

MARCH 2026

Working Paper

Artificial Intelligence and the Economy

Ricardo Paes Mamede*, Iscte – University Institute of Lisbon and Dinâmia'CET,

<https://orcid.org/0000-0001-9738-7480>

DINÂMIA'CET-Iscte WP n. 2026-02

DOI: 10.15847/dinamiacet-iul.wp.2026.02

Artificial Intelligence and the Economy

Table of contents

Introduction.....	3
1. How cheaper prediction changes firm behavior.....	4
1.1 The basic mechanism: prediction as an economic act.....	4
1.2 Business processes: substitution, augmentation and the productivity J-curve	4
1.3 Human resources, tasks and workplace design.....	5
1.4 Business models: digital servitization, data network effects and situated AI.....	6
1.5. Strategic management.....	6
1.6 Innovation activities: discovery, search and the “invention of a method of invention”	7
2. How firm level change reshapes markets	8
2.1. The vertical AI stack and sources of concentration.....	8
2.2 Startups, entry and the paradox of democratisation.....	9
2.3 Market conduct between firms: pricing algorithms	10
3. AI beyond the firm: inputs, macro-outcomes and geopolitical stakes.....	10
3.1 Hardware, semiconductors and the geopolitics of chips.....	10
3.2 Energy, data centres and the “AI energy paradox”	11
3.3 Minerals, recycling and the material footprint of AI.....	11
4. Macro aggregation: productivity, employment, and distribution	12
4.1 Translating micro gains into macro prospects.....	12
4.2 AI and Work: Substitution, Complementarity, and the Uncertain Path of Reallocation.....	12
5. The global dimension: power, policy and regional divergence.....	14
5.1. Technological leadership, industrial policy and great power competition	14
5.2 Convergence or divergence: who can leapfrog?	15
5.3 Global value chains: reshoring, dependencies and resilience.....	16
6. Governance, competition policy and democratic resilience.....	17
6.1 Competition policy for an AI economy.....	17
6.2 Corporate governance, accountability and market design.....	18
6.3 Industrial strategy and public provision of critical inputs.....	18
6.4 Labor market institutions and social insurance.....	19
Conclusion: what we know, what we don't and where to go next.....	20
References.....	21

Artificial Intelligence and the Economy

Abstract

This paper examines the economic implications of recent advances in artificial intelligence (AI), focusing on the mechanisms through which AI affects firms, markets and macroeconomic outcomes. Building on the view that modern AI primarily reduces the cost of prediction and pattern recognition, the paper analyses how lower prediction costs reshape organisational decisions within firms, including task allocation, business processes, human resource management, strategic decision-making and innovation activities. It then considers how these firm-level transformations influence market structure, highlighting the emergence of a vertically organised AI stack – from semiconductors and cloud infrastructure to foundation models and applications – and the economic forces that may lead to concentration and new forms of market power. The analysis subsequently examines the material and systemic foundations of AI, including semiconductors, energy systems, data centres and critical minerals, and discusses how these inputs interact with geopolitical competition and industrial policy. The paper also reviews the uncertain macroeconomic consequences of AI, assessing its potential effects on productivity, employment and income distribution, while emphasizing the importance of organisational complements, adoption patterns and institutional frameworks in shaping aggregate outcomes. Finally, it explores governance challenges, outlining the roles of competition policy, corporate accountability, industrial strategy and labour market institutions in shaping how AI-driven transformations affect economic efficiency, resilience and equity. The paper concludes by identifying key areas where further empirical research and policy experimentation are needed to better understand and manage the economic transition associated with AI.

Keywords: Artificial intelligence; prediction technologies; firm organisation; innovation; market structure; digital platforms; industrial policy; productivity; labour markets; global value chains; competition policy; AI governance.

Introduction

Artificial intelligence has in recent years moved from niche laboratory curiosities to broadly accessible tools that change how firms organize work, how markets form, and how entire industries assemble the hardware and energy that sustain computation. This paper sets out an account of that transformation. In line with Agrawal *et al.* (2018), we assume a simple organising insight: what modern AI fundamentally does is reduce the cost of prediction and pattern recognition at scale. When prediction costs fall, an array of economic processes is affected: firms reallocate tasks, redesign products and contracts, restructure investment in upstream inputs (chips, data centres, networks, power and minerals), and renegotiate who captures value across markets and nations. But the economic consequences are neither uniform nor inevitable. Outcomes depend on which tasks are codifiable, on firms' data and organizational capabilities, on the structure of markets and upstream supply chains, and on policy choices.

The paper proceeds in four linked parts. The first part examines the microeconomic mechanism (Section 1): how cheaper prediction translates into firm decisions about tasks, organization, business models and innovation. It reports recent evidence from field experiments and firm surveys, emphasizing the heterogeneity of effects and the importance of organizational complements. The second part traces the implications of those firm level changes for market structure (Section 2). It analyses the vertical stack that underpins modern AI – hardware, cloud, data, foundation models and applications – and explains why concentration can arise even in a technologically dynamic environment. It discusses algorithmic pricing risks and the changing economics of startup entry. The third part aggregates the micro and market effects to ask what AI means for productivity, employment and inequality at the national and global level (Sections 3 to 5). The final part considers the governance frameworks needed to shape AI's economic consequences (Section 6). It examines how competition policy, corporate accountability, industrial strategy, and labour-market institutions interact with AI-driven changes in markets and production. The section outlines emerging forms of market power, the role of public provision of critical inputs such as compute and data, and the institutional arrangements that influence how firms adopt AI and how workers adjust.

Throughout the paper we aim to summarize emerging findings and outline the mechanisms that current research suggests, while noting the limits of what can presently be inferred. In areas where evidence is thin or still evolving, we focus on plausible mechanisms and the uncertainties that remain. The concluding section synthesizes these strands, distinguishing what appears reasonably well-supported from what remains uncertain, and highlighting where further empirical work is most needed.

1. How cheaper prediction changes firm behavior

Having framed AI as a pervasive reduction in the cost of prediction, this section turns to the firm as the primary locus where those lower costs are converted into organizational change. It asks which tasks are automated or augmented, what complementary investments matter (data, governance, skills), and how these changes alter business processes, HR practices, strategy and innovation. Understanding these micro foundations is essential because firm choices determine whether algorithmic gains crystallize into productivity, fail to scale, or produce uneven outcomes across workers and firms – and it logically precedes the discussion of how those firm level shifts reshape markets.

