




From Evidence To Investment:

**AI, Data, AI, Data, and
Advanced Computational
Science in UK Healthcare**

*An Evidence-Based Assessment
of the Case for Government
Investment, National
Implementation, and Long-term
Healthcare Transformation*

In partnership with:





Purpose of This Report

A policy and delivery report informed by a focused parliamentary roundtable, wider evidence, and subsequent Curia analysis to develop a Treasury-facing investment case and practical programme for UK healthcare.

The parliamentary roundtable was convened as a focused expert challenge session within Curia's wider programme of work on AI, data, and quantum-enabled healthcare. Its purpose was not to formally validate a predetermined policy programme or establish consensus on a detailed set of recommendations. Rather, it brought together experienced contributors from healthcare, government, regulation, industry, and research to test whether the overall direction of the work was consistent with developments across the relevant communities, identify significant omissions or areas of disagreement, and challenge assumptions requiring further analysis.

Given the time available, the discussion should be understood as a light-touch expert affirmation and challenge of the direction of travel rather than a formal validation exercise. The session fulfilled this purpose by providing expert challenge to the overall trajectory of the work, broadly affirming the importance of an integrated, problem-led approach while highlighting areas requiring further development – particularly adoption, procurement, regulation, data infrastructure, and cross-government coordination. The subsequent programme design, Treasury analysis, and recommendations in this report have been developed by Curia, drawing on the roundtable and published evidence, international experience, and further analysis. They should not be interpreted as recommendations individually or as collectively endorsed by all roundtable participants.

The primary audience is HM Treasury. Secondary audiences include the Department of Health and Social Care, the Department for Business, Innovation, Science and Trade (DBIST), NHS leaders, research funders, regulators, regional authorities, and organisations responsible for infrastructure and national security.

Acknowledgements

The roundtable was hosted by Rt Hon Lord Kamall and chaired by Professor Gillian Leng CBE. It was sponsored by Cleveland Clinic Philanthropy UK and delivered by Curia. Curia thanks the speakers, contributors, and organisations represented for sharing their experience and challenging the proposed investment case.

Curia retains full editorial control.

Contents

Executive Summary and Introduction	2
Executive Summary	2
The Government Decision	5
Purpose of the Report	6
Scope of the Report	6
Cleveland Clinic London as a UK Translation Partner	7
Methodology	9
Structure of the Report	9
Reading Guide	9
Current NHS Challenges and Workforce	10
A System Under Structural Strain	10
The Workforce Shortfall	10
Demand and the Elective Backlog	11
Stalled Productivity	11
The Financial Envelope and Cost Pressures	11
The Legacy of Under-investment in Technology and Data	12
Summary of the Baseline	12
The Technology and Research Infrastructure Foundation	13
Introduction: the Foundation of the Investment Case	13
A Layered Capability at Different Stages of Maturity	14
Secure Health Data Infrastructure and Research Access	15
High-performance Computing and Advanced Computing	16
Computational Science	18
Artificial Intelligence	18
Quantum Technologies	19
The Specialist Workforce: Capability, Not Implementation Detail	21
From Infrastructure To Impact: the Research-translation Pathway	22
National and Regional Capability	23
The Regulatory Environment	24
Current Adoption Across the Stack	25
Opportunities and Current Limitations	25
Chapter Summary	26

Medical AI Use Cases and Clinical Applications	27
Introduction	27
Diagnostic and Imaging Applications	28
Clinical Decision Support and Patient Care	30
Precision Medicine and Advanced Clinical Applications	30
Operational and Administrative Applications	31
Cross-cutting Findings	31
Chapter Summary	32
International Case Studies and Global Benchmarking	33
Introduction	33
Global Progress in Healthcare AI	34
International Case Studies	34
Common Success Factors	36
Common Challenges	36
Lessons for the UK	37
Chapter Summary	37
Strategic Gap Analysis for the NHS	38
Introduction	38
Current NHS Position	39
Comparison With International Best Practice	39
Strategic Gaps	40
Opportunities for Improvement	41
Strategic Priorities	42
Chapter Summary	42
The Socio-economic Investment Case	43
The Investment Question	43
The Appraisal Framework	44
The Starting Position: Pressures, Costs, and the Funding Context	44
How and Where Value is Created	45
Three Five-year Investment Scenarios	45
A Programme-wide Socio-economic Assessment	48
National Economic Impact and Competitiveness	52
A Blended Funding Model	53
Research Translation and Commercialisation	54
Sensitivity Analysis and Delivery Risks	55
Evidence Gates and Phased Investment	57
Conclusion: Why Invest, and the Bridge to Chapter 8	57

The National Delivery Programme	58
Introduction	58
Strategic Objectives	59
Investment Principles	59
A Phased Five-year National Investment Programme	59
Build, Partner or Access? A Strategic Delivery Model	63
A Blended Public, Philanthropic, and Private Investment Model	65
National and Regional Delivery	67
Governance and Accountability	67
Intellectual Property, Commercialisation, and Public Value	68
Evidence Gates and the Release of Future Funding	68
Government Decision Required	69
The Governance, Assurance, and Public Value Framework for the National Programme	70
Why Governance is Essential for the National Programme	70
The Regulatory and Standards Environment	71
Clinical Safety and Human Oversight	72
Algorithmic Bias, Fairness, and Health Inequalities	72
Transparency and Explainability	72
Data Governance, Privacy, and Cybersecurity	72
Public Trust and Stakeholder Engagement	73
Accountability Across the Programme	73
Independent Validation and the Evidence Gates	74
Lessons From the Cautionary Evidence	74
Procurement, Partnership, and Public Value	74
Workforce Governance	75
Quantum Governance: Sensing, Computing, and Post-quantum Cryptography	75
Research Translation Governance	76
Lifecycle Governance Across the Whole Capability	76
Governance Recommendations	78
Chapter Summary	79
Conclusion and Strategic Recommendations	80
The Central Finding	80
Why Now	81
The Programme Now Available to Government	81
What Success Looks Like	82
The Decision Before Government	83
Departmental Action Plan	84
Conclusion	84
References	85



A Parliamentary roundtable that helped to inform this report was hosted by Lord Kamall, Shadow Minister for Health and Social Care and chaired by Professor Gillian Leng CBE, former Chief Executive of the National Institute for Health and Care Excellence (NICE).

Executive Summary and Introduction

Executive Summary

Key messages

The NHS faces a structural challenge, not a technological one: a workforce shortfall, a large elective backlog, and stalled productivity that no single technology can resolve on its own.

Artificial intelligence, secure health data, high-performance computing, computational science, and quantum technologies form one national capability rather than a set of separate initiatives, and value depends on building it in the right order: foundations first, applications second.

 The programme should be problem-led rather than technology-led: healthcare use cases and measurable outcomes should determine which combination of data, AI, classical compute, computational science, sensing, or quantum capability is used.

 Procurement and implementation are part of the investment model, not downstream details. Demonstrators should begin with a route to procurement and scale and should fund integration, workflow redesign, training, technical support, evaluation, and longitudinal data collection.

 The evidence of benefit is strong in specific applications, including diagnostic imaging, clinical documentation, and genomics, but remains weak, and in some cases actively cautionary, in others. Evidence should lead, and investment should follow. This is a national investment strategy, not another report on whether the NHS should adopt AI. It explains why the Government should invest, over five years, in the infrastructure underpinning NHS transformation, biomedical research, life-sciences growth, regional development, and international competitiveness together.

 No audited figure for economic return yet exists, and this report does not present modelled estimates as audited returns. It sets out low, central, and high investment scenarios using published evidence and transparent assumptions, with later funding released only as independent evidence justifies progression.

 The UK already holds much of what this programme requires: world-leading research cohorts, a rapidly maturing national computing base, a coherent regulatory architecture. The task is coordinated investment, not invention.

 The Government is asked to approve the first phase of a five-year, evidence-gated national programme, not an open-ended commitment.

This report examines what artificial intelligence, secure health data, advanced computing, computational science, and quantum technologies can realistically contribute to the UK, and the terms on which public investment in them can be justified. Prepared for HM Treasury, the Department of Health and Social Care (DHSC), and the Department for Business, Innovation, Science and Trade (DBIST), it treats health-system transformation and wider national economic objectives as a single, connected case. The question for policymakers is no longer whether the UK should invest in artificial intelligence, but how that investment should be sequenced across the national capabilities that enable it to deliver lasting public value for healthcare, scientific research, and economic growth.

The pressures on the health system are structural and well evidenced. The service faces a projected workforce shortfall of between 260,000 and 360,000 staff by the mid-2030s; an elective waiting list of some 7.22 million pathways; and a productivity record that grew by an average of only 0.6 per cent a year over more than two decades and stood some 7.8 per cent below its 2019 level in 2024,

against a committed trajectory of 2.0 per cent a year. These pressures sit within a committed budget of £204.7 billion, constrained by a maintenance backlog and rising demand. This is the baseline against which the value of any investment must be judged: the productivity paradox of roughly 19 per cent more staff since the pandemic treating only around 14 per cent more patients is a direct warning that capability delivers benefit only when it is genuinely put to use.

Delay carries a real opportunity cost; premature, unevaluated deployment carries a real and separately documented delivery risk. The UK's research cohorts, computing base, and genomics capability are genuine international assets, but they are perishable advantages. Comparable economies are investing at scale in computing, data, and quantum capability, and the country's ambition to remain among the world's leading life-sciences economies depends on maintaining pace with that investment. The evidence catalogued in this report shows what happens when deployment outruns evaluation: models that perform well in development and fail on independent validation, and strategies pursued without funded evaluation that produce no demonstrable benefit. The programme recommended here holds both risks in view: it acts now, but through evidence gates rather than ambition.

Sustainable value depends on investing first in national capability, secure health data, research infrastructure, high-performance computing, computational science, a skilled workforce, and disciplined governance with artificial intelligence and quantum technologies deployed within that capability rather than treated as the strategy itself. Value is determined by the terms on which capability is built, procured and deployed, not by the sophistication of any single technology.

Earlier assessments of NHS technology have generally asked a narrower question: should the health service adopt artificial intelligence, and where? This report asks a larger question: does the UK possess, or is it building, the national capability required to convert artificial intelligence, computing, and quantum technologies into health benefit, research leadership, and economic growth, and on what terms should the Government invest in building it? It treats secure data, computing, computational science, research translation, life-sciences growth, and regional development as returns on investment in their own right, alongside NHS transformation.

The evidence supports investment that is disciplined rather than open-ended. The strongest evidence supported by randomised or multi-trial data concerns diagnostic imaging, clinical documentation, and genomics; the evidence is materially weaker, and in several cases actively cautionary, across many operational and predictive applications. High-performance computing and computational science are mature capabilities that can be scaled with confidence now. Quantum sensing has a credible near-term pathway into NHS use; quantum computing does not, and it is treated in this report as a longer-horizon capability warranting proportionate, ring-fenced investment rather than present-day clinical reliance. No audited figure for the economic return of this programme yet exists, and this report identifies that gap rather than filling it with estimation.

Three independently derived investment scenarios are set out here: a minimum credible programme, a central programme, and a higher-ambition programme – each specifying what it would fund, what it would deliver, and the evidence gate that would justify progression to the next programme. The recommended programme is phased: proceeding from foundational investment in data, computing, and the

best-evidenced clinical applications, through evaluated expansion, to a mature, self-sustaining national capability; governed by a programme board with independent scientific, clinical, and economic assurance; funded through a blended model in which government, philanthropy, research councils, universities, industry, and the NHS each contribute what they are best placed to provide; and delivered through a defined national-to-regional architecture connecting shared national infrastructure to integrated care systems, universities, and industry.

The Government Decision

HM Treasury, the DHSC, and the DBIST are asked to approve the first phase of a five-year, evidence-gated national programme Phase 1, of the order of £250–500 million of incremental, health-directed investment over five years, front-loaded and approximately 80 per cent government-funded, together with the mechanism that would size and release its later phases. Specifically, government is asked to approve

- Phase 1 investment in the foundational data, computing, and clinical capabilities set out in this report, of the order of £250–500 million over five years, sized through the Treasury business case.
- The national governance architecture: a programme board with scientific, clinical, and independent assurance functions, accountable throughout the five-year programme.
- A standing evaluation programme sufficient to generate the audited clinical and economic evidence currently missing.
- A Phase 1 portfolio of approximately three to five problem-led demonstrators, each with a named NHS or public-sector user, funded implementation capacity, baseline measures, an independent evaluation plan, and a defined procurement or scaling decision.
- A procurement pathway that allows successful demonstrators to progress, subject to competition and evidence, from grant-supported development into public-sector customer revenue and regional or national adoption.
- Continued investment in the national research infrastructure secure data, computing, and computational science set out in Chapter 3.
- A government-philanthropic partnership framework, formalising the blended-funding and build-partner-access models set out in Chapter 8, with a defined philanthropic-eligible sub-envelope and agreed matching rule for shared research infrastructure, independent evaluation, and early translation or regional seed funding; the illustrative aggregate Phase 1 mix remains approximately 80 per cent government, 15 per cent philanthropic, and 5 per cent private.
- A national research-translation programme connecting discovery to NHS deployment and commercialisation, with the public-value and intellectual-property safeguards set out in Chapter 8.
- The evidence-gated release of future funding, so that later phases (of the order of £1–2.5 billion and £3–5 billion+ in total scale) proceed only as evidence from earlier phases justifies.

Phase 1 is expected to deliver the first audited NHS evidence, released capacity in the best-evidenced pathways, of the order of 500–2,000 net additional jobs and an NHS ten-year net present value in the region of £0–100 million; the full programme, of the order of 5,000–15,000 net additional jobs and £450 million–£2 billion of net additional annual GVA at maturity. All figures are illustrative and gate-dependent; the precise envelope is for the Treasury business case.

Purpose of the Report

This report responds to the need for an integrated national investment strategy that treats healthcare, research infrastructure, and advanced technologies as parts of a single capability rather than as separate policy initiatives. Policymakers are given an accurate account of the current evidence distinguishing what is proven, what is promising, and what remains unproven. That evidence is translated into a practical, proportionate, and affordable programme spanning both health-system transformation and wider national economic objectives: research translation, life-sciences growth, regional economic development, and international competitiveness. HM Treasury is the primary audience, with the DHSC and the DBIST alongside it; the report is constructed on the terms of public-sector appraisal proportionate to cost, honest about uncertainty, and free of claims the evidence cannot support. The ambition here is a sound investment decision across the capability this report describes, not advocacy for any single technology.

The final analysis also incorporates the July 2026 Parliamentary Policy Roundtable convened by Curia at the House of Lords. The session was designed as a focused expert challenge rather than a formal consensus exercise. It broadly affirmed an integrated, problem-led direction of travel while identifying adoption, procurement, regulation, data infrastructure, security, regional delivery, and cross-government coordination as areas requiring further development. The recommendations in this report are Curia's conclusions, informed by that discussion alongside published evidence and further analysis; they should not be interpreted as recommendations individually or collectively endorsed by all roundtable participants.

Scope of the Report

The scope spans the structural pressures facing the health system; the technology and research-infrastructure foundation secure health data, research infrastructure, high-performance computing, computational science, artificial intelligence, quantum sensing, and quantum computing that underpins any national response to them; the clinical and operational evidence across the principal medical specialties; international experience across benchmarked health systems; the NHS's strategic position; the socio-economic investment case, including low, central, and high five-year investment scenarios; a five-year national delivery programme covering governance, funding, the build-partner-access decision, and public value; a risk, governance, and public-value framework; and the conclusions and government decision that follow.

Each technology is assessed according to the maturity of its evidence, not its novelty. Quantum sensing and quantum computing are treated separately throughout because the evidence for each is materially different: quantum sensing has a credible near-term pathway into NHS use, while quantum computing remains at an early, proof-of-concept stage for healthcare applications. The same discipline applies to artificial intelligence, distinguishing mature, deployed capability from developing and cautionary applications. Where evidence is incomplete – most notably, the absence of an audited figure for economic return, and the still-developing standards for cybersecurity and clinical-AI assurance – the report states this plainly rather than substituting an unsupported estimate. It does not present any investment figure as a settled commitment: the ranges in Chapters 7 and 8 are illustrative planning envelopes to be sized and approved through the Treasury business case, and no saving or return is presented as audited where the evidence does not support that claim.

Geographic scope. The operational NHS baseline and most provider-level adoption evidence in Chapters 2, 4, and 6 relate primarily to England because the underlying workforce, waiting-list, funding, and commissioning datasets are England-specific. The research, industrial, advanced-compute, quantum, security, and regional-growth case is UK-wide. Recommendations that depend on NHS England structures should therefore be read as an England delivery model to be adapted with the devolved administrations, while UK-wide infrastructure, research, and industrial policy should be coordinated across all four nations.



Cleveland Clinic London as a UK Translation Partner

Cleveland Clinic London is uniquely well placed to help the UK convert advanced computational and quantum science into tangible patient benefit. Drawing on Cleveland Clinic's international track record of bringing clinical expertise, data, AI, classical computing and quantum capability together around real biomedical questions, it could serve as a linchpin between scientific discovery, clinical evaluation and wider healthcare adoption in the UK.

As a natural partner to the NHS, Cleveland Clinic London could provide a controlled clinical environment in which promising technologies are evaluated at an early stage against defined patient needs, clinical pathways and operational challenges. This would allow developers and clinicians to generate credible evidence on safety, performance, workflow impact and potential patient benefit before technologies progress towards wider deployment.

Its role would not be to suggest that findings from a private healthcare setting are automatically generalisable to the NHS. Rather, Cleveland Clinic London could act as a bridge between industry, private healthcare and NHS adoption: identifying which applications warrant further investigation, testing them under robust clinical governance and working with NHS partners to determine whether the evidence can be replicated across different patient populations, institutions and care settings.

Cleveland Clinic London could contribute:

- A controlled clinical setting for early evaluation and evidence generation.
- A route for connecting technology developers with specialist clinicians and meaningful healthcare problems.
- A bridge between industry, private healthcare, regulators and subsequent NHS evaluation and adoption.
- Internationally connected clinical and research leadership, drawing on Cleveland Clinic's wider expertise and partnerships.
- An anchor institution capable of attracting commercial and philanthropic capital to shared infrastructure, translational research and demonstrator programmes.

This combination of clinical capability, international experience, and convening power gives Cleveland Clinic London a distinctive potential role within a national programme.

Working alongside the NHS, universities, national laboratories, regulators, and industry, it could help the UK move beyond isolated research projects towards a disciplined pathway from scientific capability to evidence, adoption, and improved patient outcomes.



This report was funded in part by Cleveland Clinic Philanthropy and the content contributed to by Cleveland Clinic and Cleveland Clinic London

Methodology

The methodology is grounded throughout in verified evidence: official statistics, independent audit, sector authority data, peer-reviewed research, and the published record of national and international programmes. A consistent discipline runs through it, distinguishing independent evidence from

supplier or programme claims, and proven benefit from promise. Its international analysis rests on comparison across benchmarked health systems; its assessment of national readiness on domestic implementation evidence across health, research, and computing infrastructure; its economic case on appraisal consistent with the HM Treasury Green Book; and its treatment of risk on an evidence-anchored governance analysis. Where a figure or finding could not be verified, or published sources diverge, the report records the gap or the uncertainty rather than resolving it by assumption.

Stakeholder challenge and iteration. The emerging direction was tested at the Curia Parliamentary Policy Roundtable on 13 July 2026, hosted by Lord Kamall and chaired by Professor Gillian Leng CBE. Contributors from healthcare, regulation, government, research, and industry were asked to challenge the five-year investment proposition, the Treasury case, the build-partner-access choices, priority use cases, and the route from research through regulation, procurement, and NHS scale. The roundtable was a light-touch expert challenge, not formal validation; subsequent programme design and recommendations remain Curia's independent analysis.

Structure of the Report

Chapter 2 establishes the baseline, quantifying the workforce, demand, productivity, and financial pressures facing the NHS. Chapter 3 sets out the national technology and research-infrastructure foundation and explains why investment in it must precede large-scale deployment, before Chapter 4 examines the clinical and operational evidence across the medical specialties, organised by the strength of that evidence. Chapter 5 benchmarks the UK against comparable international health systems, while a strategic gap analysis in Chapter 6 identifies where the UK leads and where capability must be developed. The socio-economic investment case follows in Chapter 7, including three independently derived five-year investment scenarios and the wider economic, research, and regional returns the programme can generate. Chapter 8 presents the five-year national delivery programme itself: its phasing, governance, funding model, build-partner-access strategy, and the specific approvals sought from Government. Chapter 9 addresses the governance, assurance, and public-value framework required for safe and defensible deployment across the whole national capability, not artificial intelligence alone. Chapter 10 concludes the report, setting out the strategic recommendations and government decisions that follow from the analysis as a whole.

Reading Guide

Readers seeking the principal policy recommendations may begin with this Executive Summary and the Government Decision set out above, which is restated in full in Chapters 8 and 10. Those examining the economic rationale should refer to Chapter 7, while readers responsible for implementation will find the national delivery programme in Chapter 8 and the governance framework in Chapter 9.



The Parliamentary roundtable brought together experts from across health, technology, and policy to discuss the opportunities and for the delivery of quantum and AI in the UK ahead.

Current NHS Challenges and Workforce

A System Under Structural Strain

Any investment case rests on a clear diagnosis of the problem it is intended to address. The pressures on the National Health Service in England are not transient consequences of the pandemic. They are structural, mutually reinforcing, and, on current trends, set to intensify over the coming decade: demand is rising faster than capacity, the workforce required to meet that demand is projected to fall materially short, and the productivity gains needed to close the gap have not materialised. Together, these conditions define both the scale of the challenge and the threshold that any proposed intervention must clear if it is to represent value for money.

The Workforce Shortfall

Workforce is the binding constraint on NHS capacity. The NHS Long Term Workforce Plan, the first full fifteen-year workforce strategy for the service, projects a shortfall of between 260,000 and 360,000 staff by 2036/37 on an unchanged trajectory and recorded approximately 110,000 vacancies, close to one in ten posts, at the point of publication (NHS England, NHS Long Term Workforce Plan 2023). These projections are modelled rather than observed outturn; the underlying modelling was independently examined by the National Audit Office in 2024, and the headline figures should therefore be read as a planning estimate rather than a certainty. The Plan commits £2.4 billion to additional education and

training places over five years and sets ambitions to roughly double medical school places to around 15,000 per year by 2031/32 and to expand nurse training places by approximately 92 per cent. Its three pillars are Train, Retain, and Reform, and it is intended to be re-issued at least every two years.

Diagnostics is where that strain is most acute and most precisely quantified. The Royal College of Radiologists reports a radiologist shortage of approximately 2,300 posts, equivalent to a shortfall of around 29 to 30 per cent, with the workforce further exposed by an ageing profile: some 39 per cent of radiologists and 44 per cent of clinical oncologists are expected to retire within ten years (Royal College of Radiologists, Clinical Radiology Workforce Census 2025). For the first time, the College also records a pronounced geographic divide, with shortages roughly twice as severe in deprived and rural areas, and notes that recruitment freezes doubled over the course of a single year.

Demand and the Elective Backlog

The elective waiting list is where demand pressure shows most directly. As of April 2026, the list stood at 7.22 million pathways, corresponding to an estimated 6.11 million unique patients.

Approximately 2.53 million of those pathways had been waiting more than 18 weeks, and around 100,000 more than 52 weeks. The median wait was 11.9 weeks, against a pre-pandemic median of 7.2 weeks in April 2019 (NHS England, Referral to Treatment Waiting Times, April 2026). Against this backdrop, the Government has pledged to restore the 92 per cent 18-week constitutional standard by March 2029, supported by interim planning milestones of 65 per cent by March 2026 and 70 per cent by March 2027. Because these figures are published monthly and move accordingly, they are cited here as at the stated date.

The same pressure is visible in cancer performance, where timeliness is clinically decisive. On the most recent baseline, 76.5 per cent of patients met the 28-day Faster Diagnosis Standard, 70.2 per cent met the 62-day referral-to-treatment standard, and 91.7 per cent met the 31-day decision-to-treatment standard. Only 58.7 per cent of cancers were diagnosed at stage I or II, against a longstanding ambition of 75 per cent (NHS England, Cancer Waiting Times; DHSC and NHS England, The National Cancer Plan for England, 2026).

Stalled Productivity

Additional activity alone will not close the gap between demand and capacity; the system also needs to become more productive, and the record here is poor. NHS productivity averaged only 0.6 per cent per year between 1996/97 and 2018/19, fell after 2019 and has not recovered, standing some 7.8 per cent below its 2019 level in 2024 (Office for National Statistics, public service healthcare productivity). The committed trajectory is 2.0 per cent per year to 2029/30. What has become known as the productivity paradox captures the scale of the challenge: the service has deployed roughly 19 per cent more staff since the pandemic while treating only around 14 per cent more patients (National Audit Office, NHS Financial Management and Sustainability, 2024). The distance between recent performance and the committed target is among the most demanding requirements the system faces.

