Benefits realisation of UK AI compute investments

Benefits realisation plan and return on investment of the UK Government's investments into AI compute infrastructure

Final report

May 2026

Department for Business, Innovation, Science and Trade



London Economics

With support from IFF Research

About London Economics

London Economics is one of Europe's leading specialist economics and policy consultancies headquartered in London. We advise clients in both the public and private sectors on economic and financial analysis, policy development and evaluation, business strategy, and regulatory and competition policy.

Our consultants are highly-qualified economists with experience in applying a wide variety of analytical techniques to assist our work, including cost-benefit analysis, multi-criteria analysis, policy simulation, scenario building, statistical analysis and mathematical modelling. We are also experienced in using a wide range of data collection techniques including literature reviews, survey questionnaires, interviews and focus groups.

Head Office: Somerset House, New Wing, Strand, London, WC2R 1LA, United Kingdom.

w: londoneconomics.co.uk e: info@londoneconomics.co.uk
t: +44 (0)20 3701 7700

Acknowledgements

We would like to thank the University of Bristol, the University of Cambridge, Department for Business, Innovation, Science and Trade (BIST), and UK Research and Innovation (UKRI) for the time taken to meet with the project team both through individual consultations over the course of the project and participation in the lessons learned workshop, and for the valuable information they have provided. We would also like to thank all of the stakeholders surveyed for their informative questionnaire responses. Finally, we would like to thank the team at IFF Research for their support with the surveys and user consultations undertaken for this study. **Cover image credit:** Shutterstock/Oleksiy Mark_105784187

Disclaimer

The assessments and judgements presented in this report reflect the views of the authors. Where specific stakeholder perspectives are referenced, these are clearly identified as such. Unless otherwise stated, all views expressed are those of the authors. Responsibility for the content of this report remains with London Economics.

Authors

Moritz Godel, Partner
Daniel Herr, Associate Director
Bruno Chereque, Economic Consultant
Marina Symington, Economic Consultant

Wherever possible London Economics uses paper sourced from sustainably managed forests using production processes that meet the EU eco-label requirements.

Copyright © 2026 London Economics. Except for the quotation of short passages for the purposes of criticism or review, no part of this document may be reproduced without permission.

Table of Contents

Page

Glossary ii

Executive Summary v

PART I: BACKGROUND AND CONTEXT 1

1 Introduction 2

2 The UK Government's investments in AI Compute 3

3 UK compute landscape and international context 6

PART II: BENEFITS REALISATION 11

4 Introduction and approach 12

5 Benefits realisation plan 12

6 Summary and key takeaways 27

PART III: ECONOMIC VALUATION 30

7 Introduction and approach 31

8 Assessment of economic benefits and return on investment 32

9 Central valuations of economic benefits and RoI 38

10 Sensitivity analysis surrounding central benefit estimations 42

11 Comparison to estimated benefits of other evaluations of HPC and science infrastructure 51

12 Summary of economic valuation and key takeaways 53

PART IV: CONCLUDING REMARKS 57

13 Concluding remarks 58

INDEX OF TABLES AND FIGURES 59

Index of Tables and Figures 60

ANNEXES 62

Annex 1 Detailed evaluation approaches 63

Annex 2 Sensitivity analysis methodology 80

Glossary

Term	Description
AI	Artificial Intelligence
AI accelerators	Specialised hardware designed to perform AI computations more efficiently than general-purpose processors, typically used to train and run AI models at scale.
AI compute	Specialised High-Performance Compute capacity for AI use-cases such as training, validating, and deploying AI models
AI hyperscalers	Large technology companies that provide cloud computing services and infrastructure to support AI development and applications. Key players include Microsoft, Google and Amazon
AIRR	The AI Research Resource, a suite of publicly funded advanced supercomputers that provides AI-specialised compute capacity to researchers, academia, and industry
AIRR 20x	The 20x expansion of AIRR by 2030, committed to in the UK Compute Roadmap[1]
AIRR Phase 1	The initial investment as part of the AIRR programme into two AI supercomputers – Dawn (Cambridge) and Isambard-AI (Bristol)
AIRR Portal	A unified access point for the AIRR programme
AISI	The AI Security Institute, a UK state-backed organisation dedicated to ensuring advanced AI is safe, secure and beneficial
Alan Turing Institute	The UK's national institute for data science and Artificial Intelligence
AMD MI300X	A discrete GPU chip from AMD designed to deliver leadership performance for Generative AI workloads and HPC applications.
BriCS	The Bristol Centre for Supercomputing, which hosts Isambard-AI
Cloud-based compute	Referring to cloud-based compute offerings from commercial cloud providers (e.g., Amazon AWS, Google Cloud, Microsoft Azure, etc.)
Co-processor technologies	Specialised processors that work alongside a central processing unit (CPU) to perform specific tasks more efficiently, such as AI or high-performance computing workloads.
CPU	Central Processing Unit, a traditional general-purpose computing processor designed to handle a wide range of tasks sequentially and efficiently
Dawn	The AIRR AI supercomputer at the University of Cambridge
BIST	The UK Government Department for Business, Innovation, Science and Trade
EPCC	The Edinburgh Parallel Computing Centre, a supercomputing centre based at the University of Edinburgh

[1] DSIT (2025). UK Compute Roadmap. Policy paper. Available at: https://www.gov.uk/government/publications/uk-compute-roadmap/uk-compute-roadmap

Term	Description
EuroHPC	The European High Performance Computing Joint Undertaking, a joint initiative between the EU, European countries and private partners to develop a World Class Supercomputing Ecosystem
Federated access	An authentication approach that allows users to log in to multiple independent systems using a single, trusted identity provider
Frontier AI	Highly capable leading-edge general-purpose AI models that can perform a wide variety of tasks
Generative AI	Artificial intelligence that can create new content, such as text, images, music, or code, based on the data it has been trained on.
Gigafactories	Very large-scale computing facilities designed to train and run advanced AI models, typically using significant amounts of specialised hardware and energy.
GPU	Graphics Processing Unit, specialised processors designed to perform many calculations simultaneously
HPC	High-Performance Computing
HPE	Hewlett Packard Enterprise
Intel Data Centre GPU Max	Intel's highest performing, highest density, general-purpose discrete GPU for data centre, AI, and HPC workloads
Isambard-AI	The AIRR AI supercomputer at the University of Bristol
M&E	Monitoring & Evaluation
Machine learning	A subset of artificial intelligence that enables computer systems to learn patterns from data and make predictions or decisions without being explicitly programmed for each task.
MyAccessID	A centralised identity management layer and single sign-on across the different portals
NVIDIA A100	A high-performance GPU (Graphics Processing Unit) widely used for AI training and high-performance computing workloads. It was a leading AI compute chip prior to newer generations such as the H100 (released two years later in 2022).
NVIDIA Blackwell B200	A high-performance GPU from Nvidia based on the Blackwell architecture, designed for AI and high-performance computing workloads. It represents a newer generation of AI hardware with improved performance compared to earlier chips such as the H100 and H200.
NVIDIA Feynman	A future-generation AI compute architecture referenced in Nvidia's product roadmap, expected to succeed current architectures such as Blackwell and support advanced AI and high-performance computing workloads.
NVIDIA GH200 Grace-Hopper Superchip	A state-of-the-art compute chip from Nvidia combining CPU and GPU processing in a single chip designed for at-scale AI and high-performance compute applications.
NVIDIA H100	A previous leading discrete GPU chip for AI applications from Nvidia, now succeeded by its successor the H200.
PFlop/s	Petaflops per second, a unit of measurement for how many numerical calculations a processor or supercomputer can perform each second: One PFlop/s is equal to one quadrillion floating-point operations per second
PSRE	Public Sector Research Establishment, a diverse collection of public bodies that carry out research to support a wide range of government objectives

Term	Description
R&D	Research & Development
RAG	Red-Amber-Green, a widely used status classification system for reporting progress, performance, or risk in projects and operations
RCS	The University of Cambridge's Research Computing Service, which hosts Dawn
SMEs	Small- and Medium-sized Enterprises
Sovereign-AI	Sovereign-AI refers to the UK government's initiative to ensure that the UK can design, procure, deploy, assure, and adapt AI systems in areas that are aligned with UK Government priorities such as life sciences, health, materials discovery, and fundamental AI research without undue dependence on external actors
Tier-1	National-scale HPC services such as ARCHER2 designed to run the largest and most complex computational simulations
Tier-2	Mid-scale HPC facilities providing targeted access at a regional level or for specific disciplines
Tier-3	Local-scale HPC facilities provide HPC access to users within a specific institution or department
Top500	The Top500 list ranks, twice a year, the world's most powerful supercomputers
TRE	Trusted Research Environments, also known as secure data environments, provide secure access to sensitive or confidential data for researchers, enabling researchers to make use of this data in their research while keeping data stored within the TRE safe
TRL	Technology Readiness Level, a widely used framework for assessing a technology's stage of maturity
UKAEA	The UK Atomic Energy Authority, an executive non-departmental public body, sponsored by the Department for Energy Security and Net Zero.
UKRI	UK Research and Innovation, a non-departmental public body sponsored by BSIT, that brings together the seven disciplinary research councils, Research England, which is responsible for supporting research and knowledge exchange at higher education institutions in England, and the UK's innovation agency, Innovate UK

Executive Summary

Background and context

In December 2025, London Economics and IFF Research were commissioned by the Department for Business, Innovation, Science and Trade (BIST) to undertake a benefits realisation exercise of the UK Government's investments into AI compute infrastructure. In scope of the study are the Government's investments into Isambard-AI and Dawn, the 20x expansion of the AI Research Resource (AIRR) programme, as well as investments into related national AI compute infrastructure.

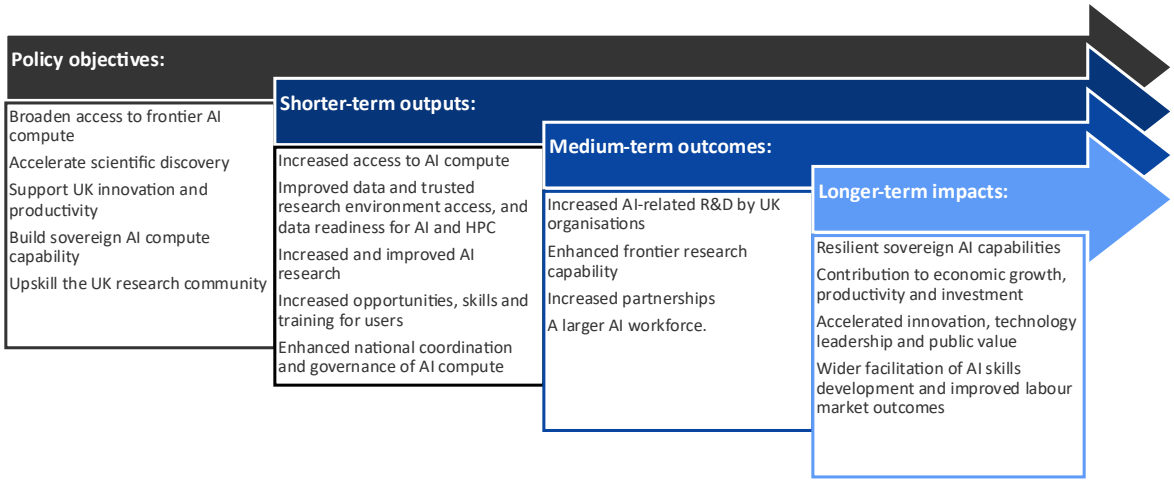
The Government's investments in AI compute are designed to strengthen the UK's national AI capability and address long-standing constraints in access to advanced AI compute. Through significant expansion in compute capacity, the investments aim to broaden access to cutting-edge AI compute, accelerate scientific discovery, and build sovereign capability in AI infrastructure. The compute capabilities are intended to support AI and AI-enabled research and innovation across academia, industry, SMEs, and the public sector, while contributing to economic growth and the upskilling of the UK research and innovation community.

The present research builds on existing studies examining the benefits of public large-scale compute infrastructure to develop a forward-looking benefits realisation plan and quantify the potential Return on Investment (RoI) that reflects the evolving landscape of AI compute in the UK. The report is designed to support the UK Government and other AI compute stakeholders to achieve a clear and shared understanding of the benefits and RoI that AI compute can deliver, the conditions and actions required for those benefits to materialise, and the roles and responsibilities for delivery and monitoring over time.

Evidence gathering combined quantitative and qualitative methods, including a user survey of existing users of the two AIRR compute facilities already available (Dawn and Isambard AI), follow-up stakeholder interviews with users and engagement with BIST, UK Research and Innovation (UKRI) and two existing host sites (the University of Bristol and the University of Cambridge), and desk research and literature reviews.

Benefits flowing from AI compute

The Government's investments in public AI compute and supporting activities are expected to deliver a broad range of benefits in line with its overarching policy objectives; these include:



OFFICIAL

Economic valuation: Return on Investment

The economic analysis sought to value the benefits of the UK's investments in AI compute, over their operational life in the next 10 years, against a hypothetical scenario where these investments would not have taken place (the "counterfactual scenario"). Under this scenario, the AIRR and related national AI compute capabilities would not have been provided, and users would have to use alternative compute infrastructure to the extent that this is possible.

Economic valuation of benefits, and comparison against expected costs, suggests a sizeable return on the UK Government's investments in large-scale public AI compute. **The Return on Investment (RoI) is estimated at between £4.2 : £1 and £20.2 : £1, with a central estimate of £6.0 per £1** of public monies invested into AIRR and related national AI compute infrastructure.

This central estimate corresponds to total benefits of approximately £12.1 billion (discounted to 2025/26) and total costs of £2 billion (also discounted to 2025/26) over the AI compute capabilities' operational life, based on investment plans committed to in the UK Compute Roadmap.[2]

Annualising these central benefits over the 10-year appraisal period implies benefits of approximately £1.2 billion per year. This is equivalent to around 10% of the current annual GVA of the UK AI sector (estimated at £11.8 billion),[3] indicating that the programme represents a material contribution to the UK's AI-driven economic activity.

Table 1 Aggregate economic valuation using central assumptions

Metric	Low estimate	Conservative central assumptions	High estimate
Total monetised economic benefits	£8.5 billion	£12.1 billion	£40.6 billion
Return on Investment (RoI)	**£4.2 per £1 invested**	**£6.0 per £1 invested**	**£20.2 per £1 invested**

Note: The low and high estimates correspond to the 5th and 95th percentiles of the results distribution from a Monte Carlo simulation conducted to capture potential outcomes under uncertainty. ***Source: London Economics analysis***

The estimated RoI range reflects the substantial uncertainty surrounding these benefit estimates (see Section 7), given the early stage of the compute investments and the inherently forward-looking nature of the analysis. Key sources of uncertainty include the timing and composition of the compute infrastructure rollout, its uptake and use across different user groups, and the timing and magnitude of benefits ultimately arising from AI research enabled by these capabilities. Therefore, conservative input assumptions have been used throughout in the central valuation of benefits.

Comparison to other evaluations of large-scale compute infrastructure shows that, while the central estimate (£6.0 : £1) is in line with existing evidence, it sits on the lower side of estimated returns, which are typically reported in the range of around £2 : £1 to £20 : £1. This is due to the conservative assumptions adopted in the central case, reflecting the significant uncertainty present in the RoI

[2] Cost figures include the initial investments into Isambard-AI and Dawn, the 20x expansion of AIRR and investments in related public AI compute committed to in the UK Compute Roadmap (DSIT (2025). UK Compute Roadmap. Policy paper. Available at: https://www.gov.uk/government/publications/uk-compute-roadmap/uk-compute-roadmap).
[3] Estimates from the 2024 UK AI sector study; DSIT (2025). Artificial Intelligence sector study 2024. Available at: https://www.gov.uk/government/publications/artificial-intelligence-sector-study-2024/artificial-intelligence-sector-study-2024.

OFFICIAL

modelling. However, the estimated RoI range of £4.2 to £20.2 per £1 invested[4] is consistent with the evidence reported in comparable evaluations of large-scale compute infrastructure.

The sensitivity analysis results further highlight that benefits could be substantially larger, with a small number (1%) of simulations suggesting benefits in excess of £64.3 billion (£32.2 : £1). These results reflect higher impact scenarios where favourable assumptions are chosen across many inputs. At the same time, around one-fifth (20%) of simulations also show benefits below the central estimates. 5% of simulations resulted in benefits of £8.5 billion (£4.2 : £1) or less. 1% of simulations resulted in benefits of £6.3 billion (£3.1 : £1) or less. However, even the lowest of these results implied a positive RoI, suggesting that even if conditions were extremely unfavourable public funds would still have delivered value for money.

In addition to the conservative assumptions, the impact of AI-enabled R&D may be understated by the existing evidence on the returns to R&D. AI has the potential to accelerate the scientific discovery and innovation process. If this leads to faster accumulation and diffusion of knowledge, the long-run returns to early R&D investment could be even higher than suggested by historical literature. However, the magnitude and persistence of such effects remain uncertain. Therefore, the central estimations do not assume additional returns and rely on existing evidence.

AI innovations enabled by public AI compute further have the potential to deliver disproportionally large, though highly uncertain, impacts such as large-scale transformation of public service delivery or the emergence of highly successful British companies comparable to leading AI firms.[5] For example, recent figures by the Alan Turing Institute show that around 40% of tasks across the public sector could be supported by AI.[6] If realised this could deliver substantial productivity savings, with ONS figures estimating employment in the public sector at 6.19 million in December 2025.[7] Similarly, the potential scale of the commercial upside is substantial as is illustrated by Open-AI's most recent valuation which, as of February 2026, was valued at $840 billion.[8]

Benefit estimates by type of benefit

The largest source of benefits in the central valuations were wider economic impacts, accounting for 70% of benefits. This group of benefits captures spillovers to the wider economy from AI research undertaken on national AI compute capabilities as well as downstream adoption of AI innovations by customers of firms developing AI products on AIRR and related public AI compute infrastructure and wider use across sectors.

[4] To enable an understanding of the uncertainty surrounding central estimates, sensitivity analysis in the form of Monte-Carlo Simulations was undertaken. In simple terms, this approach allows input assumptions to vary across a large number of model runs. For each model run, assumptions are drawn from predefined probability distributions, reflecting uncertainty informed by the literature and robustness of available evidence, and used as inputs to the economic model. Repeating this process a large number of times generates a range of possible outcomes and associated likelihoods for total economic benefits and the associated Return on Investment. In addition to uncertainty around the central input assumptions, the sensitivity analysis allows for the possibility of higher returns to AI research than suggested by historical literature estimates as well as for the potential for outsized returns associated with high-impact outcomes.

[5] While the sensitivity analysis allows for the possibility of larger returns, potential outcomes delivering extremely large benefits such as those mentioned (large-scale transformation of public services, emergence of globally leading AI companies, etc.) are difficult to capture and so could deliver benefits beyond even those indicated by the higher results from the sensitivity analysis.

[6] Alan Turing Institute (2025). Mapping the Potential: Generative AI and Public Sector Work. Using time use data to identify opportunities for AI adoption in Great Britain's public sector. Available at: https://www.turing.ac.uk/sites/default/files/2025-05/ons_tus_final_report.pdf

[7] ONS (2026). Public sector employment, UK: December 2025. Available at: https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/publicsectorpersonnel/bulletins/publicsectoremployment/december2025

[8] Tech Insider (2026). OpenAI's $110 Billion Funding Round: Inside the Largest Private Investment in History and What It Means for the AI Race. Available at: https://tech-insider.org/openai-110-billion-funding-round-2026/

Table 2 Central valuations by benefit grouping

Benefit grouping	Conservative central valuation	%
Wider economic impacts	£8,440M	70%
Access to AI compute	£1,145M	10%
Private benefits from work undertaken on public AI compute	£930M	8%
Training and skills benefits	£870M	7%
Benefits to the public sector*	£665M	6%
Total economic benefit	**£12,050M**	**100%**

Note: Figures rounded to the nearest £5 million. (*) Only includes efficiency savings among public sector staff affected by adoption. Wider benefits such as improved health-care outcomes or wider societal benefits are not captured. ***Source: London Economics analysis***

Cost savings from access to public AI compute relative to accessing commercial alternatives and private benefits to users of public AI compute (such as academic publications, the generation of intellectual property, and software products, as well as increased revenue generation among firms using public AI compute) accounted for 10% and 8% of estimated benefits, respectively.

Training and skills benefits and public sector benefits accounted for the lowest share of central benefit estimates (7% and 6%, respectively). However, this is primarily driven by the conservative scope of what is included in the valuation and the exclusion of wider, less easily monetised impacts such as improved healthcare outcomes, welfare gains, or broader societal impacts.[9]

In addition, the composition of benefits could vary significantly depending on the type of AI innovations that will be produced on AIRR and related public AI compute and their ultimate uptake in industry, the public sector, or wider society. In particular, the sensitivity analysis shows the potential for both private benefits to users and public sector efficiency outcomes to be substantially larger than indicated by these conservative central estimates (see Figure 9 and Figure 10).

Benefits realisation

The successful realisation of these benefits depends on the actions of both BIST, its delivery partners - such as UK Research and Innovation (UKRI) - and the AI compute host sites, each with specific roles and responsibilities, as well as a set of external factors largely outside of DIST's control (some of which can nevertheless be influenced by BIST).

Overall strategic oversight, funding and evidence-based allocation of resources

BIST has overarching responsibility for the continued funding and strategic oversight of the AIRR programme and investments in related public AI compute infrastructure. BIST also has responsibility for planning the composition of these infrastructures and architecture types, as well as setting the rollout timelines and refresh cycles; as continued investment is required to periodically update hardware and prevent it from becoming obsolete.

[9] In the case of the public sector, the estimates capture only efficiency savings among public sector staff arising from the adoption of AI innovations developed on public AI compute. Wider benefits from public sector adoption, such as improved healthcare outcomes, welfare gains, or broader societal impacts (e.g. safer communities), are not included in the valuation. However, qualitative evidence from interviews and the impact survey highlights emerging applications in areas such as healthcare data and diagnostic support, illustrating the types of impacts that may materialise more widely. Similarly, training and skills benefits capture the direct value of training provision and the benefits accruing to individuals and associated tax benefits. Wider system-level benefits of skills development, such as spillovers to firms, innovation capacity, or productivity improvements, are not included to avoid double counting with other benefit streams.

To allocate funding for maximum benefit, a clear and evidence-based understanding of demand is essential. This will be informed by a user demand study (already underway). Based on this evidence, BIST will develop an allocation strategy that balances the needs of multiple user groups (public sector, industry and academic) whose engagement underpins different benefit streams. BIST and its delivery partners hold overall responsibility for evaluating applications. However, developments in AI-enabled science may expand or shift the profile of users over time. BIST should monitor such changes to ensure that compute capabilities continue to meet emerging and future demand.

The host sites are responsible for capacity management and will provide BIST and delivery partners (UKRI) with data on utilisation, availability and user constraints to inform future allocation cycles.

Strategic alignment of compute investments

To ensure maximum impact, it is important that investments in the expansion of AIRR are aligned with other national investments such as national supercomputing centres and other AI-related initiatives. To this end, BIST will ensure strategic alignment between AIRR and other national AI-initiatives as well as with delivery partners such as UKRI and the host sites. This includes investments in the broader AI ecosystem such as in software programmes, people and skills development initiatives, and dataset curation and storage.

Moreover, the UK's compute investments take place against a background of vastly accelerated public and private compute investments globally. BIST's investment and service design choices should consider the broader international ecosystem so that they maximise national benefits while exploiting synergies from other parties' investments.

External communications

External communications are another critical enabler because many intended benefits depend on attracting the right users to the system. Therefore, targeted outreach to industry with a particular emphasis on engaging currently underrepresented groups such as small and medium-sized enterprises (SMEs) and start-ups will be undertaken.

While BIST will lead on the development of an overarching external communications and partnerships plan, delivery partners and the host sites will support in delivery of the plan by promoting the AI compute resources within their existing networks and by providing inputs.

User support and AIRR community

Once the appropriate user groups are accessing the systems (through effective external communications and allocation of resource), user support and community-building are essential to ensure that benefits are realised in practice. Primary responsibility in this area sits with the host sites, supported by BIST and its delivery partners where appropriate. The host sites are responsible for providing technical support, training resources, and user-facing guidance to enable researchers and innovators to use their system effectively for their intended purposes as well as in engaging in and facilitating community-building activities.

External factors influencing benefits realisation

Alongside the direct actions and responsibilities of BIST, delivery partners (e.g., UKRI), and the host sites, realisation of benefits depend on a range of wider enabling conditions and ecosystem factors that sit largely outside the direct control of the AIRR programme (though some of which BIST and the wider UK Government can influence).

These external factors include **structural factors** (such as the regulatory and legal framework, the UK's AI skills and talent pipeline, and the research, funding and commercialisation landscape); **ecosystem factors** (such as the risk appetite and innovation propensity among UK entrepreneurs and the effectiveness of commercialisation and scale-up support); **market factors** (such as end-user demand and willingness to adopt AI solutions); **geopolitical factors** (such as international competition); and **technological factors** (such as the pace/direction of change in AI globally).

The degree to which the UK's **historical difficulties in scaling start-ups and retaining high-growth firms domestically** can be addressed also presents a noteworthy external factor influencing realisation of some of the longer-term and strategic impacts of the AIRR programme and related national AI compute investments.

Monitoring of benefits

To evidence benefit realisation and track whether the programme remains on course to achieve its objectives, monitoring of benefits will also be important. A suggested set of indicators to track benefits, alongside expected timelines for benefits to emerge, are provided in Section 5.4.2.

Concluding remarks

The economic valuation shows a very positive and promising picture of the potential benefits enabled by AIRR and public investments in related national AI compute infrastructure.

While central benefit estimations indicate a Return on Investment (£6.0 : £1) on the lower end of the range indicated by comparable studies of public large-scale compute infrastructure (around £2 : £1 to £20 : £1), these estimates should be interpreted as conservative, reflecting the use of cautious assumptions throughout given the early-nature of the evaluation and the substantial uncertainty around future impacts.

The economic valuation further highlights the substantial potential upside in economic benefits that AI compute could unlock under favourable conditions. AI has the potential to accelerate the scientific discovery process, deliver substantial benefits to UK private sector productivity, and help drive broader changes in public service delivery.

However, realisation of these benefits depend on effective implementation by government including on-going investment not only in compute but also in the broader AI ecosystem to ensure that hardware delivers maximal impact; evidence-based scale-up and allocation of resources; effective engagement of potential users and user support; cross-government coordination and alignment with other national investments; as well as wider government efforts to ensure positive structural and ecosystem conditions.

Nevertheless, the analysis suggests that even if conditions were extremely unfavourable public funds would still have delivered value for money, with even the lowest estimates showing RoI figures above £1 per £1 invested in AI compute.

Alongside effective Government action, ongoing monitoring of outputs and emerging benefits will be important to ensure the investments remain on track and enable timely adjustments in response to new information or changing conditions.

PART I: BACKGROUND AND CONTEXT

Key findings:

- **Strong and growing demand for AI compute:** access to advanced, AI-optimised compute remains a key constraint across academia, industry, and the public sector.
- **Clear case for public intervention:** Existing cloud and HPC provision do not fully meet the requirements of cutting-edge AI workloads.
- **AIRR delivers a step-change in capacity:** The programme, including the planned 20x expansion, significantly increases national AI compute capability and access.
- **Aligned with wider UK strategy:** Investments support broader policy priorities, including growth, innovation, and public service transformation.
- **Global competition reinforces urgency:** Rising international investment and the UK's declining compute share strengthen the need for action.

1 Introduction

1.1 About this study

In December 2025, London Economics and IFF Research were commissioned by BIST to undertake a benefits realisation exercise of the UK Government's investments into AI compute infrastructure. In scope of the study are the Government's investments into Isambard-AI and Dawn, the 20x expansion of the AI Research Resource (AIRR) programme, as well as investments into related national AI compute infrastructure.

The Government's investments in AI compute are designed to strengthen the UK's national AI capability and address long-standing constraints in access to advanced AI compute. Through significant expansion in compute capacity, the investments aim to broaden access to cutting-edge AI compute, accelerate scientific discovery, and build sovereign capability in AI infrastructure. The compute capabilities are intended to support AI and AI-enabled research and innovation across academia, industry, SMEs, and the public sector, while contributing to economic growth and the upskilling of the UK research and innovation community.

There is a growing body of evidence on the impact of public investment in High-Performance Compute (HPC) infrastructure (see Section 11). The findings from this body of research highlight the significant returns of large-scale compute in terms of research outputs, industrial innovation, skills development, and broader economic spillovers. However, existing studies predominantly focus on HPC investments rather than AI compute specific investments.

The present research builds on the existing studies to develop a forward-looking benefits realisation plan and quantify the potential return on investment that reflects the evolving landscape of AI compute in the UK. Specifically, the aim of this research was to:

- Develop a **benefits realisation plan** setting out how benefits of investments in AI compute will be realised, who is responsible and accountable for ensuring this, define indicators to measure benefits, and set out timelines for when and why benefits may or may not appear.
- Quantify **Return on Investment** (RoI) for AI compute assessing benefits across research, innovation, industrial applications, and the public sector.

The report helps the UK Government and other AI compute stakeholders to achieve a clear and shared understanding of the benefits that AI compute is designed to deliver, the conditions and actions required for those benefits to materialise, and the roles and responsibilities associated with their delivery and monitoring over time. This will support ongoing monitoring, stakeholder coordination, and future evaluation activities.

1.2 Overarching approach

The research followed a multi-modal approach, starting with a mapping of the range of benefits expected to flow from the UK's investments in AIRR and related national AI compute infrastructure. The benefit mapping is grounded on an aggregated Theory of Change of AIRR and investments in related national AI compute, which provides the overarching structure for describing how programme activities translate into benefits over time. The Theory of Change built on, and expanded, work undertaken for the *Baseline evaluation of the AI Research Resource* study (published alongside this report).

