How the UK can become an

AI Adoption Superpower



BDO

British Chambers of Commerce



Innovation creates possibility. Adoption creates prosperity.

The future of artificial intelligence will be defined not only by the technologies we create, but by how effectively we use them.

Each chapter offers a unique perspective on the opportunities and challenges ahead, providing practical ideas for policymakers, business leaders and innovators working to build a more productive, competitive and for the UK.

Each chapter in this publication has been written independently by its author and reflects their individual views and perspectives.

The views expressed do not necessarily represent those of the British Chambers of Commerce (BCC), and inclusion in this publication should not be taken as BCC endorsement of all the opinions, proposals or recommendations contained within.

Contents

Introduction 06

Foreword: Shevaun Haviland CBE, Director General, British Chambers of Commerce 12

Foreword: Rt Hon Kanishka Narayan MP, AI Minister 16

Why Britain must act 18

01 AI's darkest hour 20
Alexandra Mousavizadeh *Chief Executive, Evident*

02 Britain won't become an AI Superpower by cheering from the sidelines 24
Lord Ed Vaizey *Business Advisor, Broadcaster and Member of The House of Lords*

03 Bringing technology and good governance together 28
Dame Chi Onwurah MP *Member of Parliament, Newcastle upon Tyne Central and West Chair, Science and Technology Committee*

04 Britain's AI opportunity begins with becoming better customers 32
Stephen Kelly *Chief Executive Officer, Cirata*

05 The state must get its own AI house in order 38
Ed de Minckwitz *Director of Public Policy, ServiceNow*

06 Closing the gap between inventing well and absorbing badly 44
Julian David *Chief Executive, TechUK*

07 Turning AI promise into UK prosperity 52
Vinay Singhvi *Head of UK & Ireland, Tata Consultancy Services*

How adoption creates economic value 58

08 The middle way: How mid-sized businesses can lead the UK's path to AI Superpower status 60
Catherine Wilks *Partner and UK lead for Data, Analytics and AI, BDO*

09 The constraint is moving 66
David Bharier *Deputy Director, Economics and Insight, The British Chambers of Commerce*

10 Depth over speed: A path to making Britain the first AI Adoption Superpower 72
Raoul Ruparel OBE *UK Chief Economist and Senior Director, Boston Consulting Group*

11 Unlocking Britain's next era of productivity: Building a nation of AI trailblazers 78
David Black *Managing Director, Google UK*

12 From AI Adoption to advantage: Turning Britain's head start into growth 84
Richard Breeden *Founder and CEO, aibl (AI For Business Leaders)*

13 Britain has adopted AI. Now it needs to go deeper 90
Rob Burlison *Director of Global Corporate Affairs and Public Policy, Intuit*

14 From AI users to AI transformers 96
Alison Wright *Small & Medium Business Director, Microsoft UK & Ireland*

15 Turning AI into everyday SME growth 102
Abdul Quershi *Managing Director, Barclays Business Banking*

Why people determine success 110

16 The real AI challenge is not technology. It is change 112
Professor Kirstine Dale *Chief AI Officer, Met Office*

17 Is the UK's AI leadership stuck in analogue mode? 118
Ann Francke OBE *CEO, Chartered Management Institute*

18 The interaction of people and technology is at the heart of AI growth 124
Rt Hon Chloe Smith *President, Chartered Institute of Personnel and Development*

19 To reimagine the future of work we need to listen to workers 128
Tabitha Goldstaub *Co-founder and CEO, Lichen AI*

20 AI skills in every apprenticeship 132
Euan Blair MBE *Founder and CEO, Multiverse*

21 AI will not shape the UK by itself. We will. 136
Baroness Martha Lane-Fox of Soho *Digital entrepreneur, Crossbench peer and former BCC President*

22 Britain's AI problem is not adoption. It is capability 142
Professor Nisreen Ameen *Professor of Artificial Intelligence and Policy, University of Reading*

How to scale with trust and resilience 148

23 Why trust comes before tech 150
Kate Hayward *UK Managing Director, Xero*

24 AI Adoption is an operations problem, not a technology problem 156
Jason Mashinchi *Founder & Managing Director, Cambridge Kinetics Board Director, Cambridgeshire Chamber of Commerce*

25 Seizing adoption advantage for business 162
Dr James Kuht MBE *CEO and Co-founder, PAIR*

26 The adoption challenge we can still win 168
Robyn Scott *CEO and Cofounder, Apolitical*

27 AI governance: The foundation of trust, growth and competitiveness 174
Helen Wylde-Archibald *CEO of Devon Chamber of Commerce Co-Founder of South West AI Solutions Limited*

28 Why assurance, auditing and governance must be central to AI leadership 182
Dr Allison Gardner MP *Member of Parliament for Stoke-on-Trent South and Chair of the APPG on Artificial Intelligence*

29 What stands in the way becomes the way 188
Ollie Whitehouse *Chief Technology Officer, National Cyber Security Centre*

References 192

Introduction

Artificial intelligence is often discussed as a race. A race to develop the most powerful models, attract the greatest investment or create the next generation of globally successful technology companies.

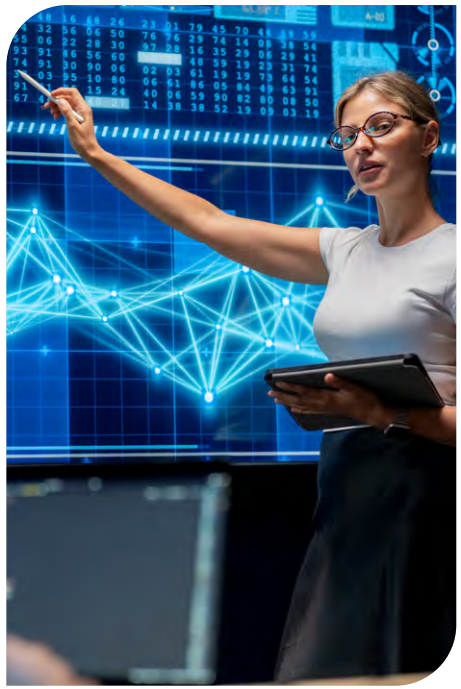
Those ambitions matter. But for the United Kingdom, another race may prove more important: becoming one of the world's most successful adopters of AI.

As a country we have formidable strengths in research, innovation and entrepreneurship. There is an openess to AI within British businesses, not found in other markets. The challenge now is to leverage this, seizing the opportunity of AI adoption to move beyond experimentation, and instead transform the way businesses operate.

Used well, it can improve productivity, create better products and services, strengthen competitiveness and support new forms of work and economic growth

This collection of essays brings together leaders from business, politics, technology and academia to examine how Britain can make that transition. Their perspectives differ, but a common argument emerges - successful adoption requires more than acquiring technology. It depends on leadership, organisational change, skills, investment and trust, plus a willingness to rethink how organisations work.

In her foreword the **BCC's Director General Shevaun Haviland** argues that the opportunity is there, but it will not wait for us. Britain now needs to seize the moment, adopting AI faster, more deeply and more widely by businesses across the economy. She argues the prize is not simply technological leadership, but greater productivity, competitiveness and growth.



Why we must act

The opening contributions examine why AI adoption has become an economic imperative and what the UK needs to do differently.

The Chief Executive of Evident, **Alexandra Mousavizadeh**, argues that adoption is a race we can win. Alexandra writes that rather than competing solely to develop the largest frontier models, the UK can excel at applying AI throughout its economy. She argues success will depend not just on invention, but on whether established companies, smaller businesses and public institutions can use the technology to solve real problems.

Lord Ed Vaizey, a former Conservative culture and digital minister, considers how the UK can build on its existing technology strengths. He describes how the challenge is to turn the quality of our technology ecosystem into wider adoption and greater economic impact.

The Labour backbencher **Dame Chi Onwurah**, who's Chair of House of Commons Science and Technology Committee, examines the conditions needed to realise that opportunity. They include regulation, enforcement capacity, digital sovereignty and inclusive growth. Her argument connects adoption to the kind of economy and society we want to build. She writes how the benefits of AI need to extend beyond a relatively small number of businesses, sectors and places.

Chief Executive Officer of Cirata, **Stephen Kelly**, brings the demand side into focus, insisting that innovative businesses need customers. Stephen describes how more confident purchasing by British organisations and government can create the markets that help technology companies grow.

Ed de Minckwitz, Director of Public Policy at ServiceNow UK, extends this argument to the public sector, asking how government can lead as an adopter as well as a policymaker. He argues that government must transform its own systems and use the state as a proving ground for AI, rather than simply encouraging businesses to adopt it.

Julian David, Chief Executive of the trade body techUK, examines the related challenge of diffusion, how innovation spreads from the technological frontier into the wider economy.

Meanwhile, **Vinay Singhvi**, Head of UK and Ireland at Tata Consultancy Services, brings the discussion back to delivery, examining how national ambition can translate into practical implementation.

How adoption creates economic value

From the national imperative, the collection of essays turns to the business case, examining how adoption can deliver measurable economic value.

The section begins with a contribution from the UK Lead for Data, Analytics and AI at BDO, **Catherine Wilks**, asks why enthusiasm for AI has not always translated into financial returns. By exploring the gap between activity and value, she makes the case for judging AI programmes against business objectives, rather than the number of experiments undertaken.

David Bharier, Deputy Director, Economics and Insight at the British Chambers of Commerce argues that realising AI's full value requires businesses to rethink how work is organised. Adding a tool to an existing process may make a task faster; redesigning the process can deliver much greater gains. He concludes that adoption is therefore about changing how a business operates, not simply what technology it buys.

Raoul Ruparel, UK Chief Economist and Senior Director at Boston Consulting Group, develops this argument by challenging organisations to move beyond pilots. Experimentation matters Raoul argues, but lasting value comes from concentrating on strategically important applications and taking them through to implementation.

Managing Director of Google UK, **David Black**, connects these decisions inside individual firms to the UK's wider productivity challenge. He describes how technology can help organisations generate more value from their people, capital and expertise.

Richard Breeden, Founder and Chief Executive Officer of AI for Business Leaders, considers how organisations can move from technical capability to genuine commercial advantage.

Meanwhile, **Rob Burlison**, Director of Global Corporate Affairs and Public Policy at Intuit, examines the importance of embedding AI into core business processes rather than treating it as an additional tool at the margins.

Alison Wright, Small and Medium Business Director at Microsoft UK and Ireland, brings the discussion to smaller businesses and the practical support they need to turn interest into transformation.

Abdul Qureshi, Managing Director of Barclays Business Banking, considers how AI can become embedded in everyday SME activity and growth. He describes how economy-wide adoption depends on these firms as much as on large corporations and technology companies. Abdul argues that for SMEs, progress often begins with a practical problem, an accessible solution and the confidence that comes from seeing tangible improvements.

Why people determine success

Realising that the value of AI depends on people, the third section examines the leadership, skills and workforce involvement needed to make adoption succeed.

Professor Kirstine Dale, Chief AI Officer and Principal Fellow for Data Science at the Met Office, examines what organisational change looks like in practice. Beyond choosing technology, Kristine suggests leaders must explain why change is necessary, bring people with them and address what new tools mean for their roles and ways of working.

Ann Francke, Chief Executive of the Chartered Management Institute, examines the responsibilities of leaders steering organisations through rapid technological change. She argues that investment in technology must be matched by stronger management capability and organisational change if businesses are to move beyond pilots and realise genuine productivity gains.

The former Conservative minister **Chloe Smith**, who's now President of the Chartered Institute of Personnel and Development, considers the relationship between people and technology. Chloe suggests that AI adoption is as much a people challenge as a technology challenge, with HR leaders central to building the skills, trust and confidence needed to use it well.

Tabitha Goldstaub, Co-founder and Chief Executive Officer of Lichen AI and former Chair of the UK Government's AI Council, makes the case for listening to workers as adoption scales. Workers, she argues, should help shape adoption rather than simply experience its consequences.

The Founder and Chief Executive of Multiverse, **Euan Blair** focuses on one of the most important foundations of adoption, skills. He writes how adoption at scale will require people throughout organisations to acquire new capabilities; it cannot rest solely with AI specialists in technology departments.

Chair of the London AI and Jobs Taskforce and Ambassador for the British Chambers of Commerce, **Martha Lane-Fox**, broadens the discussion to how people can shape the future of work. She argues that government must move faster on support, skills and labour-market intelligence so AI raises productivity without leaving smaller firms, younger workers or entry-level routes behind.

Nisreen Ameen, Professor of Artificial Intelligence and Policy at Henley Business School, University of Reading, draws these strands together by asking what makes an organisation genuinely AI-capable. She concludes that alongside technology, it needs leadership, workforce confidence, skills and structures that allow it to adapt.

How we scale with trust and resilience

As adoption grows, so does the need for confidence in how AI is used.

The final section explores how Britain can move faster while building trust, accountability and resilience.

UK Managing Director of Xero, **Kate Hayward**, explores trust as a condition for successful adoption. She argues that practical local support from Chambers, advisers and trusted platforms can help small businesses turn curiosity and caution into confident, everyday use of AI.



Jason Mashinchi, Founder and Managing Director of Cambridge Kinetics and Board Director of Cambridgeshire Chamber of Commerce, considers the discipline needed to move from enthusiasm to responsible operational use.

Meanwhile, **Dr James Kuht**, Chief Executive Officer and Co-founder of PAIR, focuses on bringing AI into real organisational settings. He argues that Britain's advantage will come from helping firms use AI faster and better, embedding it in everyday work rather than treating invention alone as the measure of success.

Robyn Scott, Chief Executive Officer and Co-founder of Apolitical, examines how government can lead by example. She argues that the public sector should set the standard by using AI to augment people, redesign workflows and turn effective adoption into a national advantage.

Helen Wylde-Archibald, Chief Executive of Devon Chamber of Commerce and Co-founder of South West AI Solutions, examines governance as a foundation for confident innovation. The challenge is to establish clear accountability without losing the ambition to experiment and improve.



The Labour MP **Dr Allison Gardner**, Chair of the All-Party Parliamentary Group on Artificial Intelligence, develops this theme through assurance and auditing. She argues as AI becomes embedded in important decisions and processes, organisations need credible ways to assess performance and identify where problems arise.

Finally, Chief Technology Officer at the National Cyber Security Centre, **Ollie Whitehouse**, closes the collection with cyber security and resilience. Security must be built into adoption from the outset if AI is to become a dependable part of critical business activity.

Together, these contributions make the case for a different approach to adoption: moving from ambition to implementation, from pilots to scale, and from access to technology to the capability to use it well. Progress depends on bringing investment, people and organisational change together, with trust, governance and resilience built in from the outset.

Britain's opportunity is not simply to be a country where world-leading AI is invented, but one where businesses of every size and sector know how to use it well.

Turning that capability into widespread adoption, and then into productivity and growth will make AI not just a technology story, but an economic one.



Shevaun Haviland CBE
Director General
British Chambers of Commerce

Foreword

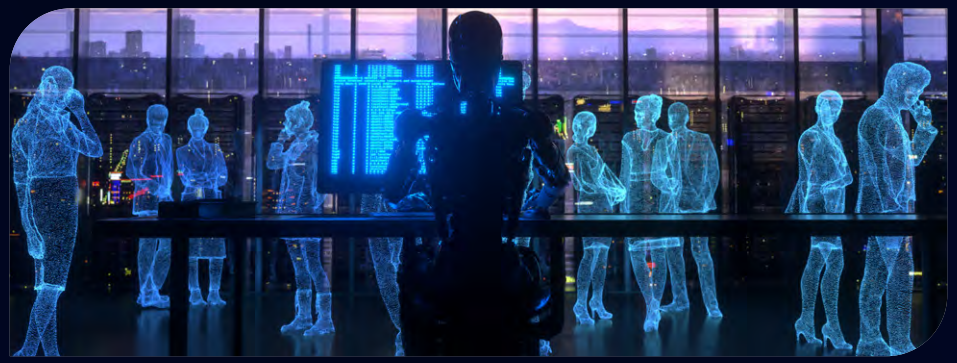
Every generation faces moments when technology reshapes the economy and redraws the boundaries of competitive advantage.

The nations that recognise those moments early and act with confidence pull ahead. Those that hesitate risk falling behind. At these moments, the challenge is not one of technology, it is of leadership.

Britain has never prospered by standing still. Throughout our history, the UK has succeeded most when we have embraced change, adapted quickly and translated innovation into commercial success. From the Industrial Revolution to the digital age, our competitive advantage has not come simply from inventing new technologies. It has come from our ability to apply those technologies to industry, scale them across the globe, and use them to create growth, opportunity and prosperity.

Today, artificial intelligence presents the UK with precisely such a moment.

AI is already beginning to reshape how businesses operate, how public services are delivered, and how economies compete. I see that in action as I travel across the UK every week, meeting businesses within the Chamber network. Much attention has understandably focused on the extraordinary pace of technological innovation, or on how to manage potential risks. However, the question that arguably matters most for the future of the United Kingdom is a different one: how can we effectively adopt AI across the whole of the economy and gain advantage from it? That question is the inspiration behind this book.



AI adoption is undoubtedly one of the defining economic opportunities of our time.

If we are serious about raising productivity, accelerating innovation, strengthening our international competitiveness, and improving living standards, then the UK must become far more ambitious about how AI is deployed.

The UK has much to be proud of. We are home to great talent through our world-class universities, pioneering researchers, globally respected technology companies, as well as some of the most innovative entrepreneurs in the world. We have and will continue to play a leading role in advancing AI itself.

But we also need to be realistic. The UK is unlikely to be able to compete head-to-head with the US or China in frontier AI infrastructure and foundation model scale. Fortunately, the UK does not need to build the largest AI model to succeed. AI invention is only part of the picture. The countries that gain the greatest economic advantage over the coming decade will not necessarily be those that create the most breakthroughs. They will be those that embed those breakthroughs most successfully throughout their economy.

In short, the prize for Britain is not only to be an important centre of AI innovation, but it is to become an AI adoption superpower.

Across the British Chambers of Commerce network, we speak to businesses of every size, in every sector and in every region of the UK. We see enormous enthusiasm for AI and a genuine appetite to embrace its possibilities. But we also see the gap that remains between enthusiasm and implementation. For a SME with limited support bridging that gap can be hard, I totally understand that.

Many organisations are still experimenting rather than transforming. Pilots remain pilots. Promising tools are being tested on an ad hoc basis by specific individuals rather than embedded across entire organisations. Leaders understand that change is coming, but often feel uncertain about where to start, how fast to move, or how to bring their workforce with them. As a result, too much of AI's potential currently remains unrealised. Closing that gap will require action from both business and government.

Businesses must be prepared to think bigger.

AI should not be viewed simply as a tool for improving efficiency around the edges of an organisation. Its transformative potential lies in reimagining products, services, processes and operating models.

Government, meanwhile, has a vital role to play in creating the conditions for success.

Skills, infrastructure, investment, regulation, access to technology and practical support for adoption all matter. If Britain is to accelerate AI deployment at scale, public policy must match the ambition of the opportunity before us.

The British Chambers of Commerce has long believed that economic growth depends upon partnership between business and government. That principle is as true for AI adoption as it has been for previous economic transitions before.

Our mission is to make the UK the best place in the world to start, grow and invest in a business. That mission has become inseparable from the opportunity of AI adoption.

A UK that successfully adopts AI will be more productive, more innovative and more competitive. It will be better able to attract investment, support entrepreneurship, and create high-quality jobs. Businesses will be able to serve customers more effectively, develop new products more rapidly and compete more successfully in global markets.

A UK that fails to adopt AI, with pace and depth, will find itself falling behind. That is why this agenda matters so profoundly.

We believe that the British Chambers of Commerce also has a responsibility to help lead the way. For more than 200 years, Chambers of Commerce have helped businesses navigate periods of economic and technological change. We have provided the practical support, trusted networks and business voice needed to turn opportunity into growth. In moments of transition, business leaders need institutions that can connect ideas with implementation, national policy with local delivery, and aspiration with action.



This publication is part of that broader effort.

The essays that follow bring together a remarkable range of contributors from business, technology, politics, academia and the wider business community. They do not all agree on every issue; that was not the objective. But together they represent a rich collection of ideas, experiences and perspectives on how we can scale AI adoption more effectively, more comprehensively, and with greatest speed, and how we can use that adoption to strengthen Britain's position in an increasingly competitive global economy.

The BCC's hope is that this book encourages a more ambitious national conversation about AI adoption. One that moves beyond surface level experimentation and basic training, and focuses instead on implementation, leadership and impact. One that recognises both the opportunities and responsibilities that accompany transformative technologies. And one that is grounded in the practical realities facing businesses every day and can overcome challenges along the way. In short, we need leadership from businesses, from organisations like the BCC, and from government.

We have the entrepreneurial spirit. We have businesses ready to embrace the opportunities ahead. But leadership is not automatic. It must be earned through ambition, investment, collaboration and execution.

The question before us is therefore not whether AI will change our economy. It is already doing this. The question is whether the UK will lead that change or simply react to it.

We must give businesses the confidence and capability to adopt AI, and we must work together to create the conditions that can drive innovation and productivity, leading to growth and shared national prosperity.

If government and business rise to that challenge together, there is no reason why the United Kingdom cannot become the world's leading AI adoption economy.

That is a prize worth striving for. And we must seize the moment before us.



Rt Hon Kanishka Narayan MP
Minister for Artificial Intelligence

Foreword

Artificial intelligence represents an extraordinary opportunity for British business. Our success in capitalising on this opportunity will depend upon businesses across the country being able to put that technology to work - in raising productivity and driving growth.

Small and medium-sized businesses are pivotal to that mission. They are the backbone of our economy and the heart of communities across the country. And if we want AI to improve the lives of British workers across the country, we need those businesses to share fully in its benefits.

Achieving this will require a more deliberative approach to supporting adoption. There is an important role for national government, but ultimately adoption happens in workplaces and across local economies. Businesses and people trust those whom they work with every day: their peers, local business networks, advisers and Chambers of Commerce. This report makes a compelling case for harnessing those relationships and bringing practical support closer to businesses.

We also know that to harness the benefits of this technology, we must bring people with us. I want AI to augment the skill, creativity and judgement of British workers, giving people greater agency over their working lives rather than diminishing it. That means giving employers and employees the confidence to use AI effectively, and involving workers in shaping how it changes their organisations. This approach is central to the government's work on adoption.

Through AI Skills Boost, government and industry are helping workers develop practical AI skills. Our AI Champions have developed sector-led plans to break down barriers to adoption. More than £200 million has been ringfenced to support worker-centred AI adoption. And we are taking forward the recommendations of the SME Digital Adoption Taskforce.

But there is more to do. This report offers a practical contribution to the next stage of that mission. It provides welcome thought leadership for government, businesses and local partners to move from experimentation to the deep adoption needed to deliver measurable improvements in productivity and growth. This report's emphasis on trusted local support, practical use cases and workforce capability is vital to ensuring businesses and workers across the country share in the benefits of AI.

I thank the British Chambers of Commerce for this important contribution.

Its network gives it a unique reach into businesses and communities across Britain, and an important role in ensuring that the AI revolution is something SMEs help to shape, not something that simply happens to them.



Why Britain must act



01

AI's darkest hour



Alexandra Mousavizadeh
Chief Executive
Evident

Alexandra Mousavizadeh is a Danish economist and CEO of the intelligence platform Evident which specialises in benchmarking and tracking AI adoption across financial services. She started her career at Moody's, then Morgan Stanley, later became CEO of ARC Ratings and then Director of the Prosperity Index.

Alexandra argues that Britain does not need to win the global race to build AI models; it can win by becoming the best place to put AI to work. That means clearer regulation, stronger talent pipelines and government policy shaped by firms already using AI on the front line.

Britain is losing the battle for the most transformative technology of our time. It's time to fight on the beaches.

America's tech giants are pushing the frontier of AI models. China is cornering the market for open-source alternatives. Canada and France have national AI champions being hardwired into their respective economies.

The prize is adoption, not invention

But Britain does not need to build the world's most advanced model to reap the biggest prize from AI.

With blood, toil, tears and sweat, the UK could become the best place in the world for businesses to put AI to work – and win a seat at the table of this new class of AI superpowers.



The foundations are already in place. It hosts four of the world's top ten universities. Its people are highly educated and digitally savvy. Its cultural 'soft power', through the history of institutions such as the BBC, gives the UK genuine global reach. There's a reason that the UK hosts such strength and expertise in professional and financial services, and why the UK boasts a legacy of applied innovation, from DeepMind to the internet and beyond.

But current setup of policy isn't setting Britain up for success. Homegrown AI talent flees to the U.S. to court deep-pocketed investors and avoid the punishingly complex set of policies that makes opening a business – let alone making it a unicorn – so hard. We see it in banking, long a bellwether for this country's economy. The UK is home to some of the most recognisable logos in the global financial system. Only one – HSBC – ranks among the top ten for AI adoption.

The adoption of this tech is already separating this industry into winners and losers. North American banks have led the field since we began benchmarking in 2023, today holding seven of the top ten places worldwide, and with an average Index score more than 20% ahead of the rest. The North American firms are leading the adoption race, driven by their proximity to the world's deepest AI talent and technology ecosystems.



