



CYBORG
ECONOMICS

CYBORG ECONOMICS | FLAGSHIP PAPER

THE TRILLION-DOLLAR OPPORTUNITY

The economics of enterprise AI readiness

From AI potential to realised value

Nicholas Herrick | Founder, Cyborg Economics

September 2026

Expert Economics for the AI era.

Contents

Executive summary.....2

1. From AI access to realised value.....4

2. AI readiness as a value-conversion capability.....6

3. Where AI value conversion breaks down.....8

4. Estimating GenAI productivity potential.....11

5. From productivity potential to value at stake.....14

6. Sensitivity and interpretation.....17

7. Wider economic effects beyond the headline estimate.....19

8. Distribution of gains and losses.....23

9. Measuring AI value conversion.....28

10. Limitations and alternative approaches.....34

11. Conclusion.....38

Notes and sources.....40

Executive summary

Powerful AI models and AI-enabled software are increasingly available across organisations. As access broadens, the economic question increasingly shifts from access to conversion: whether AI capability changes economic performance in ways that can be sustained.

Evidence from 2026 already shows the distinction. McKinsey reports widespread improvements in individual productivity alongside less widespread enterprise-level financial impact.¹ PwC finds that AI-driven returns remain highly concentrated among leading organisations, with the top-performing 20% in its study capturing 74% of measured returns and showing a much greater tendency to redesign workflows around AI.²

This paper describes the movement from AI capability to sustained organisational performance as **AI value conversion**. AI readiness is the organisational capacity to convert AI potential into sustained, measurable performance. It depends on connected capabilities across strategy and operating arrangements, people and adoption, data and technology, risk and assurance, commercial relationships, and measurement and learning.

AI access creates activity. AI readiness shapes conversion.

Readiness works as a system. Strong technology creates greater value when organisations also have worthwhile use cases, usable data, effective workflows and people capable of applying it well. Measurement allows successful applications to be identified and scaled, while organisational learning improves subsequent deployment. Weakness in one area can constrain otherwise strong capability elsewhere.

The economic value at stake is potentially substantial. Cyborg Economics estimates that incomplete GenAI value conversion could leave **around \$1 trillion per year of productivity-equivalent value at stake**. The headline estimate captures the labour-productivity channel. AI can also create value through innovation, new products and services, knowledge spillovers and improved public-sector outcomes, while implementation, adjustment and wider social effects influence the overall economic result.

Organisational readiness influences an important and actionable part of the conversion gap. The quality of the underlying technology, the economics of particular use cases, implementation constraints, capital, regulation and the wider environment in which adoption occurs also affect the amount of AI potential ultimately realised.

The paper separates four analytical questions. **Economic Opportunity** asks where AI can create worthwhile value. **AI Readiness** examines whether the organisation has the capabilities required to pursue that opportunity. **Value Realisation** establishes what economic value is subsequently generated. **Value Capture** examines where that value ultimately accrues.

Together, these questions provide a practical framework for moving from technical capability to economic outcomes.

For leaders

AI portfolios should be managed around economic value. Senior management needs to identify the strongest opportunities, understand which organisational constraints are holding them back, direct capital and management effort towards applications with strong economics, and use performance evidence to decide what should be scaled, redesigned or stopped.

As access to capable AI broadens, organisational capability becomes a more important source of advantage: the ability to redesign work, integrate new technology, learn from deployment and reproduce successful approaches across the enterprise.

For boards and investors

AI spending provides limited evidence about future returns. The stronger test is whether management can repeatedly turn improving AI capability into measurable operating performance and capture enough of the resulting value to justify investment.

Conversion capability therefore matters for capital allocation, competitive position and valuation. Firms facing similar technological opportunities can generate very different economic outcomes.

For policymakers

AI productivity across the economy will depend partly on the conditions in which firms adopt and scale the technology. Skills, infrastructure, access to capital, technology diffusion and proportionate assurance can all affect how quickly useful AI applications move into productive use.

The public sector has its own value-conversion challenge. Effective use of AI can improve the speed, quality and accessibility of public services, strengthen analysis and decision-making, and allow existing resources to deliver better outcomes. The value of these improvements can extend well beyond the labour savings captured in the headline estimate.

Regulation also influences conversion. Effective controls can reduce material harms and create confidence in higher-value deployment. Excessive constraints can suppress productive applications and the economic and social gains associated with them.

Successful AI conversion can therefore contribute both to higher economy-wide productivity and to better public services. Higher productivity expands the economic resources available to support living standards, investment and wider social priorities.

The trillion-dollar opportunity extends across firms, public institutions and whole economies. It is the opportunity to convert rapidly improving AI capability into higher productivity, better services and sustained economic and social value.

1. From AI access to realised value

As access broadens, the economic question increasingly shifts from access to conversion: whether AI capability changes economic performance.

Powerful AI models and AI-enabled software are increasingly available across organisations. An organisation can acquire licences, run pilots and encourage employees to experiment while achieving only limited change in its economic results. Value emerges when AI changes economically useful activity: increasing productive capacity, reducing costs, improving decisions or quality, or enabling additional revenue and new products or services.

The distinction is already visible in the evidence. McKinsey's 2026 global survey reports that 80% of respondents say AI has improved their individual productivity, while 37% attribute at least some positive enterprise-level contribution to earnings before interest and tax (EBIT).¹ The two measures are not directly comparable, but together they illustrate the difference between widespread local productivity effects and less widespread enterprise-level financial impact.

This paper describes the movement from AI capability to sustained organisational performance as AI value conversion.

The AI value-conversion chain



Measurement and learning also feed back into earlier stages of the chain. Evidence from deployment affects which use cases receive further investment, how applications and workflows are redesigned, and whether an approach is scaled, modified or stopped.

Technical capability creates the opportunity at the start of the chain. Economic value depends on what happens through the stages that follow.

A capable AI application aimed at a low-value problem has limited economic significance. A valuable use case may stall because the required data are unavailable or because it does not fit the way work is organised. A successful deployment may generate local productivity improvements without those gains becoming sustained organisational performance.

The paper distinguishes AI potential, value realisation and value capture within the wider process of AI value conversion. AI potential describes economically useful capability that could generate benefits. Value realisation describes the attributable economic value actually generated as AI changes organisational performance. Value capture describes where that realised value ultimately accrues—for example to firms, workers, consumers or government. The distribution of value capture is considered separately in Section 8.

Measures such as licences, user numbers, pilots and demonstrations provide evidence of AI activity. The stronger economic test is whether AI changes organisational performance in ways that can be measured and sustained.

This distinction becomes more important as AI capability becomes easier to obtain. When competing organisations can access similar models and tools, access itself provides less differentiation. Their ability to deploy valuable applications, integrate them into work and convert technical capability into results becomes more important.

That capability is AI readiness.

AI access creates activity. AI readiness shapes conversion.

The slogan captures the central proposition of the paper. In analytical terms, readiness shapes how effectively organisations move through the value-conversion chain.

2. AI readiness as a value-conversion capability

AI readiness is the organisational capacity to convert AI potential into sustained, measurable performance.

It covers the organisational capabilities required to implement worthwhile AI applications, deploy them into real work, manage the conditions needed for reliable use, scale successful approaches and learn from evidence about their performance.

Readiness works as a connected organisational system. Strength in one capability can be constrained by weakness elsewhere: technology depends on usable data, adoption depends on workable processes, and scaling depends on ownership, assurance and measurement.

The framework groups AI readiness into six capability areas.

Capability	Role in AI value conversion
Strategic and operating readiness	Establishes priorities, ownership, decision rights, operating arrangements and routes to scale for selected AI opportunities.
People and adoption readiness	Enables employees and managers to use AI effectively, exercise appropriate judgement and adapt roles and working practices around new capabilities.
Data and technology readiness	Provides the data, infrastructure, integration and deployment foundations required for reliable AI applications.
Risk, assurance and governance readiness	Provides proportionate controls, validation, accountability and oversight so that valuable applications can operate and scale with acceptable risk.
Commercial and ecosystem readiness	Ensures that models, platforms, suppliers and commercial arrangements provide suitable capability at sustainable cost while managing dependencies and switching costs.
Measurement and learning readiness	Provides responsibilities, analytical capacity and routines for generating and using evidence about AI performance and enables learning across projects.

The six categories provide a practical analytical structure. Sector-specific frameworks may divide the same capabilities differently. The economic test is whether they capture the conditions affecting movement from technical capability to realised performance.

2.1 Readiness works as a system

The capabilities reinforce and constrain one another.

Better data have greater value when organisations have worthwhile use cases and workflows capable of using them. Employee AI skills become more economically useful when suitable tools are available and managers can change how work is organised. Well-designed assurance can support adoption by creating the confidence required to deploy AI in higher-value settings. Measurement allows organisations to distinguish applications worth scaling from those that should be redesigned or stopped.

A weakness in one part of the system can therefore constrain otherwise strong capability elsewhere. Two organisations with similar levels of AI activity may have very different prospects for value conversion because their bottlenecks lie in different places. The interaction between capabilities matters more than a checklist of isolated scores.

PwC's 2026 AI Performance Study provides evidence consistent with this systems view. The study surveyed 1,217 senior executives across 25 sectors and found that the top-performing 20% of organisations captured 74% of the AI-driven returns measured in the study. Leading organisations were also twice as likely to redesign workflows to incorporate AI rather than simply adding AI tools to existing processes.²

Stronger financial outcomes were associated with broader organisational practices spanning workflow redesign, deployment, data, technology, governance and trust. PwC's analysis describes relationships within its survey; it does not identify the causal effect of any single readiness capability.

2.2 Readiness develops through deployment and learning

AI readiness evolves as organisations deploy the technology. It is a dynamic capability rather than a fixed state reached before deployment.

Applications reveal weaknesses in data, integration and operating processes. Practical experience shows where controls create confidence and where they impede useful deployment. Measurement shows whether anticipated gains survive contact with real operations. Organisations can then standardise approaches that work, improve weak foundations and redirect resources.

Learning therefore forms part of readiness itself. Organisations able to absorb evidence from deployment can improve their ability to move subsequent applications through the conversion chain.

The economic significance lies in repeatability. Isolated AI successes can create value, but sustained enterprise performance requires an organisation capable of repeatedly implementing, integrating, governing, learning from and scaling useful applications.

3. Where AI value conversion breaks down

The value-conversion chain helps explain why substantial technical potential can coexist with uneven economic results.

Four mechanisms are particularly important: weak use-case economics; incomplete workflow and organisational adaptation; constraints in data, technology and assurance; and weak measurement and learning.

These four mechanisms are not a second readiness taxonomy. They describe four points at which value conversion can fail. The first concerns the economics of the opportunity itself; the remaining mechanisms arise principally from readiness and value-realisation constraints. The mechanisms can arise separately, but their effects often reinforce one another.