1.1 The basic mechanism: prediction as an economic act

Prediction is the engine of many economic activities: forecasting demand, scoring credit risk, anticipating machine failure, synchronizing supply chains, selecting candidates in hiring, and generating text or code. Modern machine learning reduces the marginal cost of producing such predictions and extends the class of tasks to which robust algorithmic predictions can be applied (Agrawal *et al.*, 2018). That lower cost has two immediate economic consequences for firms. First, it changes the allocation of human effort: tasks with high codifiability and sufficient data are either automated or substantially augmented. Second, it alters the returns to investments in complementary assets – data infrastructure, organizational redesign, skills and governance. Firms that invest in these complements capture the productivity upside; those that do not often experience pilots that fail to scale or produce reputational risk.

This mechanisms-first view is supported by multiple strands of evidence. Large, randomized field experiments and extensive firm surveys converge on a pattern: when firms pair an accurate prediction engine with the necessary organizational complements, measured productivity often rises substantially; when those complements are absent, adoption produces friction or no improvement at all. A representative, multicompany study of customer service teams shows an AI assistant raising resolutions per hour by roughly 15 percent on average (Brynjolfsson *et al.*, 2025). At the same time, industrial AI projects are more likely to produce short run reconfiguration costs before they yield productivity gains – a causal pattern documented in manufacturing plants where sensor networks, model deployment and process redesign require substantial upfront change (McElheran *et al.*, 2025). Understanding both sides of this coin is essential for managers and policymakers.

1.2 Business processes: substitution, augmentation and the productivity J-curve

At the level of day-to-day operations, AI changes two economic margins. The first is substitution: algorithmic systems replace human effort in codifiable tasks – data extraction,

form filling, routine triage and structured decision-making. Robotic process automation and classical supervised learning accomplish many of these substitutions. The second margin is augmentation: AI supplies high-quality signals that enhance human decision-making – drafting first versions of legal memos, suggesting code snippets, summarizing complex documents, or proposing troubleshooting steps for field technicians. The empirical distinction between substitution and augmentation matters because it influences who gains and whether firms reallocate or shed labour.

Field evidence illustrates heterogeneity. In customer service settings, generative assistants have raised productivity by allowing workers to compose replies faster and to handle more interactions per hour (Brynjolfsson *et al.*, 2025). The effect is concentrated among less experienced agents who benefit most from high-quality drafts and structured templates, which compress the time required to reach the expected standard. By contrast, in industrial settings the introduction of predictive maintenance systems can produce an initial dip in measured productivity as shops install sensors, build data pipelines and redesign maintenance protocols; only after that reconfiguration does maintenance cost drop and equipment lifetime extend (Bousdekis *et al.*, 2019; Zonta *et al.*, 2020; McElheran *et al.*, 2025). Thus, many AI deployments follow a J-curve: short-run disruption followed by longer-run gains where firms invest in the required complements.

A simple way to think about when AI raises measured productivity quickly is to ask three questions. First, does the task have high codifiability – is the mapping from inputs to the desired outputs well represented in data? Second, does the firm have enough relevant data and the infrastructure to feed models reliably? Third, does the organizational model support integration – are decision rules, incentives and workflows adjusted so that the AI's outputs are acted upon? Where all three are answered affirmatively, the odds of rapid gains rise. When any of these conditions fail, the short run can be rocky.

The literature on predictive maintenance provides a useful exemplar of these dynamics. Predictive maintenance systems depend on sensor coverage, connectivity and historical failure records. When those components exist, some firms have reported maintenance cost reductions up to 38 percent and others report increases in equipment lifetime up to 60 percent in pilots (Bousdekis *et al.*, 2019). But these results come with caveats: the gains are path dependent, require cross functional teams, and often necessitate changing spare parts logistics and contractual incentives. In short, the algorithm is necessary but not sufficient.

1.3 Human resources, tasks and workplace design

Human resources is a primary channel through which AI reshapes work. Firms use AI for recruitment, training, onboarding, performance management and worker support. Natural language processing and screening tools reduce time-to-hire and administrative burden; adaptive learning platforms personalize training and may facilitate skill acquisition (Madanchian *et al.*, 2023; Aguinis, *et al.*, 2024; Verma *et al.*, 2023). Yet these same tools carry

risks: models trained on historical hiring outcomes can perpetuate past biases (the early Amazon recruiting example remains an object lesson), and opaque performance analytics can erode trust and fairness if not carefully governed (Sohani *et al.*, 2025; Özer *et al.*, 2024).

Several practical lessons follow for designers of HR systems. First, human-in-the-loop design matters. Good prompts, verification steps, and escalation rules substantially reduce hallucination risks in generative assistants and improve decision quality (Aguinis *et al.*, 2024). Second, governance must combine audits, transparency, technical fairness tools and clear human override paths. Third, training and role redesign are indispensable. AI tools can reduce some training time and understanding the best training method for individual employees, but realizing productivity gains at scale requires investments in curriculum design and managerial routines that reallocate saved time to higher value tasks (Sohani *et al.*, 2025).

1.4 Business models: digital servitization, data network effects and situated AI

Lower prediction costs do not only change tasks; they change what firms sell. Firms increasingly move from product sales to data rich services (digital servitization), using sensors, cloud analytics and outcome-based contracts to convert once discrete hardware into ongoing service relationships (Paiola & Gebauer, 2020). The classic example is precision agriculture, where sensor equipped equipment lets specialized providers sell uptime and yields rather than tractors and spare parts. These transformations require large installed bases, reliable sensor coverage, and contractual sophistication.

Generative models in turn enable new product architectures by lowering the marginal cost of customization. Personalized experiences, dynamic content and automated code generation allow firms to serve tailored solutions at scale. Firms that embed AI in strategic layers – where they collect unique and high-quality data – can build durable data network effects: better products draw more users and therefore more data, improving models and deepening lock-in (Costa Climent *et al.*, 2024). This pattern may yield three archetypal strategies. The “digital tycoon” pursues platform scale and continuous interaction to create a self-reinforcing flywheel: data → models → product → users → data. The “niche carvers” focus narrowly on domain-specific data and algorithms to deliver superior performance in specialized markets. The “asset augments” uses AI to extend traditional products into service bundles, often in capital-intensive sectors such as agricultural machinery or industrial equipment (Ruokonen & Ritala, 2024)

1.5. Strategic management

Strategic management now increasingly rests on an interaction between continuous, algorithmic sensing and intermittent, human judgement. Modern AI systems extend

managerial reach by ingesting and synthesizing streams of real-time data – transaction records, market prices, customer interactions, supply chain telemetry and public information – and by converting those inputs into scenario fragments that managers can test rapidly. Firms feed these algorithmic signals into systematic scenario exercises and counterfactual simulations to probe downside risks, to quantify the value of alternative investments, and to monitor competitors’ moves at a cadence that humans alone cannot sustain. The practical consequence is not that AI replaces strategy, but that it reshapes the decision cycle itself. AI systems now enable managers to triage opportunities and threats earlier, with recent evidence showing substantial empirical agreement between AI outputs and human expert assessments (Doshi *et al.*, 2025), and to iterate richer scenarios at greater speed. This frees scarce managerial attention for higher-order tasks such as framing, sense-making, and political judgment (Towers-Clark, 2024).