The concentration is striking: roughly half of the AI talent in the fifty largest banks tracked by Evident is concentrated in just ten firms. That concentration of talent correlates directly with outcomes — the top ten disclose twice as many public use cases as the rest of the field, and are 1.5 times more likely to publish return on investment estimates.

We also see a difference in the composition of leading banks' AI talent – the ratio of people embedding AI into live systems to those building it is roughly four times the sector average. This isn't an exercise in vanity, the top firms are starting to push this tech to the bottom line. Britain is missing out. The same, as we read every day, is happening across every other sector, too.

What government must do next

The lessons from financial services give some steer to the UK government on what must be done if the UK is to become an AI adoption superpower.

Most importantly, its priority must be to create an operating environment that gives firms the certainty and support to invest in AI implementation across their businesses.

Firstly, by reviewing the regulatory landscape to ensure it is fit for purpose. AI has the potential to disrupt all sectors, and if regulators take diverging approaches firms will face an increasingly complex series of rules. Consolidating or simplifying guidance across regulators will reduce the compliance costs facing firms, without reducing the guardrails around AI, and unlock new resource to focus on actual deployment work. Government and regulators can go even further, publishing clear examples of good and bad AI practice so that firms can more easily ensure they're meeting the standards expected.



Secondly, it is vital that regulations are 'future proofed' to keep pace with the changing AI landscape. Regulators rightly take time to pass rules to enable real engagement with firms, and this shouldn't change. But the results shouldn't be focused on a particular point in time model, given the pace at which new versions are released. Instead, regulations should be technology neutral and focused on the outcomes and risks of firms' actions, giving firms the certainty to embrace the opportunities offered by the latest AI models without waiting for a brand new rules to be agreed.

Alongside this we must also ensure that the UK is learning from international peers on how to nurture and promote a thriving AI adoption landscape.

We must remove obstacles to the best and brightest studying in the UK, and do more to facilitate people who want to set up and grow their businesses in the UK. And government must regularly review the total of the UK's AI regulation against our international peers, ensuring that we are not disadvantaging our firms on the global stage.

Ultimately there is no silver bullet to AI adoption, and no single policy that the UK can adopt to become an AI adoption superpower. But when the pace of change is unprecedented, the leadership that the UK government can provide is vital. The appointment of a Cabinet-level Minister for AI is a very positive signal of intent, and a priority must now be to ensure that businesses' voices are being heard by the government.

Superpower status won't be won by a single policy or Budget — it will be won by a government willing to keep learning from firms on the front line, and providing them with the certainty and clarity that they need to drive forward AI adoption across the UK.

02

Britain won't become an AI Superpower by cheering from the sidelines



Lord Ed Vaizey
Business Advisor, Broadcaster and Member of The House of Lords

Ed Vaizey is a member of the House of Lords, business adviser, broadcaster, speaker and event host. He was MP for Wantage between 2005 and 2019, and served as Culture and Digital Minister for 6 years.

Ed argues that Britain must back AI firms with real contracts, cheaper power, better visas, clearer copyright rules and patient capital. Otherwise, its best ideas and talent will scale elsewhere.

We have a lot to thank Demis Hassabis for. Although the founder of Deep Mind sold his company to Google, he insisted that its headquarters should remain in London. As a result, he ensured that Britain could continue to build an AI ecosystem of researchers and founders, capable of pursuing their own start ups in the AI space.

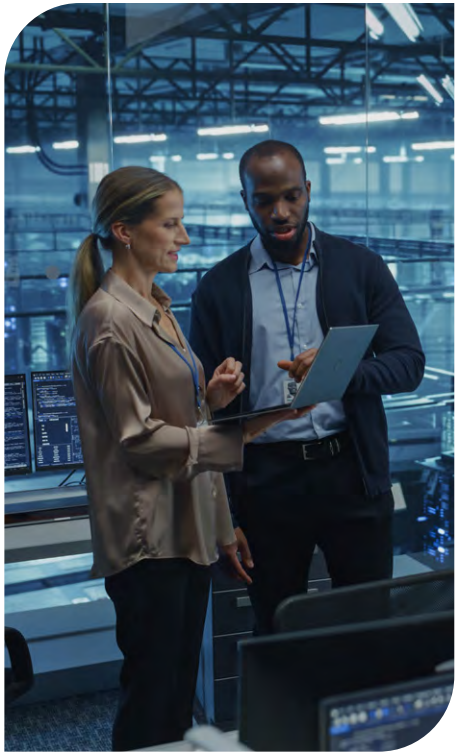
Back Britain's AI strengths

Britain will never be an AI superpower, just as she is not a military superpower.

But just as with defence, we are capable of building AI capabilities that means that we are a major player in the space, attracting investment and building companies of significance.

The trick, rather as with military procurement, is not to seek to build capabilities that ensure we cover all the vast ground of AI dominance. Britain can play a role by focusing on particular capabilities in the AI supply chain, such as building chips for niche applications, or moving forward in specific sectors such as healthcare, or indeed regulation. And again, as with the military (I realise I am stretching this comparison to breaking point), the government will play a key role.

This Labour government has extended Rishi Sunak's initial focus on AI. I confess I was a bit sceptical when Sunak decided to focus on safety in his first intervention. But he proved to be far-sighted, and the AI Security Institute is now one of the leading bodies surveying the significant risks that AI can bring. If anything the Labour government has been almost too enthusiastic. They have appointed an excellent Minister and put him in the Cabinet, but they have also created a plethora of bodies and a confusing and overlapping landscape of regulators, funders and research institutes. This needs to be simplified as a matter of urgency.



Turn support into serious contracts

Nevertheless, the Sovereign AI Fund is an important statement of intent. From a global perspective its £500 million is a drop in the ocean. But the combination of scale-up funding, and in kind support such as computing power, is better than nothing, and it's being led by smart, experienced people. But the fund must have follow up, and this means that government must start procuring from the companies it supports and funds. It cannot be emphasised how important this is. It gives start-ups with real potential a vital and reliable customer; it allows them to iterate their product in the market; and it is a powerful affirmation of capability to other customers.

The British government should set itself a target to award serious contracts – and serious here matters - not pilots, not another "innovation sandbox" – real money, to British AI companies.

Naturally you buy the best that is available, and sentimental protectionism is outdated. But a great deal of the best is, or could be, ours, and every large contract that goes to a home-grown chip designer or software firm does three things at one stroke. It adopts the new technology, it anchors the company (the capital, the personnel, the investment) in this country, and it tells global investors that Britain is both open for business and backs its own. The alternative is to restrict ourselves to summits and white papers, that is to virtue signalling.

Keep talent, power and capital here

To support this goal, one has to consider energy costs. Whether we like it or not, energy remains the binding constraint on the adoption and acceleration of new technology. For example, compute capacity in this country costs roughly four times what it does across the pond. That is why OpenAI paused its Stargate data centre here, and why serious money keeps eyeing the exit. You can have, like we do, some of the finest researchers in the world, but the mundane fact of the matter is that if the electricity to run their models is a luxury commodity, those models will be trained somewhere else. Another knock-on effect of that will be the brain drain that shall follow. Invest heavily in grid efficiency, then, unblock the connection queue, and if it comes to it, strike a proper deal with Norway for its cheap hydropower. Anything, in short, that gets the cost of power down.

We have a peculiar British genius for losing at the last hurdle. And this is not only restricted to World Cup finals. We produce and attract world-class AI researchers and then make it needlessly hard to keep them and to hire them. How to fix this? First, we need to reform non-compete clauses and notice periods, so a company can hire the person it wants next week, not next spring.

We have a real incubator economy, here and now, but we should not become incubators for others.

Second, create a genuinely fast-track visa for AI researchers from all over the world. We are in a worldwide competition for a few thousand extraordinary minds, and we should start behaving like it.

Settle copyright and unlock patient capital

There are two additional aspects to this strategy: copyright and capital. On copyright I do not accept the false choice that we either let AI hoover up creative work for nothing, or we miss out on the future – both halves are wrong. What AI companies actually want before they domicile here is, like any other business, certainty and stability of legal framework. If we finish building the licensing market – clear rules, transparency, a straightforward route to pay for it – we get two things at once; you give the technology companies the legal clarity that makes Britain a sensible home to them, and you let our creative industries sell their work. In other words, genuine growth from copyright.

Finally, capital. Another ambivalent trait of ours is our extraordinary ability at starting companies and then hopelessly failing to scale them. The missing ingredient? Patient money at the growth stage. Our pension funds are a significant lever which sit on vast reserves and put a pitiful share of them into British high-growth firms. Yet, a meaningful influx of pension capital would allow our best AI companies to raise their later funding here instead of decamping to the sunny plains of California. This is a properly patriotic act of our savings finally backing our success and our future.

I have been careful not to suggest we out-spend America or out-build China – that would be as pointless as it would be chimeric. Instead, if we start behaving like a customer, a host, and an investor in our future at the same time we shall become an AI superpower. In telegram form: Buy British where British is good. Price is our power to compete. Keep the people we train here. Settle copyright by building a market. And point our own capital at our own companies.

Do those five things and some of the best AI in the world will be built, sold and used right here. That, in the end, is what an adoption superpower actually looks like.

03

Bringing technology and good governance together



Dame Chi Onwurah MP

Member of Parliament
Newcastle upon Tyne Central and West

Chair
Science and Technology Committee

Dame Chi Onwurah is the Labour Member of Parliament for Newcastle upon Tyne Central and West. Before entering parliament in 2010 she worked as a chartered engineer, latterly as Head of Telecoms Technology at Ofcom.

Chi argues that AI offers major benefits but risks inequality and mistrust. She writes how strong assurance is vital for responsible adoption.

AI is in the news daily now. It is a huge opportunity to make the world work better and improve the quality of working life.

PwC UK have suggested that UK GDP could be up to 10.3% higher in 2030 as a result of AI. That's the equivalent of an additional £232bn, or extra spending power per household in the region of £1,800 to £2,300 across the country.

As Chair of the Science, Innovation and Technology Committee and a self-proclaimed tech evangelist, I am all for this. Instead of exhausted workers burdened with routine and repeated paperwork, many tasks can be completed with AI tools. In my own office, I can see huge potential for AI. For example I'm lucky enough to get a lot of invitations to events both in Newcastle and in Parliament, but that means a lot of work for my office to process invites, respond to them, and arrange for me to be there. That administrative work could be delegated to AI, freeing up my staff to serve my constituents more directly and creatively.

The AI-enabled world can give workers the time to focus on the real aspects of their job. A world where teachers spend more time in the classroom to inspire and teach our children.

A world where doctors and nurses spend less time staring at spreadsheets and more time delivering compassionate care to patients.

As part of my committee's inquiry, Digital Centre for government, we heard of Sutton Council's ability to use AI to detect when an elderly patient may have fallen over.

The government's AI Opportunities Action Plan aims not to frame AI as a mechanism to eliminate workers, but as a historic tool to elevate the value, productivity, and sheer satisfaction of human labour.

But before we can reach the sunlit uplands of an AI Adoption Superpower, there are three key issues to be addressed - the economics of AI, the skills necessary for AI and public trust in AI.

We must be honest about the economic challenge. There is a very real threat of a deeply polarised, two-tier labour market. Jobs directly involving AI activities are set to surge from around 158,000 to 3.9 million by 2035. But that growth is concentrated overwhelmingly in highly skilled, white-collar, professional and managerial roles.

At the same time, traditional administrative, secretarial and customer-facing roles are projected to experience a sharp decline. Without action, we could see an economic fracture open up: a highly paid elite who design, implement and own AI systems on one side, and a precarious layer of displaced workers on the other.

Small and medium-sized enterprises face a challenge too. While global multinational corporations have the deep pockets to invest heavily in AI infrastructure, our SMEs which make up 99.8% of the business community, risk being left behind - widening the productivity gap between the wealthiest firms and the rest of our economy.

Still only one in five of UK businesses have adopted it as part of the business practice, and our SMEs lag their G7 counterparts.

The government's Modern Industrial Strategy (MIS) sets out a 10-year vision to boost productivity, innovation and economic growth across key sectors. Central to this strategy is the adoption of AI, which the government rightly recognises as a transformative tool to enhance competitiveness, reduce costs and unlock new markets, particularly for SMEs.

To tackle the challenge of AI adoption among SME's the MIS establishes AI Growth Zones, including one in the North East, to foster regional innovation and high-skilled job creation.

Skills can be a critical barrier to adoption. So the MIS creates a National AI Skills Hub and partnerships with industry to train 7.5 million people in AI skills by 2030, alongside a Sovereign AI Unit to support access to data, compute resources and talent.

But AI economics and skills are closely connected to what I consider to be the biggest barrier to AI adoption: lack of public trust.

As AI disrupts existing work practices, we need to understand responsibility. If an AI system produces a misleading audit insight, who is accountable? If it exposes confidential client data, who carries the responsibility? If it discriminates in recruitment, credit, insurance, housing or access to public services, where does redress sit? And if a business adopts a system it does not understand, cannot explain and cannot monitor, is that innovation, or is it unmanaged risk?

These are not abstract questions.

We have already seen how lack of accountability can damage public trust.

The Horizon scandal and the exams algorithm showed how serious the consequences can be when automated or technical systems are trusted without sufficient transparency, accountability and challenge.

And given the lack of public understand and trust of AI means combined with its power, could lead to a whole new level of risk.

That is why I often say technology should serve people, not the other way around.

Too many of my constituents in Newcastle Central and West feel that new digital technologies are something done to them, rather than with them and for their benefit.

They feel that opaque products from powerful technology companies increasingly shape their lives, while they have little power to challenge them.

That feeling is corrosive to trust. And trust, once lost, is very hard to rebuild.

That is where AI assurance can make such a difference. AI assurance is how we move from assertion to evidence. The government's guidance describes AI assurance as a key part of wider AI governance and as a way of helping organisations operationalise responsible AI principles in practice.

It recognises that assurance can help consumers trust that products will work as intended, help industry invest confidently, and help regulators monitor compliance while enabling innovation.

We are seeing the rapid rise of "algorithmic management", where software, not human beings, is used to track keystrokes, allocate shifts, and even determine disciplinary action. When an algorithm becomes your boss, then who is accountable for your rights as a worker?

Then there is the deeply concerning issue of structural bias. AI models do not generate insights out of thin air. They are trained on historical human data. And historical data reflects historical societal prejudices.

The AI Opportunities Action Plan itself warns that without robust safety and data assurance, we are simply institutionalising discrimination. If you train a recruitment algorithm on historical corporate hiring data from a male-dominated field, the machine will automatically filter out female applicants. It doesn't eliminate bias; it automates it, disguising prejudice as mathematical certainty.

If you train a recruitment algorithm on historical corporate hiring data from a male-dominated field, the machine will automatically filter out female applicants. It doesn't eliminate bias; it automates it, disguising prejudice as mathematical certainty.

AI assurance can be the infrastructure of trust, it must be central to adoption, procurement, regulation and growth.

The future of AI is not a script written in Silicon Valley. It is in Parliament, in regulators, in professional bodies, in universities and in the businesses deciding how AI will be used.

AI assurance is not a niche technical issue. It is economic infrastructure. It is democratic infrastructure. It is trust infrastructure. And trust is the growth strategy.

As an engineer, I know the immense capability of what technology can build. As a parliamentarian, I know the immense power of what good governance can protect.

We need to bring those two worlds together to drive adoption and build AI systems that are not only innovative, but reliable and accountable.

04

Britain's AI opportunity begins with becoming better customers



Stephen Kelly
Chief Executive Officer
Cirata

Stephen Kelly is Chief Executive Officer of Cirata, a technology firm specialising in data integration and management solutions. He has a 35-year track record building technology companies, having served as the CEO of three major UK and US public tech firms.

Stephen argues that Britain has the AI talent, but not enough customers willing to back it. He says that unless government and big business buy British tech faster, the jobs, growth and investment will head to the US.

Over the past three decades, I've had the privilege of leading British technology businesses, working with hundreds of startups and scaleups, and helping organisations modernise through technology across both the public and private sectors.

From leading the UK's largest software companies, being the first Chief Operating Officer of the whole UK civil service, to working with today's generation of enterprise technology businesses, I've seen one pattern repeat itself time and again.

Britain builds the tech, but America buys it

Britain has no shortage of great technology companies but what we struggle with is creating enough great customers.

Artificial intelligence is the clearest example of this challenge. The UK has invested heavily in AI research, produced globally recognised AI companies and continues to attract exceptional entrepreneurial talent. Yet government research shows that only around 16% of UK businesses are currently using AI, despite those that have adopted it reporting significant improvements in productivity and operational efficiency.

The challenge is no longer inventing AI. It is deploying and adopting AI at scale.

This is of strategic importance to the UK as the UK lags in the G20 Productivity leagues and productivity is a key to unlocking growth.

Enterprise AI is unlike previous waves of technology. Organisations won't deploy AI into critical operations unless they trust the supplier, the data governance and the implementation. Procurement therefore becomes the gateway to AI adoption. If that gateway is too slow or too heavily weighted towards incumbents, innovation stalls before it ever reaches production.

Every month, I meet founders building world-class businesses in Britain. They have exceptional teams, innovative products and ambitions to compete globally. Yet ask them where they expect to win their biggest customers over the next five years and, often, the answer is the United States and overseas. Our own company, Cirata is a great example: 76% of Cirata's revenue came from financial services customers in the USA - the most competitive & demanding market in the world - and only 2% of sales from the UK. So, for a British company with onshore jobs in Belfast and the north of England, it makes no sense that UK potential customers have shunned UK tech companies.

That should concern us, not because Britain lacks innovation, but because too often we make it easier for our best companies to scale somewhere else. Talking to CEOs, the next step in their thinking will be where they move their HQs out of the UK, where they IPO their companies nearer their customers and where they create jobs and pay taxes - sadly, 'Going West' and heading to US.

For more than two decades, successive governments have worked to encourage entrepreneurship. The UK has built one of Europe's strongest startup ecosystems, attracted world-class talent and produced technology businesses that compete with the very best in the world. The UK has the third most valuable tech sector for investment after the USA and China. Yet there remains a missing piece in our growth strategy. We invest enormous effort helping companies start, but comparatively little help 'levelling the playing field' to them secure their first major domestic customers.

My old company, Sage, built its business by winning British customers before becoming a global software leader and created 2,000 high value jobs in the North-East around Newcastle. Kainos established its reputation through UK public sector transformation before exporting that expertise around the world. Early customers did more than generate revenue. They created credibility, reference accounts and the confidence to expand internationally.

Technology companies rarely relocate because they dislike Britain. They relocate because opportunity follows customers - Customers more than Capital. The United States is attractive not simply because of access to capital, but because innovative companies can often find larger, more willing enterprise customers much earlier in their growth journey. Customers generate revenue, growth and cash flows. Revenue growth attracts investment. Investment creates jobs. Before long, the centre of gravity shifts across the Atlantic.

If Britain wants to be a leader in the AI economy, we need to become a much easier country in which to do business.

Otherwise, the UK will become a technology colony where the growth and prosperity will be generated outside the UK, paid for by UK companies' money.

Procurement is holding back AI adoption

Ask almost any British technology founder about the hardest part of scaling their company and, after fundraising, the conversation quickly turns to scaling sales with the roadblock of procurement. Selling into UK government or one of Britain's largest organisations can mean months of procurement exercises, multiple security reviews, framework agreements, legal negotiations and repeated compliance assessments before a commercial discussion has properly begun. These processes exist for good reasons. Organisations must manage risk, protect sensitive information and spend public money responsibly. But together they create a system that often favours established non-UK incumbents with dedicated bid teams over younger companies bringing genuinely innovative solutions to market. Many US buyers foster innovation within their supply chain.

I've watched this play out throughout my career. British companies develop exceptional products, only to spend months navigating procurement processes at home while securing commercial traction overseas more quickly. That isn't because British organisations care less about innovation. It is because our systems have become better at managing risk than enabling progress and harnessing innovation. There is a way to do both.

This matters because procurement is far more than an administrative process. It is one of the most powerful industrial policy tools available to government and the UK economy. Every pound of public spending has the potential not only to deliver better public services, but also to strengthen Britain's innovation economy, create high-value jobs and accelerate AI adoption across the economy plus the multiplier effect of job creation, community cohesion, reduction of unemployment, increase in employment taxes - all goodness comes from growth.

We've shown before that procurement reform can work. In 2010, only 6% of central government procurement reached SMEs directly. By 2014-15, procurement with SMEs had reached 27.1% - going from £2bn in 2010 to £11bn in 2015 - through a combination of direct and indirect spend, exceeding the government's own ambition at the time. Procurement became a tool for creating markets rather than simply purchasing products. In parallel, innovation thrived.

That momentum has since slowed. Government has now committed to directing an additional £7.4 billion of procurement towards SMEs by 2028. That ambition is welcome, but the real measure of success should not simply be another spending target. It should be whether innovative British companies genuinely find it easier to win contracts and whether more organisations are adopting British AI solutions at scale. Sadly, the spend data shows the UK has gone backwards.

This is where the UK's AI strategy should evolve.

Customers matter more than capital

For years, the debate has focused on research funding, compute, regulation and skills. All are important. But none of them matter if innovative companies cannot sell their products and win customers.

Government should establish a fast-track procurement pathway for trusted AI suppliers that meet recognised security, governance and compliance standards.

Departments should be encouraged to move from pilot to production in weeks rather than months and years. UK government can sponsor and nurture UK Tech companies as a customer. Success should be measured not only by procurement efficiency, but by the productivity improvements AI delivers across public services. It is encouraging to hear the early 'calls to action' from the Prime Minister, Andy Burnham, but he needs to mobilise procurement policy and culture rapidly to turn the tide.

The same challenge exists beyond Whitehall. Britain's largest banks, insurers, retailers, manufacturers and healthcare providers also have a responsibility. Especially, the banks that were bailed out by the UK taxpayers in 2008 and the utilities companies supported by successive governments. Every major organisation talks about innovation, yet many still buy technology in ways that make it exceptionally difficult for British scaleups to compete. Procurement has become highly effective at managing risk, but often less effective at recognising opportunity and shutting the door on UK-led innovation.

Imagine if every FTSE 100 company made a conscious commitment to increase the proportion of technology it bought from British scaleups. Imagine if every public body measured itself not only on value for money, but also on how effectively it accelerated British innovation. More reference customers would attract more investment. More investment would create more high-value jobs. More successful technology companies would strengthen the economy and make Britain an even more attractive place to build global businesses. For UK companies, this would seem to be at the heart of a Board Director's responsibility akin to the ESG agenda.

Buy British innovation, or lose it overseas

History tells us that countries rarely become technology leaders simply because they invent great products. They succeed because they create environments where those products are adopted, commercialised and exported at scale.

Britain has no shortage of entrepreneurs, researchers or innovative ideas. What we need now is a stronger domestic market that rewards them.

This view is not partisan and calls for the balance to be reset with UK tech companies given a 'fair crack of the whip'.

Having spent much of my career working with ambitious British technology companies, I remain enormously optimistic. We have the talent, the research, universities, the language, the entrepreneurial ambition and the capability to become one of the world's leading AI economies. The next phase of the AI race will not be won by the countries that invent the most technology, it will be won by the countries that make it easiest to buy, deploy and scale it. It is a rallying cry for UK government and business leaders to buy from the best British technology companies.

The UK's next global AI champion is probably already being built. The question is whether its first major customer will be British. It is Britain's to win.



05

The state must get its own AI house in order



Ed de Minckwitz
Director of Public Policy
ServiceNow

Ed de Minchwitz is Director of Public Policy for ServiceNow UK, a technology firm helping businesses connect AI, data and workflows. Ed is also a senior policy fellow at Policy Exchange. He previously worked as Chief of Staff to the Deputy Prime Minister between 2022-2024.

Ed argues that Britain will only become an AI adoption superpower if government first transforms its own systems, using the state as a proving ground for AI rather than simply encouraging businesses to adopt it.

Many of the essays in this collection will, rightly, be about how British businesses can adopt AI faster and to greater effect. I want to make a different argument: that the most valuable thing government can do to accelerate business adoption is to get its own house in order.

If the state becomes, as the AI Opportunities Action Plan recommends, the best customer for AI, showcases what good looks like, and builds the digital foundations that private enterprise can plug into and build upon, it will do more for adoption across the economy than any grant or strategy.

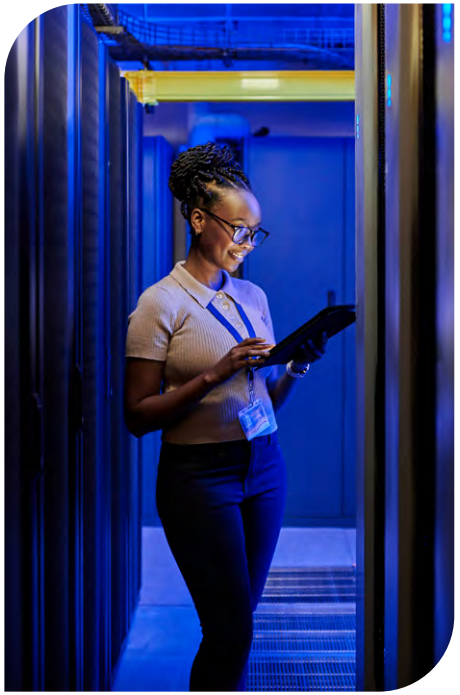
Define the destination first

But before any of us can say how the UK becomes an AI adoption superpower, we need a clear sense of what one might look like. It is only once we know the end state that we can establish what is needed to get us there.