3.1 Weak use-case economics

Value conversion begins with the economic quality of the use case.

The ease of experimenting with AI allows organisations to pursue applications because they are technically feasible or impressive rather than because they address valuable work. A functioning application can consequently generate little economic return.

A strong use case has a credible mechanism linking AI capability to an economic outcome. The scale and frequency of the affected activity matter, alongside the resources required to implement the change and the likely benefit if it succeeds. Applications aimed at costly, frequent or capacity-constrained activities can support much larger gains than technically similar applications applied to marginal work.

The first potential break in the conversion chain is straightforward: technical success applied to an economically weak problem.

3.2 Incomplete workflow and organisational adaptation

Productivity gains become economically valuable when organisations can use the productive capacity they release.

Consider a task that takes five hours before AI and three hours afterwards. The immediate result is a two-hour time saving. Those two hours may allow more work to be completed, reduce waiting times, improve quality or release skilled staff for more valuable activity. They may also remain unused or be absorbed by checking, rework and coordination.

A time saving creates potential productive capacity. Realised economic value depends on what happens to that capacity.

Workflow design therefore affects conversion. Responsibilities may need to change, tasks may need to be reordered and human effort may need to shift towards judgement, exceptions or higher-value activity. Employee adoption matters for the same reason: technical capability produces little operational change when it remains outside normal work.

McKinsey's 2026 organisational-readiness research illustrates how individual adaptation can move ahead of organisational change. In its global survey of 750 employees and leaders already incorporating AI into their work, 70% reported feeling personally prepared to adopt and use AI. Organisational-readiness questions were answered by a subset of 608 leaders, of whom only 27% believed their organisations were ready to make the wider people and cultural changes required.³ McKinsey cautions that the survey was deliberately weighted towards organisations at more advanced stages of AI maturity and should not be interpreted as an estimate of overall market prevalence.

Local productivity gains can therefore coexist with relatively limited organisational adaptation. Enterprise conversion requires the workflow and operating environment around the technology to change as well.

3.3 Data, technology and assurance constraints

AI applications operate within existing organisational systems.

Their usefulness depends on access to relevant information, sufficient data quality, appropriate technical integration and reliable deployment. Weakness in these foundations can reduce output quality or prevent otherwise useful applications from entering live processes.

Assurance influences the same economic pathway. Organisations need sufficient confidence in the reliability and control of AI applications before deploying them into important processes. Well-designed assurance can make higher-value adoption possible; cumbersome arrangements can slow deployment and learning.

Together, these conditions determine whether a technically feasible use case can become a reliable operating process.

3.4 Weak measurement and learning

The final mechanism concerns the organisation's ability to establish whether value has actually been created.

A credible assessment requires a defined baseline, a reasonable counterfactual for what would have happened without the AI intervention, and evidence on the resources required to achieve the observed change.

The sophistication of that counterfactual should be proportionate to the decision. A small operational trial may require a simple before-and-after comparison or appropriate comparator group. A large investment decision may justify more rigorous evaluation. In either case, the purpose is to distinguish the effect associated with AI from changes that would have occurred anyway.

This matters because productivity effects and realised economic value are different quantities. Time savings may create additional capacity without reducing expenditure. Quality improvements may generate value through fewer errors, better decisions or improved service rather than immediate cost reductions. Measurement therefore needs to follow the economic mechanism through which the application is expected to create value.

Measurement also affects subsequent investment. Evidence allows organisations to scale applications generating strong returns, redesign those with fixable weaknesses and stop those whose economics remain poor. Learning from these decisions improves future use-case selection and deployment.

3.5 How the mechanisms relate to the readiness framework

The four mechanisms relate to the readiness framework in different ways.

Weak use-case economics concerns the quality of the opportunity being pursued rather than organisational readiness itself. Strategic and operating readiness shapes ownership, prioritisation processes and routes to scale once opportunities have been selected. People and adoption readiness influences whether workflows and behaviour change around the technology. Data and technology readiness, together with risk, assurance and governance readiness, shapes whether applications can operate reliably. Measurement and learning readiness concerns the organisation's capacity to generate and use evidence about AI performance; evaluation of the benefits and costs of an individual investment sits within Value Realisation. Commercial and ecosystem readiness influences each stage where value depends on external models, platforms, suppliers or contractual arrangements.

The six readiness capabilities describe organisational conversion capability. The four mechanisms describe points at which the wider conversion process can break down, including weak opportunity economics as well as readiness and evaluation constraints.

COMMON BREAKDOWN POINTS

01

Weak use-case economics

02

**Incomplete workflow and
organisational adaptation**

03

**Data, technology and assurance
constraints**

04

Weak measurement and learning

Together, these mechanisms explain why technical potential can fail to become realised value. They do not tell us how large the global conversion opportunity is.

Sections 4–6 address that separate question by establishing the economic base, estimating the share of work exposed to generative AI, applying productivity evidence and estimating the share of that potential that remains unrealised.

4. Estimating GenAI productivity potential

Labour productivity is one of the main channels through which generative artificial intelligence (GenAI) can create economic value. This model isolates that channel to construct a transparent global estimate of the scale of potential value currently at stake.

The complete model is:

$$\text{Annual productivity-equivalent value at stake} = \text{world GDP} \times \text{labour income share} \times \text{GenAI occupational exposure} \times \text{effective productivity uplift} \times \text{conversion gap}$$

The first two terms establish the economic base. Three modelling layers are then applied. Exposure measures how much economically valuable work is susceptible to GenAI. Productivity estimates the additional productive capacity GenAI could generate across that work. Conversion estimates the share of that potential that remains unrealised.

Sections 4.1–4.3 establish the economic base and estimate exposure and productivity. Section 5 applies the conversion layer. Section 6 tests the three modelling assumptions in the same sequence.

All monetary values are expressed in current 2025 US dollars. Figures in the main text are rounded for readability; calculation boxes retain sufficient precision to reproduce the results.

4.1 Global labour-income base

Global gross domestic product (GDP) was approximately \$118 trillion in 2025 across the four World Bank income groups used in this analysis.¹ The International Labour Organization (ILO) estimates that labour received approximately 53% of global GDP in 2025.²

Using the underlying values:

Global labour-income base

$$\$118.36\text{tn} \times 52.6\% = \$62.26\text{tn}$$

This gives an estimated global labour-income base of approximately \$62 trillion.

Because the model focuses on labour-productivity effects, it uses labour income rather than total GDP as the valuation base. Applying the labour-income share to global GDP estimates the annual compensation associated with labour. The model then estimates how much of that labour income is associated with GenAI-exposed occupations and the additional productive capacity that GenAI could generate across that work.

In this paper, productivity-equivalent value is the monetary proxy produced by applying an assumed productivity uplift to that labour-income base. It measures additional productive capacity in labour-income terms; it is not a forecast of corporate profits, revenues or GDP growth.

4.2 Exposure: how much work is susceptible to GenAI?

GenAI exposure is unevenly distributed across the world economy.

The ILO estimates that occupations with some GenAI exposure account for approximately 34% of employment in high-income economies, 25% in upper-middle-income economies, 20% in lower-middle-income economies and 11% in low-income economies.³

Those employment shares occur in economies with very different economic weights. Cyborg Economics therefore weights each income group's occupational-exposure rate by its share of world GDP.

World Bank income group	Approx. share of world GDP	Employment with some GenAI exposure	Weighted contribution
High income	65%	34%	~22 percentage points
Upper-middle income	28%	25%	~7 percentage points
Lower-middle income	6%	20%	~1 percentage point
Low income	<1%	11%	<0.1 percentage point
Total	100%		~30%

The weighted contribution for each income group is calculated by multiplying its share of world GDP by its occupational-exposure rate. For example, high-income economies account for approximately 65% of world GDP and have an exposure rate of approximately 34%, giving a weighted contribution of around 22 percentage points. Adding the weighted contributions across the four income groups produces the global GDP-weighted exposure rate of approximately 30%.

Using the underlying unrounded values:

Exposed labour-income base

$$\$62.257\text{tn} \times 30.246\% = \mathbf{\$18.831\text{tn}}$$

The broad specification therefore identifies approximately \$19 trillion of annual labour income associated with occupations containing GenAI-exposed work.

GDP weighting remains a proxy. A more granular model would perform the calculation country by country and occupation by occupation. It could estimate each country's labour income using its own labour-income share, allocate that income across occupations using employment and compensation data, and then apply occupation- or task-specific GenAI exposure and productivity assumptions. The global estimate would be the sum of these country-occupation estimates. This would, for example, allow the exposure and economic weight of software developers in the United States to differ from those of administrative workers in Germany or accountants in Brazil.

Exposure gradients

The ILO classifies exposed occupations into four gradients using both the average level of GenAI exposure across the tasks within an occupation and the extent to which that exposure varies between tasks. Gradient 1 represents low overall exposure with substantial variation between tasks; Gradient 2 represents moderate exposure with high task-level variation; Gradient 3 represents significant exposure, again with substantial variation between tasks; and Gradient 4 represents the highest exposure, with high and relatively consistent exposure across tasks.³

The reference exposure measure counts employment falling within all four exposed ILO categories, Gradients 1–4. A conservative sensitivity test excludes Gradient 1, removing the lowest-exposure category.

This reduces the GDP-weighted occupational-exposure rate from approximately 30% to 23% and the corresponding exposed labour-income base from approximately \$19 trillion to \$15 trillion.

Exposure definition	GDP-weighted exposure	Exposed labour income
Broad: Gradients 1–4	~30%	~\$19tn
Conservative: Gradients 2–4	~23%	~\$15tn

Section 6 tests the effect of this narrower exposure definition on the final estimate.

4.3 Productivity: how much could GenAI raise productive capacity?

The second modelling layer estimates the productivity effect across exposed work.

Experimental evidence reports material productivity improvements, although the size and measurement of those effects vary across occupations and tasks. Evidence reviewed by the Organisation for Economic Co-operation and Development (OECD) includes gains ranging from roughly 5% to more than 25%. Customer-support workers in one study resolved around 14% more issues per hour. Consultants in another completed more tasks and worked substantially faster on tasks suited to GenAI.⁴ A field experiment involving 4,867 software developers found a 26% increase in completed tasks among developers given access to an AI coding assistant.⁵

These studies cover different occupations, tasks and performance measures. They do not constitute a single statistical distribution or establish a global average productivity effect.

Cyborg Economics uses 15% as the reference productivity uplift. This is a rounded central calibration within the broad range represented in the evidence reviewed here rather than a statistical mean. Lower and higher cases of 10% and 20% are tested separately.

For modelling purposes, the 15% uplift is applied to the labour income associated with exposed occupations as a whole. It should be interpreted as an average productivity uplift across that exposed labour base, rather than as a claim that every worker or every exposed task becomes 15% more productive.

