Europe's AI Innovation Wave: Who Gains, Who Decides, Who Governs?*

Tato Khundadze† Willi Semmler‡ Ekkehard Ernst§

August 18, 2026

Abstract

This paper asks how Europe should respond to the current wave of artificial intelligence. It argues that AI is best understood, at least for now, as a rapidly developing cluster of technologies within the wider information revolution rather than as a separate industrial revolution. The evidence on its economic effects remains unsettled. Productivity estimates vary widely, investment has moved ahead of what can yet be justified by measured gains, and the labour-market evidence points more towards the reorganisation of work than towards mass unemployment. These outcomes will depend less on the technology alone than on how firms deploy it, how the gains are distributed and what institutions shape adoption. Europe enters this transition from a weak position in frontier models, compute, cloud infrastructure and private investment. A credible response therefore requires more than regulation. It should combine conditional support for European firms, public investment in data and compute, stronger administrative capacity, labour force participation in deployment decisions and procurement rules designed to limit long-term dependence on foreign platforms. The paper also argues that part of any productivity gain should reach the workforce, through shorter working time and wage growth, rather than being captured only as lower labour costs and short-term capital gains. The central claim is that Europe's AI future will be decided not only by technical progress, but by ownership, institutional design and political choice.

Keywords: artificial intelligence, technological change, technological revolutions, industrial policy, productivity, labour markets, European economic sovereignty

---

*A shorter German-language version of this article appeared in *Makronom* as "Künstliche Intelligenz: Europa wird nicht gewinnen, wenn es nur Silicon Valley imitieren will."

†Co-Founder and Chief Executive Officer, SEDAI Center

‡The New School for Social Research, New York, USA; University of Bielefeld, Germany; International Institute for Applied Systems Analysis (IIASA), Laxenburg, Austria, and Co-founder of SEDAI Center

§Chief Macroeconomist at the International Labour Organization

1 Introduction

In the autumn of 1830, agricultural labourers across southern England destroyed threshing machines under cover of night and signed their warnings with a spectral name: "Captain Swing." Hobsbawm and Rudé interpreted the episode not as superstition against progress, but as a political response to a machine that removed winter work just as rural wages were being driven towards destitution. "The real name of King Ludd," they wrote, "was Swing" [Hobsbawm and Rudé, 2001, p. 298]. The conflict was not simply between workers and technology. It concerned who controlled the new technology, who received its gains and who bore the cost of adjustment.

Two centuries later, artificial intelligence has brought the same questions back in a new form. The debate is often framed in technical terms, including model capabilities, scaling laws, benchmark performance and the prospect of artificial general intelligence. Yet the central questions are economic and political. Does the present wave amount to a new industrial revolution, or is it better understood as a major technology system developing within the longer information revolution? How large are its likely effects on productivity, investment and employment? Who will capture the gains, and who will face displacement, weaker bargaining power or declining job quality? Finally, what strategy is available to governments, and to Europe in particular, given its weaker position in frontier models, cloud infrastructure, semiconductor capacity and private investment?

A further issue cuts across all four questions. AI has become a matter of economic sovereignty and national security. Model APIs are no longer peripheral software services. They are increasingly integrated into firms, public administration, research, defence and critical infrastructure. Dependence on external providers therefore creates a different kind of vulnerability from ordinary import dependence. Access may be restricted, prices may change, technical standards may be set abroad and providers may move into the same markets as their customers. In January 2025, the United States introduced export-control requirements covering certain advanced AI model weights, only for the framework to be rescinded by the incoming administration within months[Bureau of Industry and Security, U.S. Department of Commerce, 2025]. The broader lesson is that the conditions of access are set elsewhere and may shift with domestic politics in another country.

This paper is a critical, policy-oriented literature review. It brings together the neo-Schumpeterian literature on technological change with recent empirical research and policy reports on AI investment, productivity, employment and European industrial strategy. Its purpose is not to provide a systematic review of every strand of the literature, but to use the available evidence to assess the character of the current AI wave and to derive policy implications for Europe.

The argument is threefold. First, AI is at present better understood as a rapidly developing technology system within the fifth information and communications technology surge than as a separate technological revolution [Perez, 2010]. It may become more than that, but the economy-wide reorganisation required by the stronger claim has not yet been demonstrated. Second, the scale of its economic effects remains uncertain. Estimates of productivity gains vary widely, while the labour-market literature points so far more towards the reorganisation of work than towards generalised technological unemployment [Acemoglu, 2025, Aghion and Bunel, 2024, Merola et al., 2026]. These outcomes are not determined by the technology alone. They depend on how firms deploy it, how the resulting surplus is distributed and what institutions shape the transition.

Third, Europe cannot rely on market forces alone to close its technological gap. Its weakness is not limited to the number of frontier models it produces. It also lies in fragmented capital markets, dependence on foreign cloud and compute providers, weak coordination between Member States and limited administrative capacity to guide adoption [Draghi, 2024]. A European response must therefore go beyond regulation. It should combine conditional support for European champions, public investment in data and compute, stronger public capacity to plan and evaluate adoption, workforce participation in deployment decisions and procurement rules that reduce long-term platform dependence. The distribution of productivity gains should also form part of the strategy, including the possibility of converting part of those gains into shorter working time rather than job losses.

The paper proceeds in four steps. The first section places AI within the neo-Schumpeterian literature on technological revolutions, technology systems and general-purpose technologies. The second reviews the evidence on investment, productivity and employment. The third examines Europe's strategic disadvantage relative to the United States and China. The final section develops the main elements of a European policy response. The central claim throughout is that the future of AI will not be decided by technical capability alone. It will also be decided by ownership, institutional design and political choice.

2 A new industrial revolution?

Whether the current wave of artificial intelligence constitutes a technological revolution can be approached through neo-Schumpeterian analysis, which classifies technical change by defining the object in question, locating it on its trajectory, mapping the industries that carry it, and applying the criteria that distinguish a technological revolution from a technology system. In the Freeman and Perez taxonomy, a *technology system* is a constellation of technically and economically interrelated innovations that affects several branches of the economy and gives rise to new sectors, while a technological revolution is a cluster of such systems [Freeman and Perez, 1988]. Applied to AI, this analysis indicates that, at its present stage of development, the wave is better understood as a new technology system within the information revolution than as a revolution in its own right.

The classification depends first on the distinction between invention and innovation. Schumpeter distinguished innovation, the commercial introduction of a new product or "new combination", from invention, which remains within the realm of science and technology (Schumpeter, 1911, pp. 132–6, as cited in Perez, 2010, p. 186). On this distinction the transformer architecture and the underlying training methods are inventions; they become innovations once they are trained at scale, commercialised, and integrated into products such as assistants, coding tools, search and enterprise software. AI is thus wider than any single model class. Generative AI is, at present, its flagship sub-technology and the one with the widest product reach, with large language models as its most visible product class, but it is the visible expression of the system rather than the system itself.

Locating that technology on its trajectory reinforces the point. In the innovation literature, a radical innovation establishes a paradigm within which technical change proceeds as rapid incremental improvement along a trajectory, before the industry passes from its Schumpeterian phase of emergence into maturity; once a dominant design has emerged, process innovations tend to overtake product innovations [Abernathy and Utterback, 1978]. Figure 1 represents this life cycle in logistic form. Generative AI is past the initial inflection and in the phase of rapid incremental improvement, in which scaling, instruction tuning, longer context windows and multimodality extend a trajectory opened by the original breakthrough. The recent movement toward cheaper inference and smaller, distilled models corresponds to the shift from product to process innovation that this literature describes, and it is characteristic of a technology maturing within an existing system rather than one inaugurating a new one.