The benefits realisation plan expands on the links within the Theory of Change by setting out in more detail how benefits of investments in AI compute will be realised, and how these can be measured, who is responsible and accountable for ensuring this, and when and why benefits may or may not appear. The plan was informed literature review of existing studies of large-scale compute infrastructure, desk research, and London Economics' knowledge, and insights gathered from stakeholders, of the AIRR programme through interviews and workshops.

The Theory of Change also formed the basis for the economic modelling, which developed a forward-looking benefits-costs model capturing benefits identified in the Theory of Change. The evaluation approaches employed built on previous evaluations of large-scale compute infrastructure, but refines the analytical approaches and assumptions used, where this was possible, to better reflect the evolving landscape of AI compute in the UK.

Evidence gathering combined quantitative and qualitative methods, including a user survey of existing users of the two AIRR compute facilities already available (Dawn and Isambard AI), follow-up stakeholder interviews with users and engagement with BIST, UKRI and two existing host sites (the University of Bristol and the University of Cambridge), and desk research and literature reviews, to inform the evaluation of benefits.

The research has been conducted by a team of independent professional economists. Estimates of economic impacts are based on best practice and best judgment to calculate the most robust and fair estimates. The specific evaluation approaches used for each monetised benefit stream, assumptions made, and the caveats and limitations are, where appropriate, clearly set-out alongside the discussions of the results in Part III of this report. Annex 1 and Annex 2 provide further details.

2 The UK Government's investments in AI Compute

2.1 Background and context

Artificial Intelligence (AI) has seen rapid advances in recent years, reflecting AI's potential for transformative change across a wide range of sectors, including the public sector. The UK Science and Technology Framework identifies AI as one of the five most critical technologies for the UK economy.[10]

Alongside data and software that effectively leverages compute hardware, compute is fundamental for AI to unlock its transformative potential. Existing cloud-computing solutions, such as Google Cloud's Vertex AI (an AI development platform by Google), and local/regional supercomputing clusters enable researchers and businesses to train, deploy and scale AI to a degree.

However, while commercial cloud services should be seen as a key component of the UK compute infrastructure, they may be unsuitable for certain users who need specific networking or compute performance, and for those who need to port large datasets between systems.[11] Similarly, existing supercomputing offerings are not optimised to service cutting-edge AI developments and typically lack sufficient provision of hardware accelerators (Graphics Processing Units, or GPUs, specialised

[10] DSIT (2023). Science and Technology Framework. Available at: https://www.gov.uk/government/publications/science-and-technology-framework/science-and-technology-framework

[11] DSIT (2023). Independent Review of The Future of Compute: Final report and recommendations. Available at: https://www.gov.uk/government/publications/future-of-compute-review/the-future-of-compute-report-of-the-review-of-independent-panel-of-experts

processors designed to perform many calculations simultaneously), which cutting-edge AI are reliant upon.

The 2023 Independent Review of The Future of Compute found that there are significant shortages in public accelerator capacity (specialised hardware used for high-performance computing) in the UK. Facilities providing this capability are already running at capacity or being oversubscribed. At the time of the review, fewer than 1,000 NVIDIA A100 GPUs (a type of high-performance AI chip, see Glossary for details) were available to researchers, compared to at least 3,000 top-spec accelerators needed to meet the immediate needs of academic and commercial users in the AI community.[12]

In light of this AI compute gap and the findings and recommendations from the Future of Compute review, in particular to "*immediately and significantly increase compute capacity for AI research*", the AIRR programme aims to increase compute capacity for AI research through creating a national AI research resource.

2.2 The AIRR programme and related national AI compute infrastructure

In July 2025, the Department for Science, Innovation and Technology (DSIT) published the UK Compute Roadmap.[13] The roadmap sets out the Government's vision for the UK to become a global leader in AI and advanced compute. The roadmap identified the need for a step-change in national compute capacity, underpinned by a robust evidence base on user demand, access models, and barriers to adoption.

The AI Research Resource (AIRR) programme forms a central component of this roadmap. It aims to address current gaps in compute provision and to future-proof the UK's research infrastructure. AIRR will provide access to advanced compute infrastructure for AI research and innovation for a diverse user base across academia, public sector bodies, and small and medium-sized enterprises.

AIRR's mission is to ensure that the UK's research and innovation ecosystem can fully exploit the opportunities of AI, High-Performance Computing (HPC), and emerging technologies, while maintaining the highest standards of security, accessibility, and responsible use.

The initial AIRR programme comprised investments totalling £250 million into two new AI-optimised supercomputers: Dawn (hosted at the University of Cambridge) and Isambard-AI (hosted at the University of Bristol). These investments, and the benefits flowing from them, are further explored in the *Baseline evaluation of the AI Research Resource* study published alongside this report.

To further strengthen the UK's national AI capability and address long-standing constraints in access to advanced AI compute, the UK government plans to invest a further up to £2 billion in the public compute ecosystem, committed to in the UK Compute Roadmap.[14] This includes more than £1 billion in AI compute infrastructure to expand capacity of AIRR twenty-fold by 2030, up to £250 million to scale-up cloud compute capacity for AIRR and £750 million in further AIRR expansion

[12] Ibid

[13] DSIT (2025). Policy paper. UK Compute Roadmap. Available at: https://www.gov.uk/government/publications/uk-compute-roadmap/uk-compute-roadmap

[14] DSIT (2025). UK Compute Roadmap. Policy paper. Available at: https://www.gov.uk/government/publications/uk-compute-roadmap/uk-compute-roadmap

projects including £36 million to further upgrade Dawn.[15,16] In addition, the Government is also undertaking investments in related HPC and AI compute infrastructure including up to £750 million for a new national supercomputing service in Edinburgh.[17]

Together, these investments aim to broaden access to cutting-edge AI compute, accelerate scientific discovery, and build sovereign capability in AI infrastructure. The programme is intended to support AI and AI-enabled research and innovation across academia, industry, SMEs, and the public sector, while contributing to economic growth and the upskilling of the UK research and innovation community.

AI compute will support wider government priorities, including economic growth and productivity, scientific and technological leadership, transformation of public services, national security and long-term resilience of critical infrastructure. By targeting AI compute towards areas of highest potential economic value, and by shaping how it is supplied, governed, and used, AIRR supports the delivery of government missions while strengthening the UK's position as a leading science and innovation nation.

2.3 Wider policy context

The AIRR programme is one of several UK government initiatives and programmes aimed at ensuring the UK harnesses the substantial opportunities of AI, while promoting its safe and responsible development. This includes:

- the published (in July 2025) **UK Compute Roadmap** setting out the Government's plan to build a world-class UK compute ecosystem;[18]
- the **AI Opportunities Action Plan** (launched in January 2025), which sets out a roadmap for government to capture the opportunities of AI to enhance growth and productivity and create tangible benefits for UK citizens, including expanding the UK's sovereign compute capacity by at least 20x by 2030;[19, 20]
- the **AI for Science Strategy** (November 2025) setting out the UK's plans for developing frontier capability in AI-driven science and ensuring the UK retains its position of global scientific leadership;[21]
- the **AI Security Institute** (established in November 2023), the world's first state-backed organisation focused on advanced AI security for the public interest;[22]

[15] DSIT (2026). Policy paper. AI Opportunities Action Plan: One Year On. Available at: https://www.gov.uk/government/publications/ai-opportunities-action-plan-one-year-on/ai-opportunities-action-plan-one-year-on
[16] DSIT (2026). Cambridge supercomputer set to get 6 times more powerful as government backs British AI innovation. Press release. Available at: https://www.gov.uk/government/news/cambridge-supercomputer-set-to-get-6-times-more-powerful-as-government-backs-british-ai-innovation
[17] DSIT (2025). UK Compute Roadmap. Policy paper. Available at: https://www.gov.uk/government/publications/uk-compute-roadmap/uk-compute-roadmap
[18] DSIT (2025). Policy paper. UK Compute Roadmap. Available at: https://www.gov.uk/government/publications/uk-compute-roadmap/uk-compute-roadmap
[19] DSIT (2025). Artificial Intelligence (AI) Opportunities Action Plan. Available at: https://www.gov.uk/government/publications/ai-opportunities-action-plan/ai-opportunities-action-plan
[20] DSIT, PMO, 10 Downing Street (2025). Press release: Prime Minister sets out blueprint to turbocharge AI. Available at: https://www.gov.uk/government/news/prime-minister-sets-out-blueprint-to-turbocharge-ai
[21] DSIT (2025). AI for Science Strategy. Policy Paper. Available at: https://www.gov.uk/government/publications/ai-for-science-strategy/ai-for-science-strategy
[22] DSIT (2024). Introducing the AI Safety Institute. Available at: https://www.gov.uk/government/publications/ai-safety-institute-overview/introducing-the-ai-safety-institute

- the **Technology Adoption Review** (launched in December 2024), which will identify the barriers businesses face when adopting both established and novel technologies;[23]
- specific government initiatives aimed at boosting AI adoption such as the **AI Upskilling Fund**[24] (which was open for application for funding until August 2024) and the **Prime Minister's AI Exemplars programme** (announced in August 2025);[25] as well as,
- the UK's wider modern **Industrial Strategy**, Invest 2035, (published in June 2025), the government's 10-year plan to deliver the certainty and stability businesses need to invest in the high growth sectors expected to drive UK growth.[26]

3 UK compute landscape and international context

3.1 AI compute's place in the UK compute infrastructure

The UK's existing large-scale compute landscape

The UK has a long tradition of public compute provision. The UK's first national supercomputing service ("CRAY T3D") was launched at EPCC (the University of Edinburgh's supercomputing centre) in 1994, building on EPCC's earlier experience with production-scale parallel computers dating back to the 1980s.[27]

The UK's High-Performance Computing (HPC) ecosystem has traditionally been categorised into three tiers:

- **National-scale** (Tier-1) services, such as ARCHER2, the UK's previous flagship national supercomputer, designed to run the largest and most complex computational simulations.
- **Mid-scale** (Tier-2) facilities such as Isambard-3 (a High-Performance Computing system located at the University of Bristol designed to support a broad range of scientific workloads) and JASMIN (a High-Performance Computing and data analysis platform for environmental science at STFC's Rutherford Appleton Laboratory), among others, provide targeted access at a regional level or for specific disciplines.
- **Local-scale** (Tier-3) facilities provide access to HPC to users within a specific institution or department.

Beyond domestic compute capabilities, UK users benefit from **international HPC resources** via partnerships (e.g., Euro HPC and the US Department of Energy's INCITE programme) and can access compute resources via **commercial cloud providers** (Amazon AWS, Microsoft Azure, Google Cloud).

The different HPC tiers provide important and complementary functions across the full spectrum of innovation stages, known as Technology Readiness Levels[28], or TRLs. At the earliest stages (TRLs 1–3), HPC supports R&D while organisations explore new ideas and theories. At the middle stages

[23] DSIT (2024). Technology Adoption Review. Available at: https://www.gov.uk/government/calls-for-evidence/technology-adoption-review/technology-adoption-review
[24] DBT (2024). Flexible AI Upskilling Fund pilot. Available at: https://www.gov.uk/business-finance-support/flexible-ai-upskilling-fund-pilot
[25] GDS (2025). The Prime Minister's AI Exemplars programme: A portfolio approach to public sector innovation. Available at: https://gds.blog.gov.uk/2025/08/18/the-prime-ministers-ai-exemplars-programme-a-portfolio-approach-to-public-sector-innovation/
[26] DBT (2024). Invest 2035: the UK's modern industrial strategy. Available at: https://www.gov.uk/government/consultations/invest-2035-the-uks-modern-industrial-strategy/invest-2035-the-uks-modern-industrial-strategy
[27] EPCC (n.d.). History of Machines. Available at: https://www.archer2.ac.uk/community/outreach/materials/history_of_machines
[28] The Technology Readiness Level (TRL) scale is a widely used framework for assessing a technology's stage of maturity. See Heder, M. (2017). From NASA to EU: the evolution of the TRL scale in Public Sector Innovation for further information. Available at: https://innovation.cc/wp-content/uploads/2017_22_2_3_heder_nasa-to-eu-trl-scale.pdf

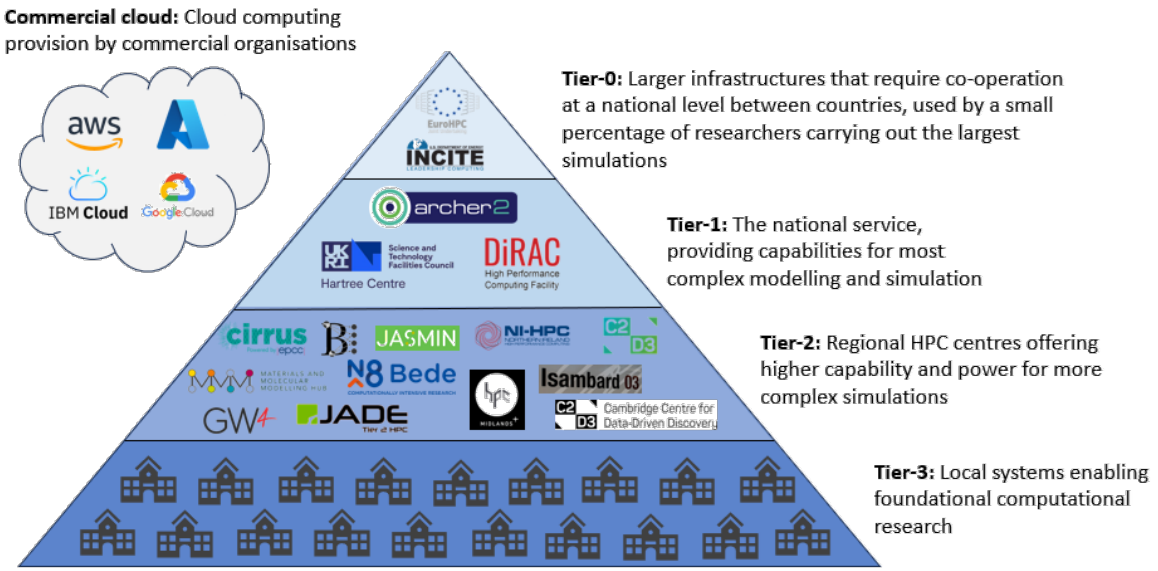
(TRLs 4–6), it supports technology prototyping and simulation, while at later stages (TRLs 7–9), it helps with product development and performance optimisation.

Small-scale theory development codes typically begin on Tier-3 or Tier-2 systems, where early models and algorithms are tested and refined. As these codes mature and are optimised, they scale up and are promoted to Tier-1 systems.

Cloud-based compute resources further extend the ecosystem, enabling users, especially in industry or those new to HPC, to access high-performance compute in a flexible, on-demand, and scalable way.

Figure 1 provides an overview of the UK's existing large-scale compute landscape.

Figure 1 Large-scale compute provision accessible by UK researchers and innovators



Source: London Economics based on the classification provided in BIST (2023). Independent Review of the Future of Compute: Final report and recommendations. Available at: https://www.gov.uk/government/publications/future-of-compute-review/the-future-of-compute-report-of-the-review-of-independent-panel-of-experts

AI computes' place in the UK compute infrastructure

Traditional High-Performance Computing (HPC) infrastructure is designed primarily for scientific research. Its primary purpose has been to run large-scale simulations and calculations that help researchers understand complex systems; for example, simulating the climate, exploring how novel materials behave under different conditions, or modelling the behaviour of complex engineering systems.

These workloads typically involve solving complex mathematical problems many times, often in parallel, using large clusters of compute processors (CPUs, see glossary). Traditional HPC systems are therefore optimised to distribute many simulation tasks across thousands of processors efficiently.

Artificial Intelligence (AI) workloads are different. Instead of running large numbers of scientific simulations, AI systems focus on training and running machine learning models (see glossary) that learn patterns from very large datasets. This process relies heavily on performing large numbers of matrix and vector calculations, which are better suited to specialised AI processors (GPUs/AI accelerators, see glossary).

As a result of these differences, compute systems designed for traditional HPC workloads do not always perform efficiently for AI tasks. AI-optimised large-scale compute systems are specifically designed for the different types of problems these systems are built to solve.

However, as highlighted in Section 2.1, UK large-scale computing systems have previously lacked a sufficient number of AI-specific processors. Investments into AI-specific compute infrastructure, such as Dawn and Isambard-AI, will fill this gap, providing AI researchers and innovators with access to the specialised public compute capabilities they need for their work.

3.2 International context and trends

AI is increasingly a strategic national priority around the world. Many governments now see advanced AI capabilities as a key component of economic competitiveness, national security, and future public service delivery. As a result, governments around the world are increasingly investing in AI infrastructure at scale.

In 2023 alone, the (Stanford) AI Index, a leading publication produced by the Stanford Institute for Human-Centered Artificial Intelligence (HAI) providing one of the most comprehensive global overviews of developments in AI, estimates that the U.S. government awarded approximately $830 million in AI-related public tenders and $4.5 billion in AI-related grants. Other nations are also increasingly investing in AI infrastructure and AI enabling technologies; for example, Canada pledged to invest $2.4 billion, China launched a $47.5 billion semiconductor fund, France committed €109 billion, India pledged $1.25 billion, and Saudi Arabia's Project Transcendence represents a $100 billion initiative.[29]

Sovereign public AI compute capabilities form a key part of governments' AI infrastructure strategies. For example, the European Commission, in 2024, pledged a €20 billion fund intended to support AI "gigafactories" (see glossary) as part of its wider "InvestAI" plans aiming to mobilising €200 billion for AI.[30] As of March 2026, the European Union has established 19 AI Factories and 13 AI Factory Antennas that offer free, customised support to SMEs and Startups.[31] More widely, among the Top 500 most powerful supercomputers in the world, more than half (a total of 255 systems) were using accelerator or co-processor technologies supporting AI workloads as of November 2025.[32]

Alongside public AI investment, private companies including the leading commercial cloud platforms (Amazon AWS, Microsoft Azure, and Google Cloud) and frontier AI developers (e.g., Open-AI, Anthropic, Alphabet) are also heavily investing in state-of-the-art AI systems. Global private AI investment recorded record highs in recent years, having reached $344.7 billion in 2025, up 127.5% from 2024 with total AI-related investment increasing approximately fortyfold since 2013.[33]

However, advanced AI compute capabilities are increasingly concentrated in certain regions. In 2025, U.S. private AI investment accounted for more than four-fifths of global AI investment ($285.9 billion of the total $344.7 billion). This is around 23 times higher than China's $12.4 billion and around 48 times the U.K.'s $5.9 billion. The gap is even more pronounced in generative AI (AI

[29] Stanford Institute for Human-Centered Artificial Intelligence (2025). Artificial Intelligence Index Report 2025. Available at: https://hai.stanford.edu/ai-index/2025-ai-index-report
[30] Innovation News Network (2024). EuroHPC unveils €1.5bn funding for Europe's first AI factories. Available at: https://www.innovationnewsnetwork.com/eurohpc-unveils-e1-5bn-funding-for-europes-first-ai-factories/53817/
[31] EuroHPC Joint Undertaking (n.d.). AI Factories. Available at: https://www.eurohpc-ju.europa.eu/ai-factories_en [Accessed 13th March 2026]
[32] Top 500 (2025). Highlights - November 2025. Available at: https://top500.org/lists/top500/2025/11/highs/
[33] Stanford Institute for Human-Centered Artificial Intelligence (2026). Artificial Intelligence Index Report 2026. Available at: https://hai.stanford.edu/ai-index/2026-ai-index-report

systems that can create content such as text, images, or code), where U.S. investment exceeded the combined total of Europe and China by $158.9 billion.[34]

Moreover, due to the vast amounts of capital required for AI accelerators and associated infrastructure (power, cooling, networking, data storage, etc), developing leading AI systems now requires very substantial financial resources.[35] As a result, access to capital has become a significant barrier to entry, with only a relatively small number of companies and organisations having the resources needed to build and operate the most advanced AI systems.[36] The high demand for AI hardware in turn puts pressure on prices and availability of these resources, making it more expensive and difficult for researchers, start-ups and smaller organisations to develop and train advanced AI systems.

Public compute investments in large-scale AI compute aim to democratise access to advanced AI computing resources, ensuring that innovative research projects across the academic, private, and public sectors can access the AI compute resources they need even without large amounts of private capital. By lowering barriers to entry, public AI compute seeks to enable a wider range of AI-related innovations and contributes to ensuring that the UK remains competitive as other nations invest heavily in large-scale AI infrastructure.

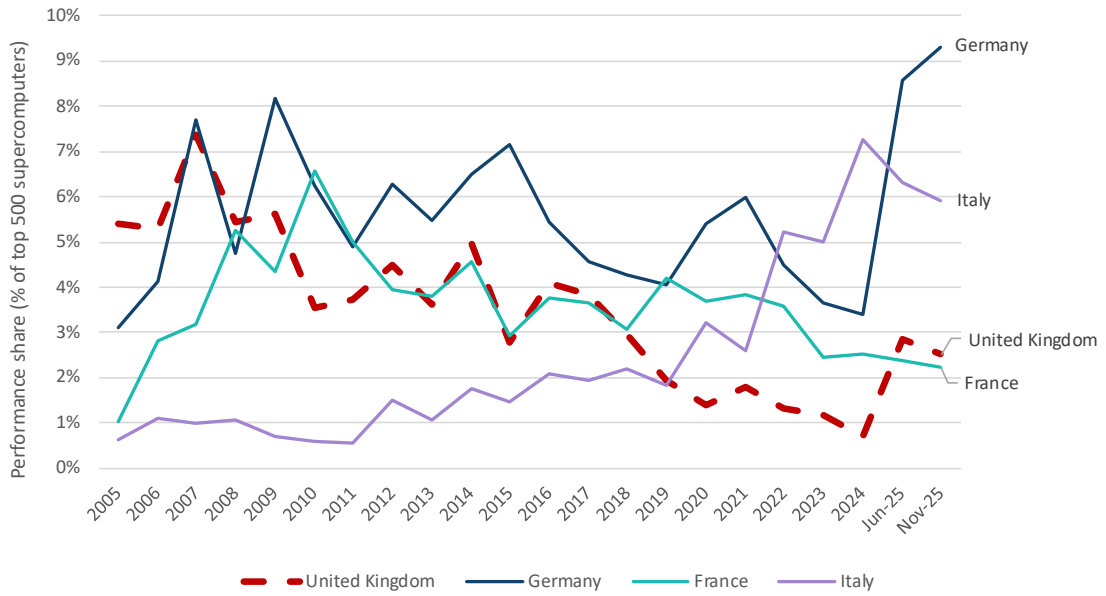
The UK's investments in AIRR and related public AI compute infrastructure also provide a much needed boost to the UK's large-scale compute provision. These investments are taking place against a backdrop of steadily falling public compute provision relative to peers over the last two decades. (see Figure 2 to Figure 4, overleaf)

In the mid-2000s, the UK's share of global large-scale compute capacity, measured by the performance of systems within the Top 500 most powerful supercomputers globally, was on par with Germany, and substantially higher than that of Italy and France.

However, as of November 2024, the latest ranking prior to the launch of the two initial AIRR systems (Dawn and Isambard-AI), the UK's most powerful supercomputers only accounted for 0.7% of global computational power among the Top 500 High-Performance Computing systems.

[34] Ibid
[35] Cottier, B et al. (2024). The rising costs of training frontier AI models. Available at: https://arxiv.org/abs/2405.21015
[36] OECD (2025) Competition in artificial intelligence infrastructure. Available at: https://www.oecd.org/en/publications/competition-in-artificial-intelligence-infrastructure_623d1874-en/full-report/component-5.html

Figure 2 UK large-scale compute performance share over time against comparator countries



Note: Shows performance share of the UK against other European comparator countries over the last decade. Performance measure = maximal LINPACK performance achieved. A country's performance share is the proportion of the total rmax GFlops (the maximum measured performance of a supercomputer running the LINPACK benchmark) of the combined Top500 computers held by that countries' computers. However, the Top500 ranking does not capture all advanced large-scale compute systems. In particular, it is well known that China is underrepresented in the ranking as it stopped submitting new entries to the Top 500 list (see TheNewStack (2024). Top500: China Opts Out of Global Supercomputer Race. Available at: https://thenewstack.io/top500-chinas-supercomputing-silence-aggravates-tech-cold-war-with-u-s/). Further, the performance measure used (the LINPACK benchmark) is unfavourable to AI-compute systems as it measures double-precision floating maths figures (FP64) while most AI compute tends to operate on lower precisions (FP34-FP8). As such, the ranking may not be truly reflective of systems' real-world comparative position.

Source: London Economics analysis based on Top500.org data

Figure 3 UK large-scale compute performance share: November 2024

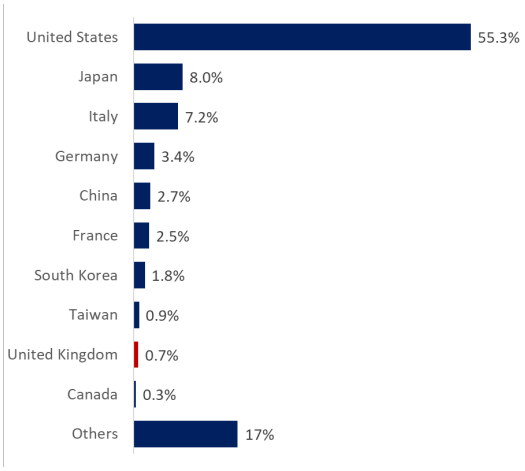
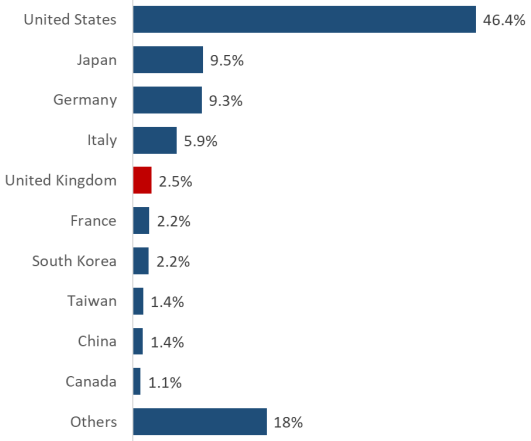


Figure 4 UK large-scale compute performance share: November 2025



Note: Shows performance share of the UK against other leading comparator countries in November 2024 (prior to launch of AIRR) and November 2025 (the latest available ranking). Performance measure = maximal LINPACK performance achieved. A country's performance share is the proportion of the total rmax GFlops (the maximum measured performance of a supercomputer running the LINPACK benchmark) of the combined Top500 computers held by that countries' computers. However, the Top500 ranking does not capture all advanced large-scale compute systems. In particular, it is well known that China is underrepresented in the ranking as it stopped submitting new entries to the Top 500 list (see TheNewStack (2024). Top500: China Opts Out of Global Supercomputer Race. Available at: https://thenewstack.io/top500-chinas-supercomputing-silence-aggravates-tech-cold-war-with-u-s/). Further, the performance measure used (the LINPACK benchmark) is unfavourable to AI-compute systems as it measures double-precision floating maths figures (FP64) while most AI compute tends to operate on lower precisions (FP34-FP8). As such, the ranking may not be truly reflective of systems' real-world comparative position.

Note 2: Excluding Dawn and Isambard-AI from the UK's November 2025 figures would entail a performance share of 0.9%.

Source: London Economics analysis based on Top500.org data

PART II: BENEFITS REALISATION

Key findings:

- **Public AI compute enables a broad range of benefits across the economy and research system:** Impacts span research outputs, innovation and adoption, partnerships, and workforce skills.
- **Benefits emerge over time through a clear pathway:** Early gains in access, data, and capability enable medium-term outcomes and longer-term economic and strategic impacts.
- **Effective delivery is critical to unlocking benefits:** Funding, allocation, coordination, communications, and user support all shape the extent to which benefits are realised.
- **Wider ecosystem factors will influence overall impact:** Adoption, access to talent and capital, and the UK's ability to commercialise and scale innovation are key external drivers.
- **A structured monitoring framework supports ongoing assessment:** Indicators and timelines have been defined to track progress and inform future evaluation.

4 Introduction and approach

This part of the report sets out the benefit realisation plan for the UK Government's investments in AI compute infrastructure. Its purpose is to identify, structure, and describe the full range of potential benefits associated with public AI compute, and to articulate how these benefits flow from the programme's activities as AIRR and related national AI compute infrastructure scale.

The benefits realisation plan is intended to provide a clear and shared understanding of the benefits that public AI compute is designed to deliver, the conditions and actions required for those benefits to materialise, and the roles and responsibilities associated with their delivery and monitoring over time. The plan provides a structured foundation that will support programme monitoring, stakeholder coordination, and future evaluation activities.

The benefit mapping presented in the plan is grounded on an aggregated Theory of Change of AIRR and investments in related national AI compute, which provides the overarching structure for describing how programme activities translate into benefits over time. The Theory of Change is used as an organising framework throughout the benefits realisation plan and wider report.

The Theory of Change and wider Benefits Realisation Plan build on and expand the Theories of Change and Monitoring and Evaluation plans developed for Dawn and Isambard-AI as part of the *Baseline evaluation of the AI Research Resource* study (published alongside this report). Work was also undertaken to review, reflect and align with BIST's own internal benefit mapping and benefit realisation work.

Further, the plan was informed by a literature review of existing studies of large-scale compute infrastructure, desk research, and London Economics' knowledge, and insights gathered from stakeholders, of the AIRR programme. Two workshops with BIST and UKRI were undertaken to refine the roles, responsibilities and expected timings of benefits. Staff from existing AIRR host sites (the University of Bristol and the University of Cambridge) also attended one of these workshops.

It is important to note that the Theory of Change and benefits realisation plan were developed in Spring 2026 during an early stage in the government's planned expansion of AIRR and investments in related public AI compute infrastructure. As such, they reflect the programme objectives at a particular point in time and may not reflect current or future direction.