This avoids introducing a separate assumption about the precise share of tasks affected within each occupation. The resulting simplification is tested through both the productivity sensitivity and the narrower occupational-exposure case.

The productivity and conversion assumptions capture different effects. The productivity parameter represents the productive gain available when GenAI is effectively used in suitable exposed work. The conversion layer captures the extent to which that productive potential is actually deployed, integrated and scaled across organisations.

Using the underlying exposed labour-income estimate:

Gross GenAI productivity-equivalent potential

\$18.83tn × 15% = \$2.82tn

The model therefore estimates approximately \$3 trillion per year of gross labour-productivity-equivalent GenAI potential.

This provides the potential value pool to which the conversion analysis is applied.

5. From productivity potential to value at stake

The approximately \$3 trillion estimate represents potential productive capacity. The final modelling layer estimates how much of that potential remains unrealised.

5.1 Defining the conversion gap

The realisation rate is the proportion of modelled productivity potential converted into realised value:

$$\text{Realisation rate} = \text{realised productivity-equivalent value} \div \text{potential productivity-equivalent value}$$

The conversion gap is the remaining share:

$$\text{Conversion gap} = 1 - (\text{realised productivity-equivalent value} \div \text{potential productivity-equivalent value})$$

A 35% conversion gap therefore describes a scenario in which approximately 65% of modelled productivity potential is realised and 35% remains unrealised.

There is currently no direct global, value-weighted measure of the GenAI conversion gap. The reference assumption is therefore informed by evidence on unrealised technology value and cross-checked against evidence on AI scaling and financial impact.

5.2 Evidence that conversion remains incomplete

The 2026 evidence shows substantial progress in AI adoption and scaling alongside a continuing gap between individual or targeted benefits and enterprise-wide financial impact.

McKinsey's 2026 global survey finds that nearly nine in ten respondents report regular AI use in at least one business function. Eighty per cent report improvements in individual productivity. Scaling is also advancing: 44% report that AI is scaling across their enterprise, up from 38% a year earlier. Enterprise-level financial impact remains less widespread, with 37% reporting a positive contribution to earnings before interest and tax (EBIT) and around 6% meeting McKinsey's definition of an AI high performer.⁶

Gartner uses a more demanding definition of scaling. Its 2026 survey of 1,303 respondents at organisations with annual revenues of at least \$50 million finds that 22% have successfully scaled AI across multiple business units or adopted an AI-first approach. The McKinsey and Gartner measures are not directly comparable, but they point in the same broad direction: scaling is advancing, while deep enterprise-wide scaling remains far from universal.⁷

Boston Consulting Group (BCG) illustrates the distinction between obtaining benefits from individual AI applications and converting them into enterprise-wide performance. Nearly nine in ten chief executive officers report some cost or revenue benefit from AI in targeted areas. Yet only 26% have embedded AI within a broader business transformation, while just 14% clearly define the profit-and-loss impact of every AI initiative. The evidence suggests that organisations can generate useful local returns long before those gains are systematically integrated, measured and scaled across the enterprise.⁸

Together, these findings support a distinction between AI value creation and AI value conversion at scale. Local productivity gains and targeted financial benefits are increasingly common. Converting those gains into sustained enterprise performance remains much less consistent.

A firm may also be generating substantial value from particular AI applications without having scaled AI across the whole enterprise. Measures of enterprise-wide scaling therefore cannot be converted directly into estimates of unrealised economic value.

5.3 Calibrating the reference conversion scenario

Deloitte's 2026 UK technology-value research provides a useful external reference point for the potential scale of unrealised value.

Its survey of 314 senior UK executives found that a majority believed 21% to 50% of the potential enterprise value in technology investments already made remained untapped. GenAI was the most commonly reported technology investment among the capabilities examined, cited by 54% of respondents.⁹

The Deloitte evidence covers technology investment more broadly than GenAI, is UK-specific and records executive assessments rather than a directly observed global value measure. It therefore does not provide an estimate of the global GenAI conversion gap.

Its underlying economic concept is nevertheless closely aligned with the question examined here: the proportion of potential value associated with technology investment that remains unrealised.

Cyborg Economics uses 35% as an evidence-informed reference conversion scenario, a round-number assumption close to the midpoint of Deloitte's 21–50% band.

The assumption is tested across a much wider 10–50% sensitivity range. It is not presented as an empirical estimate of the true global conversion gap.

5.4 Reference result

Applying the reference conversion assumption to the modelled productivity potential gives:

Reference case

$$\text{\$2.82tn} \times 35\% = \text{\$0.99tn}$$

The reference result therefore implies approximately \$1 trillion per year of productivity-equivalent value at stake.

At the level of precision supported by the analysis, this is best expressed as:

Around \$1 trillion per year

The complete calculation is:

REFERENCE CASE



The reference result of approximately \$1 trillion represents modelled incomplete conversion across the full process. AI readiness shapes an actionable part of that gap through data, workforce capability, workflow integration, ownership, governance, supplier management, organisational learning and scaling.

AI readiness is only one source of incomplete conversion. Some potential will remain unrealised because particular applications are technically or economically unattractive, or because wider constraints slow adoption. The reference result should therefore be interpreted as total modelled incomplete conversion, only part of which is addressable through improvements in AI readiness.

6. Sensitivity and interpretation

The sensitivity analysis follows the same sequence as the model:

Exposure → Productivity → Conversion

Each assumption is first varied independently. A final set of illustrative scenarios then combines the assumptions.

6.1 Exposure sensitivity

The reference case includes all four ILO exposure gradients. The conservative case excludes Gradient 1, reducing the GDP-weighted occupational-exposure rate from approximately 30% to 23%.

Holding productivity at 15% and the conversion gap at 35% gives:

Exposure case	Exposure rate	Annual value at stake
Conservative: Gradients 2–4	~23%	~\$0.8tn
Reference: Gradients 1–4	~30%	\$0.99tn

Removing the lowest-exposure category therefore reduces the result to approximately \$0.8 trillion per year, while preserving the same broad order of magnitude.

6.2 Productivity sensitivity

Holding the broad exposure definition and the 35% conversion gap constant gives:

Effective productivity uplift	Annual value at stake
10%	~\$0.7tn
15% reference	\$0.99tn
20%	~\$1.3tn

The result remains economically substantial under the lower productivity assumption.

6.3 Conversion sensitivity

Holding exposure and productivity at their reference assumptions gives:

Conversion gap	Annual value at stake
10%	~\$0.3tn
20%	~\$0.6tn
30%	~\$0.8tn
35% reference	\$0.99tn
40%	~\$1.1tn
50%	~\$1.4tn

Even a 10% conversion gap, equivalent to around 90% of modelled potential already being realised, leaves approximately \$300 billion per year of productivity-equivalent value at stake.

6.4 Combined stress tests

The individual sensitivities show the effect of changing one assumption at a time. The combined scenarios show how the result moves when the main uncertain assumptions change together.

The assumptions may not move independently. For example, activities offering the largest productivity gains may also attract faster investment and adoption, which would tend to reduce the conversion gap. Broader occupational exposure could also bring increasingly marginal applications with lower average productivity effects. There is currently insufficient global evidence to estimate these relationships reliably. The combined cases are therefore mechanical stress tests rather than probabilistic forecasts; they should not be interpreted as confidence intervals or as equally likely outcomes.

Value-at-stake scenario	Exposure	Productivity uplift	Conversion gap	Annual value at stake
Low	Conservative (~23%)	10%	10%	~\$0.15tn
Reference	Broad (~30%)	15%	35%	\$0.99tn
High	Broad (~30%)	20%	50%	~\$1.9tn

The low value-at-stake scenario combines the narrower occupational-exposure definition with a 10% productivity uplift and a 10% conversion gap. It produces approximately \$150 billion per year.

The reference scenario produces approximately \$1 trillion.

The high value-at-stake scenario combines the broad exposure definition with a 20% productivity uplift and a 50% conversion gap, producing approximately \$1.9 trillion.

These are illustrative scenarios rather than probability estimates or statistical bounds. They show how the result responds when several uncertain assumptions move simultaneously.

6.5 Interpreting the headline

The sensitivity analysis produces a wide range because the economic effects of GenAI and the rate at which organisations convert them into value remain uncertain.

The reference calculation is approximately \$1 trillion per year. The narrower exposure specification produces approximately \$0.8 trillion. A 30% conversion gap also produces approximately \$0.8 trillion. The combined illustrative scenarios extend from approximately \$0.15 trillion to \$1.9 trillion.

The central conclusion is therefore about economic scale, rather than a precisely observed global quantity:

Incomplete GenAI value conversion could leave around \$1 trillion per year of productivity-equivalent value at stake.

The structure of the model makes the assumptions behind that conclusion explicit. The economic base comes from global GDP and labour income. Exposure is derived from the ILO occupational-exposure framework. Productivity is calibrated against empirical evidence. Conversion is informed by evidence on unrealised technology value and tested across a broad range.

As AI capability becomes more widely available, access to the technology becomes less differentiating. Economic performance increasingly depends on whether organisations can turn that capability into sustained productive value.

AI readiness is the organisational capacity to convert AI potential into sustained, measurable performance.

7. Wider economic effects beyond the headline estimate

The reference estimate of around \$1 trillion per year represents the modelled annual value at stake from incomplete conversion of generative AI (GenAI) productivity potential. It isolates the labour-productivity channel. AI can also create economic value through innovation, new products and services, wider knowledge spillovers and improved public-sector outcomes. Realising AI value also requires resources and adjustment, and can generate adverse wider economic and social effects that are not captured in the core model.

These effects operate through different mechanisms, emerge over different timescales and require different approaches to measurement. They are therefore considered separately from the core estimate. The distribution of gains and losses is considered in Section 8.

7.1 Additional benefits beyond the core productivity estimate

Innovation, new products and process improvement

The core model estimates value arising from higher labour productivity within exposed work. AI may also generate value by changing what organisations produce and how production itself is organised.

GenAI can reduce the cost of activities involved in innovation, including idea generation, research, design, coding, prototyping and testing. Organisations may be able to explore more alternatives, shorten development cycles and pursue opportunities that would previously have been too costly. The OECD's review of the evidence identifies potential effects on creativity, research and development, innovation and entrepreneurship alongside the more directly measured productivity effects discussed earlier.¹

The resulting innovation can take several forms. AI may support new products, services and business models. Examples could include new AI-enabled digital services, more personalised products, new forms of professional service delivery, or business models made viable by much lower costs of producing customised analysis, software or content. It may also enable process innovation: different ways of producing existing goods and services using fewer resources or achieving better outcomes. Examples could include redesigned supply chains, improved procurement, lower inventory requirements, reduced material use or the identification of more efficient technologies and operating methods. Some of these innovations may themselves make little subsequent use of AI.

Care is needed to avoid double counting. Where process innovation ultimately raises labour productivity, part of the effect may already be captured in the core model. Gains arising through lower use of materials, capital, inventories or other non-labour inputs sit outside that channel.