The industrial structure of the wave can be examined using the distinction between the core industries of a revolution: the motive branches, which supply a cheap and pervasive input; the carrier branches, which are the visible products that bring the new possibilities into use; and the infrastructure, which extends the reach of other industries [Perez, 2010]. Table 1 maps these categories onto AI, with compute and foundation-model provision as the motive layer, assistants, copilots, agents and AI-native software as the carriers, and data centres, cloud platforms and model APIs as the infrastructure. What the mapping reveals is how much of the structure is inherited. Much of AI's infrastructure, including the cloud, the internet and the data pipelines, was built by the information revolution rather than by AI, and its key input remains the microprocessor in a specialised form. On this evidence the

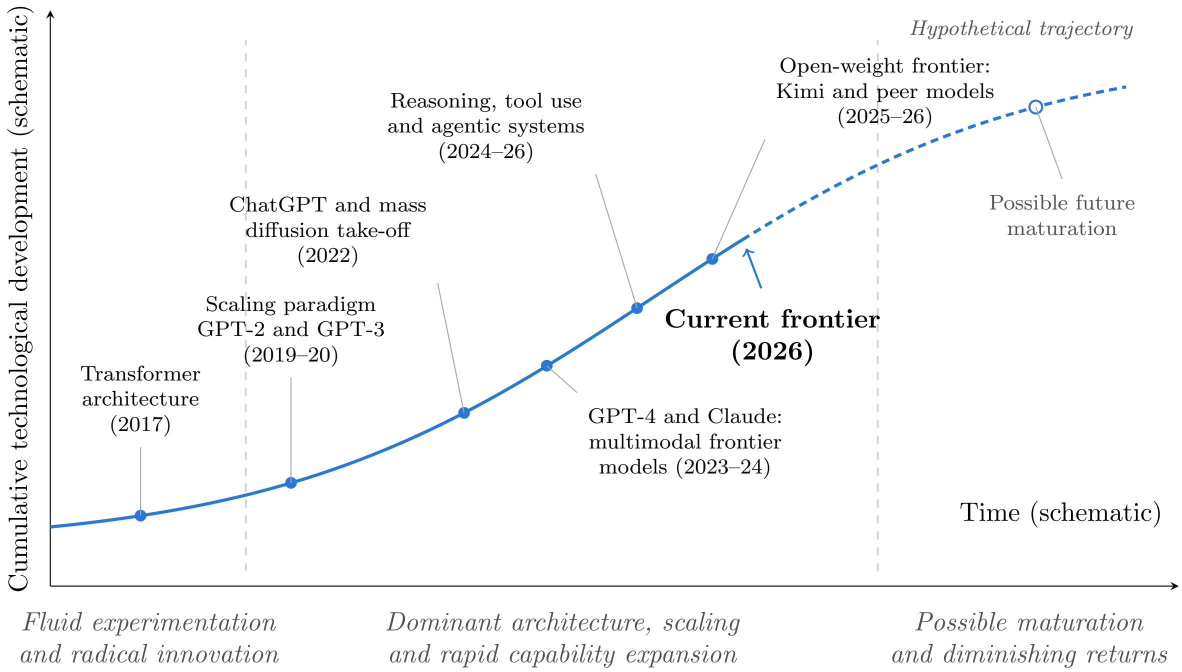

Figure 1: A stylized neo-Schumpeterian interpretation of the development of generative AI. The transformer architecture represents an enabling technological breakthrough, followed by a fluid period of experimentation and the emergence of scaling based on pretrained foundation models. ChatGPT marks the take-off of mass diffusion rather than a discrete measure of underlying capability, while GPT-4 and Claude represent the emergence of multimodal frontier systems. More recent developments in inference-time reasoning, tool use, agentic workflows and open-weight models such as Kimi suggest that the technological trajectory remains open and may be undergoing further inflections. The solid portion represents observed historical development, whereas the dashed portion presents only one possible future trajectory in which marginal capability gains eventually diminish and innovation shifts increasingly toward reliability, specialization, lower-cost inference and process improvement. The logistic form is conceptual rather than empirically estimated and should not be interpreted as evidence that generative-AI capability is approaching a fixed ceiling. Milestone positions are schematic and are not proportional to calendar time. Framework adapted from Abernathy and Utterback [1978] and Perez [2010, p. 187, Fig. 1]; placement and interpretation of generative-AI milestones by the authors.

wave presents as a technology system nested within the fifth surge rather than as a separate constellation of technologies.

Table 1: Mapping the core-industry categories onto the emerging AI technology system

Category	AI equivalent	Economic function	Relation to the fifth (ICT) surge
Motive branches	AI accelerators (GPUs and related chips), cloud compute, and foundation-model APIs	Supply compute and model access as a metered input to production	Specialised extension of the microprocessor, the fifth-surge key input, rather than a new one
Carrier branches	Assistants, copilots, agents, AI-native software, robotics, and autonomous systems	Visible products that bring the new capabilities into use and drive their diffusion	New products built on inherited software platforms and distribution channels
Infrastructure	Data centres, cloud platforms, model APIs, data pipelines, and high-speed networks	Extend the reach of AI across industries and geographical markets	Largely inherited from the information revolution, including the internet, cloud computing, and data pipelines
Induced branches	Electricity generation and grid capacity, semiconductor fabrication and advanced packaging, cooling and water systems	Pre-existing industries modernised and rescaled to enable large-scale AI deployment	Existing fourth- and fifth-surge industries expanded for AI

Categories adapted from Freeman and Perez [1988] and Perez [2010, pp. 191, 194]. The final column indicates whether each element is largely inherited from the information revolution or introduced by AI. The predominance of inherited infrastructure and inputs supports classifying AI as a technology system within the fifth surge rather than as a separate technological revolution.

The key criteria come from Perez's distinction between a true revolution and a mere assortment of systems. In her framework, a technological revolution is "a set of interrelated radical breakthroughs, forming a major constellation of interdependent technologies—a cluster of clusters, or a system of systems" [Perez, 2010, p.189]. Two features set such a revolution apart from a random collection of technological systems. First, the participating systems must be strongly interconnected and interdependent, both technologically and in their markets. Second, the revolution must be capable of profoundly transforming the broader economy and, ultimately, society. The first of these is more readily apparent, but it is the second that earns the label [Perez, 2010]. AI clearly meets the first criterion. Whether it meets the second—that is, whether it transforms industries beyond AI itself—remains an open question. The wave would count as a revolution only if it gives rise to a distinct cluster of industries, along with its own infrastructure, organisational principles, and transformative effects across the entire economy.

This classification can be tested against a rival framework from the economics of innovation: the general-purpose technology (GPT).[1] A GPT is a single, generic technology [Lipsey et al., 2005, p. 98], characterised by pervasiveness, an inherent potential for technical improvements, and the generation of 'innovational complementarities" across the sectors that use it [Bresnahan and Trajtenberg, 1995, pp. 83–84]. A recent OECD assessment applies these criteria to generative AI and finds that it 'has considerable potential to qualify as a new GPT" [Calvino et al., 2025, p. 42]: generative-AI patents span over twenty-one application areas (p. 19), training compute has grown about 4.4-fold per year since 2010 (p. 25), and forward citations show generative-AI patents feeding follow-on innovation well beyond software.

Generative AI may therefore qualify as a general-purpose technology, but that does not automatically make it the basis of a separate technological revolution. The two tests measure different things. GPT status is a property of a single technology; a technological revolution is a property of a constellation of technology systems together with an economy-wide reorganisation of production. Historically, the two have been related as part to whole: the steam engine, electricity, and the microprocessor were each general-purpose technologies, and each served as the key input of a revolution rather than constituting one. Understood in these terms, generative AI, if it passes the GPT test, is a candidate successor to the microprocessor as the key input of a deepened information revolution—which is precisely the nested classification proposed here.

The OECD's own caveats point in the same direction. Aggregate productivity gains are not yet visible; they may follow the delayed "J-curve" pattern familiar from earlier GPTs [Brynjolfsson et al., 2021]. Moreover, Bresnahan (2024, cited in Calvino et al., 2025, p. 12) notes that pre-generative AI never achieved a self-sustaining positive feedback loop with its application sectors. Such a feedback loop provides one practical indication of the second criterion: a technology transforms the wider economy when the sectors that use it begin to change in ways that feed back into its own development. Whether generative AI will produce such a loop is an empirical question that remains open.

That criterion should not be treated as satisfied merely because the industry claims it is. The large language model industry presents continued scaling as a direct path to artificial general intelligence. Yet computer tasks are not free-standing abstractions; they

[1] In this literature GPT abbreviates *general-purpose technology*; the coincidence with the model family (generative pre-trained transformer) is accidental, and the two should not be conflated.

are expressions of human purposes, and someone must still decide what is worth doing, for whom, and under what conditions. The evidence that all computer-mediated work can be automated remains limited. Nor are these reservations confined to outside critics. Yann LeCun, Meta's former chief AI scientist and a Turing laureate, regards large language models as useful but fundamentally limited, constrained by language and lacking a model of the physical world. He has left Meta to pursue an alternative research program [Heikkilä, 2026]. His wager is instructive: not that machine intelligence will fail, but that the current dominant design will not deliver it. General intelligence may be a useful horizon for research, but it should not be treated as an imminent product, nor as a justification for debt-financed overbuilding in the present. In neo-Schumpeterian terms, the capacity to transform the rest of the economy, the very feature that would make AI a revolution, is at this stage asserted rather than demonstrated. The more defensible classification is therefore the narrower one: AI is a rapidly developing technology system within the information revolution, and whether it becomes a technological revolution in its own right remains an open question.