5 Benefits realisation plan

5.1 Theory of change

The **Theory of Change (ToC)** developed for this benefits realisation plan sets out how programme inputs and activities are expected to translate into short-term outputs, medium-term outcomes, and longer-term impacts. At its core, the ToC reflects the role of AIRR and related investments in public AI compute infrastructure in expanding sovereign access to frontier AI compute within the UK, enabling high-impact research and innovation that would not otherwise be feasible at scale.

Figure 5 presents a **high-level logic model summarising AI compute's pathways to impact**. A brief elaboration of these pathways and enabling conditions is provided below. Following this overview, Section 5.2 provides further discussion on the roles and responsibilities of the key AI compute stakeholders in ensuring successful realisation of the benefits highlighted in the logic model. Section 5.3 then expands on external factors beyond the control of BIST and its partners that influence realisation of the identified benefits.

The logic model starts by summarising the key **programme inputs and activities** (leftmost column in Figure 5). These include strategic planning of AI compute provision; investment in the construction and expansion of AI compute facilities; procurement and refresh of advanced compute hardware; investment in novel and experimental compute systems; operational hosting and technical support; and ecosystem development and coordination through the AIRR community.

These inputs give rise to a set of short-term outputs, as illustrated in the second column of Figure 5. These are summarised across five main areas:

1) **Access to advanced AI compute is expanded** for eligible users across academia, the public sector, and industry, including SMEs and start-ups, helping address a key constraint on AI-enabled research, analysis, and innovation. This includes the establishment of sovereign AI compute capability located in the UK and governed under UK priorities.
2) **Access to data, Trusted Research Environments (TRE, see glossary), and data readiness** for AI and HPC workload **is improved**, enabling large-scale datasets to be used alongside compute and supporting the safe use of sensitive or restricted data through secure environments.
3) **AI research activity enabled by compute provision is increased**, supporting frontier research into innovative AI systems, AI-enabled research across disciplines, and dedicated AI safety research, including work undertaken by the AISI and related public-interest research activity.
4) **User capability is strengthened through staff support, training, and community building**, enabling users to make effective use of advanced compute in practice through onboarding, upskilling, and engagement activities such as workshops, conferences, and webinars.
5) **Strong national coordination and governance of AI compute** give rise to established shared governance arrangements, access rules, and standards that enable AI compute to function as a coherent national resource, supporting coordinated planning and stronger national visibility of UK AI compute capability.

Realisation of these short-term outputs depends on sustained public funding for procurement, operation, and refresh of compute infrastructure; effective programme oversight and allocation mechanisms; capable host sites with sufficient technical staffing; data governance and secure technical integration for data access; and clear leadership and accountability arrangements to coordinate delivery and maintain alignment with programme priorities.

Through these outputs, AI compute is expected to enable a range of outcomes over the medium term. These are summarised in four categories:

1) **AI compute supports increased applied AI R&D and adoption across UK organisations**, enabling businesses and public-sector bodies to build on AI compute-enabled research to develop, deploy, and scale AI-enabled products, services, and analytical tools. These effects extend beyond direct usage of the provided public AI compute infrastructure, supporting wider diffusion of AI capability across the economy and public sector.
2) **AI compute strengthens the UK's frontier research capability**, supporting the production of high-quality research outputs, AI research software, intellectual property, and spinouts. By enabling researchers to work at the limits of current methods and data, AIRR enhances research quality, visibility, and influence at the international frontier.
3) **AI compute enables stronger partnerships and collaboration across the research and innovation system**, including increased cross-sector collaboration between academia, industry, and the public sector, as well as improved coordination in the planning, delivery, and use of national AI compute capability.

4) **AI compute contributes to the development of a larger and more capable AI workforce**, supporting skills acquisition, career development, and retention of highly skilled AI, HPC, and research software engineering talent across research organisations, industry, and the public sector.

Realisation of these medium-term outcomes depends on effective uptake of AI compute and realisation of the short-term outputs; awareness, outreach, and community-building activities that extend learning beyond direct users; alignment between allocation decisions, skills provision, and strategic priorities; and complementary conditions beyond public AI compute itself, including organisational investment, data access, commercialisation pathways, and wider policy support.

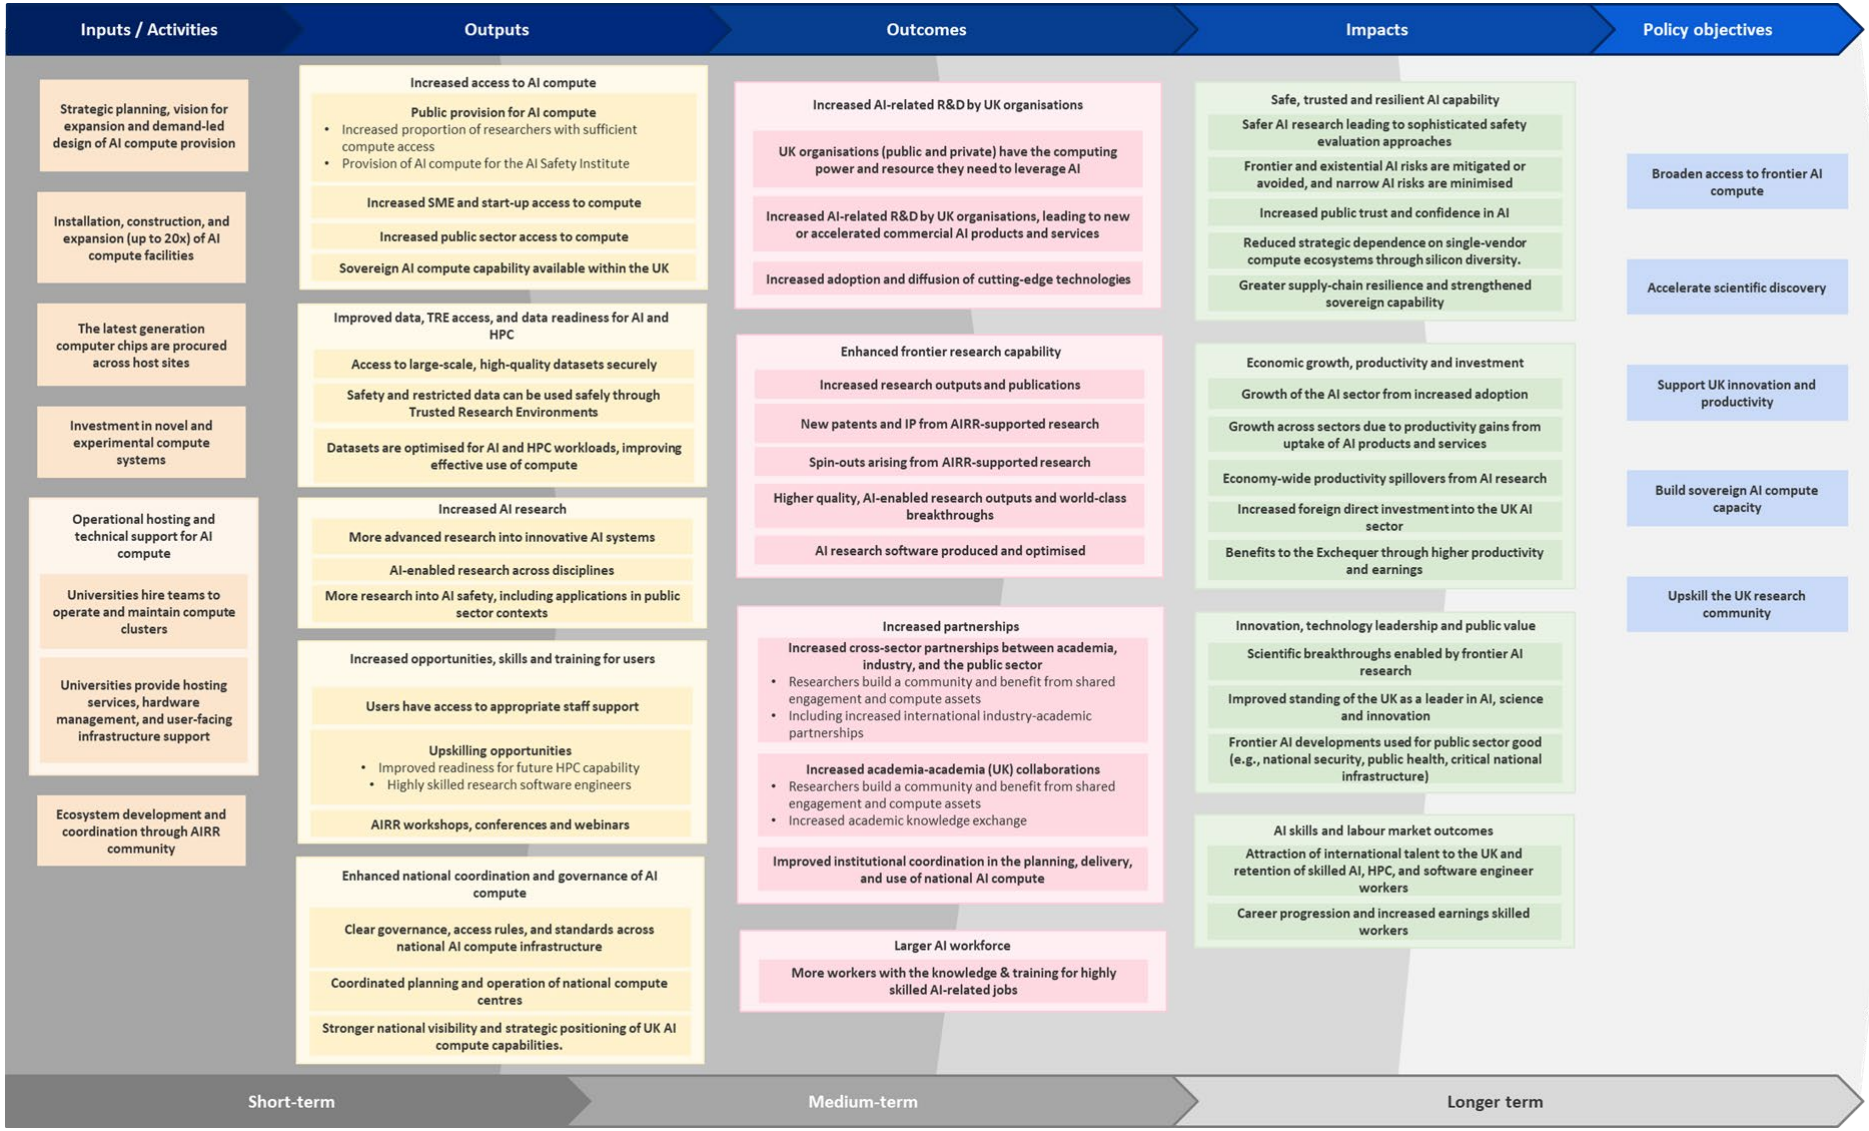
Over the longer term, these outcomes are expected to contribute to strategic impacts aligned with the Government's policy objectives. These include:

1) **The programme delivers strong sovereign AI capability and supports greater resilience of the UK's wider AI and compute ecosystem** including reduced dependence on single-vendor compute ecosystems and greater supply chain resilience.
2) **The programme enables economic growth, productivity gains and investment** by lowering barriers to advanced AI compute, supporting firm-level innovation, diffusion of AI across sectors, productivity spillovers, and increased domestic and international investment.
3) **The programme strengthens innovation, technology leadership and public value**, enabling frontier scientific breakthroughs, enhancing the UK's international standing in AI, and supporting public-sector and public-interest applications.
4) **The programme contributes to AI skills and labour-market outcomes**, helping to attract, retain and develop highly skilled AI, HPC and research software engineering talent, with longer-term benefits for career progression and earnings.

Realisation of longer-term impacts depends on a range of enabling conditions including sustained public investment in frontier AI compute; allocation models that balance frontier research, safety, public-interest and commercial use; strong coordination across BIST, delivery partners, host sites and relevant institutions such as AISI; and alignment with wider policy frameworks covering governance, skills, industrial strategy and international engagement. Beyond the programme itself, the scale and durability of impacts are shaped by effective diffusion of AI innovations, transparent communication and governance, and the ability to embed research outputs into practice, policy and markets over the medium to long term.

The precise timing for short-, medium- and long-term benefits to emerge will depend on a range of factors including, among others, the timing, scale and sequencing of infrastructure roll-out; the extent and pace of user uptake; and the rate at which follow-on innovations are developed and adopted. Given the uncertainty surrounding these factors, precise timelines are not specified in the Theory of Change. However, indicative timeframes for when different benefits are expected to emerge are provided alongside the set of indicators set-out in Section 5.4.2.

Figure 5 Investments in public AI compute - logic model



Note: The Theory of Change was developed in Spring 2026 during an early stage in the government's planned expansion of AIRR and investments in related public AI compute infrastructure. It reflects the programme objectives at a particular point in time and may not reflect current or future direction. ***Source: London Economics***

5.2 Roles and responsibilities

The realisation of the benefits from the investment involves a set of stakeholders with specific roles and responsibilities. This section sets out the roles, responsibilities and actions required to ensure that the benefit realisation plan operates effectively. The areas set out below were identified as critical in discussions between BIST, UKRI and the existing AIRR centres at the University of Cambridge and the University of Bristol. Collectively they enable the realisation of research, economic, and skills-related benefits.

5.2.1 Funding

BIST has overarching responsibility for the continued funding and strategic oversight of the AIRR programme and investments in related public AI compute infrastructure. This role is critical to enabling the realisation of programme benefits, ensuring that the UK has access to cutting-edge AI hardware at scale, develops robust software ecosystems, and builds national capability in AI research, infrastructure, and operations. By maintaining a **long-term funding commitment,** BIST provides the certainty required for infrastructure planning, resource allocation, and coordinated development across multiple sites and user communities.

BIST will invest in and oversee a range of compute infrastructure, including:

- On-premises supercomputers deployed at national supercomputing centres (NSCs);
- Cloud-based compute platforms to provide flexible access;
- A federation of compute access points (AIRR Portal) to enable seamless usage across multiple centres; and
- Experimental systems and testbed infrastructure to explore emerging AI workloads and novel architectures.

BIST also has responsibility for planning the composition of these infrastructures and architecture types, as well as setting the rollout timelines and refresh cycles. Continued investment is required to periodically update hardware and prevent it from becoming obsolete.

5.2.2 Allocation of resource

To allocate funding for maximum benefit, a clear and evidence-based understanding of demand is essential. BIST will lead on a user demand study to understand who is using the resource, how it is being applied, and which additional user groups may have viable use cases but are not yet accessing the system (already in progress).

However, developments in AI-enabled science may expand or shift the profile of users over time. BIST should monitor such changes to ensure that compute capabilities continue to meet emerging and future demand.

Using this evidence, BIST will develop an allocation strategy that balances the needs of multiple user groups (public sector, industry and academic) whose engagement underpins different benefit streams.

The host sites are responsible for capacity management. The sites will provide BIST and delivery partners (UKRI) with data on utilisation, availability and user constraints. This information is critical for informing future allocation cycles and avoiding mismatches between demand and supply.

BIST has responsibility for evaluating applications submitted through the R&D track, while delivery partners such as UKRI are responsible for evaluating all other applications to use the AIRR.

5.2.3 Strategic alignment of compute investments

BIST will also invest in the broader AI ecosystem to ensure that hardware delivers maximal impact. This includes software programmes, people and skills development initiatives, and dataset curation and storage. These responsibilities support the integration of novel hardware into research workflows and continued funding will facilitate the development of specialist skills within the workforce responsible for running and maintaining the infrastructure.

BIST also holds responsibility for aligning the AIRR programme with other national investments. This oversight includes the planning and design of national supercomputing centres, as well as strategic coordination with delivery partners such as UKRI and the host sites.

Moreover, the UK's compute investments take place against a background of vastly accelerated public and private compute investments globally. BIST's investment and service design choices should consider the broader international ecosystem so that they maximise national benefits while exploiting synergies from other parties' investments.

5.2.4 External communications

External communications are another critical enabler because many intended benefits depend on attracting the right users to the system.

Building on insights from the user demand study (discussed in Section 5.2.2), BIST will hold overall strategic responsibility for the external positioning and communications of the AIRR programme and related AI compute investments.

BIST will lead on the development of an overarching external communications and partnerships plan. This will include targeted outreach to industry, with a particular emphasis on engaging currently underrepresented groups such as small and medium-sized enterprises (SMEs) and start-ups. Such engagement is essential for unlocking benefit streams that depend on industry uptake and collaboration. Delivery partners and the host sites will support in the delivery of the external communications plan by promoting the AI compute resources within their existing networks; and by providing inputs such as research outcomes, technical achievements, and downstream impacts through surveys and case studies to support the communication strategy.

5.2.5 User support and AIRR community

Once the user groups are accessing the system (through effective external communications and allocation of resource), user support and community-building are essential to ensure that benefits are realised in practice. Users of public AI compute capabilities must be able to effectively utilise their allocated resource to enable subsequent outcomes, such as increased AI-related R&D, enhanced frontier research capabilities, stronger partnerships, and the development of a skilled AI workforce.

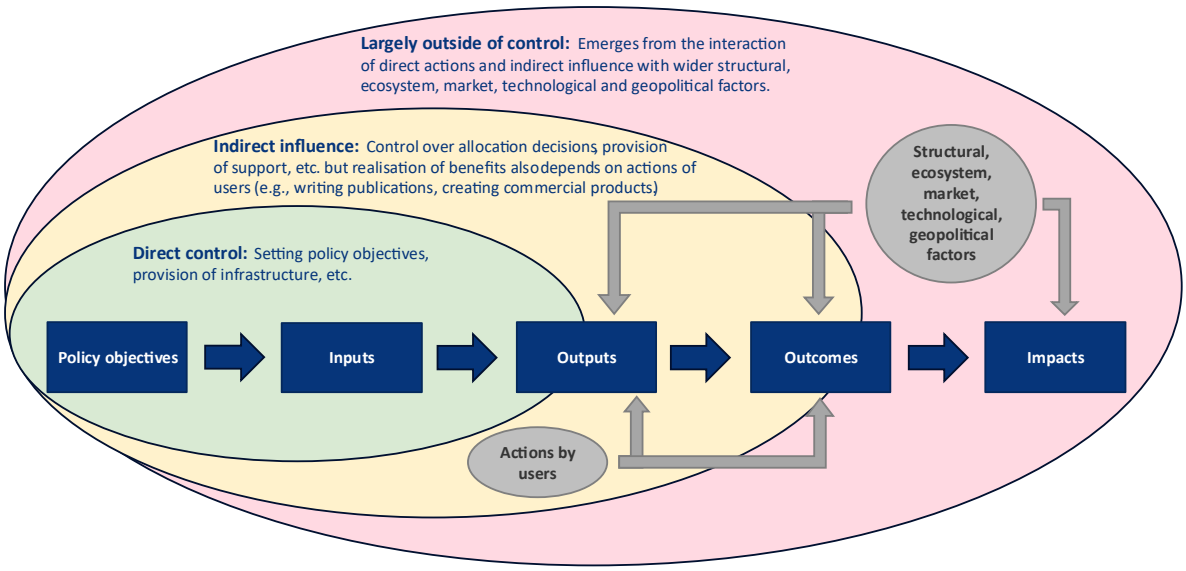
Primary responsibility in this area sits with the host sites, supported by BIST and its delivery partners where appropriate. The host sites are responsible for providing technical support, training resources, and user-facing guidance to enable researchers and innovators to use their system effectively for their intended purposes. This includes hiring and retaining specialist staff to ensure continuity of skills, as well as training and supporting users.

Community-building activities are also important, including events such as the Isambard summit, as well joint events for users of both the Dawn and Isambard-AI systems. The host sites will work collaboratively to agree on proportionate and effective approaches to community development. Rather than creating entirely new forums, this may involve building on existing networks and communities, to maximise reach and relevance while maintaining interest.

5.3 External factors influencing benefits realisation

Alongside the direct actions and responsibilities described in Section 5.2, realisation of benefits depends on a range of wider enabling conditions and ecosystem factors that sit largely outside the direct control of BIST, its delivery partners, and the host sites.

Figure 6 AIRR programme sphere of influence



Source: London Economics

By ensuring factors within the AIRR programme's control and sphere of influence, the programme can contribute to ensuring researchers, businesses and public-sector users are set-up for success. However, achieving longer term impacts such as accelerated economic growth, a flourishing UK AI ecosystem, and strong AI talent base in the UK also depend on a range of factors outside of the AIRR programme's control (but which BIST may nevertheless be able to influence). These include:

- **Structural factors** such as, among others, the UK's wider regulatory and legal frameworks including IP, data protection, and AI liability legislation; the UK's AI education and skills pipeline, including skills training by UK universities and skilled immigration and visa policies for researchers and innovation talent; the research and business funding and support landscape, such as the availability and scale of research funding, the availability and design of wider Government commercialisation and scale-up support mechanisms such as grants, tax credits, patient capital, accelerator programmes, etc.
- **Ecosystem factors** such as a thriving AI research community; the availability and effectiveness of networks connecting researchers, innovators, funders, and users; the risk appetite and innovation propensity among UK entrepreneurs; availability of commercialisation and scale-up support and expertise; etc.
- **Market factors** such as end-user demand and willingness to adopt AI solutions; access to private capital and risk appetite among funders (VC, growth equity, etc.); speed of technological diffusion; international competition with UK researchers and AI businesses

OFFICIAL

as well as global competition for AI talent; changes in hardware and energy prices; business models and supply by commercial cloud providers; etc.

- **Geopolitical factors** such as international competition and collaboration opportunities; access to other systems such as European AI factories; international standards-setting and regulatory divergence; dependence on overseas suppliers; competition between major AI-producing economies and the UK's position internationally; etc.
- **Technological factors** such as the pace and direction of global technological change in AI models and frontier AI research; developments in the AI hardware market; developments in dominant architectures or software stacks; etc.

In particular, while the UK has well-known strengths in scientific research and innovation as well as a vibrant tech and start-up scene, the UK has historically struggled with turning cutting-edge research and innovations, and early-stage spinouts into successful commercial ventures, with particular challenges in scaling start-ups and retaining high-growth firms domestically. The UK's scale-up challenges are well-documented, including in the recently published 2025 House of Lords report on AI and creative technology scaleups; they include:[37]

1) **Access to scale-up capital**, with UK scale-ups finding the capital required for growth difficult to secure domestically.
2) **Culture and risk appetites**, including a risk averse approach to investment among UK investors the lack of an ingrained, national entrepreneurial culture in the UK.
3) **Market size and conditions**, with the UK's market and exist opportunities being smaller and less dynamic, often leading to start-ups to sell up or relocate elsewhere for opportunities.
4) **Regulatory framework**, with the UK's regulatory environment often seen as more complex, costly, and onerous compared to international competitors.
5) **Access to talent**, with access to world-class AI and technical talent frequently cited as a barrier by UK businesses, as well as skills and experience gaps in managing later-stage AI companies of a global scale.
6) **Infrastructure** including access to compute.

5.4 Ongoing monitoring and evaluation

5.4.1 The M&E function

Monitoring and evaluation (M&E) is essential to ensure that the AIRR programme and related AI compute investments remain aligned with their objectives and are able to adapt as usage patterns, technologies, and user needs evolve. Given the scale and novelty of the infrastructure, monitoring arrangements must be proportionate, focused on learning as well as accountability, and sensitive to the practical constraints faced by users and host sites.

The host sites are central to the monitoring of operational performance, as they have direct visibility of system usage, allocation patterns, and utilisation levels. This includes collecting and reporting quantitative data such as node hours consumed, queue times, and system availability, and providing regular feedback to BIST and its delivery partners on how capacity is experienced in practice. These

[37] House of Lords Communications and Digital Committee (2025). *AI and creative technology scaleups: less talk, more action*. 2nd Report of Session 2024–25. Available at: https://publications.parliament.uk/pa/ld5901/ldselect/ldcomm/71/71.pdf

OFFICIAL

data are critical for informing allocation decisions, managing demand, and identifying emerging pressure points or under-utilisation.

It is important that monitoring activities are undertaken collaboratively across host sites, and in coordination with BIST and its delivery partners, to ensure consistency and avoid mismatches in the type, frequency, or interpretation of data collected. A shared and aligned approach to monitoring will support comparability across sites and enable programme-level insights to be drawn more effectively.

As noted in Section 5.2.4, BIST through its delivery partners (e.g. UKRI) and the host sites will also introduce end-of-project reporting mechanisms to capture information on outputs and outcomes in a structured and proportionate way. These mechanisms may include brief user surveys or lightweight reports focused on research achievements, technical learning, publications, partnerships, skills development, or follow-on funding, providing a more complete picture of the outcomes enabled by AI compute access beyond operational metrics alone.

5.4.2 Key performance indicators

To support the ongoing monitoring of benefits and to track whether the programme remains on course to achieve its objectives, a set of indicators have been developed. A summary of these indicators, grouped by the overarching benefit groupings set-out in the logic model, is provided in the tables overleaf. Expected timelines for benefits to emerge are also provided alongside each benefit grouping.

Table 3 Indicators to measure short-term outputs

Grouping	Output	Indicators	Expected timeline for outputs to emerge
Increased access to AI compute	Increased public provision of AI compute	Volume of compute provisioned through AIRR	Benefits are expected to arise in the short term as compute capacity is provisioned, and to increase over time as utilisation grows. Initial capacity may become available within a few months, with full operational capacity typically reached within 12–18 months, depending on delivery model. Cloud-based approaches may enable faster deployment.
	Increased SME and start-up access to compute	Volume of AIRR compute available for SME and start-ups; number and proportion of SME/start-up users/projects by industry sector.	
	Increased public sector access to AI compute	Volume of AI compute available for public sector research and analytical use cases; number of public sector organisations and projects accessing AIRR compute.	
	Sovereign AI compute capability available within the UK	Installed AI compute capacity located in the UK under UK government governance; volume of AI compute capacity within AIRR that can be prioritised for UK sovereign and public- interest use cases; number of Sovereign-AI organisations / projects accessing AIRR compute.	
Improved data, trusted research environments access, and data readiness for AI and HPC	Access to large-scale, high-quality datasets	Number and scale of datasets accessible alongside AIRR; number of projects using large-scale datasets	Benefits are expected to arise in the short to medium term as datasets and TREs are connected to AIRR compute, and to increase over time as data pipelines, governance, and optimisation mature. Initial datasets may become available within a few months, while more comprehensive provision may take several months to years.
	Sensitive and restricted data can be used safely through Trusted Research Environments (TREs)	Number of AIRR projects using TREs; volume of compute used within TRE-enabled projects	
	Datasets optimised for AI and HPC workloads	Evidence of datasets prepared or optimised for AI and HPC use; user feedback on data readiness	
Increased AI research	More advanced research into innovative AI systems	Number of projects introducing novel AI concepts, software, or new AI products and services	Benefits are expected to arise primarily in the medium term, with initial research activity beginning as compute becomes available and increasing over time as projects progress. Timelines depend on build and uptake: activity may begin within a few months under cloud delivery, or around 12–18 months for new infrastructure, followed by a period of uptake.
	AI-enabled research across disciplines	GPU hours consumed on AIRR by research organisations; number of research projects and active researchers using AIRR compute for AI-enabled research; distribution of projects by discipline	
	More research into AI safety	Number of AISI projects and publications; policy inputs from AISI research (assessed by BIST)	
Increased opportunities, skills and	Users have access to appropriate staff support	Number of dedicated AIRR support staff by host site; support response times and resolution rates; user satisfaction with support services	Benefits are expected to arise in the short term as users access support and training, and to continue into the medium term as skills deepen and the user community

Grouping	Output	Indicators	Expected timeline for outputs to emerge
training for users	Upskilling opportunities provided	Number of trained engineers and support staff at host sites; AIRR users participating in training activities; user feedback on skills development	matures. Initial benefits may be realised within a few months, depending on compute availability, uptake, and training delivery.
	AIRR workshops, conferences and webinars hosted	Number of events and attendees; user feedback on the role of AIRR in skills development and community building	
Enhanced national coordination and governance of AI compute	Clear governance, access rules, and standards across national AI compute infrastructure	Existence and implementation of a common governance framework; consistency of access rules across AIRR sites; stakeholder feedback on clarity and predictability of access decisions	Benefits are expected to arise in the short to medium term as governance and coordination mechanisms are established, and to increase over time as operations become more aligned and integrated. Initial benefits may be realised within a few months, while more mature benefits may take months to several years to materialise.
	Well-coordinated planning and operation of national compute centres	Evidence of joint planning activities; alignment of operational plans across centres; reduced duplication or conflicting investment decisions; stakeholder feedback on operational coordination	
	Stronger national visibility and strategic positioning of UK AI compute capabilities	References to AIRR in national strategy and policy documents; engagement with international partners and initiatives; stakeholder perceptions of the UK's AI compute capability and coherence	

Note: These indicators were developed in Spring 2026 during an early stage in the government's planned expansion of AIRR and investments in related public AI compute infrastructure. They reflect the programme objectives at a particular point in time and may not reflect current or future direction.