Innovation can therefore increase the overall economic value associated with AI beyond the headline estimate, although these effects are harder to quantify and may take longer to emerge.

Knowledge spillovers and positive externalities

Some benefits from AI adoption may extend beyond the organisation making the original investment.

Knowledge about effective applications can diffuse between firms, including through worker mobility and imitation. Shared technical standards, publicly available evaluation methods and other forms of accumulated knowledge can also reduce the cost and uncertainty of later adoption.

Where these benefits are not fully captured by the organisations or individuals responsible for creating them, they constitute positive externalities. The private return from an investment may consequently understate its wider economic value.

These effects can matter for the pace and scale of AI diffusion. Experience accumulated by early adopters may reduce uncertainty for later firms. Shared standards and evaluation methods can lower implementation

costs across an industry. Skills and knowledge developed in one organisation may subsequently contribute to productivity elsewhere.

Where these spillovers are significant, there may be a rationale for collective provision or public support for activities such as shared technical standards, open evaluation resources, workforce skills and research whose benefits extend across many firms.

Public-sector and institutional benefits

AI may also create value through public services and institutional capability that is only partially represented by the labour-productivity model.

For both private and public services, the value created can depend on the quality and outcomes of the service as well as the labour required to provide it. The issue is particularly important in public services, where market prices often provide little or no indication of the value of improved outcomes.

Public-sector employment forms part of the labour base used in the core calculation, so some gains from more productive public-sector work are already captured. AI applications in government may reduce administrative burdens, accelerate processing, improve analysis and forecasting, strengthen targeting or improve access to services. AI may also strengthen analytical, regulatory and administrative capability, improving the quality or timeliness of decisions whose effects extend beyond the immediate task. OECD research documents growing use of AI across government functions and highlights the importance of supporting conditions including data, skills, infrastructure, procurement and governance.²

A reduction in the time required to process a claim, planning application or regulatory case may therefore create value through more than the associated labour saving. Faster decisions, fewer errors, better allocation of public resources and improved outcomes for citizens may generate additional social benefits.

These effects illustrate a broader limitation of the core model: changes in labour productivity provide a useful monetary proxy for one channel of AI value, but they do not capture every improvement in the quality or quantity of economic and public-service outcomes.

7.2 Costs of implementing and operating AI

Realising AI productivity gains requires real resources.

At firm and organisational level, adoption may require expenditure on data preparation, systems integration, model configuration, workflow redesign, staff training, change management, governance and assurance. Deployment can also disrupt existing processes while new ways of working are tested and employees adapt. Some organisations may need to operate existing and AI-enabled systems in parallel during transition.

AI also creates ongoing operating costs. These can include model access, cloud services, compute, software licences, monitoring, maintenance and additional assurance or security requirements.

The scale of these costs will differ significantly between applications. A self-contained tool added to an existing workflow may require relatively little organisational change. Embedding AI within core operational systems can require substantial technical and organisational investment.

Beyond investment within individual firms and organisations, widespread AI adoption may require additional capacity in electricity generation and networks, data centres and related infrastructure. The International Energy Agency reported that global data-centre electricity consumption increased by 17% in 2025, with demand from AI-focused data centres growing faster still.³

The timing of these costs also matters. Some implementation and infrastructure costs are incurred upfront and generate benefits over several years, while operating costs recur. Comparison with an annual productivity benefit therefore requires costs and benefits to be placed on a consistent time basis.

The core estimate does not deduct the full implementation, operating or infrastructure resources required to realise the modelled productivity potential. These costs therefore matter when moving from productivity-equivalent value towards an assessment of net economic benefit.

7.3 Labour-market adjustment

AI adoption can also change the demand for particular tasks, skills and occupations.

The occupational exposure measure used in the core model identifies the extent to which tasks within an occupation could be affected by GenAI. The employment consequences depend on what subsequently happens to those tasks. Some may be augmented, some reorganised, some partly automated and some removed altogether.

The International Labour Organization's 2025 analysis finds that around one quarter of global employment lies in occupations with some GenAI exposure. Its central assessment is that transformation is currently more likely than wholesale replacement because most occupations continue to contain tasks requiring human input.⁴

Transformation can nevertheless generate significant adjustment costs. Workers may need to acquire new skills. Roles may be redesigned or disappear. Job search and occupational mobility take time and resources, while temporary unemployment can reduce output during the transition. Retraining, relocation and movement between sectors or occupations can also impose real costs.

These labour-market effects are economically important even where AI ultimately increases aggregate productive capacity. The distributional consequences for workers, households and particular groups are considered in Section 8.

7.4 Negative externalities and wider risks

Some costs generated by AI deployment may fall partly on people or organisations that are not responsible for the underlying decision.

Environmental effects provide one example. Electricity used by AI infrastructure can contribute to greenhouse-gas emissions where generation remains carbon-intensive and climate costs are not fully reflected in energy prices. Data-centre cooling can increase water demand, while concentrated electricity demand can add pressure to local networks. These become external costs where associated scarcity, congestion or environmental damage is not fully reflected in the prices paid by organisations deploying AI.

AI-related failures can create further external costs. Unreliable outputs, security incidents, privacy breaches, discrimination and other harms may impose losses on users, customers or third parties. OECD work on AI risks records realised harms across areas including privacy, discrimination, security and safety.⁵

The distinction between private cost and external cost matters. Compute, electricity, staffing and implementation expenditure paid directly by the deploying organisation are resource costs already faced by that organisation. Harms falling on third parties are different: the private return from an AI application may exceed its net value to society where those wider costs are not reflected in the organisation's decisions.

Assurance and regulation can help address this divergence. Weak controls may allow significant external costs to arise, while unnecessarily burdensome controls can suppress valuable adoption. The appropriate response depends on the scale, likelihood and distribution of the risks involved.

7.5 Market structure and transmission of gains

Competition and market structure differ from the preceding effects: they shape the scale, diffusion and transmission of AI's economic gains.

Falling model prices and widely available AI services can reduce the cost of capabilities that previously required large specialist teams. Smaller firms may gain access to software development, analysis, design and customer-service capabilities that were previously difficult to afford. Lower experimentation costs may also reduce barriers to entry in some markets.

At the same time, important parts of the AI value chain remain concentrated. The OECD identifies structural competition risks associated with advanced chips, cloud infrastructure, compute, data and specialised skills,

alongside the potential for vertical integration, switching costs and platform advantages to strengthen incumbent positions.⁶

Where competition is effective, lower production costs can support greater output, lower prices, improved quality and faster innovation. Productivity gains can therefore spread beyond the organisation initially adopting AI.

Where competition is weak, market power and dependency on scarce suppliers may raise adoption costs or slow diffusion. Dominant firms or suppliers may retain a larger share of the benefits, while high switching costs or restricted access to key inputs can limit the amount of AI-enabled activity that takes place.

Market structure therefore affects both the scale of the economic gains and the way those gains are transmitted. Section 8 considers the related distributional question of where the resulting benefits and costs ultimately accrue.

7.6 Why the wider effects are not included in the headline estimate

The wider effects described in this section are difficult to incorporate into a single global estimate. In many cases the first challenge is quantification: there is limited evidence on the scale of additional innovation, implementation costs, labour-market adjustment or external effects attributable to AI. Where outcomes can be quantified, monetary valuation may also be difficult, particularly for changes in service quality, environmental effects or wider social impacts.

Incorporating all of these effects into a single global monetary estimate would therefore require substantially more evidence and modelling and could create false precision. The headline estimate of around \$1 trillion per year of productivity-equivalent value at stake should instead be understood as an estimate of one important and relatively measurable channel within the wider economic impact of AI.

Section 8 turns to the final part of that assessment: where those gains and costs ultimately fall.

8. Distribution of gains and losses

8.1 From value creation to value distribution

The estimates in Sections 4–6 are global and aggregate. They indicate the scale of productivity-equivalent value that could be realised through more complete conversion of generative artificial intelligence (GenAI) potential, but they do not show where those gains would ultimately accrue.

Section 7 broadened the analysis beyond the core productivity channel to include innovation, implementation and operating costs, labour-market adjustment, externalities and market structure.

The distributional question concerns how the gains and costs arising through these different channels are ultimately spread across firms, consumers, suppliers, workers, government and different parts of the economy.

If AI enables a firm to produce a given output with fewer resources, the resulting economic benefit can be shared in different ways. The firm may retain more of it through higher profits, workers may receive some through higher pay, or customers may benefit through lower prices. The released productive capacity may also allow the organisation to produce more or improve the quality of its products or services.

Section 8 does not attempt to add these outcomes to the headline estimate of around \$1 trillion per year. Its purpose is to examine where the gains and costs associated with AI ultimately fall.

Economic actors also overlap. A household may receive income from employment, consume AI-enabled goods and services, own shares through a pension fund, pay taxes and use public services. Distributional analysis therefore traces the different channels through which economic effects reach people and organisations rather than dividing the economy into mutually exclusive groups.

A further question concerns the amount of value that different organisations and places are able to realise in the first place. Exposure to AI varies, as does the readiness required to convert technical capability into economic performance. Distribution therefore depends both on how realised value is divided and on where value conversion occurs most successfully.

8.2 Firms and consumers: product-market incidence

Productivity gains realised within firms can be divided between producers and their customers in different ways.

AI can affect firms through several related channels. Where it reduces the resources required to produce a given level of output, the immediate effect is a reduction in cost. The firm may retain some of that saving through higher margins, while competitive pressure may pass some to customers through lower prices.

AI may also increase the amount or quality of output that can be produced from existing resources. In those cases, customers may benefit through greater availability, shorter waiting times, improved reliability or more personalised services. The balance between higher margins, lower prices and changes in output or quality will depend on the nature of the productivity improvement and the market in which the firm operates.

As discussed in Section 7, competition and market power influence the degree of pass-through. Stronger competitive pressure is more likely to transmit productivity gains through prices, output, quality or service, while firms with greater market power may be able to retain a larger share through margins.

The same incidence question can arise in business-to-business markets. A supplier that becomes more productive may retain the gain or pass some of it to business customers through lower input prices or improved input quality. Those customers may then experience lower costs or better inputs into their own production.

The distributional effect therefore depends partly on market structure. Similar productivity improvements can produce different outcomes for firms and customers depending on competitive conditions and the way organisations use the capacity released by AI.

8.3 Suppliers and platforms: supply-chain capture

AI adoption also depends on inputs supplied by other organisations, including models, cloud infrastructure, compute, software, data services and specialist technical capability.

Expenditure on these inputs is principally a resource cost of AI deployment and is considered in Section 7. The distributional issue is narrower: whether scarcity, bargaining power or market concentration allows particular suppliers to capture a disproportionate share of the economic surplus created through AI adoption.