At its simplest, an AI adoption superpower would safely diffuse AI through its economy fast enough to accelerate growth, while using it to transform public services and improve the citizen experience. Adoption would relieve humans from drudgery and elevate their role in the labour market to be more fulfilling. And, because AI innovation will not stand still, it would have an eye to the technology of tomorrow, to avoid creating further legacy lock-in.

The UK starts with many advantages, and some profound structural disadvantages.

We have an open economy, world-class research, innovative businesses, and a history of embracing new technology. But we also have vast amounts of legacy infrastructure, institutional risk aversion and growing public scepticism.



Pilots will not be enough

There are encouraging signs. Every company is asking how to leverage AI. We see it applied to casework automation, notetaking and summarisation. There is no end of public sector AI pilots, and thousands of British start-ups are building at the application layer. All useful, but almost all asking how to make existing processes better. Little wonder, then, that gains have so far been incremental, and on nothing like the scale promised.

For the step-change implied by self-defining as an AI superpower, we need to ask a different question: if we were building our businesses, our economy, and our public services today, knowing what AI can already do, and in anticipation of its future capabilities, what would we build?

We cannot hope that myriad incremental improvements to the status quo will eventually add up to profound transformation, any more than adding electricity to a candle creates a lightbulb. The AI adoption superpowers will be those countries with the vision, audacity, and bold leadership to reimagine systems with AI, define the end-state they want, and then work backwards to what they need to do today.

Make government the proving ground

The UK should start with the public sector – and here government has an advantage that business lacks, because it can simply choose to lead. The monoliths of our state – our health, welfare and education systems – have been a source of great pride, envied and copied the world over. But their sprawling bureaucracies, legacy technology and institutional resistance to change risk becoming millstones around the neck of an AI-native state.

Take welfare. Lose your job due to a medical condition today and in order to secure support you will have to navigate a byzantine architecture: tell numerous people the same story, fill in countless forms, wait weeks to discover what you're entitled to. AI bolted on to that system might provide a chatbot that explains your entitlements, or points you to the right person or form. Useful, but that is incremental improvement, not transformation.

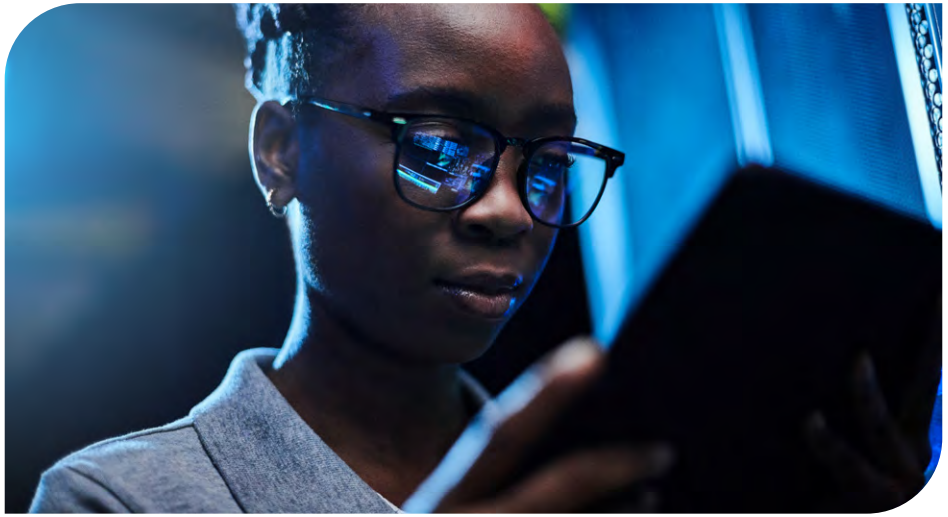
An AI-native welfare system would allow claimants to actually tell their story just once, and then use agentic technology to proactively provide them with the information and support – financial, medical or educational – they need to move forward. Unlike even the best human service, AI can be endlessly patient, continuously personalised, and relentless in its task. Done properly, it should also be more auditable and transparent than current processes,

letting citizens see exactly how their data has been used and shared, and how decisions have been reached. It could also make significant inroads into the fraud and error that currently cost taxpayers almost £10 billion a year.

Human caseworkers, meanwhile, would be freed from shuttling between systems and chasing paperwork to concentrate on the exceptions and matters of judgement that genuinely require a person.

Almost none of the difficult questions in that scenario are really about the AI model. They concern identity, data, permissions, interoperability, accountability, security and human authority. These are all challenges that can be solved today. But until we get these fundamentals right, transformative adoption – in state services, and in the wider economy that builds on them – will remain out of reach.

Today's fragmented government is inefficient partly because its systems cannot communicate. The transformative potential of an AI-native state depends on changing that: letting information move securely across institutional boundaries and, eventually, letting systems transact with one another. The step change comes when, with the appropriate authority, they can update a record, adjust an entitlement or schedule an appointment, while offering the citizen a single coherent interface to deal with the state rather than an endless series of disconnected encounters. That same connected foundation is what private business could build on and plug into – the platform, in effect, on which an AI economy runs.



Build trust into the foundations

Here, opportunity and risk become inseparable. Interoperability creates value, but increases the attack surface. Autonomous agents boost productivity, but raise the consequence of error, and personalisation, although improving services, also creates greater sensitivity. Connecting fragmented systems can transform the citizen experience, but it also means that a compromised system or identity can then reach much further.

Previous technologies let us build first and fix later. Indeed, much of today's legacy estate is the accumulated result. We should not assume AI will afford us the same luxury. Government's instinct with technological risk has been iterative: deploy, discover problems, patch, regulate, and (eventually) replace. AI compresses that cycle, and every deployment creates dependencies. Today's expedient workaround becomes tomorrow's legacy architecture. Once millions of citizens and thousands of public servants depend on interconnected AI systems, retrofitting identity, permissions, auditability and resilience will be vastly harder than designing them in now.

That is why cybersecurity and governance cannot be things we bolt onto AI adoption once the technology matures; they are part of the end-state. Nor do we have time to design bespoke approval processes for every use case.

Government should instead establish common guardrails, or governing tenets, from the outset. This points to two decisions it should make now.

The first is data. The UK should move to a presumption in favour of sharing data across the public sector, rather than the current default of siloes, following the direction taken by Singapore and, more boldly still, South Korea, whose digital infrastructure is built on precisely this premise. Without this, the ambitious devolution of powers, funding and decision-making planned by this government will lead to greater fragmentation of services, increasing the burden on citizens and business navigating the bureaucracy. Interoperability is the precondition for everything else: without it, the AI-native state is impossible, and the platform that business might build on never materialises.

The second is governance. Government should embed three tenets in a single piece of overarching regulation so that new systems are trusted and tested against them, rather than departments inventing their own approach and industry building against twenty incompatible architectures. First, identity and authority: we must reliably know who, person or machine, is acting, on whose behalf and what they are authorised to do. Second, security and interoperability: systems should exchange information through common standards without requiring indiscriminate centralisation; access should be controlled, activity logged, failures containable and components replaceable. And, third, accountability and trust: consequential actions should be auditable and contestable, with citizens able to understand how information about them is being used and humans remaining accountable for the systems they commission.

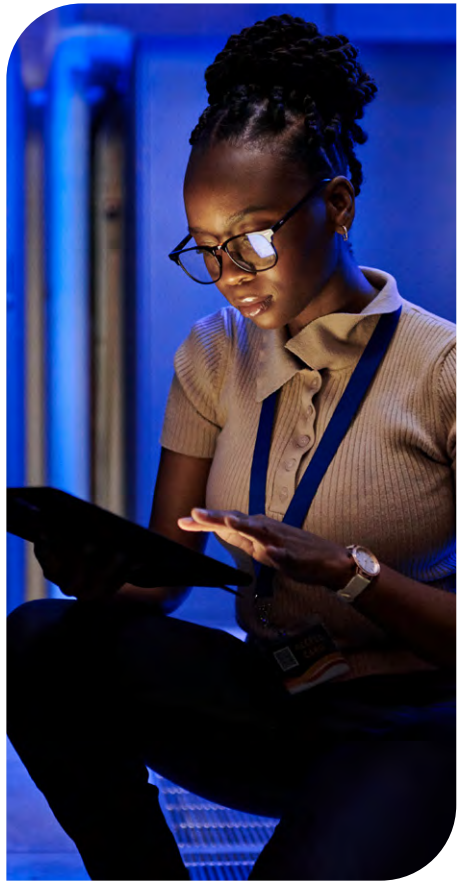
Mandating these tenets now, while the foundations of the AI state are still being poured, would accelerate innovation rather than slow it. It would also provide an answer to public scepticism: a system that can show a citizen how their data is being used and shared, and how decisions about them were reached, is far easier to trust than the black box it replaces.

We cannot know what AI will be capable of in 2030. Even if AI development were halted today, the diffusion of current capabilities would have profound consequences.

That is precisely why we should decide what we can about the end-state now. What kind of citizen experience do we want? What powers should machines wield, and what must remain human? What common infrastructure and protections need to be in place? And where must Britain retain sovereign control, while confidently embracing the very best international technology?

Uncertainty about the technology is not an excuse for uncertainty about the country we want to emerge from it. To quote the government's own AI Opportunities Action Plan, 'doing this well is not easy – it will require real leadership and radical change'.

To become an AI adoption superpower, the most important decisions we make about AI over the next few years will not be which models we build or buy. They will be the decisions about the foundations on which everything that follows will have to stand.



06

Closing the gap between inventing well and absorbing badly



Julian David
Chief Executive
TechUK

Julian David has been CEO of techUK, the leading technology trade association since 2013. He has has over thirty years of experience in the technology industry. Prior to joining techUK, Julian had a series of leadership roles at IBM including Vice President for Small and Medium Business and Public Sector.

Julian argues that the UK does not lack AI ambition, it lacks the practical machinery to help firms adopt it at scale. He explains how a coordinated national push, including, regional advice, sector use cases and trusted assurance can turn the tide.

AI could boost UK productivity by 1.5% a year, worth up to £47 billion annually. That prize is economy-wide, not one confined to specific sectors or to large London firms. Adopted responsibly, AI's benefits can reach firms and citizens in every postcode of the UK.

The challenge is that the UK lags on absorption and diffusion: how we adopt what already exists.

The Global Innovation Index ranks the UK 6th globally for innovation and 37th for knowledge absorption. That gap between inventing well and absorbing badly is the clearest international benchmark we have.

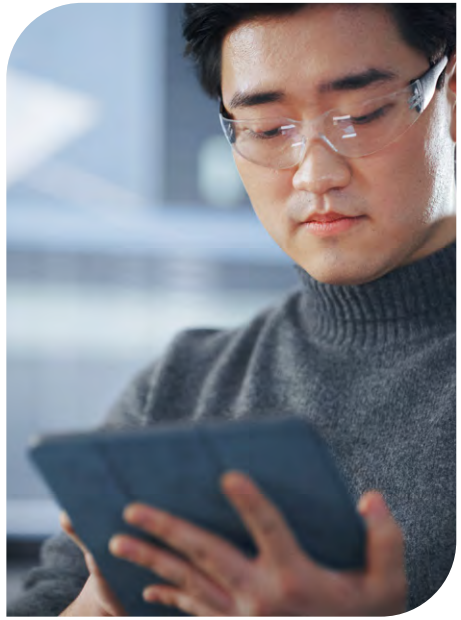
AI is following the same pattern. Just 16% of businesses currently use at least one AI technology, while 80% neither use it nor plan to, and 51% do not see AI as relevant. Preliminary research also shows a widening gap in embedding AI, with fewer than one in ten SME leaders reporting it embedded across several business functions.

Government has moved decisively, with multi-billion-pound investments and set-pieces like the AI Adoption Summit and the publication of the AI Champions' sector-focused delivery plans in June 2026. What is missing is coordination, execution, and implementation.

The barriers facing SMEs are well understood: cost, skills, leadership capability, assurance uncertainty, data access, and scaling from pilot to sustained deployment, as described in the AI Champions' plans.

Addressing them takes AI "out of the lab" and into the hands of businesses and workers. AI adoption that delivers economic growth needs to work at three levels at once – national policy that sets the incentive to adopt, regional delivery that builds the confidence to act, and sector-specific use cases that prove what works.

AI adoption in public services matters too and is a major focus of techUK's programmes. This essay focuses on the private sector, where the productivity prize is largest.





National incentives, regional delivery

62% of non-tech businesses report a growing interest in adopting new technology, according to techUK and Public First's 2026 polling. But when asked what holds them back, a third point to high upfront costs and a quarter to uncertain returns. Businesses will not spend on what they cannot price or predict.

An incentive helps, but only if it starts in the right place. Most support starts from the technology and asks a firm which tool to buy. It should start from the problem: where the firm is losing time, capacity or revenue, and what would have to change to fix it. Problem-first is the design principle, not a nicety. It is what separates a scheme that increases productivity from one that increases software sales.

techUK therefore believes that what is needed is a capped adoption voucher for SMEs, pairing financial support with impartial expert advice to support increased AI adoption.

Eligibility, the size of the cap and how the scheme is funded should be set nationally, so that whether a firm qualifies does not depend on where it happens to be based. Delivery, however, should be regional, through the local advisory networks we call for in our second recommendation below.

Singapore's Productivity Solutions Grant has supported more than 100,000 SME projects since 2017, with productivity increasing by around 4% where firms adopted sector-specific technologies. Germany's Digital Jetzt programme generated an estimated €1.2 billion in additional gross value added across more than 4,000 firms, from €163 million in public funding. Both pair money with advice.

The UK could look to introduce a simple two-stage scheme. First, an adviser helps firms identify the problem, the AI technology that answers it, and the likely cost. Second, government meets part of that cost, capped at £5,000 to start, across a broad range of AI-enabled and emerging technologies, rather than government predicting what firms should adopt.

The aim is to support investment that would not otherwise happen, not purchases firms were already planning. Where public funding covers a supplier's implementation costs, some could be linked to agreed adoption milestones. That keeps suppliers involved beyond the initial sale without making them responsible for everything affecting the firm's performance.

The scheme should be independently evaluated from the start, so government can tell whether it changes behaviour. It should begin with the Industrial Strategy sectors, but plan to expand beyond them.

The expertise that makes incentives work

Levers for AI adoption cannot just be national. They need to work for smaller firms across the country, not only the largest firms in existing clusters.

The Organisation for Economic Co-operation and Development (OECD) identifies peer learning as an important channel for AI adoption amongst UK SMEs. Firms build confidence by seeing what has and hasn't worked in their sector and region.

techUK therefore recommends that government funds regional expertise, embedded in local business networks, regional hubs and chambers of commerce.

What's needed is a vehicle to raise awareness of what adoption delivers, and direct firms to the networks they need. Private sector vehicles exist, but they should be paired with government's own incentives. People on the ground are best placed to build that awareness. The Digital Business Growth Service, launched by government alongside the 2025 Small Business Strategy to simplify how firms find and access support, should be the national front door, directing firms towards the sectoral and local support that fits them.

This should build on what each place already does, which ties

directly to the Industrial Strategy. Think advanced manufacturing in Manchester, biotechnology in Cambridge, the creative economy in Edinburgh. Regional support should be built around existing strengths rather than delivered uniformly.

This could extend to a CTO-as-a-service offer, as both the Tony Blair Institute and government's own SME Digital Adoption Taskforce have recommended. Both frame it as a national, online service, and government's 2026 update confirms it is being integrated into the Business Growth Service.

Our point here is about sequencing. An online platform works best when it is built on top of human adviser capacity. It is those advisers who bring together national, regional and local guidance, funding routes, training and vetted technologies, and who can check how responsibly adoption is working for firms.

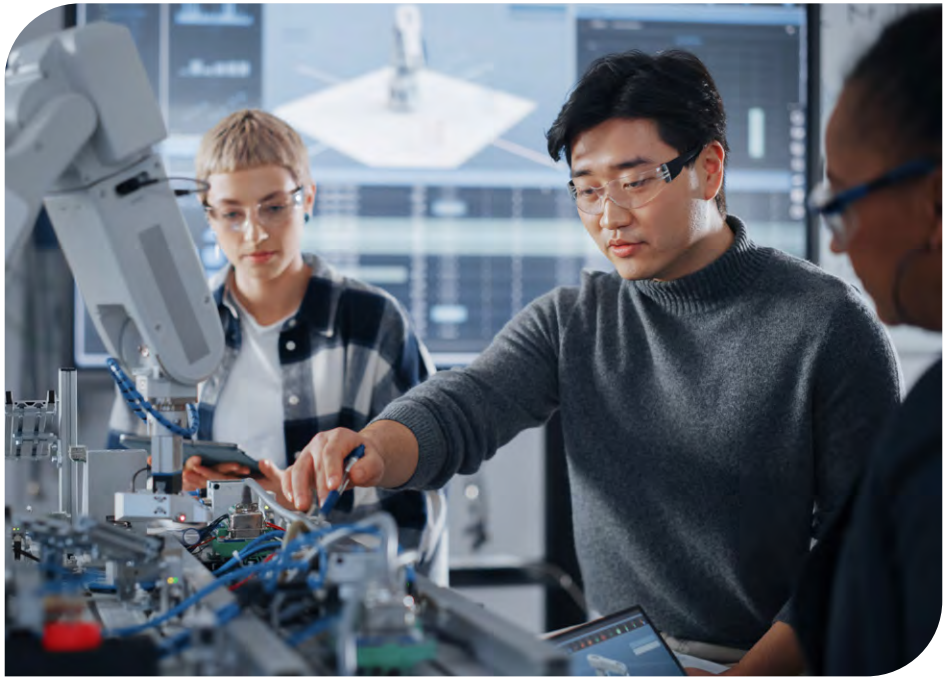
The building blocks exist but do not yet do the job together. For example, AI Growth Zones are weighted towards infrastructure supply, but could be given an explicit demand-side adoption remit to serve this purpose. There are examples of how this can work. Made Smarter, which focuses only on manufacturing, works precisely because it puts advisers alongside firms. techUK stands ready to give industry's backing, from the outset, to an initiative like this that extends the model to other sectors and businesses.

Sector-specific barriers, sector-specific answers

Sector by sector, AI adoption is patchy and uneven. Each of the June 2026 sectoral AI Adoption Plans used different statistics to measure adoption. Digital and Tech reported 39%, while the creative industries measured themselves at 51%. These figures are not comparable, which is itself part of the problem.

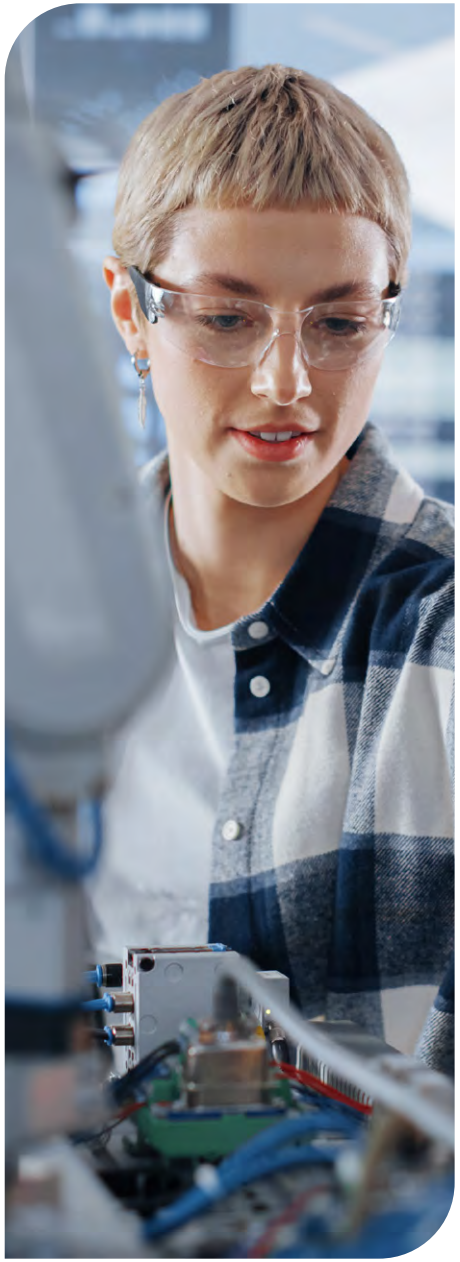
Across the Industrial Strategy eight, high adoption and unresolved barriers co-exist. The principal barriers are common to all eight – skills, confidence and trust, data, and infrastructure access – but sector-specific needs can diverge sharply. While a national offer can tackle the baseline, across-the-board issues, it will not address what actually drives adoption in our high-growth sectors.

techUK convened the AI Adoption Coalition to drive meaningful and responsible AI adoption across the UK economy. The Coalition brings together eight trade associations across the Industrial Strategy sectors, representing around seven million workers, as measured by government. techUK has already committed to helping government build the sector's AI Use Case Library, as announced in the Digital and Tech sector's AI Adoption Plan, and will use the AI Adoption Coalition to do so.



The model is replicable. techUK therefore recommends that government delivers a cross-economy AI use case library, building on the Digital and Tech sector pilot. A library funded and built through the same vehicles – the AI Champions programme and the AI Adoption Coalition – could be a co-delivered product that tackles the core issue: businesses need proven use cases, and confidence at sectoral and regional level. The AI Adoption Coalition convenes the sectors, knows the bottlenecks, and already surfaces common barriers.

There is, however, an accountability gap. Government has caveated the AI Champions' plans as independent and not formally part of government strategy. Its SME Digital Adoption Taskforce also recommended a minister accountable for SME digital and AI adoption, which government's 2026 update accepted. That post should chair a standing ministerial taskforce drawing in the sector departments, so one person is accountable and the departments that own each sector are in the room. A single minister without convening power will not move sector-specific adoption. Nor will separate remits in every department without a single owner.



Trust remains a precondition

Firms will only adopt what they can trust, and trust must be demonstrable rather than assumed. Similarly, workers' trust rises when they understand how AI is being embedded in their workflow and how it will improve their role rather than replace it. Skills, transparency, and confidence are interlinked and integral to successful, responsible AI adoption.

The UK has a clear national advantage here. Our AI assurance market was worth over £1 billion in 2024 and supported more than 12,000 people, outpacing the US, Germany and France relative to the size of our economy. Recognised accreditation and clear professional pathways would give firms in every sector and every region greater confidence to move. techUK leads the Industry Advisory Group for government's AI Assurance Stakeholder Consortium, which is taking that work forward.

A national way to drive this progress forward

Government has built a great deal since the Summit, and the Burnham government can take much of that forward, including the sectoral AI Adoption Plans, the Early Careers Jobs Alliance, AI Growth Labs, the AI Assurance Stakeholder Consortium, and the AI Economics Institute. What matters now is streamlining delivery so that these vehicles work together to help firms adopt.

techUK therefore recommends that government creates a single delivery map, connecting the UK's ambitions into one national, regional and sector-specific pipeline.

This map should tell a firm what national support exists, how local institutions are delivering it, and which use cases to rely on. This is an operating system for what exists, not a new institution or another portal. That would turn these three levels into a single flywheel, with each level reinforcing the others over time.

The UK already invents world-class AI. Becoming a responsible AI adoption superpower means getting just as good at absorbing it: putting proven, trusted tools into the hands of firms and workers in every sector and every region.

The recommendations set out here, from a capped adoption voucher paired with expert advice to funded regional expertise, a cross-economy use case library and a single delivery map connecting them, would move the UK from ambition to delivery. But they will only work if government and industry implement them together, in a coordinated and joined-up way.

Close the gap between inventing well and absorbing badly, and the prize is not just up to £47 billion a year, but growth felt in every postcode.

07

Turning AI promise into UK prosperity



Vinay Singhvi
Head of UK & Ireland
Tata Consultancy Services

Vinay Singhvi is the Head of UK & Ireland at Tata Consultancy Services (TCS), a global leader in IT services, consulting, and business solutions. Appointed to this role in September 2025, Vinay leads TCS's second-largest market worldwide, with responsibility for strategy, client partnerships, and operations across a diverse portfolio of industries.

Vinay argues that the UK can lead in AI by turning ambition into practical, trusted adoption. He writes that success depends on global partnerships, secure infrastructure and an AI-skilled workforce.

Here in the UK, we have a compelling opportunity to shape the next phase of the AI economy. The AI Action Plan has set an ambitious direction to accelerate AI infrastructure, skills, public service transformation, and business adoption.