3 The Extent of the Diffusion and Its Impacts

Ricardo had long held that labour-saving machinery was "a general good"; in the chapter "On Machinery," added to the third edition of the Principles, he announced that his opinions had "undergone a considerable change" and conceded that "the substitution of machinery for human labour, is often very injurious to the interests of the class of labourers" Ricardo [2001, pp. 282–287]. His mechanism remains instructive. Machinery may raise a nation's net produce (profit and rent) while reducing its gross produce; and since "the power of supporting a population, and employing labour, depends always on the gross produce of a nation, and not on its net produce," national income and private returns can rise while the wage fund and employment contract. Ricardo was careful, however, to claim only that machinery may have this effect: where the net surplus is large enough, or is reinvested fast enough to be reconverted into a fund for employing labour, "the situation of all classes will be improved." Labourers' fear of machinery, he held, "is not founded on prejudice and error, but is conformable to the correct principles of political economy," rational precisely because the adverse case is possible, not because it is certain [Ricardo, 2001, pp. 282–287]. The possibility of displacement thus arises from the structure of capitalist accumulation rather than from the character of any particular technology; whether it is realised depends on what happens to the surplus, and how quickly.

Ricardo's mechanism translates directly into the two macroeconomic questions the AI wave now poses. The first concerns capital and productivity: whether the scale of investment now being committed is warranted by the productivity gains the technology can plausibly deliver, a question that the history of technological revolutions, each with its own episode of overoptimism, obliges us to ask early rather than late. The second concerns employment: how large the effect on jobs will be, and what form it takes, whether elimination, expansion or reorganisation of work. The evidence reviewed below suggests a cautionary answer to the first question and a moderating answer to the second. The productivity gains on which the investment is premised remain empirically unresolved even as the capital is already committed, and the European labour-market evidence points to transformation rather than elimination.

On the investment side, the numbers already suggest excess. Hyperscalers and neoclouds are estimated to have committed roughly \$2 trillion in cumulative capital expenditure through 2026, with AI-linked spending running some \$535 billion per year above the pre-AI trend. The estimated 2026 depreciation charge approaches \$111 billion, set against generative-AI ecosystem revenue of about \$175 billion annualized, while the power of the largest data centres has grown fiftyfold in four years [Azhar et al., 2026]. But excess is not the whole story. For the second consecutive quarter, AI infrastructure revenue has exceeded the depreciation charge on the capital deployed to produce it, with headroom of 19% on hyperscaler and neocloud revenues in the first quarter of 2026, and 32% once application- and model-layer revenues are included [Azhar et al., 2026, slides 33–34]. This milestone should be read with caution: revenues now cover the ongoing quarterly cost of the buildout, yet cumulative revenue has only nearly caught up with cumulative depreciation and amounts to only about half of what a healthy margin above that depreciation would require. Still, it is enough to undermine the strongest version of the bubble thesis. The question is no longer whether paying demand exists; it is whether revenue compounds fast enough to outrun a depreciation base that will keep climbing as already-committed capacity comes online. Whether it can do so depends on productivity gains that remain deeply uncertain. The risks of the AI infrastructure boom depend heavily on how it is financed. Hyperscalers still fund most of their capital spending from operating cash flow, while neocloud companies rely far more on borrowing, and external funding has been rising as a share of the total [Azhar et al., 2026]. That distinction matters. When a firm invests retained earnings or raises new equity, poor returns fall mainly on the company and its shareholders, in the form of lower profits and falling share prices. With debt, however, repayments are still due even if expected revenues fail to materialise. Losses can then lead to defaults, weaken creditors' balance sheets, and spread beyond the firms that made the original investments.

Acemoglu [2025] estimates, from a task-based accounting, a total factor productivity gain of no more than 0.66% over ten years, roughly 0.07% per year. Using the same formula but reading the empirical literature on its components differently, Aghion and Bunel [2024] arrive at between 0.07 and 1.24 percentage points of additional annual TFP growth, with a median estimate of 0.68, excluding any effect of AI on the production of ideas. Applied to 31 European countries, the same framework yields a preferred estimate of about 1.1% cumulatively over five years, with country results ranging from near zero to above 5% across scenarios and gains concentrated in higher-income economies [Misch et al., 2025]. The disagreement is itself the policy-relevant fact: a shared accounting framework, fed with defensible readings of the same evidence, spans an order of magnitude, and trillions of dollars of already-committed capital formation rest on where within that span the truth lies. None of this implies the technology is not real. There is no contradiction between a genuine transformation and over-capitalisation: railways, electrification and the dot-com boom each outran their revenues and left financial wreckage behind, and in the history of technological revolutions the installation of a new technology system has regularly been financed through precisely such episodes of frenzied over-investment [Perez, 2002]. AI capex, leverage, valuations, energy constraints and the concentration of cloud and chips should therefore be treated as questions of macroprudential policy, not merely of industrial strategy.

On the labour-market side, the evidence so far points to transformation rather than elimination. Reviewing the empirical evidence across experiments, firm-level data, platform studies and representative worker surveys, the ILO finds that large-scale job displacement remains limited and that generative AI so far mainly reorganises tasks within jobs, with workers reallocating effort towards tasks less easily automated [Merola et al., 2026]. A framework published by OpenAI, reaches a compatible conclusion for Europe by a different route. Rather than treating AI exposure as a single measure, it sorts occupations along several dimensions at once: how much of the work AI can technically perform, whether human judgement and accountability must remain at the centre of delivery, how demand responds as costs fall, and how the technology is actually being used [Richmond, 2026]. It estimates that only 14% of mapped European employment is in jobs with higher automation potential, compared with 18% in the United States, while 27% of European employment is in jobs likely to be reorganised, compared with 24% in the United States; a further 12% of jobs may grow with AI and 47% face less immediate change. The difference reflects occupational structure: Europe employs relatively more people in manufacturing, skilled trades, transport, care, education and public-service occupations, where work is physical, place-based, regulated or tied to fixed service volumes, while the United States has relatively more managerial, sales and digitally deliverable business-service jobs. The report itself warns that these categories describe transition pressures, not forecasts of unemployment.

Firm-level European evidence points the same way. Examining some 5,000 euro area firms in the ECB's enterprise survey, Lebastard and Sondermann [2026] find no significant difference in job creation and destruction between firms that use AI and those that do not; where effects do appear, hiring rises among firms deploying AI for research, development and innovation, and falls among the minority adopting it explicitly to cut labour costs. These findings are an early snapshot, not proof that AI creates jobs: the authors themselves caution that the effects are "currently still positive" precisely because AI "has not yet significantly transformed production processes". But the European evidence, from the exposure mappings to the firm-level data, is consistent in its direction: transformation rather than elimination, reorganisation rather than displacement, with displacement concentrated where AI is deployed as a labour-cost instrument.

What gives the present wave its political charge is the social location of its effects. Earlier industrialisation battered manual and routine labour. So far AI reaches into clerical, legal, administrative and programming work, precisely the terrain on which the West built its post-industrial bargain: the tacit settlement that made the offshoring of manual work tolerable in the capitalist core rested on the survival of secure office work for a broad middle class. If those jobs are hollowed out, the bargain is exposed. The danger is not a science-fiction apocalypse of mass unemployment but a quieter corrosion, and its outline is already visible in the evidence: the ILO locates the principal risks not in job counts but in growing inequalities, the erosion of entry-level opportunities for younger workers, and the reshaping of work organisation, autonomy and job quality [Merola et al., 2026, pp. 1–4]. This is Ricardo's lesson in contemporary form: the fear is rational not because the adverse case is certain, but because it is possible, and because its realisation depends on what happens to the surplus. The two questions posed at the outset thus receive parallel answers: the productivity gains on which the investment is premised remain unresolved while the capital is already committed, and the employment effect is taking the form of reorganisation rather than elimination. Neither outcome is given by the technology; both are decided by how firms, financiers and institutions deploy it, which is what makes them questions of policy.