This can occur where important inputs are difficult to substitute or where switching between providers is costly. Control of scarce computing capacity, specialised technology, proprietary platforms or other critical inputs can strengthen the negotiating position of suppliers. The OECD identifies concentration of key inputs, vertical integration and switching costs among the structural competition risks across parts of the AI value chain.¹

Where upstream markets are competitive, prices should more closely reflect the resources required to supply the relevant inputs. Where suppliers possess substantial market power, they may be able to earn returns above those required in a more competitive market. A larger share of the economic surplus generated by downstream adoption can then accrue to those suppliers rather than to adopting firms or their customers.

Commercial and technical dependencies can therefore affect the distribution of AI-generated value even where the underlying application is productive. The importance of this channel will vary across use cases and over time as model markets, compute capacity, cloud services and other parts of the AI supply chain develop.

8.4 Workers: labour-market incidence

AI can also change the distribution of economic gains through wages, employment, hours, skills and career progression.

Where AI complements workers, higher productivity can increase the value of their work. Employees may be able to produce more, perform more complex activities or shift towards tasks requiring judgement, relationships or specialist knowledge. Some of the resulting gains may be reflected in higher wages, stronger employment prospects or increased demand for particular skills.

Other workers may experience weaker demand for tasks that AI can perform more cheaply or effectively. The effect will vary across occupations and skill groups because exposure is uneven and because workers differ in their ability to move into complementary tasks or growing areas of activity.

The effects can also arise within occupations. A profession may remain economically important while the relative value of different tasks changes. Workers whose skills complement AI may become more productive, while those concentrated in increasingly automated activities may face weaker bargaining positions or reduced demand.

Early evidence suggests that these effects may differ by career stage. Using US payroll data covering millions of workers, Brynjolfsson, Chandar and Chen find that employment among workers aged 22–25 in highly AI-exposed occupations had fallen substantially relative to less-exposed peers by mid-2026. The divergence arose mainly through weaker hiring and was concentrated in occupations where AI use was more associated with automation than augmentation. The authors describe these findings as early descriptive evidence rather than causal estimates.²

These patterns raise a wider distributional question around entry into affected occupations. Reduced recruitment at junior levels can alter opportunities for younger workers to enter particular careers, although the longer-term consequences for skills development and career progression remain uncertain.

Workers may also respond to changing wages and opportunities by altering their hours, labour-force participation, occupational choices or investment in skills.

Section 7 considered the economic resources involved in labour-market adjustment, including retraining, job search and temporary unemployment. The distributional issue is how higher productivity and the resulting changes in wages, employment, hours, labour-force participation and occupational choices affect different workers and households.

8.5 Government and public services

Government is affected through both the public finances and the provision of public services.

Changes in employment, wages, profits and consumption can affect tax receipts. Labour-market disruption may also increase expenditure on unemployment-related benefits, retraining and employment support.

Separately, AI adoption within government can raise public-sector productivity, reducing the resources required to provide a given level of service or allowing more or better services to be delivered from an existing budget. OECD analysis of AI use across government identifies automated, streamlined and tailored processes and services as the most common objective among the use cases it examined.³

The distribution of public-service gains is particularly important because service users do not necessarily benefit equally from improvements in institutional productivity.

AI may shorten processing times, improve access to information, support more consistent decisions or enable services to respond more effectively to individual circumstances. The resulting benefits will depend on which services adopt AI successfully and on which groups use those services.

Changes in service design can also affect effective access. New delivery models may improve accessibility for some users while creating difficulties for others where suitable alternatives, human support or appropriate safeguards are unavailable. The relevant distributional question is therefore whether changes in public-service quality, availability and access benefit different groups and places equally.

Public institutions consequently matter both as adopters of AI and as channels through which the wider effects of technological change reach households.

8.6 Why AI value realisation will differ across firms, sectors and places

AI value will not be realised evenly across the economy. Firms, sectors and places differ both in their exposure to AI and in their capacity to convert that exposure into sustained economic performance.

The distribution of AI value therefore depends on two distinct factors: the scale of the technical opportunity and the capacity to convert it into realised value.

8.6.1 Firms and sectors

The economic potential of AI varies between sectors because their activities contain different combinations of tasks and use cases. Some contain large amounts of information processing, analysis, software development, customer interaction or other work to which current GenAI capabilities are particularly applicable. Others depend more heavily on physical activities or technologies outside the scope of the core model. Some sectors therefore begin with substantially greater technical potential than others, but exposure alone does not determine how much value is ultimately realised.

Firms also differ in the readiness capabilities described in Section 2: their ability to integrate selected AI opportunities into workflows, manage data and technology, adapt their workforce, govern deployment, manage suppliers and learn from performance evidence. Firms facing similar technological opportunities can therefore realise very different economic outcomes.

Evidence discussed earlier in the paper is consistent with substantial concentration in realised AI returns. PwC's 2026 AI Performance Study found that the top-performing 20% of organisations in its sample captured 74% of the measured AI-driven returns. Leading firms also differed in areas including workflow redesign, deployment, data, technology and governance. The study does not establish the causal effect of individual readiness capabilities, but it provides evidence that measured AI returns were highly concentrated across organisations in its sample.⁴

Firm size can affect conversion capacity. Across OECD countries, larger firms currently adopt AI at substantially higher rates than smaller firms, while SMEs report barriers including skills, finance, infrastructure and access to AI-enabling inputs.⁵

The falling cost and increasing accessibility of AI tools may nevertheless reduce some of these disadvantages faced by smaller firms. OECD survey evidence from more than 5,000 SMEs across seven countries found GenAI in use in 31% of SMEs surveyed and noted that lower costs and data requirements have widened access to capabilities that previously required greater technical and financial resources.⁶

Firm size therefore does not determine the outcome by itself. Large firms can benefit from greater financial, technical and organisational resources, while increasingly accessible general-purpose tools can allow smaller firms to obtain capabilities without building extensive AI infrastructure internally.

Sector structure can reinforce these differences. Regulation, workforce composition, data availability, the economics of particular use cases and the concentration of firms within a market can all affect the speed and extent of conversion.

A sector can therefore contain substantial AI-exposed work while realising only a limited share of its technical potential. Another may have lower exposure but convert a greater proportion because its organisational and market conditions are more favourable.

The same point applies within sectors. Firms facing similar technological opportunities can achieve very different results where their readiness capabilities differ.

Sector-level outcomes will therefore depend on both the scale of technical opportunity and the ability of firms within the sector to convert that opportunity into realised performance.

8.6.2 Countries and regions

The same distinction applies geographically.

Three dimensions are relevant: economic exposure, conversion capacity and value-chain income capture.

Economic exposure reflects the extent to which the activities undertaken within an economy are susceptible to productive use of AI. Section 4 showed that current GenAI occupational exposure is substantially higher in high-income economies because their employment structures contain larger shares of exposed occupations.

Conversion capacity concerns the ability to turn that exposure into economic performance. Digital and energy infrastructure, workforce skills, capital, management capability, data systems, public institutions and regulatory arrangements can influence the speed and effectiveness with which AI is adopted and integrated into economic activity.

Value-chain income capture concerns where income generated through AI adoption ultimately accrues. Economies hosting major model developers, semiconductor producers, cloud providers or other technology suppliers may receive income from AI adoption taking place elsewhere. Economies relying heavily on imported AI inputs may capture more of the gains through downstream productivity and consumption than through ownership of the upstream technologies themselves. OECD analysis finds that important elements of AI investment, compute and upstream technology capacity are already highly concentrated across firms and countries.¹

These dimensions can vary independently. A country may have high occupational exposure but relatively weak conversion capacity. Another may deploy AI effectively while relying heavily on technology supplied from abroad. Countries with strong domestic AI industries may receive gains from both local adoption and external demand for technology and services.

Differences can also emerge within countries. AI-sector activity already shows substantial geographic concentration on some measures. In the UK, London, the South East and East of England accounted for approximately 75% of registered AI-company office locations in 2024, although AI-firm numbers were increasing across all UK regions.⁷

Exposure also varies sharply between regions. OECD analysis finds that the most exposed region within a country is, on average, around 1.6 times as exposed to GenAI as the least exposed. Across OECD regions, prospective high exposure ranged from 16% of workers in Guerrero, Mexico, to 77% in Greater London. In Colombia, Bogotá was more than three times as exposed as La Guajira on the same prospective measure.⁸

These differences do not determine where future productivity gains will occur. They show, however, that occupational structure and AI-related economic activity are already distributed unevenly between places. Skills, firms, investment, research capability, infrastructure and institutional capacity can further affect the ability of regions to convert technological opportunity into economic activity.

Public-sector capability can reinforce these patterns. National and local institutions influence skills provision, infrastructure, technology diffusion and the effectiveness of AI adoption in public services. Institutional readiness can therefore affect the capacity of places to convert wider technological change into local economic and public value.

AI readiness may consequently become a sorting mechanism across firms, sectors and places. Organisations and economies with stronger conversion capabilities may capture a disproportionate share of the gains available from increasingly capable and widely accessible technology.

Successful conversion may also strengthen future readiness by generating experience, organisational learning and resources for further investment. This creates a plausible mechanism through which initial readiness differences could widen over time, although the strength of that effect remains uncertain. PwC's finding that leading organisations are already more likely to redesign workflows and deploy AI more extensively is consistent with such a mechanism, but does not establish that divergence will persist or compound.⁴

These differences mean that a large global productivity opportunity can coexist with highly uneven outcomes across firms, sectors and places.

AI can therefore increase aggregate productive capacity while producing markedly different outcomes across the economy. AI readiness shapes how much value is realised and where it is generated. Who ultimately captures the gains depends on the wider economic and institutional conditions through which that value is distributed.

9. Measuring AI value conversion

The preceding sections have examined the scale of the economic opportunity from AI, the organisational conditions affecting its conversion, the wider effects that may result and the distribution of the resulting gains and costs. These questions can also provide the basis for practical measurement.

AI value conversion can be examined through four related analytical questions:



These questions perform different roles over the life of an AI investment. Economic Opportunity is principally an appraisal question. AI Readiness concerns organisational capability before and during implementation. Value Realisation compares actual performance with the economic case. Value Capture considers the extent to which realised value translates into returns for the adopting organisation and other parties.

9.1 A four-part measurement framework

Each question requires a different combination of measures and evidence.