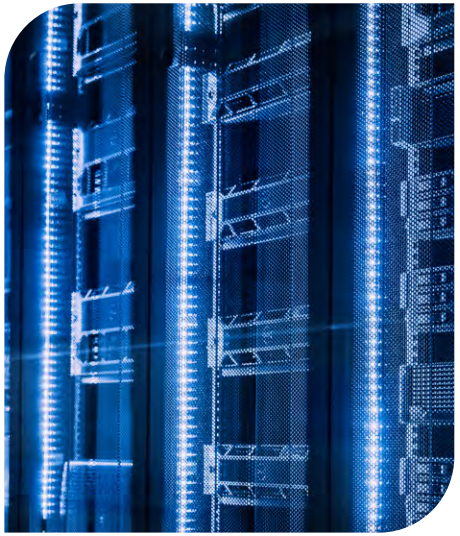
Combined with the launch of the Prime Minister's AI Taskforce, AI is now at the centre of economic and public service transformation.

However, to truly become an AI adoption superpower, we must turn policy momentum into practical implementation.

Businesses should be at the vanguard and access the right technologies, scale beyond pilots, embed trust, and equip people to work productively with AI. This means integrating AI into businesses and people's working lives in a way that is regulated and safe, while placing a strong focus on developing the skills needed to use it effectively and responsibly.

Drawing on our decades of experience in the UK and globally, we know the country already has strong foundations: credibility in responsible AI from hosting the first AI Safety Summit, and strong international partnerships. It also has the assets to turn that ambition into adoption; globally respected universities, a vibrant technology sector, and strengths across key industry sectors. For example, with our own National Health Service, the UK holds a unique opportunity to accelerate AI adoption across life sciences and healthcare, both in medicine advancement, treatment delivery, and patient experiences.

As a proud partner of the UK's tech ecosystem, TCS sees these strengths through our work with clients across sectors including financial services, retail, energy, telecommunications, transport, healthcare and public services. To become an AI leader, the UK must build on these foundations to combine global innovation with trusted local implementation — helping British businesses choose the right technologies, build the skills and foundations to use them well and safely, and turn ambition into measurable outcomes.



Build a globally connected AI economy

Choice and interoperability must sit at the centre of the UK's AI strategy. AI is not one technology, nor will it be shaped by one company or country. Models differ in capability, cost, transparency, and suitability for different sectors. While some tasks need frontier-model performance, others may benefit from smaller, domain-specific models that can be deployed efficiently and governed closely. Different organisations will need different models, controls, and deployment choices depending on their data, risk, and operational needs. The UK should therefore encourage organisations to avoid single-provider dependence and develop multi-model, multi-partner strategies.

TCS' approach is built around open, outcome-led choice, working with more than 3,600 partners across the technology ecosystem. Recognising that AI is only as powerful as the robust data foundations that feed it, our partnerships with OpenAI, Mistral, Anthropic, Databricks and Snowflake, alongside those with leading hyperscalers, give UK businesses access to different capabilities rather than a single prescribed route.

Our recently established Interactive Creative Engineering Studio in London helps bring this approach to life, enabling clients to draw on global AI-led engineering capabilities, while developing solutions grounded in local markets and customer needs. The UK Government's AI Opportunities Action Plan and expanded BridgeAI programme support the de-risking of adoption, by combining domestic public data assets with technology partnerships. A globally connected approach will help British enterprises deploy and commercialise AI, rather than build it in isolation.

Turn innovation into operational transformation

To scale AI, businesses need data readiness, cyber resilience, modern architecture, governance, and a willingness to redesign work.

AI should become an operating foundation—helping businesses make better decisions, improve services, identify risk earlier, and free people to focus on work requiring judgement, creativity, and empathy

That opportunity is particularly significant for small and medium-sized businesses, accounting for more than 99% of all UK businesses. AI could help boost productivity, innovation and competitiveness among UK SMEs, generating an estimated £78.1 billion in additional economic value over the next decade.

Organisations of all sizes must take a "giant leap" from digital enterprise to AI enterprise, reimagining workflows around outcomes rather than automating isolated tasks. As enterprise AI moves from assistants that answer questions to agentic systems that can plan and act across entire workflows, this redesign becomes more important, not less. Agentic AI acts as a decision coach, presenting options and creating feedback loops for continuous improvement, shifting the focus from mere automation to true autonomy

British industry offers one of the clearest proving grounds. AI-first manufacturing, where quality inspection, maintenance and production planning are continuously optimised by AI while people oversee the system, could help the UK reindustrialise on a more productive and resilient footing. TCS recently launched a lights-out factory lab to advance AI-first manufacturing, and as the Official AI partner to Jaguar TCS Racing, which competes in the all-electric ABB FIA Formula E World Championship, we are exploring how AI can unlock new efficiencies across racing operations and engineering.

Government initiatives such as AI Advisory Growth Labs and the Tech Town programme, can extend opportunity beyond the largest companies and established technology clusters. It will be important that such support reaches businesses across the UK, including those outside the largest cities and established technology centres.

Make trust, sovereignty and security engines of growth

The UK has already established an important foundation for responsible AI. Trust is built across the AI lifecycle, from data quality and provenance, security controls and model selection, to testing, human oversight and continuous monitoring. For regulated industries and public services, these are prerequisites for adoption.

Assurance and ethics are core to any responsible deployment approach. Governance, observability, and human oversight must be designed from the outset, with client data privacy protected through clear permissions and controls.

The AI Taskforce and the Government's £18 million Hardware Security R&D Programme can help make security integral to national AI capability, rather than an issue addressed after deployment.

Sovereignty is equally important, but it should not mean isolation. The UK's investment in Sovereign AI capacity building will allow businesses to not only depend upon models from other countries. Organisations also need confidence over where critical data resides and how essential services remain resilient amid cyber, commercial, or geopolitical disruption.

Make AI efficient, sustainable and affordable

AI adoption at national scale will be limited as much by power as by models. This is where the UK can set the pace by leading on efficient AI, matching every workload to the right compute. Much of everyday enterprise AI, from document processing to customer service, can run on smaller, domain specific models using energy efficient hardware, including air-cooled systems located closer to where data is generated.

Making energy per outcome a design measure alongside accuracy, cost and risk would lower costs for businesses of every size, ease pressure on the grid and keep AI growth in line with the UK's climate goals. AI Growth Zones, which aim to speed up grid connections and reduce energy costs, can anchor this approach.



Create opportunity in a Human + AI economy

The UK will not lead by seeking to outmatch every larger economy in every part of the AI stack. Our distinctive advantage can be to become one of the world's most AI literate nations, where people across businesses, public institutions, and communities have the confidence and capability to use AI productively.

The focus must be reskilling, not replacement.

The most effective Human + AI models will reduce cognitive load, helping people retrieve information, identify patterns and prepare recommendations, while retaining human judgement, context, empathy, and accountability where they matter most.

The biggest risk for all of us in the UK today is not AI, rather it is not preparing people for AI, with one of the most cited barriers for UK businesses being a lack of skills.

This talent equation must also be balanced with democratised access to AI infrastructure and compute power for academic institutions, entrepreneurs, and citizens in general, to ensure the country as a whole can benefit from the technology.



Government, employers, and technology partners must therefore do more to help people across every sector build the confidence and skills to use and access AI responsibly, while ensuring that workforce transformation creates opportunity rather than leaving people behind.

The government's AI Skills Boost programme has already recorded more than 1.7 million completed courses. Its investment in AI Growth Zones, scholarships, industry placements and an AI Economics Institute recognises that skills and evidence must accompany adoption.

People must have a voice in how AI changes their workplaces, and employers must make a credible commitment to learning and opportunity.

A globally connected route to leadership

The UK has the ingredients to become an AI adoption superpower: research excellence, globally significant industries, policy momentum, an international outlook and deep talent.

By pairing global partnerships with strong domestic foundations, responsible governance, and an AI-literate workforce, together we can help turn Britain from a source of AI promise into durable productivity, resilience, and opportunity. If the UK can make AI adoption safe, skills led and accessible to businesses of every size, it will be better placed to turn technological innovation into broad-based economic growth.

How adoption creates economic value



08

The middle way: How mid-sized businesses can lead the UK's path to AI Superpower status



Catherine Wilks
Partner and UK lead for Data, Analytics and AI
BDO

Catherine Wilks is a Partner at BDO, where she leads the firm's Data, Analytics and AI capability. She is responsible for driving commercial growth, developing strong client relationships, and building an engaged, high-performing team.

AI is spreading rapidly through British business – everywhere, it seems, except the bottom line. Catherine Wilks of BDO argues the answer is a middle way between innovation and control, and that mid-sized firms are best placed to find it.



Even for those of us who work with data and AI every day, the last few years have felt giddying at times.

With new models arriving almost weekly, the debate exploding into the mainstream and adoption cascading through the UK's businesses, it couldn't be clearer that the AI era is arriving at warp speed.

Except that it isn't paying off – at least, not yet.

Tooling turns out to be the easy part. In a recent report, the BCC found that over half the firms it surveyed were actively using AI. But working out where and how to capture the value is proving persistently difficult: in the government's own research, more than three-quarters of firms using AI have yet to see any change in revenue.

The cause isn't hard to diagnose, especially when it comes to the UK's mid-sized firms. Under pressure to keep up, many firms have tried to roll out AI like any other enterprise software, without recognising the deep organisational and cultural change needed for AI to make a dent in their bottom line.

Don't write off the middle

This oversight is understandable, particularly as mid-sized businesses are often too small for dedicated in-house AI teams and overlooked by the support which flows to start-ups.

However, we must not underestimate the role the mid-market can play in making Britain an AI superpower. These companies are agile enough to experiment, yet substantial enough to scale quickly when pilots turn into measurable progress. Their size makes cross-functional collaboration easier, too – a necessary condition for AI initiatives to flourish. And the will is clearly there: our recent Innovate to Grow report found that 56% of high-growth businesses are investing in AI, more than in any other technology.

To turn that investment into transformative value, however, leaders need to adopt the right mindset, strategy and governance – working together.

Firms are using a fraction of their AI brainpower

At the heart of the AI rollout lies a growing cognitive dissonance. For many of our clients, expectations of what AI should deliver are outpacing their organisational reality. They see their investment rewarded by modest efficiency gains as staff speed up routine tasks such as research, note-taking and email drafting, but those gains stay limited at the individual level. Few firms have gone deeper by embedding AI in the processes that actually run the business.

Investment has been lopsided, too, with firms pouring money into per-seat licences and token budgets while neglecting the leadership, skills and culture change that is critical to driving lasting results. In other words, too many firms have bought the tools, changed nothing else – and nothing's really changed.

When it comes to raw AI capabilities, however, the picture is reversed. Even if the big labs pulled the plug on frontier development tomorrow, today's models would still be capable of far more than most firms ask of them. GDPval, a benchmark of economically valuable tasks, has found that the leading models now match or beat industry experts around eight times in ten. And these are the least capable models we will ever use, improving significantly with each new release over weeks and months, not years.

Governance sets the ceiling

So why aren't businesses acting more boldly? The short answer is that the governance of AI systems is highly complex. Unlike the last digital revolution, today's AI agents cannot merely see but do things – taking actions across systems in a far less predictable way than their algorithmic ancestors, and constantly changing after go-live. While automation in itself is nothing new, current controls and assurance frameworks were simply not built for general purpose, multimodal systems that don't stand still long enough to be checked.

The risks are particularly acute in the mid-market, where the basics are often still being worked out – from data foundations to accountability and the unchecked spread of 'shadow AI' outside approved systems. In professional services, despite enthusiastically embracing AI, three-quarters of firms are not yet ready on data, orchestration and monitoring (according to the government's AI Adoption Plan). Data residency is also becoming a live issue in boardrooms: many firms are turning down US-based frontier tools because they can't guarantee their data will be processed in the UK.

Proportion beats paralysis

All this means that many mid-market leaders understandably find themselves caught between the drive to innovate and the need to control. Standing still is not an option while capabilities keep climbing. Rivals that rebuild their processes around AI may lower their costs and raise customers' expectations, leaving slower firms behind, with margins and market share that are harder to defend.

Yet the risks of moving too fast are just as real – and compounding.

Across many leadership teams we advise, the emerging tensions are plain to see: CEOs want growth at speed, CFOs want to know where to start and what it will cost, while COOs are left to make it all work on messy data and broken processes. This is a recipe for paralysis.



The choice, however, is a false one. If firms throw caution to the wind and lead with innovation alone, their data, controls and skills will struggle to keep pace, and risk will snowball. But if they focus too heavily on control, they equally risk being left behind by bolder competitors – not to mention driving their people to circumvent overly restrictive controls with shadow AI. Appropriately enough for the mid-market, the road to value runs between those two extremes. Set proportionate controls for each use case – light where AI drafts an email, heavy where an agent acts on a customer's behalf – and you can move at speed while staying protected.

When risk is well understood and managed, businesses have a licence to be bolder.

This is where the mid-market is right-sized to win in AI by design. Leaders who are closer to their people, processes and customers can act quickly to rewire their organisations as the technology evolves. Mid-sized firms may lack the scale and in-house expertise of larger rivals, but their relatively streamlined org charts make for faster decision-making. By contrast, large companies tend to be weighed down by legacy technology, unintegrated acquisitions and layers of complexity, making change slower and costlier.

While Whitehall waits, customers won't

Outside the office walls, the regulatory picture isn't helping either. Keen to spur innovation, the government has opted to 'wait and see', leaving businesses unsure whether to hold out for domestic rules or follow the EU's AI Act. But, in reality, the guardrails are arriving anyway, through customers, supply chains and international operations. It is the worst of both worlds: little clarity at home, and a patchwork of rules arriving from abroad.

Instead, policymakers need to adopt a middle way of their own: rules that are targeted and proportionate, giving businesses clarity without holding back lower-risk innovation. In our Innovate to Grow research, 21% of the businesses we surveyed said complying with regulation is a major driver of their innovation plans, with a slightly larger number agreeing that better regulation would help them innovate. Read together, those numbers point to an appetite for clearer, more enabling rules.

With the right external and internal conditions, established mid-sized businesses can turn AI capability into productivity and growth. Despite comprising fewer than 1% of UK businesses, mid-market businesses generate more than 40% of private-sector revenues. So a mid-market that learns to transform its operations with AI could become the proving ground the rest of the economy needs.

It's time to experiment – with discipline

A proving ground is a place for experiments, and almost every business is still experimenting – ourselves included. Nobody has the universal playbook yet, and the technology is moving too fast for anyone to write one. But the firms making progress are bold enough to experiment widely, while keeping a close eye on the constraints that matter. Our work with mid-sized businesses suggests that designing your approach around three dimensions should help set you on the right track.

1. **Mindset: ask the uncomfortable questions**
   Real value from AI means being prepared to rethink the whole organisation, starting with questions most leadership teams would rather avoid. If you were building your business from scratch today, what would it look like? How big should each team be, and how can you empower the people within them? And while you grapple with these high-level questions, be sure to keep hiring juniors. As agents take on routine work that humans oversee, your people need to keep developing the knowledge and experience that supervisory roles require.

2. **Strategy: think big, start small**
   Start with the biggest prize you can imagine, then work backwards. Transformative gains can come from rethinking how you serve customers or run a core process end to end, and mid-sized firms are close enough to their customers to spot the opportunities. Then start small, with a first step you can test and an outcome you can measure. Scale what works, stop what doesn't, and treat token costs like any other operating cost.

3. **Governance: make bold experiments safe**
   Experimenting widely only works if you know where the hard limits are. Aim for what MIT Sloan researchers call 'minimum viable governance': enough structure to understand and own the risk, without enterprise-scale bureaucracy. Too little and AI cannot scale; too much and the value never arrives.

Keep track of every AI use, shadow AI included. Name one person who owns AI risk – vital for mid-sized firms with no chief risk officer – and have a protocol for switching off any system. Be clear whether people are in the loop (approving before an agent acts) or on the loop (monitoring with the power to step in). Either way, humans stay accountable.

The middle sets the pace

For all the frothy headlines, the real measure of an AI superpower is the value its businesses generate – and the capability is already here.

What's missing is the organisational change to put it to work, and from our perspective that means taking the middle way. Firms need to be bold enough to experiment widely and disciplined enough to govern each use of AI by what it can do – a proportionate approach that policymakers would do well to mirror.

The mid-market is where that balance is easiest to strike.

The leaders who ask the uncomfortable questions, and then act on the answers, will be the first to see AI pay off – and with a fair wind, they will set the pace others can follow.

That's why Britain's AI future will be decided in the mid-market.

09

The constraint is moving



David Bharier
Deputy Director, Economics and Insight
The British Chambers of Commerce

David Bharier is Deputy Director at the British Chambers of Commerce where he leads the organisation's national research programme, including on AI and technology. He's a Policy Fellow at the University of Essex and sit on advisory boards across government, universities, and major external research programmes.

David argues that AI is changing the main business constraint from execution to imagination and organisational freedom. He shows how UK firms are adopting AI fast, but most are still using generic tools rather than redesigning how work is done.

Electricity was the paradigm-shifting technology of the early twentieth century. Yet it took roughly forty years for factories to reorganise their operations around it. During the steam age, one steam engine was the only source of power in the building, so its motion had to be carried to every machine by a shaft and belts. When electricity came, firms simply swapped the steam engine for one big electric motor and left the belts in place[1].

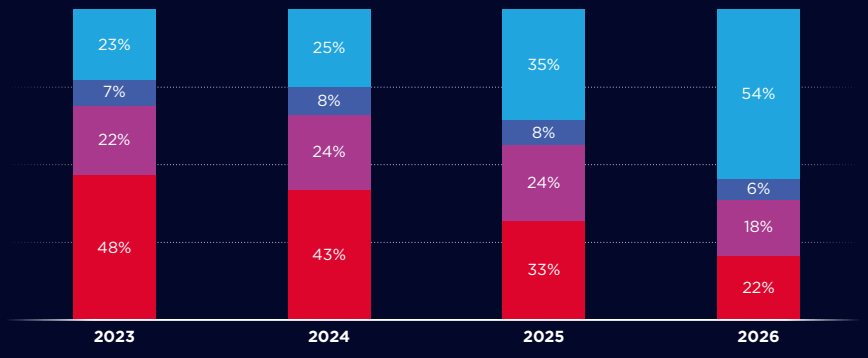
Electricity had dissolved the productivity-limiting constraint that tied a machine to its place in the line. But despite this, firms left the old layout untouched. Only decades later did they give each machine its own motor and rebuild the floor around the work itself.

The AI industrial revolution may be repeating the pattern. A firm that buys everyone a ChatGPT licence and changes nothing else is the factory with a new motor bolted to the old shaft.

The power source has changed; the layout has not.

AI adoption among UK firms is clearly accelerating. British Chambers of Commerce (BCC) surveys have shown usage rising from around 25% in 2024, to 35% in 2025, and 54% in 2026. But of the 668 firms examined in the 2026 BCC study with the University of Essex[2], 45% used generic tools alone, while only 10% said they had built something bespoke.

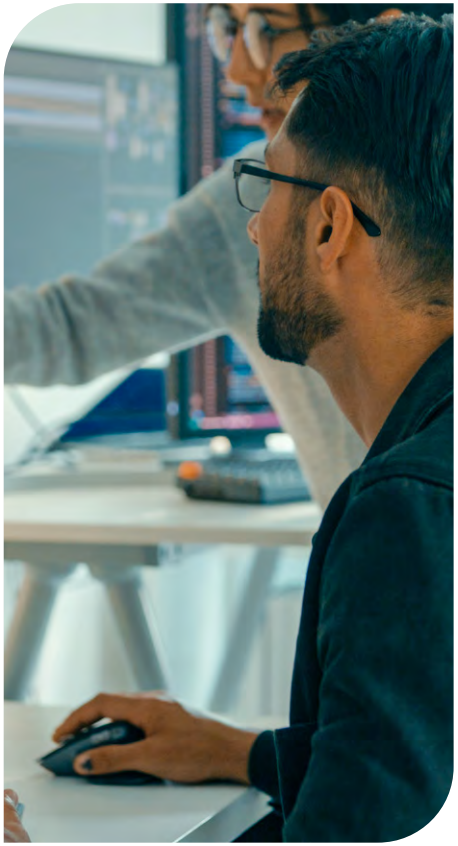
AI adoption rates are up compared to previous years



Q: Is your organisation currently using any specific AI technology? (n=668 for 2026).
Source: British Chambers of Commerce and University of Essex

The constraint that AI is removing is the cost and time of execution.

For many firms, two other constraints will now matter more: imagination – whether the person using AI can conceive of a new use for it, and freedom to act – whether the organisation lets them try it.



The perennial constraint: the cost of execution

To execute an idea, a firm traditionally must move through several costly stages – research, recruitment, procurement, production, approval. In 1937, Ronald Coase argued that firms exist because using the market is expensive: every deal means finding a supplier, negotiating and enforcing the contract[3]; for Jay Galbraith[4], writing in 1973, managerial hierarchy was the structure to process that information. AI challenges that model by performing parts of that chain directly, and at far lower cost. A worker can describe an output and receive analysis, copy, code, a prototype, or even the fully functional application within minutes.

The evidence that AI speeds up execution is growing. Developers given GitHub Copilot completed a task 55.8% faster than a control group with full internet access[5]. Consultants at Boston Consulting Group using GPT-4 completed 12.2% more tasks, 25.1% faster, within the technology's capability frontier[6].

But doing an existing task faster doesn't change the fundamental value of the completed task. What would raise the ceiling is the creation of new ideas: a firm doing something it could not or would not have previously attempted. An analyst who once produced a single static report can now build a suite of interactive simulators, test the scenarios facing the business, and

	Limited freedom to act	Freedom to act
Limited imagination	Shallow adoption: old tasks done faster	Experiments with no clear value, or no experiments at all
Strong imagination	Ideas that get trapped inside the hierarchy	Experiments that could lead to new systems, new products, or fundamental organisational change

see the results directly. The evidence that AI is producing more of this type of change is still thin.

The new competitive advantage: imagination and freedom to act

If AI continues to cut the cost of execution, execution alone becomes less of an advantage. What remains scarce is knowing what to build and being able to test it.

Imagination is the ability to conceive of a use for a technology that nobody has tried yet. Bresnahan and Greenstein called this work co-invention: the effort of adapting a general-purpose technology to one's own operations, which they found governed how quickly the gains from computing actually spread, more than the technology itself did[7]. Acemoglu and Restrepo[8] argue that automation removes workers from the tasks it takes over, leaving growth dependent on the invention of new tasks for people to do instead.

The BCC survey bears this out: bespoke users were around three times more likely than generic-only users to have restructured job roles. In open-text responses, generic users tend to name a tool or a familiar activity: ChatGPT, writing, research. Bespoke users are more likely to describe a workflow: quality control, maintenance, graphic design. So far, though, the evidence points more to displacement rather than new tasks: 21% of bespoke users reported an AI-related reduction in staffing, against 3% of generic-only users. The bespoke segment is small, but the direction is striking.



Freedom to act is the second constraint. Some firms may withhold authority to experiment with AI for good reason. A regulated business, or one handling sensitive data, has real obligations to control what goes into a system, and an unsupervised experiment could pose a genuine risk. Other firms take a limit-first approach without a specific restriction in mind: in BCC's 2024 AI Survey, 53% of firms said they were not confident they could apply AI technology, compared with just 28% for ordinary digital tools such as CRM or accountancy software.

Even where firms are open to experimentation, staff may not be aware of this freedom or have no time set aside to pursue an idea beyond their existing workload. Employees may leave ideas on the sidelines because they do not know whether the rules allow them to proceed or who would need to approve it. The owner may conclude that staff lack ideas, when the real problem is that nobody knows they are allowed to act. Aghion and Tirole called this the gap between formal and real authority: a firm can welcome initiative on paper without giving anybody the standing to start something.



Where policy can help

Governments cannot supply imagination or dictate how a company delegates authority. But policy can help widen the supply of examples firms can learn from and help organisations give staff a clear route to test an idea.

1. Spread practical examples

Firms are unlikely to attempt an application they have never seen. The gains from a general-purpose technology spread as awareness of use-cases spread, not at the speed the technology improves. The BCC's Powering Productivity[9] report recommends an AI Roadmap Toolkit, an AI Champion for SMEs and sector-specific adoption support. These should show businesses what comparable organisations have actually built, what work changed and whether the new approach lasted. The BCC AI Academy is already taking this approach, using role-specific exercises to help employees identify where AI could improve a real process, rather than simply teaching them how to use a chatbot.

Government already holds over 150 examples of its own AI uses in its Algorithmic Transparency Records repository, including tools used by HMRC to check tax returns and by the DVSA to target MOT inspections. A selection could be translated into short business case studies and distributed through Chambers of Commerce.

2. Give staff a clear route to test ideas

The Powering Productivity report argues that employees need the freedom and confidence to experiment, supported by leadership, HR and managers. government should work with businesses and Chambers to turn that principle into a simple template that firms can adopt: a small, reversible AI trial, limited to one process, requiring little to no budget, a named owner and a clear review point. Any publicly supported trial should then publish a short account of what was attempted, what happened, and whether it continued. That would give employees permission to try while creating further examples for other firms to learn from.

Whether the constraint has really moved remains an empirical question. The next wave of BCC research will begin testing it across thousands of firms, including whether employees below owner or director level can start a small AI trial without formal approval.

