechnique Insights, March 2026.

4

Two Trajectories: Proactive vs Delayed Training (2025-2035)

United Kingdom — Windfall Trust / Gresham

Productivity gains concentration index (base 100 in 2025) — forward-looking scenarios



Scenario A — Proactive Training

- Broadly distributed productivity gains
- Stable inequality (+7 pts by 2035)
- Strengthened middle class

Scenario B — Delayed Training

- Concentrated gains (top 20%)
- Sharply rising inequality (+28 pts)
- Accelerated social divide

Source: Buerkli (Windfall Trust) & Susskind (Gresham College) — Securing Humanity's AI Future, London, Feb. 5, 2026. Forward-looking scenarios — indicative values.

Exhibit 4 — Two trajectories: proactive versus delayed training. Source: Windfall Trust / Susskind scenario work (2026); chart prepared for Polytechnique Insights, March 2026.