Source: London Economics

Table 4 Indicators to measure medium-term outcomes

Grouping	Outcome	Indicators	Expected timeline for outcomes to emerge
Increased applied AI R&D by UK organisations	UK organisations have the computing power and resources they need to leverage AI	User feedback on adequacy of compute access; proportion of users reporting sufficient access to AI compute to meet project requirements	Benefits are expected to arise over the medium term as applied AI projects progress, with wider productivity and organisational impacts emerging over the medium to long term. Activity may begin within a few weeks of compute access, with projects typically lasting a few months, though more complex efforts may take several years.
	Increased applied AI R&D by UK organisations, leading to new or accelerated AI-enabled products, services, and operational tools	Number of organisations undertaking applied AI R&D using AIRR, by sector and organisation type; compute time on AIRR utilised by user type; evidence of new or accelerated AI-enabled outputs	
	Increased adoption and diffusion of cutting-edge technologies	Measures of AI adoption among organisations	
Enhanced frontier research capability	Increased research outputs and publications	Number of research outputs and publications based on AIRR-supported research	Benefits are expected to arise over the medium to longer term. Early outputs, such as working papers, publications, and software, may appear earlier, while impacts related to research quality, intellectual property, and spinouts typically materialise over longer time horizons.
	New patents and intellectual property from AIRR-supported research	Number of patents and disclosures linked to AIRR-supported research	
	Spinouts arising from AIRR-supported research	Number of spinouts from AIRR-supported research; turnover, employment, or investment information where available	
	Higher-quality, AI-enabled research outputs and world-class breakthroughs	Citation-based metrics; peer review assessments; high-impact publications and awards; user feedback on the role of AIRR in improving research quality	
	AI research software produced and optimised	Number of AI research software tools produced; reuse, uptake, and citations	
Increased partnerships	Increased cross-sector partnerships between academia, industry, and the public sector	Number of cross-sector collaborative projects; mapping of collaboration networks; user feedback on the role of AIRR in enabling partnerships	Benefits are expected to arise over the medium term as collaborative projects and networks form, and to deepen over time as relationships mature. More sustained partnerships and coordination benefits are likely to emerge

Grouping	Outcome	Indicators	Expected timeline for outcomes to emerge
	Increased academia-academia collaboration within the UK	Number of academic collaborations; mapping of collaboration networks; user feedback on the role of AIRR in enabling collaboration	over the medium to longer term, with clearer patterns of collaboration visible within the first few years.
	Improved institutional coordination in the planning, delivery, and use of national AI compute	Feedback from stakeholder engagement; evidence of coordinated planning and delivery processes	
Larger AI workforce	More workers with the knowledge and training for highly skilled AI-related jobs	Workers with the knowledge and training for highly skilled AI-related jobs	Benefits from the development of a larger AI workforce are expected to arise over the medium to long term. Initial skills gains may occur as users and staff engage with AIRR, while wider workforce effects, including increased supply and retention of AI talent, are likely to materialise over longer time horizons. Meaningful employment impacts are expected to emerge after a few years, as organisations scale activity and translate research into growth.

Note: These indicators were developed in Spring 2026 during an early stage in the government's planned expansion of AIRR and investments in related public AI compute infrastructure. They reflect the programme objectives at a particular point in time and may not reflect current or future direction.
Source: London Economics

Table 5 Indicators to measure longer-term impacts

Grouping	Impact	Indicators	Expected timeline for impacts to emerge
Resilient AI capabilities	Reduced strategic dependence on single-vendor compute ecosystems	Diversity of compute architectures within AIRR Evidence of reduced vendor concentration	Increased resilience benefits primarily materialise over the long term, as a diverse vendor ecosystem is established and sufficient, and diverse, sovereign AI compute has been provided. However, resilience should be seen as an ongoing consideration in planning and decision making throughout the AIRR programme
	Greater supply-chain resilience and strengthened sovereign capability	Share of compute capability that is sovereign; continuity and reliability of compute availability; evidence of supply chain diversity	
Economic growth, productivity gains and investment	Growth of the AI sector from increased adoption	Growth in AI sector employment and revenues, for example from AI sector studies	Impacts on economic growth, productivity, and investment are expected to materialise primarily over the medium to long term. Early effects may be observed in firm growth and investment decisions, while broader productivity gains, spillovers across sectors, and fiscal benefits are likely to emerge gradually as AI-enabled products and services are adopted across the economy. These impacts depend on the progression from research to commercialisation and adoption, and may take several years to materialise, particularly for wider spillover effects. Frontier Economics (2024) in their well-known estimation of Returns to UK public R&D use an average lag between university research and spillover productivity impacts taking place of six years, though note that impacts could be greater if a longer lag was chosen.[38]
	Growth across sectors due to productivity gains from uptake of AI products and services	Productivity and GDP indicators Modelling-based estimates of AI-enabled impacts	
	Economy-wide productivity spillovers from AI research	Estimated spillover[39] multipliers from AI-related R&D Qualitative case studies	
	Increased foreign direct investment into the UK AI sector	FDI inflows into the UK AI sector	
	Benefits to the Exchequer through higher productivity and earnings	Estimated tax receipts and fiscal impacts from long-term evaluations	
Innovation, technology	Scientific breakthroughs enabled by frontier AI research	High-impact publications and citations; Novel methods or discoveries attributed to AI; Patents and spinouts; Major research prizes or recognition	Impacts are expected to emerge over the medium to long term. Early effects may be seen as compute capacity enables initial research and applications, while broader impacts on

[38] Frontier Economics (2024). Boosting productivity: Returns to UK public R&D. Available at: https://www.frontier-economics.com/uk/en/news-and-insights/news/news-article-i21101-boosting-productivity-returns-to-uk-public-rd/

[39] Indirect productivity and innovation benefits arising beyond the original research performers.

Grouping	Impact	Indicators	Expected timeline for impacts to emerge
leadership and public value	Improved standing of the UK as a leader in AI	International benchmarks and rankings, such as global AI indices and peer recognition	international standing and strategic capability are likely to materialise over longer time horizons.
	Frontier AI developments used for public sector good	Research outputs and applications with clear public sector or public-interest relevance	
AI skills and labour-market outcomes	Attraction of international talent to the UK and retention of skilled AI, HPC, and software engineer workers	Size and composition of skilled AI workforce; Labour-market statistics; Number of inward researchers and specialists using AIRR; Survey/interview evidence on the role of AIRR in location decisions	Impacts are expected to emerge over the medium to long term. Early effects may be seen in talent attraction and skills development, while broader labour-market impacts, including retention, progression, and earnings, are likely to materialise over longer time horizons. While early labour market effects may materialise over a few years, more meaningful ecosystem impacts - particularly the expansion of AI skills across the wider UK labour market - will take many years or even decades to materialise.
	Career progression and increased earnings of skilled workers	Career tracking and earnings proxies from surveys or longitudinal data	
	Skilled AI workforce in the UK	Number of skilled workers with AI skills in the UK	

Note: These indicators were developed in Spring 2026 during an early stage in the government's planned expansion of AIRR and investments in related public AI compute infrastructure. They reflect the programme objectives at a particular point in time and may not reflect current or future direction.
Source: London Economics

6 Summary and key takeaways

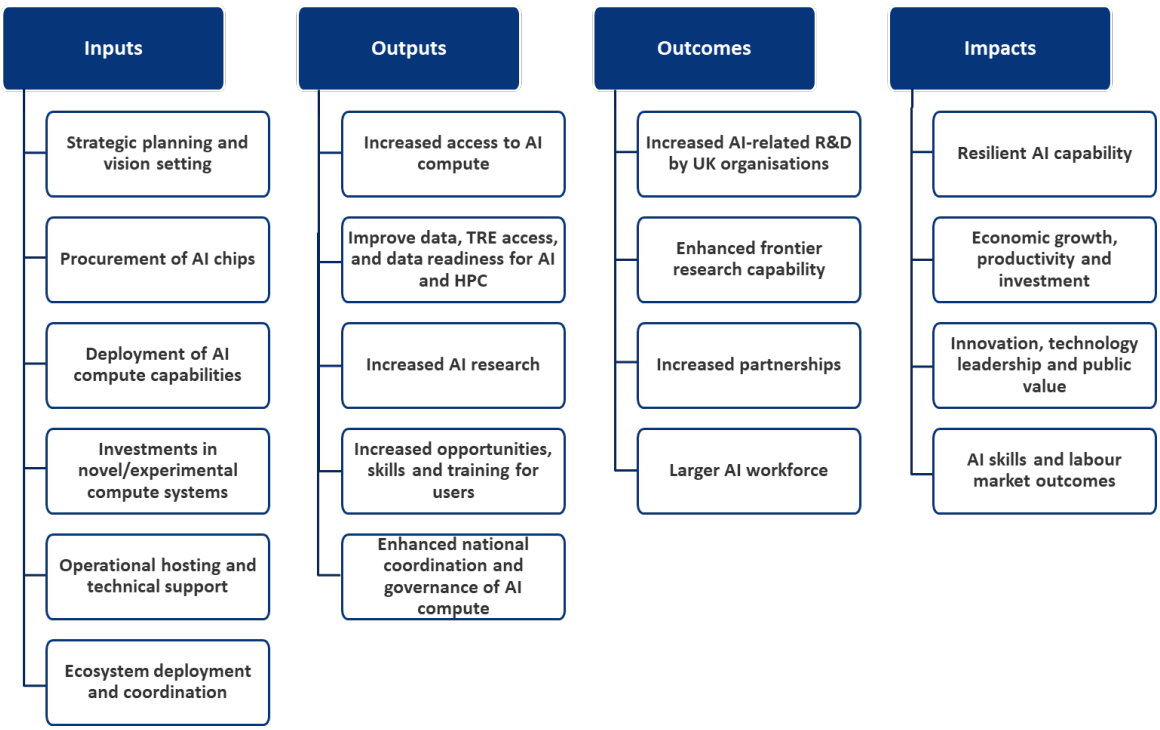
This part of the report set out the Benefit Realisation Plan for the AI Research Resource (AIRR) programme, including its 20x expansion and investments in associated AI compute infrastructure. The benefits realisation plan identifies, structures, and describes the full range of potential benefits associated with AIRR, and articulates how these benefits flow from the programme's activities as AIRR scales. It is intended to provide a clear and shared understanding of the benefits that AIRR is designed to deliver, the conditions and actions required for those benefits to materialise, and the roles and responsibilities associated with their delivery and monitoring over time.

Theory of Change

The benefit mapping presented in this report is grounded on an aggregated Theory of Change (ToC) of AI compute provision, which provides the overarching structure for describing how programme activities translate into benefits over time. This Theory of Change is aligned with internal work on AI compute benefits by the BIST policy team but is expressed at a more aggregated level to support systematic benefit identification, monitoring and future evaluation.

The ToC sets out the programme's pathway to change, linking activities and inputs to short-term outputs, intermediate outcomes, and long-term impacts in line with its overarching policy objectives: broaden access to frontier AI compute; accelerate scientific discovery; support UK innovation and productivity; build sovereign AI compute capability; and upskill the UK research community. Specifically, the ToC identifies the following broad groups across AIRR's benefits realisation pathway; the full logic model is provided in Section 2.

Figure 7 Overview of AI compute benefits



Note: The benefit mapping was undertaken in Spring 2026 during an early stage in the government's planned expansion of AIRR and investments in related public AI compute infrastructure. It reflects the programme objectives at a particular point in time and may not reflect current or future direction. ***Source: London Economics***

OFFICIAL

Programme-level actions and responsibilities

In a workshop with BIST and UKRI, undertaken in March 2026, a set of key cross-cutting actions and responsibilities required to collectively enable the realisation of the AIRR programmes' benefits were identified. These actions and responsibilities bring together themes that intersect multiple benefit streams and were identified as critical in discussions between BIST, UKRI, the University of Cambridge and the University of Bristol. Identified programme level actions and responsibilities are briefly summarised below and further elaborated on in Section 5.2. It should be noted that the actions and responsibilities listed here reflects the programme objectives at the time the research was undertaken, in Spring 2026, and may not reflect current or future direction.

BIST

- Provide continued funding for the AIRR programme to ensure long-term stability and strategic oversight.
- Invest in on-premises supercomputers, cloud compute, federated compute access, experimental systems, testbed infrastructure and wider ecosystem development.
- Set the mix of infrastructures and architecture types, rollout timelines and refresh cycles.
- Co-ordinate investment into other national supercomputing centres (NSCs).
- Undertake a user demand study to understand current and potential users/use cases and monitor changes where appropriate.
- Develop an allocation strategy that balances the needs of multiple user groups.
- Lead the strategic positioning and external communications of the AIRR programme.
- Dedicate resource to developing an external communications and partnerships plan.
- Lead targeted industry outreach, including engagement with SMEs and start-ups.
- Lead international communications to position the UK as a global leader in AI research and infrastructure.
- Determine which user behaviours should be incentivised to support programme objectives.

Delivery partner (e.g. UKRI)

- Evaluate applications submitted through the rapid access and gateway tracks.
- Support BSIT in the delivery of the external communications and partnerships plan.
- Assist host sites and BIST in user support and community-building activities.
- Work with the host sites to introduce end-of-project user surveys to capture research outcomes, technical achievements, and downstream impacts.
- Work with BIST and host sites to align monitoring and evaluation approaches.

Host sites

- Provide BIST and UKRI with timely, accurate data on cluster utilisation, capacity, and operational constraints.
- Support capacity management by reporting on user feedback, allocation and utilisation.
- Support users to make effective use of their allocated resource.
- Provide technical support, training resources, and user-facing guidance to AIRR users.
- Hire and retain specialist staff to ensure skills retention and effective system operation.
- Identify, develop, and promote case studies and success stories where possible.
- Work collaboratively to deliver community-building activities and events.
- Coordinate with each other, UKRI, and BIST to ensure consistent and proportionate monitoring and evaluation.

OFFICIAL

External factors influencing benefits realisation

Alongside the direct actions and responsibilities of BIST, delivery partners (e.g., UKRI), and the host sites, realisation of benefits depend on a range of wider enabling conditions and ecosystem factors that sit largely outside the direct control of the AIRR programme (though some of which BIST and the wider UK Government can influence). These include:

- **Structural factors** including the regulatory and legal framework (e.g., IP, data protection and AI liability), the UK's AI skills and talent pipeline (including university training and immigration/visa policy), and the research, funding and commercialisation landscape (e.g. grants, tax credits, patient capital and accelerator support).
- **Ecosystem factors** such as a thriving AI research community; effective networks connecting researchers, innovators, funders, and users; the risk appetite and innovation propensity among UK entrepreneurs; effectiveness of commercialisation and scale-up support.
- **Market factors** such as end-user demand and willingness to adopt AI solutions; access to private capital and risk appetite among funders; speed of technological diffusion; international competition; changes in hardware and energy prices; business models and supply by commercial cloud providers.
- **Geopolitical factors** such as international competition and collaboration opportunities; access to other compute systems; international standards-setting and regulatory divergence; dependence on overseas suppliers; competition between major AI-producing economies and the UK's position internationally.
- **Technological factors** such as the pace/direction of change in AI globally; developments in the AI hardware market and dominant architectures or software stacks.

The degree to which the UK's historical difficulties in scaling start-ups and retaining high-growth firms domestically can be addressed also presents a noteworthy external factor influencing realisation of some of the longer-term and strategic impacts of the AIRR programme.

Monitoring of benefits

To support the ongoing monitoring of benefits and to track whether the programme remains on course to achieve its objectives, a set of indicators have been developed for each group of benefits set-out in the logic model. Expected timelines for benefits to emerge are also provided alongside each benefit grouping.

PART III: ECONOMIC VALUATION

Key findings:

- **The economic valuation suggests strong value for money across a range of scenarios:** The analysis estimates a Return on Investment in the range of **£4.2 to £20.2 per £1 invested**, with a central estimate of **£6.0 per £1 invested**.
- **The central estimate is conservative:** It is based on cautious assumptions and excludes a number of strategic, societal, and capability-building benefits that could not be robustly monetised.
- **Wider economic spillovers account for most of the estimated value:** The largest share of monetised benefits comes from wider spillovers linked to research and innovation, alongside avoided compute costs, private user benefits, public sector efficiency gains, and skills benefits.
- **There is meaningful upside beyond the central case:** The sensitivity analysis indicates that higher-return scenarios are plausible, with a median estimated benefit **(£9.1 : £1)** above the central case and a pronounced upper tail of potential outcomes.
- **The results are broadly in line with wider evidence:** The estimated returns are consistent with the range reported in other evaluations of large-scale compute and research infrastructure **(£2 : £:1 to £20 : £1)** in the UK and internationally.

OFFICIAL

7 Introduction and approach

This part of the report presents the economic valuation of the Return on Investment (RoI) associated with the UK Government's investments into Isambard-AI and Dawn[40], the 20x expansion of the AI Research Resource (AIRR), and related national AI compute infrastructure. The analysis assessed the potential benefits to the UK across the academic, public, and private sectors that AI compute is likely to deliver over their operational life over the next 10 years under the central scenario. Estimated benefits are compared to the expected costs of delivery of the AI compute services to evaluate the return on these investments for the UK.

The analysis built on the Theory of Change and on previous evaluations of large-scale compute infrastructure to develop a forward-looking benefits-costs model, but refines the analytical approaches and assumptions used, where this was possible, to better reflect the evolving landscape of AI compute in the UK. The modelling was conducted by a team of independent professional economists. Estimates of economic impacts are based on best practice and best judgment to calculate the most robust and fair estimates. All monetary estimates are presented in 2025/26 real discounted prices, using the HM Treasury Green Book discount rate of 3.5%.[41]

Annex 1 describes in detail the methodological approach including the specific evaluation approaches used for each monetised benefit stream and the assumptions made, as well as the caveats and limitations of the analysis, where appropriate. A high-level summary of the methods used and caveats and limitations are provided in the next section (Section 8).

Evidence gathering combined quantitative and qualitative methods, including a user survey of existing users of the two AIRR compute facilities already available (Dawn and Isambard AI), follow-up stakeholder interviews with users and engagement with BIST, UKRI and two existing host sites (the University of Bristol and the University of Cambridge), and desk research and literature reviews, to inform the evaluation of benefits.

Given the forward-looking nature of the AI compute investments and the uncertainty around both the magnitude of potential benefits and associated timelines for benefits to materialise, a modelling framework based on Monte Carlo simulations was developed. This framework incorporates key sources of uncertainty to generate a range of possible benefit outcomes, along with their associated likelihoods, enabling a probabilistic assessment of return on investment.

In simple terms, this approach runs many different scenarios using a range of possible assumptions about hardware costs, benefits, and timing. This produces a spread of possible outcomes, showing not just what might happen, but also how likely each outcome is, thereby providing a better understanding of the range within which the return on investment is likely to fall, and the likelihood of different outcomes within that range.

[40] Note, that Dawn was already being procured and was subsequently brought into the AIRR programme with AIRR only funding the operational expenditure of Dawn. As such, in line with HMT Green Book Guidance on additionality, only a share of benefits expected to emerge from Dawn was attributed to AIRR. This share was assumed to be equivalent to the (estimated) share of operational costs relative to total costs in the central case. Sensitivity analysis allowed for the attributable share to vary. While attribution was allowed to vary down, given the relatively low magnitude of operational costs vs. capital costs for compute investments, the central assumption was judged to be already on the conservative side. Therefore, the main direction of variation in attribution under the sensitivity analysis was upwards in order to allow for additional impacts resulting from the decision to bring Dawn into AIRR. These additional benefits could emerge, for example, from better policy integration and alignment with national priority areas, improved utilisation of the service, easier access pathways, e.g., for industry who may be less familiar with public large-scale compute, via standardised access mechanisms, etc.

[41] In line with HMT Green Book guidance a reduced discount rate of 3% is applied to long-running benefits, such as lifetime earnings of people trained through AIRR, in years 31 onwards.

OFFICIAL

Further the estimates of Return on Investment are benchmarked against RoIs from similar large-scale compute infrastructure in the UK and abroad, providing context on how the UK's AI compute investments compare and helping to validate the robustness of the results.

8 Assessment of economic benefits and return on investment

8.1 Counterfactual scenario

The analysis considered the benefits of the UK's investments in AI compute against a hypothetical scenario where these investments would not have taken place (the "**counterfactual scenario**"). Under this scenario the AIRR and related national AI compute capabilities would not have been provided and users would have to use alternative compute infrastructure where this is possible.[42]

To understand the real-world implications of this scenario, for example the extent to which users would have been able to use alternative compute such as commercial cloud services, the user survey and interviews collected feedback from users of the existing AIRR capabilities (Dawn and Isambard-AI). These results also informed assumptions on additionality (i.e., the share of benefits that may have occurred anyway in the absence of public AI compute), further discussed in the Annex.

Overall, among respondents to the survey, 30% said they would have been able to access alternative compute resources had they been unable to access AIRR, while only 23% (three respondents) considered commercial cloud computing as a plausible alternative.[43] Moreover, all users responding to the survey, including those with access to alternative compute resources, noted that not being able to access AIRR would have at least some impact on their research or work with 30% saying the impact would have been catastrophic and a further 60% saying there would have been a major impact.[44] Further discussions on these results are provided in Section 9.4 in the *Baseline evaluation of the AI Research Resource* report accompanying this study.

8.2 Overarching modelling framework

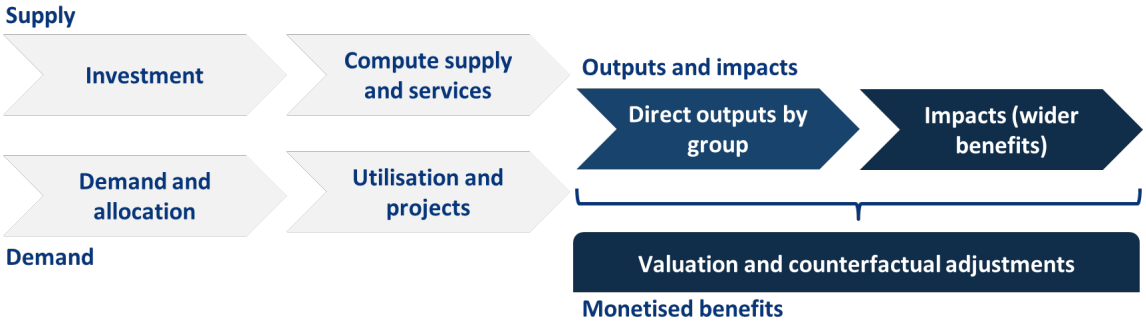
The analysis is forward looking and the vast majority of AI compute facilities have not yet been procured or built. As such, there is significant uncertainty around the final capabilities including the volume of compute that will ultimately be provided; the uptake of compute by future users, the composition of users, and their usage of the compute resources; the volume of innovations and other outputs materialising from this use; and the timing of provision, uptake, and for benefits to materialise.

To enable valuation of future benefits in light of this uncertainty, a step-by-step modelling framework that expresses direct outputs (such as publications, innovations, etc.) and wider benefits (such as contributions to economic growth) as a result of the interaction of (i) AI compute supply and enabling services and (ii) demand and allocation across user groups has been developed.

[42] Note, it is assumed that Dawn would have existed in the absence of the AIRR programme. As such, only a share of benefits emerging from Dawn was attributed to the AIRR programme. See footnote 40 for further details.
[43] Based on 46 survey responses to the question: "If you had not been able to access the AIRR, would there have been an alternative compute resource you could access?".
[44] Based on 43 survey responses to the question: "What would the impact on your research/work have been in this [no access to AIRR] scenario?".

OFFICIAL

Figure 8 RoI modelling – overarching modelling chain



Source: London Economics

The model first uses information on the planned investments including the amount of capital invested in AI compute infrastructure, the phasing of this investment over time, and high-level investment choices (e.g., the split between hardware, maintenance, and ancillary offerings such as trainings and architectural choices) as well as associated costs (e.g., hardware purchase prices) and operational constraints (e.g., plausible uptime, asset life of hardware, etc.), to derive estimates of the amount of usable AI compute capacity likely to be available per year.

In parallel to estimating supply, the model derives estimates of the usage of the available compute capacity by different user groups by taking into account information such as the allocation model to different user groups and the size of allocations; user demand and growth expectations over time; as well as usage constraints such as retention of surge capacity to enable flexible access or reserved allocations.

Together, these estimates of supply and demand provide an understanding of the plausible future usage of available AI compute capacity for each different user group (academia, industry, the public sector) over time.

Following this, the model then translates the estimates of compute usage into a set off expected outputs in line with the Benefits Realisation Plan (see Part II of this report). This includes public AI compute enabled research outputs and innovations such as publications, citations, the development of software, intellectual property, as well as spin-out companies. It also includes estimates of benefits to firms using AI compute in terms of revenue increases. To do this, the model uses assumptions, with appropriate uncertainties, from the literature as well as best judgement where robust assumptions are not available.

Finally, the model translates direct outputs into wider benefits such as the knowledge spillovers from AI research, and wider economic and growth effects again using appropriate assumptions and taking account of uncertainties. Benefits from ancillary services such as the value of training and skills benefits are also estimated.

Where possible the valuation methods account for attribution (how much of the estimated outcomes is due to the AIRR programme and related national AI compute) and apply appropriate adjustments to the estimated outputs and wider benefits to account for additionality (what outcomes would have happened anyway in the absence of the provision of public national AI compute capabilities). A detailed elaboration of the methods employed and adjustments made to value benefits is provided in Annex 1. The next section (Section 8.3 provides a brief overview).

OFFICIAL

8.3 Overview of assessed benefits and valuation methods

Building on the Theory of Change, benefits have been estimated across outputs, outcomes, and impacts. Monetised benefits fall into the following categories:

- **Direct benefits of access to users** capturing the avoided costs of accessing public AI compute compared to alternative commercial compute resources. These benefits are counted only for the share of users that would have accessed commercial compute under the counterfactual scenario.
- **Private benefits from work undertaken on public AI compute:** These benefits capture direct outputs of work such as academic publications, software, IP, as well as private benefits to firms, in terms of revenue increases, through channels such as efficiency savings, productivity improvements or commercialisation of new innovations. Benefits in this category are counted where the research would not have taken place under the counterfactual scenario.
- **Training and skills benefits:** Capturing the direct value of training provision as well as skills benefits to AI compute learners in terms of higher lifetime graduate premia and associated benefits to the UK exchequer. Benefits in this category were only assumed for a subset of users among research teams new to AI compute.
- **Benefits to the public sector:** Capturing efficiency savings from adoption of public sector-focused AI research projects undertaken on AIRR and related national AI compute in the public sector. Benefits in this category are counted where the public-sector focused research would not have taken place under the counterfactual scenario.
- **Wider economic impacts**: Capturing productivity spillovers to the private sector from academic research and downstream productivity effects among customers adopting the innovations produced using public AI compute and the wider economy. Benefits in this category are only counted where the research would not have taken place under the counterfactual scenario.

Table 6 provides an overview of the valuation methods employed to value economic benefits for each benefit stream that could be monetised.

In addition to the benefits presented in this table, there are significant unmonetiseable benefits. For example, this includes the strategic value of sovereign AI compute capability to the UK, benefits to the UK's international standing in AI and large-scale computing and international competitiveness in AI research and innovation, and value of retention and development of domestic AI/HPC infrastructure expertise.

Further, some of the monetised benefits are likely to be underestimates. For example, the value of software was estimated using a production cost approach quantifying the time-cost of developing the software. Actual benefits to the scientific community and private organisations using the software are likely to be substantially higher, though very difficult to monetise. Moreover, given the significant uncertainty surrounding benefits flowing from future AI compute infrastructure, conservative central assumptions were used throughout the analysis.

The central valuations therefore represent a conservative estimate of the overall benefits of AIRR and related national AI compute infrastructure. Further details on the approaches used to value each of these benefit streams, as well as assumptions used and important considerations underlying the economic assessments, are provided in the methodological annex (Annex 1).

OFFICIAL

Table 6 Valuation methods employed for individual benefit streams

Benefit stream	Benefit	Description	Valuation approach	Data / assumptions used
Direct benefits of access to users	Provision of AI compute (avoided cost)	GPU hours accessed via public AI compute, segmented by user type	Avoided cost of equivalent commercial cloud compute	Estimates of GPU hours used, market prices for comparable cloud GPU services. Adjustment for share of users that would have accessed alternative compute under the counterfactual.
Private benefits from work undertaken on public AI compute	Value of publications	Academic publications generated through use of AI compute	Value per publication based on literature proxies	Number of publications per project; value per publication; time lag for benefits, additionality assumptions.
	Value of citations	Downstream impact and influence of research outputs	Value per citation and citation-weighted publication value	Number of citations per publication; valuation per citation; Field weighted citation index for AI research; time lag for benefits; additionality assumptions.
	Value of software	AI Software tools and developed through research activities	Production cost approach to value AI research software (average time spent on developing software)	Data on time spent developing AI software; salary benchmarks, software products expected per project; time lag for benefits; additionality assumptions.
	Value of intellectual property (IP)	Protectable intellectual property generated through AI compute-enabled research	Value per patent/IP based on literature proxies	Expected IP outputs per project; average value per IP; time lag for benefits; additionality assumptions.
	Value of spinouts	Companies created from research enabled by AI compute	Lifetime GVA contribution of spinouts, based on academic literature	Probability of spin-out formation; survival rates; net present value of lifetime value per spinout; additionality assumptions; adjustment factor for research spillovers.
	Contribution of industry impacts to UK output	Productivity and innovation gains for firms using AI compute	Uplift in firm-level GVA aggregated across firms	Number of firms; average turnover per firm size, average uplift per firm; delay until impacts, duration of impact; additionality assumptions.

OFFICIAL

Benefit stream	Benefit	Description	Valuation approach	Data / assumptions used
Training and skills benefits	Direct value of provision of training courses	Value of training delivered through host sites	Market-equivalent cost of training provision	Number of training days or participants; commercial day rates; share of learners participating per training day.
	Premium and associated tax benefits of training users	Additional earnings and productivity gains for trained individuals, and associated fiscal benefits (e.g. income tax, NICs)	Earnings uplift over time, with associated tax revenues derived using tax rates	Number of trainees; estimated earnings premium; duration of impact; employment assumptions; tax rates; additionality assumptions.
Benefits to the public sector	Contribution of public sector projects to UK output	Efficiency and productivity gains from AI-enabled public sector use cases	Monetised time savings using wage-based cost equivalents	AI projects deployment rate; average FTE affected per deployment; average hours saved per FTE, wage rates; additionality assumptions.
Wider economic impacts	Spillover impact of scientific R&D (academic)	Wider economic spillovers generated by AI-enabled research activity	R&D spillover multiplier applied to research investment	Implied R&D funding for academic projects (e.g. AI compute cloud cost); spillover multipliers from literature; attribution to public AI compute.
	Wider societal returns from private gains	Broader spillover effects from industry innovation	Application of wider economic multipliers from literature	Base industry impact; multiplier values from literature.

Source: London Economics

8.4 Caveats and limitations

The results presented in this report are based on a forward-looking economic model and should be interpreted in light of a number of important limitations. These limitations primarily relate to the evolving nature of AI compute investments, the availability of evidence, and the assumptions required to translate access to compute into economic impacts.