The Financial Envelope and Cost Pressures

The budget within which these pressures must be managed is finite and already heavily committed. The DHSC's departmental budget for England was £204.7 billion in 2024/25, of which around £187 billion was allocated to NHS England, with capital funding of approximately £11.5 billion. The system carried a maintenance backlog of £15.9 billion, a provider deficit of £787 million, and agency spending of around £2.1 billion, against an efficiency target of £9.3 billion. Health accounts for approximately 9 per cent of GDP (DHSC and NHS England, Annual Report and Accounts 2024/25).

Forward pressures are similarly weighty. Analysis by the Health Foundation's REAL Centre attributes the principal cost pressures to demography (£6.8 billion), pay (£6.2 billion), and waiting-time recovery (£7.3 billion). Sustaining the service on its historical basis is estimated to require real-terms funding growth of around 3.8 per cent per year and roughly 4.5 per cent per year over the first five years, equivalent to some £46 billion of additional funding by 2029/30 (Health Foundation REAL Centre, health-care funding projections). Administrative inefficiency compounds the position: of approximately 124.5 million outpatient appointments each year, around 8 million, or 6.4 per cent, are missed, at an estimated annual cost in the order of £1.2 billion.

The Legacy of Under-investment in Technology and Data

Alongside these pressures sits the state of the system's technological and data infrastructure, which the independent record shows to be a longstanding weakness. The National Audit Office found the NHS's digital-transformation track record to be poor: the predecessor National Programme for IT (2002–2011) cost approximately £7.3 billion and was largely unsuccessful, and the “paperless by 2018” target was missed. Trust spending on information technology has run at around 2 per cent of expenditure, against a recommended level of approximately 5 per cent, with named risks including legacy systems, weak interoperability, and digital-skills shortages (National Audit Office, Digital transformation in the NHS, 2020). This history is a necessary reference point. It sets both the realism required of any implementation plan and the cost of continued under-investment – a point this report returns to directly when it sets out the technology foundation in the next chapter.

Summary of the Baseline

The evidence points to a system in which demand is rising, the workforce is projected to fall well short of need, productivity has stalled below both its historical average and its committed trajectory, and the financial headroom to respond is constrained by existing deficits, backlogs, and cost pressures. The diagnostic workforce provides the sharpest illustration: a shortage of roughly 2,300 radiologists, an ageing workforce, and £362 million spent in a single year on outsourcing, overtime, and locum cover to keep pace a sum the Royal College of Radiologists notes would fund more than 3,000 full-time consultant posts (Royal College of Radiologists, Clinical Radiology Workforce Census 2025). At root, the problem the NHS is seeking to solve is structural rather than technological. The contribution of the technologies examined in this report, and the value of any proposed investment, must be judged against that baseline.

Inaction is not neutral, but degrading. On present trends, the waiting list, the workforce gap, and the productivity shortfall all widen, and sustaining the service on its current model alone is estimated to require some £46 billion of additional funding by 2029/30. The relevant question for any intervention is not whether it is affordable in isolation, but whether it can relieve these pressures more cost-effectively than continuing as now.

None of these pressures can be closed by additional activity or additional funding alone while the underlying technology and data foundations remain as constrained as section 2.6 describes. The next chapter sets out that foundation in full – the secure data, computing, computational science, artificial intelligence, and quantum capabilities on which any credible response now depends – and explains why investment in it must come first.



Dr Lara Jehi, Chief Research Information Officer at Cleveland Clinic, joined the roundtable virtually to share insights from Cleveland Clinic's research.

The Technology and Research Infrastructure Foundation

Introduction: the Foundation of the Investment Case

The technologies relevant to this investment case are not a single innovation but a connected system of capabilities: the secure data that describes patients and populations, the computing on which that data is processed, the computational science and artificial intelligence that turn it into insight, the quantum technologies that may in time extend what computing can do, and the specialist workforce and regional ecosystems that make the whole system usable. This chapter establishes that foundation. Its placement, before the clinical evidence (Chapter 4), the economic case (Chapter 7), and the implementation programme (Chapter 8), is deliberate, because the central argument of the report rests on it: value is realised not by acquiring a technology but by building the infrastructure, capability, and governance that allow it to be deployed, evaluated and translated into benefit. Investment in that foundation is therefore not preparatory to the strategy; it is the first and most consequential part of it.

The organising judgement of this chapter, for a Treasury audience, is that readiness differs sharply from one capability to the next, and that public money returns most where it is directed at what is mature and enabling rather than at what is merely promising. Some capabilities of secure data infrastructure, high-performance computing, and computational science are proven and can be scaled now and are the precondition for everything else. Others, chiefly in artificial intelligence, are effective in specific tasks but uneven and, in some cases, cautionary. Quantum computing belongs to a longer horizon and is treated as a strategic bet rather than a present pillar, with the important nearer-term exception of quantum sensing. What follows sets out what is which, on the evidence, and for each layer, why it matters strategically: to NHS transformation, to biomedical research and its translation, to the UK's life-sciences sector and its commercialisation, to regional economic development, and to international competitiveness.

Every statistic in this chapter can be sourced to an official, peer-reviewed or otherwise verifiable publication; supplier and programme claims are distinguished from independent evidence; established evidence, developing evidence, and future potential are kept explicitly separate; and where funding announcements overlap or projections differ, the uncertainty is acknowledged rather than resolved by assumption.

A Layered Capability at Different Stages of Maturity

The foundation is best pictured as a set of layers that build upon one another, together with the workforce that operates them. Secure data infrastructure is the base: without coherent, accessible, and well-governed data, nothing above it can be built safely. High-performance computing and computational science form the engine that converts that data into discovery, and are mature, widely used capabilities. Artificial intelligence is the applied layer that learns from the data using that

computing power, and its maturity is uneven and task dependent. Quantum technologies sit at the frontier. Running through all of these is the specialist workforce in data science, computing, computational biology, and quantum – without which, none of the layers deliver. Comparator countries increasingly describe the ambition as a single, integrated data-computing-AI-quantum ecosystem rather than a set of separate procurements, and this report adopts the same framing, because the returns come from the system, not from any one component.

International evidence supports this framing. Across member countries of the Organisation for Economic Co-operation and Development, artificial intelligence is now used universally in administration, yet national-level clinical scale-up remains limited, reaching only around a tenth of systems even for medical imaging – the most advanced clinical application held back by fragmented data foundations, regulatory uncertainty, and shortfalls in governance and workforce capacity (OECD, 2026). The binding constraint is rarely the algorithm and almost always the data, the infrastructure, and the capacity to deploy and monitor safely. That is the insight around which this chapter, and the investment strategy, are organised.

The table below summarises the maturity of each layer and its strategic role; it is the map for the sections that follow, in which each judgement is supported by the underlying evidence.

Layer/Capability	Maturity	Strategic Role in the Investment Case
Secure health data infrastructure and research cohorts	Foundational; strong national assets, fragmented access	The binding constraint; the precondition for research translation; the first investment priority
High-performance and advanced computing	Mature	The engine; underpins genomics, imaging AI, and simulation; can be scaled now
Computational science (modelling, simulation, multi-omics)	Mature	Accelerates biomedical discovery and drug development; supplies and amplifies AI
Artificial intelligence computer vision/imaging	Mature	Deployed; scale where the evidence is strong
Artificial intelligence clinical documentation (ambient)	Developing	First randomised evidence in 2025; deploy with evaluation
Artificial intelligence predictive analytics	Cautionary	Independent validation required before use
Quantum sensing	Developing (nearer-term)	Government mission commits NHS benefit by 2030; imaging and diagnostics
Quantum computing	Proof-of-concept (medium- to longer-term)	Potential in molecular simulation and drug discovery; not yet clinically relevant
Post-quantum cryptography	Actionable now	Protects long-lived medical data against future decryption
Specialist workforce (AI, data science, HPC, computational biology, quantum)	Developing; a recognised capacity gap	National capability that determines whether every other layer delivers

Secure Health Data Infrastructure and Research Access

Data is the foundation on which every other layer depends, and the UK's position is simultaneously a genuine national asset and its principal constraint. The asset is a single-payer system with a national patient identifier, longitudinal records covering the whole population, and a set of research cohorts without equal internationally. The constraint is that the data estate remains fragmented, interoperability is weak, and access for research and development has historically been slow. The scale of historical under-investment is set out in the independent record: trust spending on information technology has run at around 2 per cent of expenditure against a recommended level of approximately 5 per cent, with named risks including legacy systems, weak interoperability, and digital-skills shortages (National Audit Office, 2020). No advanced capability above this layer can be scaled on data that is not coherent, accessible, and well governed; data infrastructure is therefore the first-order investment, not a supporting detail.

Terminology used in this report: two terms recur throughout the remaining chapters and are used with a single, consistent meaning here. A Secure Data Environment (SDE), the term used in current NHS England and DHSC policy, and a Trusted Research Environment (TRE), the term more commonly used across the wider UK research community and by individual cohort custodians such as UK Biobank, describe the same underlying model: a secure platform to which approved researchers bring their analysis, rather than one from which data is extracted and moved. This report treats the two terms as equivalent and uses “Secure Data Environment (SDE)” as its primary label throughout, reserving “Trusted Research Environment (TRE)” only for the specific instances below where an individual cohort or programme is conventionally described that way in its own published materials.

Secure, interoperable data and Secure Data Environments: the mechanism now accepted across the UK for using health data safely is the SDEs. The Health Data Research Service (HDRS), funded at up to £600 million by government and the Wellcome Trust, is intended to provide a single secure access point to NHS datasets, complemented by SDEs as the route for safe secondary use (DHSC and NHS England, 2026). This model matters strategically because it resolves the apparent tension between protection and access: the international evidence shows that well-designed data governance is an enabler rather than a brake. Finland’s Findata authority delivers both rigorous protection and rapid access, turning around data-access permits in weeks rather than the many months seen elsewhere, and Denmark’s universal identifier and linked registries provide a research substrate built over decades (Government of Finland, 2020).

Research-grade cohorts: the UK’s cohorts are a distinctive national advantage and a direct input to research translation. UK Biobank has followed 500,000 volunteers recruited between 2006 and 2010, combining genomic, biological, lifestyle, and imaging data; more than 20,000 researchers worldwide have used its data, and it is relocating to a new purpose-built facility supported by £127.6 million from the UK Research and Innovation Infrastructure Fund (UK Biobank, 2025). Our Future Health, established with funding from government, charities, and industry, is now the largest programme of its kind: it has recruited well over two million consented participants against an ambition of up to five million, holds biobanked samples for future multi-omic analysis, is embedded within the NHS to support translational research and makes its data available to registered researchers exclusively through a Trusted Research Environment (UK Research and Innovation, 2024; Our Future Health, 2025). Alongside these sit Genomics England and the NHS Genomic Medicine Service, which, since 2018, has offered whole-genome sequencing in routine care (NHS England, 2026). Together, these cohorts constitute research infrastructure of a scale and linkage that few systems can match, and they are the reason the UK is well placed to convert data into discovery, and discovery into life-sciences growth.

Cloud and national research infrastructure: these data assets are increasingly accessed through secure cloud infrastructure and connected to the national compute capability described in the next section, so that research-grade cohorts, SDEs, and high-performance computing function as one system rather than as separate holdings. The strategic significance for this report is that the data layer is not only the enabler of NHS transformation but the foundation of the research-translation pathway set out in section 3.9, and the reason the economic case in Chapter 7 treats data and interoperability as the true rate-limiting investment.

High-performance Computing and Advanced Computing

If data is the foundation, high-performance computing is the engine, and this is a mature capability, not an experimental one. Unlike quantum computing, discussed below, it does not depend on a scientific breakthrough that has not yet occurred; it depends on sustained investment in capacity, access, and skills. For a Treasury audience, this places it in a different category from frontier technologies: it is a capability in which the return is a function of how much is invested and how well access is organised, rather than a bet on future feasibility. It underpins the genomic analysis, the training of medical-imaging models, and the molecular simulation examined in the next section, and it is therefore the common engine beneath both NHS transformation and biomedical research.



Contributors included Dame Chi Onwurah MP, Chair of the House of Commons Science, Innovation and Technology Committee, alongside senior leaders from across healthcare, science, research and technology.

The UK has moved decisively in this area, and the current position is the platform the strategy builds upon. The Government's Compute Roadmap, published in July 2025, set out a long-term plan for national capacity, and the Government has committed over £2 billion to build out compute infrastructure, including over £1 billion to expand the AI Research Resource, and up to £750 million for a new national supercomputer in Edinburgh (DBIST, 2025). Launched in July 2025, the AI Research Resource gives researchers, small businesses, and start-ups access to compute of a scale previously available only to the largest technology companies. Its principal facility, Isambard-AI at the University of Bristol, is the UK's most powerful public compute facility, built around 5,448 Nvidia Grace-Hopper superchips and delivered notably for a strategy concerned with delivery risk in under two years; the Dawn system at the University of Cambridge received a £36 million upgrade in January 2026 to increase its capacity roughly sixfold (UK Research and Innovation, 2026; DBIST, 2026). Isambard-AI also provides the compute underpinning the Government's Sovereign AI programme, linking public capacity directly to commercial innovation. The trajectory can be measured in the Government's own analysis, which records national AI compute rising roughly tenfold between 2024 and 2025 as these facilities came online, against a stated commitment to increase sovereign compute at least twentyfold by 2030 (DBIST, 2026).

Two implications follow. First, the enabling compute for health and life-sciences research is a capability the UK can invest in now with confidence, because the technology is proven and the delivery model has been demonstrated. Second, the design question for the health system is less whether such capability should exist than how it should be accessed, and whether the NHS and its research partners build dedicated capacity, partner with national facilities such as the AI Research Resource or procure capacity as a cloud service. The strategic question is therefore not simply which technologies should be adopted, but which capabilities the UK should build domestically, which should be accessed through shared national infrastructure, and which are best delivered through commercial or international partnerships. Different technologies will require different delivery models, and this distinction becomes a central consideration in the investment programme developed in Chapter 8.

Computational Science

Computational science, the use of large-scale modelling and simulation to solve problems that cannot be solved analytically or by experiment alone, sits between raw computing power and applied artificial intelligence, and it is the capability that most directly connects this investment case to the UK's life-sciences ambitions. It is not a future prospect; it is an established method that already underpins the most advanced work in biomedical research, and its significance for this report is that it both accelerates discovery in its own right and supplies the structured knowledge on which artificial intelligence is trained. A brief survey of its principal biomedical applications makes the point concrete.

Protein structure prediction is the clearest illustration. AlphaFold, developed by Google DeepMind in London, predicts the three-dimensional structure of proteins from their amino-acid sequence – a problem that had resisted solution for half a century. The associated public database, built with the European Molecular Biology Laboratory's European Bioinformatics Institute at Hinxton, now provides predicted structures for over 200 million proteins, used by millions of researchers worldwide; the work was recognised with the 2024 Nobel Prize in Chemistry (Jumper et al., 2021; Nobel Prize in Chemistry, 2024). It is a British-anchored achievement that has already reshaped structural biology and drug design. Around it sit the wider methods of the field: molecular simulation and computational drug discovery, which model how candidate molecules bind to biological targets and so narrow the search before laboratory work begins; systems biology, which models the interactions of genes, proteins, and pathways as networks rather than in isolation; computational pathology and imaging analysis, which apply large-scale computation to digitised tissue and scans; genomics at scale, which is only tractable because the cohorts described above can be analysed on high-performance computing; and the early development of digital twins computational models of an organ or physiological system used to simulate disease progression or treatment response, a promising but still-maturing technique that should be described as developing rather than established.

Two strands of strategic value follow. Computational science accelerates biomedical research directly, compressing work that once took years into days and lowering the cost of discovery – the mechanism by which investment in this layer supports the UK's life-sciences sector and its commercialisation. It is also symbiotic with artificial intelligence: computational methods generate the structured, high-quality data on which models are trained, and AI, in turn, extends what computational science can attempt. This is why data, computing, computational science, and AI are treated as one stack rather than four separate investments in this report, and why the life-sciences returns identified in Chapter 7 depend on the whole stack rather than on AI alone.

Artificial Intelligence

Artificial intelligence is the layer that has attracted the most attention and the most exaggerated claims, and it is treated with particular care here: its maturity is uneven, and it maps closely to the nature of the task. It is greatest where the problem is one of pattern recognition on well-defined data and weakest in operational and predictive applications, where the underlying evidence is frequently non-randomised, single-site, or supplier-reported. Computer vision is the most mature clinical application, deployed in diagnostic imaging; natural language processing and generative methods underpin the automation of clinical documentation, a newer application with the first randomised evidence appearing only in 2025 and effects that appear heterogeneous rather than uniformly positive; and machine learning and deep learning are the underlying methods on which these applications rest. Predictive analytics use such models to anticipate clinical risk or where demand warrants particular caution, because the literature contains prominent instances in which models with strong development performance failed on external validation or produced no benefit in practice.

One finding carries throughout this report and is the most important of them: strong performance in development does not guarantee benefit in deployment. A widely cited evaluation of an artificial intelligence system for breast screening – developed and tested in part on a UK dataset of 25,856 women – demonstrated diagnostic accuracy comparable to expert readers, though on a retrospective basis and with reproducibility concerns, subsequently raised (McKinney et al., 2020) an example that illustrates both the genuine capability of the technology and the need to distinguish a controlled study from a deployed outcome. The candid conclusion matches the prevailing international assessment: artificial intelligence is real and, in specific domains, effective, but its wholesale readiness for clinical deployment should not be overstated, and clinical scale-up is constrained less by the algorithms than by data quality, regulatory clarity, and the capacity of systems to deploy and monitor them safely (OECD, 2026). Strategically, artificial intelligence is the layer at which the health system realises the most immediate operational return in the diagnostic and documentation applications where the evidence is strongest, but only when the data, computing, and workforce beneath it are in place. The detailed clinical evidence for each application is assessed in Chapter 4; this chapter establishes only that AI is one applied layer of a larger system, effective where well-founded and to be treated cautiously where it is not.

Quantum Technologies

Quantum technologies are addressed on a strictly evidence-based footing, distinguishing what is real today from what is a roadmap target. Over-claiming is easier here than anywhere else in the stack, and more damaging to a serious investment case. Two distinct technologies must first be separated, since they are frequently conflated: quantum sensing, which is nearer to use in health, and quantum computing, which is not.

Quantum sensing (nearer-term): quantum sensing uses quantum properties to measure physical quantities, such as magnetic fields, gravity, and time, with far greater sensitivity than conventional instruments. Its health applications are nearer-term and more concrete than those of quantum computing: including quantum-improved optical coherence tomography, magnetoencephalography, histopathology, and developing quantum magnetometry techniques for highly sensitive biomedical imaging and measurement. The UK's National Quantum Strategy includes a mission that, by 2030, every NHS trust will benefit from quantum sensing-enabled solutions (Department for Science, Innovation and Technology (DSIT), 2023). This should be presented as a developing capability with a credible near-term pathway, not as a mature one, but it is the quantum-related area most relevant to NHS transformation within the strategy's timeframe.

Quantum computing (medium- to longer-term): quantum computing is presently in what is termed the noisy intermediate-scale era, in which machines have on the order of tens to a thousand physical qubits and lack the fault tolerance that practical algorithms require; physical error rates remain far above what error-corrected computation needs. A landmark result demonstrated error rates falling as more qubits were added, strong evidence that error-corrected machines are ultimately buildable, but this is not yet a fault-tolerant computer (Google Quantum AI, 2025). In healthcare specifically, the technology remains at the proof-of-concept stage: the most rigorous synthesis of quantum machine learning for digital health found no theoretically provable quantum advantage for the arbitrary-structure datasets typical of medicine and noted that some methods have been matched by classical algorithms (npj Digital Medicine, 2025). No clinically validated quantum-derived therapeutic or diagnostic yet exists.

One UK line of research directly relevant to health data is blind quantum computing, which allows a user to run a computation on a remote quantum processor without that processor or its operator being able to see the data or the computation itself – a potential safeguard for any future health application – though it remains a research capability, not a deployed one.

Realistic applications and their horizons: the credible medium-term application of quantum computing in health is molecular simulation for drug discovery modelling molecular interactions more accurately than classical methods can, but molecules of pharmaceutical interest require computational resources beyond current devices, so practical advantage awaits fault-tolerant machines (Chemical Reviews, 2025). A distinct and nearer-term consideration is security: post-quantum cryptography, the migration of encryption to standards resistant to future quantum attack, matters now because long-lived medical data captured today could be decrypted once capable machines exist. Internationally recognised post-quantum encryption standards were finalised in 2024, and this is treated as a security action in Chapter 9 rather than a computing opportunity. In short, quantum sensing is developing and nearer-term; secure communications and post-quantum cryptography are actionable now; and quantum computing's clinical applications advanced diagnostics, molecular simulation, and drug discovery are medium- to longer-term and contingent on fault tolerance that does not yet exist.



Dame Chi emphasised that successful innovation depends not just on the technology itself, but on the people and processes needed to embed it effectively.

The investment context, and transparency about it: quantum has become an arena of substantial and accelerating public investment internationally, which is relevant here as a competitive context rather than as evidence of clinical readiness. The UK's ten-year National Quantum Strategy commits £2.5 billion from 2024. This builds upon the National Quantum Technologies Programme established in 2013, through which the Government invested approximately £1.4 billion in UK quantum technologies, including the National Quantum Computing Centre at Harwell. That earlier investment forms the foundation of the current strategy and should not be interpreted as additional to the new ten-year commitment. Separately, in March 2026, the Government announced a package of up to £2 billion focused on scaling and procurement, including £1 billion for the "ProQure" programme, a first-of-its-kind advanced procurement intended to inform the future purchase of large-scale quantum computers beyond 2030, alongside over £1 billion for technology development, skills, and facilities across computing, sensing, and networking (DBIST, 2026). In the interests of transparency, the relationship between the ten-year Strategy figure and the March 2026 package is not made fully explicit in the published sources, whether the latter is additional to, or partly encompassed within, the former, and the figures are therefore reported here as distinct announcements rather than summed.

Internationally, the United States has moved comparably: its Department of Commerce announced approximately \$2 billion in incentives to nine quantum companies in 2026, and its Department of Energy launched a programme aiming to demonstrate a scientifically relevant, fault-tolerant quantum computer, with logical qubits numbering in the low hundreds, for chemistry, materials science, and physics by 2028 (US Department of Commerce, 2026; US Department of Energy, 2026). That goal should be read precisely: it is a scientific milestone at a national laboratory, not a commercial or clinical capability, and independent physicists regard even that timeline as highly ambitious, with leading industry roadmaps targeting fault-tolerant systems around 2029. Public projections of the economic value of quantum, such as the frequently cited estimate of a contribution of around £200 billion to the UK economy by 2045, are economy-wide, long-range, and subject to wide uncertainty, and are not used in this report as health-sector estimates. The implication for the strategy is consistent throughout: quantum warrants proportionate, ring-fenced investment in the enabling foundations sensing, post-quantum security, skills, and research access for drug discovery, but is not yet a pillar of the clinical or economic case, and no claim of near-term clinical benefit is made.

Delivery philosophy: own, collaborate, access. For more than a decade, the UK's quantum programme has organised its strategy around a three-way distinction between what the country should own outright, what it should collaborate on with international partners, and what it should simply access through trusted supply chains described by the programme's own strategy leadership as an "own, collaborate, access" (OCA) approach. This is, in substance, the same build, partner, access distinction this report applies to the wider AI, data, and computing capability in Chapter 8, and the quantum programme's decade of practical experience running it is a direct, existing UK precedent for the framework recommended there.

The Specialist Workforce: Capability, Not Implementation Detail

The layers above deliver nothing without the people to build and operate them, and it is a deliberate judgement of this report that specialist skills are part of national capability rather than a downstream implementation issue. A supercomputer without computational scientists, an SDE without data engineers, or a quantum programme without quantum specialists is an asset that cannot be used to its purpose. The skills at stake span data science and data engineering; high-performance and research computing; computational biology and bioinformatics;

artificial-intelligence and machine-learning engineering; and, at the frontier, quantum science and quantum error correction. Several of these are scarce internationally: the pool of specialists in quantum error correction in particular remains very small, relative to the demand that a scaled quantum programme would create, and this is a recognised constraint on the pace at which the frontier layer can advance.


The wider skills challenge has been recognised at the level of national policy, and the scale of the commitment is significant. Through the AI Opportunities Action Plan, the Government has set an objective to support the upskilling of up to 10 million workers over the decade to 2030, reports that around 1 million AI training courses had been delivered by June 2025, ahead of the intended trajectory, and records some 200,000 people studying AI-related higher-education programmes (DBIST, 2025; 2026). These figures describe a genuine national effort in general AI capability; they also indicate the distance still to travel in the most specialised disciplines of computational biology, research computing, and quantum, where the numbers are far smaller and the lead times to train a specialist are considerably longer.

Workforce development is, for the purposes of this report, a foundation layer in its own right, not a training line to be added once the infrastructure is built. It conditions how quickly the data, computing, computational-science, and artificial-intelligence investments can be absorbed and used;

it is a precondition for safe clinical deployment, as Chapter 9 sets out in relation to clinical-AI safety and model-monitoring roles; and it is itself a source of the high-value employment and regional growth that form part of the economic case in Chapter 7 and connect directly to the innovation ecosystems described in section 3.10. Skills should therefore be funded, planned and sequenced as part of the capability itself.