The evidence reviewed here relates mainly to cognitive and clerical work, where generative AI has so far been most widely deployed, and it describes a period in which, as the firm-level studies themselves note, AI has not yet significantly transformed production processes. It therefore tells us little about what might happen if AI were combined with robotics and began to transform physical work. Such a development does not appear imminent. The scaling strategy that produced large language models does not transfer easily to embodied systems: language models benefited from an enormous pre-existing body of text, whereas robots must generate much of their training data through physical interaction with the world, and combining datasets collected from different robots and environments can produce negative transfer rather than the relatively smooth gains from scale seen in language models. Adoption would also be slower, because robots are physical capital, so their diffusion depends on investment, production, installation and organisational change rather than on the rapid distribution of software. Nevertheless, if AI were eventually integrated with robotics in a way that transformed physical work across the economy, it would bring about precisely the kind of economy-wide reorganisation required by the second criterion, the capacity to transform the rest of the economy. Under those conditions, the current wave would qualify as a technological revolution rather than as a technology system developing within an existing one. The classification advanced here is therefore conditional, and this is the main circumstance under which it would need to be revised

4 Strategic Disadvantage of Europe

Europe entered the AI wave with a structural disadvantage inherited from the previous digital revolution. EU productivity began to diverge from that of the United States in the mid-1990s, and the Draghi report identifies digital technology as the main source of the gap: "the key driver of the rising productivity gap between the EU and the US has been digital technology". Europe produced fewer major technology firms and diffused digital technologies more slowly across the economy. Once the technology sector is excluded, however, European productivity growth over the past two decades is broadly comparable to that of the United States [Draghi, 2024]. The weakness is therefore concentrated precisely in the sectors that now shape the AI transition. No EU company with a market capitalisation above EUR 100 billion has been created from scratch in the past fifty years, while all six US companies valued above EUR 1 trillion were founded during that period [Draghi, 2024, p. 6].

AI is widening this existing divide. In 2025, the United States produced 59 notable AI models and China 35, while Europe produced only two [Sajadieh et al., 2026, p. 17]. US private AI investment reached $285.9 billion, compared with $20.9 billion in Europe and $12.4 billion in China. The United States also recorded 1,953 newly funded AI companies, compared with 172 in the United Kingdom, 92 in Germany and 84 in France. US private AI investment grew by 160% over 2024, while European investment increased by only 7% [Sajadieh et al., 2026, pp. 182–183]. The Chinese figure, however, needs to be interpreted differently. It excludes a large channel of state-directed investment. Government guidance funds and related vehicles are estimated to have deployed $912 billion across industries between 2000 and 2023, including roughly $184 billion to AI companies [Sajadieh et al., 2026, p. 184]. The comparison therefore points to more than a difference in the amount of private capital available. It reflects different institutional systems for mobilising investment.

The consequences are increasingly visible at the technological frontier. As of July 2026, Moonshot's Kimi K3 ranks first on the WebDev Arena leaderboard, while Z.ai's GLM-5.2 ranks fourth overall; no European model appears in the top ten [Arena.ai, 2026]. One explanation for China's ability to convert capital into technological capability is the "platform state" described by Yang and Zhang [2026] and formalised in companion work by one of us [Ernst, 2026].[2]

Europe's position is particularly striking because it already possesses important capabilities within the AI value chain. As García-Herrero and Martens [2026, pp. 2–3] document, ASML controls the critical technology for extreme ultraviolet lithography, while Belgium's IMEC is a global leader in advanced semiconductor research; yet much of the downstream value enabled by these technologies is captured outside Europe. A similar dependence appears at the model layer: Mistral relies on Microsoft Azure infrastructure, works with Nvidia on compute capacity and has expanded into Palo Alto to access American talent and capital. Europe's constraint is therefore "not primarily talent or ideas, but access to compute and capital" [García-Herrero and Martens, 2026, p. 3]. The central weakness is not the absence of technological capabilities, but Europe's limited ability to connect them into an integrated AI ecosystem.

European policy has increasingly recognised these weaknesses, first through regulation and more recently through investment. The AI Act, in force since August 2024, established the world's first comprehensive risk-based framework for AI [European Union, 2024]. Yet its implementation soon became more difficult. Before the high-risk obligations had begun to apply, the Commission proposed a "Digital Omnibus" amending the Act, and by June 2026 the co-legislators had postponed the high-risk rules until the end of 2027 and 2028 and softened several obligations, citing delays in designating national authorities and the absence of harmonised standards [Niestadt, 2026]. Civil-society organisations described the changes as "a rollback of AI safeguards before they even apply" [European Digital Rights (EDRi) et al., 2026]. The episode illustrates the difficulty of sustaining regulatory ambition when implementation capacity is uneven and concerns about competitiveness are growing.

The investment response has sought to address Europe's technological weakness more directly. The InvestAI initiative aims to mobilise EUR 200 billion, while the AI Continent Action Plan includes up to EUR 20 billion in public investment and a network of AI factories and future gigafactories under the EuroHPC Joint Undertaking [Kyosovska and Renda, 2025]. Yet both the scale and organisation of this investment remain problematic. The commitment is modest relative to the US buildout, and many AI factories are located outside

[2] Ernst [2026] is an unpublished working paper by one of the present authors. We cite it as companion work rather than as independent corroboration of the argument made here.

Europe's strongest AI clusters. Their consortia are dominated by research institutions, while their mandates are often only loosely connected to surrounding industrial capabilities [Kyosovska and Renda, 2025]. They may support research and the development of medium-sized models, but are "not sufficient to boost commercial AI innovation across the EU at scale" [Lemke and Schneider, 2025, p. 3].

These problems are reinforced by fragmentation. European cloud policy shows how national interests can undermine a common technological strategy. Calcara [2026] shows how initiatives such as Gaia-X evolved into national "sovereign cloud" arrangements involving the same American hyperscalers whose dominance Europe initially sought to reduce. France's Bleu combines Capgemini and Orange with Microsoft, while S3NS combines Thales with Google. Fragmentation also weakens Europe's demand-side power. Public purchasing remains divided across twenty-seven member states, "each acting largely alone", preventing Europe from using procurement to create the large and predictable markets that domestic technology firms need to scale [García-Herrero and Martens, 2026, pp. 3–4].

The Airbus comparison shows what a more coherent European strategy would require. Airbus succeeded not simply because governments spent money, but because production was organised around existing industrial capabilities, governance became increasingly integrated, and the project was sustained over decades. French avionics, German manufacturing and British wing design were combined within a structure whose decisions eventually became "subordinated to commercial logic" [García-Herrero and Martens, 2026, pp. 4–6]. An AI strategy built on the same principle would allocate investment according to technological capabilities rather than geographical balance, integrate public procurement with industrial development, and coordinate infrastructure, finance and demand at the European rather than national level.

From a broader macroeconomic perspective, the policy debate therefore concerns not whether the state should intervene, but how. One approach emphasises integrated capital markets, joint funding and public mechanisms for de-risking private investment [Draghi, 2024, pp. 18, 63]. Evidence from the green bond market suggests that such interventions can lower financing costs and reduce volatility, making entry easier for new firms [Braga et al., 2021]. A second approach goes further by arguing that public support should come with stronger conditions, including risk and reward sharing and a greater public role in directing investment [Mazzucato et al., 2026, p. 8]. A third regards both approaches as insufficient and calls for public ownership of key financial and productive sectors [Roberts, 2026]. These positions differ sharply in ambition, but they share a common premise: market forces alone are unlikely to overcome Europe's structural disadvantage.

The central challenge is therefore not simply to spend more on AI. Europe needs institutions capable of connecting its technological strengths with capital, compute, firms and demand. Without that coordination, upstream capabilities will continue to generate much of their value elsewhere, while European policy remains caught between regulatory ambition, fragmented investment and dependence on foreign infrastructure.

5 What strategy? Elements of a European answer

Europe will not win by dressing itself up as a second Silicon Valley, nor should it want to. The analysis above points elsewhere: toward a strategy of human-centred AI, in which the technology is governed by social needs rather than by the impatience of quarterly returns, and in which the state acts deliberately across the whole chain, from the allocation of productivity gains to the ownership of infrastructure [Ernst, 2022].