The UK's AI ambitions rest on whether a company of five or fifty people can redesign how its work is done and create new and valuable products. More than half of firms are now using AI, but only around one in ten have built something bespoke. Like the factories of the early twentieth century, the rest have a new power source but the same floor plan.

10

Depth over speed: A path to making Britain the first AI Adoption Superpower



Raoul Ruparel OBE
UK Chief Economist and Senior Director
Boston Consulting Group

UK Chief Economist and Senior Director at the management consulting firm Boston Consulting Group. Raoul's previous roles include Director of Trade and Investment at Deloitte and Special Adviser to Prime Minister Theresa May from 2016-19.

Raoul argues that Britain must pursue deep AI adoption through better policy, skills and business transformation. He explains how government and firms must act together.

Britain has a strong record of inventing new technologies but a much weaker one when it comes to diffusing them across the economy. If this pattern is repeated, Britain risks missing out on significant economic gains.

More than that, though, Britain could see its competitive position in key services eroded and inequality between different sectors and regions grow further.

So far, the evidence suggests there is a real risk that this will happen. AI adoption remains narrow and shallow. According to ONS surveys, around two-thirds of firms with ten or more employees aren't using AI at all, rising to more than seven in ten among the smallest firms. The BCC's own survey shows that while more than half of firms now use AI, only one in ten has gone beyond off-the-shelf tools to bespoke ones. All of this points to the fact that relying on natural diffusion of AI across the UK economy is unlikely to be sufficient. There are several market failures that mean this will require a concerted effort from both the public and private sectors.

Why Depth Matters More Than Speed

How, then, can we help ensure Britain becomes an AI adoption superpower? We should first be clear about what this requires. The previous Government set the ambition to have the "fastest AI adoption in the G7". While this is welcome, it is far from sufficient to make Britain an adoption superpower.

What really matters is depth of adoption.

This means moving beyond using off-the-shelf AI tools within existing business structures. It means integrating bespoke AI tools alongside the reshaping of workflows and investment in other forms of intangible capital, such as staff upskilling and retraining. The prize here is potentially significant. Our modelling suggests moving from broad use of off-the-shelf tools to deeper, workflow-integrated adoption could raise the productivity boost to GDP over the next decade from around £400 billion to around £1 trillion.

How government Can Unlock Deeper Adoption

With this in mind, there are four areas where government action can help meaningfully unlock deeper AI adoption.

First, there are some government dependencies, specifically regulation. As it stands, uncertainty over how existing regulations and guidance should be interpreted in an AI-driven world is having a chilling effect on adoption.

A clear set of guardrails, translating existing regulations and guidance into practical steps for firms using AI, can help reduce the current risk aversion.

It should also be backed by reassurance that firms staying broadly within these guardrails will not face overly harsh enforcement.

The Monetary Authority of Singapore has led the way here, with its set of principles for the financial sector that outline how to ensure AI usage is fair, ethical, accountable and transparent.

Second, there are areas where businesses want to act but are constrained by a lack of skills, capital or time. Here, the government can help to ease these resource bottlenecks. Skills is the most important of these. But it is also the one where the current approach is most likely to fall short.

While the previous Government took the positive step of creating the AI Skills Hub, the courses available so far look unlikely to achieve what is needed. The approach is both broad, aiming at 10 million workers trained by 2030, and shallow, with some courses taking under 20 minutes and focused on "the use of AI for tasks like drafting text, creating content and completing administrative tasks".

The aim should be less about the technicalities of using AI and more about ensuring workers understand how to work with and integrate AI. This means understanding how to design the right inputs, manage workflows and critically evaluate outputs.

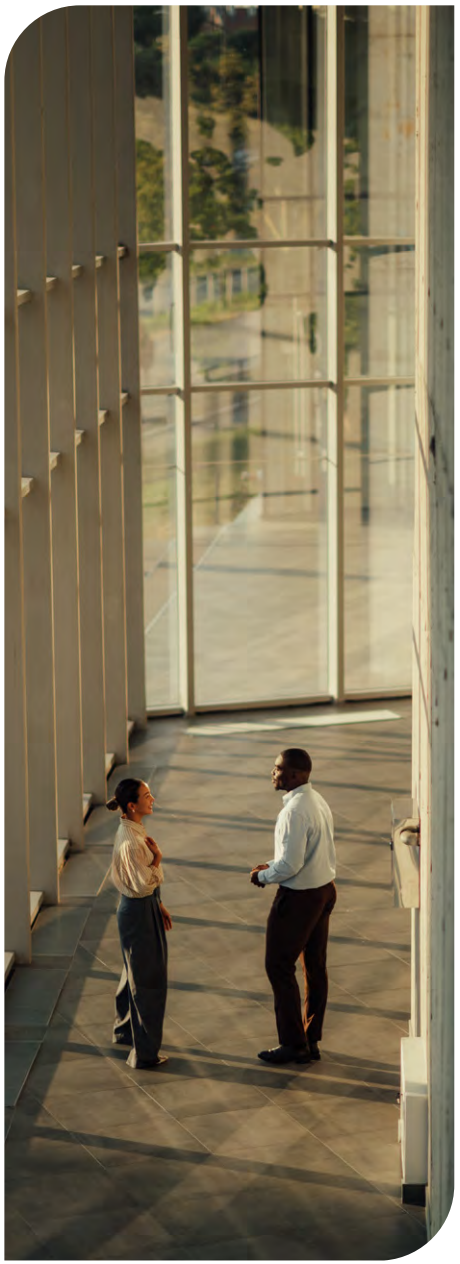
The new Government would do well to reassess the targeting of the AI Skills Hub and related spending. Ireland's Skillnet model is a useful example. Around 70 business networks, organised by sector and region, design training with providers and co-fund it with the state, so courses are grounded in the skills needed by businesses.

Third, many businesses simply don't know where to start. This is particularly true for small and medium-sized enterprises (SMEs), which do not have the time and resources to trial lots of different uses for AI to find ones that deliver lasting returns.

Furthermore, the speed of development in AI means many feel they can never keep up, so they shouldn't try. Here, the Government can play an important convening role. This means bringing together those who have technical AI expertise with those who understand day-to-day business operations. Together, they can identify and then share proven use cases that pay off.

"Show, don't tell" has long been the mantra for spreading new general-purpose technologies.

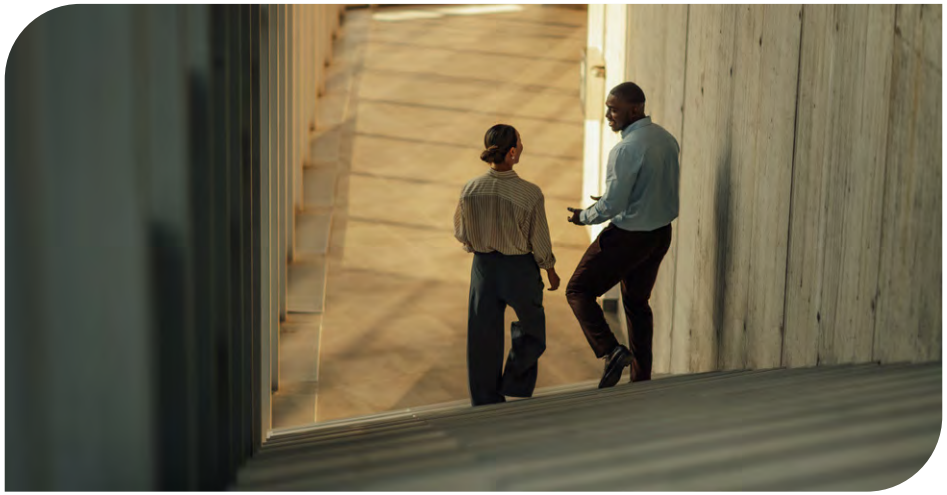
Indeed, the UK already has a model to build on here. The Made Smarter programme pairs SME manufacturers with technology advice, leadership training and match funding, and is now being rolled out across the UK. The Government should give it a more explicit AI focus and extend it to the services sectors with the most to gain.



Fourth, there are potentially significant national economic and social externalities that come with being an AI adoption superpower. These range from systemic security risks to labour displacement. All need to be considered, but I will focus here on the potential economic reliance and resilience angle.

As it stands, the UK owns little of the AI stack. When it comes to large-scale AI adoption, it is reliant on models, data centres, chips and energy from others. While significant adoption can bring economic gains, it can also sharply increase this reliance, to the point where it becomes a serious resilience and security risk. This means that any push on AI adoption should go hand in hand with developing the tools to manage this. The Government has the right organisation in the new AI Economics Institute (AIEI), but its remit should be widened to cover economic dependence and resilience. It must also have teeth.

The AIEI (along with the already world-leading AI Security Institute) should make clear recommendations for how to develop resilience across the AI stack, from diversifying the models' firms rely on to deciding which operations need to stay in the UK. They should also set out the infrastructure needed to make that possible. There have already been steps in this direction in one sector. In July, HM Treasury brought Amazon Web Services, Google Cloud, Microsoft and Oracle under direct Bank of England and FCA oversight as critical third parties to the financial system. Commissioning an in-depth benchmark of AI dependencies and resilience across the economy would be a show of intent.



What Businesses Must Do

Finally, it needs to be said plainly that this cannot and should not all fall to the Government. Ultimately, most adoption decisions will be made at firm level. At BCG, we see a few common elements that make AI adoption successful and lasting:

- Too often, firms treat AI as a tool to boost an individual's productivity, looking at how a single person can operate more efficiently or effectively. But the more valuable gains come at the collective level. Capturing that value means redesigning a few core processes end to end.
- This only works if people are at the heart of it. This means supporting people in upskilling, giving them the time, tools and direction to do so effectively. We consistently find that roughly 10% of the effort required to succeed relates to the algorithms themselves, 20% to technology and data infrastructure and 70% to people and processes.
- Leaders need to set a clear aim and drive it consistently. Employees with clear direction and limited tools outperform those with strong tool access and no direction.

Together, these highlight that AI adoption is not simply a case of integrating a new tech tool. Achieving the full effect often requires redesigning workflows and shifting the organisation's culture.

No country is yet an AI adoption superpower. The race is still open. The UK can win it by aiming for depth as well as speed. The prize goes beyond productivity gains and keeping pace with competitors. Deep adoption would give Britain's AI-first firms what many of them need: customers in their home market to help them scale.

A country that uses AI well creates the demand that lets its own AI businesses grow. That is how the UK becomes a superpower in both building and adopting AI.

11

Unlocking Britain's next era of productivity: Building a nation of AI trailblazers



David Black
Managing Director
Google UK

David Black has worked at Google since 2011, starting as director of online sales for UK & Ireland. Prior to Google, he spent more than eleven years in digital and sales roles at Trinity Mirror Group, now Reach.

David argues that UK AI use is widespread but shallow, limiting productivity gains. He explains how deeper adoption, training and supportive policies could unlock major economic growth.

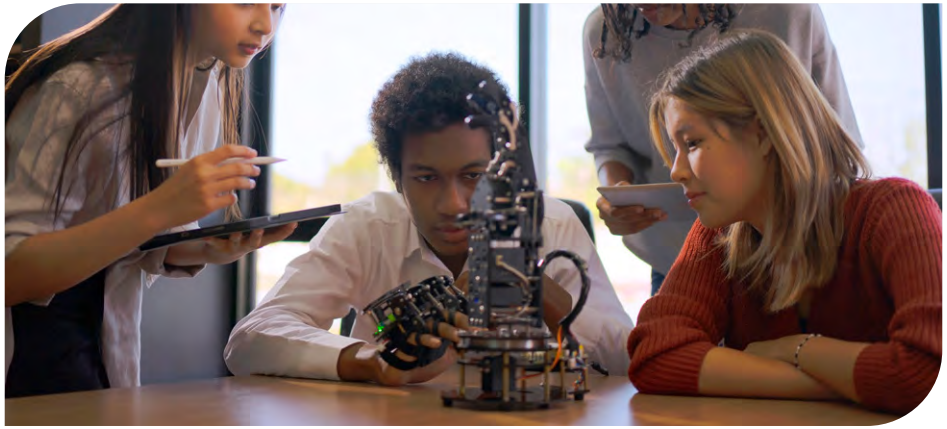
This is a pivotal moment for the UK. Long celebrated as a cradle of scientific discovery and industrial enterprise, the country now faces its next defining test: transitioning from basic and surface-level use of AI tools to the type of deep widespread adoption needed to unlock the country's next era of productivity.

The foundation is already laid, but our focus must now shift to ensuring that every business, across every sector and region, can harness the power of this transformative technology and drive meaningful economic growth.

The scale of the opportunity is extraordinary. In 2025 alone, Google products and services supported £140 billion in UK economic activity - a 19% increase from £118 billion in 2023, and an impact equivalent to the entire economy of Greater Manchester. Small and medium-sized enterprises (SMEs) accounted for £60 billion (43%) of this total, demonstrating how these products and features empower a host of growing businesses to grow and streamline their operations. Furthermore, Google tools supported over £23 billion in UK exports. From specialised manufacturers in the Midlands, boutique retailers in Wales, or creative agencies in Scotland, these tools provide a direct line to global consumers and effectively dissolve geographical barriers.



Yet, as substantial as that £140 billion baseline is, it represents merely the launchpad for what comes next. Research from Public First reveals that the widespread deployment of advanced digital tools and artificial intelligence could inject £400 billion into the UK economy by 2030, an annual productivity boost of 2.6% to the UK's GDP. With the right policies AI can serve as a catalyst that reinvigorates regional business corridors, elevates export competitiveness, and funds public services for years to come. Leaving AI adoption to chance risks stalling that growth as £200bn of these predicted gains depend on the workforce adopting AI and using it productively.



The Adoption Paradox: Wide, But Shallow

To capture the UK's full £400 billion growth opportunity, British business leaders and policymakers must take action to address our current landscape. Over the past year, British workers have embraced AI at a rate rarely matched in industrial history. Between 2025 and 2026, workplace AI adoption in the UK effectively doubled, climbing from 34% to 73%.

However, a closer look at the data reveals a fundamental structural challenge: British workplace AI adoption is currently wide, but shallow.

While nearly three-quarters of workers use AI, most are only scratching the surface.

Many workers still only use basic AI for quick copy edits, initial research, or scheduling assistance. Only 15% of the UK workforce have crossed the threshold into becoming truly advanced "AI Trailblazers." These power users don't use AI occasionally; they deeply integrate generative models and agentic workflows into core tasks and use AI tools for the 'heavy lifting', helping them to completely transform the way they work for the better.

This uneven adoption - with the bulk of the workforce still stuck in first gear, and only the top 15% unlocking significant productivity boost and time savings - lays bare the core factors holding the UK back from realising the full potential of AI-driven growth. If Britain is to realise its potential as a genuine AI superpower, our shared priority must be transforming casual users into highly productive advanced users.

The Trailblazer Dividend: Quantifying Real Productivity Gains

The performance gap between standard users and AI Trailblazers is vast, demonstrating that the real return on investment occurs when workers transition from basic capability to mastery. Our AI That Works for Britain report reveals the impact that power-user status has both on individual careers and the organisations that they work in. For example, AI Trailblazers report:

- **Significant Time Reclaimed**
  Saving an average of 7.7 hours a week.
- **Accelerated Career Progression**
  They are 84% more likely to be promoted than their non-power-user peers.
- **Superior Work Quality**
  This group is 88% more likely to receive top performance ratings from executive management.
- **Shared Prosperity**
  Advanced users are 53% more likely to secure pay rises, reinforcing the link between technical mastery and wage growth.

When multiplied across Britain's thirty-million-strong workforce, moving even a modest percentage of casual users into the Trailblazer category creates a compounding macroeconomic flywheel.

Paving the way for greater use

Unlocking this national advantage requires deliberate action from individuals, business leaders and policymakers. We can't afford to leave it to chance.

Advanced AI use doesn't require deep technical knowledge or coding expertise.

Practicing skills like writing clear prompts, iterating, and automating tasks are easy habits that can help people to extract more value from AI tools.

At Google, we recognise that upskilling requires concerted effort, and we are deeply committed to helping British businesses navigate this transition. Since 2015, the Google Digital Garage has trained over 1.2 million people across the UK in free digital and AI skills. In 2026, we accelerated this commitment with the launch of "AI Works for Britain," a multimillion-pound programme bringing practical AI training directly to high streets, universities, businesses and Jobcentres, alongside the new "AI Works for Businesses" training initiative in partnership with the British Chambers of Commerce and Enterprise Nation.

Creating the Conditions for Success: A Call to Action for Business Leaders

The workplace is the primary gateway for developing advanced AI skills. People are more confident and frequent users when their organisation provides clear guidance, permission, and support. I urge British business leaders to take the following practical steps:

1. **Approve and Encourage Safe Tools**
   Give your teams permission to innovate. Provide a set of trusted AI tools and clearly communicate that experimentation is encouraged.

2. **Simplify Policies**
   Most employees don't need a complex legal document. A one-page, plain-English policy outlining what is encouraged and what is restricted can significantly increase AI use.

3. **Invest in Practical Training**
   Offer hands-on training focused on real, role-specific tasks. Our experience shows a 10:1 return on investment for AI training.

4. **Foster a Culture of Experimentation**
   Encourage employees to share successful use cases and celebrate the wins.

5. **Develop a Deep Integration Plan**
   Work to embed AI deeply into your business operations, reviewing how data is stored, what can be automated, and how to maximise your return on investment.



Laying the Groundwork: Government Priorities for the AI Era

While businesses must drive adoption internally, the Government plays a crucial role in shaping the environment for success. To see the full benefits of AI, the UK needs a coordinated approach.

Drive Public Sector Adoption

The Government must lead from the front on AI adoption. Currently, less than half of public sector workers have received AI training. By demonstrating effective AI use in public service delivery, the Government can normalise safe, confident AI adoption across the wider economy.

Support AI infrastructure

Building on the UK AI Hardware and Action Plans, the UK must secure reliable, affordable energy to maintain competitive infrastructure investment and unlock high-quality public datasets for innovators.

Give confidence to businesses to seize AI's opportunities

A pro-innovation, sector-tailored regulatory approach will provide clear guardrails and practical guidance, instilling the confidence companies need to safely adopt and scale AI solutions.

Amplify best practice

The Government should amplify AI adoption guidance and best practice for businesses. Business.gov.uk is a great platform to signpost available AI adoption guidance and best practice guides from industry groups.

Empower local delivery

The Government should support devolved administrations, mayoral combined authorities, local councils, and trade unions to deliver tailored AI training locally to ensure the benefits of AI are seen across the UK economy.

Realising Our Ambition

The UK possesses all the necessary ingredients to be an AI leader - the talent, the research institutions, a vibrant startup ecosystem, and the ambition to be an AI leader. But unlocking the full economic potential of this technology requires a collective effort. By focusing on practical skills, creating supportive workplace cultures, and implementing forward-thinking government policies, we can ensure that AI truly works for Britain.

Let's move beyond basic adoption and build an economy where every business and individual is empowered to thrive in the AI era.

12

From AI Adoption to advantage: Turning Britain's head start into growth



Richard Breeden
Founder and CEO
aibl (AI For Business Leaders)

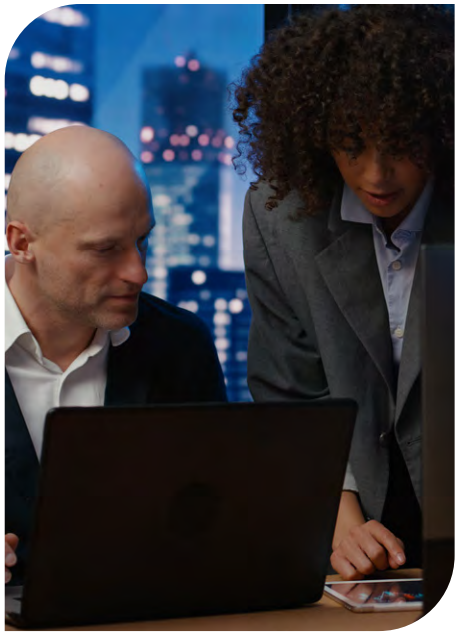
A veteran of value creation and digital transformation, Richard founded aibl in 2025. He has over 20 years of experience in B2B media and events.

Richard argues that Britain is ahead on AI adoption, but now it needs to make AI pay. He explains how the prize will go to the firms, and the country, that turn pilots into productivity, profit and growth.

Britain is ahead in the race to switch AI on. The race to make it pay is still open, and that is the one we should be focused on.

In aibl's survey of 755 UK mid-market business leaders this year, nearly 70% called themselves all-in on AI or building momentum. Across the wider economy the Office for National Statistics finds business use of AI up from about 12% of firms to 35% in two years. Even on that measure the UK leads its peers, against roughly 20% of EU enterprises (Eurostat, 2025) and one in ten US firms (US Census, 2026).

Fewer than half of those mid-market firms can yet point to a measurable return. That is true almost everywhere: a Massachusetts Institute of Technology (MIT) study last year found roughly 95% of enterprise generative-AI pilots produced no measurable return, and Gartner expects a third of AI projects to be dropped at proof of concept. No country has cracked this, which is what makes it an opportunity.



The prize does not go to the country that adopts AI fastest, but to the one that first turns adoption into measurable return.

That should decide what the government does next. The UK will not become an AI superpower by chasing the fastest adoption rate in the G7, but by helping its mid-market turn activity into returns it can show. What the state funds and rewards, the governance and workforce capability that make the difference, is what matters.

The target should be proof, not speed

"The fastest rate of AI adoption in the G7" was the inherited ambition, and it does not survive contact with the arithmetic. It carries no baseline, method or deadline, so it cannot be tracked. The underlying number is not even stable: the US Federal Reserve found UK business AI adoption reading anywhere between 9% and more than half. And it measures breadth, not depth or return. A better target is the share of firms that can show AI has paid.

Governance and capability are what make AI pay

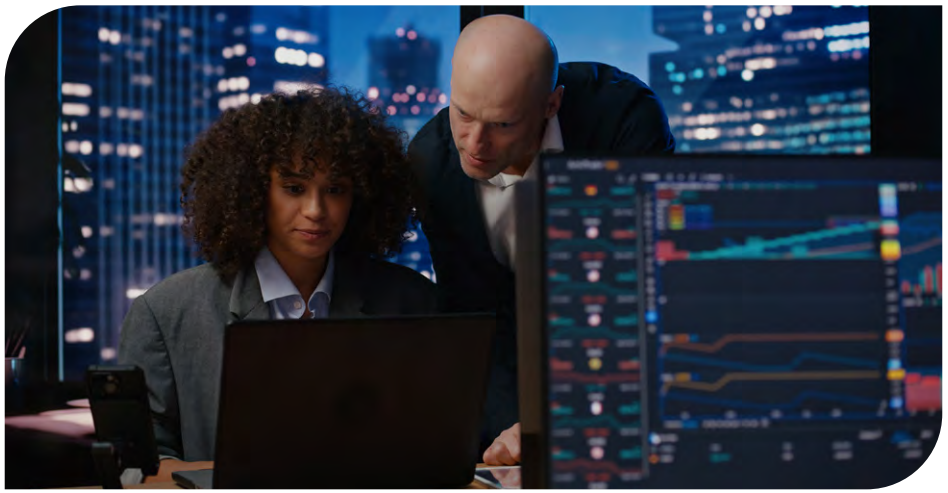
We looked for what separates the firms that can show a return from those that cannot. It was not better tools, cleaner data or talent, but two things.

The first is governance. Rank firms by how mature their AI governance is and the share that can show a measurable return climbs almost in a straight line, from 22% where there is none to 85% where it is mature and embedded, on much the same tools. It is an association, not proof of cause, but a strong one that holds across the sample. Governance here is not a compliance department. It is three plain habits: someone owns the outcome, the measure of success is written down before the money is spent, and the numbers hold up under Finance Director scrutiny. An independent Xero survey of 520 UK accountancy practices found the same: 87% of daily AI users keep documented, maintained processes, against 18% of non-adopters, and process, not software, separated the firms saving 7.1 hours a week from those saving nothing.

The second is capability. Where people are well equipped and the firm is well governed, about 81% can show a return; with one of the two, about 40%; with neither, 15%. It starts at the top: 71% of firms whose leaders use AI in their own day can show a return, against 31% where they barely touch it. A survey of 500 UK and Ireland mid-market firms found 91% rated their internal AI expertise highly, yet only one in ten had scaled every project beyond a pilot, with a shortage of that expertise the most common reason pilots stalled.

Enthusiasm alone is not the differentiator: only 39% of firms all-in on AI but with no governance can show a return, against 54% that are well governed but not especially keen. A policy nobody follows can be worse than none: in HR, governance defined but applied inconsistently returns just 16% of the time, against 33% for teams with no policy at all.

Yorkshire Building Society built the foundations first: a governed data platform and an internal AI academy that put 634 colleagues through more than 1,500 training modules, then narrow, well-scoped agents that save up to 26 minutes on a complaint and around 40% on internal risk testing, with people making the final call. The tools were on sale to every competitor; the return came from what was built around them. Most of the mid-market still sits on the lower half of that curve, which is exactly where the opportunity lies.