From Infrastructure To Impact: the Research-translation Pathway

The layers described in this chapter are valuable individually, but their strategic purpose is realised only when they operate as a pathway from foundation to patient benefit to economic return. That pathway can be stated simply:



Each stage depends on the one before it. Research-grade cohorts and SDEs (section 3.3) provide the data; high-performance computing and computational science (sections 3.4 and 3.5) turn that data into models and candidate discoveries; artificial intelligence (section 3.6) is developed and refined on that foundation; the resulting tools must then pass independent clinical validation before deployment in the NHS; and successful deployment, in turn, creates the evidence, the intellectual property, and the industrial capability that support commercialisation and economic growth. The UK's distinctive advantage is that it holds strong positions at both ends of this pathway – world-class research

cohorts and computational science at the front, and a single-payer health system able to deploy and evaluate at the back – and the strategic task is to connect them deliberately rather than leave the connection to chance.

The pathway is also why investment in foundational capability is an economic policy rather than simply a technology policy. The same infrastructure that enables NHS transformation also strengthens biomedical research, accelerates research translation, supports commercialisation, enables new UK companies to grow and improves the UK's attractiveness for international research collaboration and inward investment. The value of the investment therefore extends well beyond healthcare delivery alone.

The pathway is established here as a foundation; the chapters that build on it are not pre-empted. The clinical evidence that determines what may pass validation and deployment is assessed in Chapter 4. The economic logic of how value is created and captured along the pathway and why the return depends on the terms of deployment rather than on the technology is developed in Chapter 7. The programme and delivery mechanisms that make the pathway operate at national scale, including the partnership and commercialisation models, are set out in Chapter 8. The purpose here is only to make explicit that infrastructure investment is the entry point to a chain that ends in NHS transformation, life-sciences growth, and economic return.

National and Regional Capability

Capability of this kind is built at two levels, and both matter to the investment case. At the national level sit the shared foundations: the national compute facilities and the Compute Roadmap, the HDRS and SDEs, the national cohorts, and the National Quantum Computing Centre. These are proper national assets, because their scale and cost make duplication wasteful and because their value increases with the breadth of the data and users they serve.

Regionally, capability is realised through innovation ecosystems in which universities, NHS research organisations, national infrastructure, and commercial partners co-locate and collaborate. The UK's AI Growth Zones launched in 2025 and expanded through 2026, with designated zones including Oxfordshire, North and South Wales, the North East, and North Lanarkshire, are the clearest current expression of this approach, and they are explicitly intended to spread the economic benefits of advanced computing beyond the established south-eastern cluster and to attract private investment into regions with available land and energy (DSIT, 2025; 2026). The same approach is evident within quantum technologies, where the UK Quantum Technology Programme supports national research hubs, including programmes led by University College London and the University of Cambridge that focus on healthcare applications of quantum technologies. The existing life-sciences and data clusters, for example around the genomics and data institutions of the Oxford-Cambridge corridor and the research campuses of the north of England, show the same pattern of anchor institutions drawing in academic, clinical, and commercial activity.

International experience offers a useful illustrative model of how such ecosystems can be organised and held to account, though it must be read with care for the difference in health-system structure. The Cleveland Innovation District in the United States brings together anchor hospitals and universities with state government and commercial partners under a single public-private initiative, with shared investment, defined job-creation and economic targets, and interim accountability reporting against those targets including, within it, dedicated advanced-computing and quantum research in partnership with industry (Cleveland Clinic, 2024). It is cited here as one illustration of a model anchor institution, with blended public, philanthropic, and commercial funding, co-located research and clinical capability, and transparent accountability for outcomes relevant to how the UK might structure regional innovation around its own national assets; it is not the only such model internationally, and the delivery and partnership implications are developed further, alongside other comparators, in Chapter 8.



The illustration is instructive not because it should be replicated directly, but because it shows how coordinated investment across government, healthcare providers, universities, research organisations, and industry can create both scientific capability and measurable regional economic development. The strategic lesson for the UK is that national infrastructure delivers greatest value when combined with regionally distributed innovation ecosystems capable of translating research into clinical and commercial outcomes.

The Regulatory Environment

Regulation of these technologies in the UK rests on several complementary regimes, and the architecture is coherent even as it continues to change. Where a tool functions as software as a medical device, it falls within the remit of the Medicines and Healthcare products Regulatory Agency; the National Institute for Health and Care Excellence maintains an Evidence Standards Framework for digital health technologies, updated in 2022 to bring adaptive, data-driven algorithms explicitly into scope and to require evidence of economic impact, real-world performance, and sound data practices (NICE, 2022). Alongside these sit service regulation by the Care Quality Commission, the Digital Technology Assessment Criteria applied before NHS deployment, and the requirements of the Data Protection Act 2018 and UK General Data Protection Regulation.

Two developments extend this picture and bear on the widened scope of the strategy. First, international experience is converging on lifecycle regulation for adaptive artificial intelligence mechanisms that allow a model to be updated in the field without restarting approval, an approach

developed most fully in the United States and offering a ready reference as domestic regulators develop their own. Second, and looking to the frontier layer, the migration to post-quantum cryptography is now a live standards question: internationally recognised post-quantum encryption standards were finalised in 2024, and national cyber-security guidance is matched to them, providing the basis for protecting long-lived medical data. Any quantum-derived clinical tool that eventually appeared would fall under the existing software-as-a-medical-device and cost-effectiveness routes; the governance architecture already built for artificial intelligence would apply; a point developed in Chapter 9.

Current Adoption Across the Stack

NHS adoption has been concentrated and centrally directed rather than uniform, and it is best understood as sitting across the layers described above rather than in artificial intelligence alone. At the applied layer, the principal vehicle has been the AI Diagnostic Fund, established in 2023 with

£21 million for diagnostic tools addressing cancer, stroke, and cardiac conditions and expanded by approximately £30 million in June 2026, with roughly a third of NHS chest radiographs (by volume read, not by number of trusts) now reported to be AI-enabled and artificial intelligence used to support stroke diagnosis, though the associated outcome claims are programme-reported and await independent corroboration (DHSC, 2026). At the data layer, the HDRS and SDEs represent the foundational investment; at the compute layer, the AI Research Resource provides national capacity that health and life-sciences research can draw upon; and at the research layer, the national cohorts and the Genomic Medicine Service are already in operational use. Adoption, nonetheless, remains uneven across the provider setup, with a marked division between digitally advanced trusts and the remainder, a distribution question that bears directly on the regional dimension of the investment case and that the implementation strategy addresses in Chapter 8.

Opportunities and Current Limitations

Real and well-founded opportunities exist in certain domains – particularly in imaging and diagnostics, where the evidence is strongest – and the operational pressures set out in Chapter 2 are most acute and extend, at the foundational layers, to the compute, data, and computational-science capacity on which the wider life-sciences and research economy depends. Set against these, the limitations are equally material and are stated plainly. No audited system-wide figure for artificial-intelligence-driven productivity or cashable savings within the NHS yet exists; such macroeconomic estimates that are available describe the economy as a whole rather than the health system and cannot properly be read across to it (OECD, 2026). The data foundations remain fragmented, interoperability is weak, and the historical under-investment in technology documented earlier continues to constrain deployment (National Audit Office, 2020). The specialist workforce is a recognised capacity gap. And at the frontier, quantum computing is not yet operationally relevant to the delivery of care. A further consideration is the growing energy and environmental demand associated with advanced computing infrastructure. Both large-scale artificial intelligence and developing quantum technologies require increasingly significant computing resources, making energy efficiency and sustainability important factors in long-term infrastructure planning.

Although solid programme-level estimates are not yet available, future investment decisions should consider environmental sustainability alongside performance, security, and economic value. Where the evidence for a given capability is thin, single-site, or supplier-reported, or where published figures overlap or diverge, this report records the gap or the uncertainty rather than filling it with estimates – as applied above to the distinction between mature computing and unproven quantum computing, and to the reporting of quantum funding.

Chapter Summary

The technology and research foundation for this investment case is a connected system of layers at different stages of maturity, operated by a specialist workforce and realised through national and regional ecosystems, and the strategy must be sequenced accordingly. Secure health data infrastructure and the UK's exceptional research cohorts are the binding constraint and the first-order investment. High-performance computing and computational science are a mature engine, already substantially funded, that the country can scale now and that already underpins genuinely world-leading achievements such as protein-structure prediction. Artificial intelligence is an applied layer that is effective in specific, well-evidenced tasks – chiefly imaging, documentation, and genomics – but uneven and, in its predictive applications, actively cautionary. Quantum technologies are a frontier layer: quantum sensing is developing with a credible near-term NHS pathway, post-quantum cryptography is actionable now, and quantum computing is a longer-horizon strategic capability, not a present pillar. Above all, these capabilities reinforce one another: data feeds computation, computation and computational science feed and amplify artificial intelligence, and the whole system, operated by a skilled workforce, generates both clinical benefit and economic value.

The essential conclusions for a Treasury reader are five. The UK must build an integrated ecosystem, not procure isolated technologies. Each layer matters strategically to NHS transformation, biomedical research and translation, life-sciences growth and commercialisation, regional development, and international competitiveness. Investment in the foundational infrastructure must precede large-scale AI deployment, because deployment on weak foundations does not deliver. That infrastructure is the entry point to a research-translation pathway that ends in patient benefit and economic return. And the whole is achievable only if the specialist workforce is treated as part of the capability. This layered, evidence-led foundation is the basis for the clinical evidence assessed in Chapter 4, the socio-economic investment case developed in Chapter 7, and the national implementation programme set out in Chapter 8.



The Parliamentary roundtable discussion considered priority healthcare use cases where emerging technologies could deliver meaningful clinical benefit and demonstrate real-world value.

Medical AI Use Cases and Clinical Applications

Introduction

The clinical evidence for artificial intelligence across the medical specialties is the subject of this chapter. Each application is tested against a consistent standard: what it is intended to do, what independent evidence shows it achieves, how ready the NHS is to adopt it, and what limitations remain. Clinical use cases matter to the investment case because they are where the structural pressures set out in Chapter 2 – workforce shortfall, diagnostic backlog, and stalled productivity – meet the capabilities described in Chapter 3. The strength of the evidence, however, is highly uneven and maps closely to the nature of the task. It is strongest where the problem is one of pattern recognition on well-defined data and weakest in operational and administrative applications, where much of the available evidence is non-randomised, single-site, or supplier-reported. One theme runs throughout: strong performance in development does not guarantee benefit in practice, and the distinction is drawn consistently in what follows.

Application	Evidence	NHS Readiness	Recommended Posture
Breast screening (mammography)	Strong (RCT)	High	Scale now
Endoscopy/Colonoscopy	Strong (RCT + UK trial)	High	Scale now
Genomics variant prioritisation	Strong	High	Scale (decision-support only)
Mental-health self-referral/triage	Strong (multi-site)	High	Scale now
Cardiology ECG (AF detection)	Strong (RCTs)	Moderate	Deploy with evaluation
Chest X-ray interpretation	Moderate (claims programme-reported)	High	Scale with evaluation
Radiotherapy auto-contouring	Moderate (efficiency gain)	High	Deploy with evaluation
Ambient documentation/scribes	Developing (first RCTs 2025)	High	Deploy with evaluation
Stroke imaging	Contested (company-cited)	Deployed	Evaluate outcomes
Pathology	Moderate (no outcome trial)	Moderate	Conditional
Ophthalmology (diabetic retinopathy)	Moderate	High	Evaluate before scaling
Operational/Admin (flow, coding, scheduling)	Weak, single site	Low	Hold
Pharmacy/Nursing support	Weak, few trials	Low	Co-design + evaluate
Predictive analytics/Population health	Cautionary (UK RCT found harm)		Validate before any use
Sepsis prediction (ICU)	Cautionary (failed validation)		Do not deploy without validation
Drug discovery	Early (no AI drug approved)		Strategic/Longer-horizon

Diagnostic and Imaging Applications

Diagnostic imaging is where clinical artificial intelligence is most mature and best evidenced. In radiology, the strongest evidence concerns AI-supported breast screening: a randomised trial reported a substantial reduction, in the order of 44 per cent in the screen-reading workload, while maintaining cancer-detection performance, alongside evidence of fewer interval cancers. National deployment is already under way in the interpretation of chest radiographs, where roughly a third of NHS chest X-rays (by volume read, not by number of trusts, see the trust-level adoption figure in Chapter 6) are now reported to be AI-enabled, though the associated outcome claims remain programme-reported (DHSC, 2026). Readiness is high; the existing diagnostic infrastructure, the acute workforce shortage documented in Chapter 2, and the screening programmes ready to adopt such protocols create a clear operational pull.



The new No.10 administration will need to drive a coordinated approach across government, bringing together the Treasury, Department for Business, Innovation, Science and Trade (DBIST) and wider departments to support the infrastructure needed to turn innovation into measurable healthcare outcomes. (Photo: Simon Dawson / No 10)

Pathology has a strong body of evidence on diagnostic accuracy, multi-reader, multi-institution studies, and first-of-kind regulatory clearances, but no patient-outcome trial has yet been completed (NICE, 2022). The first regulatory authorisation for artificial intelligence in digital pathology was granted in 2021, and further tools have since been cleared. NHS adoption is gated by the completion of laboratory digitisation, into which the Government has invested more than £100 million; readiness is therefore moderate to high but dependent on scanning and storage capacity.

Cardiology carries some of the strongest evidence of any specialty, supported by multiple randomised trials and a systematic review of cardiovascular AI. The clearest application is in electrocardiography, where a single routine reading detects paroxysmal atrial fibrillation before a thromboembolic event in only around two-thirds of cases, and AI-supported interpretation offers a means of improving detection.

In oncology, the best-evidenced application is radiotherapy auto-contouring, which the National Institute for Health and Care Excellence has recommended for use with evidence generation, on the basis that it is quicker than manual contouring and may reduce inter-observer variation. While not being expected to change patient outcomes, it offers an efficiency and consistency gain rather than a clinical one (NICE, 2024). The economics are instructive and are examined in Chapter 7; the Institute records tool costs of between £4 and £50 per plan and time savings reported to it across a range of 3 to 80 minutes per plan, and stresses that further evidence is required for a full cost-benefit assessment (NICE, 2024). All AI-generated contours must be reviewed and edited by a trained professional.

In ophthalmology, autonomous diagnostic systems for diabetic retinopathy represent one of the earliest regulated applications of clinical AI, and screening programmes elsewhere provide a direct analogue for the NHS Diabetic Eye Screening Programme. Readiness is high, though the priority in several of these areas is now solid prospective outcome evaluation rather than further rollout.

Clinical Decision Support and Patient Care

Where artificial intelligence supports clinicians directly in the delivery of care, the principle of human oversight is essential. In emergency medicine, AI for the interpretation of stroke imaging has been deployed in every stroke centre in England since 2024, with the Department reporting use in more than 100,000 suspected-stroke patients (DHSC, 2026). The evidence is presented transparently: favourable single-centre before-and-after studies report marked improvements in functional independence, while more cautious prospective evaluations confirm improved detection and reduced time to decision. Given that some two million neurons are lost per minute in an acute stroke, the time-critical nature of the pathway makes this a high-value application, though the more striking outcome figures are company-cited and warrant independent confirmation.

Ambient documentation systems and AI-supported triage are growing in NHS deployment in primary care, with national guidance published and a supplier registry established for ambient voice technology (NHS England, 2026). Reported productivity effects of the order of two to three minutes saved per consultation are recorded with their provenance, and a validated UK-specific effectiveness figure is not yet available. Much of the published effectiveness evidence originates in other health systems and does not transfer directly.

Mental health provides the strongest evidence for AI-enabled self-referral and triage, in the form of a multi-site evaluation reporting improved access and clear equity gains, with the first such tool holding medical-device status in the UK (NICE, 2022). A large service evaluation across many services reported reduced assessment time and high diagnostic reliability. Readiness is high for triage and self-referral, where NHS deployment is proven, but earlier-stage for AI-delivered therapy, where distinct safety risks arise.

Intensive care is deliberately treated as a cautionary case. A widely deployed commercial sepsis-prediction model failed independent external validation, achieving an area under the curve of only 0.63 and missing approximately two-thirds of cases. The principle it illustrates runs throughout this report: predictive accuracy claimed in development does not equate to benefit in practice, and independent external validation must precede deployment.

In nursing support, named systems are appearing and systematic reviews report benefits, but this is limited by heterogeneity and implementation gaps, with few nursing-specific trials. The evidence shows how much front-line and nursing co-design matters, a workforce that has historically been under-represented in the design of such tools.

Clinical documentation through ambient scribes carries strong evidence, with the first randomised trials published in 2025 and large real-world deployments in place, though effect sizes remain heterogeneous. It addresses directly the documentation burden and associated clinician burnout identified as drivers of workforce pressure.

Precision Medicine and Advanced Clinical Applications

In genomics and precision medicine, the NHS possesses a genuinely world-leading asset in its Genomic Medicine Service, within which, AI-based prioritisation of genetic variants is already embedded; a sub-study reported a 23 per cent increase in new diagnoses (NHS England, 2026). Readiness is high, and the frontier is the safe adoption of newer AI classifiers, which must remain decision-support rather than autonomous, with attention to ancestry and equity bias.

Drug discovery is an early clinical-stage opportunity of considerable strategic significance. Landmark milestones have been achieved, but no AI-designed drug is yet approved, and the evidence base is largely industry-generated. Traditional discovery takes more than a decade and in excess of US\$2

billion per approved drug; that is the baseline against which the potential of computational methods and, prospectively, the quantum-chemistry applications discussed in Chapter 3 should be understood.

In surgery, the strongest evidence concerns computer-aided detection in endoscopy, supported by a meta-analysis of randomised trials and a UK NHS trial that assessed AI-assisted colonoscopy against standard practice with an NHS-perspective economic evaluation. Increased polyp detection is associated with reduced subsequent cancer incidence, though the long-term mortality effect of the technology itself is not yet demonstrated, and endoscopy capacity pressures create a clear pull for adoption.

Pharmacy evidence rests largely on reviews and operational studies, with systematic reviews flagging methodological weakness and dependence on local data, and few randomised trials. The applications prescribing support, clinical screening, and medicines reconciliation are operationally relevant but not yet firmly evidenced.

Operational and Administrative Applications

Operational and administrative applications hold intuitive appeal but rest on a materially weaker evidence base and are presented here with corresponding caution. In medical coding and hospital administration, tools for patient flow, bed management, and discharge planning address genuine operational needs, but outcome evidence is weak, models are frequently localised and non-generalisable, and benefit depends heavily on the human coordination that acts on the predictions; direct outcome evidence has been assessed as unclear.

The underlying problem in appointment scheduling is substantial of approximately 124.5 million outpatient appointments each year, around 8 million (6.4 per cent) are missed, at an estimated cost of £1.2 billion, but the evidence for AI-based approaches is single-trust and programme-reported, with an explicit independent-reproducibility caveat.

In patient triage and remote patient monitoring, applications match strategic priorities such as virtual wards, and some randomised evidence exists for particular components; but studies are frequently small and short-term, the AI contribution can be difficult to distinguish from the surrounding care model, and digital exclusion is a material risk.

Predictive analytics is treated as a cross-cutting cautionary theme. Models frequently discriminate well in development yet fail on external data or fail to change outcomes; the recurring lesson is that prediction does not equal improved outcomes, and the NHS should require independent external validation, calibration, and, ideally, trial-level outcome evidence before deployment. Population health management exemplifies the point, where risk-stratification tools have long been embedded in the NHS, but where the definitive UK randomised trial found no benefit, and the possibility of harm a risk tool that increased emergency admissions without patient or system benefit.

Cross-cutting Findings

Several conclusions appear consistently across the specialties. Randomised or multi-trial evidence of the strongest kind sits in AI-supported mammography, AI-assisted colonoscopy, ambient documentation, and specific cardiology screening applications, complemented by the NHS's world-leading position in genomics. The specialties most ready for adoption today combine credible evidence, a plausible deployment route enabled by existing regulatory mechanisms, and a clear problem to solve: radiology, clinical documentation, radiotherapy auto-contouring in oncology, endoscopy in surgery, genomics, mental-health self-referral, and primary-care documentation.

Those for which deployment can be justified only after further evaluation include the predictive and population-health tools, much of the operational and administrative cluster, and AI-delivered therapy.

Implementation challenges recur across these applications: dependence on data quality and infrastructure, the need for clinician review to guard against transcription and automation error, the difficulty of isolating the technology's contribution from the surrounding care model, and the risk that released time or flagged risk yields no benefit unless a resourced clinical response follows.

Regulatory considerations recur in much the same way: where a tool informs a clinical decision, it is regulated as software as a medical device, subject to oversight by the relevant regulator, the digital assessment criteria for NHS deployment, and data-protection requirements, though a notable gap is that many models embedded within electronic health record systems have had limited independent scrutiny.

Chapter Summary

The clinical evidence points in a clear direction. Artificial intelligence is demonstrably effective in specific, well-defined tasks, chiefly in diagnostic imaging, documentation, and genomics where the NHS is both ready to adopt and under acute operational pressure. Elsewhere, particularly in operational, administrative, and predictive applications, the evidence is thinner, more frequently supplier-reported, and, in the most important cases, actively cautionary. The prudent course is not wholesale adoption but the targeted deployment of the best-evidenced applications, each accompanied by independent evaluation and human oversight.



Experts from quantum technology and the healthcare sector were in attendance at the Parliamentary roundtable.

International Case Studies and Global Benchmarking

Introduction

For a Treasury audience, international benchmarking is not an academic exercise. It is a means of learning, at others' expense, which approaches to funding, regulation, data, and evaluation have generated demonstrable value and which have produced activity without benefit, and of identifying where the UK already leads and should build on existing strengths, rather than importing models designed for structurally different systems.

One caveat governs the comparison throughout. Health systems differ fundamentally in the way they are financed and organised, between tax-funded single-payer models, social-insurance systems, and predominantly private arrangements, and headline figures are therefore rarely like for like. The analysis below concentrates on transferable lessons in strategy, governance, and implementation, and flags explicitly where a country's structure limits what can be read across to the NHS.

Global Progress in Healthcare AI

Cross-country evidence points to a consistent pattern: artificial intelligence has advanced furthest in administration and least in frontline clinical care. Across the member countries of the Organisation for Economic Co-operation and Development, AI is now used universally in administrative functions. National-level clinical scale-up, by contrast, remains limited, reaching only around a tenth of systems even for medical imaging, the most advanced application. It is constrained by fragmented data foundations, regulatory uncertainty, and shortfalls in governance and workforce capacity (OECD, 2026). The World Health Organization's global framework, which most countries meet at an intermediate level of maturity, sets shared objectives around collaboration, national strategy, governance, and people-centred design, providing the normative reference against which national approaches can be read (World Health Organization, 2021).

National approaches within this common trajectory sit along a recognisable spectrum. At one end is a market-led model, exemplified by the United States, in which the private sector drives innovation and the federal government acts as an enabler. At the other is a state-directed model, such as China's, in which centralised data and industrial policy accelerate development. Several systems have deliberately sought a middle path. Finland, for instance, explicitly frames its strategy as a route between the enterprise-driven American and the administration-driven Chinese models, oriented towards a human-centred data economy (Government of Finland, 2020). The UK's own approach, pro-innovation and regulator-led, sits closer to this centre ground.

International Case Studies

Country	Model	Key Lesson	Relevance to NHS
United States	Market-led	Most mature adaptive-AI regulation (1,250+ devices)	Reference for MHRA adaptive regulation
Japan	Hybrid	Adaptive approval without restarting	Reference for adaptive regulation
South Korea	Single-payer	Reimbursement, not approval, drives adoption	Maps to NHS tariff/commissioning
Germany	Social insurance	First digital reimbursement pathway; low uptake shows payment ≠ adoption	Outcome-linked pricing, with caution
France	Social insurance	Reimbursement route little used; Health Data Hub	Data-hub model; reimbursement caution
Denmark	Tax-funded	Universal ID + linked registries over decades	Long-run data-foundation model
Finland	Tax-funded	Fast and safe data access (~51-day permits)	Data-governance-as-enabler model
Singapore		AI embedded in existing screening pathway	Direct analogue for Diabetic Eye Screening
Canada	Tax-funded	Federally funded scribes with built-in evaluation	National funding + local delivery + evaluation
Australia	Tax-funded	Strategy without funded evaluation → no evidence of benefit	The failure mode to avoid
China	State-directed	Fast, but data governance incompatible with UK norms	Contrast on speed vs data rights
UK	Tax-funded	Home/baseline case being assessed	



Dr Michael Cuthbert, Director of the National Quantum Computing Centre (NQCC), was among the senior experts contributing to the roundtable discussion.

It is more instructive to group the 12 benchmarked countries by the distinctive lesson each contributes than to survey each system in isolation.