- **Forward-looking nature of impacts:** The analysis assesses the expected economic impacts of the AIRR programme and investments in related AI compute infrastructure at an early stage. To date, realised benefits are largely concentrated around the construction, operation, and initial use of existing AI compute infrastructure (Dawn and Isambard-AI). Many of the anticipated benefits are expected to emerge gradually over time as capacity expands and usage matures.

  As a result, the evidence base on realised impacts remains limited, and a significant share of the benefits modelled in this analysis relate to future outcomes that are inherently uncertain.
- **Uncertainty in the magnitude of benefits:** Estimating the economic impact of AI-enabled research and innovation presents inherent challenges. In particular, many of the activities supported by AI compute involve early-stage or foundational research, where pathways to impact are often indirect, long-term, and highly diffused. For example, while there is strong evidence that academic research generates substantial spillovers to the wider economy, robust methods for quantifying these effects remain limited. The magnitude and timing of these impacts depend on how knowledge diffuses and is translated into productivity gains across firms and sectors.

  More generally, traditional cost-benefit analysis approaches, such as Net Present Value, are not always well suited to capturing the full extent of uncertainty associated with R&D investments, particularly in rapidly evolving technological domains such as AI. In this context, future impacts are not only uncertain in magnitude but also depend on a range of evolving policy and delivery choices and wider ecosystem conditions. These include decisions on the scale and location of AI compute infrastructure, the choice of hardware, the allocation of compute across different user groups, and the nature of supported activities (e.g. fundamental research, foundational AI development, or applied use cases).

  As a result, the model represents a structured approximation of potential outcomes under a given set of assumptions, rather than a precise prediction of future impacts.
- **Use of general assumptions where AI-specific evidence is limited:** In a number of benefit areas, AI-specific evidence on impacts is still emerging. As a result, the model draws on broader evidence from the literature on public R&D, innovation, and productivity to inform key parameters.

  While this approach ensures that assumptions are grounded in established evidence, it may not fully reflect the specific characteristics of AI-related research and its potential impacts. The model has been designed to allow for the refinement of these parameters as more AI-specific evidence becomes available. The sensitivity analysis presented in Section 10 has been used to explore how results vary under alternative assumptions, including scenarios that reflect stronger AI-related impacts.
- **Simplified representation of adoption and impact pathways:** The model necessarily simplifies the process through which access to AI compute translates into economic impacts. In particular, it assumes a relatively direct relationship between compute access, research outputs, and downstream economic benefits. In practice, the realisation of these

OFFICIAL

impacts is likely to depend on a range of additional factors, including the availability of relevant skills, organisational capacity within firms and public sector bodies, and complementary investments in data, infrastructure, and implementation. These factors may affect both the timing and scale of realised benefits.

These enabling conditions are reflected in the programme's Benefit Realisation Plan and will need to be effectively harnessed to maximise the impacts of AI compute capabilities.

- **Independence of modelled parameters in the sensitivity analysis:** The probabilistic sensitivity analysis varies key model parameters independently. While this allows for a transparent exploration of uncertainty across a wide range of inputs, it does not explicitly capture potential relationships between variables. In reality, some parameters may be correlated. For example, stronger research outcomes may be associated with higher spillover effects or greater commercialisation success. The absence of such correlations in the modelling may result in an under- or over-estimation of combined effects in certain scenarios.
- **Partial coverage of total benefits:** While effort has been made to quantify economic benefits wherever possible, not all impacts can be robustly monetised. The analysis focuses on those benefits that can be expressed in monetary terms, such as productivity gains, avoided costs, and earnings premia. However, a range of additional benefits are not fully captured in the quantitative estimates. These include strategic and capability-building impacts, such as strengthening the UK's position in AI research, enhancing resilience, and supporting long-term innovation ecosystems. As a result, the estimates presented in this report should be interpreted as a conservative estimate of the total value generated by the programme.

Taken together, these limitations reflect the inherent uncertainty associated with modelling the impacts of emerging technologies. Results presented in Sections 9 and 10 should therefore be interpreted as indicative estimates of potential economic value, providing a structured and transparent basis for assessment rather than precise forecasts of future outcomes.

9 Central valuations of economic benefits and RoI

9.1 Aggregate benefits and Return on Investment

Table 7 provides the main estimates of the aggregate economic benefits expected to flow from the UK Government's investments in AIRR and related national AI compute capabilities. Estimated aggregate benefits are compared to expected total costs to the public purse. Implied Return on Investment (RoI) figures, derived by comparing total aggregate benefits to total public costs, are also provided.

Table 7 Aggregate estimates of benefits and RoI (£$_{2025/26}$ prices)

Metric	Valuation using conservative central assumptions
Total monetised economic benefit	£12.1 billion
Total expected costs accruing to public purse	£2.0 billion
Return on Investment (RoI)	**£6.0 per £1 invested**

Note: The total expected costs over AIRR and related public AI compute capabilities operational life used in the RoI estimation are based on UK government investment plans in the public compute ecosystem, committed to in the UK Compute Roadmap. *Monetised values are discounted to 2025/26 terms.* ***Source: London Economics analysis***

OFFICIAL

The estimates presented in Table 7 are based on conservative central assumptions used as inputs to the economic model. As such, they should be interpreted as realistic but cautious baseline estimates of the economic value generated by AIRR and related national AI compute capabilities. In particular, these estimates do not capture the potentially substantial benefits that could arise from low-probability, high-impact outcomes, such as the emergence of highly successful AI firms or breakthrough innovations enabled by access to public AI compute.

Using these assumptions, total monetised economic benefits are estimated at £12.1 billion in real prices over the appraisal period. This compares to total costs to the public purse of approximately £2.0 billion in real prices, implying a Return on Investment of £6.0 per £1 invested.

The estimated RoI highlights the strong return that public investment in AI compute infrastructure is expected to deliver to the UK economy, even when assessed under conservative assumptions. Annualising these benefits over the 10-year appraisal period implies benefits of approximately £1.2 billion per year. This is equivalent to around 10% of the current annual GVA of the UK AI sector (estimated at £11.8 billion),[45] indicating that the programme represents a material contribution to the UK's AI-driven economic activity.

While these central estimates provide a robust baseline, they do not reflect the full range of potential outcomes associated with uncertainty in key assumptions. The implications of this uncertainty are explored in more detail in Section 10.

9.2 Central valuations by benefit stream

Table 8 provides a breakdown of the central valuations of economic benefits for each benefit stream that could be monetised.

Table 8 Valuations by benefit stream (£$_{2025/26}$ prices)

Benefit stream	Valuation using conservative central assumptions
Access to AI compute	**£1,146m**
Direct benefit of avoided costs of access	£1,146m
Private benefits from work undertaken on public AI compute	**£932m**
Value of publications	£40m
Value of citations (average-level research)	£0.5m
Value of citations (quality premium)	*£0.5m*
Value of software	£44m
Value of IPs	£38m
Value of spinouts	£117m
Contribution of industry impacts to UK GVA	£692m
Wider economic impacts	**£8,438m**
Spillover impact of scientific R&D	£6,361m

[45] Estimates from the 2024 UK AI sector study; DSIT (2025). Artificial Intelligence sector study 2024. Available at: https://www.gov.uk/government/publications/artificial-intelligence-sector-study-2024/artificial-intelligence-sector-study-2024.

OFFICIAL

Benefit stream	Valuation using conservative central assumptions
Spillover impact industry contributions	£2,077m
Benefits to the public sector	**£663m**
Public sector efficiency savings	£663m
Training and skills benefits	**£872m**
Direct value of provision of training courses	£5m
Premium and associated tax benefits of training users	£868m
Total monetised economic benefit	**£12,052m**

Source: London Economics analysis

Wider economic impacts

The largest share of monetised benefits is generated through wider economic impacts, estimated at £8.4 billion under conservative central assumptions, accounting for 70% of total monetised benefits. These impacts are primarily driven by spillover effects associated with scientific R&D (£6.4 billion), alongside spillovers linked to industry contributions (£2.1 billion).

These benefits arise through knowledge diffusion, follow-on innovation, and productivity improvements across the wider economy. For example, early deep learning research conducted in universities led to breakthroughs such as AlexNet, a deep learning model for image classification that significantly advanced the performance of computer vision systems, which subsequently led to increased use of AI imaging technologies in the medical domain.[46]

While these wider spillovers are inherently indirect and difficult to observe directly, qualitative evidence highlights the mechanisms through which they may arise. Interviewees pointed to the translation of AI compute-supported research into downstream applications, including its use in ongoing contractual work and early-stage commercial activities. One interviewee noted that the ability to demonstrate full end-to-end control over model development can strengthen competitive positioning, “giving an edge over other organisations bidding for similar opportunities”.

Other users emphasised the role of open research outputs, such as models and software, in enabling wider uptake across the research community and in other domains. One interviewee highlighted plans to develop open-source tools to support more efficient workflows, noting that “the aim is to build software that others can use and adapt, contributing to the wider community, with potential for commercialisation in the future.” Together, these examples illustrate how AI compute-supported activity can extend beyond individual projects to generate broader economic value through diffusion and reuse.

The wider economic impacts category also captures downstream adoption of AI innovations by customers of firms developing AI products on AIRR and related public AI compute infrastructure and wider use across sectors. For example, this could include translation of AI research into a commercial product that is then taken up by other firms, thereby supporting wider diffusion of AI capabilities across the economy.

[46] Singh, P. et al (2020). 3D Deep Learning on Medical Images: A Review. Available at: https://arxiv.org/abs/2004.00218

OFFICIAL

Direct benefits from access to AI compute

These wider impacts are underpinned by more immediate and directly observable benefits to users of public AI compute. Direct benefits from access to AI compute, measured as avoided costs relative to commercial provision, are estimated at £1.1 billion, accounting for 10% of total monetised benefits. These represent immediate and tangible economic gains to users of public AI infrastructure.

The impact survey evidence strongly supports the importance of these cost savings. Only 30% of respondents indicated that they would have been able to access alternative compute resources in the absence of AIRR, and even in those cases, alternatives were typically smaller-scale institutional resources with limited capability. Commercial cloud provision was considered a viable alternative by only a small minority of users, reflecting both cost and access constraints. Indicative cost comparisons further highlight the scale of these benefits, suggesting that replicating AIRR allocations through commercial cloud could cost between £28,000 and £1.2 million per project, depending on scale.[47]

Private benefits from work undertaken on public AI compute

Building on this, private benefits from work undertaken on public AI compute are estimated at £0.9 billion (8% of total monetised benefits). Of this total, approximately 74% (£0.7 billion) is estimated to accrue through industry impacts to UK GVA, reflecting increased productivity, commercialisation, and revenue generation among firms using public AI compute. The remaining benefits are associated with research and innovation outputs, including publications, software, intellectual property, and spin-outs, with spin-outs alone contributing £0.1 billion.

The impact survey evidence indicates that access to AIRR is a key enabler of these private benefits. Around 59% of respondents reported that they would have been unlikely or unable to achieve similar outcomes without AIRR, with a further 28% indicating that only some outcomes would have been achieved.[48] This is reflected in concrete examples from user interviews. For instance, one research team reported producing 8 to 10 academic papers within a short period of access to AIRR, alongside the development of software tools, code, and open-source models, and identified potential pathways towards commercialisation and spinout activity.

Evidence from industry users further highlights the commercial relevance of these outputs. One start-up reported that AIRR access enabled the creation of a proprietary dataset described as a significant competitive advantage, underpinning ongoing discussions with industrial partners and future commercial applications. In another case, AIRR-supported research for computer vision tools has already informed ongoing contractual work, demonstrating early-stage translation of research outputs into applied use. These examples illustrate how public AI compute enables both the generation of high-value research outputs and their translation into commercial and applied contexts.

Benefits to the public sector and training and skills benefits

Training and skills benefits and benefits to the public sector accounted for the lowest share of central benefit estimates accounting for 7% and 6% of total benefits, respectively. However, this

[47] These estimates consider individual cloud prices, costs for large-scale compute provision via cloud (e.g., Met Office supercomputing deal with Microsoft Azure) are different and would likely incur substantial discounts. Further analysis is provided in Section 9.4 in the *Baseline evaluation of the AI Research Resource* report.
[48] Impact survey results are provided in Section 9.4 in the *Baseline evaluation of the AI Research Resource* report.

primarily reflects the conservative scope of what is included in the valuation. Public sector benefits are limited to efficiency savings among staff arising from the adoption of AI innovations developed on public AI compute, and do not capture wider outcomes such as improved healthcare, welfare gains, or broader societal impacts.

Further, AI compute projects targeting the public sector could deliver substantially higher benefits for example through large-scale transformation of public service delivery. However, such scenarios are highly uncertain and depend on a variety of factors outside of AI compute provision itself, and therefore very difficult to capture in the analysis in a robust way. As such these scenarios were not modelled explicitly, meaning public sector benefits could conceivable be substantially higher than indicated by the benefit estimates presented in this study.

Qualitative evidence provides examples of how such public sector impacts may arise in practice. Interviewees highlighted the development of AI models and workflows applied to healthcare data, with potential relevance for improving diagnostic processes and supporting NHS workflows. While these impacts are not directly monetised in the central estimates, they illustrate the types of applications that underpin the modelled efficiency savings.

Similarly, training and skills benefits capture only the direct value of training provision and the benefits to individuals in terms of higher lifetime earnings. Wider system-level impacts of skills development were not estimated to avoid double counting with other benefit streams. Qualitative evidence suggests that public AI compute contributes to skills development and capability building within the UK AI ecosystem, including through training researchers, attracting high-skilled talent, and enabling students and early-career researchers to develop advanced technical skills through hands-on use of large-scale compute infrastructure.

Non-monetised benefits and uncertainty surrounding composition of benefits

As highlighted in Section 9.1, these estimates capture only those benefits that could be robustly monetised. A range of additional impacts, including strategic, societal, and capability-building benefits, remain unquantified. This suggests that the central estimates presented in this section should be interpreted as conservative lower-bound estimates of the total economic value generated by investments in public AI compute infrastructure, with the extended estimates illustrating the potential scale of additional benefits arising from wider public R&D spillovers.

In addition, the composition of benefits could vary significantly depending on the type of AI innovations that will be produced on AIRR and related public AI compute and their ultimate uptake in industry, the public sector, or wider society. In particular, the sensitivity analysis shows the potential for both private benefits to users and public sector efficiency outcomes to be substantially larger than indicated by these conservative central estimates (see Section 10.2).

10 Sensitivity analysis surrounding central benefit estimations

The central estimates presented in Section 9 are based on a set of conservative assumptions reflecting the expected economic impacts of the AIRR programme and investments in related public AI compute infrastructure. However, given the forward-looking nature of the analysis, there is inherent uncertainty surrounding many of the key parameters. In addition, the central scenario does not fully capture several sources of potentially significant economic value.

OFFICIAL

First, the estimation of wider economic spillovers from academic research is limited to the level of research activity implicitly captured through the value of compute allocated to academic users. In practice, public AI compute is expected to enable and attract additional research funding, which may generate further spillovers through knowledge diffusion, innovation, and productivity gains. While some of this activity may occur in the counterfactual, there is a strong rationale to expect that AI-enabled research supported by public AI compute generates higher returns than more conventional forms of investment.[49] To reflect this, the sensitivity analysis incorporates an uplift to selected research-related benefit streams, of between 0% and 40% with a central value of around 20%, capturing the potential for AI-driven innovation to deliver higher productivity gains than standard R&D[50]. Taken together, these factors suggest that spillover impacts in the central case are likely to be understated.

Second, the central analysis does not capture the potential for outsized returns associated with high-impact outcomes. These include the emergence of highly valuable spin-outs, breakthrough industry applications, and large-scale public sector deployments that generate transformational efficiency gains. Evidence from AI and digital innovation suggests that such outcomes, while rare, can account for a disproportionate share of total economic value.

These effects are inherently uncertain and difficult to quantify within a deterministic central scenario. To account for this, a probabilistic sensitivity analysis was undertaken. This approach allows multiple assumptions to vary simultaneously, rather than testing individual parameters in isolation. For each model run, values are drawn from predefined probability distributions and used as inputs to the full economic model. Repeating this process across a large number of simulations generates a distribution of possible outcomes for total economic benefits and Return on Investment (see Annex 2 for a more detailed explanation on the sensitivity analysis methodology).

The sensitivity analysis is specifically designed to capture these additional sources of uncertainty and upside potential. In particular, skewed, long-tailed distributions are applied to selected parameters to allow for the possibility of low-probability, high-impact outcomes. This includes modelling the potential for high-impact outcomes, such as valuable spin-outs, significant industry applications, and public sector deployments that may generate substantial returns.

However, while this approach allows for the possibility of higher returns, it remains challenging to fully capture the most extreme outcomes within a tractable modelling framework. In particular, scenarios involving the large-scale transformation of public services, or the emergence of globally leading AI firms generating exceptionally high economic value, are subject to very high levels of uncertainty and are not fully reflected in the modelled distributions. As a result, even the upper bound estimates from the sensitivity analysis may not fully capture the potential upside.

[49] For reference see Chen et al. (2024), this study identified that AI patents exhibit a 9% value premium, and 26% higher forward citations relative to non-AI patents. Available at: https://www.hbs.edu/ris/Publication%20Files/24-069_e5bcc300-d7f3-43b7-af9f-945b5937495e.pdf.

[50] A large body of literature characterising the sizeable productivity impacts AI could deliver across a wide range of economic sectors exists. As such, it appears likely that AI research would deliver additional economic impacts than those captured by traditional spillover estimates covering all forms of research. However, economic literature on the impact of AI research vs. other forms of research is scarce and significant uncertainty surrounds this assumption. Nevertheless, a number of indicative sources exist. For example, AI patents typically receive 26% more forward citations than non-AI patents (Harvard Business School (2024). The Value of AI Innovations. Available at: https://www.hbs.edu/ris/Publication%20Files/24-069_e5bcc300-d7f3-43b7-af9f-945b5937495e.pdf), while, in climate science, AI in climate patents were associated with around 30–100% more subsequent inventions when counting all technologies (Verendel V. (2023). Tracking artificial intelligence in climate inventions with patent data. Available at: https://www.nature.com/articles/s41558-022-01536-w). However, due to the limited explicit evidence on spillover impacts, no additional impact was assumed in the central case and a conservative range centred around a conservative value of 20% was chosen for the sensitivity analysis.

OFFICIAL

The results of this analysis are presented in the following sections. Section 10.1 summarises the overall distribution of results, while Section 10.2 examines how benefits are distributed across individual benefit streams and highlights the key drivers of uncertainty underpinning the results.

10.1 Results of the sensitivity analysis

The sensitivity analysis generates a distribution of possible outcomes for total economic benefits and Return on Investment (RoI) estimates, reflecting uncertainty across the full set of modelled assumptions. The results not only highlight variations in magnitude of estimated benefits, but also the likelihood with which benefits within certain ranges may emerge.

Table 9 presents summary statistics reflecting the distribution of estimated economic benefits and the associated Return on Investment (RoI) figures across the simulations. The table shows the central valuations discussed in Section 9, alongside the mean and median estimates from the sensitivity analysis and key cut-off points (percentiles) in the range of the estimated benefits. These percentiles reflect the percentage of simulation results that fall below a given point.

Table 9 Distribution of total economic benefits and Return on Investment (£$_{2025/26}$ prices)

Metric	Total benefits	Return on Investment (RoI)
Central estimates	***£12.1bn***	***£6.0 : £1***
Mean	£20.8bn	£10.4 : £1
99th percentile	£64.3bn	£32.2 : £1
95th percentile	£40.5bn	£20.2 : £1
80th percentile	£27.4bn	£13.7 : £1
Median (50th percentile)	£18.3bn	£9.1 : £1
20th percentile	£12.2bn	£6.1 : £1
5th percentile	£8.5bn	£4.2 : £1
1st percentile	£6.3bn	£3.1 : £1

Source: London Economics analysis

The mean estimated total economic benefit is £20.8 billion, significantly higher than the central estimate of £12.1 billion presented in Section 9. The median estimate is somewhat lower, at £18.3 billion, reflecting the skewed distribution of outcomes. A similar pattern is observed for the associated Return on Investment metrics. The mean RoI figure is approximately £10.4 per £1 invested, higher than the central estimate of £6.0, with a median RoI of £9.1 per £1 invested.

Across the central range of simulations, reflecting the results of the majority (60%) of simulations, total economic benefits are estimated to lie between £12.2 billion (20th percentile) and £27.4 billion (80th percentile). A wider range of £8.5 billion to £40.5 billion is observed across the 5th to 95th percentiles, reflecting the range in which 90% of estimated benefits fall across the simulations.

Figure 9 shows that the distribution of outcomes is characterised by a concentration of simulations around the central estimates, alongside a pronounced upper tail. This means that while the majority of simulations fall within a relatively bounded range, there is a possibility for substantially higher benefits to emerge in a range of higher-impact scenarios.

Figure 10 provides a detailed view of the upper tail of the distribution (the top 5% of simulations), illustrating the scale and variation of these high-impact outcomes. While 5% of simulations result in

OFFICIAL

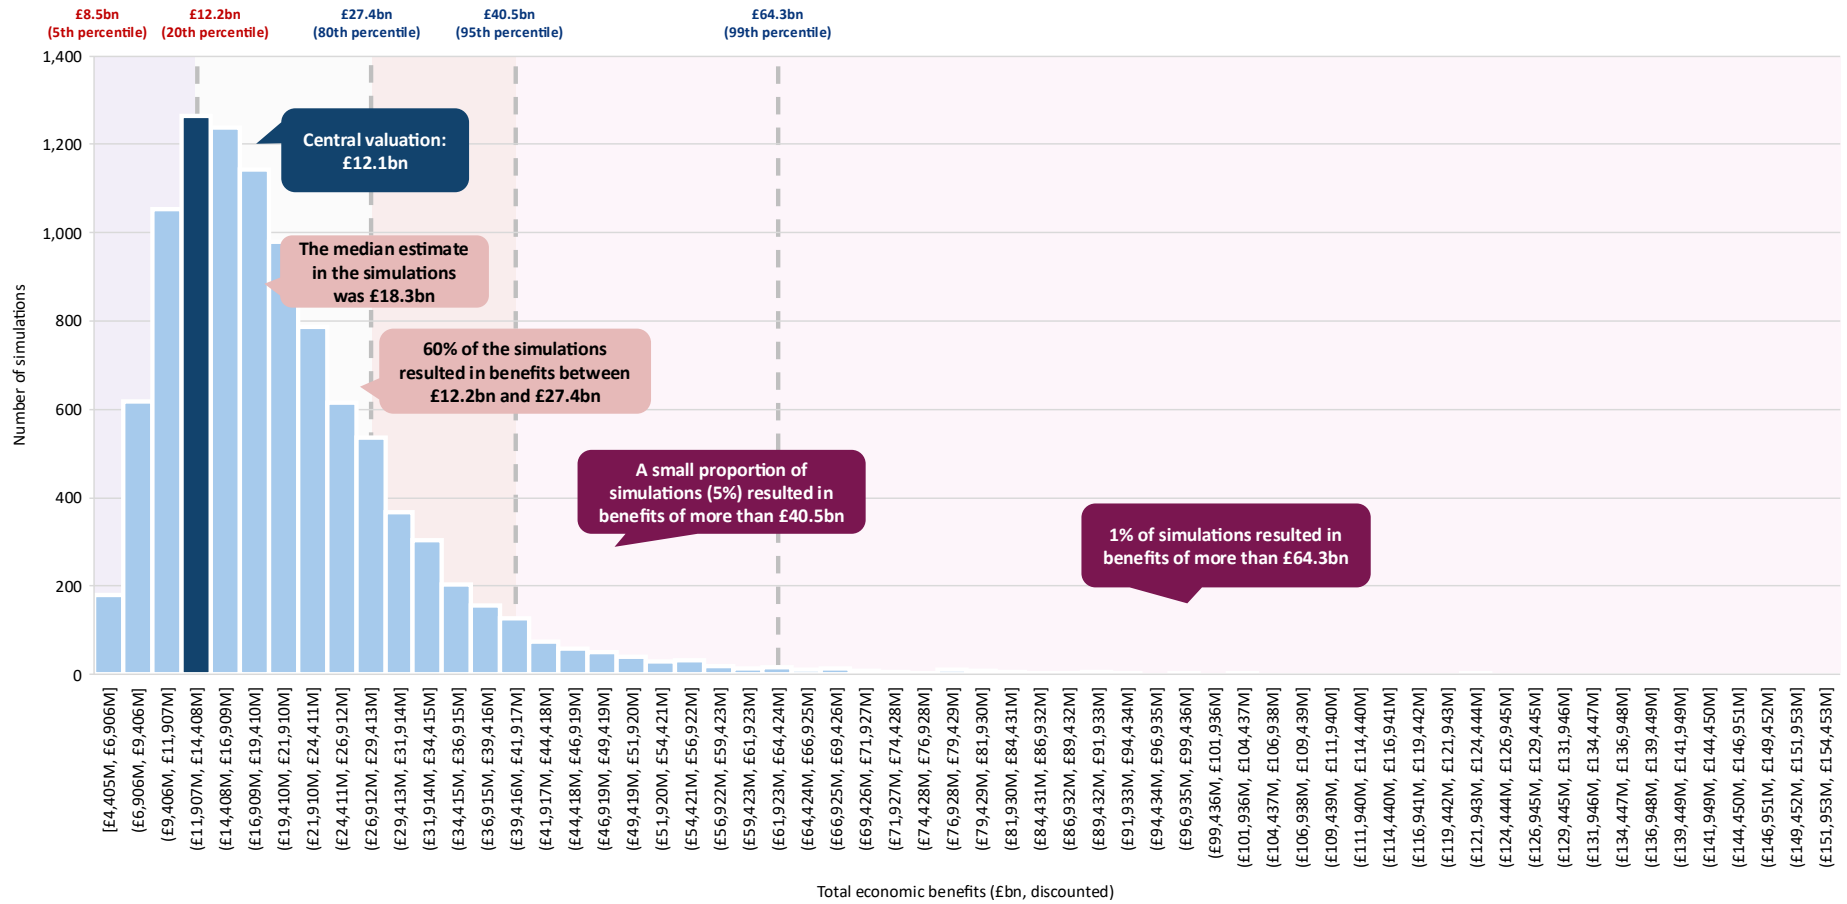
benefits above £40.5 billion, a substantial proportion of these outcomes extend well beyond this threshold. The upper tail includes simulations with benefits exceeding £100 billion and, in extreme cases, above £150 billion (equivalent to a RoI of £75 per £1 invested). However, benefit estimates of this magnitude were extremely rare and only emerged in a very small number of simulations.

These upper-tail outcomes reflect scenarios in which public AI compute-enabled research and innovation lead to highly impactful breakthroughs, including the emergence of high-value firms, transformative technologies, and large-scale productivity gains. While such outcomes are inherently uncertain and are not included in the central estimates, the analysis demonstrates that the programme has the potential to generate substantially outsized economic returns.

Overall, the results indicate that the central estimates presented in Section 9 sit within the lower end of plausible outcomes and should be interpreted as conservative baseline estimates, with significant upside potential.