Analytical question	Measurement focus	Illustrative measures and data
Economic Opportunity	Establish the expected economic case for a potential AI application and compare it with available alternatives.	Current problem and value pool • task or transaction volumes • staff time and labour cost • unit cost • errors and rework • backlogs, service problems or revenue leakage AI-task fit and feasibility • suitability of the task for AI • required performance and accuracy • specific technical feasibility • pilot or testing evidence where available Expected benefits • productivity and capacity • cost reduction • quality and service improvement • revenue or demand effects

Analytical question	Measurement focus	Illustrative measures and data
		• risk reduction where measurable Expected costs • implementation and integration • technology and supplier costs • staff and organisational change • assurance • ongoing operation and maintenance Alternatives and uncertainty • status quo and non-AI alternatives • scalability and reuse • key assumptions and sensitivities Expected net value • NPV, BCR, payback or another proportionate decision measure <i>Data: finance and operational systems, process data, technical testing, pilot evidence, supplier quotations and relevant external evidence</i>
AI Readiness	Assess whether the organisation possesses the capabilities required to implement worthwhile AI applications and sustain their performance.	Strategic and operating • ownership and decision rights • prioritisation processes • workflow and operating-model adaptation • routes from testing to scale People and adoption • workforce and managerial capability • role clarity • user support • adoption and feedback mechanisms Data and technology • usable data • integration capability • deployment foundations • monitoring and operational support Risk, assurance and governance • accountability • proportionate controls • validation • escalation and continuing assurance Commercial and ecosystem • supplier selection and management • contracting • interoperability • dependency and switching management Measurement and learning • defined responsibility for measurement and learning • access to analytical capability • established review processes • use and transfer of evidence across projects <i>Data: governance, HR, architecture, operational, procurement, risk and portfolio records</i>
Value Realisation	Test actual performance against the economic case and estimate the net economic value attributable to the intervention.	Measurement plan • appraisal measures to be tracked • baseline • comparator or counterfactual • success criteria and targets • data requirements and collection frequency • attribution approach Implementation and use • deployment • adoption • usage intensity • workflow coverage Actual benefits • task or cycle time • productivity and throughput • capacity

Analytical question	Measurement focus	Illustrative measures and data
		• quality • errors and rework • service performance • attributable revenue or demand effects Actual costs • implementation • technology • staff and change costs • assurance • ongoing operation and maintenance Realised value • attributable net benefit • variance from the original appraisal • persistence, decay or growth of benefits over time <i>Data: workflow systems, management accounts, telemetry, CRM, quality data and evaluation datasets</i>
Value Capture	Assess how much of the value created is retained by the adopting organisation and, where the evidence permits, how the remainder is distributed.	Expected value capture and appropriability conditions • competitive intensity • bargaining power • ease of imitation • switching costs • complementary assets • supplier concentration Organisation • share of realised value retained • retained cash savings • contribution and profit effects Customers or service users • evidence of pass-through through prices, service quality, availability or waiting times Workers • wages • employment • hours and task changes Suppliers and platforms • evidence of supplier rents or bargaining effects • dependency and contractual position Government and wider incidence • fiscal effects • other material distributional effects <i>Data: management accounts, pricing and payroll data, contracts, procurement records, customer data and relevant market evidence</i>

The relationship between Economic Opportunity and Value Realisation is a forecast-and-outturn relationship. Opportunity appraisal establishes the expected benefits, costs and net value of an intervention. Value Realisation subsequently measures actual performance against those expectations and estimates the net value attributable to the intervention.

Value Capture addresses a different economic question. Realised value may be retained by the adopting organisation, passed through to customers, shared with workers, captured by suppliers or distributed elsewhere through the economy. Expected capture depends partly on appropriability conditions, including ease of imitation and control of complementary assets, and partly on wider market conditions such as competition and bargaining power. Where expected capture materially affects the private investment case, those assumptions should feed back into the Economic Opportunity appraisal.¹

The four questions also support learning over time. Implementation can change the estimated economics of the opportunity, reveal previously hidden readiness constraints, provide evidence on realised performance and alter expectations about how future gains will be captured.

9.2 Levels of analysis

AI value conversion can be assessed at different organisational levels:

- Use case or workflow — a particular AI application and the process through which value is expected to arise.
- Business unit or function — a group of applications and organisational capabilities within a defined operating area.
- Enterprise — AI capability, investment and economic effects across the organisation as a whole.

The appropriate evidence changes with the level of analysis. A use-case assessment may examine the economics, data requirements, operating constraints and performance of a specific application. A functional assessment can examine common capabilities and the performance of a portfolio of applications. An enterprise assessment can examine organisation-wide capability, investment allocation and aggregate value.

Evidence from the different levels can inform the wider diagnosis, but the measures should retain the meaning associated with the level at which they were collected. An organisation may have strong enterprise-wide AI infrastructure and governance, while a particular customer-service use case still lacks suitable data, clear workflow ownership or effective integration with the case-management system.

Measurement effort should also reflect the economic significance, complexity and risk of the decision. A bounded internal productivity application may require a relatively light assessment. A major deployment affecting core operations, substantial expenditure or high-consequence decisions warrants correspondingly stronger evidence.

The same principle applies across the four analytical questions. A small use case may require a straightforward appraisal, a proportionate readiness review, a small number of performance measures and a basic assessment of retained value. Larger strategic investments can justify deeper economic modelling, capability assessment, evaluation and analysis of value capture.

9.3 Operationalising AI readiness

Each of the four analytical questions would require its own detailed measurement approach. The remainder of this section uses AI Readiness as a worked example because readiness is the organisational-capability component developed most fully in this paper.

Section 2 identified six connected readiness capabilities: strategic and operating; people and adoption; data and technology; risk, assurance and governance; commercial and ecosystem; and measurement and learning. These provide a practical structure for gathering evidence about organisational conversion capability.

A readiness assessment can bring that evidence together into a capability profile. An illustrative example is shown below.

Capability	Illustrative evidence	Diagnostic interpretation
Strategic and operating	• clear accountable ownership • established prioritisation process • defined route from testing to scale	Operating arrangements support progression from selected opportunities into implementation
People and adoption	• strong specialist capability • uneven managerial capability • inconsistent adaptation of roles and workflows	Adoption may become a constraint as deployment expands
Data and technology	• suitable AI tools available • fragmented access to required data • difficult integration with core systems	Material constraint on deployment and scaling
Risk, assurance and governance	• clear accountability • established assurance route • proportionate controls	Current arrangements support deployment at the present level of risk
Commercial and ecosystem	• suitable suppliers and contracts • established supplier management • significant dependence on one platform	Strong near-term capability accompanied by increasing supplier dependency
Measurement and learning	• defined responsibility for reviewing performance	Material constraint on organisational learning

Capability	Illustrative evidence	Diagnostic interpretation
	• analytical capability available • inconsistent sharing of evidence between projects • weak use of portfolio evidence in subsequent decisions	and improvement

The evidence required would vary with the purpose and level of the assessment. At use-case level, data and technology readiness might concern access to the information and integration required by a particular application. At enterprise level, it might concern the organisation's ability to make governed, usable data and deployment infrastructure available repeatedly across multiple applications.

Measurement and learning readiness is similarly a capability question. It concerns whether the organisation has responsibilities, analytical capacity and routines for generating and using evidence about AI performance. The design of the evaluation for an individual investment, and the resulting estimate of actual benefits and costs, sit within Value Realisation.

A practical assessment needs clear indicators, examples of the evidence that should be considered and simple guidance for interpreting the findings. More formal standardisation could be developed later where benchmarking across organisations or repeated scoring would add value.

9.4 Interpreting readiness as a connected system

A readiness profile should be interpreted both in terms of individual capabilities and the way those capabilities operate together. Four patterns are particularly useful: reinforcing relationships, tensions, constraints and uneven capability development.

Reinforcing relationships occur where capability in one area increases the effectiveness or value of capability elsewhere. Better governed and more accessible data, for example, can improve technical deployment while also strengthening the evidence available for subsequent evaluation. Stronger workforce capability can create greater value when roles and workflows are redesigned around the technology. These relationships correspond to what the economics literature describes as complementarities.

Tensions arise where strengthening one part of the readiness system creates costs or competing demands elsewhere. Stronger assurance can reduce risk while increasing the time and resource required for deployment. Greater enterprise standardisation can improve efficiency and reuse while reducing flexibility for particular workflows.

Constraints arise where weakness in a capability materially restricts conversion despite strength elsewhere. Fragmented data integration, for example, can prevent otherwise strong strategic, workforce and technical capability from producing value at scale. Several constraints can operate at the same time.

Uneven capability development occurs where connected capabilities advance at materially different rates. An organisation may build strong technology and procurement capability faster than workflow redesign, managerial capability and organisational learning. It may then become able to deploy AI faster than it can absorb, govern or evaluate it effectively.

In the illustrative profile above, fragmented data is a material constraint on deployment, while weak organisational learning limits the organisation's ability to carry evidence from one project into subsequent decisions. Technical and supplier capability have also developed further than managerial adoption and organisational learning, creating unevenness in the readiness profile.

The value of further investment in any one readiness capability therefore depends partly on the surrounding system. Investment should be directed towards the constraints and reinforcing relationships with the greatest economic significance, while taking account of tensions and areas of uneven capability development. This is consistent with evidence from earlier information technologies showing that returns to technology depend partly on complementary investment in skills, work organisation and wider organisational capability.²

Deployment also provides new information about readiness. Data limitations, workflow problems, assurance requirements, adoption difficulties and supplier dependencies can become clearer once an application is

operating in real work. Reassessment allows the organisation to update its readiness profile and redirect capability investment accordingly.

Taken together, the four analytical questions provide a practical structure for examining AI value conversion: identify worthwhile economic opportunities; establish whether the organisation has the capabilities required to pursue them; measure the net economic value subsequently realised; and assess how much of that value is retained and, where the evidence permits, where the remainder accrues.

This creates a basis for identifying where AI value is available, what may prevent its conversion into sustained performance and where further intervention is most likely to improve economic outcomes.

10. Limitations and alternative approaches

The headline estimate is designed to show the order of magnitude of productivity-equivalent value at stake under the reference assumptions. Its main advantage is transparency: the economic base, GenAI occupational exposure, productivity uplift and conversion gap are explicit, and the effect of changing the principal assumptions can be tested separately.

The calculation nevertheless simplifies a more complex economic process. It combines global economic aggregates, occupational exposure estimates, productivity evidence from particular tasks and occupations, and an assumed conversion gap. The resulting estimate should therefore be interpreted as an annual scenario of productivity-equivalent value at stake from incomplete GenAI conversion.