The primacy of regulation as an enabler of adaptive AI is demonstrated by a first group of countries. The United States operates the most mature regulatory architecture for clinical artificial intelligence, having authorised more than 1,250 AI-enabled medical devices, around three-quarters of them in radiology. This is supported by predetermined change-control plans and good machine-learning practice for adaptive algorithms (Food and Drug Administration, 2025). Japan and South Korea have developed comparable mechanisms, allowing AI software to improve in the field without restarting approval; South Korea had authorised 139 such products by 2022, concentrated in radiology, pathology, and electrocardiography (Ministry of Food and Drug Safety, 2022). The lesson for the UK is that regulatory design for adaptive AI is a solved problem elsewhere, and these frameworks offer a direct reference for the Medicines and Healthcare products Regulatory Agency.

A second group of countries speaks to reimbursement and adoption support, separate from regulatory approval. Germany built the world's first national reimbursement pathway for digital health, moving to outcome-linked pricing from 2026. Yet only around 12 per cent of eligible clinicians have prescribed a reimbursed application, evidence that a payment pathway alone does not drive uptake (Federal Institute for Drugs and Medical Devices, 2024). France established a temporary reimbursement route for new digital devices but has seen minimal use to date. South Korea's single-payer experience yields the clearest statement of the principle: reimbursement design, not approval, drives adoption. For a tax-funded system, this maps directly onto NHS tariff and commissioning decisions.

A third group makes the case for long-term investment in data infrastructure. Denmark's universal personal identifier and complete linked registries provide a substrate for health AI built over decades. Finland's Findata authority delivers both rigorous protection and rapid access, turning around data-access permits in approximately 51 days, against timescales measured in many months elsewhere (Government of Finland, 2020). France's Health Data Hub enables population-scale health-technology assessment across records covering tens of millions of individuals. The common lesson is that data governance, well designed, is an enabler rather than merely a constraint.

A fourth group shows the value of embedding AI within existing pathways and evaluating it as it goes. Singapore's national diabetic-retinopathy screening programme succeeded by slotting an AI system into an established pathway with human oversight, achieving strong performance and freeing graders to focus on complex cases. This provides a direct analogue for the NHS Diabetic Eye Screening Programme (Ministry of Health Singapore, 2023). Canada's federally funded scribe programme, extending to as many as 10,000 primary-care clinicians with evaluation built in, exemplifies the pairing of national funding with local delivery and designed evaluation.

Two cases serve principally as cautions. Australia's own government review concluded that there was a lack of evidence to support the benefits of AI in its health system, precisely because strategy and guidance were not accompanied by funded deployment and evaluation, the clearest illustration of the failure mode this programme is designed to avoid. China demonstrates the speed achievable when centralised data, industrial policy, and expedited approval match, but at the cost of a data-governance model incompatible with UK norms. Its usefulness is therefore chiefly as a contrast, illustrating the trade-off between deployment speed and data-rights protection.

Common Success Factors

Synthesised across the 12 cases, the factors consistently associated with successful national adoption are strikingly uniform. The clearest successes pair central funding with local delivery and built-in evaluation, as in Canada and the UK's own diagnostic fund. They embed AI within an existing clinical or screening pathway rather than bolting it on, with human oversight retained by design, as in Singapore. They treat data governance as an enabler, investing in the infrastructure and permit systems that make safe access fast, as in Finland and Denmark. They rest on mature, adaptive regulation of the kind developed in the United States, Japan, and South Korea. Sustained engagement with clinicians, patients, and the wider public was a further common characteristic among successful implementations: countries that embedded transparency, clinical involvement, and clear communication into AI deployment generally experienced stronger acceptance and greater confidence in the use of AI-enabled healthcare services, so that public trust arose not as a secondary consideration but as a fundamental enabler of long-term implementation success. And they recognise that reimbursement and active adoption support, not approval alone, determine whether tools are actually used.

Common Challenges

The barriers recur with equal consistency. Data fragmentation and weak interoperability are near-universal constraints, cited from Australia to Japan, and are the principal reason clinical scale-up lags administrative use across the OECD (OECD, 2026). Regulatory uncertainty persists even in advanced systems, with several countries not expecting AI-specific legislation for years. Workforce capability and adoption recur repeatedly, with Germany's low prescriber uptake being the clearest illustration that capability and engagement matter more than the availability of a pathway. Evidence generation is perhaps the most consistent gap of all: across almost every country examined, system-level outcome, productivity, and cost-savings figures are not published in comparable form, and this report records those absences plainly rather than estimating around them. To these are added persistent concerns over procurement complexity, ethics, algorithmic bias, and privacy, the last of which, China's model throws into sharp relief.

Lessons for the UK

For the NHS, the international experience produces a realistic and encouraging assessment. Several practices can be adopted directly. The pairing of a national fund with local delivery and mandatory evaluation is already partly established through the AI Diagnostic Fund and should be extended.

Following Singapore's example, the embedding of AI within existing screening pathways is directly applicable to the Diabetic Eye Screening Programme and to the chest-radiograph rollout already under way. The adaptive-regulation mechanisms of the United States, Japan, and South Korea offer a ready reference for domestic regulators, and Germany's move to outcome-linked pricing informs how the NHS might reimburse AI tools on the basis of demonstrated value.

Other lessons require adaptation to the UK's scale and federated structure. The data-governance architectures of Finland and Denmark are portable in principle, and the Health Data Research Service and SDEs represent the NHS's response, but the UK's larger and more fragmented data estate makes coordination harder than in a small, centrally registered population. Reimbursement lessons from single-payer comparators translate more cleanly than those from insurance-based systems, whose financing and administrative-cost structures cannot be read across.

In specific respects, the UK already demonstrates international leadership. Its genomics infrastructure is genuinely world-leading, its centrally directed diagnostic-AI programme is more coordinated than the diffuse American model, and its data-foundation investments position it well against comparators. The recurring international lesson is that roadmaps and principles do not, by themselves, generate outcome evidence, which is precisely the gap that mandatory, funded evaluation (set out in Chapters 7 and 8) is designed to close.

Chapter Summary

A clear conclusion arises from the international record. The health systems that have realised value from artificial intelligence have done so not through the ambition of their strategies but through the rigour of their implementation: funding paired with evaluation, technology embedded in existing pathways, data governed to enable rather than obstruct, and reimbursement designed to drive adoption. Those that have fallen short, Australia chief among them, have generally produced strategy and guidance without funded, evaluated deployment. The UK holds genuine advantages in genomics, data infrastructure, and coordinated diagnostic programmes, and faces the same barriers of fragmentation, workforce capability, and evidence generation as its peers. This is consistent with the OECD's own conclusion: across high-performing healthcare systems, implementation capability, not technological availability, is the defining factor in translating AI innovation into measurable improvements in health service delivery.



Strategic Gap Analysis for the NHS

Introduction

The distance between the NHS's current position and the demonstrated potential of artificial intelligence, and what must be in place before deployment at national scale can be justified, is the subject of this chapter. It is deliberately sequenced before the economic case, because the size and nature of the implementation gaps determine both the cost and the risk of any investment.

Identifying these gaps is not criticism but realism: the systems that realised value did so by understanding their own constraints and addressing them in order, while those that proceeded on strategy alone generated activity without benefit. The analysis distinguishes between foundations that are mature and areas that require development.

Current NHS Position

The NHS starts from a position of genuine and specific strength. It possesses a national imaging base, with picture archiving and communication systems established across trusts, and electronic health records in widespread use, providing the substrate on which the highest-value diagnostic applications depend (National Audit Office, 2020). It holds a world-leading asset in its Genomic Medicine Service, within which AI-based prioritisation is already embedded, and it has demonstrated the capacity for coordinated national deployment, having rolled out artificial intelligence for stroke imaging to every stroke centre in England and extended AI-assisted chest radiography to around half of trusts (by number of trusts offering it, not by volume of radiographs read, see the volume-based figure in Chapters 3 and 4) (DHSC, 2026).

Policy direction is clear and supportive. The Ten-Year Health Plan identifies artificial intelligence as one of five transformative technologies and requires NHS bodies to reserve at least 3 per cent of annual expenditure for one-off transformation investment (DHSC and NHS England, 2026). The principal delivery vehicle, the AI Diagnostic Fund, has committed

£21 million from 2023 and a further £30 million in 2026, complemented by substantial investment in data infrastructure through the HDRS, funded at up to £600 million (DHSC, 2026). The regulatory foundations are coherent, resting on a pro-innovation, regulator-led model that coordinates medical-device regulation, digital assessment criteria, and cost-effectiveness appraisal. These are not the attributes of a system starting from a standing position; they are the assets on which further progress can realistically be built.

Comparison With International Best Practice

Against the leading systems examined in Chapter 5, the NHS's standing is mixed rather than uniform, and the comparison is instructive precisely because it varies by dimension.

Dimension	NHS Position	Verdict
National strategy	Coherent, centrally coordinated (more so than the US)	Leading
Clinical implementation	National deployment achieved in imaging	Developing deployment has outpaced evaluation
Data infrastructure	NHS number as identifier, but fragmented estate; HDRS a response, not a resolution	Behind (vs Denmark, Finland)
Workforce capability	Same adoption challenges as peers, plus structural shortfall	Developing
Funding	Central-fund model fits with best practice	Developing early-stage scale
Regulation	Coherent framework	Developing no adaptive-AI lifecycle mechanism yet
Evaluation	Strong exemplars in commissioned studies	Developing not yet a universal condition
Procurement	Traditional models	Behind not yet outcome-linked (vs Germany)
Public trust	Strong information-governance regime	Leading (contrast to China)
Innovation ecosystem	World-leading genomics and life sciences	Leading

The “Leading” assessment above reflects the coherence of the proposed strategy rather than the durability of cross-government coordination. Previous UK technology programmes have shown that coordination can weaken as institutional priorities and leadership change. The governance architecture proposed in Chapters 8 and 9 has therefore been designed around enduring functions and defined responsibilities rather than individual office holders, reducing the risk that strategic coordination diminishes over time.

Strategic Gaps

From this comparison, a set of principal gaps can be drawn out, and it matters to distinguish those that reflect mature foundations requiring extension from those that require more substantial development.

Economic evaluation is the most consequential gap. No audited figure for cashable savings from artificial intelligence exists for any specialty within the NHS, and closing this through NHS-specific economic evaluation feeding the appraisal framework examined in the next chapter is the precondition for scale-up rather than an afterthought (NHS England, 2026). Closely related is the gap in clinical validation: the cautionary cases set out in Chapter 4 demonstrate that models must be independently validated in situ before deployment, and that a resourced clinical response must accompany any predictive tool.



Attendees heard contributions from Professor Sir Peter Knight, Chair of the UK Quantum Technologies Strategic Advisory Board, who shared his perspective on the opportunities presented by quantum technologies.

The true rate-limiting constraint sits in digital infrastructure and interoperability. Maturity is uneven between digitally advanced and lagging trusts, the estate is burdened by a £15.9 billion maintenance backlog, and interoperability remains weak – the international evidence from Denmark and Finland showing precisely the payoff that long-run data-foundation investment can yield (National Audit Office, 2020; OECD, 2026). Data quality is a corollary, since the more advanced applications cannot be scaled on fragmented or poorly governed data.

Regional capability is a further strategic gap. Although the UK possesses internationally recognised research strengths across multiple regions, technology infrastructure, innovation investment, and adoption capacity remain unevenly distributed. Addressing these differences will require stronger coordination between national infrastructure, regional innovation ecosystems, and place-based investment programmes, reinforcing the role of AI Growth Zones and the regional capability model described in Chapter 3.

On workforce, development is needed in AI literacy across clinical staff, new roles such as clinical AI safety officers and model-monitoring analysts, and genuine clinical co-design, the international evidence being clear that neglecting front-line and nursing involvement is a recognised failure mode. Governance foundations exist but require standardisation, notably through a national model registry and a single accountable assurance function. Without these, oversight will fragment across multiple bodies, with no clear owner. Procurement needs to move towards outcome-based, anti-lock-in models. Public trust, while supported by a strong governance regime, requires active maintenance through transparency and engagement. National implementation consistency also remains uneven, with adoption varying markedly across the provider setup.

Some of these imaging infrastructures, including genomics, the regulatory architecture, and the central-fund mechanism, are mature foundations to be extended. The rest – economic evaluation, interoperability, standardised assurance, and workforce capability – require development before national deployment can proceed safely.

Within government itself, there is a related capability gap. Although the UK possesses internationally recognised expertise across universities, research organisations, and industry, comparatively fewer technical specialists are embedded within government departments and public institutions responsible for procurement, regulation, and strategic investment. Strengthening internal technical capability therefore matters not only for policy development but also for intelligent commissioning, independent evaluation, and long-term programme oversight.

Opportunities for Improvement

For the NHS, the realistic opportunities lie in building on existing strengths while importing the specific practices that international experience has validated. The most immediate is to extend the established central-fund-plus-local-delivery model so that every funded deployment carries a mandatory, independent evaluation converting the NHS's strong evaluation exemplars into universal practice and beginning to close the economic-evidence gap that constrains everything else.

Alongside this sits a case for prioritising the applications for which the evidence is already strongest and the infrastructure most ready: radiology represents the highest-value entry point, given the national imaging base and the acute workforce pressure, followed by clinical documentation, radiotherapy contouring, endoscopy, genomics, mental-health triage, and primary-care documentation.

Data infrastructure and interoperability should be treated as the foundational investment they demonstrably are, with the HDRS and SDEs emulating the fast, safe data access achieved elsewhere. Explicit adaptive-AI regulation should be adopted, drawing directly on the mature mechanisms available internationally, and the assessment-plus-validation gate before NHS deployment should be standardised. Recent developments within the UK's regulatory environment also indicate movement towards more adaptive lifecycle regulation, reinforcing the direction proposed in this report and discussed further in Chapter 9. Procurement, in turn, should be reformed towards outcome-linked pricing and away from supplier lock-in. Sustained public and clinical engagement should accompany each stage of implementation: international experience demonstrates that transparency, clinician

involvement, and clear communication strengthen public confidence and support the long-term adoption of AI-enabled healthcare services. Each of these is a practical improvement grounded in existing capability rather than an aspiration, and each addresses a gap identified above.

Strategic Priorities

These opportunities resolve into a logical sequence of priorities that must precede safe deployment at national scale

1. **Establish the data and interoperability foundations:** the true rate-limiters; no advanced application can scale without them.
2. **Build workforce capability and governance in parallel:** AI literacy, clinical-safety roles, co-design, a national model registry, and a single accountable assurance function.
3. **Make independent evaluation mandatory on every deployment:** generating the first audited NHS outcome and cost evidence the economic case requires.
4. **Scale the proven applications within existing pathways:** gated by regulation, not wholesale adoption.
5. **Modernise regulation and procurement** so adaptive tools can be adopted and paid for on the basis of demonstrated value.

The ordering reflects the international lesson that foundations, governance, and evaluation must be in place before, not after, adoption.

Chapter Summary

The strategic gap analysis points to a measured and constructive conclusion. The NHS holds genuine international advantages in imaging infrastructure, genomics, coordinated national programmes, and a coherent regulatory model that place it well to realise value from artificial intelligence. It also faces material and specific gaps, chief among them being the absence of audited economic evidence, the weakness of data foundations and interoperability, and the need to standardise clinical validation, governance, and workforce capability. The distance between current position and demonstrated potential is therefore neither trivial nor insurmountable: it is a defined set of foundations to be built in a logical order, consistent with the implementation-capability lesson set out at the close of Chapter 5. The single largest of these gaps, the absence of any audited figure for the economic return on investment, is precisely the question to which the analysis now turns in the economic case for AI investment.



The Socio-economic Investment Case

The Investment Question

This chapter answers a single question for HM Treasury, the DHSC, and the DBIST: why should government invest, over the next five years, in a national programme spanning artificial intelligence, secure health-data infrastructure, high-performance computing, computational science, quantum technologies, and the research infrastructure that connects them, and what would different levels of investment deliver?

How the Treasury should appraise artificial-intelligence investment still matters, and the appraisal method required for that is retained in full below; but it is not the subject of this chapter. The subject is the investment case itself.

The argument advanced here is deliberately conditional rather than promotional. Investment in this programme is justified on the evidence of the pressures set out in Chapter 2, the capabilities set out in Chapter 3, and the clinical evidence set out in Chapter 4, but its value depends on how it is structured, sequenced and governed, not on the technology alone. Three tasks follow. The chapter establishes why the investment is justified; it sets out, through three independently constructed five-year scenarios, what different levels of ambition could achieve and at what risk; and it explains how value is created and captured across the health system, the research base, and the wider economy. The evidentiary standard is the same as the rest of the report: established evidence, developing evidence, illustrative assumptions, and future potential are kept explicitly separate; no audited forecast is

fabricated; and where figures are uncertain or sources diverge, that is stated rather than concealed. Consistent with the commissioning brief, the chapter does not wait for a predetermined budget: the scenarios are derived independently from published evidence and transparent assumptions, and the precise investment envelope remains a matter for the Treasury business case that Chapter 8 describes.

The Appraisal Framework

The framework for the analysis is the HM Treasury Green Book, applied in full but treated here as the supporting method rather than the subject. The model developed for this report adopts the Five Case Model, discounts future costs and benefits over a ten-year horizon, applies explicit optimism-bias uplifts to costs and, where credible patient-outcome effect sizes are available, values health gain through quality-adjusted life years against the relevant cost-effectiveness threshold. Where effect sizes are not yet sufficiently robust, patient-outcome benefits are carried as quantified but unmonetised, pending the Phase 1 evaluations, rather than converted into unsupported monetary values (HM Treasury, 2022).

NICE's standard threshold has been £20,000–£30,000 per quality-adjusted life year since the early 2000s; the Government confirmed in December 2025, as part of the UK-US Economic Prosperity Agreement on pharmaceutical pricing, that this will rise to £25,000–£35,000 from April 2026 (NICE, 2022, on the underlying appraisal methodology this report otherwise draws on; DHSC/NICE, December 2025, on the confirmed threshold change). This is a government-led policy change tied to a trade agreement, not an independent NICE methods revision, and is noted as such here for precision.

Two methodological points are carried transparently throughout. The empirical opportunity cost of NHS spending is estimated to sit closer to £13,000–£15,000 per quality-adjusted life year than to the decision threshold, so health-gain benefits are shown against both. The standard discount rate is itself under review, which re-bases any long-horizon calculation; this is flagged wherever it is material. The purpose of the framework is not to generate a headline number but to ensure that the case is constructed on terms a Treasury audience will recognise and accept.

The Starting Position: Pressures, Costs, and the Funding Context

The case begins from the financial and demand position established in Chapter 2, which defines both the size of the prize and the constraint on realising it. The departmental budget for health and social care in England was £204.7 billion in 2024/25, constrained by a maintenance backlog of £15.9 billion, a provider deficit, and an efficiency requirement of £9.3 billion; sustaining the service on its historical basis is estimated to require real-terms funding growth of around 3.8 per cent a year, some £46 billion of additional funding by 2029/30 (DHSC and NHS England, 2025; Health Foundation, 2024). Public-service healthcare productivity grew by an average of only 0.6 per cent a year over more than two decades and stood some 7.8 per cent below its 2019 level in 2024, against a committed trajectory of 2.0 per cent a year (Office for National Statistics, 2024). Against this, an elective waiting list of 7.22 million pathways and a projected workforce shortfall of between 260,000 and 360,000 staff by the mid-2030s define the scale of unmet demand (NHS England, 2023; 2026). A transparency note on denominators: several different “NHS budget” totals are in circulation – the departmental total, the NHS England allocation, and broader ONS and public-expenditure measures – and this chapter states which is used wherever a proportion is quoted, because mixing them materially distorts any ratio.

The implication is not that the service is unaffordable, but that capital for this programme competes directly with pressing alternative demands and must therefore clear a correspondingly high bar. That is the standard against which the rest of the chapter is held.

How and Where Value is Created

Sizing any scenario requires a precise account of the mechanisms through which value arises, because the central finding of the economic analysis is that value is determined by the terms of deployment, not by the technology. The mechanisms are several: released clinician time that can be redeployed against the workforce shortfall; increased diagnostic capacity per clinician-hour; administrative efficiency through reduced waste; waiting-list reduction; improved outcomes through earlier or better detection; the substitution of cheaper capacity for expensive outsourcing; and, at the research and economic level, the acceleration of discovery, the attraction of inward investment, and the growth of the life-sciences and AI sectors. Each of the health-system mechanisms is supported by verified efficacy evidence from Chapter 4; each of the economic mechanisms operates along the translation pathway established in Chapter 3 from research infrastructure, through computational science and AI development, to clinical validation, NHS deployment, commercialisation, and economic growth. Value is created at every stage of that pathway, but it is captured only where the infrastructure, workforce, and governance exist to convert capability into outcome. This is why foundational investment in data, computing, and skills is treated as preceding, not following, large-scale AI deployment.

A worked example: radiotherapy auto-contouring. The clearest demonstration of the central finding is the one application for which real unit prices are known. Taking radiotherapy auto-contouring priced at between £4 and £50 per plan (NICE, 2024), the illustrative ten-year net present value swings from around minus £21.5 million to plus £33.2 million, and the benefit-cost ratio from well below one to nearly eight, across the plausible range of assumptions – on a basis of 75,000 plans, a fully-loaded clinician value of about £60 per hour, a £2.0 million one-off implementation cost, and a present-value factor of 8.32. Critically, that swing is driven almost entirely by three parameters: the procurement price of the tool, the minutes actually saved in NHS workflow as opposed to in trials, and whether released time is genuinely redeployed. These are illustrative worked consequences of stated assumptions, not verified savings. The lesson scales directly to the programme level: the return on a national investment will be decided by the price at which capability is procured, the effect it achieves in real workflow, and whether the capacity it releases is put to use and the programme must therefore be structured to control those three levers.

Three Five-year Investment Scenarios

The commissioning brief asks for three evidence-based five-year scenarios, a minimum credible investment, a central programme, and a higher-ambition programme developed independently rather than around a predetermined figure, and to set out what each level of investment would cost, how it would be funded, what it would purchase, and what it would deliver. This section provides them in costed form. The figures are illustrative and are built up from published commitments and unit costs rather than asserted; they are expressed as ranges, not point estimates, and the precise envelope remains a matter for the Treasury business case. What has changed from a general statement of ambition is that the ranges are now explicit, so that Government can see the order of investment involved, the funding mix required, and the returns that section 7.6 then quantifies.

Key assumptions. The scenarios rest on a small number of stated assumptions, applied consistently and transparently

- **Horizon.** Investment is profiled over five years; benefits are appraised over ten years, consistent with the Green Book and with the present-value factor of 8.32 used in section 7.4.
- **Discounting and optimism bias.** Future costs and benefits are discounted at the standard 3.5 per cent rate, and costs carry explicit optimism-bias uplifts of the order of +40 per cent for early-stage, less-evidenced capability, falling toward +10 to +20 per cent as audited NHS evidence accumulates through the evidence gates.

- **Incrementality.** The figures describe the incremental, health-directed investment layered on top of infrastructure already committed and largely cross-sectoral: the Compute Roadmap (over £2 billion), the HDRS (up to £600 million), the National Quantum Strategy (£2.5 billion over ten years), and the Life Sciences Sector Plan (over £2 billion). They are not additional to those totals where capability is shared.
- **Anchoring.** Scenario scales are anchored to the AI Diagnostic Fund (£21 million in 2023, expanded by approximately £30 million in 2026) and the HDRS, scaled to the breadth of each programme.
- **Phasing and adoption.** Investment is front-loaded on foundations and released in evidence-gated tranches (section 7.11); adoption is assumed to ramp progressively rather than instantaneously, and released capacity is treated as value only where workflow is redesigned to redeploy it.
- **Status of the figures.** All monetary values are illustrative planning ranges to be confirmed through the Treasury business case and, for the co-funded elements, through negotiation with philanthropic and private partners.

Scenario A: minimum credible investment. Strategic objective: to establish the enabling national foundations and scale only the best-evidenced applications, generating the first audited NHS evidence on which every subsequent decision depends. Illustrative envelope: of the order of £250–500 million of incremental, health-directed investment over five years, a few times the scale of the AI Diagnostic Fund and well within the already-committed HDRS envelope. Cost profile: foundation-weighted and therefore capital-heavy at the outset, approximately 60 per cent capital (interoperability, SDE build-out, national compute access) and 40 per cent operating (independent evaluation, workforce, monitoring), shifting toward operating expenditure as the programme matures. Funding mix: predominantly government (of the order of 80 per cent), which properly funds the foundational public good, with modest philanthropic co-funding (~15 per cent) on data infrastructure and evaluation following the roughly even Wellcome Trust/HDRS precedent and limited private contribution (~5 per cent). What it purchases: completion of the HDRS and SDEs; the interoperability programme; scale-up of the Tier 1 diagnostic and documentation applications identified in Chapter 4, each with mandatory independent evaluation; and the core evaluation and clinical-AI-safety workforce. What it delivers: released diagnostic and documentation capacity in the best-evidenced pathways, and critically, the first audited NHS effect sizes and unit costs, which convert the whole case from illustrative to evidenced. The quantified return is set out in section 7.6. Each funded demonstrator should carry a dedicated implementation and adoption budget covering systems integration, workflow redesign, protected staff time, training, technical support, independent evaluation, and longitudinal data collection, together with a defined procurement or scaling decision if the evidence gate is met.