5.1 Convert productivity gains into time and general well-being

Whether AI's productivity gains materialise at all depends on a factor that rarely appears in the investment arithmetic: whether employees cooperate with the technology's deployment. The early evidence suggests that cooperation cannot be assumed. In a 2026 survey of 2,400 executives and employees across the US, the UK and Ireland, Benelux, France and Germany, conducted by an enterprise-AI vendor and restricted to workplaces already using AI, 29% of employees admitted to undermining their company's AI strategy in at least one way, from refusing to use the tools to deliberately degrading their output, a figure that rises to 44% among Gen Z workers [WRITER, 2026, p. 24]. A parallel survey of 3,750 leaders and employees across 14 countries finds that in any given month more than half of workers abandon their enterprise tools and complete the work manually, that 37% avoid AI because it disrupts their workflow, and that 45% used unapproved AI tools in the past 30 days [WalkMe, 2026, pp. 12, 24]. Workers, in short, are simultaneously resisting the AI their employers prescribe and adopting the AI their employers prohibit. Both behaviours point to the same underlying condition: deployment is proceeding without the workforce's consent, and productivity gains extracted from a workforce that expects to be displaced by them will be partial, contested and fragile. The question is therefore how to give workforce a visible stake in the gains.

The most direct such stake is time: the immediate policy response to AI-driven productivity gains should be the reduction of working hours rather than the reduction of employment. Since reorganisation rather than elimination dominates the European exposure profile, as the labour-market evidence above showed, the gains can be obtained as shorter hours instead of labour-cost reductions. The proposal has historical precedent, since earlier industrial revolutions were followed by a gradual decline in working time, and a classical pedigree in Keynes's argument that technological progress could eventually translate into leisure rather than longer toil [Keynes, 2010]. It is also politically decisive: the reception of the AI wave will depend on whether people perceive it as enriching only firms and upper-income groups or as improving well-being broadly, and shorter hours make the gains visible by converting productivity into time. Survey evidence supports the social legitimacy of the instrument: asked what they hope an economy shaped by AI will look like in ten years, over half of respondents to the Anthropic Economic Index survey expressed a desire for meaningful human–AI collaboration, just over half hoped that AI would automate the tedious parts of their jobs so that they could have more free time, and about one third hoped that the economic gains would be widely shared [Massenkoff et al., 2026, p. 30].

The empirical record on working-time reduction offers no unified verdict. Studies of the French experience reach opposing conclusions: Estevão and Sá [2008] find no positive employment effects from the 35-hour week, Crépon and Kramarz [2002] associate the 1982 reduction with job losses among affected workers, while Askenazy [2013, pp. 339, 340] reports official evaluations attributing some 300,000 to 350,000 jobs to the reform and cautions that its effects cannot be separated from the subsidies packaged with it. The design of any AI-era working-time policy would therefore need to be negotiated, conditional on productivity, and evaluated as it proceeds.

The favourable evidence comes from reductions that were negotiated, conditional and reversible. Pencavel's analysis of output and working hours shows that the productivity of an hour declines as weekly hours lengthen, so that reductions from long hours are partly self-financing [Pencavel, 2015]. Iceland's public-sector trials found productivity and service provision maintained or improved across the majority of trial workplaces, and were followed by negotiated contracts under which roughly 86% of the country's working population either moved to shorter hours or gained the right to them [Haraldsson and Kellam, 2021]. The largest multi-country trial to date, covering 2,896 employees across 141 organisations in six countries, found improvements in burnout, job satisfaction and mental and physical health that persisted at twelve-month follow-up, with pay held constant and work reorganised to maintain output, though among self-selecting firms [Fan et al., 2025]. The design implications for an AI-era working-time policy follow directly. Reductions should be bargained at sectoral and firm level through the co-determination channels discussed below, tied to realised productivity gains rather than legislated uniformly, so that unit labour costs and hence competitiveness are held constant by construction; the public sector, as Europe's largest employer and itself a major AI adopter, should move first, on the Icelandic model; and pilot programmes with built-in evaluation should precede any generalisation. Working-time reduction, so designed, is not a transfer against productivity but a rule for sharing its growth: firms keep the gains they generate, and workers receive their share in the currency the survey evidence says they want most, which is time. Each of these presupposes institutions capable of striking and monitoring the bargain. Works councils, board-level representation and comparable structures give the workforce standing not only over hours but over how AI is introduced, how restructuring is handled and how the income generated by higher productivity is distributed. Without them, a productivity-contingent agreement has no forum in which it could be negotiated or enforced.

5.2 Picking winners: conditional public support for European champions

On the macro-policy level Europe needs a picking-winners strategy, and it should run through the industrial base. AI adoption in most European countries competes US levels, over three quarters of European AI funding targets vertical applications, and Europe shows "genuine competitive strength in robotics, AI manufacturing, and deep tech applications such as autonomous driving and AI drug discovery" [Prosus, 2026, p. 4]. Germany and France, for example, retain deep engineering capabilities that position them as second movers in AI applications rather than as competitors at the model frontier. Strength of this kind grows out of production capabilities and learning in production [Chang and Andreoni, 2020]. The capital for such a strategy is already partly public: the EU AI Champions Initiative groups more than sixty companies around AI adoption in manufacturing, energy and defence [General Catalyst, 2025, EU AI Champions Initiative, 2026], and France's public investment bank Bpifrance, through its Digital Venture fund, participated in the \$1.03 billion seed round of LeCun's AMI Labs [Heim, 2026]. The design of such support matters as much as its volume. The formal analysis of the platform state introduced in Section 4 shows that screening and market-based selection are complements rather than substitutes [Ernst, 2026]: a state's advantage from engineering broad entry depends on its ability to lower entry costs, so "a state that cannot lower entry costs should not engineer over-entry", and the choice between picking winners ex ante and letting the market select ex post turns on the ratio of the state's information to the market's selectivity — an empirical question, not one of ideology. For Europe, whose fragmented single market weakens exactly the ex post selection mechanism the platform regime relies on, this cuts in favour of combining conditional support with measures that deepen market integration.

Picking winners is defensible only if the public picks up more than the risk, and the champions' own agenda shows that reciprocity will not be offered voluntarily: the initiative's report asks for risk-absorbing public finance, lighter regulation and publicly supported infrastructure, while offering no commitment on equity, employment or worker voice in return [General Catalyst, 2025, pp. 43–46]. Two conditions should therefore be attached by policy. The first concerns the distribution of gains: socialised risk should carry socialised reward, through public equity stakes in the firms taxpayers help create [Mazzucato, 2018, Mazzucato et al., 2026]. The second concerns decision-making: firms receiving public equity, subsidies, preferential procurement or publicly financed compute should be required to give their workforce a formal voice in how AI is deployed, following the co-determination practice established in the German corporate sector. The justification for this condition is economic rather than social, because the employment effect of AI is decided at the point of deployment: firm-level evidence from France indicates that even among the roughly 10% of workers in highly exposed, highly substitutable jobs, employment rises where AI is deployed in production or IT security and falls where it is applied to administrative cost-cutting [Aghion, 2025]. Co-determination steers deployment toward the former, and workers with a say in how AI enters their workplace have little reason to work against it.

5.3 Invest in the public inputs of AI production

A picking-winners strategy presupposes inputs. The two decisive inputs of AI production, data and compute, are where public action is both feasible and legitimate. The first priority is a European public-data infrastructure: selected public-sector datasets made available to European AI labs, universities and SMEs for training and research, under clear rules on privacy, security, accountability and public value. The precedents are established, from the Human Genome Project and the Protein Data Bank to Transport for London's open-data programme and Taiwan's immediate publication of non-personal public data [Mazzucato et al., 2022, pp. 6–7]. The terms of access matter as much as access itself: public data should not be aggregated in ways that entrench dominant positions, as the commercial exploitation of public meteorological and transport data illustrates [Mazzucato et al., 2022, p. 7].

Policies that support national or European champions usually face two objections. The first is that governments should not be in the business of picking winners. Yet, as Chang argues, public policy is never entirely neutral. When resources are limited, decisions about funding, regulation and infrastructure inevitably favour some activities over others, even when the policy is presented in general terms[Chang et al., 2013, p. 7]. The real issue is therefore not whether governments target particular sectors or firms, but whether they do so transparently and can be held accountable for those choices.