OFFICIAL

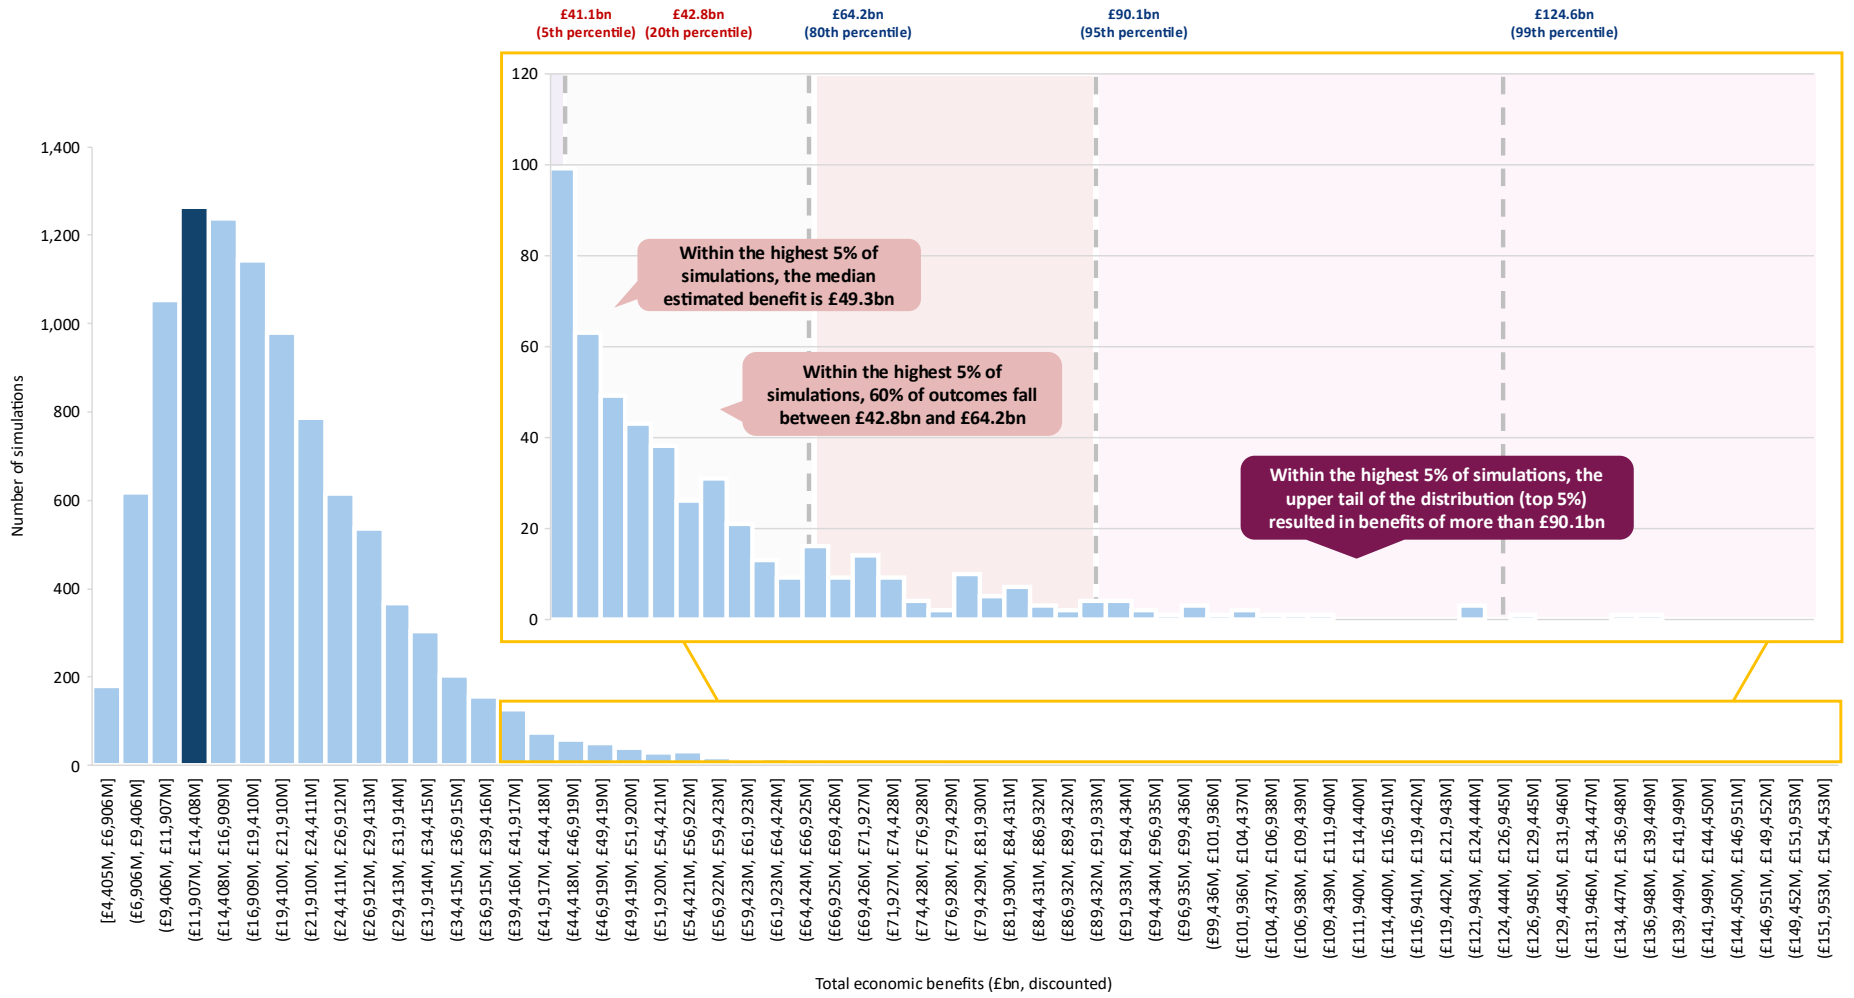
Figure 9 Sensitivity analysis results: Distribution of total economic benefits



Source: London Economics Analysis

OFFICIAL

Figure 10 Sensitivity analysis results: Deep dive on the upper tail (top 5%) of the distribution of total benefits



Source: London Economics Analysis

OFFICIAL

10.2 Variation in benefit estimates across benefit groupings

The sensitivity analysis also provides insight into how total economic benefits are distributed across individual benefit streams and how this distribution varies across simulations.

The overall contribution of benefit streams to overall impacts remains broadly consistent with the central estimates presented in Section 9. In particular, wider economic impacts remain the dominant source of value, across simulations. Similarly, private benefits from work undertaken on public AI compute represent the second largest category across a large number of simulations. Within this category, industry impacts to UK GVA account for the majority of value.

However, the sensitivity analysis shows the potential for both private benefits to users and public sector efficiency outcomes to be substantially larger than indicated by the conservative central estimates (see Figure 11). In particular, the impact on public sector efficiency gains is highly uncertain and depends on the scale, effectiveness and breadth of adoption of AI innovations across the public sector. As a result, the conservative assumptions chosen result in a relatively smaller share of total benefits for public sector impacts but show large variations in this share with more optimistic simulation runs showing public sector benefits accounting for more than one-third of total benefits and the most optimistic simulations showing shares in excess of 50% (see Figure 13 and Figure 14).

In practice, the composition of benefits could also vary significantly depending on the type of AI innovations that will be produced on AIRR and related public AI compute, their uptake in industry, the public sector, or wider society, and the magnitude and effectiveness of their ultimate impact. Moreover, the analysis does not capture wider benefits from public sector adoption, such as improved healthcare outcomes, welfare benefits, or the value of safer streets. Similarly, training and skills benefits only account for the direct value of training provision and the benefits to trainees themselves (in terms of higher lifetime earnings). Wider benefits of skills development and training were not estimated to avoid double counting with other benefit streams. These may significantly affect the share of total benefits resulting from these groups.

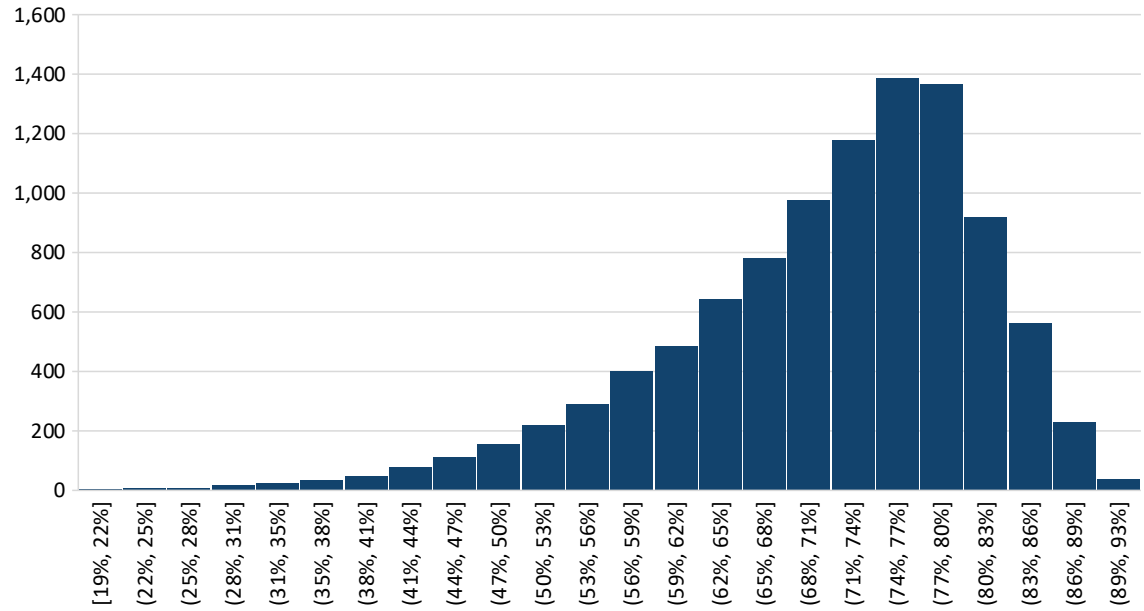
Figure 11 Sensitivity analysis: Variation of largest benefit groupings



Note: These plots show the results of the sensitivity analysis for the top three benefit groupings with the widest variation in estimated impacts. The boxplots summarise the distribution of sensitivity results. Each box shows the middle range of values resulting from the simulations (from the 25th to the 75th percentile), with the line inside the box indicating the median (middle value). The "whiskers" extend to show the range of most results, while any dots represent individual results that fall outside this typical range. Note the different scale among wider economic benefits due to the higher magnitude of monetised benefits. ***Source: London Economics***

OFFICIAL

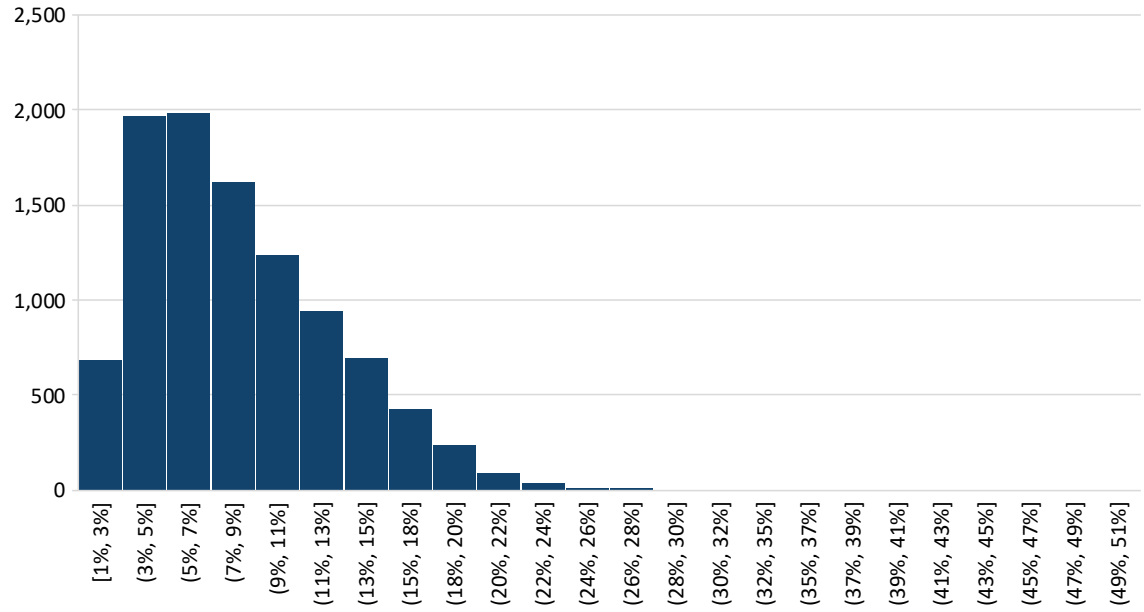
Figure 12 Share of wider economic impact in total benefits



Note: These plots show the distribution in the shares of total benefits across simulation runs for the three highest benefit groupings (wider economic benefits in this case). Each bar shows the number of simulations that resulted in benefits within the indicated range with the height of the bars representing the frequency with which shares in this range were observed.

Source: London Economics

Figure 13 Share of private benefits in total benefits

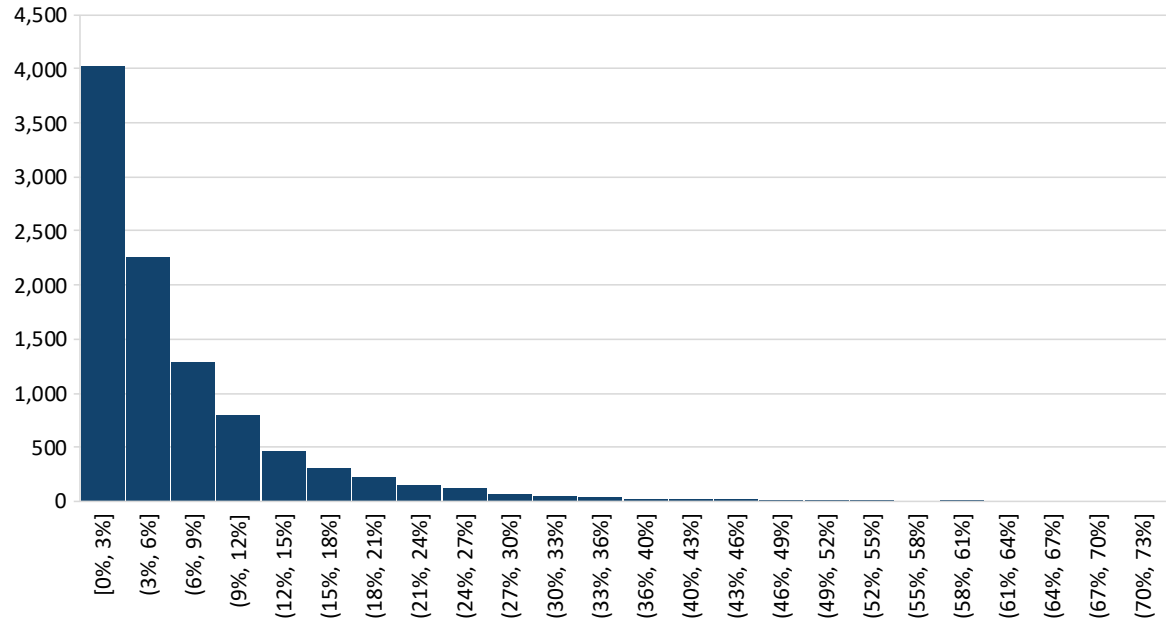


Note: These plots show the distribution in the shares of total benefits across simulation runs for the three highest benefit groupings (private benefits in this case). Each bar shows the number of simulations that resulted in benefits within the indicated range with the height of the bars representing the frequency with which shares in this range were observed.

Source: London Economics

OFFICIAL

Figure 14 Share of public sector benefits in total benefits



Note: These plots show the distribution in the shares of total benefits across simulation runs for the three highest benefit groupings (public sector benefits in this case). Each bar shows the number of simulations that resulted in benefits within the indicated range with the height of the bars representing the frequency with which shares in this range were observed.

Source: London Economics

To reflect these uncertainties, the model assigns probability distributions to some inputs that are skewed and allow for a higher likelihood of extreme values (i.e., have a "heavy-tail"). In particular, lognormal distributions are used where appropriate to capture the potential for uncertain higher-than-expected outcomes while maintaining conservative central assumptions (e.g. higher spillovers, highly successful commercialisation by firms, higher value spinouts, and larger-scale public sector adoption).

It is also important to note that differences in the estimated contribution of individual benefit streams are driven by a combination of underlying evidence, modelling assumptions, and allocation choices within the framework. As such, these estimates should not be interpreted as precise indicators of where investment should be prioritised. In practice, realised benefits will depend on a range of contextual factors, including the nature of projects supported, user demand, and the evolving AI ecosystem.

Consequently, while the analysis provides useful insight into the relative scale and drivers of different benefit streams, it should not be used in isolation to inform allocation decisions. The results instead highlight that the overall value for money of the programme is driven by a combination of complementary benefit channels, with wider economic impacts playing a particularly important role in shaping the potential for higher-impact outcomes.

OFFICIAL

11 Comparison to estimated benefits of other evaluations of HPC and science infrastructure

A number of existing studies evaluate the benefits of investments in large-scale compute infrastructure. UK evidence includes the 2010 evaluation of HPC Wales, the 2018 baseline evaluation of the Hartree Centre, the 2019 evaluation of EPSRC's investments in High Performance Computing, and the 2025 evaluation of ARCHER2. In addition, evaluation studies of the National Collaborative Research Infrastructure Strategy (NCRIS) Australia's investments in science infrastructure (from 2021), including the Pawsey Supercomputing Research Centre, and the Finnish IT Center for Science (CSC)'s investments in large-scale compute infrastructure (2024) provide international context.

However, evidence on AI-compute specific investments is scarce. To date, the only relevant evaluation is the 2024 assessment of CSC's high-performance computing services, which—while not limited to AI—includes early evidence from systems such as LUMI, an AI factory based in Finland, and provides indicative estimates of economic returns (see below). Moreover, among evaluations of public compute infrastructure that provide quantitative estimates of return on investment, reported results vary widely. Aggregate benefit-cost ratios or RoI figures are typically reported in the range of around £2 : £1 to £20 : £1 when assessed over an appropriate programme or asset lifetime. This variation reflects differences in evaluation scope, timing, and methodological choices, including which benefits are included, how counterfactuals are treated, and the maturity of the infrastructure at the point of assessment.

Early-stage or baseline evaluations tend to report lower returns, reflecting the fact that many benefits materialise only over the longer term. Some studies also report even higher returns. The 2024 evaluation of CSC's high-performance computing services (capturing two national supercomputers, Puhti (2019) and Mahti (2020), and the pan-European LUMI (2022)) is particularly noteworthy, reporting high return figures of €25–€37 return per euro invested.[51] However, the full study has not been published, making it difficult to understand what drives these high valuations.

In contrast, a limited set of firm- and project-level studies, report substantially higher financial returns, in some cases exceeding hundreds of pounds of revenue per pound invested. These figures are derived from reported revenues, profits, or cost savings associated with successful compute-enabled projects and reflect a much narrower scope of benefits. While not directly comparable to public-sector evaluations, they provide indicative evidence of the scale of commercial value that compute infrastructure can unlock in specific application contexts.

While not directly comparable to AI compute, these studies nevertheless provide useful context and validation for the estimated return on investment figures derived for AIRR and related national AI compute capabilities. In particular, it is important to note that many of the higher return estimates reported in the literature (see Table 10) are based on **ex-post evaluations using realised data** from mature infrastructure and established user bases. By contrast, the central estimates of £6.0 per £1 invested presented in this study are **ex-ante,** reflecting assumptions about how benefits materialise over time for a programme that is still at an early stage of development.

The sensitivity analysis indicates an aggregate return on investment in the range of £4.2 to £20.2 per £1 invested (covering 90% of simulations). This range is consistent with the evidence reported

[51] CSC (2024). SROI of CSC's high-performance computing services studied: €25–37 return per euro invested. Media Release. Available at: https://csc.fi/en/media-release/sroi-of-cscs-high-performance-computing-services-studied-e25-37-return-per-euro-invested/

OFFICIAL

in comparable national and international evaluations of compute and research infrastructure. While differences in scope, methodology, and timing limit direct comparability, the results provide additional confidence that the estimated returns for AIRR and related public AI compute infrastructure are within a plausible range when set against the wider evidence base. Further, more optimistic scenarios suggest that realised benefits could extend beyond the RoI figures commonly reported in the evaluation literature of large-scale compute infrastructures.

Table 10 Headline findings from evaluations of HPC compute infrastructure

Study	Scope	Study period	Headline findings
London Economics (2019) – Impact of EPSRC HPC	UK national HPC infrastructure: ARCHER (2013) + Tier-2 regional computing centres (2012) + Tier-2 centres (2016) + continuation impacts from HECToR (2007)	2008-2018	Aggregate ROI estimated between 6.5:1 and 19.5:1, with total economic impacts of £3.0bn-£9.1bn, driven by research productivity, skills, and spillovers
London Economics (2025) – Impact of ARCHER2	UK national HPC service: ARCHER 2 (2020)	2020-2026	Aggregate ROI estimated at 8.3:1 and 5.2:1 excluding research spillover benefits. 90% confidence range around central estimates of between 4.6:1 and 24.1:1.
Innovation Partnership (2015) – HPC Wales	Wales; ERDF-funded regional HPC programme: HPC Wales (2010)	2010-2015	Estimated ROI from employment impacts of 2.1:1 (at the time of publication), projected to rise to a total ROI 8.4:1. Additional smaller ROI from training impacts estimated to be 0.1 at the time of publication, rising to a total of 0.6 in the long term.
CSC / Taloustutkimus (2024)	Finland; national HPC services: Puhti (2019) and Mahti (2020), and the pan-European LUMI (2022)	2018-2023	Social return on investment estimated at €25-€37 per €1 invested, based on stakeholder benefits; wider spillovers not monetised
Technopolis (2018) – Hartree Centre	UK; baseline evaluation of (national) Hartree Centre Phase 1 and 2 (2013), excl. Phase 3 (2016)	2013-2016	The baseline evaluation estimates £27.5m in commercial net economic impact and a further £7.1m from operational expenditure, described as already close to the £37.5m initial Phase 1 capital investment within four years of opening. Excludes longer-term scientific impacts.
Hyperion Research for Ansys (2023)	Global; various (71) completed firm level HPC-enabled projects and success stories (mid-2010s through early-2020s)	Mid 2010s through early 2020s	Average $151.90 in revenue and $34.90 in profit generated per dollar of HPC invested[52]
Hyperion Research (2019/2020)	Global; various (763) firm-level completed/successful HPC projects (early 2000s through late 2010s)	Early 2000s through late 2010s	Average $507 revenue and $47 profit or cost savings per $1 invested, excluding outliers; direct financial returns only
IDC (2013) – Pilot ROI study	Global; various (208) HPC-enabled project case studies prior to 2013	Early-2000s through early-2010s	Average $356 revenue and $38 profit or cost savings per $1 invested; returns typically realised around two years post-investment
Lateral Economics (2021) – NCRIS	Australia; National Collaborative Research Infrastructure Strategy (NCRIS) investments in national research infrastructure since 2005-06 including in Pawsey Supercomputing Research Centre (2013)	2005/06 - 2022/23	Applies an assumed ROI of approximately 7.5:1, informed by prior infrastructure studies, to project aggregate community benefits of NCRIS funding

Source: London Economics

[52] These figures reflect project-level commercial performance in a specific application domain and are not directly comparable to public-sector value-for-money assessments or to evaluations of national compute infrastructure.

OFFICIAL

12 Summary of economic valuation and key takeaways

The economic valuation assessed the Return on Investment (RoI) associated with the UK Government's investments in the AI Research Resource (AIRR) and related national AI compute infrastructure.

The economic analysis sought to value the benefits of the UK's investments in AI compute, across the academic, public, and private sectors, against a hypothetical scenario where these investments would not have taken place (the "counterfactual scenario"). Under this scenario AIRR and related national AI compute capabilities would not have been provided, and users would have to use alternative compute infrastructure to the extent that this is possible.

Estimated benefits are compared to the expected costs of delivery of the AI compute services to evaluate the RoI of the investments.

Given the forward-looking nature of the AI compute investments and the uncertainty around both the magnitude of potential benefits and associated timelines for benefits to materialise, a modelling framework based on Monte Carlo simulations was developed. This framework incorporates key sources of uncertainty to generate a range of possible benefit outcomes, along with their associated likelihoods, enabling a probabilistic assessment of return on investment.

Further the estimates of Return on Investment are benchmarked against RoIs from similar large-scale compute infrastructure in the UK and abroad, providing context on how the UK's AI compute investments compare and helping to validate the robustness of the results.

Aggregate valuation of benefits and RoI

Economic valuation of benefits, and comparison against expected programme costs, suggests a sizeable return on the UK Government's investments in large-scale public AI compute. **The Return on Investment (RoI) is estimated at between £4.2 : £1 and £20.2 : £1, with a central estimate of £6.0 per £1** of public monies invested into AIRR and related national AI compute infrastructure (see Table 1).

This central estimate corresponds to total benefits of approximately £12.1 billion. Sensitivity analysis suggests that a plausible range for total benefits is £12.2 billion (£6.1 : £1) to £27.4 billion (£13.7 : £1), with around 60% of simulations falling within this interval. A wider 90% credible range indicates benefits are likely to lie between £8.5 billion (£4.2 : £1) and £40.5 billion (£20.2 : £1).

The sensitivity analysis results further highlights that benefits could be substantially larger, with a small number (1%) of simulations suggesting benefits in excess of £64.3 billion (£32.2 : £1). These results reflect higher impact scenarios where favourable assumptions are chosen across a combination of inputs. These higher estimates reflect the possibility of higher returns to AI research than suggested by historical literature estimates as well as for the potential for outsized returns associated with high-impact outcomes. Conversely, even the lowest simulation outcomes implied a positive RoI, suggesting that even if conditions were extremely unfavourable public funds would still have delivered value for money.

However, while the sensitivity analysis allows for the possibility of larger returns, potential outcomes delivering extremely large benefits, large-scale transformation of public services or the emergence of globally leading AI companies are difficult to capture and could deliver benefits beyond even those indicated by the higher results from the sensitivity analysis.

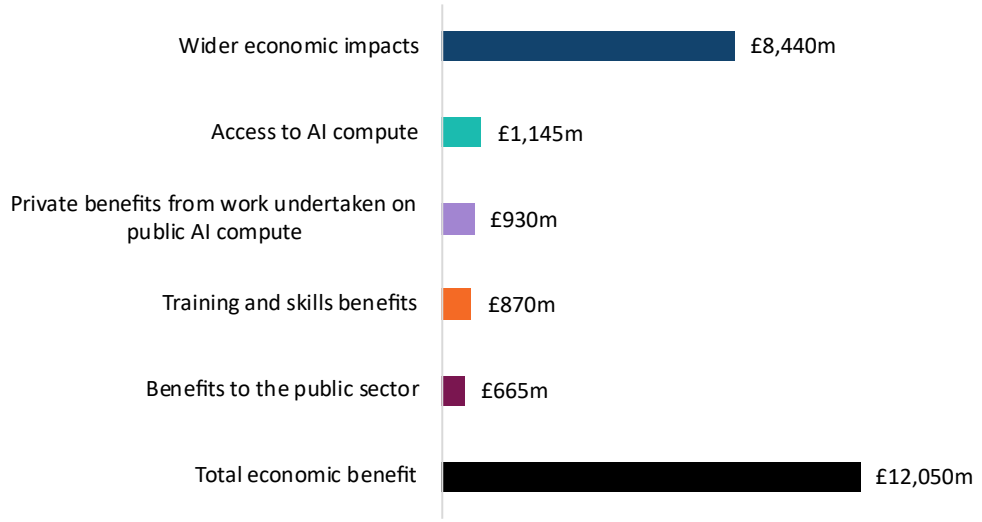
Benefit estimates by type of benefit

The largest group of benefits in the central valuations were **wider economic impacts**, accounting for 70% of benefits.

This group of benefits captures spillovers to the wider economy from AI research undertaken on national AI compute capabilities, which arise through activities such as knowledge transfers and the commercialisation of research. For example, early deep learning research conducted in universities led to breakthroughs such as AlexNet, a deep learning model for image classification that significantly advanced the performance of computer vision systems, which subsequently led to increased use of AI imaging technologies in the medical domain.[53]

The wider economic impacts category also captures downstream adoption of AI innovations by customers of firms developing AI products on AIRR and related public AI compute infrastructure and wider use across sectors. For example, this could include translation of AI research into a commercial product that is then taken up by other firms, thereby supporting wider diffusion of AI capabilities across the economy.

Figure 15 Central valuations by benefit grouping



Note: Figures rounded to the nearest £5 million.

Source: London Economics analysis

Cost savings from access to public AI compute relative to accessing commercial alternatives and **private benefits to users** accounted for 10% and 8% of estimated benefits respectively. The first group of benefits captures cost savings for the share of firms that plausibly would have used commercial compute alternatives such as commercial cloud in the absence of public AI compute. These estimates were informed by the impact survey evidence, which indicates that only a minority of users would have been able to access comparable alternatives, and that replicating AIRR usage through commercial cloud would incur significant costs for users (see Annex A1.1).

Private benefits capture direct outputs that are expected to not have materialised in the absence of public AI compute such as academic publications, the generation of intellectual property, and software products, as well as increased revenue generation among firms using public AI compute,

[53] Singh, P. et al (2020). 3D Deep Learning on Medical Images: A Review. Available at: https://arxiv.org/abs/2004.00218

OFFICIAL

through productivity improvements or commercialisation of AI innovations. Both impact survey and interview evidence reinforce the additionality of these outputs. Impact survey responses indicate that a majority of users would have been unable or only partially able to achieve similar outcomes without AIRR, while interviews provide concrete examples such as research teams producing multiple academic papers and developing reusable models and software, as well as firms creating proprietary datasets and progressing towards commercial applications (see Section 9.2 for more detail).

Training and skills benefits and **benefits to the public sector** accounted for the lowest share of central benefit estimates; 7% and 6% of total benefits, respectively. This is primarily driven by the conservative scope of what is included in the valuation and the exclusion of wider, less easily monetised impacts.

In the case of the public sector, the estimates capture only efficiency savings among public sector staff arising from the adoption of AI innovations developed on public AI compute. For example, this could include AI software developed aimed at reducing administrative support time or increasing operational efficiency among public sector organisations such as government departments, the NHS, or the wider public sector.

Wider benefits from public sector adoption, such as improved healthcare outcomes, welfare benefits, or the value of safer streets as well as highly uncertain scenarios that could deliver substantially larger impacts (such as large-transformation of public service delivery) are not captured by these estimates. However, qualitative evidence from interviews and the impact survey highlights emerging applications in areas such as healthcare data and diagnostic support, illustrating the types of impacts that may materialise more widely.

Training and skills benefits capture the direct value of training provision and the benefits accruing to individuals and associated tax benefits. That is these benefits capture the direct value to individuals, in terms of higher lifetime earnings, due to their increased skills and the associated higher tax income received by the UK Exchequer. Impact survey evidence highlights AI compute's role in developing advanced technical skills through hands-on use of large-scale compute infrastructure and dedicated training, with 69% of respondents using existing AIRR capabilities agreeing that AIRR has improve their level of AI-related skills[54]. However, similar to public sector benefits, wider system-level benefits of skills development, such as spillovers to firms, innovation capacity, or productivity improvements, are not included to avoid double counting with other benefit streams.

In addition, the composition of benefits could vary significantly depending on the type of AI innovations that will be produced on AIRR and related public AI compute and their ultimate uptake in industry, the public sector, or wider society. In particular, the sensitivity analysis shows the potential for both private benefits to users and public sector efficiency outcomes to be substantially larger than indicated by these conservative central estimates (see Figure 13 and Figure 14).

[54] Based on based on 46 survey responses to the question: "To what extent do you agree or disagree with the following statements? Access to the AIRR has…".

OFFICIAL

Comparison to estimated benefits of other evaluations of HPC and science infrastructure

Comparison to other evaluations of public large-scale compute infrastructure shows that, while the central RoI estimate of £6.0 : £1 is in-line with existing evidence, it sits on the conservative side of estimated returns, which are typically reported in the range of around £2 : £1 to £20 : £1.

Particularly noteworthy, the only other evaluation of public AI compute identified, the 2024 evaluation of CSC's High-Performance Computing services,[55] which included the early operational life of LUMI (a powerful European supercomputer based in Finland, used for AI and scientific research), reports substantially higher return figures of €25–€37 return per euro invested. However, the full study has not been published, making it difficult to understand what drives these high valuations.