Scenario B: central investment programme. Strategic objective: to build the integrated capability described in Chapter 3 and connect it deliberately to research translation. Illustrative envelope: of the order of £1–2.5 billion of health-directed investment over five years, a material share of it co-funded. Cost profile: a more even capital/operating balance (of the order of 50/50) as recurring evaluation, workforce, and monitoring costs grow alongside continued infrastructure investment. Funding mix: government share falling to around 55 per cent as private capital is crowded in at the deployment and translation layers (~30 per cent), with philanthropy sustained at around 15 per cent on evaluation and early cluster seed funding. What it purchases: everything in Scenario A, plus broader clinical deployment on conditional, evaluated terms; research-translation platforms linking the national cohorts and SDEs to NHS deployment and industry; a ring-fenced seed for health-relevant computational science and quantum capability (sensing, post-quantum security, and drug-discovery research access); and a funded workforce programme in data science, computational biology, and clinical-AI safety roles. What it delivers: broader capacity and outcome gains across evaluated pathways, a measurable contribution to waiting-list recovery, and the beginnings of research-translation and commercialisation activity quantified in section 7.6.

Scenario C: higher-ambition programme. Strategic objective: to operate the system as a national capability that also drives regional growth and life-sciences leadership. Illustrative envelope: larger again, of the order of £3–5 billion or more in total programme scale but with a materially greater share crowded in from philanthropy and industry, so that the marginal public cost is moderated even as total scale rises. Cost profile: operating expenditure predominates as the programme reaches a steady state of continuous monitoring, re-validation, and reinvestment. Funding mix: government share falling further to around 35–40 per cent, private contribution rising toward 45–50 per cent, and philanthropy sustained at around 15 per cent. What it purchases: everything in Scenario B, plus regional innovation-cluster build-out matched to the AI Growth Zones, accelerated life-sciences translation and commercialisation infrastructure, and a deliberate national-to-regional distribution of capability. What it delivers: system-wide capacity uplift with the NHS operating as a learning health system, sustained research-productivity gains, regional employment and innovation clusters, and a measurable contribution to the Life Sciences Sector Plan’s ambition of a top three global life-sciences economy by 2035, quantified in section 7.6.

The scenarios also recognise that different technologies are supported by different levels of evidence maturity. Artificial intelligence, secure data infrastructure, and advanced computing already possess substantial implementation evidence, whereas quantum technologies remain at an earlier stage; quantum investment is therefore treated throughout as a longer-term, ring-fenced strategic capability rather than a near-term deployment, and forms only a small, clearly bounded component of each envelope.

Illustrative annual profile (Scenario A, at the mid-point of the range, ≈ £375 million over five years). The profile below shows the foundation-front-loaded shape; Scenarios B and C scale the same shape proportionately, with the later years increasingly co-funded. Figures are indicative and for illustration of profile only.

Year	Focus	Indicative Spend	Capital / Operating
1	Interoperability and Secure Data Environment foundations; evaluation capacity stood up	~£90m	~70%/30%
2	Foundations completed; Tier 1 application scale-up begins with evaluation	~£95m	~65%/35%
3	Tier 1 scale-up at pace; first audited effect sizes reported	~£75m	~55%/45%
4	Consolidation; monitoring and workforce embedded	~£60m	~45%/55%
5	Steady-state operation; evidence gate to Scenario B assessed	~£55m	~40%/60%

Costed scenario summary. The table below consolidates the three scenarios across the dimensions the brief specifies. Investment and funding figures are illustrative ranges; benefit statements describe the type, direction, and indicative scale of expected effect, with the full quantification of employment, Gross Value Added, cashable and non-cashable NHS benefit, and public return set out in section 7.6.

Dimension	A Minimum Credible	B Central Programme	C Higher Ambition
Illustrative 5-year envelope	£250–500m	£1–2.5bn	£3–5bn+ total scale
Capital/ Operating balance	~60%/40%	~50%/50%	Operating predominates
Government share	~80%	~55%	~35–40%
Philanthropic share	~15%	~15%	~15%
Private share	~5%	~30%	~45–50%
Primary focus	Generate audited NHS evidence; complete enabling infrastructure	Connect infrastructure, research translation and evaluated deployment	National competitiveness, life-sciences growth, regional innovation
What it purchases	HDRS/SDE completion; interoperability; Tier 1 scale-up with evaluation; core safety workforce	Scenario A plus conditional wider deployment; translation platforms; computational-science and quantum seed; full workforce programme	Scenario B plus regional cluster build-out; commercialisation infrastructure; national-to-regional distribution
NHS benefit	Released diagnostic/documentation capacity in best-evidenced pathways; first audited effect sizes	Broader capacity and outcome gains; measurable waiting-list contribution	System-wide capacity uplift; learning health system
Wider economic benefit	Improved data foundations attract trials	Translation shortens discovery-to-deployment; more UK-anchored R&D	Regional jobs, GVA and commercialisation at scale
Evidence gate	Independent evaluation of every funded deployment	Gate B→C on audited NHS effect and cost evidence	Gate on demonstrated translation and regional outcomes

The recommended posture is to lead with Scenario A establishing the evidence and foundations first, at an illustrative order of £250–500 million over five years, and to release the increment to Scenarios B and C through the evidence gates set out in section 7.11, rather than committing the full ambition in advance. This keeps the low case survivable and the high case reachable, and it is the profile the Government decision in Chapter 8 asks the Treasury to approve.

A Programme-wide Socio-economic Assessment

Section 7.5 set out what each scenario would cost, how it would be funded, and what it would purchase; the task here is to set out what it would return. In line with the commission, the analysis moves beyond the application-level worked example of section 7.4 to a programme-wide assessment, and quantifies that return in productivity and capacity, patient outcomes, employment, Gross Value Added, research and inward investment, commercialisation, and public return on investment as transparent ranges rather than as omitted potential.

A note on method and on why figures are given. No audited NHS dataset yet exists for a programme of this shape, but the purpose of a socio-economic investment case is to estimate the prospective return using transparent assumptions where audited data are not yet available, and to present the result as a range. That is what this section does. Every figure below is **illustrative and modelled, not audited**; each is built from a stated assumption anchored to a published comparator; each is expressed as a low-central-high-range corresponding to the three scenarios in section 7.5; and each will be firmed up as Scenario A's mandatory evaluations report the first audited NHS effect sizes and unit costs (section 7.11). The figures are constructed to be conservative, are discounted at 3.5 per cent and carry optimism-bias uplifts, and are netted for additionality, displacement and leakage as set out below. They are offered to inform the Treasury business case, not to pre-empt it.

Appraisal parameters. The estimates below apply the following parameters consistently. Each is an illustrative planning assumption, to be confirmed in the business case; the ranges are deliberately wide because the underlying evidence is still maturing.

Audit trail and double-counting control. The Treasury business case should maintain a parameter register showing, for every material assumption, the source, calculation, range, sensitivity, and resulting output. Monetised welfare benefits should be counted once. Employment, Gross Value Added, inward investment, company revenue, and exports are important wider-economic indicators but should not be mechanically added to NHS productivity or health benefits in the formal Green Book benefit-cost ratio unless a distinct additional welfare effect is evidenced. The screening ranges below should therefore be reconciled through a non-overlapping benefits map before final approval.

Parameter	Illustrative Value	Basis
Discount rate	3.5% (10-year appraisal)	HM Treasury Green Book
Optimism-bias uplift on costs	+40% (early) falling to +10–20% (evidenced)	Green Book; reduces as evidence gates are passed
GVA per direct high-value job	~£90,000–£130,000	Derived from Office for Life Sciences and House of Commons Library sector measures (\$7.7)
Employment multiplier (Type II, incl. induced)	1.6–2.2	Standard range for high-value R&D and advanced-technology activity
Additionality (share genuinely caused by the programme)	50–70%	Net of deadweight activity that would have occurred anyway
Displacement (activity displaced from elsewhere in the UK)	10–25% deduction	Standard appraisal adjustment
Leakage (benefit accruing outside the UK/target region)	10–20% deduction	Standard appraisal adjustment



The Chancellor of the Exchequer, Rt Hon John Healey MP, will be central to decisions on the public investment given the October Budget. (Photo: Kirsty O'Connor/HM Treasury)

How the netting works (worked illustration). To make the method auditable, the employment estimate is derived through the standard gross-to-net chain, shown here for the higher-ambition Scenario C. Attribution: the employment and GVA channel is attributed principally to the crowded-in private investment and regional-cluster activity, of the order of £2–3 billion, of the Scenario C blended envelope rather than to the whole programme, because the core data and compute infrastructure generates comparatively few direct jobs. Gross direct jobs: at an illustrative 2–6 jobs per £million of cluster-generating investment (anchored to the Cleveland Innovation District’s reported ~2,600 jobs on ~\$500 million and its 20,000-job/\$3 billion decade ambition, and to the Life Sciences Innovative Manufacturing Fund’s ~1,300 jobs on ~£700 million crowded in), this yields of the order of 6,000–18,000 gross direct jobs. Total gross jobs: applying the 1.6–2.2 employment multiplier gives roughly 10,000–40,000 gross direct, indirect, and induced jobs. Net additional jobs: applying additionality ($\times 0.5$ – 0.7), displacement ($\times 0.75$ – 0.9), and leakage ($\times 0.8$ – 0.9), a combined net factor of roughly 0.30–0.57 yields of the order of 5,000–15,000 net additional jobs over the programme’s build-out and early operation. The same chain, scaled to the smaller cluster component of Scenarios A and B, produces the ranges in the summary table below. The width of these ranges is genuine and is itself an argument for the evidence-gated phasing: the figures narrow as audited data replace assumption.

Health-system value: productivity, capacity, outcomes, and cashable versus non-cashable benefit. The core NHS return is generated by the mechanisms of section 7.4 – released clinician time, increased diagnostic capacity, reduced outsourcing, and improved outcomes scaled across the Tier 1 applications in Chapter 4. Scaling the radiotherapy worked example (a ten-year net present value of –£21.5 million to +£33.2 million for a single application) illustratively across the six-to-seven best-evidenced Tier 1 applications gives an indicative programme-level ten-year NHS net present value

of, conservatively, close to zero or modestly negative; centrally, of the order of £100–500 million; and, under optimistic redeployment, higher. The decisive qualification, carried from section 7.4, is the split between cashable and non-cashable benefit: released time is a cashable saving only where workflow is redesigned to redeploy or release the resource, and a non-cashable capacity or quality gain otherwise. On the evidence of the productivity paradox (roughly 19 per cent more staff since the pandemic treating around 14 per cent more patients), this report assumes only 30–60 per cent of gross released-time value is cashable in the central case, with the remainder counted as non-cashable capacity benefit. On patient outcomes, the verified application-level effects fewer interval cancers in AI-supported mammography, a 23 per cent increase in new diagnoses in a genomics sub-study are real but not yet convertible into an audited programme-level quality-adjusted-life-year total; they are therefore carried as a quantified but unmonetised outcome, valued against the £25,000–£35,000 (and empirical £13,000–£15,000) thresholds once Scenario A generates the effect sizes.

Economic value: employment and Gross Value Added. Applying the netting chain above, the programme generates direct employment (in data, computing, research, and clinical-AI roles), indirect employment (in the supply chain), and induced employment (through wider spending), together contributing to regional Gross Value Added. Net additional employment and the associated GVA (net additional jobs × £90,000–£130,000 GVA per job) are set out by scenario in the summary table. These are the figures that most directly answer the commission’s central interest in the programme’s GVA potential; they are illustrative and netted, and they should be read as the prospective scale of the opportunity rather than as a forecast.

Research and inward investment. Faster, safer data access and greater compute capacity raise research productivity along the translation pathway of Chapter 3. The scale of the addressable activity is evidenced by the sector’s momentum: in its first year, the Life Sciences Sector Plan reported attracting around £3 billion of investment and cutting clinical-trial set-up times from 169 to 122 days, while investment into UK AI firms ran at the order of £4.7 billion in 2025 (\$7.7). The programme is assumed to contribute to, not wholly cause, these flows; consistent with the additionality parameter, the case attributes only the additional inward investment the programme’s foundations can reasonably be said to release, which the business case should estimate against a defined counterfactual.

Commercialisation and export potential. Discovery generates intellectual property that universities and industry convert into spinouts, products, and exports (\$7.9). The evidenced pattern, for example the Life Sciences Innovative Manufacturing Fund’s reported crowding-in of over £700 million and creation or safeguarding of more than 1,300 jobs, indicates the order of multiplier achievable, and the commercialisation return is treated as a real but partly capturable benefit, quantified conservatively and net of leakage, rather than through speculative product valuations.

Public return on investment. The ranges below are best treated as screening estimates for the Treasury business case rather than as a final Green Book benefit-cost ratio. The formal appraisal should build the ratio from mutually exclusive monetised welfare benefits and costs, with employment, Gross Value Added, inward investment, and commercialisation reported alongside the BCR unless a separate additional welfare effect is evidenced. This discipline is necessary to avoid double counting while preserving the wider growth case.

On the present illustrative assumptions, the screening public BCR is of the order of 0.8–2.5 for Scenario A, 1.5–4.0 for Scenario B, and 2.0–6.0 for Scenario C, with the lower bound of each range at or below one to reflect the genuine possibility, evidenced in Chapter 4, that individual components deliver no net benefit. These ratios are illustrative and gate-dependent: they improve in expectation as evidence accumulates and as private and philanthropic co-funding lowers the public cost denominator, and they are precisely the figures Scenario A’s evaluations are designed to convert from modelled to audited. These ranges should be recalculated in the formal business case after the non-overlapping benefits map and parameter register are agreed.

Programme-wide summary. The table consolidates the modelled return by scenario. All figures are illustrative, discounted, netted ranges, to be confirmed through the Treasury business case.

Return Dimension	A Minimum Credible	B Central Programme	C Higher Ambition
Net additional jobs (over build-out and early operation)	~500–2,000	~2,000–7,000	~5,000–15,000
Net additional GVA (annual, at maturity)	~£50–250m	~£200–750m	~£450m–2bn
NHS 10-yr net present value (illustrative)	~£0–100m	~£100–400m	~£300m–1bn+
Cashable share of NHS benefit (central)	30–60%	30–60%	30–60%
Patient-outcome benefit	Quantified, unmonetised (QALYs pending Scenario A)	As A, broader coverage	As B, system-wide
Research/Inward investment	Foundations attract trials	Additional UK-anchored R&D	Preferred site for data-intensive research
Commercialisation/Export	Minimal	IP and spin-out pipeline builds	Active licensing and export pipeline
Illustrative public BCR (conservative-led)	0.8–2.5	1.5–4.0	2.0–6.0

What this establishes. The assessment shows that the programme’s principal economic justification lies not in any single artificial-intelligence application but in the cumulative effect of strengthened data infrastructure, advanced computing, research capability, improved clinical productivity, and accelerated research translation, and that this cumulative return, on transparent and conservative assumptions, is material at every scenario level and rises with ambition and co-funding. It should be appraised as a national capability investment rather than as a collection of technology procurements. The figures are modelled rather than audited, and this report states that plainly; but they are given, as ranges with their assumptions exposed, because estimating the prospective return, not deferring it, is the purpose of the exercise the commission set.

National Economic Impact and Competitiveness

The case for this programme reaches well beyond the health system, and this is a central argument rather than a concluding aside. Investment in AI, data, computing, computational science, and quantum is investment in the foundations of national competitiveness, and the international context makes the stakes concrete: comparable economies are committing at scale, from the United States’ multi-billion-dollar quantum and compute programmes to sustained investment across Europe and Asia (Chapter 3). Standing still is therefore a relative decline.

The programme connects to national economic impact through several channels. It strengthens UK competitiveness directly, supporting the Life Sciences Sector Plan’s stated ambition for the UK to become the world’s third-largest life-sciences economy by 2035 (DSIT and DHSC, 2025). It underpins biomedical innovation and research infrastructure through the cohorts, computing, and computational science described in Chapter 3 capabilities, such as protein-structure prediction, anchored in the UK, that have already reshaped drug discovery. It contributes to advanced

manufacturing and the data economy, both substantial parts of the national economy in their own right. It develops skills as a durable national asset. It catalyses private investment: the Life Sciences Innovative Manufacturing Fund, for example, was reported to have crowded in over £700 million of investment and created or safeguarded more than 1,300 jobs. A similar pattern has been reported within the UK's quantum programme, where Innovate UK funding has been associated with significant levels of additional private-sector co-investment, illustrating how strategic public investment can stimulate wider commercial participation. And through regional innovation clusters matched to the AI Growth Zones, which are deliberately distributed beyond the established south-eastern cluster, it supports the long-term productivity and regional rebalancing that are objectives of the Modern Industrial Strategy. The strategic point is that a health-anchored programme is also an industrial-strategy instrument: the same infrastructure that transforms NHS care builds national capability, and the two returns reinforce rather than compete with one another.

Investment in these capabilities also strengthens national resilience by reducing dependence on fragmented digital infrastructure, improving sovereign capability in strategically important technologies, and providing secure research and computational capacity able to respond more effectively to future public-health emergencies. These resilience benefits are difficult to monetise but represent an important strategic consideration within the wider investment case.

A Blended Funding Model

A programme of this scale and risk profile should not be funded by the Government alone, and the international and domestic evidence shows that it need not be. The model proposed is a blended one, in which each partner contributes what it is best placed to provide, in sequence.

The Government provides the foundational, high-risk, and public-good investment data infrastructure, interoperability, evaluation, and the enabling compute that the market will not fund but on which all else depends. Philanthropy co-invests in research infrastructure and mission-driven programmes, as in the HDRS, funded at up to £600 million jointly by government and the Wellcome Trust (DHSC and NHS England, 2026). Industry funds development, deployment, and manufacturing at scale, crowded in by the public foundation, as evidenced by the private commitments attracted under the Life Sciences Sector Plan. Research institutions and universities contribute discovery, talent, and intellectual property. And the NHS provides the deployment environment, the data (under governance), and the evaluation capability that together make the UK an attractive place to develop and prove health technologies.

Philanthropy is best placed to co-fund the areas where public money alone is slowest to move and market money will not yet go: shared data and research infrastructure, in the manner of the HDRS; the evaluation and outcomes research that has no direct commercial return but underpins the evidence gates in section 7.11; and early-stage regional innovation-cluster seed funding of the kind explored in Chapter 8. The HDRS' roughly even government/Wellcome Trust co-funding of its up-to-£600 million commitment is the closest existing UK precedent for how a philanthropic partner might be structured into this programme; Chapter 8 sets out a fuller partnership framework building on it.

Matched philanthropic funding within the blended model. The original commission envisaged government investment matched by philanthropy. The evidence supports retaining that principle for the parts of the programme where philanthropy has a genuine comparative advantage rather than applying a one-for-one match to the entire sovereign infrastructure envelope. The Treasury business case should therefore define a philanthropic-eligible sub-envelope – shared research infrastructure, independent evaluation and outcomes research, and early translation or regional seed funding – and seek an agreed matching ratio for those lines, with 1:1 matching used where partners and project economics support it. The aggregate Phase 1 funding mix may still be approximately 80 per cent government, 15 per cent philanthropy, and 5 per cent private because the Government also funds foundational public goods that are not appropriate for matching.



Blending is economically rational because it reduces public risk while accelerating innovation: public money de-risks the foundational layer and thereby draws in private and philanthropic capital at the deployment layer, so that each public pound supports more total activity than it could alone. The principles of the international innovation-district model anchor institutions, blended public, philanthropic and commercial funding, co-located research, and clinical capability, and transparent accountability against agreed targets (Cleveland Clinic, 2024) are one illustration of how the UK might structure this, and are developed as delivery mechanisms, alongside other models, in Chapter 8. The precise balance between government, philanthropic, and private funding should be determined through HM Treasury business case development, informed by independent evaluation, affordability, delivery readiness, and the evidence generated during each phase of the national investment programme. The Cleveland model also illustrates the need for coordinated governance across government, healthcare providers, universities, research institutions, and industry, ensuring that investment, accountability, and delivery stay coordinated throughout the programme lifecycle.

Research Translation and Commercialisation

The economic pathway established in Chapter 3 – investment, research infrastructure, discovery, clinical validation, NHS deployment, commercialisation, economic growth – is where health investment becomes economic return, and it is worth being precise, without exaggeration, about where value is created and captured. Discovery generates intellectual property; universities and research institutions convert that into spinouts and licence it to industry; university-industry collaboration and access for small and medium-sized enterprises and AI companies turn it into products; and successful products generate exports and inward investment. The UK holds genuine strengths along this pathway – world-class cohorts and computational science at the discovery end,

and a single-payer health system able to validate and deploy at the other – and the strategic task is to connect them so that value discovered in the UK is also captured in the UK, rather than developed here and commercialised elsewhere.

Two cautions apply. The translation pathway is long and uncertain: not every discovery reaches deployment, and the report does not assume otherwise. And commercialisation benefits should be described in terms that the evidence supports the sector's demonstrated ability to attract investment and create high-value jobs rather than through speculative valuations of future products. The clinical evidence that determines what may pass validation is in Chapter 4; the point here is that the commercialisation return is real, is partly capturable through deliberate policy, and is a legitimate part of the investment case.

Procurement as a route from evidence to scale. The July parliamentary roundtable reinforced that grant funding alone does not create a market. Each demonstrator should therefore be designed from the outset as a potential pre-procurement or early-procurement programme: the NHS or relevant public body identifies the problem and required evidence; the demonstrator generates independent clinical, workflow and economic evidence; and, if the gate is met, the user department or NHS system progresses to competitive procurement and wider adoption. This creates a clearer economic pathway from public R&D and evaluation to customer revenue, private investment, manufacturing, and export potential without guaranteeing contracts to individual suppliers.

Sensitivity Analysis and Delivery Risks

The decisive parameters remain uncertain, so the case should be presented through sensitivity analysis rather than point estimates. In line with the Green Book, the analysis should lead with the most conservative scenario. This section shows how the modelled returns in section 7.6 move as the key assumptions vary, which parameters matter most, and what delivery risks sit behind the low case. The low, central, and high outcomes below correspond to the three scenarios, but they are driven as much by delivery as by technology: the same investment can land anywhere across the range, depending on how well it is implemented.

The parameters that decide the answer. The programme-wide model in section 7.6 is sensitive, in descending order of influence, to a small number of levers. The first is the cashable fraction of released capacity – whether released clinician and diagnostic time is redeployed through workflow redesign (cashable) or merely freed (non-cashable); moving this from 30 to 60 per cent is the single largest swing on the NHS net present value. Next is additionality net of displacement and leakage – the share of employment and GVA genuinely caused by the programme rather than deadweight or activity drawn from elsewhere; the combined net factor of roughly 0.30–0.57 used in section 7.6 roughly halves the gross jobs and GVA figures, and small changes in it move the headline materially. Third comes the procurement price and in-workflow effect size of deployed tools – the levers identified in the radiotherapy worked example, which scale directly to the programme. Private and philanthropic co-funding is the fourth, lowering the public-cost denominator and therefore raising the public benefit-cost ratio without changing the gross benefit.

Finally, adoption pace and optimism bias – how quickly capability is absorbed, and how far real costs exceed plan. A one-way sensitivity on each of these, holding the others at their central values, produces results consistent with the ranges presented in Section 7.6; it is only when the pessimistic ends of several parameters combine that the low case turns negative, which is the scenario the phasing is designed to contain.

Low implementation outcome. Under a low outcome, benefits are delayed or diminished by identifiable and evidenced risks: slower adoption than modelled, as the modest uptake of reimbursed digital tools in some comparator systems illustrates; workforce shortages, particularly in the specialist

computational and quantum disciplines where the talent pool is small; procurement delays and the risk of supplier lock-in; regulatory uncertainty, especially around adaptive AI; and infrastructure constraints where data foundations or compute access remain incomplete. In this case, and consistent with the cautionary clinical evidence in Chapter 4, benefits may be zero or negative for individual applications, released time may never become cashable, and the public BCR sits at the lower bound of each scenario's range, at or below one. This is not a tail risk to be noted and dismissed; it is the outcome that materialises if foundations are underfunded or evaluation is skipped, and it is the reason the report recommends leading with Scenario A rather than committing the full envelope.

Central implementation outcome. Under a central outcome, the foundations are in place, adoption proceeds on evaluated terms, workflow is redesigned to capture a material share of released capacity, and co-funding is crowded in broadly as assumed. Benefits accrue steadily rather than immediately, and the programme returns the central ranges set out in section 7.6 – a positive NHS net present value, net additional employment and GVA at the scenario scale, and a public BCR above one.

High implementation outcome. Under a high outcome, strong delivery, effective workforce development, and successful private and philanthropic multiplier effects produce accelerating returns across both the health system and the wider economy: a higher cashable fraction, faster adoption, greater additionality as regional clusters mature, and a lower public-cost share as co-funding rises. This is the upper bound of the section 7.6 ranges, and it is reachable only if the delivery risks below are actively managed rather than assumed away.