The second concern is that successful champions may become entrenched monopolies. But size alone is not the problem. Large firms can often undertake expensive and uncertain investments that smaller firms cannot, and Schumpeter saw the large corporation as an important driver of technological progress. What matters is how firms use their market power: whether they continue to innovate or use their position mainly to keep rivals out and collect rents [Schumpeter, 1976, ch. 8, pp. 87–106]. The conditions discussed above would provide some restraint, but they would not be enough on their own. Strong competition policy would still be needed. Support for European AI champions should therefore be tied to measures that preserve competition, including support for open-source models, data-sharing requirements and the extension of Digital Markets Act principles to cloud services and AI. These measures would help offset the advantages already enjoyed by the small number of firms that dominate cloud infrastructure and GPU access [Aghion, 2025]. Europe could also create mission-oriented agencies modelled on DARPA. A coalition of willing Member States could take the lead rather than waiting for agreement across the entire Union [Aghion, 2025].

The second priority is publicly financed compute for European labs, universities and SMEs. The cost of computation has concentrated frontier research in a handful of firms: industry produced over 90% of notable AI models in 2025 [Sajadieh et al., 2026, p. 14], and the champions initiative's own report identifies GPU access as the key enabler Europe lacks, with European providers offering no operational sovereign cloud and firms forced into continued reliance on US infrastructure [General Catalyst, 2025, p. 45]. A European research cloud is the corresponding remedy, subject to one caveat: it must not "simply end up empowering large cloud companies without democratising access to cloud resources" [Mazzucato et al., 2022, p. 10]. These investments also serve the sovereignty problem directly: public data, public compute and European champions reduce dependence on external systems whose availability is no longer guaranteed.

5.4 Build the administrative capacity to plan adoption

None of the above executes itself, and Europe's competitiveness problem is institutional as much as technological: it needs better coordination, stronger governance and faster implementation capacity [Draghi, 2024, pp. 67–69]. An EU AI Public Sector Planning Taskforce should identify where and how AI can be responsibly incorporated into public-sector planning, administration and service delivery, moving governments beyond fragmented pilots toward common standards in infrastructure planning, energy systems, transport, healthcare administration, procurement and labour-market services. Governments need stronger public-sector capabilities to govern, procure, deploy and evaluate AI, rather than relying excessively on outsourcing [Mazzucato et al., 2022, pp. 13–14]. Such a body would address directly the three coordination failures Draghi identifies: between Member States, whose uncoordinated national policies produce duplication and incompatible standards; among financing instruments, split along national lines and therefore unable to form the large capital pools that breakthrough innovation requires; and across policies, at a time when industrial strategies elsewhere combine fiscal, trade and foreign economic policy into a single package [Draghi, 2024, pp. 15–16].

Institutions capable of this kind of coordination are not without precedent, and Europe can learn from their design. In *Governing the Market*, Wade shows that Taiwan's industrial upgrading was neither left entirely to market forces nor directed through central planning. It relied instead on a practical division of labour between firms and public agencies [Wade, 1990]. The Industrial Development Bureau was central to this system[3]. Before 1983, it employed around 180 professional staff, including roughly 130 engineers. Officials regularly visited firms, developed a detailed understanding of their production capabilities, assessed which companies were ready to move into more advanced activities, helped arrange finance and connected domestic suppliers with foreign affiliates [Wade, 1990, pp. 202, 207]. The renewed interest in industrial policy has made the case for such institutions more, rather than less, relevant [Wade, 2012]. A similar approach could be applied to AI in Europe. The coordinating body would not simply regulate firms after problems arise. It would identify strategically important applications, improve access to data and computing resources, support experimentation across Member States and guide adoption towards broader public goals. Its role, in other words, would be to shape the development of the market rather than merely correct its failures [Mazzucato et al., 2022].

5.5 Procurement and the problem of platform dependence

One instrument cuts across all five elements and deserves separate treatment: public procurement. Procurement has long served as an instrument of industrial policy, as a guaranteed first market, a demand-side subsidy and a mechanism for building domestic capability; the government-procurement channel was among the standard tools of successful industrial policy in the United States, Japan and Finland alike [Chang and Andreoni, 2020]. Procurement is also the most direct demand-side mechanism for achieving the paradigm shift that one of us has argued the AI wave requires Ernst [2022]. This requires moving away from the current supply-push regime, where the direction of AI development is largely shaped by the investment priorities of a small number of large firms, towards a demand-pull regime

[3] A parallel account for Korea is given by Amsden [1992], who stresses the reciprocal character of state support: subsidies were tied to performance standards that the state monitored and enforced.

in which technology evolves "toward applications beneficial from a societal perspective." By deliberately designing public procurement strategies, governments can steer AI development towards applications with high social returns, such as improving resource efficiency, complementing human labour, and enhancing public services, rather than reinforcing the labour-displacing applications encouraged by the current incentive structure. But AI chatbots and APIs are not conventional procurement objects. They are platform technologies: procured once, they embed themselves in workflows, data flows and institutional routines, generating switching costs that grow with use. The Cloud and AI Development Act, proposed in June 2026 as the centrepiece of the Commission's tech-sovereignty package, gestures at this by requiring public authorities to weigh "Union added value" as a non-price criterion in cloud and AI procurement [European Commission, 2026b]. Yet the criterion is explicitly ancillary and non-decisive, with an indicative ceiling of roughly 15 of 120 award points suggested in the recitals [Maynard et al., 2026], which is unlikely to shift outcomes where the incumbent offering is cheaper and more capable. The question European policy must confront is whether a marginal scoring adjustment is an adequate instrument for a technology whose defining property is lock-in.

This dependence may be deeper than in earlier forms of procurement. The leading AI laboratories are already developing vertical applications, including in law, that may compete directly with the firms building on their models [Azhar et al., 2026, slide 61]. Once an AI provider becomes closely involved in a firm's operations, it is no longer simply an outside supplier. It may also be developing products that compete with that firm, while learning a great deal about how the business works. Alex Karp has expressed the concern in especially sharp terms. His warning is not disinterested, since Palantir offers a competing model, but the underlying point remains important. Firms that pay for access to external models may also be exposing the proprietary routines, workflows and decision-making processes that make up their competitive advantage to a provider that serves their competitors as well [Subin, 2026, Kahn, 2026]. Contractual protections do exist, and enterprise API terms generally state that customer data will not be used for training by default. But this does not remove the broader structural problem. Firms are not buying a discrete product once and moving on. They are entering into a continuous relationship with a provider that is also expanding into the industries in which its customers operate. The issue is therefore not only market share, but a deeper form of dependence: European firms and public institutions may reach a point where a foreign provider understands their internal operations better than they can reproduce those capabilities without it. Europe does have a deployment strategy on paper. The Commission's Apply AI Strategy calls for an "AI first" approach in public administration based on European solutions and seeks to mobilise around EUR 1 billion from existing programmes [European Commission, 2026a]. But this is modest when compared with an infrastructure buildout measured in the trillions and a procurement preference worth no more than 15 out of 120 award points. The mismatch matters because dependence will build gradually through routine procurement decisions that current policy tools are too weak to change.

6 Conclusion

AI is advancing quickly, but its economic meaning is still unsettled. The evidence reviewed here suggests that it is best understood, for now, as a powerful technology system within the wider information revolution rather than as a separate technological revolution. Investment has moved faster than the productivity evidence, while the labour-market effects so far point more towards the reorganisation of work than mass unemployment.

For Europe, the main risk is not only falling behind in model development. It is becoming dependent on foreign providers for compute, cloud infrastructure, data access and innovations increasingly used in firms and public administration. Regulation alone though it is needed, will not solve that problem. Europe needs a more active strategy built around public investment in data and compute, conditional support for European firms, stronger administrative capacity, workforce participation in form governance and procurement rules that reduce long-term lock-in.

The distribution of the gains also matters. If AI raises productivity, part of that gain should reach the workforce through better jobs, greater security, wage growth and, where possible, shorter working time. The alternative is that returns are retained as cash flow and paid out as dividends, feeding asset markets in which large foreign investors hold a substantial share. Whether that happens depends on corporate governance. Systems that give workers a formal voice in how AI is introduced, as German worker's councils and codetermination do, are better placed to turn productivity gains into wages and conditions than systems where those decisions rest with owners alone. The outcome will not be determined by the technology itself. It will depend on who owns it, how it is deployed and which institutions shape its use.