Algorithmic augmentation takes three concrete forms that matter for strategic practice. First, integration: AI links internal data (sales, costs, inventories) with external signals (prices, reviews, patent filings) to produce near real-time indicators of market momentum and margin pressure. Second, experimentation: generative and simulation tools let firms instantiate hundreds or thousands of plausible futures. Third, surveillance and synthesis: automated monitors detect competitor pricing updates, product changes or shifts in platform ranking. Together these capabilities allow firms to act proactively rather than reactively, shortening the interval between information and strategic response.

Strategic managers must recognize that algorithmic foresight changes what strategy looks like without changing its core epistemic demands. AI widens the range of plausible scenarios and accelerates learning from experiments; it does not remove the need for normative judgments about trade-offs, stakeholder interests and long run commitments.

1.6 Innovation activities: discovery, search and the “invention of a method of invention”

AI is already transforming R&D and innovation. Deep learning methods accelerated protein structure prediction and compressed parts of early drug discovery (Jumper *et al.*, 2021), illustrating a broader point: AI can act as an “invention of a method of invention” (IMI). In many domains, algorithmic search and pattern detection expand the effective design space; firms can explore thousands of design variants or screen millions of candidate molecules at far lower marginal cost than before (Cockburn *et al.*, 2019). Empirical manager surveys show large expected and realized gains in trend spotting, virtual prototyping and the analytical stages of innovation, while human comparative advantages persist in creativity (Füller *et al.*, 2022).

Successful R&D integration follows a common pattern. AI contributes to discovery, exploration of large parameter spaces and optimization. Humans contribute to problem

framing, interpretation, boundary crossing and moral judgment. Organizational complements – data pipelines, cross functional teams, and executive sponsorship – are decisive. Firms that build these complements in design and governance become frontrunners; those that fail to invest remain occasional experimenters. The micro foundations of innovation with AI therefore rest not only on algorithms but on managerial capabilities that convert algorithmic outputs into experiments and products.

2. How firm level change reshapes markets

This section examines the next causal layer: how widespread firm reorganization and new AI based business models change competitive dynamics and the vertical structure of the economy. It analyses the AI value chain – from chips and cloud to models and applications – and explains the forces that generate concentration, platform power and new forms of tacit coordination. Having established how individual firms reconfigure tasks and products, we now ask how those reconfigurations aggregate into market power, entry dynamics and strategic risks that matter for competition policy.

2.1. The vertical AI stack and sources of concentration

The economic and political significance of AI is shaped in no small part by the vertical stack that underpins modern capabilities. Analytically useful layers are hardware (chip design and foundry fabrication), cloud infrastructure (IaaS/PaaS/SaaS and developer tools), training data, foundation models (large pretrained models that are finetuned), and end-user applications. Each layer has different competitive dynamics and policy problems (Hagiú & Wright, 2025; Gambacorta & Shreeti, 2025).

At the upstream layer, chip fabrication is capital intensive and geographically concentrated. Advanced node foundries require enormous capital, proprietary process knowledge and an ecosystem of equipment makers; a few companies dominate advanced fabrication and accelerator design (Gambacorta & Shreeti, 2025). The U.S. and Europe industrial policies (U.S. CHIPS and Science Act; European Chips Act) and export controls on these technologies illustrate how governments view semiconductor capacity as strategic.

Cloud infrastructure concentrates as well. Hyperscalers provide not only compute but developer tools, distribution channels and billing infrastructures. These platforms internalize economies of scale in datacentre operation, network peering and software ecosystems (Hagiú & Wright, 2025; Gambacorta & Shreeti, 2025). The economic significance of control at the cloud layer lies not only in low marginal costs but also in the bundling possibilities: cloud providers can tie compute provision to data services, model hosting, and distribution, which can foreclose rival model developers and application builders if not carefully regulated.

Foundation models sit in the middle of this stack. Training frontier models is costly – estimates range from a few million to hundreds of millions of dollars – but open-source releases and model distillation could partly counterbalance concentration (Härlin *et al.*, 2023; May, 2024). Still, when firms integrate foundation models into platforms and combine them with proprietary data and user interfaces, they can extract rents by monetizing unique data flows and by setting API terms (Hagiu & Wright, 2025).

The economic forces that produce concentration are not exotic: economies of scale in hardware, network effects in platforms, data feedback loops in models, and the fixed costs of training create persistent returns to scale. Unlike classic scale economies that can sometimes be overcome by managerial ingenuity, the compute/data stack's structural features – from semiconductor physics to platform effects – make contestability difficult without public policy and market design interventions. When core inputs are costly, rival, and geographically concentrated, firms controlling them extract rent, exercise market power, and create strategic dependencies that are hard to unwind. Moreover, companies with substantial market power at one level of the vertical stack may leverage their position to limit or distort competition at other levels (Hagiu & Wright, 2025).

Ultimately, concentration in these layers matters because it shapes who captures value and who remains dependent. The structure of technological supply chains not only determines the distribution of profits but also constrains innovation, resilience, and strategic autonomy.

2.2 Startups, entry and the paradox of democratisation

Notwithstanding the tendency towards market concentration, start-ups remain a crucial route to innovation, bringing domain-specific focus and agility to AI application development. Practitioners and surveys show that many AI startups rely on cloud access, open models and finetuning to deliver value to mid-sized customers (Bessen *et al.*, 2024). This dynamic has two faces. On one hand, cloud and open frameworks materially lower entry costs for many applications and permit niche carvers to innovate rapidly. On the other hand, the very largest applications and services – those that require frontier training compute or massive proprietary datasets – remain accessible mainly to large incumbents or to firms with privileged platform partnerships (Taneja & Zakaria, 2024).