Principal delivery risks and their mitigations. The width between the outcomes above is genuine, and it is itself the argument for evidence-gated phasing. The principal delivery risks, each evidenced elsewhere in this report, are paired here with the mitigation built into the programme:

Delivery Risk	Effect on the Case	Mitigation in the Programme
Released capacity not redeployed (non-cashable)	Largest single downside to NHS NPV	Fund workflow redesign; count only cashable benefit in the business case (\$7.6, Ch.8)
Low additionality/High displacement	Halves net jobs and GVA	Estimate against a defined counterfactual; target genuinely additional regional activity (\$7.6)
Slow adoption/Weak absorption	Delays and diminishes all benefits	Phased roll-out; workforce programme; adoption tracked as a KPI
Optimism bias and cost overrun	Erodes BCR	Explicit optimism-bias uplifts; evidence-gated tranches (\$7.5, \$7.11)
Co-funding not crowded in	Raises public-cost share; lowers public BCR	Blended model with defined partner roles; public foundation de-risks private capital (\$7.8, \$8.7)
Specialist workforce shortage (quantum, computational biology)	Constrains pace, especially Scenarios B–C	Skills treated as a foundation layer, funded and sequenced (\$3.8, Ch.8)
Individual tool delivers no benefit	Zero or negative return on that component	Independent validation before deployment; go/no-go gates (Ch.4, Ch.9)

The overriding conclusion is that the programme should be structured so that the low case is survivable and the high case is reachable, rather than betting the whole envelope on the central assumption. That is exactly what the phased, evidence-gated design in section 7.11 and Chapter 8 delivers: it concentrates early commitment where the evidence is strongest, generates the audited data that narrows every range above and releases further capital only as that evidence justifies it.

Evidence Gates and Phased Investment

Funding should therefore be released in tranches, each gated by evidence, rather than committed in advance. This is the mechanism that converts the analysis above into a defensible investment programme. Progression from one phase to the next and from Scenario A to B to C should depend on: independent evaluation of every funded deployment; measurable outcomes against pre-agreed metrics; deployment evidence that effect sizes achieved in trials and the transfer to NHS workflow; NHS-specific economic evaluation that begins to close the audited-savings gap; demonstrated research translation along the pathway of Chapter 3; and governance milestones, including the assurance and validation functions set out in Chapter 9. Tying future funding to demonstrated evidence reduces implementation risk, protects public money against the failure modes documented throughout the report and ensures that the programme scales only as fast as the evidence justifies. This gated, phased approach is the practical expression of the whole chapter's argument.

Conclusion: Why Invest, and the Bridge to Chapter 8

The chapter has set out why investment is justified. The pressures on the health system are structural and costly; the capabilities to help address them are real, though uneven, and at the frontier, immature; and the return in health gain, released capacity, research acceleration, and wider economic growth is genuine but conditional on the terms of deployment rather than on the technology. Three independently constructed scenarios show what different levels of ambition could achieve and at what risk, without resting on a predetermined budget or a fabricated forecast. A programme-wide socio-economic assessment establishes the structure of the return and, honestly, the limits of what can yet be quantified. A national economic-impact analysis shows that a health-anchored programme is also an instrument of competitiveness, regional growth, and life-sciences leadership. A blended funding model shows how public, philanthropic, and private investment can share the cost and the risk. And a framework of evidence gates shows how the programme can be scaled safely. What this chapter does not do is specify how the investment should be delivered, the phasing, the delivery bodies, the procurement and partnership mechanisms, the governance, and the workforce programme that turn a justified case into a working national programme. That is the task of Chapter 8. In short, this chapter demonstrates why the Government should invest; Chapter 8 sets out how.

Accordingly, this report does not recommend that the Government immediately commit to the highest level of investment. Rather, it recommends approval of the minimum credible programme required to establish the national foundations, generate independent evidence and demonstrate measurable value of the order of low hundreds of millions of pounds of incremental, health-directed investment over five years, per the Scenario A range set out in section 7.5. Progression to the broader investment scenarios should then depend upon the evidence gates established in this chapter, ensuring that future public expenditure is released only where independently evaluated outcomes justify further expansion. The question is therefore no longer whether the UK should invest in these capabilities, but how quickly it can build them in a disciplined, evidence-gated manner that maximises public value.



The National Delivery Programme

Introduction

Chapter 7 established why the Government should invest in a national programme spanning artificial intelligence, secure health data infrastructure, high-performance computing, computational science, and quantum technologies, and set out what different levels of ambition could achieve over five years. This chapter answers the delivery question that follows directly from it: how should the Government fund, govern and deliver that programme? It is written as an implementation strategy for HM Treasury, the DHSC, and the DBIST, not as an NHS deployment plan; the NHS is one delivery partner within a wider national programme that also serves research translation, life-sciences growth, and regional economic development. Consistent with the approach applied throughout this report, no specific funding figure is recommended, and no economic return is asserted beyond what the evidence supports. The appropriate investment envelope remains a matter for the Treasury business case, informed by the appraisal model in Chapter 7 as its evidence gaps are closed.

Strategic Objectives

The strategic objectives follow directly from the analysis in Chapters 2, 3, and 7. The programme aims to relieve the structural pressures identified in Chapter 2 – workforce shortfall, diagnostic backlog, and stalled productivity – by deploying artificial intelligence and the infrastructure that supports it where the evidence shows benefit, treating technology as one evidence-led lever rather than a panacea. It aims to generate, for the first time, audited evidence of clinical, research, and economic effect, closing the gap that currently prevents any credible statement of return. It aims to build the enabling national foundations of data, computing, computational science, workforce, and governance, without which no application can scale safely. It aims to establish a durable capability for continuous evaluation, monitoring, and reinvestment. And, reflecting the widened scope of this report, it aims to convert that national capability into research translation, life-sciences growth, and regional economic development, so that the programme delivers value beyond the health system as well as within it. These objectives are mutually reinforcing and are achievable within the existing institutional setup, adapted as set out below.

Investment Principles

Seven principles govern the investment. Start with the healthcare problem and measurable outcome rather than a technology category. Sequence by evidence: deploy and scale the best-evidenced applications now and hold weaker or less mature capabilities to explicit validation gates. Fund the enabling foundations – usable data, interoperable systems, secure compute, evaluation, and workforce – before assuming applications will scale. Build procurement into the route from demonstrator to adoption rather than relying on grants alone. Fund implementation capacity as part of every demonstrator. Choose deliberately what the UK should build, what it should deliver through partnership, and what it should access as a service. Finally, lead investment decisions with the conservative case and release later capital only when independent evidence gates are met.

These principles are drawn from the international evidence examined in Chapter 5, and match precisely with the Green Book's emphasis on proportionate, evidence-led appraisal.

A Phased Five-year National Investment Programme

This section is the centrepiece of the delivery programme. It translates the three costed scenarios developed in Chapter 7 into a working five-year sequence, and the priority order of Chapter 4 into a delivery timetable. Progression between phases is governed by the evidence gates set out in section 8.11, not by the calendar: the year ranges below indicate an expected working assumption, not a fixed commitment, and a phase does not advance until its evidence gate is satisfied, whatever the elapsed time. The figures carried through from Chapter 7 are illustrative planning ranges, to be confirmed through the Treasury business case; they are reproduced here so that the delivery sequence, its cost, and its funding mix can be read together.

How to read the investment envelopes. Scenarios A, B, and C are cumulative alternative programme scales, not amounts to be added sequentially. If the programme progresses from A to B, the additional commitment is the difference between expenditure already approved or incurred under A and the total B envelope; Scenario C likewise represents total programme scale rather than a further £3–5 billion on top of B. Within each envelope, cash should be released in annual or milestone-based tranches, subject to the evidence gates. The Treasury business case should convert these ranges into a year-by-year profile before approval.

Phase 1 foundations and proof (Years 1–2). Strategic objective: to establish the national data, computing, and governance foundations, and to generate the first audited NHS and research evidence, consistent with Scenario A in Chapter 7. Illustrative investment: of the order of £250–500 million over five years, front-loaded into Years 1–2 (indicatively ~£90 million and ~£95 million in the first two years, on the profile in section 7.5), predominantly government-funded (~80 per cent), capital-weighted (~60–70 per cent capital). Priority investments: completion of the HDRS and SDEs; the interoperability programme; continued Phase 1 should also approve a small, regionally distributed portfolio of approximately three to five demonstrators selected against patient or system need, technical fit, maturity, data readiness, workflow and adoption, UK advantage, economic value, equity, and trust. Every demonstrator must name the NHS or public-sector user, establish baseline measures, include funded implementation capacity and define the procurement or scaling decision that follows a successful evaluation.

Build-out of national compute capacity through the AI Research Resource and the Compute Roadmap; scale-up of the Tier 1 clinical applications identified in Chapter 4 (radiology, clinical documentation, radiotherapy auto-contouring, endoscopy, genomics, mental-health triage), each with funded independent evaluation; establishment of the national governance architecture described in section 8.9; and a funded workforce programme in data science, clinical-AI safety, and evaluation capability. Programme leadership and delivery partners: HM Treasury and the DHSC set the strategic envelope; the national assurance function (section 8.9) and the DBIST lead on infrastructure and compute; NHS providers and the HDRS deliver deployment and data access; the National Institute for Health and Care Excellence and the Medicines and Healthcare products Regulatory Agency lead validation and regulation. Illustrative return (from section 7.6): of the order of 500–2,000 net additional jobs, an NHS ten-year net present value in the region of £0–100 million, and first audited effect sizes. Evidence gate to Phase 2: independent evaluation confirming effect sizes achieved in NHS workflow (not only in trials); a first audited cost and outcome baseline for the priority applications; and confirmation that the data and interoperability foundations are operating as intended.

Illustrative candidate portfolio. Final demonstrators, delivery partners, and budgets should be confirmed through the Treasury business case; the table does not pre-select suppliers.

Candidate Phase 1 Demonstrator	Primary Problem/Outcome	Capability Mix	Decision at Evidence Gate
Diagnostic and oncology imaging	Diagnostic fidelity, throughput, waiting time, and avoidable follow-up	AI, imaging data, classical compute; quantum sensing where sufficiently mature	Competitive procurement/regional scale if clinical and workflow value is demonstrated
Clinical documentation and administrative productivity	Clinician time, administrative burden, and error	AI, secure data integration, workflow redesign	Scale only where released time is evidenced and redeployed
Quantum-enabled sensing/brain imaging	More sensitive or patient-friendly imaging in defined pathways	Quantum sensing, clinical imaging, AI interpretation	Progress from demonstrator to procurement only after regulatory and pathway evidence
Drug discovery/multi-omics translation	Faster, higher-quality analysis and target identification	Research data, HPC, computational science, AI, and selective quantum access	Continue or expand only where benchmark advantage and translation value are evidenced
Complex NHS and social-care optimisation	Scheduling, flow, diagnostic capacity, discharge, and workforce constraints	Classical optimisation, AI, and hybrid advanced computing	Adopt where operational benefit exceeds leading classical baseline and is reproducible



Phase 2 evaluated expansion (Years 2–4). Strategic objective: to extend deployment on conditional terms, connect the national research infrastructure to clinical translation, and seed the enabling computational-science and quantum foundations, consistent with Scenario B in Chapter 7. Illustrative investment: of the order of £1–2.5 billion over five years, a material share co-funded government share falling to around 55 per cent as private capital (~30 per cent) is crowded in at the deployment and translation layers, with philanthropy sustained at around 15 per cent; capital/operating balance moving toward 50/50. Priority investments: capital released in tranches, gated by the Phase 1 evidence, to extend deployment to the Tier 2 applications identified in Chapter 4; research-translation platforms linking the national cohorts and SDEs described in Chapter 3 to NHS deployment and to industry; a ring-fenced seed for health-relevant computational science and quantum capability sensing, post-quantum security, and drug-discovery research access; modernised, outcome-linked procurement; and the workforce programme extended to computational biology and quantum-adjacent skills. Programme leadership and delivery partners: the Programme Board (section 8.9) takes on full oversight of the expanded portfolio; regional delivery boards, universities, and industry partners join as delivery partners, particularly for research-translation and computational-science investments; philanthropic co-funders participate on the blended-funding terms set out in section 8.7. Illustrative return (from section 7.6): of the order of 2,000–7,000 net additional jobs, ~£200–750 million of net additional annual GVA at maturity, an NHS ten-year net present value in the region of £100–400 million, and the beginnings of research-translation and commercialisation activity. Evidence gate to Phase 3: audited NHS-specific economic evaluation demonstrating net benefit at the applications deployed; evidence of research translation into industry or NHS practice; and confirmation that regional delivery partnerships are functioning as intended.

Phase 3 learning health system and national capability (Years 4–5 and beyond). Strategic objective: to operate the programme as a self-sustaining national capability continuously monitored, re-validated, and reinvested, and to realise the regional and commercialisation returns set out in Chapter 7, consistent with Scenario C. Illustrative investment: larger again in total programme scale of the order of £3–5 billion or more but with a materially greater share crowded in from philanthropy and industry (government share ~35–40 per cent, private ~45–50 per cent, philanthropy ~15 per cent), so the marginal public cost is moderated even as total scale rises; operating expenditure predominates at steady state. Priority investments: continuous monitoring, re-validation, and retirement of underperforming tools; reinvestment funded from evaluated savings; regional innovation-cluster build-out matched to the delivery model in section 8.6; accelerated commercialisation and export support; and revisiting, on the evidence then available, the longer-horizon quantum-computing opportunity. Programme leadership and delivery partners: full national-to-regional delivery architecture as set out in section 8.6, with the Programme Board accountable to HM Treasury for the programme's overall public value. Illustrative return (from section 7.6): of the order of 5,000–15,000 net additional jobs, ~£450 million–£2 billion of net additional annual GVA at maturity, an NHS ten-year net present value above £300 million, and an active commercialisation and export pipeline. Evidence gate: sustained delivery against agreed key performance indicators, published equity and outcome monitoring, and demonstrated regional and commercialisation impact, before any further national scale-up is committed.

Summary of the phased programme. The table reproduces the phase sequence with the costed figures carried through from Chapter 7. All figures are illustrative ranges, to be confirmed through the Treasury business case.

Phase	Timeframe (working assumption)	Illustrative Investment	Funding Mix (gov/phil/private)	Primary Investment	Illustrative Return (net)	Evidence Gate
1 Foundations and proof	Years 1–2	£250–500 m, front-loaded	~80/15/5	Data, interoperability, compute, Tier 1 applications, governance	~500–2,000 jobs; NHS NPV ~£0–100m; first audited evidence	Evaluated NHS effect sizes and cost baseline; foundations operating
2 Evaluated expansion	Years 2–4	£1–2.5bn, part co-funded	~55/15/30	Tier 2 applications, translation platforms, computational science/quantum seed, workforce	~2,000–7,000 jobs; GVA ~£200–750 m/yr; NHS NPV ~£100–400 m	Audited economic evaluation; evidence of translation; regional partnerships functioning
3 Learning health system	Years 4–5+	£3–5bn + total scale	~37/15/48	Continuous monitoring, regional clusters, commercialisation, quantum reassessment	~5,000–15,000 jobs; GVA ~£450m–2 bn/yr; NHS NPV >£300m	Sustained KPI delivery; published equity/outcome monitoring; demonstrated regional/commercial impact

The programme is designed to be entered at Phase 1 and scaled through the gates, not committed in full in advance. This keeps the low case survivable and the high case reachable, and it is the profile that the Government decision in section 8.12 asks Treasury, DHSC, and DBIST to approve.

Build, Partner or Access? A Strategic Delivery Model

A recurring question for a programme of this breadth is not only what to fund but how to deliver it: which capabilities the Government should build as sovereign national infrastructure, which are best delivered through partnership with universities, industry, or philanthropy, and which should simply be accessed as a service rather than owned. Answering this well matters directly to public value, because building what could be accessed wastes capital, while accessing what should be built risks dependency on capability the country does not control. The table below sets out the report's assessment for the principal capabilities examined in Chapter 3, on the evidence available; where the balance depends on decisions not yet taken, this is stated rather than resolved by assumption.

Capability	Build	Partner	Access	Strategic Rationale
Secure health data infrastructure (HDRS, national identifier)	✓			A sovereign public good; cannot be safely outsourced given the sensitivity and permanence of the data.
Secure Data Environments	✓	✓		Core platform built nationally; individual environments may be delivered in partnership with academic and NHS host institutions.
High-performance computing (national facilities)	✓			Already a proven national build (AI Research Resource, Isambard-AI, Dawn); the returns to sovereign capacity are established.
AI Research Resource	✓	✓		National asset; access arrangements for health and life-sciences research delivered in partnership with universities.
Computational science capability	✓	✓		Best delivered through university and industry partnership (for example, the AlphaFold model); routine capacity can be accessed via national and cloud compute.
AI deployment (clinical applications)		✓	✓	Procured from industry and delivered through NHS partnership; not a sovereign build.
Cloud infrastructure			✓	A mature, competitive commercial market; access on well-governed, portable terms is the efficient choice.
Quantum sensing		✓		Developing capability best developed through partnership between national programmes, universities, and industry.
Quantum computing		✓	✓	Early-stage capability. Access to national facilities such as the National Quantum Computing Centre and international research partnerships is the appropriate near-term approach while evidence continues to mature. Decisions regarding sovereign national capability should remain evidence-led and be revisited as technologies develop and strategic requirements become clearer.

The pattern that appears is consistent with the report's wider approach. National capabilities that should be built are those that are foundational, long-lived, and where sovereign control is itself the point: secure data infrastructure and the core national compute estate. Capabilities best delivered through partnership are those where the UK's comparative advantage lies in bringing together anchor institutions universities, NHS organisations, and industry, rather than in government acting alone: research translation, computational science, quantum sensing, and clinical AI deployment. Capabilities that should simply be accessed are those available in a mature, competitive market where ownership adds cost without adding public value: general cloud infrastructure, and, for now, most quantum-computing capacity.

This three-way distinction is not proposed in isolation. The UK's National Quantum Technologies Programme has organised its strategy around the same principle for more than a decade through its "Own, Collaborate, Access" (OCA) framework. The Build-Partner-Access model recommended throughout this report therefore extends an approach already tested within UK quantum policy to the wider AI, data, advanced computing, and healthcare research ecosystem, rather than introducing an entirely new delivery framework.

Two illustrative international models show how the partnership tier can be organised as a delivery structure, not a source of funding, which is addressed separately in section 8.7. The Cleveland Innovation District in the United States brings together anchor hospitals and universities, government, and industry to co-locate research and clinical capability under a single delivery structure with shared governance and published targets, including dedicated advanced-computing and quantum research delivered in partnership with industry (Cleveland Clinic, 2024). Medicon Valley, the cross-border life-sciences cluster spanning eastern Denmark and southern Sweden, illustrates a different but comparable pattern: universities, some 32 hospitals, and several hundred life-sciences companies operate through a shared cluster organisation, jointly fund PhD programmes split between hospital, university, and company partners and, most relevant to this report, have co-invested in shared advanced-computing infrastructure, including the Gefion supercomputer, which supports drug discovery, genomics, and AI research across the whole cluster rather than any single institution. Medicon Village, one constituent site, was itself formed when a former industry facility closed and local universities, the municipality, and healthcare region stepped in alongside private investors to keep the site operating as shared research infrastructure – a smaller scale, but instructive precedent for rescuing and repurposing capability through exactly the blended, multi-party funding model this report proposes.

The lesson for delivery, common to both models, is that partnership works best where it is structured, accountable, and anchored to durable institutions rather than to a single procurement. Neither is presented as a template for direct replication; the NHS's single-payer, single-national-identifier structure differs materially from both the fragmented US system and the bi-national Danish-Swedish arrangement, but together, they demonstrate that anchor-institution partnership, shared computing infrastructure, and blended funding are a proven pattern internationally, not a single-country idiosyncrasy. This differentiated approach reflects both international experience and recent expert discussion, which concluded that national capability should combine sovereign investment, strategic partnerships, and selective access, depending on the maturity and strategic importance of each technology. The recommendation throughout this report is therefore to apply all three approaches deliberately rather than adopting a single delivery model across the programme. This balanced build-partner-access approach is intended to maximise public value while reducing delivery risk: it directs scarce public capital to sovereign foundations, uses partnership to mobilise capability the country does not need to own outright, and reserves market access for capacity that is already efficiently supplied.

A Blended Public, Philanthropic, and Private Investment Model

A programme of this scale should not be funded by the Government alone, and it need not be. The funding model set out here describes how different sources contribute, and in what proportion; the delivery arrangements that sit alongside it are addressed in sections 8.6 and 8.8. The figures below are illustrative planning ranges, carried through from the scenarios in Chapter 7; the precise contributions are a matter for the Treasury business case and for direct negotiation with philanthropic and private partners, and are offered here to show the shape and order of the funding mix rather than to fix it.

The Government funds the foundational, high-risk infrastructure that the market will not fund but on which everything else depends: secure data, national compute, and the governance and evaluation architecture. Philanthropy co-invests in research infrastructure and mission-driven programmes where its risk appetite and time horizon complement public funding, as already demonstrated by the Wellcome Trust's co-funding of the HDRS. Research councils, principally UK Research and Innovation and the National Institute for Health and Care Research, fund discovery science, the national research cohorts, and the translational studies that connect research to clinical practice. Universities contribute talent, computational-science capability, and the intellectual property that underpins commercialisation. Industry funds development, deployment, and manufacturing at scale, crowded in by the public foundation rather than substituting for it, as evidenced by the private investment attracted under the Life Sciences Sector Plan. And the NHS provides the deployment environment, the governed data, and the evaluation capability that make the UK an attractive place to develop and prove health technologies.

How the funding mix shifts across the programme. The defining feature of the model is that the public share falls as the programme matures, because the early public investment in foundations is precisely what crowds in private and philanthropic capital at the later deployment and translation layers. Government carries the great majority of the cost in Phase 1, where the investment is foundational and pre-commercial; by Phase 3, private and philanthropic sources together carry the majority, and the marginal public cost of each additional pound of programme activity is at its lowest. Each public pound therefore supports more total activity than it could alone, and the proportions are stated here as illustrative ranges to be confirmed in the business case.

Phase	Illustrative Total	Government	Philanthropy	Private/Industry
1. Foundations and proof	£250–500m	~80% (~£200–400m)	~15% (~£40–75m)	~5% (~£15–25m)
2. Evaluated expansion	£1–2.5bn	~55% (~£550m–1.4bn)	~15% (~£150–375m)	~30% (~£300–750m)
3. Learning health system	£3–5bn+ total scale	~37%	~15%	~48%

The Government contribution is therefore front-loaded and foundational: of the order of £200–400 million in Phase 1, concentrated on the data infrastructure, interoperability, national compute access, evaluation, and governance that no other partner will fund but on which the whole programme depends. This is the investment the Government decision in section 8.12 asks Treasury, DHSC, and DBIST to approve first.

The philanthropic contribution is not a general-purpose top-up. It is directed at three specific categories of funding gap, where its characteristics – patience, mission fit, tolerance of pre-commercial risk – are a genuine comparative advantage over both government and industry money, at an illustrative order of £40–75 million in Phase 1, rising with programme scale.

- **Shared research infrastructure with no direct commercial return**, following the HDRS precedent, in which the Government and the Wellcome Trust co-funded a commitment of up to £600 million on a roughly even basis. This is the closest existing UK precedent for how a philanthropic partner should be structured into Phases 1 and 2 of this programme, and the same co-funding architecture, joint governance, published targets, and a named philanthropic partner alongside government should be the default model rather than a one-off arrangement.
- **Independent evaluation and outcomes research**, which has no commercial return but is the single evidence gate on which every phase transition in section 8.5 depends. Philanthropic funders have historically been willing to fund evaluation capacity that neither government (facing near-term budget pressure) nor industry (with no direct interest in an independent finding) will reliably prioritise.
- **Early-stage regional cluster and translation seed funding**, of the kind illustrated by Medicon Village's formation in section 8.6, where a small, patient tranche of philanthropic capital at the outset can enable a self-sustaining cluster that government and industry then scale.

The private and industry contribution rises steeply as the programme matures from a modest ~£15–25 million in Phase 1 to the largest single share by Phase 3 because private capital follows evidence and de-risked foundations rather than leading them. The evidenced pattern of co-investment under the Life Sciences Sector Plan and the Life Sciences Innovative Manufacturing Fund (which reported crowding in over £700 million of investment) indicates the order of private co-investment a well-structured public foundation can mobilise.

The precise financial ratio between government, philanthropic, and private contributions across the five-year programme is, correctly, a matter for the Treasury business case and for direct negotiation with philanthropic partners. This report does not assert a fixed proportion but the ranges above are intended to give government and prospective partners a concrete, adoptable starting structure rather than a general aspiration to “blend” funding. Within that overall mix, the business case should define the philanthropic-eligible sub-envelope and the matching rule for it, so that the original matched-funding proposition is visible and contractible rather than implied.