AI Usage Statement

AI agents and large language models (LLMs), including ChatGPT, Claude, and Kimi, were used during the preparation of this article to assist with drafting, editing, grammatical correction, and the generation and refinement of code used to produce figures. All AI-assisted output was reviewed and revised by the authors, who take full responsibility for the accuracy, arguments, interpretations, and final content of the article.

References

William J. Abernathy and James M. Utterback. Patterns of Industrial Innovation. *Technology Review*, 80(7):40–47, 1978.

Daron Acemoglu. The Simple Macroeconomics of AI. *Economic Policy*, 40(121):13–58, 2025. doi: 10.1093/epolic/eiae042. URL https://academic.oup.com/economicpolicy/article/40/121/13/7728473.

Philippe Aghion. Should We Fear AI?, December 2025. URL https://bcf.princeton.edu/events/philippe-aghion-on-should-we-fear-ai/. Bendheim Center for Finance, Princeton University.

Philippe Aghion and Simon Bunel. AI and Growth: Where Do We Stand? Policy Note, 2024. URL https://www.frbsf.org/wp-content/uploads/AI-and-Growth-Aghion-Bunel.pdf.

Alice H. Amsden. *Asia's Next Giant: South Korea and Late Industrialization*. Oxford University Press, June 1992. ISBN 978-0-19-507603-5. doi: 10.1093/0195076036.001.0001. URL https://doi.org/10.1093/0195076036.001.0001.

Arena.ai. WebDev AI Leaderboard - Best AI Models for Web Development, 2026. URL https://arena.ai/leaderboard/code/webdev.

Philippe Askenazy. Working time regulation in France from 1996 to 2012. *Cambridge Journal of Economics*, 37(2):323–347, March 2013. ISSN 0309-166X. doi: 10.1093/cje/bes084. URL https://doi.org/10.1093/cje/bes084.

Azeem Azhar, William Gildea, Hannah Petrovic, Nathan Warren, and Marija Gavrilov. The State of the AI Economy — Exponential View, 2026. URL https://intelligence.exponentialview.co/.

Joao Paulo Braga, Willi Semmler, and Dieter Grass. De-risking of green investments through a green bond market – Empirics and a dynamic model. *Journal of Economic Dynamics and Control*, 131:104201, October 2021. ISSN 0165-1889. doi: 10.1016/j.jedc.2021.104201. URL https://www.sciencedirect.com/science/article/pii/S0165188921001366.

Timothy F. Bresnahan and M. Trajtenberg. General purpose technologies 'Engines of growth'? *Journal of Econometrics*, 65(1):83–108, January 1995. ISSN 0304-4076. doi: 10.1016/0304-4076(94)01598-T. URL https://www.sciencedirect.com/science/article/pii/030440769401598T.

Erik Brynjolfsson, Daniel Rock, and Chad Syverson. The Productivity J-Curve: How Intangibles Complement General Purpose Technologies. *American Economic Journal: Macroeconomics*, 13(1):333–372, January 2021. ISSN 1945-7707. doi: 10.1257/mac.20180386. URL https://www.aeaweb.org/articles?id=10.1257/mac.20180386.

Bureau of Industry and Security, U.S. Department of Commerce. Framework for Artificial Intelligence Diffusion, January 2025. URL https://www.federalregister.gov/documents/2025/01/15/2025-00636/framework-for-artificial-intelligence-diffusion.

Antonio Calcara. European cloud computing policy: failing in Europe to succeed nationally? *West European Politics*, 49(4):994–1018, June 2026. ISSN 0140-2382. doi: 10.1080/01402382.2025.2491962. URL https://doi.org/10.1080/01402382.2025.2491962.

Flavio Calvino, Daniel Haerle, and Sarah Liu. Is generative AI a General Purpose Technology?: Implications for productivity and policy. *OECD Artificial Intelligence Papers*, June 2025. doi: 10.1787/704e2d12-en. URL https://www.oecd.org/en/publications/is-generative-ai-a-general-purpose-technology_704e2d12-en.html.

Ha-Joon Chang and Antonio Andreoni. Industrial Policy in the 21st Century. *Development and Change*, 51(2):324–351, 2020. ISSN 1467-7660. doi: 10.1111/dech.12570. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/dech.12570. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/dech.12570.

Ha-Joon Chang, Antonio Andreoni, and Ming Leong Kuan. International Industrial Policy Experiences & the Lessons for the UK. *Working Papers*, December 2013. URL https://ideas.repec.org//p/cbr/cbrwps/wp450.html. Number: wp450.

Bruno Crépon and Francis Kramarz. Employed 40 Hours or Not Employed 39: Lessons from the 1982 Mandatory Reduction of the Workweek. *Journal of Political Economy*, 110(6):

1355–1389, December 2002. doi: 10.1086/342807. URL https://www.journals.uchicago.edu/doi/10.1086/342807.

Mario Draghi. The Future of European Competitiveness. Technical report, European Commission, 2024. URL https://commission.europa.eu/topics/competitiveness/draghi-report_en.

Ekkehard Ernst. The AI trilemma: Saving the planet without ruining our jobs. *Frontiers in Artificial Intelligence*, 5, October 2022. ISSN 2624-8212. doi: 10.3389/frai.2022.886561. URL https://www.frontiersin.org/journals/artificial-intelligence/articles/10.3389/frai.2022.886561/full.

Ekkehard Ernst. The platform state A model of industrial policy by market design, 2026.

Marcello Estevão and Filipa Sá. The 35-hour workweek in France: Straightjacket or welfare improvement? *Economic Policy*, 23(55):417–463, July 2008. ISSN 0266-4658. doi: 10.1111/j.1468-0327.2008.00204.x. URL https://doi.org/10.1111/j.1468-0327.2008.00204.x.

EU AI Champions Initiative. EU AI Champions Initiative, 2026. URL https://aichampions.eu.

European Commission. Apply AI Strategy | Shaping Europe's digital future, 2026a. URL https://digital-strategy.ec.europa.eu/en/policies/apply-ai.

European Commission. Proposal for the Cloud and AI Development Act (CADA) | Shaping Europe's digital future, 2026b. URL https://digital-strategy.ec.europa.eu/en/library/proposal-cloud-and-ai-development-act-cada.

European Digital Rights (EDRi), ARTICLE 19, Access Now, AlgorithmWatch, Amnesty International, Danes je nov dan, European Center for Not-for-Profit Law (ECNL), Lafede – justícia global, and Politiscope. European Digital Rights (EDRi) and ARTICLE 19 and Access Now and AlgorithmWatch and Amnesty International and Danes je nov dan and European Center for Not-for-Profit Law (ECNL) and Lafede – justícia global and Politiscope. Joint position paper, European Digital Rights (EDRi), 2026. URL https://algorithmwatch.org/en/wp-content/uploads/2026/06/20260618_AI_Omnibus_Analysis_EDRI_et_al.pdf.

European Union. Regulation (EU) 2024/1689 of the European Parliament and of the Council laying down harmonised rules on artificial intelligence, June 2024.

Wen Fan, Juliet B. Schor, Orla Kelly, and Guolin Gu. Work time reduction via a 4-day workweek finds improvements in workers' well-being. *Nature Human Behaviour*, 9(10): 2153–2168, October 2025. ISSN 2397-3374. doi: 10.1038/s41562-025-02259-6. URL https://www.nature.com/articles/s41562-025-02259-6.

Christopher Freeman and Carlota Perez. Structural Crises of Adjustment, Business Cycles and Investment Behaviour. In Giovanni Dosi, Christopher Freeman, Richard Nelson, Gerald Silverberg, and Luc Soete, editors, *Technical Change and Economic Theory*, pages 38–66. Pinter Publishers, London, 1988.

Alicia García-Herrero and Bertin Martens. Europe Needs a Strategy to Close the Artificial Intelligence Compute Gap, May 2026. URL https://papers.ssrn.com/abstract=6871658.

General Catalyst. An Ambitious Agenda for European AI. Technical report, General Catalyst, 2025. URL https://files.elfsightcdn.com/eafe4a4d-3436-495d-b748-5bdce62d911d/2f2982fc-bd45-44cb-8a54-7c0ee138e11e/An-Ambitious-Agenda-for-Europe-Apr-2025.pdf.

Gumundur D. Haraldsson and Jack Kellam. Going Public: Iceland's Journey to a Shorter Working Week. Technical report, Autonomy and Alda, 2021. URL https://autonomy.work/portfolio/icelandsww/.