This juxtaposition explains why the startup landscape looks “democratized” at the application layer but concentrated at the foundation layer. Many startups can build differentiated products by finetuning open models or by combining domain data with cloud compute. Yet where training frontier models or controlling massive inference fleets matters for distribution and margins, scale advantages reassert themselves (Gambacorta & Shreeti, 2025). Policy, therefore, has a delicate role: it should preserve entry points for startups (through data portability standards, API access, public compute commons) without undercutting incentives for the investments that produce frontier capabilities.

2.3 Market conduct between firms: pricing algorithms

One specific and interesting issue concerns how algorithms change the nature of interactions among competing firms. Pricing algorithms can respond continuously to market conditions; recommender systems can alter demand patterns; automated decision rules can change conduct in ways that are hard to observe. Laboratory and simulation studies find that learning agents can converge to price paths that resemble human tacit collusion under particular algorithmic choices (Bichler *et al.*, 2025). Yet a careful interpretation is required. For instance, Xu *et al.* (2024) results show that many collusive outcomes in labs are fragile: the tacit collusion is commonly driven by failure to learn. This matters for regulators: instead of chasing every case of expensive prices, they should look more closely at how the learning systems themselves are built and used (Abada *et al.*, 2025).

3. AI beyond the firm: inputs, macro-outcomes and geopolitical stakes

Market level changes depend on concrete upstream inputs and systemic constraints. This section digs into the material and institutional foundations of modern AI: semiconductors, datacentres, energy systems, minerals, and telecommunications. It links firm and market behaviour to supply chain vulnerabilities, grid impacts and national strategies, and it shows why technological diffusion is as much a question of atoms and kilowatt-hours as of algorithms. This perspective naturally extends the prior market analysis toward the global arena, highlighting how input chokepoints and strategic policies shape both economic opportunity and geopolitical rivalry.

3.1 Hardware, semiconductors and the geopolitics of chips

The hardware layer is a chokepoint. Advanced chips require cutting-edge foundries, tightly integrated design tools, and specialized process gases and materials. The high capital intensity and long lead times for new fabs combine with geopolitical tensions to make chip supply both economically and strategically sensitive. The result is the emergence of a geopolitics of chips: export controls on accelerators, CHIPS style subsidy programmes, and national efforts to secure regional capacity. These measures reshape where value accumulates.

Quantitatively, industry scenarios show the scale of potential demand. McKinsey's modelling suggests generative AI workloads could drive very large additional wafer demand at advanced nodes under plausible adoption paths, implying the need for several additional advanced fabs to avoid shortages (Burkacky *et al.*, 2024). Those projections are scenario

based and sensitive to assumptions about algorithmic efficiency and architectural shifts, but they illustrate the magnitude of investment required to sustain frontier compute.

3.2 Energy, data centres and the “AI energy paradox”

AI is not immaterial. Training and inference consume electricity at scale. The IEA's Energy and AI analysis shows that datacentre electricity demand is projected to rise markedly under plausible adoption scenarios and that AI workloads will represent a growing share of that demand (IEA, 2025). At the same time, AI offers tools that can reduce energy use across systems – better grid forecasting, predictive maintenance, and faster materials discovery that can accelerate clean-energy technologies. That duality is the “AI energy paradox” (Greene-Dewasmes *et al.*, 2025).

Estimates of future energy demand are volatile and depend on two moving parts: how fast AI workloads scale and how rapidly model and hardware efficiency improves. The IEA presents scenario ranges in which datacentre electricity roughly doubles to 2030 and AI workloads are a key driver of the increase (Spencer *et al.*, 2025). These dynamics are not confined to domestic electricity systems but are also reshaping international energy markets, affecting cross-border trade in electricity, global gas demand, and the geopolitics of energy security.

The policy problem becomes how to supply this growing demand with low-carbon, dispatchable resources and how to exploit AI itself to reduce net emissions.

3.3 Minerals, recycling and the material footprint of AI

Modern AI hardware depends on a wide basket of minerals – from silicon and copper to gallium, germanium, rare earths, lithium and cobalt. Clean energy transitions and AI hardware buildouts together press demand on these minerals and create geopolitical leverage where processing and refining is concentrated. The IEA documents sharply rising demand for critical minerals and the concentration of processing capacity in a few countries (IEA, 2025). Importing economies can pursue diversification, recycling, and industrial policies that foster local processing and material substitution.

These pressures also expose an important temporal risk: mineral supply responses operate on long investment cycles, while AI-driven demand can rise much faster. Opening new mines, expanding refining capacity, and developing recycling infrastructure often require a decade or more, whereas surges in AI hardware demand can occur within a few product cycles. This mismatch creates the possibility of short-run bottlenecks, price volatility, and strategic vulnerabilities, particularly for countries with limited bargaining power in upstream supply chains.

4. Macro aggregation: productivity, employment, and distribution

This section aggregates the micro and market mechanisms to assess plausible macroeconomic pathways: how much productivity can AI deliver, how jobs are reallocated across tasks and sectors, and what distributional consequences may follow. It lays out scenario logic, reviews evidence on both augmentation and displacement, and explains why the aggregate payoff depends crucially on adoption patterns, complements and reallocation frictions. The analysis ties together the firm level J-curves and upstream constraints discussed earlier to show how short-run disruptions can diverge from longer-run gains.

4.1 Translating micro gains into macro prospects

A perennial policy question is the size of AI's macroeconomic payoff. Micro evidence shows meaningful task-level gains in many settings, yet macro estimates vary widely. The heterogeneity arises because the micro-to-macro translation requires four conditions: exposure (how many tasks can be automated or augmented), adoption (how many firms adopt at scale), complements (skills, data infrastructure, organizational redesign), and reallocation/spillovers (whether cheaper intermediate inputs raise final output rather than simply lowering prices). OECD exercises that specify plausible ranges for these conditions produce estimates of labour productivity contributions on the order of 0.4 to 1.3 percentage points per year for advanced economies under medium to fast adoption scenarios (Filippucci *et al.* 2025).

Historical experience with general-purpose technologies suggests that organizational change, regulation, and infrastructure often lag technical possibilities, delaying aggregate gains and widening cross-country dispersion. At the same time, model-based exercises rarely incorporate upstream constraints such as energy, compute, and critical minerals, or potential feedback from financial and geopolitical shocks. As a result, current macro estimates are best seen as indicative of the scale that AI could reach under certain conditions, rather than as predictions of what economies will necessarily achieve.

4.2 AI and Work: Substitution, Complementarity, and the Uncertain Path of Reallocation

Much confusion about AI arises from conflating task automation with net employment outcomes. Modern evidence points to three simultaneous forces, though their relative strength remains uncertain and context-dependent. First, AI substitutes for codifiable tasks

within jobs, reducing time spent on those activities. Second, it augments other tasks and, in many settings, may raise the marginal productivity of remaining tasks, potentially increasing demand for complementary labour. Third, firm-level productivity gains can expand output and labour demand, offsetting substitution to an extent that is still difficult to quantify.