Funding Source	Principal Role	Illustrative Contribution	Illustrative Example
Government	Foundational infrastructure; governance; evaluation	~£200–400m in Phase 1; falling share thereafter	HDRS; Compute Roadmap; AI Diagnostic Fund
Philanthropy	Shared infrastructure, evaluation, cluster seed	~£40–75m in Phase 1, rising with scale	Wellcome Trust co-funding of HDRS (~50/50); Medicon Village formation
Research councils	Discovery science; national cohorts; translational research	Matched existing budgets	UK Research and Innovation; NIHR
Universities	Talent, computational science, and intellectual property	In-kind and co-investment	AlphaFold (Google DeepMind, EMBL-EBI)
Industry	Development, deployment, and manufacturing at scale	Rising to the majority share by Phase 3	Private investment under the Life Sciences Sector Plan
NHS	Deployment environment, governed data, evaluation	In-kind; deployment and evidence generation	NHS trusts and the national clinical-AI evidence base

Blending is economically rational because it reduces public risk while accelerating innovation: public investment de-risks the foundational layer, which in turn crowds in private and philanthropic capital at the deployment and translation layers, so that each pound of public money supports more total activity than it could alone.

National and Regional Delivery

Delivery operates at two levels, and the programme depends on the connection between them. At the national level, the Government sets strategy and funds the shared foundations, the data infrastructure, the national compute estate, the governance, and assurance architecture, because their scale and cost make duplication wasteful and their value increases with the breadth of data and users they serve. National capability is then translated into use through a chain of regional and institutional delivery: regional innovation centres, matched to the AI Growth Zones described in Chapter 3, bring together universities, NHS organisations, and industry around shared compute and research infrastructure; integrated care systems, through their integrated care boards, commission and coordinate deployment across their local NHS providers; NHS trusts deliver care and generate the deployment evidence that closes the audited-evidence gap; and universities and industry partners supply the research, computational-science, and commercial capability that connect local deployment back to national research translation and commercialisation.

A UK portfolio with clear territorial scope. National research and industrial capability should be organised as a UK-wide portfolio, drawing on evidenced strengths across the nations and regions – including photonics capability in Scotland, compound semiconductors in South Wales, national infrastructure at Harwell, established life-sciences clusters, and the wider university and company network highlighted at the parliamentary roundtable. Regions should act as delivery and learning partners, with defined roles and common outcome measures, rather than competing for disconnected pilots. Where recommendations rely on NHS England commissioning or provider structures, equivalent delivery arrangements should be agreed with the devolved administrations rather than be assumed to apply automatically.

This chain is presented with an important caveat about institutional change. The structures through which regional delivery operates are themselves being reformed during the period this programme covers. The Government announced in March 2025 that NHS England will be abolished, with its functions absorbed into the DHSC; DHSC’s own published timetable targets legal completion by April 2027, subject to the necessary primary legislation.

Separately, and on its own timetable, NHS England is consolidating the existing 42 integrated care boards into a smaller number of larger footprints, with mergers taking effect on 1 April 2026 and 1 April 2027 (DHSC, 2025; NHS England, 2026). This transition affects accountability during the roadmap period and is noted here as a live delivery risk rather than a settled fact: the programme’s national-to-regional delivery chain should be designed to remain workable as these institutional changes proceed, and the governance arrangements in section 8.9 are framed accordingly, with lines of accountability to the DHSC rather than to NHS England specifically.

Governance and Accountability

Safe and accountable delivery of a programme this size requires a governance architecture with clear roles and a single line of accountability throughout the five years, not a set of parallel committees. A programme board, accountable to HM Treasury and the DHSC, holds overall responsibility for the programme’s strategic direction, its budget, and its public value, and is the body that takes the go/no-go decisions at each evidence gate described in section 8.5. A scientific advisory board provides independent expert scrutiny of the computational-science, artificial-intelligence, and quantum evidence underpinning the programme, ensuring that technical claims are tested rather than taken on trust. A clinical advisory board, drawing on the expertise embedded in the National Institute for Health and Care Excellence and the Medicines and Healthcare products Regulatory Agency, oversees clinical safety, validation, and the human-oversight requirements set out in Chapter 9. Economic evaluation is commissioned as a standing function, not a one-off exercise, and is responsible for generating the audited NHS-specific evidence that closes the gap identified in Chapter 7. Independent assurance, sitting outside the delivery chain, audits the programme against its evidence gates and reports directly to the programme board, so that the body making funding decisions is never solely reliant

on self-reported delivery performance. And regional delivery boards, matched to the integrated care systems and regional innovation centres described in section 8.8, are accountable for local delivery and feed performance and evaluation data upward to the national functions.

Accountability runs in both directions throughout the five years: the programme board sets the conditions under which capital is released, the advisory and assurance functions test whether those conditions are met, and the regional boards deliver and report against them. This structure is designed to survive the institutional transition described in section 8.8, because it is built around functions rather than around named bodies that are themselves subject to reform.

Intellectual Property, Commercialisation, and Public Value

Where public and philanthropic capital funds discovery, and where NHS data underpins the development of clinical tools, the terms on which resulting intellectual property is owned, licensed and shared are a matter of public value, not merely of contract law, and this report addresses them as a policy question. Four principles govern this. On ownership and licensing, the default position should ensure that intellectual property arising substantially from public investment carries licensing terms that protect a public return, whether through direct licence income, reduced future procurement cost, or agreed benefit-sharing with the NHS bodies whose data or clinical partnership contributed to the discovery. On NHS and research access, any tool developed using NHS data or NHS clinical collaboration should carry contractual guarantees of continued NHS access on reasonable terms, and continued research access for approved investigators through the SDE route described in Chapter 3, so that public contribution to discovery is not later priced out of public use. On safeguards against vendor lock-in, procurement should require open standards, data portability, and clear exit rights as a condition of purchase, following the procurement principles set out for clinical AI here and extended to computing and data infrastructure contracts, so that the programme's dependence on any single supplier is limited by design. And on commercialisation, the Government should support spinouts, licensing, and university-industry collaboration as legitimate and intended returns on this investment, consistent with the pathway from research infrastructure to economic growth set out in Chapter 7, while ensuring that the terms of that commercialisation reflect the public origin of the underlying capital and data. Taken together, these principles are intended to ensure that taxpayer investment in discovery generates a taxpayer return, whether that return is realised in patient benefit, in licence income, in reduced future cost, or in the broader economic growth described in Chapter 7.

Evidence Gates and the Release of Future Funding

Consistent with Chapter 7, future funding under this programme should be released only as evidence justifies it, and the evidence required is specific rather than general. Progression through the phases in section 8.5 and any expansion beyond the initial Phase 1 commitment should depend on: independent evaluation of every funded deployment, conducted at arm's length from the supplier and, ideally, from the delivering body; economic evidence, generated through the standing evaluation function in section 8.9, that closes the audited-savings gap identified in Chapter 7; clinical evidence, meeting the validation standards set out in Chapter 9, including external validation of any predictive or risk-stratification tool before deployment; operational performance against the key performance indicators the programme establishes, including adoption, capacity, and equity measures; research translation, evidenced by discoveries or tools moving from the national research infrastructure into clinical use or commercial development; and public value, assessed against the intellectual-property and benefit-sharing principles in section 8.10. This is a deliberate design choice: it makes future government funding decisions the product of demonstrated evidence rather than of a predetermined spending profile, and it is the practical expression of the evidence-gated discipline that runs throughout this report.

Government Decision Required

This report asks HM Treasury, the DHSC, and the DBIST to approve the first phase of a five-year, evidence-gated national programme, together with the initial demonstrator, procurement, and implementation architecture, and to endorse the mechanism through which its later phases would be sized and released. The recommendation is to commit to Phase 1, of the order of £250–500 million, of incremental, health-directed investment over five years, front-loaded into the first two years and approximately 80 per cent government-funded (with philanthropic and private co-funding of around 15 and 5 per cent, respectively) and to approve the governance and evidence-gate mechanism that would determine any further investment. Specifically, government is asked to approve

- **Phase 1 investment** in the foundational data, computing, and clinical capabilities set out in this report, of the order of £250–500 million over five years, sized precisely through the Treasury business case to complete the HDRS and SDEs, deliver the interoperability programme, secure national compute access and scale the Tier 1 clinical applications of Chapter 4, each with mandatory independent evaluation.
- **The national governance architecture:** a programme board with scientific, clinical, and independent assurance functions, accountable throughout the five-year programme.
- **A standing evaluation programme** sufficient to generate the audited clinical and economic evidence currently missing, and so to convert the illustrative returns of Chapter 7 into audited ones.
- **Continued investment in the national research infrastructure:** secure data, computing, and computational science set out in Chapter 3.
- **A government-philanthropic partnership framework**, formalising the blended-funding and build-partner-access models set out in this chapter, with a defined philanthropic-eligible sub-envelope and an agreed matching rule for shared research infrastructure, independent evaluation, and early translation or regional seed funding.
The illustrative aggregate Phase 1 mix remains approximately 80 per cent government, 15 per cent philanthropic, and 5 per cent private, while matched funding is applied within the eligible philanthropic sub-envelope rather than across all foundational public-good expenditure.
- **A national research-translation programme** connecting discovery to NHS deployment and commercialisation, with the public-value and intellectual-property safeguards set out in section 8.10.
- **The evidence-gated release of future funding**, so that later phases (of the order of £1–2.5 billion for Scenario B and £3–5 billion or more in total scale for Scenario C, with a rising private and philanthropic share) proceed only as evidence from earlier phases justifies.

What the Phase 1 investment is expected to deliver. On the illustrative, conservative, and netted basis set out in Chapter 7, the Phase 1 commitment is expected to generate: the first audited NHS effect sizes and unit costs for the best-evidenced applications; released diagnostic and documentation capacity in those pathways; of the order of 500–2,000 net additional jobs; an NHS ten-year net present value in the region of £0–100 million; and the enabling data, compute, and governance foundations on which Phases 2 and 3 depend. Progression to the broader scenarios would be expected to deliver, cumulatively, of the order of 5,000–15,000 net additional jobs and £450 million–£2 billion of net additional annual Gross Value Added at maturity, alongside research-translation, commercialisation, and life-sciences returns. Each figure is illustrative, gate-dependent, and to be firmed up as the standing evaluation reports.

No specific funding figure, saving, or return is asserted here as audited, because none yet is; the figures above are illustrative planning ranges built on transparent assumptions (Chapter 7), and the precise Phase 1 envelope remains a matter for the Treasury business case. What is asked for now is a bounded and specific commitment approval of Phase 1 at the order of magnitude set out, together with the disciplined, phased, and evidence-gated mechanism through which the full national investment can subsequently be sized, committed and safely scaled, consistent throughout with HM Treasury Green Book principles, and with the evidence-led approach this report has applied to every claim within it.



Lawrence Tallon, Chief Executive of the Medicines and Healthcare products Regulatory Agency (MHRA), was among the senior leaders contributing to the Parliamentary roundtable discussion.

The Governance, Assurance, and Public Value Framework for the National Programme

Why Governance is Essential for the National Programme

The delivery programme set out in Chapter 8 depends, for its safety, its public value, and its credibility, on a governance framework of corresponding scope. Governance in this report is not a constraint applied after the fact to an otherwise complete strategy. It is the mechanism that makes investment defensible, that protects patients, research participants, and taxpayers from the specific harms the evidence has identified, and that sustains the public trust without which the programme cannot endure. This chapter governs the whole national capability described in Chapter 3 – secure health data, research infrastructure, high-performance computing, computational science, artificial intelligence, and quantum technologies – not artificial intelligence alone, because the risks and the safeguards against them differ by layer, and each requires its own proportionate treatment.

The clinical evidence surveyed in Chapter 4 contains a decisive lesson that applies equally across the wider programme: a capability that performs well in development may deliver no benefit, or active harm, in practice. Governance is the discipline that stands between that risk and the patient, the researcher, and the public. It is proportionate rather than uniform, calibrated to the autonomy of the tool or capability and the severity of the harm it could cause, and it operationalises, rather than duplicates, the delivery, funding, and accountability arrangements set out in Chapter 8.

The Regulatory and Standards Environment

The UK possesses a coherent, if still-developing, regulatory architecture spanning the layers of the national capability, resting on the coordinated responsibilities of several bodies rather than a single regulator for the whole programme. For artificial intelligence functioning as software as a medical device, the Medicines and Healthcare products Regulatory Agency retains oversight of conformity marking and is developing its approach to adaptive systems. The National Institute for Health and Care Excellence maintains the Evidence Standards Framework for digital health technologies and, through its conditional appraisal routes, permits use accompanied by evidence generation; it also sets the cost-effectiveness threshold of £25,000 to £35,000 per quality-adjusted life year from April 2026 (see Chapter 7, section 7.2, for the basis of this figure and its provenance) that governs commissioning. The Care Quality Commission provides provider-level assurance; the Information Commissioner's Office oversees the lawful processing of personal data; and the national data infrastructure described in Chapter 3 operates through the HDRS and SDEs as the governed route for research access.

This distribution of responsibility is a strength, but it carries a corresponding risk that oversight fragments across bodies, with no single accountable owner for the programme as a whole – a risk this chapter addresses directly in section 9.8. It is also a picture in transition: the Government's plan to abolish NHS England and absorb its functions into the DHSC, alongside the consolidation of integrated care boards described in Chapter 8, will change where specific regulatory and assurance functions sit during the life of this programme, and the governance model recommended here is deliberately built around functions rather than named bodies, so that it survives that transition.

Beyond clinical regulation, two further standards regimes bear on the widened scope of this programme. The first is the lifecycle regulation of adaptive artificial intelligence. This has historically been more developed internationally than domestically, but the position is changing: the Medicines and Healthcare products Regulatory Agency has signalled three specific reforms, due to be published via a dedicated commission report in September 2026, that would bring UK practice into closer fit with the adaptive, post-market model recommended throughout this report. The reforms comprise an explicit reweighting of regulatory risk, so that the risk of withholding a beneficial technology through inaction is weighed against the risk of approving one that later causes harm, rather than only the latter being visible to a regulator; a shift from a single, high pre-market barrier toward a series of proportionate checks conducted after market entry, described by the regulator as a hurdles race rather than a high jump, using sandboxes and controlled real-world test-beds; and a move toward context-specific regulation, which assesses the producer, the deploying institution, and the professional using a tool, not only the product itself, an approach with a direct analogy in financial-services regulation. The second standards regime relevant here is post-quantum cryptography, where internationally recognised encryption standards were finalised in 2024 and now form the basis of national cyber-security guidance, addressed in full in section 9.6.

A related governance issue concerns the statutory powers available to the Medicines and Healthcare products Regulatory Agency. The regulator has indicated that further legislative change may be required to support the full implementation of future adaptive regulatory approaches. The programme should therefore monitor both regulatory reform and the associated legislative changes as part of its governance framework.

Clinical Safety and Human Oversight

Clinical safety rests on two foundations: a documented safety case and the retention of human judgement. Every clinical deployment should carry a clinical safety case and hazard log, owned by a named clinical safety officer and prepared under the established NHS clinical-risk-management standards for both manufacturers and deploying organisations, alongside a documented clinical risk assessment proportionate to the tool's autonomy and its potential for harm. Human oversight is not discretionary. For any application that informs a clinical decision, a clinician must remain in the loop, as the National Institute for Health and Care Excellence already requires for the review of AI-generated radiotherapy contours, and tools should be used as an adjunct where the evidence supports assistance rather than autonomy, following the pattern of second-read pathology and stroke-imaging support. Levels of autonomy should be defined per tool, and autonomy permitted only where a system has been independently validated and its outcomes evidenced. A recurring and specific control is the calibration of alert thresholds, since poorly calibrated alerts cause fatigue and, paradoxically, undermine the very oversight they are meant to support.

Algorithmic Bias, Fairness, and Health Inequalities

Fairness is both a clinical-safety and an ethical requirement, and the evidence identifies specific, not hypothetical, risks. Genomic variant classifiers exhibit ancestry bias, with performance varying by genetic ancestry and reference datasets over-representing European populations; an administrative risk tool mis-flagged around half of the patients it assessed, creating a clear potential to disadvantage particular groups where such flags drive the uneven allocation of resources or transport. The controls follow directly: the representativeness of training data must be assessed before deployment, sub-group performance and false-flag rates must be monitored across ancestry, deprivation, and protected characteristics after deployment, and equity should be assessed within the impact-assessment and clinical-risk process, with equity metrics published. Fairness controls must also be designed to preserve positive gains, not merely to prevent harms: the evidence includes a case in which an AI self-referral tool substantially increased referrals among under-served groups, and such equity benefits should be actively protected. The same representativeness discipline applies to the national research cohorts described in Chapter 3: their value as a national asset depends on their continued breadth and representativeness, and this should be monitored alongside the deployed-tool metrics above.

Transparency and Explainability

Transparency serves both clinical scrutiny and public trust. Where a system informs a clinical decision, its output should be interpretable and explainable to the clinician relying upon it, supporting the exercise of the human oversight described above; interpretable diagnostic systems already in use demonstrate that this is achievable in practice. Transparency extends beyond the individual tool to the honest communication of evidence at programme level, distinguishing what is proven from what is merely promised across every layer of the capability described in Chapter 3, and reporting failures alongside successes, since over-claiming against thin evidence is itself a governance risk that erodes the confidence of clinicians, researchers, and the public alike.

Data Governance, Privacy, and Cybersecurity

Data governance should be approached as an enabler rather than a brake, on the principle borne out by international experience that protection and usable access are complementary rather than opposed. Development and training on NHS data should take place within SDEs, accessed through the HDRS as the governed route, with data-access turnaround benchmarked against the most efficient international comparators. This applies equally to the national research cohorts described in Chapter 3

– UK Biobank, Our Future Health, and the Genomic Medicine Service – which operate under the same SDE principle: researchers bring their analysis to the data, rather than extracting it. Data minimisation, purpose limitation, and linkage governed through the national patient identifier are requirements under UK data-protection law, and data quality must be assured, since fragmentation remains a binding constraint on the whole programme. A data protection impact assessment is mandatory for any processing of patient data, the Caldicott Principles apply to identifiable information, and the National Data Guardian provides independent oversight of data use.

Cybersecurity governance must now extend beyond the baseline NHS cyber posture to the specific risk that long-lived medical and research data captured today could be decrypted in future once sufficiently capable quantum computers exist – the “harvest now, decrypt later” risk described in Chapter 3. Internationally recognised post-quantum cryptography standards were finalised in 2024, and the recommended posture is to begin migrating the encryption of long-lived NHS and research data to these standards now, treating this as a scheduled security programme rather than an open-ended aspiration, with clear milestones for cataloguing at-risk systems, prioritising the longest-lived data, and completing migration in a phased manner consistent with national cybersecurity guidance. Beyond this specific migration, AI-specific standards addressing model theft, data poisoning, adversarial inputs, and the integrity of the model supply chain are not yet fully established, and the development of an NHS AI-cyber standard and assurance process remains a genuine and acknowledged gap. Governance should also recognise the physical and intellectual-property security of organisations developing advanced AI and quantum technologies. Protecting strategic research capability requires attention not only to data security but also to the security of research infrastructure, early-stage technology companies, and nationally important intellectual property. This consideration further supports the recommendation that strategically important quantum computing capability serving UK healthcare research should remain under trusted national governance arrangements. Secure model supply chains, updated integrity, and appropriate data-residency controls should be treated as the direction of travel, moving from recommended to mandatory as the standard matures.

Public Trust and Stakeholder Engagement

Public trust is a precondition of sustainable adoption across the whole programme and must be actively earned and maintained, extending to research participants in the national cohorts as well as to patients. This requires demonstrating, through SDEs and the governed access routes, that data is both protected and usable; treating equity as a trust issue through the visible publication of sub-group performance and open management of bias; and communicating data rights, lawful basis, and opt-out arrangements clearly, so as to avoid repeating the trust failures of earlier data programmes. Independent oversight, including genuine patient, participant, and public involvement of the kind embedded in the NHS evaluations already under way, should be a standing feature rather than an occasional consultation. Trust, once lost, is difficult to rebuild, and the framework treats its maintenance as an ongoing obligation across every layer of the capability, not only its clinical applications.

Accountability Across the Programme

Clear accountability is essential wherever a machine, a dataset, or a computational model contributes to a decision affecting patients, research participants, or public money, and this report matches that accountability directly to the governance bodies established in Chapter 8, rather than to an abstract or unnamed function. The Programme Board holds overall accountability for the programme’s public value and its go/no-go decisions at each evidence gate. The Scientific Advisory Board is accountable for the technical scrutiny of computational-science, AI, and quantum claims. The Clinical Advisory Board is accountable for clinical safety and validation, drawing on the Medicines and Healthcare products Regulatory Agency and the National Institute for Health and Care Excellence. Independent

assurance, sitting outside the delivery chain, is accountable for auditing the programme against its evidence gates and for reporting directly to the Programme Board. And regional delivery boards are accountable for local delivery and for feeding performance and evaluation data upward. Intellectual-property, data-use, and clinical-liability terms should be clarified at the point of procurement or partnership agreement, consistent with the public-value framework in section 9.11, rather than left to be resolved after an adverse event. The governing principle is that this structure, not any single named individual or body, must be able to answer for any element of the programme: whether it is validated, whether it is monitored, whether it is equitable, and whether a resourced response exists to act on its outputs. Accountability that cannot answer those questions is accountability in name only.

Independent Validation and the Evidence Gates

Independent validation is the single most important governance control in this framework, and it is the mechanism that operationalises the evidence gates set out in Chapter 8. It applies with particular force to predictive and risk-stratification tools. Validation must assess calibration, not merely discriminatory accuracy; it must be conducted at the deploying site and population rather than relied upon from the development setting; and it must include sub-group validation for equity. For predictive tools specifically, validation must demonstrate outcome benefit or the existence of a resourced intervention to act on the prediction, rather than accuracy alone. The operative rule, applied consistently across every phase of the programme, is that accuracy reported in a paper is not the same as safety and value in a given deployment: a capability advances to the next phase of Chapter 8's programme only on independent evidence, ideally generated within the NHS or the relevant national research infrastructure using the health-economics evaluations already in train as templates for equivalents across the priority applications.

Lessons From the Cautionary Evidence

The necessity of these controls is established not in the abstract but by the evidence discussed earlier in this report. A widely deployed commercial sepsis-prediction model, on independent external validation, achieved discrimination little better than chance and missed approximately two-thirds of cases – the direct justification for mandatory external validation and calibration before any predictive tool is used. A UK randomised trial of a population-health risk-stratification tool found that it increased emergency admissions without patient or system benefit – the direct justification for requiring a resourced downstream clinical response, since a prediction acted upon without a considered intervention can convert an intended benefit into a harm. These two cases are the empirical foundation of this framework: they demonstrate that the governance controls set out here are responses to observed failure modes, not precautionary formalities, and the same principle – independent scrutiny before scale-up – is applied to every layer of the wider programme below.

Procurement, Partnership, and Public Value

Procurement and partnership are governance instruments as much as commercial ones, and this section governs them consistently with the build-partner-access model and the public-value framework set out in Chapter 8. Where a capability is built as sovereign national infrastructure, procurement should embed the assessment gate and independent validation as pre-conditions of purchase, with pricing outcome-linked where applicable, and contracts should contain anti-lock-in terms open standards, application interfaces, data portability, and exit rights to prevent dependence on a single supplier, since interoperability remains the binding constraint on the wider system.

Where a capability is delivered through partnership, as recommended in Chapter 8 for research translation, computational science, quantum sensing, and clinical AI deployment, governance must additionally address the terms of that partnership: intellectual-property ownership and licensing arrangements that protect a public return; benefit-sharing with the NHS bodies and research participants whose data or clinical collaboration contributed to a discovery; guaranteed continued NHS and research access to any tool developed substantially using public data or public clinical partnership; and clear safeguards against commercialisation terms that fail to reflect the public origin of the underlying capital and data. Where a capability is simply accessed as a service, as recommended for general cloud infrastructure and, for now, most quantum-computing capacity, governance should focus on data residency, portability, and exit terms rather than ownership.

Vendor accountability is reinforced by transparency: reliance on proprietary, non-peer-reviewed performance claims is a recognised risk, to be countered by open standards and contractual transparency. Suppliers and delivery partners should be contractually required to support ongoing performance monitoring, provide clinically or scientifically necessary updates throughout the product or platform lifecycle, and cooperate fully with independent post-deployment evaluation. These obligations ensure that responsibility for safety, performance, and public value continues beyond initial procurement or partnership agreement, and they are the mechanism by which the taxpayer-value principles set out in Chapter 8 are made enforceable in practice.