Britain's edge is adoption, not the model race

The Chancellor delivered his first Budget into weak growth, and the new AI minister has disowned "tech boosterism". Tight money and a preference for rigour over hype are the conditions now: back what can be shown to pay.

Britain will not win the race for general-purpose foundation models, which is decided by capital it cannot match, but it can compete on specialist, domain-specific models and on the application layer that puts AI to work in a real business, where sector expertise and proprietary data count for more than raw compute. The single ask of this essay is the demand side: helping scale-ups and mid-market firms build the governance and capability to turn AI into growth and profit, which is also the home market British builders need to sell into.

We would not be inventing the model. Germany funds regional, vendor-neutral AI centres and around a hundred "AI trainers" who work through local Chambers of Commerce. The United States has run a manufacturing extension service for decades, judged by the jobs and sales its firms report rather than the pilots they start: According to the National Institute of Standards and Technology, in 2024, its clients reported $15 billion in new or retained sales and 108,000 jobs. Both carry smaller firms over the gap and count an outcome, not a launch. Britain already has the pieces, in its skills funding and trusted local bodies; what they lack is a shared mandate and a metric, both within the government's gift.

One change the government can make

The ask is a single move with three parts, each building on something that already exists.

Change what the government counts. Retire the G7 adoption target and report, in the industrial strategy's existing annual update, the share of firms that can show a measurable return from AI.

Back the step that works. Tie the adoption support already running, BridgeAI, Made Smarter and June's package, to a light "AI in production" standard: a named owner, a written measure of success, an approved route quicker than the workaround. Reaching it should take a mid-sized firm a quarter, and meeting it is what unlocks the support. Public money now rewards firms for starting an AI project; this would reward them for finishing one that works.

Fund capability, and route it through what firms already use. Confirm the AI courses now fundable through the Growth and Skills Levy, and give firms one entry point to the skills money already devolved across twenty mayoral authorities, worked through local bodies like the Chambers rather than twenty separate rulebooks.



The prize is large, and the mid-market decides whether it lands

One estimate, modelled by Public First, puts the gain from AI at around £550 billion of UK output over the decade to 2035, close to a fifth of a single year in a roughly £3 trillion economy.

Growth is weak: the economy grew 0.4% in the second quarter of 2026, and the Office for Budget Responsibility expects a little over 1% for the year. Productivity, output per hour, grew by around 2% a year before the 2008 crash and has barely moved since, leaving the UK around 18% below the United States.

AI is one of the few levers that could lift that trend, but the modelling is clear the gain depends on diffusion, not invention.

It puts the cost of a five-year delay at about £150 billion, and £5 of value for every extra £1 firms put into tools and skills.

The prize only lands if AI reaches beyond the largest firms and the frontier. The mid-market is about a quarter of national output and eight million jobs, and it has done the hard part, which is starting. Move the average firm up the governance and capability curve and the national productivity number moves with it; leave it switched on but not paying, and the £550 billion stays a forecast.

An AI adoption superpower is not the country where the most businesses have opened ChatGPT. It is the country where the most businesses can show a board, a lender, an investor or a customer what AI has actually done for them.

A firm that has switched AI on is easy to see; one that has made it pay is not, so what the state does not measure it will not reward. Where the mid-market finishes now depends on what the government chooses to count and fund.

Please refer to the Reference section at the end for further information.

13

Britain has adopted AI. Now it needs to go deeper



Rob Burlison
Director of Global Corporate Affairs and Public Policy
Intuit

Rob Burlison joined Intuit as Director for Global Corporate Affairs and Public Policy, in 2018. He is responsible for Intuit's government and Regulatory strategy across all markets outside of the US.

Rob argues UK SMEs have embraced AI, but too many tools are disconnected to deliver real productivity gains. He says government should fix data-sharing and interoperability so existing technology can unlock growth across the country.

In the space of about three years, without a mandate, a subsidy scheme or much in the way of a national programme, British businesses have adopted a genuinely new technology at remarkable speed.

Research published this autumn by the Social Market Foundation (SMF) found that 79% of UK small and mid-sized businesses now use AI in some form.

Our own AI Impact Report, drawing on anonymised payment records from more than half a million UK businesses alongside survey data, tells the same story from a different direction: daily use among UK small businesses more than tripled in eighteen months, and the share saying AI has made them more productive went from under four in ten to more than three in four.

That has happened from the bottom up. A sole trader in Oldham discovered it could write her quotes. A plumber in Swansea found it could chase his invoices. Nobody asked permission. For a country that has spent years trying to figure out why its firms were technology-shy, this has been a quiet success.

So the question in front of the government isn't how to get British business interested in AI. That's done. The question is how to turn an adoption success into an economic one.

Turning AI adoption into economic success

Picture a manufacturing business in the North West with 120 staff. A hard-working finance director, a growing sales team, a decent year behind it. It also runs an invoicing system, a CRM, a payroll platform, a project tool, a stock system and a spreadsheet that quietly holds the whole thing together. Each bought to solve a real problem. Each one did.

Ask the finance director what her week actually contains and the answer is rarely finance. It's reconciliation. Lifting a number out of one system to put it into another. Checking whether two reports that should agree actually do.

She has tried AI, of course. She uses a chatbot most days - to draft a board summary, tidy a supplier email, make sense of a contract clause. It's useful. It saves her an hour here and there.

But it hasn't touched the thing that actually limits her company growth.

That's the national picture.

Nearly every AI-using small business in Britain is using a general-purpose chatbot. Only about half have moved to AI built into the software they genuinely run the business on.

84% already use accounting software; only one in five uses accounting software with AI inside it.

The SMF put a number on what's sitting in that gap - up to £47 billion in additional revenue if firms embedded AI properly across their workflows - finance, sales, marketing, HR. Not from technology that doesn't exist yet. From tools they could buy today.



Size matters

Part of what's holding that back is that we keep treating five million businesses as one population.

The smallest firms genuinely are under-digitised. More than a quarter still don't use basic digital tools at all, despite clear evidence that connected businesses are substantially more productive and more likely to grow. For them, the task is access, confidence and a reason to start.

The firms in the middle are the mirror image. They aren't under-digitised. They're over-digitised.

Our research has found that small and medium-sized businesses are running as many as 25 separate software tools simply to operate. That isn't transformation, it's overhead. And it's expensive. UK business leaders lose around 6.6 hours a week — close to a full working day — just moving between disconnected systems. 40% say they're still stuck in daily operations rather than leading. Only one in twenty feels confident delegating. Three in four have abandoned a growth idea in the past year, not for want of ambition but for want of capacity to act on it.

Add it up and more than half of this country's small and medium-sized business growth potential goes unrealised - a six-figure sum, for the average firm, left on the table every year. Their problem isn't conviction, it's integration.

Greater Manchester already proved the point with buses

For decades, a passenger in Greater Manchester faced a transport network that existed only in theory. The buses were there. The trams were there. But they ran as separate businesses with separate tickets, separate information and separate maps. Getting across the city meant knowing which operator ran which route, buying more than one fare, and holding the connections in your head. Every operator was behaving rationally. The experience was incoherent.

What the Bee Network did was not add vehicles. It connected the ones already in transit. One network across bus and tram, and real-time information in a single place. The physical assets barely changed. What changed was that a collection of disparate services finally behaved like a system - and once it did, people used more of it.

That is exactly the choice now facing British businesses. There is a country of companies with excellent individual tools, but with no means of connecting them.

The lesson from Manchester is that you don't fix a disconnected system by buying more buses.

Fix the plumbing

This mattered less when AI was a tool a person used as a prompt. You opened a tool, asked a question, did the work yourself. In that world fragmentation was an irritation, because you were the integration layer: you copied the answer out and pasted it somewhere else.

That era is ending. The AI arriving in businesses now behaves less like a tool and more like a team. AI comes in for parts of the job; matching a vendor name spelled three different ways across systems, categorising and reconciling transactions and flagging an anomaly before it becomes a problem, suggesting the optimal chart of account structure for consolidated reporting, drafting recurring entries, detecting anomalies, and turning it all into a forecast the business owner can act on.

Go back to our finance director in the North West with her many different business systems. The reason her AI use can't give her back her week isn't a shortage of capability. It's that the customer record sits somewhere the accounting system can't see, and the stock position sits somewhere neither can reach.

AI that isn't plugged into the data, can't see the customer record and can't chase the invoice. It can't read the pipeline and can't forecast the cash position.

AI cannot automate across a workflow that is fragmented.

Growth in every postcode

Here's why this is a national question and not a software one.

The SMF research shows that British SMEs could benefit from up to £47 billion in extra revenue if this plumbing was fixed. The most striking finding is how evenly this falls. The variation rate of SME growth sits within a whisker of 2% in every region of Great Britain. Expressed as a share of the whole regional economy, London comes last - while the East Midlands, the East of England and the West Midlands come top.

This is one of the few growth opportunities on the table that genuinely delivers in every postcode, and delivers hardest where small and mid-sized firms make up more of the local economy. It doesn't require a new industrial site or a decade of construction. It requires the businesses already trading in every part of the country to be able to get the most out of what they've already bought.

What government should do

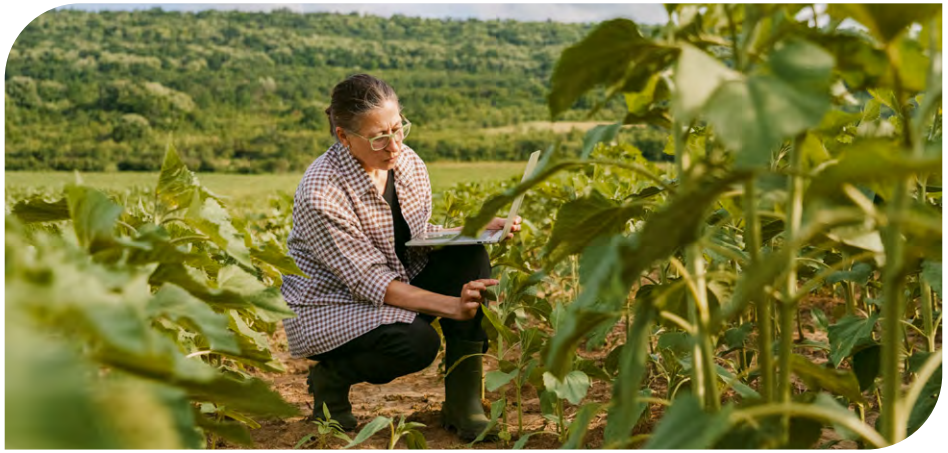
My proposal is narrow and unglamorous:

make interoperability a condition of the market, and give firms back control of their own data.

Set neutral open API and data-portability standards for business software, so a company's data moves between its accounting, payroll, CRM and banking systems without anyone retyping it. The market won't deliver this unaided — fragmentation is commercially comfortable for anyone benefiting from lock-in. A standard rewards the platforms that connect and penalises app sprawl.

Extend open finance to business data properly. The FCA published its open finance vision in April and named SME credit as a priority use case, with a framework to be developed with the Treasury by the end of 2027. That work should recognise accounting and business software platforms as data holders and recipients alongside banks, with better access to credit as an explicit, measured outcome. Financial data is the richest asset a scaling company holds, and right now it's trapped.

Build trust into the plumbing rather than bolting it on afterwards — a government-backed marketplace of assured AI tools with baseline requirements on security, accuracy and reliability, tied to the HMRC



recognised software lists. It must cover AI embedded in mainstream accounting, payroll and CRM software, because that's how firms actually adopt. A quarter came to AI because it arrived in something they already paid for. 8% came through a government programme.

And use the channel that already exists. A third of firms adopted AI because someone they trusted recommended it, and for a scaling business with no chief technology officer, that person is its accountant. Make AI implementation part of professional development with the accountancy bodies and you have a national delivery network overnight.

Finish the job

Britain has already done the hard, unpredictable part. Millions of businesses decided, on their own initiative, that this technology was worth their time.

What they're waiting on now is not encouragement. It's a system that lets the tools they've already bought talk to each other - so the hour saved here and there becomes a day, and the day becomes growth that shows up in the accounts.

Our finance director doesn't need another application. She needs the ones she has to work as one. Do that, and the £47 billion is reachable with technology that already exists.

An AI adoption superpower isn't the country where the most businesses have tried AI. We're virtually there. It's the country where their data can move freely.

14

From AI users to AI transformers



Alison Wright
Small & Medium Business Director
Microsoft UK & Ireland

Alison Wright is the Small and Medium Business (SMB) Director for Microsoft UK, where she leads business strategy, customer engagement, revenue outcomes, and digital transformation for the UK mid-market and small business segment

Alison argues in this chapter that the UK must help SMEs move from trying AI to using it to transform how they work. That means clearer use cases, local support and practical programmes that prove AI's value.

I spend a lot of time talking to businesses about AI, and one thing is increasingly clear:

Interest in AI is not our problem. Turning that interest into meaningful business results is.

Britain has many of the ingredients needed to lead: world-class research, a thriving start-up ecosystem, renowned universities and entrepreneurs[1]. But the next phase of the AI race will not just be about who can build the most sophisticated models, but by those who embed AI across the everyday economy to improve productivity, competitiveness and resilience.

There are encouraging signs: the United Kingdom ranks eighth globally for AI diffusion[2], while British Chambers of Commerce research found that SME adoption rose from 25% to 54% in 2026[3]. Yet only 11% used technology extensively, and adoption varied sharply by sector[4]. ONS data shows that adopting businesses used an average of just 1.6 AI technologies in 2026[5], while Deloitte found that only 24% of British workers used generative AI daily and 31% did so without their employer's knowledge[6]. The challenge is, therefore, not interest in AI – but the depth of adoption.



Using AI is not the same as transforming with AI

Using an AI tool is easy to count as adoption: an employee drafts an email, summarises a meeting, creates content or finds information faster. These uses save time and build confidence – but they are the beginning of the journey, not the destination.

Transformation begins when a business asks what AI could enable it to do differently and what measurable value that would create.

That could mean improving customer service, helping salespeople win more deals, reducing factory downtime or enabling better decisions through faster analysis.

I think of adoption as a ladder: awareness, experimentation, scale, redesign and transformation. government research supports this distinction by examining not only whether businesses use AI, but whether they scale it, what prevents adoption and whether it improves productivity and revenue[7]. The challenge is to help many more British businesses climb beyond the first two rungs.

Why getting this right for SMEs matters

SMEs represent 99% of British businesses[8], so the United Kingdom cannot become an AI adoption superpower if transformation is confined to large enterprises and a small group of digitally mature firms. Moreover, between 2023 and 2025 uptake among large businesses (250+) nearly doubled to 44%, vs 26% for firms under 50 employees (Bennett School, 2026)[9].

Research from Microsoft and WPI Strategy estimated that wider AI adoption among British SMEs could generate an additional £78 billion in economic value over the decade to 2035, with gains spread across regions[10]. AI adoption is, therefore, a growth, productivity and regional prosperity issue.

Most SMEs do not have an AI strategy team, data scientists or a transformation office. The person exploring AI may also be managing customers, rising costs, cash flow and today's deliveries. They need practical answers to three questions: Where could AI make the biggest difference? How can we use it safely? How will we know it is creating value? Policy must help businesses answer them.

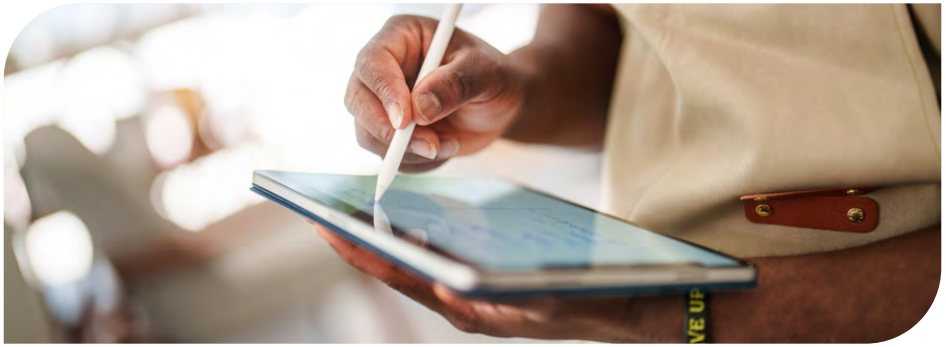
This is as much about people as technology

Microsoft's 2025 Work Trend Index highlights a growing capacity gap: 53% of leaders said productivity needed to increase, while 80% of the global workforce said they lacked the time or energy to do their work. Employees were interrupted by an email, meeting or notification every two minutes on average[11]. AI adoption must address how work is designed, not just which tools are deployed.

The businesses making the greatest progress are thinking about people, skills and processes alongside technology. They are asking which work should remain fundamentally human, where AI can remove friction and where people and AI working together can produce an outcome neither could achieve as effectively alone.

That requires leadership. AI skills should not just mean teaching people how to prompt. Business leaders and managers need enough understanding to identify opportunities, redesign work, support people through change, govern AI responsibly and measure whether the investment is working.





Three practical steps towards an AI adoption superpower

1. Help SMEs identify clear use cases and benefits

The Department for Business, Innovation, Science and Trade has already invested in business.gov.uk and the Backing your Business campaign. BIST should leverage these investments to make visible AI use cases organised around practical SME challenges, such as using month-end reconciliation to improve cash-flow management. Each use case should set out indicative costs, expected benefits, potential productivity gains and how to measure value, creating a shared evidence base for AI adoption.

BIST could work with industry and professional services firms to ensure relevance and credibility, with intermediaries adapting content for sector-specific use. The library should allow segmentation by demographic, ensuring it can be leveraged easily by differing industries and organizational sizes.

2. Make AI adoption local

National ambition matters, but adoption happens in real companies, sectors and places. A manufacturer in Birmingham is unlikely to have the same starting point as a professional-services business in Reading or a tourism company in Cornwall.

Regional Growth Hubs could convene organisations that businesses already know and trust, including Chambers of Commerce, combined authorities, industry bodies and responsible technology partners. The lesson from Germany's Mittelstand model is that SME innovation is best accelerated through a long-term ecosystem of practical skills, impartial local support and targeted innovation programmes that help businesses translate new technology into commercial value[12].

Growth Hubs should leverage existing programmes and investments to drive increased take up of practical advice, demonstrations, and short sector-specific programmes across priority regions. Businesses want to see companies like them using AI and ask straightforward questions: What did you do? What did it cost? What went wrong? What changed? Was it worth it? AI diffusion needs national ambition, but it will depend on local trust.

One model already showing promise is the use leveraging the partnerships formed through AI Skills Boost to deliver regional AI agent-building workshops, where small businesses bring a repetitive, time-consuming process and, in a single morning, learn how to apply AI to improve efficiency, response times and consistency.

3. Test, learn and scale AI transformation

BIST should establish an AI Transformation Challenge Programme focused on priority productivity challenges identified with sector departments. Regional partnerships would bring together SMEs facing similar challenges and provide practical support with implementation, workflow redesign, skills and responsible governance.

The programme should build on proven approaches. Britain's Made Smarter Adoption programme, for example, combines expert advice, transformation roadmaps, skills training and implementation support, increasing adoption of industrial digital technologies[13].

Businesses should establish a baseline and report a small set of outcomes, such as time saved, errors reduced, capacity released or revenue increased. Successful approaches should feed into helping government identify trends, understand what works and scale lessons across industries.

The opportunity is ours to take

The United Kingdom already has the research base, businesses and ambition to lead. The task now is to turn that strength into practical capability across the economy. Highlighting AI use cases would help SMEs see where AI can create value; a locally convened network of partners would give them trusted support to act; and an AI Transformation Challenge Programme would identify, prove and scale the interventions that deliver measurable gains.

Together, these policies would move Britain from awareness to implementation – and from isolated experiments to sustained business transformation.

The goal is not to have the most businesses trying AI; it is to become the country where every business can turn AI into measurable growth.

15

Turning AI into everyday SME growth



Abdul Quershi
Managing Director
Barclays Business Banking

Abdul Quershi is Managing Director of Barclays Business Banking. He leads a division of over 4,000 Barclays employees who support more than one million SME founders across the UK.

Abdul argues that the real productivity prize for the UK will come from helping SMEs move from experimentation with AI to confident, embedded use across everyday business tasks.



The UK is in an enviable position. It has a thriving and diverse small and medium enterprise (SME) ecosystem, with 5.7 million private-sector businesses – around 5.42 million of them are micro (fewer than 10 employees) – as well as world-class research and a strong pipeline of high-growth technology companies[1].

But to be an AI adoption superpower we need AI to be routinely improving decisions, increasing capacity and driving growth across the small businesses that underpin the UK economy.

It's no secret that the UK is facing a long-standing productivity challenge. The government's SME Digital Adoption Taskforce estimates that even a 1% productivity uplift across SMEs could add £94 billion a year to UK GDP[2]. Used well, AI can give an owner or a small team access to capabilities in analysis, marketing, customer service and administration – it democratises tools and resources at a marginal cost.

The aim is not to replace jobs, but to release scarce capacity to focus on value-add work, deepening relationships with clients and identifying new sources of revenue.

The UK begins from a position of strength. The Barclays Eagle Labs Ones to Watch AI:100 report profiles 100 AI businesses across all four devolved nations, spanning software-as-a-service, data and analytics, robotics and automation. Collectively, they have raised more than £11.3 billion, generate turnover exceeding £734 million and employ more than 8,500 people[3]. The technology is being created, but the bigger opportunity is diffusing it across the millions of small businesses across all regions and all sectors from agriculture and manufacturing to the creative industry.

EMBEDDED ADOPTER

James Wiltshire, Freemasons (Artist), Music Industry

James' "eureka moment" came just over 18 months ago, when he used AI to solve an IT issue. A task that would usually take a full day was completed in 20 minutes – and the solution remains in use.

Since then, James has embedded AI across several workflows: troubleshooting technical problems, building internal applications and automating time-consuming data conversions. He's developed logical code frameworks around large language models so they can be applied consistently and safely to practical business tasks.

Crucially, the output isn't always another AI tool; it can be a reliable application that continues to run without further AI assistance.

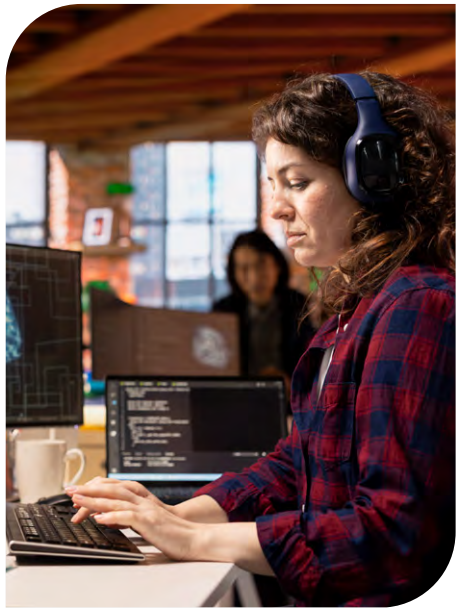
The impact goes far beyond efficiency. AI has saved his two-person company thousands of pounds and released enough time and capacity to pursue projects James says he "would never have even considered possible a year ago".

AI hasn't just made James' business more efficient – it's fundamentally changed what they can achieve, enabling a micro-business to compete at a higher level, unlock lucrative new opportunities and pursue growth that was previously out of reach.

Adoption is growing – but depth and reach matter more

According to the latest data from the Barclays Business Prosperity Index, nearly half of SMEs (47%) have invested in AI, with 14% describing their investment as significant. More than two in five SMEs say AI and automation have improved productivity or operational efficiency, while businesses are also using or planning to use AI for data analysis, administration, customer experience and cost reduction[4].

The challenge is no longer simply to persuade SMEs that AI matters, but to help them move from occasional use to sustained adoption.



When done well, founders – like James Wiltshire – have accessed opportunities that previously felt unattainable.

James' experience shows us what can be achieved through deep adoption and integration of AI at every level, but most founders only access a fraction of its full potential. Whilst overall adoption is building there are still challenges - in 2023, 39% of firms cited difficulty identifying relevant activities or use cases, 21% cited cost and 16% pointed to insufficient expertise and skills[5]. There are practical issues at play here - a business owner managing finance, people and sales may understand the opportunity but lack the time to compare products, redesign a workflow, manage data risks and calculate the return.

General encouragement won't close the gap. We need different policy objectives - counting licences, course completions or chatbot users won't reveal whether adoption is succeeding. We should measure outcomes: hours released from administration, faster customer responses and the proportion of SMEs that progress from experimentation to embedded use. This should be grounded in simple, responsible habits that work for small businesses – keeping data safe, checking AI-generated work and keeping security front of mind.

Small businesses can be the UK's adoption advantage

SMEs are the engine of growth, which is why Barclays has launched Backing Britain's Future: a commitment to support 1.5 million entrepreneurs by 2030 and drive investment, innovation and growth across the UK.

This ambition is underpinned by continued investment in Business Relationship Managers, Eagle Labs and the Innovation Bank. AI adoption is an important part of this support, with tailored regional masterclasses, networks and the Eagle Labs Academy helping businesses build their knowledge and capabilities.