Large scale surveys and administrative studies present a nuanced and evolving picture. Vacancy and job posting evidence suggests changing demand for skills – employers increasingly seek managerial, business process, and digital competencies in AI-exposed occupations (Green, 2024) – but it is not yet clear how persistent or widespread these shifts will be. Representative employer–worker linked administrative data currently show small average time savings and no large immediate wage declines at the national level (Humlum & Vestergaard, 2025). In contrast, randomized trials and field pilots report larger gains for targeted tasks and for less experienced workers (Brynjolfsson *et al.*, 2025), although these results are context-specific and may not generalize. Complementary firm-level evidence from manufacturing indicates J-curve dynamics with short-run productivity dips and long-run gains for firms that make substantial complementary investments (McElheran *et al.*, 2025), but the durability of these patterns across sectors is still uncertain.

Despite the sizable gains observed in targeted experiments, a key analytical challenge is explaining why national-level studies find only modest time savings and limited wage effects. The small average reductions in work hours reported in representative employer–worker data (Humlum & Vestergaard, 2025) contrast with the much larger improvements seen in controlled or early-adopter settings (Brynjolfsson *et al.*, 2025). This divergence suggests that the translation of AI capabilities into realized productivity depends heavily on the task and organizational complements – training, workflow redesign, data quality, and managerial priorities – whose adoption remains uneven and only partially understood. In many firms, these complements have not yet been made, leaving them on the downward slope of the productivity J-curve (McElheran *et al.*, 2025), but it remains uncertain how rapidly organizations will adapt or whether these investments will deliver the expected returns.

Evidence also complicates standard assumptions about who benefits from new technologies, though the patterns are still emerging. Some experimental settings show that AI tools disproportionately assist less experienced workers, narrowing performance gaps (Brynjolfsson *et al.*, 2025). If this holds more broadly, it will challenge the expectation that technological change primarily favours high-skill workers. However, these findings come from specific occupational contexts, and it is not yet clear whether AI will compress skill differentials across the full labour market. Vacancy data showing rising demand for managerial and digital competencies (Green, 2024) may reflect temporary transition dynamics rather than long-term structural shifts.

These dynamics unfold in an environment where firm-level capacity to adopt and integrate AI varies widely. Early adopters that invest in integration, oversight, and data infrastructure appear to capture productivity gains sooner (McElheran *et al.*, 2025), but it is uncertain how many firms can or will follow this trajectory. Combined with the absence of broad wage gains

in national data (Humlum & Vestergaard, 2025), this raises concerns – but not certainties – about a scenario in which productivity improvements accrue disproportionately to early adopters and their suppliers. How quickly diffusion will occur, and under what conditions it leads to broader prosperity, remains an open empirical question.

The implication is that policy responses must remain flexible and adaptive. Rapid reskilling and portable credentials matter, but their effectiveness will depend on the actual pace and direction of technological change. Active labour market programs tied to employer needs may be essential, yet their design must accommodate significant uncertainty about future skill demands. Wage insurance could help where adjustment costs are concentrated, though the magnitude and distribution of such costs are still hard to forecast. Above all, the distributional challenge is political: without deliberate institutional action, early adopters may capture a disproportionate share of the gains, while others bear adjustment burdens. However, the scale and persistence of these asymmetries remain to be seen as the technology and its uses continue to evolve.

5. The global dimension: power, policy and regional divergence

Building on the macro picture, this section asks who gains and who falls behind across countries and regions. It contrasts strategies of technological leadership, industrial policy and openness, explores possibilities for leapfrogging in lower income contexts, and explains how supply chain choices (chips, minerals, energy) produce asymmetric advantages. By connecting domestic adoption paths to international policy and finance, this part clarifies the global distributional stakes of AI and how policy choices either mitigate or magnify divergence.

5.1. Technological leadership, industrial policy and great power competition

AI's strategic stakes are geopolitical because the inputs that enable frontier capabilities – compute, chips, energy and large datasets – are concentrated and because models influence information and military systems. Different states pursue different mixes of private led innovation, state directed scaling and regulatory orientation: the United States relies on frontier private innovation with targeted industrial support; China pursues rapid scale and system integration with tight state coordination; the EU emphasizes regulatory norms, public goods and measured industrial policy (Schmid *et al.*, 2025; Huq, 2024).

These differences matter for comparative advantage. States that secure low-cost power and reliable access to advanced compute and minerals create an industrial environment

favourable to hosting data centres and AI firms (Härlin *et al.*, 2023). Export controls on accelerators and design tools change where and how models can be trained. The emergence of lower cost architectures – for example the reported efficiency claims of some Chinese firms’ models – can blunt the leverage of export controls (García-Herrero & Krystyanczuk, 2025).

Taken together, these dynamics underscore that national AI strategies operate under conditions of structural scarcity and asymmetric control over critical inputs. States must therefore navigate trade-offs between security, openness, and competitiveness: policies that restrict access to key technologies or heavily subsidize domestic capacity may strengthen short-term strategic positioning but risk weakening long-run innovation networks and scientific exchange. At the same time, the rapid evolution of model architectures and deployment patterns introduces uncertainty into the durability of current advantages. As technological trajectories shift, so too may the balance of power in the global AI landscape, making policy choices today consequential not only for industrial capacity but for the broader configuration of geopolitical influence in the years ahead.

5.2 Convergence or divergence: who can leapfrog?

Will AI converge national productivity paths or deepen divides? The answer is both – but under sharply different conditions. Low- and middle-income countries (LMICs) can achieve pockets of rapid progress by deploying AI in sectors where digital tools substitute for missing infrastructure or weak institutions. As Khan *et al.* (2024) note, targeted applications in health, agriculture, and logistics often address binding constraints directly. The Zipline case in Rwanda – where autonomous drones deliver medical supplies to remote areas – shows how focused public-private interventions can generate rapid, localized gains by overcoming geographic and administrative bottlenecks.

Yet such successes remain exceptions rather than evidence of broad convergence. The foundational inputs of the AI frontier – compute, advanced chips, energy-intensive data centres, engineering talent, and algorithmic research capacity – remain heavily concentrated in a small number of countries. Huq (2024) emphasizes that under the “darkening shadow of geopolitical competition,” access to these resources is shaped by strategic rivalry rather than cooperation. Without coordinated international support, finance, and technology transfer, LMICs risk being locked into structural dependency, with early pilots failing to scale into sustained productivity growth.