The sensitivity analysis indicates an aggregate return on investment in the range of £4.2 to £20.2 per £1 invested (covering 90% of simulations). This range is consistent with the evidence reported in comparable national and international evaluations of compute and research infrastructure. While differences in scope, methodology, and timing limit direct comparability, the results provide additional confidence that the estimated returns for AIRR and related public AI compute infrastructure are within a plausible range when set against the wider evidence base.

[55] CSC (2024). SROI of CSC's high-performance computing services studied: €25–37 return per euro invested. Media Release. Available at: https://csc.fi/en/media-release/sroi-of-cscs-high-performance-computing-services-studied-e25-37-return-per-euro-invested

OFFICIAL

PART IV: CONCLUDING REMARKS

13 Concluding remarks

The economic valuation shows a very positive and promising picture of the potential benefits enabled by AIRR and public investments in related national AI compute infrastructure.

Economic valuation of benefits, and comparison against expected costs, suggests a sizeable return on the UK Government's investments in large-scale public AI compute. **The Return on Investment (RoI) is estimated at between £4.2 : £1 and £20.2: £1, with a central estimate of £6.0 per £1** of public monies invested into AIRR and related national AI compute infrastructure (see Table 1). This central estimate corresponds to total benefits of approximately £12.1 billion (discounted to 2025/26) and total operational costs of £2 billion (also discounted to 2025/26), in line with commitments made in the UK Compute Roadmap[56], over the compute capabilities' operational life.

While this central estimate sits on the lower end of the range indicated by comparable studies of public large-scale compute infrastructure (around £2: £1 to £20 : £1), these estimates should be interpreted as conservative, reflecting the use of cautious assumptions throughout given the early-nature of the evaluation and the substantial uncertainty around future impacts. However, the estimated RoI range for AIRR and related national AI compute infrastructure of £4.2 to £20.2 per £1 invested (covering 90% of simulations) is consistent with the evidence reported in comparable national and international evaluations of compute and research infrastructure.

The economic valuation further highlights the substantial potential upside in economic benefits that AI compute could unlock under favourable conditions, with a small number (1%) of simulations suggesting benefits in excess of £64.3 billion (£32.2 : £1). These results reflect the potential of AI to accelerate the scientific discovery process, deliver substantial benefits to UK private sector productivity, and help drive broader changes in public service delivery.

However, while the sensitivity analysis allows for the possibility of larger returns, potential outcomes delivering extremely large benefits such as large-scale transformation of public services or the emergence of globally leading AI companies are difficult to capture and could deliver benefits beyond even those indicated by the higher results from the sensitivity analysis.

Moreover, realisation of these benefits depend on effective implementation by Government including on-going investment not only in compute but also in the broader AI ecosystem to ensure that hardware delivers maximal impact; evidence-based scale-up and allocation of resources; effective engagement of potential users and user support; cross-government coordination and alignment with other national investments; as well as wider Government efforts to ensure positive structural and ecosystem conditions.

Nevertheless, the analysis suggests that even if conditions were extremely unfavourable public funds would still have delivered value for money, with even the lowest estimates showing RoI figures above £1 per £1 invested in AI compute.

Alongside effective Government action, ongoing monitoring of outputs and emerging benefits will be important to ensure the investments remain on track and enable timely adjustments in response to new information or changing conditions.

[56] DSIT (2025). UK Compute Roadmap. Policy paper. Available at: https://www.gov.uk/government/publications/uk-compute-roadmap/uk-compute-roadmap

OFFICIAL

INDEX OF TABLES AND FIGURES

Index of Tables and Figures

Tables

Table 1	Aggregate economic valuation using central assumptions	vi
Table 2	Central valuations by benefit grouping	viii
Table 3	Indicators to measure short-term outputs	21
Table 4	Indicators to measure medium-term outcomes	23
Table 5	Indicators to measure longer-term impacts	25
Table 6	Valuation methods employed for individual benefit streams	35
Table 7	Aggregate estimates of benefits and RoI (£$_{2025/26}$ prices)	38
Table 8	Valuations by benefit stream (£$_{2025/26}$ prices)	39
Table 9	Distribution of total economic benefits and Return on Investment (£$_{2025/26}$ prices)	44
Table 10	Headline findings from evaluations of HPC compute infrastructure	52

Figures

Figure 1	Large-scale compute provision accessible by UK researchers and innovators	7
Figure 2	UK large-scale compute performance share over time against comparator countries	10
Figure 3	UK large-scale compute performance share: November 2024	10
Figure 4	UK large-scale compute performance share: November 2025	10
Figure 5	Investments in public AI compute - logic model	15
Figure 6	AIRR programme sphere of influence	18
Figure 7	Overview of AI compute benefits	27
Figure 8	RoI modelling – overarching modelling chain	33
Figure 9	Sensitivity analysis results: Distribution of total economic benefits	46
Figure 10	Sensitivity analysis results: Deep dive on the upper tail (top 5%) of the distribution of total benefits	47
Figure 11	Sensitivity analysis: Variation of largest benefit groupings	48
Figure 12	Share of wider economic impact in total benefits	49
Figure 13	Share of private benefits in total benefits	49
Figure 14	Share of public sector benefits in total benefits	50

OFFICIAL

Figure 15 Central valuations by benefit grouping 54

OFFICIAL

ANNEXES

Annex 1 Detailed evaluation approaches

A1.1 Direct benefits of access to users

To monetise the direct benefits of access to public AI compute-supported compute, a single strand of analysis was undertaken: the value of compute provision using an avoided cost approach.

This captures the immediate value of access by estimating the cost that users would otherwise incur to obtain equivalent GPU capacity through commercial cloud providers. The approach links the volume of compute supplied to market-based benchmarks, providing a transparent estimate of the value of access.

This benefit is limited to the value of compute itself and excludes downstream impacts arising from its use, which are estimated separately in subsequent sections to avoid double counting.

A1.1.1 Provision of AI compute

This calculation was undertaken separately by GPU generation and user group. The model distinguishes between four GPU generations. In the current analysis, three generations are populated: Generation 1 (benchmarked to NVIDIA Hopper H200 and equivalent GPUs from other manufacturers), Generation 2 (benchmarked to NVIDIA Blackwell B200 and equivalent GPUs from other manufacturers), and Generation 3 (benchmarked to a future NVIDIA Feynman-type or similar equivalent). Generation 4 is included as a placeholder to allow for future extensions but is not populated.

For each GPU generation, the model first estimates the theoretical annual GPU-hours available from the AI compute infrastructure. These GPU-hours are then adjusted to reflect reserved contingent capacity and expected uptake over time (e.g. allocation of GPU hours through projects), producing an estimate of the GPU-hours effectively available for use. The adjusted GPU-hours are allocated across the main user groups included in the model: academia, industry, public sector users, and AISI. For each group, the value of access is estimated by multiplying the allocated GPU-hours by the assumed commercial cloud price per GPU-hour for an equivalent GPU.

In simplified terms, the calculation is:

Value of AI compute access	=	*Adjusted GPU-hours supplied*	x	*Assumed user group allocation share*	x	*Commercial cloud price per GPU-hour*

Commercial cloud prices are specified separately for each GPU generation. In the central case, the starting price assumptions are £6.0 per hour for Generation 1 GPUs, £8.1 per hour for Generation 2 GPUs, and £20.0 per hour for Generation 3 GPUs.[57] These values are intended to represent the cost of procuring equivalent compute capacity through commercial cloud providers, while no value is currently assumed for Generation 4. It should be noted that commercial cloud pricing varies with some providers applying GPU-only pricing and billing separately for compute memory and storage,

[57] Gen 1 GPU prices based on Subramanian (2026) Available at: https://docs.jarvislabs.ai/blog/h200-price. Gen 2 prices based on Adeboye (2025) Available at: https://northflank.com/blog/how-much-does-an-nvidia-b200-gpu-cost. Gen 3 cloud prices (Nvidia Feynman) are not available yet; the research team considers the cost of £20 per hour to be a reference figure.

OFFICIAL

while others apply a bundled price for GPU access, memory and storage. The central case assumes a fully bundled model with no costs beyond the GPU-hour rate. The sensitivity analysis relaxes this assumption, allowing for variation in total costs, both upward (e.g., additional storage or memory charges) and downward (e.g., due to competitive pricing pressures).

A real-terms depreciation over time, to reflect a reduction in prices as newer GPU generations emerge, is also built into the analysis. Annual real-term depreciation in prices of approximately 5%-7.5% are assumed[58]. The sensitivity analysis further allows for lower real-terms price deflation of only 2.5%, reflecting a nominal increase in prices to allow for recent spikes in some cloud-based services[59]. The model projects the corresponding price path over time for each generation and allows for additional shocks to be applied where required.

The resulting annual values are discounted to present values using the HM Treasury Green Book discount rate. An additionality adjustment of 20% based on the impact survey is applied to reflect the share of AI compute that could plausibly have been obtained without AIRR and related national AI compute capabilities through alternative sources, such as commercial cloud providers.[60] This ensures that only the avoided expenditure attributable to public AI compute is monetised.

This benefit stream is limited to the direct value of access to compute itself. It does not capture the downstream benefits enabled by that access, such as research outputs, industry impacts, public sector gains, or wider spillovers, which are estimated separately elsewhere in the model to avoid double counting.

A1.2 Private benefits from work undertaken on public AI compute

To monetise the private benefits arising from work undertaken using public AI compute, six strands of analysis were undertaken: (i) the value of publications resulting from AI projects undertaken on public AI compute, (ii) the value of citations associated with those publications, (iii) the value of software outputs, (iv) the value of intellectual property (IP), (v) the value of spinouts, and (vi) the contribution of industry impacts to UK GVA.

Each of these benefit streams captures a distinct channel through which access to AI compute translates into economic value, reflecting both research outputs and their subsequent commercialisation or application.

In addition to the quantified benefit streams above, AIRR and related public AI compute capabilities are also expected to support collaboration across researchers and institutions, which can enhance both scientific outputs and their wider economic value. However, these effects are not estimated as a standalone benefit. This is because collaborative activity is already reflected within other monetised outcomes. For instance, estimates of publication value incorporate the effort of all contributors, while collaborative research typically leads to outputs of greater impact, which is captured through higher citation values. On this basis, collaboration is not valued separately in order

[58] The actual price evolution is highly uncertain. However, over multi-year horizons, the cost of GPU compute tends to decline on a performance-adjusted basis. Evidence from industry analyses indicates reductions on the order of ~10%per year in nominal terms for prior-generation hardware, implying ~5%-7.5% real annual declines after inflation. (E.g., see McKinsey (2025). The evolution of neoclouds and their next moves. Available at: https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/the-evolution-of-neoclouds-and-their-next-moves)

[59] E.g., see TechRadar (2026). OVH boss predicts major cloud price rises are coming - so get ready. Cloud prices could rise 5-10% next year. Available at: https://www.techradar.com/pro/ovh-boss-predicts-major-cloud-price-rises-are-coming-so-get-ready

[60] Impact survey results are provided in Section 9.4 in the *Baseline evaluation of the AI Research Resource* report. Results indicate that only a minority of users would have been able to access comparable alternatives if they did not access AIRR. Additionally, replicating AIRR usage through commercial cloud would incur significant costs for users.

OFFICIAL

to avoid double counting, and its contribution is instead embedded within the other quantified benefit streams.

A1.2.1 Value of publications

The value of publications was estimated to capture the economic value associated with academic outputs generated through projects supported by public AI compute.

The starting point of the analysis is the estimation of the number of projects enabled by AIRR and related national AI compute capabilities. Project volumes are derived by dividing total GPU-hours allocated to each user group by assumptions on the average GPU-hours required per project (the central scenario assumes an average value of 100,000 hours per project over the lifetime of the compute assets).

Publications are then estimated by applying assumptions on the number of publications generated per project. In the central case, academic users are assumed to produce one publication per project, while public sector, industry and AI safety projects are assumed to generate 0.66 publications per project. These assumptions are based on the impact survey responses and adjusted for optimism bias and low sample size.[61]

A time lag of one year[62] is applied to reflect the delay between the use of compute and the production of research outputs. In the central case, publications are assumed to occur one year after project activity. The total number of publications in each year is therefore given by:

Number of publications	=	*Number of projects*	*x*	*Publications per project (by user group, shifted by the assumed time lag)*

The economic value of publications is then estimated by applying a unit value per publication. This value is informed by a range of evidence from the literature, drawing on three main approaches: production cost and marginal production cost (MPC) methods, publishing market approaches (e.g. open access costs), and willingness-to-pay (WTP) estimates. These estimates are converted into a consistent price base and used to define a plausible range of values. In the central case, a value of £8,168 per publication (in 2025 prices) is applied, with lower and upper bounds of £2,541 and £19,830, respectively, used for sensitivity analysis. It is important to note that these values are not specific to AI publications but was based on Rousseau et al (2021)[63], Morretta et al (2022)[64] and

[61] This is based on responses to the impact survey, where all users, except one industry user, said that their work either has already or is expected to lead to academic outputs such as publications. While some projects may result in multiple publications, the central case assumes only one publication per project to remain conservative. Further, due to the low sample size for non-academic projects, the one respondent from industry who does not expect their work to lead to publications, as well as the differences in focus (academic research vs. commercial outputs / AI products for the public sector) a 25% reduction was applied to the assumed publication figures for non-academic projects to avoid overcounting.

[62] Publication timelines can vary widely. However, one years was judged as a reasonable central estimate in-line with the literature. The sensitivity analysis allowed for lags between no delay (i.e., a few weeks in practice) and up to two years. See, for example, Runde, B. (2021). Time to publish? Turnaround times, acceptance rates, and impact factors of journals in fisheries science. Available at: https://pmc.ncbi.nlm.nih.gov/articles/PMC8459974 or Björk, B. and Solomon, D. (2013). The publishing delay in scholarly peer-reviewed journals. Available at: https://doi.org/10.1016/j.joi.2013.09.001

[63] Rousseau, S., Catalano, G., & Daraio, C. (2021). Can we estimate a monetary value of scientific publications? Research Policy, 50(1).

[64] Morretta, V., Vurchio, D., & Carrazza, S. (2022). The socio-economic value of scientific publications: The case of Earth Observation satellites. Technological Forecasting & Social Change.

OFFICIAL

Giffoni et al (2025).[65] Evidence on the value of AI-specific publications is not available. However, recent research found that AI publications receive more than twice the expected number of citations within their respective subject area[66]. This suggests, the applied publication values are likely to be conservative estimates.

The total value of publications is calculated as:

Value of publications	=	*Total number of publications*	x	*Value per publication*

Annual values are discounted to present values using the HM Treasury Green Book discount rate. An additionality adjustment is applied to reflect the share of research activity that would have taken place in the absence of public AI compute. In the central case, this is set at 20% based on the impact survey[67], ensuring that only the additional research outputs attributable to public AI compute are included.

This benefit captures the direct value associated with producing academic outputs. The wider impact of these publications, such as their influence on subsequent research or economic activity, is captured separately in other benefit streams to avoid double counting.

A1.2.2 Value of citations

The value of citations was estimated to capture the additional impact generated by research outputs produced using public AI compute. Citations are used as a proxy for the scientific relevance and influence of publications, reflecting the extent to which research contributes to subsequent knowledge creation.

The starting point of the analysis is the total number of publications generated through public AI compute-supported projects, as estimated in the previous section. The number of citations is derived by applying an assumption on the average number of citations per publication. In the central case, this is set at 14.9 citations per publication, based on evidence from the literature. It is important to note that this estimate is specific to AI publications in the UK in 2020 based on Al-Marzouqi & Harabi (2024).[68]

A time lag is applied to reflect the delay between publication and the accumulation of citations. In the central case, citations are assumed to arise one year after publication.

The total number of citations in each year is therefore given by:

[65] Giffoni, F., Colnot, L., & Sirtori, E. (2025). The link between large scientific collaboration and productivity. Rethinking how to estimate the monetary value of publications. Scientometrics, 3773-3811.
[66] Erfanmanesh, A. (2025). The citation advantage of AI-related publications. Sci Ed. 2025;48. Available at: https://doi.org/10.36591/SE-4804-04
[67] Additionality was informed by the impact survey evidence on the counterfactual impact of not accessing AIRR (question B7). Responses were mapped to indicative output retention rates (catastrophic = 0%, major = 20%, some = 50%), implying that ~20% of research work would have occurred without AIRR. A central assumption of 1-20% = 80% is adopted. Further detail is provided in Section 9.4 in the Baseline evaluation of the AI Research Resource report.
[68] Al-Marzouqi, A., & Arabi, A. (2024). A comparative analysis of the performance of leading countries in conducting artificial intelligence research. Wiley.

OFFICIAL

Total number of citations	=	*Total number of publications*	x	*Citations per publication (adjusted for time lag)*

To estimate the economic value of citations, a time-based approach is used. This approach values citations based on the time researchers spend engaging with cited publications. Specifically, each citation is assumed to correspond to a fixed amount of reading time, which is then monetised using an estimate of researcher labour costs.

In the central case, each citation is assumed to correspond to 30 minutes of reading time. This time is valued using an estimate of researcher salary (including an uplift associated to research roles related to AI), converted into an hourly rate based on assumed weekly working hours. This results in an estimated value per citation, which is applied uniformly across all citations.

The total value of citations is therefore calculated as:

Value of citations	=	*Total number of citations*	x	*Value per citation*

Annual values are discounted to present values using the HM Treasury Green Book discount rate. An additionality adjustment is applied to reflect the share of research activity that would have taken place in the absence of public AI compute. In the central case, this is set at 20% in-line with the additionality assumption used for publications, ensuring that only the additional citation impact attributable to public AI compute is included.

In addition, a field-weighted citation index (FWCI) is applied to reflect the relative impact of publications compared to global averages in the relevant research fields. This adjustment captures the extent to which public AI compute-enabled research is expected to generate higher-than-average citation impact. In the central case, an FWCI of 2.0 is assumed[69], and the additional citation value associated with above-average impact is estimated separately.

A1.2.3 Value of software

The value of software was estimated to capture the economic value associated with software outputs generated through projects supported by AIRR and related national AI compute infrastructure.

As with publications, the starting point of the analysis is the estimation of the number of projects enabled by public AI compute. Software outputs are then derived by applying assumptions on the number of software products generated per project. In the central case, the academic sector is

[69] Erfanmanesh, A. (2025). The citation advantage of AI-related publications. Sci Ed. 2025;48. Available at: https://doi.org/10.36591/SE-4804-04

OFFICIAL

expected to produce 0.56 software products per project, while the industry is expected to produce 0.75 software products per project (based on results from the impact survey[70]).

A time lag is applied to reflect the delay between the use of compute and the development of software outputs. In the central case, software outputs are assumed to materialise one year after project activity.

The total number of software outputs in each year is therefore given by:

Total number of software outputs	=	*Total number of projects*	x	*Software outputs per project (adjusted for time lag)*

The economic value of software outputs is estimated using a production cost approach, where the value is proxied by the resources required to develop the software. In the central case, each software output is assumed to require 180 days of development time[71], valued using an estimate of researcher salary (using ONS data on salaries in the education sector[72] including an uplift of 40% associated to research roles related to AI[73]). The total value of software is therefore calculated as:

Value of software	=	*Total number of software outputs*	x	*Time to develop software (days)*	x	*Researcher daily cost*

To avoid double counting with other benefit streams, the valuation is restricted to software outputs generated through academic projects. Software outputs from industry projects are assumed to be more closely linked to commercialisation pathways and are therefore captured under other benefit streams.

[70] Specifically, the impact survey asked users whether their work on AIRR has already led, or is it expected to lead to the creation of any commercial products (F2). 56% of academic users expected their work to lead to software products. For industry, all respondents said they expected their work to lead to software products. However, due to the low sample size for industry users (8 responses), an 25% adjustment for optimism bias was applied to this figure. As with all other assumptions, a sensitivity range (of 25%) was applied to these assumptions in the sensitivity analysis.

[71] The exact development time will vary substantially from project to project depending on the complexity of the software developed. However, six months is assumed to be a reasonable central estimate based on various sources. For example:
- Forbes estimates that it usually takes anywhere from three to six months to develop software with simple functionality but six to nine months to make an average-complexity product (Forbes (2022). Software Development Time Estimation: How Long Should It Take To Develop A Product?. Available at: https://www.forbes.com/councils/forbesbusinesscouncil/2022/12/02/software-development-time-estimation-how-long-should-it-take-to-develop-a-product).
- InnoTech puts the average time for custom software development between 4-12 months (InnoTech (n.d.). Software Development: How Long Does It Take To Build A Software?. Available at: https://innotech-vn.com/software-development-timeline.
- KVY Technology suggests medium-time projects require around 3-7 months, while larger scale projects could take from 6 months to several years (KVY Technology (2024). How Long Does It Take to Build Your Software? A Comprehensive Timeframe Guide. Available at: https://kvytechnology.com/blog/software/estimate-development-time)

[72] ONS (2026). Average weekly earnings by industry. Available at: https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/earningsandworkinghours/datasets/averageweeklyearningsbyindustryearn03

[73] ONS industry earnings data captures the education sector as a whole, rather than salaries for academic researchers. Therefore, desk research on current academic job postings and salary bands was undertaken to verify the applicability of this estimate. This exercise suggested a significant underestimation. Therefore, a 20%-60% uplift was applied, to bring salary estimates in line with early-career researchers (e.g., research assistant) salaries on the lower end and research associates at the higher end. A central value in the middle of this range, 40%, was chosen for the central estimates.

OFFICIAL

Annual values are discounted to present values using the HM Treasury Green Book discount rate. An additionality adjustment is applied to reflect the share of software outputs that would have been developed in the absence of public AI compute. In the central case, this is set at 20% based on the impact survey[74], ensuring that only the additional outputs attributable to public AI compute are included.

A1.2.4 Value of IPs

The value of intellectual property (IP) was estimated to capture the economic value associated with patentable outputs generated through projects supported by public AI compute.

As with previous benefit streams, the starting point of the analysis is the estimation of the number of projects enabled by public AI compute. IP outputs are then derived by applying assumptions on the number of IP disclosures generated per project. In the central case, the academic sector is assumed to produce 0.5 IP disclosures per project, while the industry sector is assumed to produce 0.66 IP disclosures per project (based on the impact survey).

A time lag is applied to reflect the delay between the use of compute and the generation of patentable outputs. In the central case, IP outputs are assumed to materialise one year after project activity.

The total number of IP outputs in each year is therefore given by:

Total number of IP outputs	=	*Total number of projects*	x	*IP outputs per project (adjusted for time lag)*

The economic value of IP is estimated using literature-based proxies for the value of patents. Evidence from Bessen (2006)[75] and Serrano (2005)[76] provides a range of estimates for the economic value of patents, which were converted into 2025 prices. These suggest a lower bound of approximately £9,666, a central estimate of £29,631, and an upper estimate of £92,150, reflecting the skewed distribution of patent values. The total value of IP is therefore calculated as:

Value of IP	=	*Total number of IP outputs*	x	*Value per patent*

To avoid double counting with other benefit streams, the valuation is restricted to IP generated through academic projects. IP generated by industry projects is assumed to be more closely linked to commercialisation and firm-level impacts, which are captured separately under industry benefits.

[74] Additionality for research outputs was informed by the impact survey evidence on the counterfactual impact of not accessing AIRR (question B7). Responses were mapped to indicative output retention rates (catastrophic = 0%, major = 20%, some = 50%), implying that ~20% of research work (such as software development) would have occurred without AIRR. A central assumption of 1-20% = 80% is adopted. Further detail is provided in Section 9.4 in the *Baseline evaluation of the AI Research Resource* report.
[75] Bessen, J. (2006). The Value of U.S. Patents by Owner and Patent Characteristics. Boston: Boston University School of Law.
[76] Serrano, C. (2005). The Market for Intellectual Property: Evidence from the Transfer of Patents. University of Minnesota and Federal Reserve Bank of Minneapolis.

OFFICIAL

Annual values are discounted to present values using the HM Treasury Green Book discount rate. An additionality adjustment is applied to reflect the share of IP outputs that would have been generated in the absence of public AI compute. In the central case, this is set at 20% based on the impact survey, ensuring that only the additional IP flowing from public AI compute is included.[77]

A1.2.5 Value of spinouts

The value of spinouts was estimated to capture the economic value associated with new companies created from research undertaken on public AI compute.

The starting point of the analysis is the number of academic projects enabled by public AI compute. Spinouts are then estimated by applying an assumption on the probability that an academic project gives rise to a new spinout company. In the central case, this probability is set at 1.0%[78].

As not all new spinouts are expected to remain active, a survival rate is then applied to estimate the number of spinouts that survive and generate economic value. In the central case, the survival rate is set at 40%[79].

The number of surviving spinouts in each year is therefore given by:

Number of surviving spinouts	=	*Number of academic projects*	x	*Probability of spinout formation*	x	*Spinout survival rate*

The economic value of spinouts is estimated by applying an assumed lifetime GVA contribution per surviving spinout. This is based on literature evidence on the real net GVA generated over the lifetime of a spinout company. Specifically, in the central case, a value of £15.3 million per surviving spinout (in 2025 prices) is used based on CPC (2024) Economic impact assessment of Biotechnology and Biological Sciences Research Council (BBSRC) attributable spinouts.[80]

The total value of spinouts is therefore calculated as:

[77] Additionality for research outputs was informed by the impact survey evidence on the counterfactual impact of not accessing AIRR (question B7). Responses were mapped to indicative output retention rates (catastrophic = 0%, major = 20%, some = 50%), implying that ~20% of research work would have occurred without AIRR. A central assumption of 1-20% = 80% is adopted. Further detail is provided in Section 9.4 in the Baseline evaluation of the AI Research Resource report.

[78] Existing evidence from the recent evaluation of ARCHER2 reports that 2% of ARCHER2 users reported that their work led to the creation of spin-outs. (London Economics (2025). The scientific and economic impacts of ARCHER2. Available at: https://www.ukri.org/wp-content/uploads/2026/03/ARCHER2-Supercomputer-Evaluation-Nov-2025-1.pdf). However, this value may not directly translate to AI compute. Moreover, multiple users could have been involved in the same spin-out and so this figure could overestimate the probability of a spin-out Emerging. Therefore, the central scenario assumed 50% reduction in this figure to account for optimism bias.

[79] This is based on a study by the Royal Academy of Engineering that examined 2,064 UK university spinouts, of which 441 (21.4%) were no longer operational. (Royal Academy of Engineering (2025). Spotlight on Spinouts. Available at: https://raeng.org.uk/media/qropz5hv/spotlight-on-spinouts-2025-final-18_03_25.pdf). However, as the study was undertaken at a point in time, it will not capture spin-outs that ceased to exist following the study. Therefore, the 79% (100%-21%) survival rate was used as an upper bound estimate. The central assumptions adjusts this estimate by 50% to consider optimism bias, resulting in a central estimate of 40% while allowing for wide variation in survival rates. This is in-line with ONS five-year survival rates for businesses, which sits at just under 40% (ONS (2025). Business demography, UK: 2024. Available at: https://www.ons.gov.uk/businessindustryandtrade/business/activitysizeandlocation/bulletins/businessdemography/2024)

[80] Cambridge Policy Consultants. (2024). Economic impact assessment of BBSRC attributable spin-outs. This is in-line with ONS five-year survival rates for businesses, which sits at just under 40% (ONS (2025). Business demography, UK: 2024. Available at: https://www.ons.gov.uk/businessindustryandtrade/business/activitysizeandlocation/bulletins/businessdemography/2024)

OFFICIAL

Value of spinouts	=	*Number of surviving spinouts*	x	*Lifetime GVA value per spinout*

Annual values are discounted to present values using the HM Treasury Green Book discount rate. An adjustment is then applied to account for overlap between the value of spinouts and the wider spillover impacts of research estimated elsewhere in the model. This is intended to avoid double counting, as part of the economic value generated by spinouts may already be captured within the R&D spillover estimates. In the central case, 25% of spinout GVA is assumed to overlap with wider spillover effects and is therefore excluded from the standalone spinout valuation.

A further additionality adjustment is applied to reflect the share of spinouts that would have been created in the absence of public AI compute. In the central case, this is set at 20% in-line with additionality assumptions on research outputs based on the impact survey[81], ensuring that only the additional value attributable to public AI compute is included.

A1.2.6 Contribution of industry impacts to UK output

The contribution of industry impacts to UK output was estimated to capture productivity and revenue gains experienced by firms using AIRR and related national AI compute infrastructure.

The starting point of the analysis is the number of industry projects enabled by public AI compute. These are converted into the number of participating firms using an assumption on the average number of projects undertaken per firm. In the central case, each firm is assumed to undertake 1.5 projects.

The number of participating firms is then distributed across firm sizes using the observed composition of the UK AI sector. In the central case, this corresponds to 74% micro firms, 16% small firms, and 10% medium firms.[82] Large firms are not included, reflecting the assumption that public AI compute is not primarily targeted towards large firms.

For each firm size, baseline turnover assumptions are applied using UK business population statistics from the Department for Business and Trade (2024).[83] In the central case, average turnover is estimated at approximately £0.6 million for micro firms, £3.6 million for small firms, and £25.1 million for medium firms (in 2025 prices). Productivity impacts are then estimated by applying an assumed increase in annual revenue. In the central case, firms are assumed to experience a 33%

[81] Given the early nature of AIRR, it is too early for a sufficiently large number of spin-outs to have formed in order to robustly assess AIRR's additionality. As such, an additionality assumption in-line with that for research outputs was used. Specifically, additionality was informed by the impact survey evidence on the counterfactual impact of not accessing AIRR (question B7). Responses were mapped to indicative output retention rates (catastrophic = 0%, major = 20%, some = 50%), implying that ~20% of research work would have occurred without AIRR. A central assumption of 1-20% = 80% is adopted. Further detail is provided in Section 9.4 in the Baseline evaluation of the AI Research Resource report.