Through Barclays LifeSkills, we are also equipping young people with entrepreneurial skills and introducing them early to the opportunities created by AI. Together, these initiatives can help SMEs turn their agility into an advantage. Despite having limited resources, smaller businesses often have shorter decision chains, fewer legacy systems and closer links between owners, employees and customers.

This enables them to experiment with AI quickly and practically – as Shair Bashir found (below) when initial testing revealed wider opportunities for her business.

Shair's experience shows how testing, trying, even dabbling with AI capability can free up capacity for more important tasks. It also illustrates that early success can be a starting point rather than the end goal, with scope to deepen adoption by applying AI across more workflows as confidence grows. Policy should frame AI as everyday productivity and growth infrastructure, not a specialist innovation programme.



EARLY-STAGE ADOPTER

Shair Bashir, founder and master artist of The Make-Up Box, luxury mobile hair and make-up studio, based in London

Shair wanted more time for the personal relationships at the heart of her make-up business. Initial enquiries, planning and group emails took time away from understanding clients and building trust.

Her first low-risk use was shaping a website rebrand. AI helped her visualise ideas and explore options, while people finalised the work. It saved considerable back-and-forth and gave her greater clarity and confidence.

Shair remains careful about how she uses AI. Authenticity and the human touch matter to her, so she's wary of impersonal AI-generated content. However, she sees potential for a chatbot linked to her CRM to handle initial enquiries and identify genuine opportunities, freeing her to deepen client relationships.

Her experience is practical and relatable – she's used AI for small, clearly defined tasks retaining human judgement but also building her confidence and unveiling further opportunities for efficiency and automation.

A policy compact to drive adoption

Government could build on their strong network of business groups, technology providers, banks, universities and local growth organisations to give small businesses a practical journey from business problem to tested solution.

Four interventions would be helpful.

1. Have a trusted adoption front door (i.e. via Business Growth Service)
Businesses could complete a short, practical diagnostic based on the outcome they want – more sales, quicker administration, better forecasting or improved customer service – and be directed to relevant advice, approved support and local expertise.

2. Build sector-specific use-case libraries
As advocated for in our Barclays 2025 report – Unlocking Small and Medium-sized Enterprise Investment – publication of concise, independently tested examples showing the problem solved, implementation effort, safeguards, cost and measurable benefit. Peer evidence from similar businesses will be more persuasive than generic claims about AI.

3. Make skills and responsible use inseparable
Owners and employees need practical skills to identify use cases, evaluate outputs and redesign work. Proportionate guidance on data protection, cyber security, accountability and human oversight should be built into every journey rather than treated as an additional compliance exercise.

4. Measure and share business outcomes
The programme should publish comparable evidence on time saved, productivity, revenue, resilience and movement from trial to sustained use. This would provide accountability for public investment and build the confidence that comes from seeing businesses like yours succeed.

The compact should mobilise existing networks that have trusted relationships with SMEs, like the Chambers of Commerce, accountants and banks. Connecting networks would give policy both reach and credibility.

AI is the next all-purpose technology

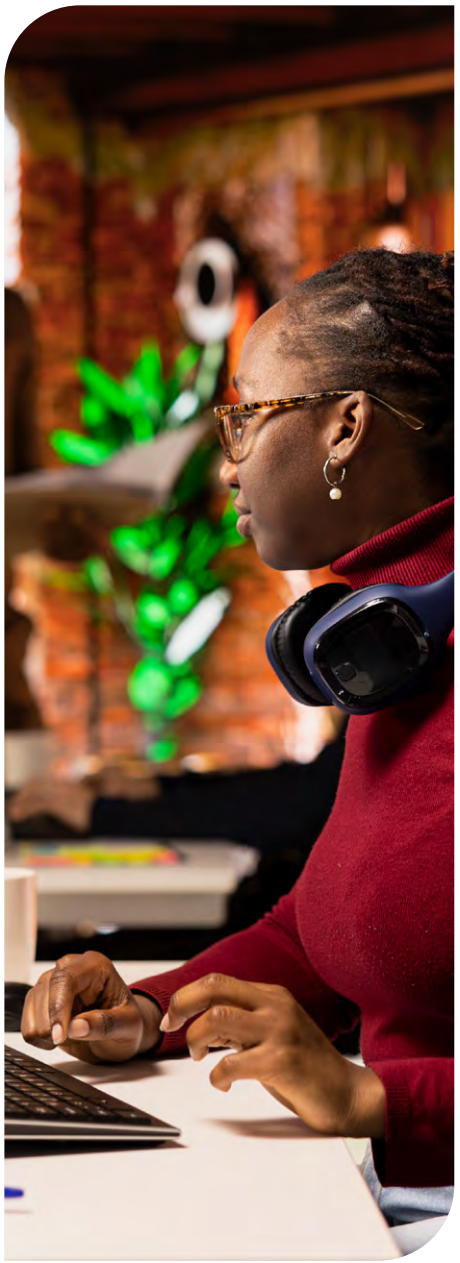
AI can transform every sector but growth depends on helping SMEs find valuable use cases, test them safely and prove the return. Policy should help businesses start small, learn quickly and deepen adoption as value emerges.

Recent economic growth offers an early signal of that potential: the UK economy grew by 0.4% in July 2026, against expectations of no growth, supported by strong services performance, particularly computer programming[6]. AI and related technologies are already contributing to growth; the task now is to broaden their impact.

A policy compact would turn the UK's strong AI foundations into higher productivity, stronger growth and more resilient businesses.

The UK will become an AI adoption superpower when business owners move beyond uncertainty, build confidence through experience and become excited by what AI can make possible.

As James put it: "AI became a tool, not a threat. The moment I turned fear into excitement, everything changed." – James Wiltshire, Founder



Why people determine success



16

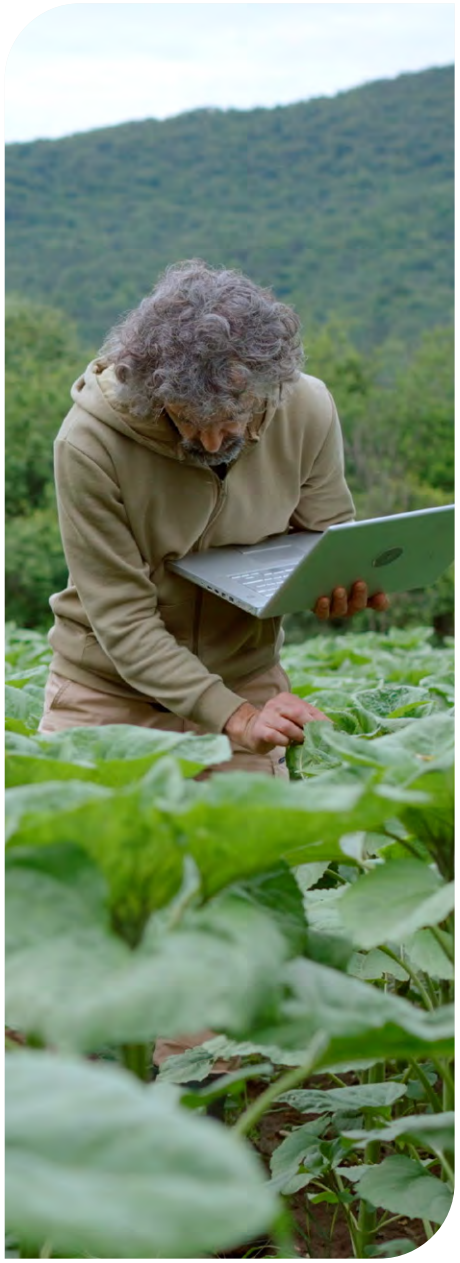
The real AI challenge is not technology. It is change.



Professor Kirstine Dale
Chief AI Officer
Met Office

Professor Kirstine Dale is the Met Office's Chief AI Officer and Principal Fellow for Data Science. Kirstine is Co-Director of the Natural Environment Theme of the Turing Research and Innovation Cluster on Digital Twins. She is also an Honorary Professor at the University of Exeter.

Kirstine argues that the UK's real challenge in becoming an AI adoption superpower is not a lack of technology or ambition, but the ability of organisations to manage change well. She writes how leadership, culture, communication and workforce confidence must be at the centre of AI change.



The United Kingdom has no shortage of AI ambition or potential. We have world-leading universities, vibrant technology companies, strong public institutions and a growing ecosystem of AI innovators.

Yet there is a gap between AI experimentation and AI transformation. Organisations pilot AI, but relatively few are embedding it deeply enough to change how they operate, serve customers, develop products or generate value.

This distinction matters. AI adoption is not fundamentally a technology challenge. It is a change challenge.

More specifically, how can change be better managed, more responsible, effective and sustainable?

When I was asked to draft an article on 'How The UK Can Become An AI Adoption Superpower', I initially wondered how I could achieve this in just two thousand words. Upon reflection, I think just four are sufficient: 'lean into change management'.

Adaptability is now the real competitive advantage

The environment in which organisations operate today is characterised by continuous disruption. AI is perhaps the most visible example of this phenomenon today, but it is unlikely to be the last. In such a context, competitive advantage comes not from mastering a single technology, but from developing the organisational capability to repeatedly and rapidly adapt. That is certainly my experience as the Met Office's Chief AI Officer where the AI Revolution has necessitated one of the biggest changes in our organisation's 170-year history. AI, and the technological opportunities that it presents, are potentially game-changing, but I propose that its power to disrupt will not move from a challenge to an opportunity unless we embrace the change that it necessitates.

If we want to become an 'AI adoption superpower', we may need to first become a 'change management superpower'.

This is not simply about implementing today's AI tools. It is about creating organisations that are agile enough to embrace continuous transformation. I'm a dyed-in-the-wool scientist but when navigating this period of technological disruption, I find myself increasingly leaning into the Change Management studies from my MBA and courses on Leading Digital Transformation and Innovation.

There is a mountain of advice available on change management, I like to pick and choose based on the challenge/opportunity that we face. My first port of call tends to be Kotter (arguably a 'Godfather of Change Management'), and in particular his 2014 book Accelerate[2].

As Kotter noted in a Harvard Business Review article[3] that preceded his book, 'perhaps the greatest challenge business leaders face today is how to stay competitive amid constant turbulence and disruption'. That was in 2012. It would be nice to say the world has calmed down. It hasn't. Kotter observed that organisations increasingly need the ability to sense opportunity, mobilise people, learn rapidly and adapt continuously- that was as true fourteen years ago as it is today.

At the Met Office, our experience embedding AI has reinforced a simple lesson: successful AI adoption depends far more on people than technology. The most important investments are not always in algorithms or infrastructure. They are in culture, leadership, engagement and behavioural change. Successful change starts and ends with people.

I'll highlight two change management concepts that have been particularly effective for us.

Create a Sense of Urgency Around a Big Opportunity

Kotter identifies creating a sense of urgency as the foundation of successful transformation. Importantly, this urgency must be linked to a compelling opportunity rather than simply a fear of being left behind. It must be both strategically rational and emotionally meaningful. We've taken this advice to heart.

At the Met Office, AI is not positioned as a standalone technology initiative. Instead, it is framed within a much larger opportunity with porous edges that everyone can access.

The challenge facing the Met Office is profound. Demand for trusted weather and climate intelligence continues to grow, while the volume and complexity of environmental data expand exponentially. At the same time, new machine learning approaches are fundamentally reshaping weather prediction systems. AI models can produce forecasts in seconds that previously required vast supercomputing resources[4,5]. It can help us identify patterns within enormous datasets, support decision-making during extreme weather events and better understand risks[6,7,8,9,10]. The same opportunities for transformation can also be applied for the use of AI in climate modelling[11,12,13,14].

The question was not whether AI would affect how we predict the weather. The question was whether we would actively shape that future or simply react to it. This was the urgency needed to the change of embedding AI.

We launched an enterprise-wide transformational change programme: 'AI4Everyone' and focused on demonstrating how AI could accelerate scientific discovery, enable us to focus on areas where human expertise adds most value, enhance services and unlock new forms of value for customers and citizens. A keystone principle of AI4Everyone was that AI adoption is a workforce challenge more than a technology one. With a people-centred approach, the programme is a central part of the Met Office's ongoing AI journey taking staff from cautious to curious, and curious to confident.

Critically, the narrative is not one of replacement but augmentation. AI is positioned as a capability that could help scientists, technologists and operational teams achieve more, rather than a threat to their expertise. I'd be lying if I said it has been easy. There have been, and continue to be, many twists and turns, including concerns over the environmental implications[15] of using AI and the perceptions of accuracy and confidence using AI generated forecasts[16]. Our commitment to responsible AI, transparency and taking 'everyone with us' has meant that we faced tough decisions head on.

For the UK, and in particular public services such as the weather forecast, I suggest that a key component of change management is moving the narrative from one of promoting AI as a technology towards the benefits and outcomes that it can support. Organisations, need a clear and inspiring narrative about how AI can enrich our working environment, strengthen services, competitiveness, productivity and growth. Adoption accelerates when people understand not just what is changing, but why it matters.

Head and the Heart, not just Head.

A crucial lesson for change management comes directly from Kotter's principle of appealing to both "head and heart". People support transformations when they connect emotionally with the purpose behind them and resonate with their values.

Some of the strongest momentum emerged not from formal strategy documents but from stories. Stories about scientists accelerating research. Stories about colleagues eliminating repetitive tasks. Stories about teams creating new solutions for customers. Stories about people rediscovering time for creativity and innovation.

These stories transformed AI from an abstract concept into something tangible and personal, something human.

The experience aligns strongly with emerging thinking about organisations operating in highly complex environments. For example, Werner and Le-Brun[17] argue that organisations need to move beyond traditional command-and-control models towards more adaptive structures built on ownership, curiosity and distributed innovation. They describe "Octopus Organisations" that thrive through continuous learning and experimentation rather than centralised control.

This resonates strongly with AI adoption.

The organisations making the greatest, sustainable progress are rarely those where AI has been imposed from the centre, or even worse, the top. They are likely to be the organisations where employees have been empowered to explore, experiment and innovate. At the Met Office, the most valuable use cases have often emerged from frontline teams closest to real operational challenges.

Our lesson has been clear, AI adoption strategies should not focus exclusively on technology infrastructure. Equal attention should be given to leadership capability, workforce confidence and organisational culture. AI adoption is ultimately a human endeavour.

A Call to Action

The UK has every opportunity to become an AI adoption superpower.

But this will not happen through technology adoption alone.

I believe successful AI adoption hinges on change management, and change begins and ends with people. We have every reason to be optimistic, this is well within our wheelhouse. We'll need to see more intentional change management, more human-centred design, and greater empowerment through the right environments, skills, principles and guardrails. Most of all, it will require better communication so that people understand the opportunity, feel part of the journey, and have confidence in the future we're building together. My plea to leaders wanting to adopt AI is to dust off your change management books - there are some corkers that have stood the test of time and can help us find a successful route through the labyrinth, whatever the weather.



17

Is the UK's AI leadership stuck in analogue mode?



Ann Francke OBE
CEO
Chartered Management Institute

Ann Francke OBE is Chief Executive of the Chartered Management Institute. She started her career at Procter & Gamble and has held senior executive positions at Mars, Boots, Yell and BSI. In 2020, Ann was awarded an OBE for services to workplace equality.

Ann explains how AI will only boost UK growth if firms improve leadership and management, not just buy new tech. SMEs are showing the way, but wider adoption needs better training and clearer guardrails.

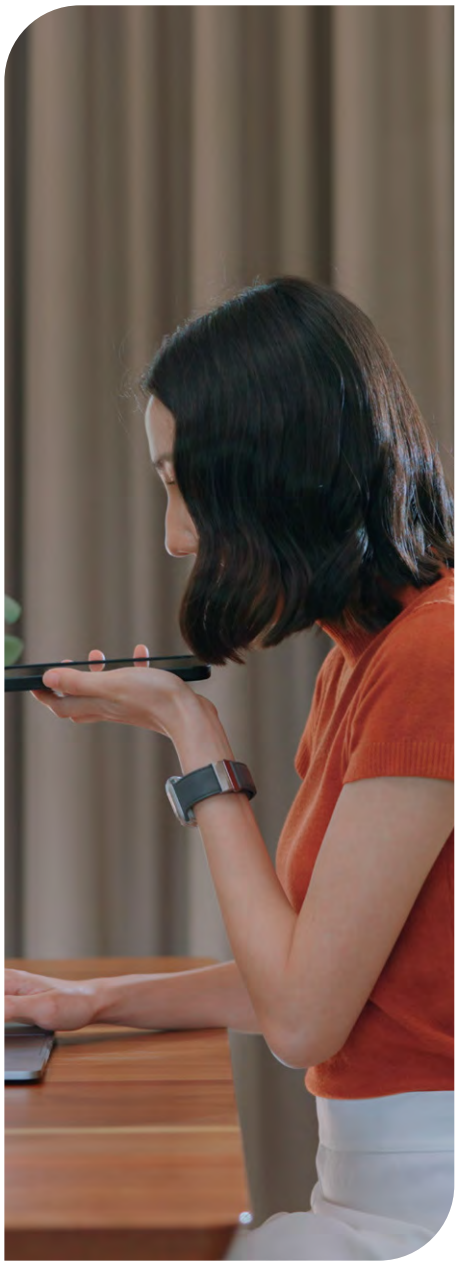
The UK government has set an admirable goal: establish Britain as the fastest-adopting AI economy in the G7 and unlock up to £140 billion a year in economic growth.

This is a challenging ambition. But as leaders across business and government, we need to ask ourselves a frank question: is this blind ambition, or are we acting with clear intent and laying the proper foundations for success?

Right now, business leaders are throwing money at new technology. An impressive 86% of employers expect AI to fundamentally transform their operations by 2030. Yet, beneath the boardroom hype and the mounting software invoices, British enterprise faces a sharp reality check.

Spending billions on AI tools will not deliver meaningful shifts in productivity if the people charged with implementing them are operating in analogue mode.

Research from the RAND Corporation lays it out in plain English. An extraordinary 84% of AI implementation failures are not caused by the lack of investment in technology. They are driven by leadership failures and weak workplace cultures. In short, it is a human problem, not a technical one.



The view from the top (and the bottom)

At CMI, our recent research report, Artificial Intelligence; Real Leadership, bears this out. We found a glaring gap between what organisations are spending on AI and where they are seeing an actual impact on productivity. More than two-thirds (68%) of British organisations remain trapped in "pilot purgatory", experimenting or running early-stage AI trials without moving into scale-up mode.

While 70% of UK managers report basic efficiency gains from AI tools, a tiny 5% have achieved truly transformational productivity improvements. Worse still, more than a quarter (26%) report seeing no gains whatsoever. This productivity bottleneck persists because technology is only one piece of the puzzle. The rest comes down to workplace culture, employee trust and effective leadership.

Right now, British enterprise is suffering from a double capability gap spanning senior executives and frontline management. At the executive level, senior leaders champion AI loudly in theory, but fail to model it in practice. While 64% of senior leaders encourage staff to experiment, only 13% of managers strongly agree that C-suite executives actively use the tools themselves. It is hard to lead a digital revolution from the sidelines. Furthermore, just 18% of managers believe senior leadership understands AI's commercial benefits, and a worryingly low 8% agree that executives are actually monitoring ROI.

On the frontline, middle managers serve as the vital bridge between strategy and execution, yet only 12% feel very confident leading AI-enabled teams. When asked about managing autonomous agentic AI, confidence drops to just 10%.

There's another piece to this puzzle that we cannot ignore. The base level of management training is stagnant with only a quarter of managers rating their own management training as sufficient. Employees are turning to AI for the guidance that they are not getting from their leaders and managers - 70% of managers told us they turn to AI chatbots for workplace guidance instead of their own bosses. They find the responses they receive come quicker, are clearer and are non-judgmental compared to what their manager can give them. Add to that the acknowledgement of a growing use of so-called 'grey AI' - unauthorised use or use of AI without clear guardrails in place. More than one in three managers (37%) say they know it's happening and that it poses an active risk in their business.

SMEs on the front foot

There are several bright spots emerging in the UK's business landscape. Amid big corporate hurdles, it is small and medium-sized enterprises that are giving the rest of the country some insights into how to turn technological potential into commercial performance.

SMEs are currently outpacing larger corporations in active experimentation, 51% compared to 34%. As a result, they are extracting higher early productivity gains (33% versus 25%). Why? Because SME leaders are more hands-on, they don't have the luxury of farming out AI adoption to someone else. The result is that they are demonstrating significantly higher personal leadership confidence (31% versus 23%).

Without endless committees or enterprise-wide working groups, they are unencumbered and agile. SME leaders prove that active, personal involvement is the real secret to unlocking immediate business value. They show us that agility and genuine leadership engagement can turn technical tools into bottom-line performance.

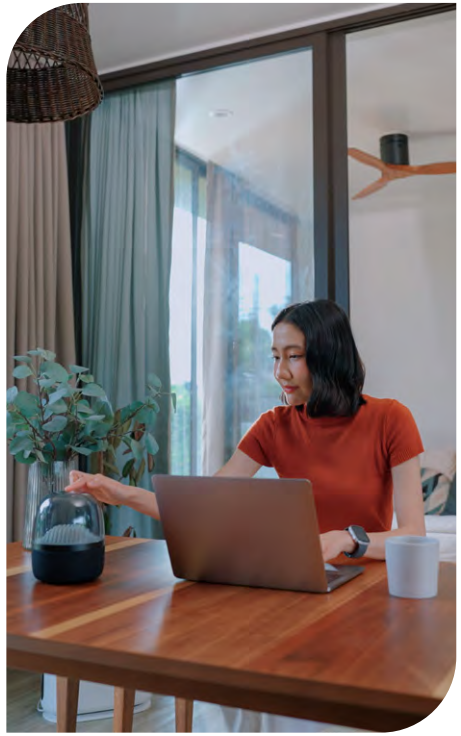


Ambition to action

If we want to support adoption at scale and turn the UK into a true AI superpower, government policy must evolve and leaders across business and in the public sector need to act to address the management skills gap when it comes to adoption. CMI is proud to actively support this national vision as a strategic partner in the government's AI Skills Boost initiative. Through this work, we are able to advise on where the gaps lie if we are to realise these gains, escape that loop of endless pilot projects and start to capture evidence of real wins and true productivity. At every level of the economy, there needs to be an explicit focus on building human capital and management capability across the workforce.

We need to understand what is holding us back. A winning national strategy must rebalance technological expenditure with organisational change management, treating AI adoption as a major workplace transformation project rather than a new piece of IT infrastructure to be procured and rolled out. government needs to partner with professional bodies to scale contextualised management qualifications nationally. We need an 80/20 rule embedded across every career stage, where 80% of training focuses on leadership, change management, and ethical oversight, paired with 20% practical technical literacy.

Other realities that we cannot allow to fall off the priority list include the need to collectively address how we protect early-career development and bridge the middle-management gap as routine entry-level tasks become automated. Incoming Gen Z talent must be supported as accelerated "citizen developers" through structured mentoring and redesigned roles, while managers must be trained to transition into "agent bosses" and "intelligent commissioners" who exercise critical human judgment over hybrid workflows.



Bringing everyone with you

Employers need to be thinking about how to future proof their adoption, looking at opportunities such as the new Leadership of AI Apprenticeship units developed by Skills England and asking themselves if they are giving today's employees the best possible chance of supporting AI-fueled growth through their own learning and development plans. CMI's own suite of leading AI regulated qualifications, alongside flexible 'masterclass style' programmes, were designed to support that need from early career line managers all the way through to middle management and executive leadership.

Finally, government must ensure that SMEs, the backbone of the UK economy, are not left behind. Be it practical strategy toolkits, frameworks for identifying and mitigating risks, sponsoring regional peer-learning networks or funding dedicated training in how to lead AI adoption, we need to move at pace to deliver equitable AI that reaches into every area of the country and every business, regardless of size.

Competitive advantage is being reshaped right before our eyes. The risk for British enterprise is not merely becoming slightly less efficient, but falling materially behind competitors who successfully integrate human leadership with technological know-how.

To deliver sustainable growth across the UK, our national mantra must be clear: leadership first, technology second.

It may be artificial intelligence, but turning investment into real economic growth requires real leadership.

18

The interaction of people and technology is at the heart of AI growth



Rt Hon Chloe Smith
President
Chartered Institute of Personnel and Development

Chloe Smith is a senior leader with expertise in technological change, workforce transformation, global economics, regulation, risk and governance. As an MP and government minister she was the first to head both the technology and labour market departments in the UK and she now works across the public, private and charitable sectors.

Chloe argues that AI adoption is not just a technology challenge but a people challenge. She writes that HR leaders are central to building the skills, trust and confidence firms need to use AI well.

AI adoption is the challenge of our age. The technology offers enormous opportunity to almost all businesses, their customers and their workforce, as well as to citizens and their public services.

This industrial revolution holds out a prize in productivity likely to be in the billions. And it's happening fast. But people need skills, confidence, understanding and trust[1]. So adoption is key, and that depends on effective decisions by leaders and action throughout organisations. Those who focus on people, alongside technology, will enable progress. In the best businesses, HR may lead successful adoption of AI.