Workforce Governance

Governance in this framework covers people as well as technology, and this is treated as a distinct and necessary discipline rather than an incidental concern. The specialist workforce described in Chapter 3 in data science and engineering, high-performance and research computing, computational biology, artificial-intelligence engineering, and quantum science requires its own governance because capability without accountable, competent people to operate it is not safely deployable capability at all. This has four practical elements. Clinical and scientific accountability requires named roles wherever a tool or model informs a decision: a clinical safety officer for clinical applications, and equivalent named scientific leads for computational-science and quantum-research activity, consistent with the accountability structure in section 9.8. Competency and training standards should be established for the specialist roles this programme creates or expands – clinical-AI safety officers, model-monitoring analysts, computational biologists, and quantum specialists – so that capability expansion does not outpace the assurance that the people operating it are appropriately qualified. Workforce diversity and front-line co-design, particularly with nursing and allied health professions, should be a governance requirement of every deployment, not an aspiration, since the international evidence is clear that neglecting the front line is a recognised failure mode. And the acute scarcity of specialist skills in quantum science and quantum error correction, noted in Chapter 3, is itself a governance risk to be tracked at programme level: a national capability cannot be assured as sustainable if it depends on a workforce pool too small to provide redundancy or independent scrutiny.

Quantum Governance: Sensing, Computing, and Post-quantum Cryptography

Quantum governance must, consistently with the treatment in Chapter 3, separate three distinct activities that carry different risks and different levels of maturity, and this framework does not apply a single governance standard across all three.

Quantum sensing is a developing but nearer-term capability, with the National Quantum Strategy committing that every NHS Trust will benefit from quantum sensing-enabled solutions by 2030.

Governance here should follow the same medical-device and clinical-safety pathway as any other early-stage diagnostic technology – MHRA regulation where applicable, NICE evidence standards, and the human-oversight principles in section 9.3 introduced proportionately as sensing applications approach clinical use, and not assumed to require a novel regulatory regime.

Quantum computing remains, on the evidence in Chapter 3, at a proof-of-concept stage for healthcare applications, and its governance should reflect that status: it is a strategic research investment, not an operational deployment, and this framework recommends against any governance approach that implies imminent clinical use. Research access, whether through national facilities or international partnership, should be governed as any other research-infrastructure investment: subject to the Scientific Advisory Board’s technical scrutiny, transparent about the gap between demonstrated capability and clinical application, and explicit that any future clinical tool arising from this research would, in due course, fall under the existing software-as-a-medical-device and cost-effectiveness routes described in section 9.2. No exception to those routes should be created in anticipation of quantum computing’s maturity.

Post-quantum cryptography, by contrast, is an actionable governance requirement now, and is addressed in full in section 9.6. It is a security migration, not a computing capability, and should be resourced and monitored as a scheduled programme rather than treated as part of the longer-horizon quantum-computing research agenda.

Research Translation Governance

Governance must also cover the pathway by which discoveries move safely from research to public benefit, set out in Chapter 3 and developed as an economic case in Chapter 7: research infrastructure, discovery, clinical validation, NHS deployment, commercialisation, and economic growth. Each transition along this pathway requires a distinct governance check. The transition from discovery to clinical validation requires the independent-validation standard set out in section 9.9, applied before any research finding is treated as a candidate for NHS deployment. The transition from validation to NHS deployment requires the clinical safety and human-oversight controls in section 9.3, and confirmation that the evidence gate for the relevant phase of the Chapter 8 programme has been met. The transition from deployment to commercialisation requires the intellectual-property, benefit-sharing, and public-value safeguards in section 9.11, so that public and philanthropic investment in the earlier stages of the pathway is reflected in the terms on which later commercial value is captured. Governing this pathway explicitly rather than assuming that governance applies only to the deployed clinical tool at the end of it is what allows the national research infrastructure described in Chapter 3 to generate public benefit safely and with public value protected throughout, and it is the reason this chapter, and not only Chapter 8, must speak to research translation.

Lifecycle Governance Across the Whole Capability

Governance in this framework applies across the full lifecycle of every capability in the national programme, not only before deployment: research, development, validation, deployment, monitoring, independent evaluation, and retirement or replacement. A capability that has passed validation and been deployed is not thereby exempt from further scrutiny; it enters continuous monitoring for drift, calibration, and equity, recorded in a national model and capability registry; it is subject to defined triggers for re-validation or retirement where it underperforms or drifts; and any resulting incident is captured through existing patient-safety and, where relevant, device-vigilance reporting, with a named accountable owner. This lifecycle view applies equally to research infrastructure and computational-science capability, where the relevant “monitoring” is of continued data quality, cohort representativeness, and computing capacity, and to quantum-sensing tools as they move toward clinical use, and it should be built into the evidence-gate assessments described in Chapter 8 rather than treated as a separate, later concern.

Running through the whole lifecycle is a risk-management approach in which each identified risk is paired with a specific, evidence-anchored control and an accountable owner, matched to the Chapter 8 governance bodies. Table 9.1 summarises the principal governance risks, their evidence base, the proposed control, and the accountable owner within the national programme.

Risk	Evidence Anchor	Control	Owner
Predictive failure (works in trial, not in NHS)	Sepsis model, AUC 0.63, missed ~2/3 cases	Independent external validation + calibration before go-live	Clinical Advisory Board/MHRA
Prediction-induce d harm	UK RCT increased emergency admissions	Resourced downstream clinical response for every prediction	Clinical Advisory Board/ Regional delivery boards
False positives at scale	Risk tool ~50% mis-flag rate	Post-deployment false-flag monitoring; calibrated alert thresholds	Deploying trust/ Independent assurance
Bias against under-served groups	Genomic ancestry bias	Pre-deployment representativeness check; sub-group monitoring; published equity metrics	Scientific Advisory Board/NICE
Automation bias	Alert fatigue undermines oversight	Human-in-the-loop for all clinical decisions; defined autonomy tiers	Clinical safety officer
Deployment outpacing evidence	Stroke, AIDF claims programme-reported	Assessment gate + funded evaluation as a condition of deployment (Chapter 8 evidence gates)	Programme Board
Vendor lock-in	Proprietary, non-peer-reviewed claims	Anti-lock-in terms: open standards, portability, exit rights	Programme procurement function
Cyber compromise, including future decryption of long-lived data	AI-specific standards not yet established; quantum decryption risk	NHS AI-cyber standard; scheduled post-quantum cryptography migration	Independent assurance/DBIST
Unlawful data processing	Trust failures of past data programmes	DPIA + lawful basis for every deployment	ICO/National Data Guardian
<i>Workforce exclusion or capability gap</i>	<i>Neglecting the front line is a known failure mode; scarce quantum/computational skills</i>	<i>Mandatory clinical co-design; workforce governance (section 9.12)</i>	<i>Regional delivery boards/ Programme Board</i>
<i>Research-translation failure or loss of public value</i>	<i>Discovery commercialised without public benefit-sharing</i>	<i>IP and benefit-sharing terms set at procurement/partnership stage</i>	<i>Programme Board/ Independent assurance</i>

Governance Recommendations

The framework resolves into a concise set of recommendations, addressed to the bodies established in Chapter 8 rather than to institutions in the process of reform. Government and the Programme Board should

 Make independent external validation and calibration a hard evidence-gate requirement for all predictive and risk tools and require a resourced downstream clinical response.

 Mandate human oversight for all clinical decisions, with autonomy levels defined by risk tier.

 Require equity assurance: pre-deployment representativeness assessment and post-deployment sub-group and false-flag monitoring, extended to the national research cohorts.

 Confirm the Programme Board, Scientific Advisory Board, Clinical Advisory Board, and independent assurance function as the single accountable governance architecture for the whole programme, designed to remain stable through the NHS England transition described in Chapter 8.

 Adopt adaptive AI lifecycle regulation with mandatory post-market surveillance, building on the direction of regulatory reform already signalled by the UK's regulator, while applying the same lifecycle governance principles to research infrastructure, computational science, and future AI capability.

 Treat data governance as an enabler, using SDEs, and begin the scheduled migration to post-quantum cryptography for long-lived data.

 Govern quantum sensing, quantum computing, and post-quantum cryptography separately and proportionately, treating quantum computing as strategic research rather than operational deployment.



Establish workforce governance as a standing function, covering competency, accountability, and diversity across clinical, computational, and quantum roles.



Apply the build-partner-access and public-value principles in Chapter 8 to every procurement and partnership agreement, with intellectual-property and benefit-sharing terms set at the outset.



Govern the research-translation pathway end to end, from discovery through validation and deployment to commercialisation.

Chapter Summary

Safe, ethical, and accountable delivery of the national programme set out in this report is achievable, but only within a governance framework proportionate to risk, anchored in the evidence, matched to the bodies established in Chapter 8, and extending across the whole capability rather than to artificial intelligence alone. The essential guardrails are consistent throughout: human oversight of decisions, independent validation before progression through the evidence gates, a resourced response behind every prediction, explicit equity monitoring, workforce competency and accountability, proportionate and separated quantum governance, protected public value through the research-translation and commercialisation pathway, and lifecycle oversight that continues long after initial deployment. The cautionary evidence in this chapter demonstrates that these are responses to real and observed failures rather than precautionary formalities, and the framework openly acknowledges where standards in cybersecurity, in quantum-specific assurance, and in workforce capability remain to be completed. Taken with the delivery programme in Chapter 8, this framework provides the means to realise the benefits of the national programme while protecting patients, research participants, and public value, and it prepares directly for the conclusions and government decision set out in Chapter 10.



Conclusion and Strategic Recommendations

The Central Finding

The principal conclusion of this report is not that artificial intelligence should be adopted across the National Health Service. It is that sustainable value in health, in research, and in the economy depends on investing first in national capability: secure health data, research infrastructure. The programme should therefore be organised around priority healthcare problems and measurable outcomes, with technologies assembled according to the requirements of each use case rather than treated as separate ends in themselves.

High-performance computing, computational science, a skilled workforce, and disciplined governance, with evidence-led AI deployment as one application of that capability rather than the strategy itself. This is the thread that runs through every chapter of this report. The NHS's underlying problem, set out in Chapter 2, is structural: a projected workforce shortfall of hundreds of thousands, an elective waiting list in the millions, and productivity below both its historical average and its committed trajectory. The technology foundation set out in Chapter 3 shows why capability, not any

single tool, is the proper unit of investment: data, computing, and computational science are the enabling layers on which artificial intelligence, research translation, and quantum technologies alike depend. And the socio-economic case in Chapter 7 shows precisely why this distinction matters for public money: value is determined by the terms on which capability is built, procured and deployed, not by the sophistication of the technology acquired.

This reframing has a direct consequence for how the recommendations that follow should be read. Where earlier assessments of NHS technology have asked, “Should the health service adopt AI?” this report asks a different and more consequential question: does the UK possess, or is it building, the national capability required to convert artificial intelligence, computing, and quantum technologies into health benefit, research leadership, and economic growth? The answer, on the evidence, is that the country holds substantial existing assets and a credible path to build the rest, provided investment proceeds through the evidence-gated mechanism this report sets out, rather than through enthusiasm for any particular technology.

Why Now

Delay and premature deployment carry different but comparable risks, and the timing case for action rests on weighing them together rather than treating caution as costless. On one side, the NHS pressures described in Chapter 2 are intensifying rather than stabilising, and the research opportunity described in Chapter 3 – world-leading cohorts, a mature computing base, and a genomics service without international equal – is a perishable advantage: capability under-used today is capability that other systems will build first. The international evidence in Chapter 5 and the competitiveness analysis in Chapter 7 are consistent on this point: comparable economies are investing at scale in computing, data, and quantum capability, and the UK’s life-sciences ambition to become one of the world’s leading life-sciences economies depends on maintaining pace with that investment, not merely matching its own past trajectory. Regional development, too, is a time-sensitive opportunity: the innovation clusters and AI Growth Zones described in Chapter 3 are attracting private investment now, and a national programme that arrives after that investment has located elsewhere secures less regional benefit than one that shapes it from the outset.

On the other side, the cautionary evidence assembled in Chapter 4 – models that perform well in development and fail on independent validation, deployment that outpaces evidence and produces no benefit – is a direct warning against large-scale, unevaluated rollout. The clearest international failure mode, examined in Chapter 5, is strategy and guidance pursued without funded, evaluated deployment. The programme recommended in this report is designed precisely to hold both risks in view at once: it does not wait for a complete evidence base before acting, because the cost of delay is real; and it does not commit to national scale before the evidence justifies it, because the cost of premature deployment is equally real and better documented. This balance – acting now, but through evidence gates rather than ambition – is the central judgement of the report and the organising principle of the programme in Chapter 8.

The Programme Now Available to Government

The Government is no longer choosing whether to have a strategy; it has one, set out in full in Chapter 8. That programme provides a phased structure proceeding through foundations, evaluated expansion, and a mature learning-health-system capability; evidence gates governing the release of funding between phases, tying capital to independent evaluation rather than to the calendar; a governance architecture with clear lines of accountability from a Programme Board through scientific, clinical, and economic assurance to regional delivery; a delivery architecture that distinguishes what national government should build, what is best delivered through partnership with universities, industry, and philanthropy, and what should simply be accessed as a service; an evaluation framework

generating, for the first time, audited NHS-specific evidence of clinical and economic effect; a blended funding model in which government, philanthropy, research councils, universities, industry, and the NHS each contribute what they are best placed to provide; and a national-to-regional delivery chain connecting shared national infrastructure to integrated care systems, NHS trusts, universities, and industry, designed to remain workable through the institutional changes now under way in the NHS's own architecture. What remains is not to design this programme but to approve its first phase and hold it to the evidence-gated standard by which it is meant to proceed.

What Success Looks Like

Government should judge the programme's success across the full range of its strategic objectives rather than through technology adoption alone. In healthcare, it means measurably improved patient care in the applications with the strongest evidence, and higher NHS productivity where released capacity is genuinely redeployed rather than merely reported as saved. In research, it means stronger research translation discoveries moving from the national cohorts and computational-science base described in Chapter 3 into clinical validation and NHS use and the commercialisation this translation makes possible, in the form of licensing, spinouts, and university-industry collaboration. Regionally, it means the arrival of genuine innovation clusters beyond the established south-eastern concentration, contributing to regional economic development in the manner set out in Chapter 7. Nationally and internationally, it means a demonstrable contribution to the UK's life-sciences ambitions and its standing in international competitiveness, particularly relative to the scale of investment other systems are now committing to computing and quantum capability. And throughout, it means governance that is trusted by clinicians, by researchers, by industry partners, and by the public, because it has been transparent about uncertainty, has reported failure alongside success, and has protected public value in the terms on which intellectual property and data access are shared. Economic outcomes should not be given false precision where no audited baseline yet exists; establishing those baselines is among Phase 1's first outputs. Delivery success can nevertheless be specified in advance through observable programme milestones

- At least three healthcare applications have moved beyond pilot stage into sustained NHS procurement or deployment, with independent evidence on clinical, workflow, and economic effect.
- The UK has a recognised route for evaluating, regulating, and continuously assuring novel AI and quantum-enabled health technologies, with clear triggers for revalidation after material change.
- National laboratories, research infrastructure, and NHS users can benchmark competing systems and access appropriate compute securely without premature lock-in to one hardware pathway.
- Private and philanthropic capital has been crowded in on transparent terms that preserve public value, and participating companies show measurable progress in revenue, employment, manufacturing, investment, or exports in the UK.
- Regional delivery partnerships have defined roles and measurable gains in adoption, patient access, jobs, skills, research infrastructure, and company growth, with evidence that successful models transfer between different settings.
- The programme publishes unsuccessful as well as successful results, uses evidence gates to stop or redirect underperforming projects, and has begun an organised transition towards post-quantum security.

The Decision Before Government

HM Treasury, the DHSC, and the DBIST are being asked to approve a bounded Phase 1 commitment and the mechanism through which later investment would be released. The decision includes the following commitments:

- Phase 1 remains of the order of £250–500 million of incremental, health-directed investment over five years, front-loaded and predominantly government-funded, with matched philanthropic co-funding applied to defined eligible components and private co-investment increasing as evidence and procurement opportunities mature.
- Phase 1 investment in the foundational data, computing, and clinical capabilities set out in this report, of the order of £250–500 million over five years, sized through the Treasury business case.
- The national governance architecture – a Programme Board with scientific, clinical, and independent assurance functions, accountable throughout the five-year programme.
- A standing evaluation programme sufficient to generate the audited clinical and economic evidence currently missing.
 - An initial portfolio of approximately three to five problem-led demonstrators, with funded implementation capacity, common baseline measures, and a defined procurement or scaling decision at the evidence gate.
 - A procurement pathway enabling validated technologies to move from demonstrator to competitive public-sector purchasing and scale, rather than relying on grants alone.
- Continued investment in the national research infrastructure secure data, computing, and computational science set out in Chapter 3.
- A government-philanthropic partnership framework with a defined philanthropic-eligible sub-envelope and agreed matching rule, within the wider blended-funding and build-partner-access models in Chapter 8.
- A national research-translation programme connecting discovery to NHS deployment and commercialisation, with the public-value and intellectual-property safeguards set out in Chapter 8.
- The evidence-gated release of future funding, so that later phases proceed only as evidence from earlier phases justifies.

This is a bounded and specific ask, not an open-ended commitment. The recommended Phase 1 envelope is of the order of £250–500 million over five years; the illustrative returns of the order of 500–2,000 net additional jobs and an NHS ten-year net present value in the region of £0–100 million in Phase 1, rising to 5,000–15,000 net additional jobs and £450 million–£2 billion of net additional annual GVA at maturity are set out in Chapter 7, are illustrative and gate-dependent, and are precisely what the standing evaluation is designed to convert from modelled to audited. The precise investment envelope remains, correctly, a matter for the Treasury business case. What is requested

is approval of Phase 1 and of the phased, governed, and evidence-gated mechanism through which the full programme can be sized, committed, and safely scaled.

Departmental Action Plan

Lead	First Action After Approval	Primary Accountability	Early Success Measure
HM Treasury	Approve Phase 1 envelope, evidence gates, and benefits-accounting rules	Value for money; release of later tranches	Non-overlapping benefits map, parameter register, and agreed gate criteria
DHSC and NHS	Select initial demonstrators and fund implementation/evaluation capacity	Clinical value, adoption, and procurement	Named users, baselines, workflow plans, and procurement route for every demonstrator
DBIST, UKRI, and national laboratories	Secure compute, technical benchmarking, application support, and UK-wide capability coordination	Research/industrial capability and technical independence	Benchmark access in place; UK-wide portfolio roles agreed
MHRA and assurance bodies	Define iterative assurance, testbed, and revalidation requirements	Safety, effectiveness, and lifecycle assurance	Clear entry, surveillance, and revalidation triggers
Regional and devolved partners	Agree delivery and learning roles linked to national outcomes	Place-based adoption, skills, and growth	Common metrics and demonstrator participation beyond leading centres

Conclusion

The UK already possesses many of the national assets this programme requires: a single-payer health system with a national patient identifier and cohorts without international equal; a rapidly maturing national computing base; a coherent, if still-developing, regulatory architecture; and early, credible investment in quantum sensing, computing, and data infrastructure. The remaining challenge is not invention but coordination connecting these assets deliberately, through sustained and evidence-gated investment, rather than leaving their value to accumulate by chance or dissipate through delay.

This report has set out a mechanism for making that connection: national capability built and governed first; artificial intelligence and quantum technologies deployed within it where the evidence justifies; value captured across the health system, the research base, and the wider economy; and funding released only as independent evidence confirms that value is being realised. This is the proposition put to the Government. Investment should proceed through evidence, governance, and evaluation, not through optimism about any single technology or unmanaged enthusiasm for rapid deployment. Approached this way, the programme offers the UK a genuine and proportionate opportunity: to convert national capability already substantially in place into lasting benefit for patients, for scientific and economic leadership, and for the public finances that must ultimately answer for the investment. That is the basis on which this report recommends the Government proceed. If delivered in this way, the programme represents not only an opportunity to modernise the NHS but also a long-term investment in the UK's scientific capability, economic resilience, and international competitiveness. The next step is not another broad strategy but a Treasury business case that confirms the initial demonstrator portfolio, delivery partners, year-by-year cash profile, matched philanthropic commitments, implementation budgets, procurement routes, and measurable gate criteria.

References

The references below compile, in alphabetical order, all published sources cited throughout this report. Each reference corresponds to a source cited in the text and is included to support the evidence, analysis, and recommendations presented.

- Curia (2026). AI, Data and Quantum in UK Healthcare: From Evidence to Investment. Parliamentary Policy Roundtable, 13 July 2026, House of Lords. Commissioned by Cleveland Clinic Philanthropy UK.
- Cleveland Clinic (2024). Cleveland Innovation District partnership overview and reported outcomes.
- Department for Science, Innovation and Technology (2023). National Quantum Strategy.
- Department for Science, Innovation and Technology (2025). AI Opportunities Action Plan.
- Department for Science, Innovation and Technology (2025). Compute Roadmap.
- Department for Science, Innovation and Technology (2025; 2026). AI Growth Zones designation and expansion.
- Department for Science, Innovation and Technology (2026). Isambard-AI and Dawn compute upgrades; sovereign compute commitments.
- Department for Science, Innovation and Technology (2026). March 2026 quantum scaling package, including the “ProQure” programme.
- Department for Science, Innovation and Technology (2026). UK AI-sector investment figures.
- Department for Science, Innovation and Technology (2026). Life Sciences Innovative Manufacturing Fund.
- Department for Science, Innovation and Technology and Department of Health and Social Care (2025). Life Sciences Sector Plan.
- Department of Health and Social Care (2025). Announcement on the abolition of NHS England and absorption of functions into DHSC (March 2025); DHSC integration timetable targeting completion by April 2027.
- Department of Health and Social Care (2026). AI Diagnostic Fund funding and deployment update, including chest radiography and stroke-imaging adoption.
- Department of Health and Social Care and NHS England (2025). Departmental budget and annual accounts, 2024/25.
- Department of Health and Social Care and NHS England (2026). HDRS and SDEs funding and design.
- Department of Health and Social Care and NHS England (2026). Ten-Year Health Plan.
- DHSC/NICE (December 2025). Confirmation of the NICE cost-effectiveness threshold change to £25,000-£35,000 per QALY from April 2026, as part of the UK-US Economic Prosperity Agreement.
- Federal Institute for Drugs and Medical Devices [Germany] (2024). Digital health reimbursement pathway prescriber uptake data.
- Food and Drug Administration [United States] (2025). AI-enabled medical device authorisations.
- Google Quantum AI (2025). Error-correction scaling result.
- Government of Finland (2020). Findata and the Finnish national data strategy.
- Health Foundation REAL Centre (2024). NHS funding projections and cost-pressure analysis.

- House of Commons Library (2026). Life sciences sector SIC-based GVA and employment measure.
- Jumper, J. et al. (2021). AlphaFold protein structure prediction.
- Ministry of Food and Drug Safety [South Korea] (2022). Adaptive AI medical device authorisations.
- Ministry of Health [Singapore] (2023). National diabetic retinopathy screening programme.
- McKinney, S. M. et al. (2020). International evaluation of an AI system for breast cancer screening.
- National Audit Office (2020). Digital transformation in the NHS.
- National Audit Office (2024). NHS Financial Management and Sustainability; examination of NHS Long Term Workforce Plan modelling.
- NHS England (2023). NHS Long Term Workforce Plan.
- NHS England (2026). Referral to Treatment Waiting Times, April 2026.
- NHS England (2026). Cancer Waiting Times.
- NHS England (2026). Genomic Medicine Service variant prioritisation sub-study.
- NHS England (2026). Ambient voice technology national guidance and supplier registry.
- NHS England (2026). Integrated Care Board mergers and boundary changes, effective April 2026 and April 2027.
- Nobel Prize in Chemistry (2024). Prize citation AlphaFold / protein structure prediction.
- npj Digital Medicine (2025). Systematic review of quantum machine learning for digital health.
- Office for Life Sciences (2025). UK life sciences sector turnover and employment.
- Office for National Statistics (2024). Public service healthcare productivity.
- Organisation for Economic Co-operation and Development (2026). Healthcare AI adoption across member countries.
- Royal College of Radiologists (2025). Clinical Radiology Workforce Census.
- UK Biobank (2025). Cohort overview and relocation funding.
- UK Research and Innovation (2024; 2026). Our Future Health; Isambard-AI and national compute facilities.
- US Department of Commerce (2026). Quantum-sector incentive programme.
- US Department of Energy (2026). Fault-tolerant quantum computing demonstration programme.
- World Health Organization (2021). Global framework for the ethics and governance of AI in health.
- Copenhagen Capacity (n.d.). Medicon Valley the leading Life Science Cluster of the Greater Copenhagen Region. copcap.com/medicon-valley.
- Labiotech (2024). Medicon Valley shapes life sciences in Denmark and Sweden. labiotech.eu.
- Medicon Valley Alliance (n.d.). About MVA; Medicon Village formation. mva.org.
- NICE (2025-2026). Changes to NICE's cost-effectiveness thresholds confirmed. nice.org.uk.
- Civil Service World / HSJ (January 2026). DHSC timetable for completion of the NHS England merger, targeting April 2027.
- Chapter and section cross-references throughout this report (e.g. "NICE, 2024" for radiotherapy auto-contouring costs, "NHS England, 2026" for the Genomic Medicine Service sub-study) point to the specific named publication in this list; where the same organisation is cited multiple times in different years or for different publications, the in-text citation and this list distinguish them by year and subject.