[82] This distribution excludes large firms and is based on DSIT (2025) Artificial Intelligence sector study 2024. Available at: https://www.gov.uk/government/publications/artificial-intelligence-sector-study-2024/artificial-intelligence-sector-study-2024

[83] DBT (2024). Business population estimates for the UK and regions 2024: statistical release. Available at: https://www.gov.uk/government/statistics/business-population-estimates-2024/business-population-estimates-for-the-uk-and-regions-2024-statistical-release

OFFICIAL

increase in revenue, with 44% of firms realising this effect (based on information from previous HPC evaluations[84]). This implies an effective average uplift of 15% per participating firm.

A time lag is applied to reflect the delay between firms accessing public AI compute and the point at which productivity gains materialise. In the central case, impacts begin after 6 years[85] and persist for 3 years[86]. The total increase in firm turnover is therefore calculated by applying the effective revenue uplift to the number of firms affected, accounting for both the delay and duration of impacts. Turnover impacts are then converted into Gross Value Added (GVA) using firm size-specific ratios based on ONS data. In the central case, these are 0.42 for micro firms, 0.35 for small firms, and 0.30 for medium firms.[87]

The total value of industry impacts is therefore calculated as:

Value of industry impacts	=	*Number of firms*	x	*Average turnover per firm*	x	*Effective revenue uplift*	x	*GVA-to-turnover ratio*

Annual values are discounted to present values using the HM Treasury Green Book discount rate. An additionality adjustment is applied to reflect the share of productivity gains that would have occurred in the absence of public AI compute. In the central case, this is set at 20% based on the impact survey report, ensuring that only the additional impact attributable to public AI compute is included.[88]

This benefit captures firm-level productivity improvements associated with the use of AIRR and related national AI compute infrastructure. Wider economy-wide spillovers arising from these impacts are not included here and are instead captured separately under wider economic impacts to avoid double counting.

A1.3 Training and skills benefits

To monetise the training and skills benefits arising from public AI compute-supported activities, two strands of analysis were undertaken: (i) the direct value of training provision, and (ii) the longer-term earnings premium and associated tax benefits accruing to trained users.

[84] Specifically, the 44% figure is an average of the value reported in the evaluation of the Hartree Centre (where 50% of companies reported an increase in profit; Technopolis (2018). Hartree Centre Phase 1 & 2 Baseline Evaluation. Available at: https://www.hartree.stfc.ac.uk/wp-content/uploads/2022/10/Hartree-Centre-Baseline-Impact-Evaluation-Full-Report.pdf) and the 2019 evaluation of EPSRC's investments in HPC (which assumed 38% of companies experienced increased profits; London Economics (2019). The impact of EPSRC's investments in High Performance Computing infrastructure. Available at: https://www.ukri.org/wp-content/uploads/2022/07/EPSRC-050722-ImpactEPSRCInvestmentsHighPerformanceComputingInfrastructure.pdf).

[85] Hirschey and Weygandt (1985) have observed that R&D expenditure demonstrates a deferred return in 5-10 years. (Hirschey and Weygandt (1985). Amortization Policy for Advertising and Research and Development Expenditures. Journal of Accounting Research Vol. 23, No. 1 (Spring, 1985), pp. 326-335) A central value in line with Frontier Economics findings on lag of return of productivity impacts from R&D research (6 years) was assumed. (7 Frontier Economics. (2024). Returns to public R&D: Report for the Department for Science, Innovation and Technology (DSIT)).

[86] The economic literature suggests returns to R&D persist for "a number of years". A range of between 1 and 5 years was assumed as a conservative range in the absence of concrete evidence. A central value of three years was chosen as the baseline.

[87] ONS (2025) Non-financial business economy, UK: employment size-band. Available at: https://www.ons.gov.uk/businessindustryandtrade/business/businessservices/datasets/uknonfinancialbusinesseconomyannualbusinessurveyemploymentsizeband

[88] Impact survey results are provided in Section 9.4 in the *Baseline evaluation of the AI Research Resource* report.

OFFICIAL

These benefit streams capture both the immediate value of training delivered through the programme, as well as the longer-term economic and fiscal impacts associated with improvements in skills and human capital. The direct value of training reflects the market-equivalent cost of training provision, while the earnings premium captures the sustained increase in productivity and wages of individuals benefiting from public AI compute-supported training.

In addition to the quantified benefit streams above, public AI compute is also expected to generate wider skills-related benefits, including knowledge diffusion, improved matching between skills and labour market needs, and spillovers across organisations and sectors. However, these effects are not estimated as standalone benefits. This is because they are already reflected, at least in part, within the earnings premium applied to trained users, which captures the overall improvement in productivity associated with enhanced skills.

A1.3.1 Direct value of provision of training courses

The direct value of training provision captures the economic value of training courses delivered through public AI compute host sites. This is estimated using a market price approach, where the value of training is proxied by the equivalent cost of comparable training in the market.

The starting point of the analysis is the total number of individuals receiving training through public AI compute supported activities. These include learners from academia, industry, the public sector, and AI safety-related activities. The number of learners is derived from the volume of projects supported by public AI compute, combined with assumptions on the number of learners per project across each user group, as well as participation from AISI staff.

Training provision is modelled explicitly through the number of training days delivered each year. In the central case, the programme delivers 10 training days in the first year, growing annually at 10% over time. The number of participants per year is estimated by applying an assumption on the share of total learners attending each training day (5% in the central case).

The total number of training participants in each year is therefore calculated as:

Number of participants	=	*Total number of learners*	x	*Share of learners attending training*	x	*Number of training days*

This approach ensures that participation scales both with the size of the user base and the expansion of training provision over time. The economic value of training is then estimated by applying a market price per participant. In the central case, this is assumed to be £600 per participant[89], based on typical costs of comparable technical training courses (IFF Research, 2020).[90] The total value of training provision is therefore calculated as:

[89] This is broadly consistent with UK benchmarks for the "Digital" training route, which exhibits some of the highest monthly delivery costs (around £503 per learner) and total eligible costs (approximately £9,975), driven by higher staff salaries and the need for specialised instruction. A central assumption of £600 is adopted to reflect variation between more resource-intensive one-to-one technical training and lower-cost, standardised group delivery.

[90] IFF Research. (2020). Cost of delivering apprenticeship standards. The Institute for Apprenticeships and Technical Education and the Education and Skills Funding Agency. London: IFF Research.

OFFICIAL

Value of training provision	=	*Number of participant days*	x	*Market price per participant day*

Annual values are discounted to present values using the HM Treasury Green Book discount rate. This benefit reflects the direct value of training delivered through public AI compute host sites. It does not capture wider productivity or wage effects arising from improved skills, which are considered separately under longer-term skills and labour market impacts to avoid double counting.

A1.3.2 Premium and associated tax benefits of training users

This benefit captures the longer-term labour market gains associated with AI compute training, as well as the associated tax benefits to the Exchequer. It is estimated by applying an earnings premium to public AI compute-trained users and then calculating the additional tax revenues associated with the resulting increase in earnings.

The starting point of the analysis is the total number of public AI compute-trained users across academia, industry, the public sector, and AI safety-related activities. An adjustment is then made for the proportion of learners assumed to remain in the UK labour market, as only these users are assumed to generate benefits to the UK economy. In the central case, 80% of trained users are assumed to remain in the UK.

Average salary assumptions are then applied separately by user group. For academic and AI safety users, salaries are proxied using researcher pay in the education sector with an uplift. For industry users, salaries are based on earnings in the Professional, Scientific and Technical Activities sector.[91] For public sector users, salaries are based on earnings in public administration. A training premium is then applied to these baseline salaries. In the central case, AI compute training is assumed to raise annual earnings by between 5% and 30%, with a central estimate of 11%[92].

The annual skills premium is therefore calculated as:

[91] ONS (2026) EARN03: Average weekly earnings by industry. Available at: https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/earningsandworkinghours/datasets/averageweeklyearningsbyindustryearn03

[92] 5% reflects a conservative estimate equivalent to literature findings of the AI skills premium within the same job title (see Alekseeva, L. et al (2021). The demand for AI skills in the labor market. Labour Economics Volume 71, August 2021, 102002. Available at: https://doi.org/10.1016/j.labeco.2021.102002). The same paper finds a higher premium of 11% for job postings that require AI skills compared to similar postings within the same firm. This mirrors findings by CEPR (2020), which also found an average wage premium of 11% for jobs requiring AI skills compared to similar jobs that do not require knowledge of AI, and a premium as high as 18% for some sectors. (CEPR (2020). The demand for AI skills in the labour market. Available at: https://cepr.org/voxeu/columns/demand-ai-skills-labour-market). More recent studies such as PwC's 2025 Global AI Jobs Barometer have found that the AI skills premium has risen from 7% between 2018-2022 to 27% between 2018-2024, and reaching a high of 56%, in 2024, since GenAI's proliferation. (PwC (2025). Global AI Jobs Barometer. Available at: https://www.pwc.com/gx/en/news-room/press-releases/2025/ai-linked-to-a-fourfold-increase-in-productivity-growth.html) Other evidence suggests skills premia in the region of 30% (e.g., AI Invest (2025). AI Skills Command 30% Salary Premium as Demand Surpasses Supply. Available at: https://www.ainvest.com/news/ai-skills-command-30-salary-premium-demand-surpasses-supply-2508/; Fortune (2025). AI talent comes at a 30% salary premium: 'If you try to play catch up later, this is going to cost you even more'. Available at: https://fortune.com/2025/08/11/ai-talent-salary-premium/) Given these findings, 5% was assumed to be a relatively conservative estimate and was used as the lower bound. 11% was used as the central estimate, while the sensitivity analysis allowed for skills premia as high as 30%. 56% was judged too high as the upper bound due to the significant rise in the premium in recent years as well as given that there is no guarantee that these high premia will persist over the lifetime of a worker's life.

OFFICIAL

Annual AI compute skills premium	=	*Number of trained users remaining in the UK*	x	*Baseline salary*	x	*Skills premium (%)*

To estimate the lifetime value of this premium, the model assumes that the earnings uplift accrues from the average age at which users complete AI compute training until retirement. In the central case, learners are assumed to complete training at age 35 and retire at age 67, implying 32 years in the labour force. Future earnings premia are discounted to present values using the HM Treasury Green Book discount rate.

The premium is estimated separately for academic and AI safety users, industry users, and public sector users, reflecting differences in salary levels across sectors. The resulting values represent the net earnings premium accruing to trained users over their remaining working lives.

In addition to the benefit accruing to individuals themselves, the model estimates the associated tax benefits arising from higher earnings. These are calculated by comparing the tax and contribution profile of an average worker with that of a public AI compute-trained user receiving the training premium. The model includes income tax, employee National Insurance contributions, employer National Insurance contributions, and VAT.

For VAT, the model assumes that a share of additional net income is consumed rather than saved, and that only a proportion of consumption is subject to VAT. In the central case, the model assumes a household savings rate of 10%, implying a marginal propensity to consume of 90%, with 48% of consumption assumed to be standard rated for VAT. This results in 43% of total expenditure being subject to VAT, which is then taxed at the standard rate of 20%. The total tax benefit is therefore calculated as:

Tax benefit	=	*Additional employee income tax and NICs*	+	*Additional employer income tax and NICs*	+	*Additional VAT*

These fiscal benefits are estimated separately for each user group and discounted to present values over the same working-life horizon used for the earnings premium.

The final value for this benefit stream is calculated as the sum of: (i) the total NPV of the net AI compute skills premium accruing to trained users; and (ii) the total NPV of the associated tax benefits. In addition, the model adjusts these values for delivery timing, including the lead time between investment and GPU installation, so that benefits only begin to accrue once training enabled by public AI compute investments can take place.

This benefit captures the longer-term value of skills development through AI compute training. It is separate from the direct value of training provision, which is estimated using a market price approach, and is intended to capture the downstream economic and fiscal returns associated with improved skills.

OFFICIAL

A1.4 Benefits to the public sector

To monetise the benefits to the public sector arising from public AI compute-supported activities, one strand of analysis was undertaken: the contribution of public sector projects to UK output through efficiency gains.

This benefit captures the value of productivity improvements generated by AI-enabled tools and systems developed using public AI compute. These improvements arise from time savings and more efficient use of public sector labour, which are valued as reductions in the cost of delivering public services.

The approach therefore focuses on quantifying the operational efficiency gains associated with deployed public sector projects. Wider benefits, such as improvements in service quality, user experience, or longer-term system transformation, are not estimated as standalone benefits. This is because they are more difficult to robustly quantify and may overlap with measured efficiency gains.

On this basis, the estimated benefit reflects a conservative measure of the contribution of public AI compute-supported public sector projects to UK output, based on observable productivity improvements and avoided labour costs.

A1.4.1 Contribution of public sector projects to UK output

The contribution of public sector projects to UK output captures the efficiency gains generated by public AI compute-supported projects that are successfully deployed as operational tools or systems within the public sector.

The starting point of the analysis is the number of public sector projects enabled through public AI compute. A deployment rate is applied to reflect the share of projects that result in live tools. In the central case, 25% of projects are assumed to be successfully deployed. This assumption is informed by evidence from real-world AI deployment programmes. For example, an initiative launched by Île-de-France in 2024 developed 40 AI use cases, of which 30 progressed to proof of concept and 10 reached full production, implying a deployment rate of 25%. [93]

For each deployed tool, the model estimates the number of public sector staff affected and the associated time savings. In the central case, each deployment is assumed to affect on average 150 full-time equivalent (FTE) staff. This assumption reflects evidence from existing public sector AI deployments, which range from small process-level tools to large-scale implementations affecting thousands of users. For instance, some tools operate at a departmental level or small teams (affecting few staff),[94] while others operate across government (affecting thousands).[95] A conservative central estimate of 150 FTEs is therefore adopted to reflect a typical deployment that is more limited in scope.

Each affected FTE is assumed to save an average of 80 hours per year. This is based on evidence from the Department for Work and Pensions (2026),[96] where AI-enabled tools have been found to

[93] EY. (2025). From ideas to impact: a government leader's guide to responsible AI implementation Turning AI ambition into real-world value.
[94] DWP. (2026). An Evaluation of DWP's Microsoft 365 Copilot Trial. London: Department for Work & Pensions.
[95] DSIT & Rt Hon Peter Kyle MP. (2025). Landmark government trial shows AI could save civil servants nearly 2 weeks a year. Retrieved March 15, 2026, from Gov UK: https://www.gov.uk/government/news/landmark-government-trial-shows-ai-could-save-civil-servants-nearly-2-weeks-a-year
[96] DWP. (2026). An Evaluation of DWP's Microsoft 365 Copilot Trial. London: Department for Work & Pensions.

OFFICIAL

save approximately 19 minutes per working day, equivalent to around 80 hours annually assuming standard working patterns.

The total time savings are therefore calculated as:

Total hours saved	=	*Number of deployed tools*	x	*FTE affected per deployment*	x	*Hours saved per FTE*

These time savings are then monetised using an estimate of the average hourly wage in the public administration sector. In the central case, this is assumed to be £20 per hour, based on ONS salary data[97]. This salary is then uplifted to consider additional non-wage labour costs (30% is assumed aligned with ratios used in other research).[98] This provides a proxy for the value of productivity gains arising from improved efficiency.

To reflect the persistence of benefits over time, savings from each deployment are assumed to accrue over a fixed duration. In the central case, benefits persist for 3 years. A time lag of 1 year is also applied between project completion and the realisation of efficiency gains. The total value of public sector efficiency gains is therefore calculated as:

Value of efficiency gains	=	*Total hours saved*	x	*Hourly wage (including non-wage labour costs)*	x	*Duration of impact (adjusted for time lag)*

Annual values are discounted to present values using the HM Treasury Green Book discount rate. An additionality adjustment of 10% is applied to reflect the share of public sector projects that would have been implemented in the absence of public AI compute. This ensures that only the additional efficiency gains attributable to public AI compute are included.

A1.5 Wider economic impacts

To monetise the wider economic impacts arising from public AI compute-supported activities, two strands of analysis were undertaken: (i) the spillover impact of scientific R&D generated through academic use of public AI compute, and (ii) the wider societal returns associated with private sector gains enabled by public AI compute-supported innovation.

These benefit streams capture the broader effects of AI-enabled research and innovation that extend beyond the immediate users of AIRR and related public AI compute infrastructure. In particular, they reflect how knowledge created through academic research diffuses across the economy, and how firm-level innovation generates additional benefits through spillovers such as supply chain linkages, competition, and knowledge transfer.

[97] ONS (2026). Average weekly earnings by industry. Available at: https://www.ons.gov.uk/employmentandlabourmarket/peopleinwork/earningsandworkinghours/datasets/averageweeklyearningsbyindustryearn03

[98] See for reference DfT TAG Guidance (2022) Available at: https://assets.publishing.service.gov.uk/media/63a32a8d8fa8f539198d9bf3/TAG_Unit_A4.1_-_Social-impact-appraisal_Nov_2022_Accessible_v1.0.pdf.pdf?

OFFICIAL

Wider economic impacts are estimated using established parameters from the R&D literature, including productivity spillover multipliers and ratios of social to private returns (described further below).

A1.5.1 Spillover impact of scientific R&D (academic)

The spillover impact of scientific R&D captures the wider economic benefits generated by academic research enabled through public AI compute. These benefits arise from knowledge diffusion, innovation, and productivity improvements in the wider economy beyond the original research activity.

The starting point of the analysis is the value of compute allocated to academic users. This is used as a proxy for the level of R&D investment enabled by public AI compute, reflecting the cost of undertaking equivalent research using commercial compute resources. A time lag is applied to reflect the delay between research activity and the materialisation of wider economic impacts. In the central case, spillover effects are assumed to begin 6 years after the initial research activity.

The model then applies a spillover multiplier to the implied R&D investment to estimate the wider economic impact. In the central case, a multiplier of 5.7 is used. This is informed by evidence from the literature on returns to public R&D, including Frontier Economics (2025),[99] which finds that public research investment can generate substantial and sustained increases in private sector productivity. Specifically, the central multiplier is based on an assumed 40% uplift in private sector productivity.

While parts of the literature assume that such productivity effects persist indefinitely, the model adopts a more conservative approach by limiting the duration of impacts to 20 years. In addition, consistent with the empirical evidence based on Frontier Economics (2025), a 6-year lag between R&D investment and realised productivity gains is incorporated. The resulting stream of benefits is converted into a present value equivalent to derive the implied multiplier.

This approach remains grounded in the wider R&D literature while being conservative relative to higher-end estimates, such as those reported for research council funding (around 12.7 : 1[100]) or high-performance computing investments (30 : 1[101]).

The total spillover impact is therefore calculated as:

Spillover impact (academic R&D)	=	*Spillover multiplier*	*x*	*Implied R&D investment (adjusted for time lag)*

Annual values are discounted to present values using the HM Treasury Green Book discount rate. An additionality adjustment of 20% based on the impact survey report is applied to reflect the share of research activity that could have accessed alternative compute resources in the absence of public

[99] Frontier Economics. (2025). The value of public R&D. Retrieved March 15, 2026, from GOV UK. Department for Science, Innovation & Technology: https://www.gov.uk/government/publications/the-value-of-public-rd/the-value-of-public-rd

[100] Haskel, J., & Wallis, G. (2010). Public support for innovation, intangible investment and productivity growth in the UK market sector.

[101] IDC (2014). EESI-2 Special Study To Measure And Model How Investments In HPC Can Create Financial ROI And Scientific Innovation In Europe

OFFICIAL

AI compute.[102] This ensures that only the additional spillover benefits attributable to public AI compute are included.

This benefit captures the wider economic impacts of academic research enabled by public AI compute. It is distinct from the direct value of research outputs, which are estimated separately, and reflects the broader productivity gains that arise through knowledge spillovers across the economy.

A1.5.2 Wider societal returns from private gains

This benefit captures the wider societal returns generated by private sector gains enabled through public AI compute-supported activities. The starting point of the analysis is the additionality-adjusted contribution of industry projects to UK output. This represents the private returns accruing to firms through increased productivity, revenues, or cost savings associated with public AI compute-enabled innovation.

To capture wider societal benefits, the model applies an uplift based on the ratio of social to private returns to R&D. This reflects the well-established evidence that the total economic benefits of innovation extend beyond the firm undertaking the activity, through channels such as knowledge diffusion, supply chain effects, and competition. In the central case, a ratio of 3 is applied, consistent with estimates from the literature (e.g. Bloom et al., 2020[103]; Frontier Economics, 2023[104]). This implies that for every £1 of private return, an additional £3 of wider societal benefit is generated.

However, these estimates are not AI specific and cover all forms of R&D. Given the significant body of evidence highlighting the substantial benefits of AI-related innovations, it is plausible that AI innovations would lead to higher societal benefits than indicated by the existing literature evidence. As such, the sensitivity allows for the potential of higher returns (see Annex 2). Nevertheless, due to the limited explicit literature on AI-specific societal returns compared to other forms of R&D, and to remain conservative, the central analysis assumes no additional impacts from AI R&D.

The total wider societal return is therefore calculated as:

Wider societal returns	=	*Private returns*	x	*Ratio of social to private returns*

The resulting values are expressed in present value terms, consistent with the underlying industry contribution estimates. This benefit captures the additional economic value generated beyond the firms directly benefiting from public AI compute-supported activities. It is presented separately from private returns to highlight the broader societal impact of innovation, while avoiding double counting within the overall valuation framework.

[102] Additionality was informed by the impact survey evidence on the counterfactual impact of not accessing AIRR (question B7). Responses were mapped to indicative output retention rates (catastrophic = 0%, major = 20%, some = 50%), implying that ~20% of research work would have occurred without AIRR. A central assumption of 1-20% = 80% is adopted. Further detail is provided in Section 9.4 in the Baseline evaluation of the AI Research Resource report.

[103] Bloom, N., Lucking, B., & Van Reenen, J. (2020). Have R&D spillovers declined in the 21st century? Fiscal Studies. The journal of applied public economics, 4(40), 561-590.

[104] Frontier Economics. (2023). Rate of return to investment in R&D. A report for the Department for Science, Innovation and Technology. London: Frontier Economics.

OFFICIAL

Annex 2 Sensitivity analysis methodology

The central estimates presented in the main report are based on a set of conservative point assumptions for key model parameters. While this provides a clear baseline estimate of the expected economic benefits associated with the AIRR programme and investments in other public AI compute infrastructure, many of the underlying parameters are inherently uncertain. This is particularly the case for assumptions relating to future use of the infrastructure, the scale of outputs and impacts, the timing of benefits, and the magnitude of wider spillover effects.

To reflect this uncertainty, a sensitivity analysis was undertaken using a Monte Carlo simulation approach. The purpose of this analysis was not to change the structure of the cost-benefit model, but to test how the estimated benefits and return on investment vary when key uncertain assumptions are allowed to vary simultaneously within plausible ranges.

A2.1 General approach

The sensitivity analysis used the same underlying model structure as the central scenario. The difference is that, instead of holding selected assumptions fixed at a single value, the model allowed those assumptions to be drawn from probability distributions reflecting uncertainty, both in terms of magnitude as well as robustness, around the central assumptions.

For each simulation, the model generated one draw for each input parameter and then recalculated the full set of results using that combination of assumptions. Repeating this process across a large number of model runs (10,000 runs) produced a distribution of possible outcomes for each benefit stream and for the aggregate benefit-cost metrics. This approach allows uncertainty to be reflected consistently across the model, rather than by varying individual assumptions in isolation.

In addition to varying key parameters, the sensitivity analysis also incorporates a limited number of extensions to the central modelling assumptions in order to capture potential additional sources of impact that are not explicitly included in the baseline estimates due to uncertainty.

In particular, the central scenario estimates spillover benefits from academic research based on the implied funding value of research activity associated with access to public AI compute only (proxy from GPU cloud compute prices). However, this approach does not capture the potential contribution of additional research funding that may be enabled or attracted through public AI compute-supported projects.

To reflect this, the sensitivity analysis includes an additional funding channel based on observed levels of research and innovation funding per project. Specifically, an estimate of the median research funding per project (based on UKRI data)[105] is applied to the number of academic projects supported, generating an estimate of total associated research funding. This additional funding is then incorporated into the spillover calculation using the same multiplier framework as in the central case.

However, it is recognised that, in the counterfactual, a proportion of this research funding may have been allocated to other fields rather than AI-related activity. To account for this, only the incremental value associated with AI research is included. This is captured through an assumption

[105] UKRI (2025). UKRI investment and outputs data 2015 to 2025. Available at: https://www.ukri.org/publications/ukri-investment-and-outputs-data-2015-to-2025/

OFFICIAL

on the relative value of AI research compared to standard research.[106] This parameter is modelled as a distribution in the sensitivity analysis, with a central estimate and plausible lower and upper bounds, with the lower bound implying no additional value. The resulting uplift is applied to the additional research funding to estimate the incremental spillover effects attributable to public AI compute.

A similar approach is applied across other research-related benefit streams. In particular, the assumed uplift associated with AI research is used to adjust estimates of industry contributions, research outputs such as publications, citations, software, intellectual property, spinouts; as well as the academic spillovers from the implied funding value associated with the access to public AI compute. In each case, the sensitivity analysis captures both the baseline level of activity and the potential incremental value associated with AI-enabled research relative to the counterfactual.

The next subsection describes which distributions were used in the analysis and a general overview of how the parameterisation was selected.

A2.2 Selection of distributions and parameterisation

Different distribution types were used depending on the nature of the parameter being modelled:

- **Triangular distributions** were used for many assumptions where a central estimate and plausible lower and upper bounds could be specified. This was the default approach for a large number of operational and valuation assumptions, including project-level outputs, additionality assumptions, time lags, utilisation assumptions, and a number of cost and pricing parameters. This distribution was used where evidence was limited but it was possible to define a reasonable range around the central value (based on the literature or the research team assumptions).
- **Lognormal distributions** were used for a small number of assumptions where uncertainty was expected to be positively skewed, meaning that most values were expected to cluster around more modest outcomes but there remained a possibility of occasional high-value outcomes. This was particularly relevant for assumptions such as the value of spinouts, firm-level revenue effects, and public sector savings, where the downside is naturally bounded but the upside may be more open-ended.
- **Beta distributions** were used for selected assumptions expressed as proportions or probabilities bounded between 0 and 1, such as survival rates or the share of firms expected to realise an impact. Beta distributions are useful in these cases because they respect the natural bounds of the variable and can be shaped to reflect differing degrees of uncertainty.
- **Truncated normal distributions** were used where it was appropriate to allow variation around a central value in a more symmetric way, while still preventing implausible values outside a specified range. This approach was used for selected multiplier assumptions

[106] A large body of literature characterising the sizeable productivity impacts AI could deliver across a wide range of economic sectors exists. As such, it appears likely that AI research would deliver additional economic impacts than those captured by traditional spillover estimates covering all forms of research. However, economic literature on the impact of AI research vs. other forms of research is scarce and significant uncertainty surrounds this assumption. Nevertheless, a number of indicative sources exist. For example, AI patents typically receive 26% more forward citations than non-AI patents (Harvard Business School (2024). The Value of AI Innovations. Available at: https://www.hbs.edu/ris/Publication%20Files/24-069_e5bcc300-d7f3-43b7-af9f-945b5937495e.pdf), while, in climate science, AI in climate patents were associated with around 30–100% more subsequent inventions when counting all technologies (Verendel V. (2023). Tracking artificial intelligence in climate inventions with patent data. Available at: https://www.nature.com/articles/s41558-022-01536-w). However, due to the limited explicit evidence on spillover impacts, no additional impact was assumed in the central case and a conservative range centred around a conservative value of 20% was chosen for the sensitivity analysis.

OFFICIAL

where a central estimate, lower bound, upper bound, and implied spread could be specified.

The parameters of each distribution were set using the central assumptions embedded in the core model, together with additional inputs defining the range or dispersion of outcomes. In practice:

- The **central scenario value** was used as the anchor for the triangular, beta and truncated normal distributions.
- **Minimum and maximum values** were specified where appropriate to define plausible lower- and upper-bound outcomes (maximum values were used as the anchor for lognormal distributions).
- **Dispersion parameters** such as standard deviations were introduced only where required by the distributional form.

For triangular distributions, this involved specifying lower and upper bounds around the central estimate based on available evidence or judgement. For bounded share variables, beta distributions were parameterised to reflect both the central value and an appropriate degree of variation. For truncated normal distributions, lower and upper bounds were used to constrain draws while allowing variation around the centre. For lognormal distributions, parameters were chosen to allow for right-skewed outcomes, such that high values are possible but occur infrequently while the expected value of each draw remained consistent with the central assumptions.

Overall, the choice of distributions and parameterisation was designed to remain consistent with the conservative central case while allowing the simulations to capture plausible variation and, where appropriate, expected upside potential.

A2.3 Simulation process and interpretation

The sensitivity analysis was implemented through Monte Carlo simulation. In each model run, the selected uncertain assumptions were redrawn from their assigned distributions, and the full model was recalculated using those values. The resulting outputs were then recorded for each run.

The simulations generated distributions for the main outputs of interest, including:

- individual benefit streams;
- total monetised economic benefits;
- total public costs where relevant assumptions were varied; and
- implied Return on Investment (RoI) metrics.

The simulation results were then summarised using distributional statistics such as the mean, median, selected percentiles, and minimum and maximum values. These statistics were used to illustrate both the central tendency of the simulated results and the extent of uncertainty around them. Charts were also produced to show the spread of outcomes visually.

The sensitivity analysis should be interpreted as an assessment of the range and distribution of plausible outcomes conditional on the assumptions and parameter ranges used in the model. It does not provide a single revised point estimate. Rather, it complements the conservative central case by showing how estimated benefits and returns may vary when uncertainty across multiple parameters is considered jointly.

OFFICIAL

In this sense, the central scenario remains the main reference point for the appraisal, while the simulation results provide additional context on downside and upside risk, as well as on the relative importance of uncertainty in different parts of the model.

OFFICIAL



LE London Economics

Somerset House, New Wing, Strand,
London, WC2R 1LA, United Kingdom
info@londoneconomics.co.uk
londoneconomics.co.uk
𝕏 @LondonEconomics
+44 (0)20 3701 7700