For developing countries, the strategic challenge is therefore dual. Short-term pilots must demonstrate clear public value – improving crop yields, streamlining medical logistics, detecting disease, or automating administrative tasks – to build political support and mobilize local ecosystems. But these efforts must be paired with long-term investments in absorptive capacity: education and technical skills, digital and energy infrastructure,

institutional quality, and legal frameworks for data governance that balance privacy with responsible sharing.

External actors play an essential enabling role. High income countries, development banks, and multilateral institutions can accelerate catch-up by providing public compute, shared datasets, finance for local processing capacity, and technical assistance. Without cooperation the risks of divergence remain high.

5.3 Global value chains: reshoring, dependencies and resilience

AI is reshaping where value is captured along global value chains (GVCs). As advanced robotics and increasingly capable AI systems erode the cost advantages of low-wage labour, parts of production that were previously offshored – especially routine assembly and certain supervisory tasks – can become economically viable to reshore. This shift reflects a broader transition in which computational resources and autonomous machines substitute for human labour, reducing the significance of cross-border wage differentials (Korinek, 2024).

At the same time, AI introduces new structural dependencies in the most technologically intensive layers of the supply chain. Highly concentrated markets for crucial inputs – such as specialised chips fabricated overwhelmingly by TSMC and GPUs dependent on Nvidia's CUDA ecosystem – expose economies to single points of failure. Similar concentration characterises cloud computing, where a small group of providers controls most global infrastructure essential for training and deploying AI models (Gambacorta & Shreeti, 2025). The net effect is a mixed pattern: partial reshoring of labour-intensive segments combined with tighter dependency on a narrow set of hardware and cloud suppliers.

These dynamics complicate national strategies for resilience. Governments increasingly combine industrial finance, strategic stockpiles of advanced chips, support for domestic fabrication capacity, and the creation of public compute resources to ensure broader access to critical inputs. Policies aimed at interoperability, open standards, and diversified supplier bases attempt to mitigate lock-in and reduce switching costs. Yet each intervention carries trade-offs in terms of fiscal cost, efficiency loss, or potential distortions in innovative ecosystems.

The result is that GVC restructuring has become a geopolitical question as much as an economic one. Decisions about which parts of the AI stack to develop domestically versus through trusted alliances shape long-term technological sovereignty. Likewise, standards for cross-border data flows and rules for accessing foundation models influence how countries position themselves within an increasingly bifurcated technological landscape. In this context, the strategic choices made today determine not only resilience but also future distribution of economic returns.

6. Governance, competition policy and democratic resilience

The economic consequences of AI will depend as much on institutions as on algorithms. Governance in this sense is the set of public and private rules, incentives and organizations that shape how firms invest, how markets allocate resources, and how society shares gains and absorbs dislocations. The paragraphs below outline an economically grounded agenda for governance that emphasizes competition policy, corporate accountability and market design, industrial strategy and public provision of critical inputs, labour market institutions and social insurance, and regulatory experimentation and evaluation. Where possible we draw on recent empirical and theoretical work to identify feasible levers and the trade-offs they imply.

6.1 Competition policy for an AI economy

Competition policy remains an important instrument for shaping how AI affects market structure and consumer welfare, but its effectiveness now hinges on recognising forms of market power that arise before overt exclusion occurs. Vertical foreclosure, tying, exclusive contracts and the removal of nascent rivals still matter, yet in AI markets these behaviours are often embedded in technical governance decisions – access rules, API design, data-usage terms – rather than in explicit contractual restrictions (Hagiú & Wright, 2025; Gambacorta & Shreeti, 2025).

A central challenge is that the boundaries of competitive harm become harder to define when firms organise production through tightly integrated data-model-distribution bundles. These arrangements can shape market trajectories without any single action clearly amounting to abuse, making it essential for authorities to develop the technical capability to audit the incentives and constraints embedded in AI development pipelines. Market power may thus be exercised through architectural control rather than through traditional exclusionary strategies.

This also complicates remedy design. Many competitive risks in AI arise from dependency and lock-in – high switching costs, opaque integration layers, or the inability to port models and data – rather than from explicit forms of coercion. Addressing such structural vulnerabilities requires policy tools oriented toward contestability rather than only policing misconduct, including obligations around portability, interoperability and transparency (Gambacorta & Shreeti, 2025).

For these reasons, rigorous merger review and abuse investigations must be complemented by remedies tailored to the platformed, vertical character of AI markets. Effective competition policy for an AI economy requires shifting attention from isolated violations to the institutional and technical conditions that determine who can enter, who can switch, and

who ultimately controls the pathways through which AI capabilities are accessed and deployed.

6.2 Corporate governance, accountability and market design

Firms that build and deploy AI combine engineering choices with business models and governance practices. Economic governance should therefore attend to corporate structures, contracting practices, and incentives that influence both the creation and the diffusion of AI (Cordeiro *et al.*, 2026). Boards and senior management must be accountable for technology risk and for the design choices that shape market outcomes; regulators and investors can align these incentives by clarifying fiduciary duties with respect to long-run systemic risks, requiring disclosure of model risk management, and mandating incident reporting for high impact systems. Market design tools – from standard contracting templates that specify provenance and licensing of model training datasets to conditional procurement rules that tie public contracts to interoperability commitments – can shift market incentives toward openness where social gains justify. At the same time, corporate governance should preserve entrepreneurial experimentation; a pragmatic regulatory stance combines targeted disclosure and auditability requirements for high impact uses with permissive regimes for low-risk innovation.

6.3 Industrial strategy and public provision of critical inputs

Markets alone are unlikely to deliver the large-scale infrastructure, coordination, and knowledge spillovers needed for an innovative and broadly diffused AI ecosystem. Frontier model development depends on inputs – compute, data, and skills – with high fixed costs, strong economies of scale, and substantial externalities, all of which tend to be underprovided by private actors. Public policy therefore has a role in underwriting these lumpy investments and creating shared inputs that lower effective entry costs.

A first priority is access to compute. Training and deploying advanced models requires expensive, scarce GPU/TPU clusters, and compute access remains highly concentrated among firms with capital and preferential access to hardware (Chardon-Boucaud *et al.*, 2024). Public provision of non-commercial compute – through supercomputers and subsidised access for research teams and SMEs – is designed to ease these financial and technical barriers. This implies that expanding publicly accessible compute reduces advantages based purely on capital-intensive access to raw compute and increases contestability.