10.1 Main methodological limitations

Limitation	Implication for the analysis
The monetary estimate covers GenAI rather than all forms of AI	The model uses the ILO's GenAI occupational-exposure framework and empirical evidence on GenAI-assisted work. It therefore does not estimate value arising from conventional machine learning, optimisation systems, computer vision, robotics, embodied AI or future AI capabilities outside that framework. The AI readiness framework used elsewhere in the paper has a broader scope than the monetary estimate.
Labour income is used as the valuation base	The productivity uplift is applied to labour income rather than total output. This provides a monetary proxy for the additional productive capacity associated with affected labour input. It is not equivalent to additional GDP, revenue, profit or cash saving. The value ultimately realised depends on how released capacity is used and on associated changes in output, quality and costs.
The global economic base and weighting approach are aggregated proxies	The calculation applies a global labour-income share and uses income groups' shares of world GDP to weight GenAI occupational exposure. This captures the greater economic weight of exposed work in higher-income economies, but abstracts from differences in labour shares, wages and occupational composition between countries. Within occupations, it also abstracts from differences in compensation and task mix. A more granular model would estimate labour income and exposure by country and occupation and, where possible, by task.
Occupational exposure does not measure adoption or effective use	The ILO framework identifies occupations containing tasks with potential GenAI exposure. It does not show whether AI has been adopted, whether a particular application performs sufficiently well, whether workers use it effectively or whether an exposed task represents an economically worthwhile use case. Exposure also varies substantially between tasks within the same occupation. The narrower Gradients 2–4 sensitivity reduces reliance on the least-exposed category but does not remove this variation.
The productivity parameter is transferred from specific studies to a much wider labour base	The 15% reference uplift is a rounded central calibration within empirical evidence spanning roughly 5% to more than 25% across the studies considered. Those studies cover different occupations, tasks, users, AI systems and measures of performance and do not establish a global average productivity effect. Effects observed on suitable tasks may also change when AI is embedded into complete workflows and organisational processes. The 15% assumption is therefore a scenario parameter rather than an estimated global mean.
Task productivity and realised economic value are different quantities	Faster task completion or greater output can release productive capacity without creating an equivalent economic benefit. Capacity may be redeployed to additional output or higher-value activity, absorbed by checking and coordination, or remain unused. Quality changes can also raise or reduce the value of an apparent productivity gain. The conversion from task performance into economic value therefore depends on what happens to the capacity released.
The conversion gap is not directly observed	There is no global, value-weighted dataset comparing realised GenAI productivity with matched potential. The 35% reference

Limitation	Implication for the analysis
	conversion gap is therefore an evidence-informed scenario. Evidence on unrealised technology value, AI scaling, enterprise financial contribution and related outcomes helps calibrate judgement, but these sources measure different populations, concepts and units and cannot be combined statistically into a single observed conversion rate. The analytical distinction between productivity potential and conversion is useful, but the two components are not separately observed at global level.
The main modelling assumptions may be related	Exposure, productivity and conversion need not vary independently. High-value applications may attract faster investment and adoption and therefore have smaller conversion gaps. Extending GenAI into more marginally exposed activities may reduce the average productivity effect. The model does not estimate these relationships. The combined cases in Section 6 are therefore mechanical stress tests rather than probability-weighted forecasts or confidence intervals.
Incomplete conversion has several causes	AI readiness shapes how successfully organisations convert AI potential into sustained performance, but other factors also matter. Technical limitations, weak use-case economics, implementation costs, capital constraints, regulation, risk appetite and normal diffusion lags can all leave economically exposed work below the modelled potential. The reference result therefore represents total modelled incomplete conversion, only part of which would necessarily be addressable through improvements in AI readiness.
The headline estimate does not deduct the full costs of realising the modelled productivity potential	The calculation does not deduct all implementation, integration, compute, supplier, training, organisational change, assurance, maintenance and transition costs required to realise the modelled productive capacity. The headline figure is therefore an estimate of productivity-equivalent value at stake rather than net economic benefit.
The core model isolates the labour-productivity channel	It does not monetise other effects examined in Section 7, including innovation, changes in non-labour inputs, knowledge spillovers, labour-market adjustment, environmental effects, risk-related externalities and market-structure effects. These effects can increase or reduce overall economic value and are therefore considered separately from the headline calculation.
The estimate is global, aggregate and current-period	The model does not allocate the headline result across firms, sectors, countries or places, and it does not model the future path of AI capability, adoption, investment or the conversion gap. The analysis combines a 2025 economic base and exposure framework with evidence from 2025–26 during a period of rapid technological change. The estimate will therefore require periodic updating as the underlying evidence changes.

Some simplifications could raise the estimate relative to net economic benefit. Transferring task-level productivity evidence to a wider exposed labour base and excluding implementation costs are two examples. Other omissions exclude potentially material value, including innovation and improvements in non-labour productivity. Current evidence does not support a robust estimate of the net direction or magnitude of these effects.

A further distinction concerns the type of uncertainty tested in Section 6. The sensitivity analysis examines parameter uncertainty within the selected model structure by varying exposure, productivity and the conversion gap. It does not quantify structural uncertainty around the choice of labour income as the valuation base, the use of GDP-weighted exposure or the use of an annual snapshot rather than a dynamic diffusion model.

The sensitivity results therefore show how the estimate responds to changes in key parameters. They are not probability bounds and do not establish that the selected model structure is uniquely correct.

10.2 Alternative approaches

The current model is intentionally transparent and relatively simple. More granular methods could improve particular components of the estimate or answer related economic questions as the evidence base develops.

Alternative approach	What it could add	Evidence requirements and appropriate use
Country–occupation–task model	Country-specific labour income, occupational compensation and task exposure could replace the global labour-share and GDP-weighting approximations. Productivity effects could also vary across occupations and tasks.	This is the most direct technical extension of the present model. It would require internationally comparable data linking labour compensation, occupations and task-level GenAI exposure at sufficient granularity.
Sector-by-sector bottom-up model	Sector-specific exposure, use cases, productivity effects, implementation costs and conversion rates could capture substantial differences in the economics of AI adoption between industries.	Useful for sector strategy and policy. A robust global aggregation would require comparable evidence on sector-level productivity effects and conversion gaps, which remains limited.
Dynamic diffusion model	Adoption, investment, organisational learning, improving AI capability and the closing of the conversion gap could be modelled through time.	Appropriate for forecasting the trajectory of value realisation. Current evidence does not yet provide stable transition or diffusion parameters for a technology evolving as rapidly as GenAI.
Firm- and use-case-level evaluation	Expected benefits and costs could be compared with subsequent performance, allowing more direct measurement of realised value and the relationship between AI readiness and conversion. Comparisons between organisations with different readiness profiles could also test whether particular capabilities are associated with stronger outcomes.	Requires longitudinal operational and financial data, credible baselines or counterfactuals and sufficiently representative samples. Causal interpretation must also account for wider differences in management quality, digital maturity, investment and organisational characteristics.
Macroeconomic or general-equilibrium model	Productivity changes could feed through into output, prices, wages, employment, investment, consumption and resource reallocation.	Appropriate for estimating economy-wide GDP or welfare effects. It addresses a broader question than the productivity-equivalent value-at-stake estimate used here and requires additional behavioural and structural assumptions.
Net firm-level economic appraisal	Benefits could be compared directly with implementation and operating costs to estimate NPV, BCR, payback or another decision measure.	Appropriate for a defined organisation, programme or use case. Section 9 shows how this fits within Economic Opportunity and Value Realisation, but the required data cannot currently be populated consistently at global level.
Full social cost-benefit appraisal	Wider economic, social, environmental and distributional effects could be assessed alongside direct productivity benefits and costs.	Appropriate for a defined intervention and counterfactual. Applying the approach globally would require extensive valuation assumptions and could create apparent precision without commensurate evidence.
Top-down external value benchmark	Published estimates of GenAI or AI value can provide an independent cross-check on the broad economic scale suggested by the bottom-up result.	Useful for triangulation. Published estimates often measure different concepts, including salary-valued task augmentation, enterprise cost and revenue potential, macroeconomic GDP effects or wider AI impacts, so they should not be treated as interchangeable denominators.

The country–occupation–task model is the clearest technical extension of the current calculation. It would replace several global approximations with country-specific labour income and occupational compensation and would allow exposure and productivity assumptions to differ across types of work. This could materially improve measurement of the potential value pool.

The more difficult evidence problem concerns conversion. A more precise estimate requires comparable evidence linking expected AI benefits, actual deployment and realised economic outcomes. Firm- and use-case-level evaluation could progressively build that evidence. The strongest designs would establish credible baselines or counterfactuals so that changes in performance could be distinguished from changes that would have occurred without the AI intervention.

As the evidence base develops, observed distributions of productivity and realised value could replace some of the scenario assumptions used here. More complex modelling would then become worthwhile because additional structure would be supported by additional evidence.

For the present analysis, greater complexity would not necessarily improve accuracy. The current model keeps the main assumptions visible, allows their effects to be tested separately and provides a transparent estimate of the economic scale of incomplete GenAI value conversion.

Taken together, these limitations define the appropriate use of the estimate. The model provides a transparent global benchmark for the scale of incomplete GenAI value conversion. More granular firm-, sector- and country-level analysis would require evidence and methods appropriate to those levels. The value of the model lies in making its economic assumptions visible, testable and updateable as the evidence improves.

11. Conclusion

AI capability is improving rapidly and becoming easier to access. The economic consequences will increasingly depend on the ability of firms, institutions and economies to convert that capability into useful performance.

Much of the early AI race has centred on models, compute, platforms and access. Those dimensions remain important. As access broadens, conversion becomes more consequential: identifying worthwhile applications, integrating them into real work, changing how work is organised, measuring results and scaling approaches that deliver value.

AI readiness shapes that process. Data, technology, workforce capability, operating arrangements, governance, supplier relationships and organisational learning can reinforce one another or become constraints. Firms with stronger conversion capability can turn successive improvements in AI into better operating performance. Public bodies face a similar challenge in services, administration and decision-making. Economies differ in the infrastructure, skills, institutions and investment required to translate technical opportunity into productive value.

The potential scale is substantial. The analysis in this paper suggests that incomplete GenAI value conversion could leave **around \$1 trillion per year of productivity-equivalent value at stake**. The estimate captures one important productivity channel within a wider set of economic effects. Innovation, new products and services, knowledge spillovers and improvements in public-sector outcomes can add further value beyond the core calculation.

The opportunity cost of weak conversion is easy to overlook. Investment expenditure, infrastructure, training, organisational change and the risks associated with deployment are visible. Foregone productivity appears differently. Output that could have been produced never enters the accounts. Capacity that could have been released remains embedded in existing processes. Improvements in products, services or public administration leave no obvious trace when they fail to occur.

Economic appraisal should account for these foregone benefits when productive AI applications are delayed, constrained or abandoned. This is particularly relevant to regulation and public policy. Effective regulation can reduce harmful externalities, protect rights and create confidence in deployment. Restrictions can also reduce productive activity. Where regulation or public policy materially constrains economically productive AI use, it carries an opportunity cost. The benefits of a restriction should be assessed alongside the productive value forgone. **A smaller economic pie may be a policy choice. It should not be mistaken for a costless one.**

Higher productivity expands the amount and quality of goods and services that can be produced from available labour, capital and time. Over sustained periods, this creates greater scope for higher real incomes, investment, stronger public services and other social priorities.