Melissa Heikkilä. Computer scientist Yann LeCun: 'Intelligence really is about learning'. *Financial Times*, January 2026.

Anna Heim. Yann LeCun's AMI Labs raises $1.03B to build world models, March 2026. URL https://techcrunch.com/2026/03/09/yann-lecuns-ami-labs-raises-1-03-billion-to-build-world-models/.

Eric J. Hobsbawm and George F. E. Rudé. *Captain Swing*. Phoenix, 2001. ISBN 978-1-84212-235-8.

Jeremy Kahn. Palantir CEO Alex Karp is paranoid about Anthropic and OpenAI. That doesn't mean he has nothing to fear., 2026. URL https://fortune.com/2026/07/07/palantir-ceo-alex-karp-cnbc-rant-wrong-about-anthropic-and-openai-but-he-still-has-reason-to-fear/.

John Maynard Keynes. Economic Possibilities for Our Grandchildren. In *Essays in Persuasion*, pages 321–332. Palgrave Macmillan UK, London, 2010. ISBN 978-1-349-59072-8. doi: 10.1007/978-1-349-59072-8_25. URL https://doi.org/10.1007/978-1-349-59072-8_25.

Nicoleta Kyosovska and Andrea Renda. EU Plans for AI (Giga)Factories: Sanctuaries of Innovation, or Cathedrals in the Desert? Technical report, Centre for European Policy Studies (CEPS), 2025. URL https://cdn.ceps.eu/2025/11/251027-Sanctuaries-or-Cathedrals.pdf.

Laura Lebastard and David Sondermann. Artificial Intelligence: friend or foe for hiring in Europe today?, March 2026. URL https://www.ecb.europa.eu/press/blog/date/2026/html/ecb.blog20260304~d9e34fc95f.en.html.

Nicole Lemke and Catherine Schneider. The European Union's AI Factories: Lessons for Public Investment in AI Infrastructure in Europe. Technical report, Interface, 2025. URL https://www.interface-eu.org/publications/ai-factories.

Richard G Lipsey, Kenneth I Carlaw, and Clifford T Bekar. *Economic Transformations: General Purpose Technologies and Long-Term Economic Growth*. Oxford University Press, November 2005. ISBN 978-0-19-928564-8. doi: 10.1093/oso/9780199285648.001.0001. URL https://doi.org/10.1093/oso/9780199285648.001.0001.

Maxim Massenkoff, Eva Lyubich, Szymon Sacher, Zoe Hitzig, Shaoyi Zhang, Ryan Heller, and Peter McCrory. Anthropic Economic Index report: Cadences. Technical report, Anthropic, 2026. URL https://www.anthropic.com/research/economic-index-june-2026-report.

Paul Maynard, Dita Charanzová, Jadzia Pierce, Fredericka Argent, and Wolfgang Maschek. The EU Cloud and AI Development Act in Depth, June 2026. URL https://www.insideglobaltech.com/2026/06/11/the-eu-cloud-and-ai-development-act-in-depth/.

Mariana Mazzucato. *The Entrepreneurial State: Debunking Public Vs. Private Sector Myths*. Penguin Books, 2018. ISBN 978-0-14-198610-4.

Mariana Mazzucato, Marietje Schaake, Séb Krier, and Josh Entsminger. Governing artificial intelligence in the public interest. Working Paper, Stanford University, Cyber Policy Center and UCL Institute for Innovation and Public Purpose, July 2022. URL https://cyber.fsi.stanford.edu/publication/governing-artificial-intelligence-public-interest.

Mariana Mazzucato, Anna Hope Emerson, Samuel Fleischman, and Lara Merling. Labour Squeezed, Investment Stalled: The Hidden Cost of Financialisation in Europe's Corporate Sector. POLICY BRIEF POLICY BRIEF NO. 35, UCL Institute for Innovation and Public Purpose, 2026. URL https://www.etuc.org/en/document/report-labour-squeezed-investment-stalled.

Rossana Merola, Ekkehard Ernst, Daniel Samaan, Maria del Rio-Chanona, and Ole Teutloff. The impact of GenAI on jobs, productivity and work organization: a review of the empirical evidence, June 2026. URL https://www.ilo.org/publications/impact-genai-jobs-productivity-and-work-organization-review-empirical.

Florian Misch, Ben Park, Carlo Pizzinelli, and Galen Sher. AI and Productivity in Europe. IMF Working Paper Working Paper No. 2025/067, International Monetary Fund, 2025. URL https://doi.org/10.5089/9798229006057.001.

Maria Niestadt. Digital Omnibus on AI. EU Legislation in Progress PE 782.651, European Parliamentary Research Service, Brussels, June 2026. URL https://www.europarl.europa.eu/thinktank/en/document/EPRS_BRI(2026)782651.

John Pencavel. The Productivity of Working Hours. *The Economic Journal*, 125 (589):2052–2076, 2015. ISSN 1468-0297. doi: 10.1111/ecoj.12166. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/ecoj.12166. _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/ecoj.12166.

Carlota Perez. *Technological Revolutions and Financial Capital*. Edward Elgar Publishing, Inc, 2002. URL https://www.e-elgar.com/shop/usd/technological-revolutions-and-financial-capital-9781840649222.html.

Carlota Perez. Technological revolutions and techno-economic paradigms. *Cambridge Journal of Economics*, 34(1):185–202, 2010. ISSN 0309-166X. URL https://www.jstor.org/stable/24232030.

Prosus. State of AI Europe. Technical report, Prosus, 2026. URL https://www.prosus.com/state-of-ai-europe.

David Ricardo. *On The Principles of Political Economy and Taxation, by David Ricardo, 1821*. Batoche Books Kitchener, 52 Eby Street South, Kitchener, Ontario. N2G 3L1 Canada, 2001.

Alex Martin Richmond. The AI jobs transition framework for the EU. Research report, OpenAI Economic Research, July 2026. URL https://openai.com/index/mapping-ai-jobs-transition-eu/.

Michael Roberts. EU companies: a 'capitalism of rent'?, July 2026. URL https://thenextrecession.wordpress.com/2026/07/20/eu-companies-a-capitalism-of-rent/.

Sha Sajadieh, Loredana Fattorini, Raymond Perrault, Yolanda Gil, Vanessa Parli, Lapo Santarlasci, Juan Pava, Nestor Maslej, Russ Altman, Erik Brynjolfsson, Carla Brodley, Jack Clark, Virginia Dignum, Vipin Kumar, James Landay, Terah Lyons, James Manyika, Juan Carlos Niebles, Yoav Shoham, Elham Tabassi, Russell Wald, Toby Walsh, and Dan Weld. Artificial Intelligence Index Report 2026, June 2026. URL http://arxiv.org/abs/2606.15708. arXiv:2606.15708 [cs.AI].

Joseph A. Schumpeter. *Capitalism, Socialism, and Democracy*. Allen and Unwin, 5 edition, 1976. ISBN 978-0-04-335031-7.

Samantha Subin. Palantir's Karp bashes OpenAI, Anthropic token model: 'Something has gone completely wrong', July 2026. URL https://www.cnbc.com/2026/07/01/palantir-karp-open-ai-anthropic-tokens.html. Section: AI Age.

Robert Wade. *Governing the Market: Economic Theory and the Role of Government in East Asian Industrialization*. Princeton University Press, 1990. doi: 10.2307/j.ctv346sp7. URL https://www.jstor.org/stable/j.ctv346sp7.

Robert H. Wade. Return of industrial policy? *International Review of Applied Economics*, 26(2):223–239, March 2012. ISSN 0269-2171. doi: 10.1080/02692171.2011.640312. URL https://doi.org/10.1080/02692171.2011.640312. _eprint: https://doi.org/10.1080/02692171.2011.640312.

WalkMe. The State of Digital Adoption 2026: The AI Reality Check. Technical report, WalkMe, 2026. URL https://www.walkme.com/the-state-of-digital-adoption-2026/.

WRITER. AI adoption in the enterprise — How to overcome adoption hurdles, reduce risk, and turn AI ambition into enterprise impact. Technical report, WRITER and Workplace Intelligence, 2026. URL https://go.writer.com/ai-adoption-enterprise-2026.

Alex Yang and Angela Zhang. The Rise of the Chinese Platform State, March 2026. URL https://www.project-syndicate.org/commentary/china-innovation-model-could-give-it-an-advantage-in-ai-race-by-s-alex-yang-and-angela-huyue-zhang-2026-03.