A second priority is strengthening absorptive capacity through investments in skills and data commons. Public compute and curated public datasets complement private investment by mitigating key bottlenecks and enabling a wider range of actors to engage in AI

development and adoption, provided they are accompanied by fair access and governance standards. An effective industrial strategy should therefore pair targeted subsidies for strategic upstream capacity (e.g., semiconductor foundries, advanced fabrication) with policies that secure broad access to the inputs that amplify competitive entry and support economy-wide diffusion and social value capture.

Therefore, states can become a prime mover that coordinates an ecosystem of agents (public research, firms, SMEs, civil society) toward shared societal objectives – thereby increasing the potential for broad diffusion, collective learning and absorptive capacity. Targeted applications of AI offer a powerful tool for addressing societal challenges – such as health, environment, or social inclusion – by unlocking new capabilities and overcoming institutional bottlenecks (e.g., limited public-sector capacity, fragmented infrastructures, chronic budget constraints, and the limited information and coordination failures that hinder effective responses to complex problems). By deploying shared compute, curated datasets and open innovation platforms under public direction, AI can mobilize heterogeneous actors to converge on common problems, reduce setup costs, accelerate experimentation and enable systemic transitions.

6.4 Labor market institutions and social insurance

The distributional consequences of AI hinge on how labour markets reallocate workers and how institutions smooth the adjustment. Economic governance should combine active labour market policies – rapid reskilling, portable credentials, and employer linked apprenticeship models – with social insurance that insulates workers during transitions. Wage insurance, time limited income support linked to retraining, and subsidies for on-the-job learning can reduce resistance to productive reallocation while preserving incentive compatibility. Collective bargaining institutions and firm level codetermination can also play a constructive role: where workers participate in deployment decisions and receive training and profit-sharing, adoption can be less disruptive and more productive.

Yet even well-designed labour-market tools face practical limits when adoption is uneven and adjustment frictions differ across regions, sectors and firms. Many employers – particularly SMEs – lack the resources to anticipate skill needs or provide structured training, and regional disparities in training systems can widen gaps in who benefits from new technologies. These patterns underscore that the capacity of the state and other institutions to coordinate programs consistently, and to narrow differences in institutional quality and support across the economy, is central to whether adjustment mechanisms translate into more balanced outcomes rather than reinforcing existing divergence.

Conclusion: what we know, what we don't and where to go next

What do we know with confidence? First, modern AI reduces the cost of prediction and information processing in many tasks. Field experiments and firm surveys show substantive gains in specific settings – customer service, software engineering, certain R&D activities – when firms invest in complementary data pipelines and organizational change. Second, the distribution of gains is uneven. Early gains concentrate in digitally mature firms, knowledge intensive sectors and workers with access to training and tools. Third, the upstream stack matters: concentration in chips and cloud compute creates strategic bottlenecks and influences market power and geopolitical rivalry.

What remains uncertain? Long run aggregate magnitudes. Micro gains seem real; the translation into national productivity growth depends on adoption, complements and spillovers. Second, the net employment and inequality effects depend on firm behaviour, labour market institutions and policy. Early evidence suggests modest short-run wage effects but large heterogeneity in task and occupation outcomes. Finally, energy and mineral demands of scaled AI are substantial and must be integrated into energy and industrial policy to avoid lock-in to carbon intensive paths.

What should researchers and policymakers prioritize? Five research priorities stand out. First, causal, longitudinal micro evidence on how generative AI affects firm productivity, wages and firm survival across sectors and countries. Second, better measurement of AI intensity. Researchers should standardize measures of task exposure, adoption intensity and model reliance to compare across countries. Third, energy and material accounting. We need transparent data on model training and inference energy footprints, and robust lifecycle analyses for hardware. Fourth, competition and governance tools need empirical validation. We must move beyond theory to evaluate remedies in case studies and pilot regulations. Fifth, international policy experiments on distributed public compute and shared data commons for public research should be evaluated for their effects on innovation diffusion and inequality.

As a set of practical priorities for managers and policymakers, the mechanisms documented in this paper suggest sensible actions. Managers should scan their workflows for tasks with high prediction value, invest first in reliable data pipelines and human-in-the-loop governance, pilot narrow use cases designed to scale, and treat AI deployments as organizational change programs rather than technology procurement. Policymakers should invest in public compute for research and small firms, fund complementary clean energy and mineral processing projects where strategic dependencies exist, build standards for data portability and algorithmic auditability, and design labour market policies that facilitate reskilling and portability of credentials.

Taken together, these findings point toward a pivotal but contingent technological moment. AI's economic promise is neither automatic nor evenly distributed: it is mediated by institutional capacity, market structure, energy systems, and the choices that firms and governments make now. The challenge ahead is to convert localized productivity gains into broad-based, sustainable growth while managing distributional risks and geopolitical constraints. Doing so will require a blend of rigorous empirical research, strategic public investment, and adaptive governance frameworks that evolve with the technology itself. If researchers, managers, and policymakers align around evidence-based experimentation and long-term capacity building, AI can become not just a source of efficiency, but a driver of more resilient, inclusive, and strategically secure economic systems.

Funding

DINÂMIA'CET-Iscte Working Papers are funded by Fundação para a Ciência e a Tecnologia (UID/3127/2025)

References

Abada, I., Harrington, J. E., Jr., Lambin, X., & Meylahn, J. M. (2025). Algorithmic collusion: Where are we and where should we be going? *SSRN*. <https://doi.org/10.2139/ssrn.4891033>

Aguinis, H., Beltrán, J. R., & Cope, A. (2024). How to use generative AI as a human resource management assistant. *Organizational Dynamics*, 53(1), 101029. <https://doi.org/10.1016/j.orgdyn.2024.101029>

Agrawal, A., Gans, J., & Goldfarb, A. (2018). *Prediction machines: The simple economics of artificial intelligence*. Harvard Business Review Press.

Bessen, J. E., Impink, S. M., & Seamans, R. (2025). The business of AI producing startups: Evidence from a worldwide survey. In N. Jia & F. Csaszar (Eds.), *Handbook on artificial intelligence and strategy (forthcoming)*. *SSRN working paper*. <https://doi.org/10.2139/ssrn.3293275>

Bichler, M., Durmann, J., & Oberlechner, M. (2025). Algorithmic pricing and algorithmic collusion. *Business & Information Systems Engineering*. Advance online publication. <https://doi.org/10.1007/s12599-025-00965-z>

Bousdekis, A., Lepenioti, K., Apostolou, D., & Mentzas, G. (2019). Decision making in predictive maintenance: Literature review and research agenda for Industry 4.0. *IFAC PapersOnLine*, 52(13), 607–612. <https://doi.org/10.1016/j.ifacol.2019.11.226>

Brynjolfsson, E., Li, D., & Raymond, L. (2025). Generative AI at work. *The Quarterly Journal of Economics*, 140(2), 889–942. <https://doi.org/10.1093/qje/qjae044>

Burkacky, O., Pototzky, K., Patel, M., Tang, D., Vrijen, R., & Zhu, W. (2024, March 29). *AI in semiconductor manufacturing: The next S curve?* McKinsey & Company.