The scale also matters for public services and wider social objectives. For perspective, the **around \$1 trillion** of productivity-equivalent value at stake is more than five times the **\$174.3 billion of official development assistance provided by OECD Development Assistance Committee members and associates in 2025¹**. Realising more of that productive potential would expand the economic base available to support living standards, healthcare, education, social protection, international development and other social priorities.

AI can also improve social outcomes directly. In government and public services it can reduce administrative burdens, accelerate processing, improve analysis and forecasting, strengthen targeting and widen access to services. Faster decisions, fewer errors and better allocation of public resources can create value beyond the associated labour saving.

The distribution of AI-generated value will have important economic and social consequences. Productivity helps determine the scale of the gains available, while competition, ownership, wages, taxation and public policy influence where those gains ultimately accrue. These choices affect household living standards, public finances and the resources available to pursue wider social objectives.

Conversion capability may itself become an important source of economic advantage. Firms and economies that are better able to absorb, apply and scale AI can capture a larger share of the productivity gains, while weaker conversion capability risks widening existing differences in performance.

The central economic question is therefore how much of AI's expanding capability will become realised value. For firms, the answer affects returns and competitive position. For public institutions, it affects service quality and institutional capability. For economies, it affects productivity, living standards and the resources available to pursue wider social objectives.

There will be legitimate arguments about how those gains should be shared. **They have to be created first.**

Notes and sources

Executive summary

1. McKinsey & Company (2026), *The state of AI in 2026: On the road to ROI*. (McKinsey & Company)
2. PwC (2026), *AI Performance Study*, 13 April 2026; and *How leading companies generate ROI from AI*, 27 May 2026. (PwC)

Sections 1–3

1. McKinsey & Company (2026), *The state of AI in 2026: On the road to ROI*. McKinsey reports that 80% of respondents say AI has improved their individual productivity, while 37% report a positive contribution to enterprise earnings. The measures capture different dimensions of AI impact and should not be interpreted as components of a common conversion-rate calculation.
2. PwC (2026), *AI Performance Study*, 13 April 2026; and *How leading companies generate ROI from AI*, 27 May 2026. The global study surveyed 1,217 senior executives across 25 sectors, primarily at large publicly listed companies. PwC measured AI-driven financial performance as reported revenue and efficiency gains derived from AI, adjusted against industry medians. The top 20% of organisations captured 74% of measured AI-driven returns and were twice as likely as other organisations to redesign workflows around AI.
3. De Smet, A., Goldstein, D., Price, H., Catlin, T. et al. (2026), *From adoption to impact: Three horizons of AI transformation*, McKinsey & Company, 8 July 2026. The research surveyed 750 English-speaking employees and leaders who reported incorporating AI into their work. All 750 answered personal-readiness questions; 608 executives, senior leaders and middle managers answered organisational-readiness questions. Recruitment deliberately targeted respondents from organisations in more advanced AI-maturity categories, so McKinsey cautions that estimates may not reflect overall market prevalence.

Sections 4–6

1. World Bank, *World Development Indicators: GDP (current US\$)*, 2025 income-group aggregates. The values used in the calculation are approximately \$76.88tn for high-income economies, \$33.53tn for upper-middle-income economies, \$7.31tn for lower-middle-income economies and \$0.64tn for low-income economies.
2. ILOSTAT reports that the global labour-income share fell to 52.6% in 2025.
3. Gmyrek et al. (2025), *Generative AI and Jobs: A Refined Global Index of Occupational Exposure*, ILO Working Paper 140. The study uses task-level information to classify occupations into four progressively increasing exposure gradients and reports substantial variation in exposure by country income level.
4. Calvino, Reijerink and Samek (2025), *The effects of generative AI on productivity, innovation and entrepreneurship*, OECD Artificial Intelligence Papers No. 39. The review synthesises experimental evidence across occupations including customer support, consulting and software development.
5. Cui et al. (2026), *The Effects of Generative AI on High-Skilled Work: Evidence from Three Field Experiments with Software Developers, Management Science*. Across three field experiments involving 4,867 developers, access to an AI coding assistant was associated with a 26.08% increase in completed tasks.
6. Tinkoff, Van der Veken, Chui and Balakrishnan (2026), *The state of AI in 2026: On the road to ROI*, McKinsey & Company, 25 August 2026. The survey covered 1,719 respondents across 97 countries, with results weighted by countries' contribution to global GDP.
7. Gartner (2026), *Gartner Survey Finds Only 22% of Organizations Have Successfully Scaled AI Across Multiple Business Units*, 1 September 2026. Survey of 1,303 respondents from organisations with annual revenues of at least \$50 million.
8. Boston Consulting Group (2026), *CEOs Are Starting to See Value from AI. Now Comes Execution*, 22 July 2026. The research is based on a global survey of 152 chief executives at companies with revenues of at least \$500 million.
9. Deloitte UK (2026), *Bridging the tech value gap: Maximising returns on investment in digital transformation*. Deloitte surveyed 314 senior UK executives in September 2025 and supplemented the survey with 10 qualitative interviews conducted between October 2025 and January 2026.

Section 7

1. Calvino, F., Reijerink, J. and Samek, L. (2025), *The effects of generative AI on productivity, innovation and entrepreneurship*, OECD Artificial Intelligence Papers No. 39. The review considers experimental evidence on productivity alongside potential effects on creativity, innovation, research and development, and entrepreneurship.
2. OECD (2025), *Governing with Artificial Intelligence: The State of Play and Way Forward in Core Government Functions*. The OECD examines AI use across core government functions and identifies potential gains in productivity, responsiveness and decision-making, alongside requirements around data, infrastructure, skills, procurement and governance.
3. International Energy Agency (2026), *Key Questions on Energy and AI*. The IEA reports that data-centre electricity demand increased by 17% in 2025, with AI-focused data centres growing faster than overall data-centre demand.
4. Gmyrek, P. et al. (2025), *Generative AI and Jobs: A Refined Global Index of Occupational Exposure*, ILO Working Paper 140. The ILO estimates that around one quarter of global employment is in occupations with some GenAI exposure and concludes that transformation is currently the most likely effect for most exposed occupations because many tasks continue to require human input.
5. OECD (2026), *Trends in AI incidents and hazards reported by the media*; OECD (2025), *Governing with Artificial Intelligence*. The OECD identifies realised AI-related risks including privacy, discrimination, security and safety harms and discusses risk-based governance approaches to trustworthy deployment.
6. OECD (2026), *Artificial Intelligence Markets*. The OECD identifies declining quality-adjusted model prices and continuing dynamism in parts of the AI market, alongside structural competition risks associated with concentration in compute, data, chips and skills, vertical integration and switching costs.

Section 8

1. OECD (2026), *Artificial Intelligence Markets*. The OECD identifies concentration of key inputs, economies of scale, vertical integration and switching costs among structural competition risks across the AI value chain. It also documents substantial concentration of AI investment and compute capacity across firms and countries.
2. Brynjolfsson, E., Chandar, B. and Chen, R. (2026), *Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence*, Stanford Digital Economy Lab, revised 12 August 2026. Using administrative payroll data through June 2026, the authors report that employment among workers aged 22–25 in AI-exposed occupations was 19% below the level implied by keeping pace with less-exposed peers. The divergence was concentrated in hiring and in occupations where AI usage was more automation-oriented. The authors explicitly characterise the results as descriptive indicators rather than causal estimates.

3. OECD (2025), *Governing with Artificial Intelligence: The State of Play and Way Forward in Core Government Functions*. The study analyses 200 government AI use cases; 57% supported automated, streamlined or tailored processes and services. It also identifies skills, legacy information technology systems, data, budgets and governance requirements among constraints on government adoption.
4. PwC (2026), *AI Performance Study*, 13 April 2026; and *How leading companies generate ROI from AI*, 27 May 2026. The study surveyed 1,217 senior executives across 25 sectors, primarily at large publicly listed companies. PwC measured AI-driven performance using reported revenue and efficiency gains attributable to AI, adjusted against sector medians. The top-performing 20% captured 74% of the measured returns.
5. OECD (2025), *AI adoption by small and medium-sized enterprises*; OECD (2026), *Skills in the AI age*. OECD evidence shows substantially lower AI adoption among SMEs than large firms and identifies skills, finance, infrastructure and access to AI-enabling inputs among important barriers.
6. OECD (2025), *Generative AI and the SME Workforce: New Survey Evidence*. The representative survey covered more than 5,000 SMEs in Austria, Canada, Germany, Ireland, Japan, Korea and the UK. It found GenAI in use in 31% of SMEs and notes that lower costs and data requirements have broadened access to AI capabilities.
7. Department for Science, Innovation and Technology (2025), *Artificial Intelligence Sector Study 2024*. London, the South East and East of England accounted for approximately 75% of registered office locations of AI companies identified in the study in 2024. The number of AI firms nevertheless increased across all UK regions. Registered-office locations are a measure of the geography of the UK's AI sector, rather than a direct measure of regional AI adoption or productivity.
8. OECD (2024), *Job Creation and Local Economic Development 2024: The Geography of Generative AI*. The OECD estimates substantial within-country variation in prospective GenAI exposure. The most exposed region within a country is on average around 1.6 times as exposed as the least exposed; the cross-regional range extends from 16% in Guerrero to 77% in Greater London, while Bogotá is more than three times as exposed as La Guajira on this measure.

Section 9

1. Teece, D.J. (1986), *Profiting from technological innovation: Implications for integration, collaboration, licensing and public policy*, *Research Policy*, 15(6), pp. 285–305; Lepak, D.P., Smith, K.G. and Taylor, M.S. (2007), *Value Creation and Value Capture: A Multilevel Perspective*, *Academy of Management Review*, 32(1), pp. 180–194.
2. Brynjolfsson, E. and Hitt, L.M. (2000), *Beyond Computation: Information Technology, Organizational Transformation and Business Performance*, *Journal of Economic Perspectives*, 14(4), pp. 23–48; Bresnahan, T.F., Brynjolfsson, E. and Hitt, L.M. (2002), *Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence*, *Quarterly Journal of Economics*, 117(1), pp. 339–376.

Section 10

1. Gmyrek, P. et al. (2025), *Generative AI and Jobs: A Refined Global Index of Occupational Exposure*, ILO Working Paper 140. The framework measures potential occupational and task exposure to GenAI rather than observed adoption or realised productivity.
2. Calvino, F., Reijerink, J. and Samek, L. (2025), *The effects of generative AI on productivity, innovation and entrepreneurship*, OECD Artificial Intelligence Papers No. 39. The evidence reviewed shows substantial variation in measured productivity effects across tasks, occupations and empirical settings; the 15% reference assumption is therefore used as a rounded scenario calibration rather than an estimated global mean.

Section 11

1. OECD (2026), *International aid fell sharply in 2025, says OECD*, 9 April 2026. Preliminary data report official development assistance by Development Assistance Committee members and associates of \$174.3 billion in 2025. (OECD)