Cockburn, I. M., Henderson, R., & Stern, S. (2019). The impact of artificial intelligence on innovation. In A. Agrawal, J. Gans, & A. Goldfarb (Eds.), *The economics of artificial intelligence: An agenda* (pp. 115–146). University of Chicago Press. <https://doi.org/10.7208/chicago/9780226613475.003.0004>

Cordeiro, C. M., Adomaitis, L., & Huang, L. (2026). The AI-policy-governance nexus: How regulation and AI shift corporate governance toward stakeholders. *Technology in Society*, 84, Article 103117. <https://doi.org/10.1016/j.techsoc.2025.103117>

Costa Climent, R., Haftor, D. M., & Staniewski, M. W. (2024). AI enabled business models for competitive advantage. *Journal of Innovation & Knowledge*, 9, 100532. <https://doi.org/10.1016/j.jik.2024.100532>

Doshi, A. R., Bell, J. J., Mirzayev, E., & Vanneste, B. S. (2025). Generative artificial intelligence and evaluating strategic decisions. *Strategic Management Journal*, 46(3). <https://doi.org/10.1002/smj.3677>

Filippucci, F., Gal, P., Laengle, K., & Schief, M. (2025). Macroeconomic productivity gains from artificial intelligence in G7 economies. *OECD Artificial Intelligence Papers*, 41. OECD Publishing. <https://doi.org/10.1787/a5319ab5>

Füller, J., Hutter, K., Wahl, J., Bilgram, V., & Tekic, Z. (2022). How AI revolutionizes innovation management – Perceptions and implementation preferences of AI based innovators. *Technological Forecasting & Social Change*, 178, 121598. <https://doi.org/10.1016/j.techfore.2022.121598>

Gambacorta, L., & Shreeti, V. (2025). The AI supply chain. *BIS Papers*, No. 154. Bank for International Settlements.

Green, A. (2024). Artificial intelligence and the changing demand for skills in the labour market. *OECD Artificial Intelligence Papers*, No. 14. OECD Publishing. <https://doi.org/10.1787/88684e36en>

Hagi, A., & Wright, J. (2025). Artificial intelligence and competition policy. *International Journal of Industrial Organization*, 83, 103134. <https://doi.org/10.1016/j.ijindorg.2025.103134>

Härlin, T., Björnsson Rova, G., Singla, A., Sokolov, O., & Sukharevsky, A. (2023, April 26). *Exploring opportunities in the generative AI value chain*. McKinsey Digital.

Huq, A. Z. (2024). The geopolitics of digital regulation (*Public Law and Legal Theory Working Papers No. 868*). University of Chicago Law School.

Humlum, A., & Vestergaard, E. (2025). The unequal adoption of ChatGPT exacerbates existing inequalities among workers. *Proceedings of the National Academy of Sciences*, 122(1), e2414972121. <https://doi.org/10.1073/pnas.2414972121>

- International Energy Agency (IEA). (2025). *Energy and AI: World Energy Outlook Special Report*. IEA.
- Madanchian, M., Taherdoost, H., & Mohamed, N. (2023). AI based human resource management tools and techniques: A systematic literature review. *Procedia Computer Science*, 229, 367–377. <https://doi.org/10.1016/j.procs.2023.12.039>
- May, R. (2024). Artificial intelligence, data and competition (*OECD Artificial Intelligence Papers No. 18*). OECD Publishing. <https://doi.org/10.1787/e7e88884>
- McElheran, K., Yang, M. J., Kroff, Z., & Brynjolfsson, E. (2025). The rise of industrial AI in America: Microfoundations of the productivity J curve(s) (*Working Paper No. 25 27*). U.S. Census Bureau.
- Özer, M., Perc, M., & Suna, H. E. (2024). Artificial intelligence bias and the amplification of inequalities in the labor market. *Journal of Economy Culture and Society*, 69, Article e1415085. <https://doi.org/10.26650/JECS2023-1415085>
- Paiola, M., & Gebauer, H. (2020). Internet of things technologies, digital servitization and business model innovation in B2B manufacturing firms. *Industrial Marketing Management*, 89, 245–264. <https://doi.org/10.1016/j.indmarman.2020.03.009>
- Ruokonen, M., & Ritala, P. (2024). How to succeed with an AI first strategy? *Journal of Business Strategy*, 45(6), 396–404. <https://doi.org/10.1108/JBS-08-2023-0178>
- Schmid, S., Lambach, D., Diehl, C., & Reuter, C. (2025). Arms race or innovation race? *Geopolitical AI development. Geopolitics. Advance online publication*. <https://doi.org/10.1080/14650045.2025.2456019>
- Sohani, S. S., Pandey, J., Varma, A., & Ray, P. (2025). Is it necessary? A framework for assessing the utility of A.I. in HRM practices. *Acta Psychologica*, 254, 104816. <https://doi.org/10.1016/j.actpsy.2025.104816>
- Taneja, H., & Zakaria, F. (2024, October 28). Can startups thrive in an age of AI? *Harvard Business Review*.
- Towers Clark, C. (2024, April 19). *3 ways to use AI for strategic decision making*. Medium.
- Verma, P., Islam, M., Patel, P., Malik, A., Budhwar, P., & Gupta, S. (2023). AI augmented HRM: Literature review and a proposed multilevel framework for future research. *Technological Forecasting and Social Change*, 193, 122645. <https://doi.org/10.1016/j.techfore.2023.122645>
- Xu, Z., & Zhao, W. (2024). On mechanism underlying algorithmic collusion. *arXiv*. <https://doi.org/10.48550/arXiv.2409.01147>
- Zonta, T., da Costa, C. A., Righi, R. da R., de Lima, M. J., da Trindade, E. S., & Li, G. P. (2020). Predictive maintenance in the Industry 4.0: A systematic literature review. *Computers & Industrial Engineering*, 150, 106889. <https://doi.org/10.1016/j.cie.2020.106889>