For twenty years I've worked in different fields of human capital. I'm one of very few who have led the labour market and technology ministries of a major economy. I believe profoundly in the importance of people. And I also believe HR is hardly for the fainthearted when an erstwhile human resources manager runs NATO[2].



Human Resources Manager moment to lead

This is a pivotal moment for the people profession. At CIPD, the professional body for HR and people development, our new leadership team is united in our view that people professionals are uniquely placed to lead on productivity and performance[3].

For new CIPD CEO Neil Carberry, AI is a significant opportunity. HR professionals should play a leading role in shaping how new technologies are introduced, ensuring they support organisational goals and help people adapt[4].

As CIPD Chair Chris Bones said earlier this year, "In an AI-enabled, global, flexible economy, knowing who your workforce are, their skills and motivations... poses challenges. Operating models will change as business models and service delivery models change. This is functional expertise that is critical to future performance. This sits with HR."[5]

And as we listen to our members about taking the lead on AI, Emma Parry of Cranfield University adds: "the skills businesses need are changing daily. Who else is going to do this? We've never needed HR more."[6]

Skills, trust and confidence will decide success

Globally, reskilling and upskilling is a top priority for 9 in 10 companies. This is unsurprising, given the World Economic Forum's estimate that 60% of all workers will need to shift their skills substantially within this decade. Yet data from CIPD shows that as many as around 45% of UK organisations are still in a stage of ad hoc and experimental use of AI. They've yet to build the strategy or organisational capabilities required[7].

Businesses everywhere in Britain, already facing immense pressure, will know this feeling. As BCC showed in its recent work with small firms across the country, "we're at a pivotal moment in the digital transformation of the UK's business landscape. Artificial intelligence has the power to be truly transformative for our SMEs – those companies who are the backbone of our economy. It's truly exciting, but understandably for some firms it is either daunting or difficult, or both."[8] Shevaun Haviland was right to argue this summer that productivity comes when businesses have both the tools and talent they need to succeed, so people and skills are central to embracing the AI opportunity[9].

Human capital is economically essential. That's why the HR profession is chartered. The CIPD's Royal Charter is a commitment to public benefit, because the impact of good people management reaches well beyond individual organisations and into the economy and society.

HR is already leading technology transformation successfully. I found it thought-provoking meeting Steph Coward, MD at software firm IRIS, who explained their internal AI transformation to me and 12,500 attendees at our summer conference: "We learned that making AI work isn't an IT project — it's a people project." At IRIS, she illustrated, AI adoption has been led by its people and talent function, not IT – a deliberate choice that offers a blueprint. Rather than mandating new tools, IRIS invited colleagues to pitch their own ideas for how AI could improve their day-to-day work as one of its first steps in rolling out AI across the business. That momentum was backed by real investment in people: company-wide initiatives give people hands-on experience and ongoing capability-building, with senior leaders taking part alongside their teams. I see this demonstrating trust with transparency.

The insight is simple, and it's one HR is well-placed to show: give employees a genuine say in how AI gets used, and build capability - adoption follows naturally[10].

Today's HR leaders have a big job, then. Rapid advances in technology, changing workforce expectations, pressure on productivity, and increasing demands for commercial capability are fundamentally reshaping the HR profession. The strongest Chief People Officers increasingly combine deep people expertise with technology, strategy, the balance sheet and business transformation. The people function is playing an increasingly strategic role in shaping organisational resilience and business growth.

As the boundaries between people and technology increasingly converge, much greater collaboration between HR and IT is critical, bringing together technology expertise with an understanding of workforce capability, organisational design, governance, culture and behavioural change[11]. The organisations that succeed will see a fundamental truth: that the interaction of people and technology, rather than either one on their own, is at the heart of growth.

So the best in HR will step up to lead AI adoption because it's overwhelmingly about people, and business cannot succeed or grow without a solid approach to the resilience of its people.

Yes, this might be one of the biggest challenges for many companies. Yet it's where the competitive edge is.

Firms that embrace this leadership challenge will enable AI adoption, and will get ahead. Firms that don't, will fall behind.

Government can help firms unlock the AI prize

Now, it's business not governments that deliver growth. But government can support in at least two key areas. Firstly, government should improve the incentives for business to invest in skills, across a range of central and devolved levers and addressing the overall burden of cost and regulation. Secondly, government should use its responsibility for the jobs and careers service and welfare to raise digital inclusion and match hidden British talent with the vacancies available.

Together, without doubt, leaders in government and across the economy should be talking about how to lead change with AI. I'm pleased to work with British Chambers of Commerce to help leaders take the essential first step of building trust, investing in skills and securing organisational capability.

AI adoption matters enormously to the UK's growth, productivity and competitiveness – and it all comes down to people.

19

To reimagine the future of work we need to listen to workers



Tabitha Goldstaub
Co-founder and CEO
Lichen AI

Tabitha Goldstaub is the Co-founder of Lichen AI, helping companies to reimagine the future of work by listening to their people. She is a former Chair of the UK government's AI Council and member of the Digital Economy Council for which Tabitha was awarded and MBE for services to Artificial Intelligence.

Tabitha argues that AI will only boost productivity if workers shape how it is used. She urges ministers to make "pro-worker" AI the norm, so firms grow faster while staff come with them.

The UK is explicit in its ambition to be a global AI Superpower, with a desire to 'lead the world in AI adoption, innovation, and economic growth'. On many fronts, we're making good progress.

Britain is already a leader in many critical ecosystem components, from AI talent to investment in AI research, venture and infrastructure[1]. But lift the lid on the evidence around business adoption of AI, and you'll find a less advanced picture.

Microsoft's latest AI Diffusion Report puts the UK eighth in the world behind European neighbours Ireland, France, Norway and Spain[2].

While technologists focus on 'build it and they will come' mentality, UK plc should focus AI adoption plans around the strengths and frailties of humans, with ways for them to voice problems, propose solutions, agree alternative work and claim a share of the productivity gains.

The government can speed this up by making worker consultation best practice, backing regional models and championing pro-worker adoption.

Research from Deloitte[3] shows that over 60% of companies are using GenAI to some degree but only 12% do so multiple times a day. Understanding why they haven't felt the value should be more important than mandating more. The vast majority of people are using AI to enhance personal productivity, like writing emails, not to enhance the work of teams. The culture of 'what can it do for me' rather than 'what can it do for us' lacks capturing real revolutionary company wide gains. Only half of users actually think AI is saving time and even in cases where time is clearly being saved, many companies don't have a robust plan for how to productively deploy this new found additional capacity. This causes fear and uncertainty which is no wonder the Glassdoor[4] report shows a growing swell of negativity amongst UK employees towards the deployment of AI in the workplace.

As we understand more about the reality of the adoption curve for business, there is mounting evidence that the most effective strategies are those that take a people-first approach. Momentum is gathering around applying best practice that focuses on working in partnership with every business's greatest asset, its staff. The people who uniquely know the most about the job, its components, what needs to remain entirely human versus what could be automated.

Sarah O'Connor's new book, We Are Not Machines[5], makes a clear case: how new technology lands depends less on the tool than on whether the people doing the

job had a say before it arrived. At Boliden's Renström mine in Sweden, management had to negotiate automation with the unions. When the company proposed tracking exactly where every miner was underground, union representatives ensured a collective voice was heard. The result was a compromise: workers would appear as anonymous number tags, not names. The mine now produces roughly double the ore it once did with about the same number of people. Compare that with UK truck drivers whose employers installed AI cameras to watch their faces for signs of tiredness. Even with the best of intentions, many drivers felt the cameras made the job worse and accidents increased. Similar technology, very different experiences, and the difference wasn't the software. It was whether anyone talked to the people who would live with it, and understood what they value and what worried them, before deciding how to roll it out.

Only a third of UK employees say they were consulted before new technology arrived in their workplace. Where they were, 70% felt positive about its impact on their job. Where they weren't, just 20% did[6]. The Pissarides Review, led by the Institute for the Future of Work, reached the same conclusion: "firms that practice high-involvement HR and prioritise employee wellbeing secure better outcomes for both workers and firms"[7]. Microsoft's own guidance agrees that "centring worker voice in AI design boosts productivity, satisfaction, and skill growth"[8].

Consulting people makes the technology stick and makes the work better.

It's a win-win, so the task now is to make consultation standard practice and simple to do.

Leaving people out doesn't just cost goodwill, it stalls adoption and reduces the potential to increase revenues. IKEA's AI adoption demonstrates the potential gains. Their "Billie" AI chatbot was launched in 2021, and by 2023 Billie was resolving around 47% of customer enquiries. That's over 3 million interactions, delivering €13 million in operating savings. This gave IKEA the opportunity to make redundant or retrain the 8,500 call centre agents. They chose listening to the individuals, understanding the opportunity and deciding to retrain them to become remote interior designers. This new service is now responsible for an additional revenue of €1.3 billion[9].

In addition to the benefits of a people-first approach for both businesses and workers alike, there are also huge wins for the government - namely an AI economy that takes the UK's workers with it rather than leaving them behind. Not to mention, the near term growth benefits of UK businesses scaling despite the challenges of the wider global economy.

What can the government do to encourage more Pro Worker[10] AI adoption?

In addition to keeping up the pace on the fundamentals, the government should consider the following areas of focus:

Lead by example
Government should apply pro-worker adoption to its own workforce and its own spending. The government Digital Service Standard[11] already treats civil servants as users whose needs shape internal tools, and this should extend to every AI deployment. The AI buying framework already scores suppliers on ethics and social value but could, as the TUC suggests, ensure "Government procurement frameworks embed worker voice throughout digital transformation"[12].

Back regional models
Support mayoral authorities to run their own AI reviews, so each region understands its workers, challenges and opportunities, with taskforces like London's AI and Jobs Taskforce[13] and the Liverpool City Region AI Taskforce[14].

Champion pro-worker adoption
Create a common standard businesses can sign up to, use the AI Economics Institute[15] to track and highlight good practice, and celebrate businesses that take a people-first approach. The Pro-Worker Prize[16] is a good start. Make worker consultation an eligible cost in AI adoption grants and vouchers, and build best practice into the AI for Growth Hub[17], AI Skills Hubs[18] and other programmes that already work for small businesses.

Today, it's clear which jobs are at the mercy of the algorithm; in time, it will be less clear, and so it's vital that enough is done to recognise and protect livelihoods and what brings joy at work. Dehumanisation of work is not inevitable, but a different outcome will require different choices and a different distribution of power in the workplace. There is a world in which AI enables us to spend more time on the things we love, but that will require a concerted effort, eyes wide open and bold organisations ready to put people at the heart of reimaging work.

AI is not a force of nature: making it work is a series of choices.

There is an opportunity for British businesses to help define how the UK makes sure AI works for us, not us for AI.

20

AI skills in every apprenticeship



Euan Blair MBE
Founder and CEO
Multiverse

Euan Blair is the co-founder and chief executive officer of Multiverse, an educational technology company set up in 2016 that provides employer-led apprenticeships in data, technology and business roles.

Euan argues that AI will boost UK productivity only if ordinary workers are trained to use it. His answer is simple, put practical AI skills into every apprenticeship.

Few things matter more to Britain's economic future than how quickly and how deeply we adopt AI.

The OBR estimates that economy-wide labour productivity could be increased by around 2.3% over the next decade as a result of this technology. Against the backdrop of two decades of flatlining, the effective deployment of AI could finally be the breakthrough we need to jumpstart the economy and restore creaking public services.

But those gains depend on people. Specifically, a workforce that is trained in how to actually use the technology to their benefit, and the benefit of their customers and employers.

There is a practical way to build and train that workforce, and by extension to spread the benefits across the whole economy: embed AI skills in every apprenticeship.

Apprenticeships have a larger and more cross-cutting reach than almost any other educational mechanism, bar universities, with more than 735,000 people currently participating in one. If we use them effectively, we can reach hundreds of thousands more people in an immediate boost to the UK's AI adoption ambitions.

The tools without the talent

The good news for AI adoption is that businesses are waking up to the opportunity. ONS data suggests that the proportion of companies that have bought AI tools for their teams has almost tripled since 2023.

But the gains have not kept pace with these investments: 56% of CEOs report no return on investment from their AI spending, according to PwC.

It's telling that the frontier AI labs are now going beyond building the tools themselves to spinning up multi-billion dollar ventures that focus on helping organisations deploy them. They recognise that the barrier holding their customers back is human and organisational capability.

But if we want AI's benefits to be distributed economy-wide, we can't allow the gains to accrue only to the companies developing the technology. To capitalise on this revolution, we must build that deployment and adoption capability within every organisation, in every postcode.



Building capability requires time and structured learning

The will is there: our latest skills intelligence data shows that 97% of the workforce want to learn how to use AI.

Yet currently, the skills are missing. One of the reasons is that we are asking people to become fluent in this new technology in the gaps around their existing workload.

By including AI in every apprenticeship, we would enable hundreds of thousands more people to learn critical skills within time that is already protected.

Apprenticeships by nature offer on- and off-the-job learning: apprentices learn new skills in structured teaching time and then apply them to their real job. It is in part this feedback loop that embeds new skills and capabilities so effectively.

We know apprenticeships work. There is a reason why someone who completes a higher apprenticeship earns over £5,000 more than the median graduate five years after finishing their studies, according to the DfE. In our case, we've trained over 40,000 learners and nearly three in five get a promotion or pay rise by the time they complete their apprenticeship.

AI skills can't be left to AI specialists

The benefits of AI won't be delivered by brilliant technicians alone. They will come from nurses, teachers, accountants, factory workers, and everyone in between becoming more productive, inch by inch and then metre by metre.

That's because AI is only as good as the context you give it. So we have to put it in the hands of the people who have workplace context: frontline workers who know where the inefficiencies are and are empowered to fix them.

Today, AI and tech roles are heavily tilted towards male graduates based in London. A small handful of apprenticeships are now dedicated to AI, but relatively few people have the time, or an employer willing to release them to train in AI specifically. If the AI revolution is to lift productivity across the country, it needs a much broader base.

Apprenticeships can provide a vital part of that base. More than three quarters of a million people are already doing one in England. Nearly a third of new starts are in STEM subjects, and one in ten in digital technology. But the vast majority of apprentices are training for jobs where AI will matter enormously, yet no one is teaching them how to use it.

So the answer is to pair a system we already know works with the ambition to reach beyond AI and tech teams, and into frontline roles in every sector.

Government should embed a module on safe and productive AI use into every desk-based apprenticeship, and offer it voluntarily for all others. AI now runs through almost every job in the country, and the learning already taking place should reflect that.

Hundreds of thousands of people are already learning on the job.

Making every apprenticeship an AI apprenticeship would turn them into the AI champions of tomorrow, and give Britain one of its most practical routes to the productivity gains and growth the government is reaching for.



21

AI will not shape the UK by itself. We will



Baroness Martha Lane-Fox of Soho

Digital entrepreneur, Crossbench peer and former BCC President

Martha Lane-Fox is a digital entrepreneur and former BCC President. She became a crossbench peer in the UK House of Lords in 2013. She served as Chair of the London AI and Jobs Taskforce this year.

Martha argues government must move faster on support, skills and labour-market intelligence so AI boosts productivity without hollowing out entry-level routes or leaving smaller firms and younger workers behind.

The United Kingdom has spent years asking whether we can become an artificial intelligence superpower. I think there is a more impactful question. Can we become the country that is exceptionally good at putting this technology to work across the economy, while making sure that people have some say in how it changes their lives?

We should celebrate our brilliant researchers and ambitious founders. But inventing technology is not the same as making good use of it.

The real economic prize will come when organisations of every size are able to use AI to improve what they do and create opportunities.

There is still a considerable gap between having access to AI and knowing what to do with it.

ONS data show that reported AI use among UK businesses with ten or more employees rose from around 12% in late 2023 to around 35% in June 2026. The headline is encouraging. Look underneath it and the picture is less impressive. The average number of AI technologies being used has barely moved and only one in ten businesses says it uses AI extensively. Much of what is called adoption is still using AI to make existing processes a bit quicker or easier.

The problem is no longer that businesses have not heard about AI. It is whether they have the confidence and capability to work out where it is genuinely useful and then make it work. Smaller businesses will find this particularly difficult. Most do not have a chief technology officer or a team of consultants standing by.

We have been here before, I have seen these themes playing out over the last thirty years.

I remain an optimist about technology, but decades of working around the internet has cured me of believing that its benefits magically spread themselves around. They do not. Power can concentrate very quickly and people can be left behind, even as the technology races ahead.

The answer is not to slow innovation down. It is to get much better at spreading its benefits.

This was one of the starting points for the London AI and Jobs Taskforce, which I chaired for the Mayor of London. We were asked to look at how AI was already changing jobs across the capital, who might be most affected and what London could do now rather than wait for the consequences.

We tried to avoid the extremes of believing AI will solve almost everything or predicting an imminent jobs apocalypse.

AI is coming. But the way it changes work is not predetermined.

Employers will decide which tasks they automate and what happens to the productivity gains. They will decide whether they invest in their people. government will decide whether it notices what is happening quickly enough to do anything useful about it.

The Taskforce's work found that just over half of London's employment is in occupations with relatively high exposure to AI but also a strong capacity to adapt. Around 600,000 Londoners are in jobs with high exposure and much less capacity to adjust. Many are in clerical and administrative roles. Women are disproportionately represented and pay tends to be below the London average.



I do not interpret this as 600,000 jobs about to disappear. I interpret it as an opportunity to see where some of the pressure may come before it arrives.

In fact, it is the less dramatic consequences that worry me more.

What happens if companies simply stop recruiting as many people at the bottom? What happens to career ladders if many of the tasks through which young people traditionally learn disappear? How do you acquire judgement if you never get the chance to do the work that teaches it?

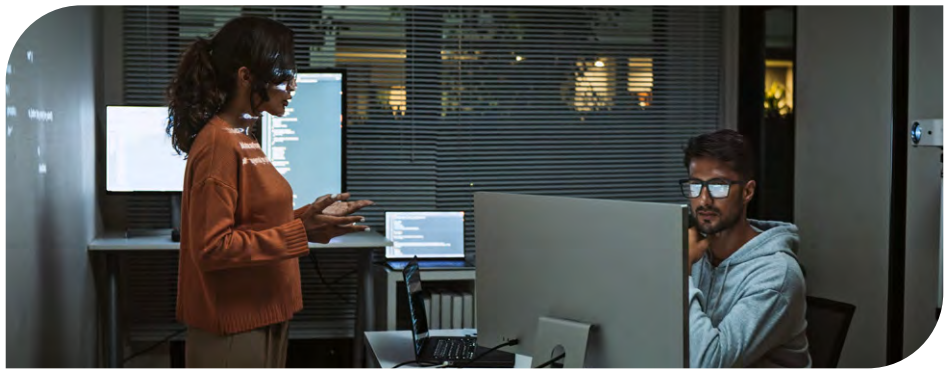
This is not only a problem for young people looking for their first job. Businesses may discover ten or fifteen years from now that they have hollowed out the generation that should by then be running things.

The debate about skills has also become too narrow. AI makes human capabilities more important, not less.

When I questioned ministers in the House of Lords about the evidence from our Taskforce, I focused on critical thinking and judgement. I dislike the phrase "soft skills". There is nothing soft about being able to work out when a machine is wrong or persuade another human being of something. Nor is there anything soft about having an original idea.

So what should government actually do?

I would start in a few places.



Make AI adoption much easier for ordinary businesses.

I would create a National AI at Work Partnership with a serious offer for SMEs.

Government should use the levers it has to encourage companies to invest in their people at the same time as they invest in AI. Businesses receiving support should involve employees in significant AI deployments and be clear about the problem they are trying to solve. We should encourage companies to redesign work rather than simply remove tasks.

For SMEs, the support needs to be extremely practical. Help them work out where AI could be useful. Give them access to independent advice and people who have actually implemented these systems.

And please do not build another website and expect millions of businesses to find it.

Use organisations they already know. Chambers of Commerce are an obvious place to start, alongside sector bodies. Government already has useful programmes in the Business Growth Service and Made Smarter. BridgeAI is another piece of the puzzle. The job now is to make these things feel like one journey rather than a collection of government initiatives.

A business owner needs to be able to work out what problem they are trying to solve and whether AI is actually the right answer. They need to understand what it will mean for their employees. And they need some way of knowing whether the investment has worked.

If we cannot help a small business answer those questions, talking about becoming an AI superpower is fairly meaningless.

Get much faster at seeing what is happening to jobs.

Our labour-market statistics were not designed for a technological shift that can change individual tasks long before it eliminates an occupation.

We need a UK AI Early Action System that combines labour-market data with what people are actually seeing inside workplaces.

I would publish a simple State of AI and Work dashboard and, crucially, agree in advance what would cause government to act. A sustained fall in entry-level vacancies should get our attention. A sudden change in recruitment in a particular sector should too.

If the evidence changes, government should be able to move training money or change careers advice. It should be able to intervene before a problem has become entrenched.

This was the thinking behind the Early Action System we proposed to the Mayor. We did not want another observatory producing interesting reports about things that had already happened. We wanted an early-warning system that might allow London to intervene while there was still time to make a difference.

Give young people a better chance of succeeding in an AI economy.

Qualifications will remain important, but they cannot be the whole answer.

I would like to see a Future Ready Guarantee built around the capabilities people will need to work well alongside AI. Judgement and critical thinking are central. So is knowing how to use AI effectively without becoming dependent on it.

But this should not become another certificate.

Young people should earn it by doing real things. They could work on problems with employers or take part in AI-enabled apprenticeships. The experience should count for something because the work itself counted for something.

Young people should also help design it. We should not spend the next decade telling them to become "AI ready" while giving them no influence over the workplaces being created around them.

AI adoption is not about buying software licences. It is about changing organisations, which means changing how people work.

Workers need to understand the tools they are being asked to use and have some influence over how they are deployed. Leaders need to take responsibility for the consequences. Smaller businesses need access to expertise that would otherwise be available only to large companies.

People are more likely to trust this technology if they can see it improving their working lives. They are more likely to trust it if they have some agency over how it is being used.

Government cannot decide which AI tool every company should buy. Nor should it try to preserve every existing job indefinitely. Its job is to make it easier for people and businesses to adapt, while making sure the gains do not flow only to those who are already best placed to seize them.

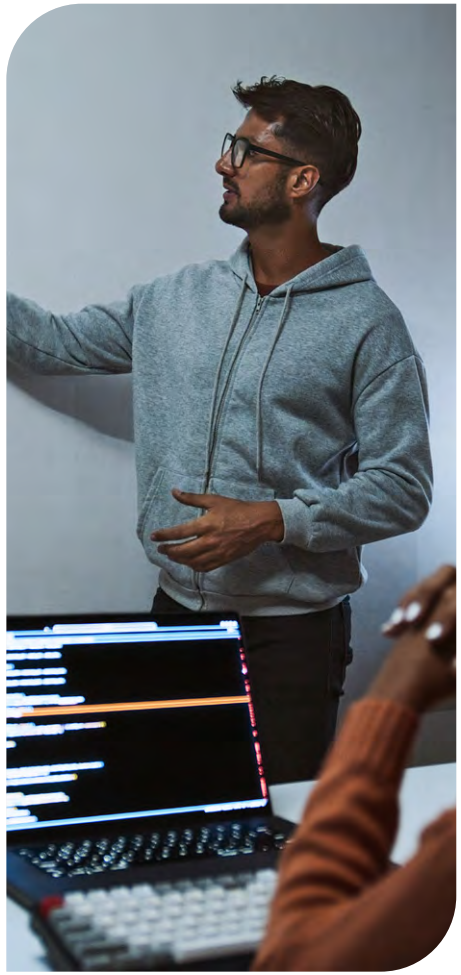
This cannot sit neatly inside one department. I have made the same argument for years about digital inclusion. Somebody at the centre of government has to care enough to keep asking whether the benefits are actually reaching people.

Britain has extraordinary advantages. We have world-class research and some brilliant entrepreneurs. Our universities are formidable. Millions of businesses are capable of adapting remarkably quickly when given the right tools.

But we need to stop confusing announcements with progress.

The countries that succeed over the next few years will not necessarily be those that build the biggest AI models. They will also be the ones that work out how to make this extraordinary technology useful to the greatest number of people.

Britain must be one of them.



22

Britain's AI problem is not adoption. It is capability



Professor Nisreen Ameen
Professor of Artificial Intelligence and Policy
University of Reading

Professor Nisreen Ameen is Professor of Artificial Intelligence and Policy at Henley Business School, University of Reading. She is an internationally recognised scholar and policy expert whose work focuses on AI skills, workforce transformation, responsible innovation, digital inclusion and the governance of emerging technologies.

Nisreen argues that Britain needs AI-ready organisations, not just AI users. That means better leadership, skills, governance and work design.

The UK is right to be ambitious about accelerating Artificial Intelligence adoption. AI can raise productivity, strengthen competitiveness, improve decision-making and create new products and services. But the number of businesses using AI is a poor proxy for whether the technology is generating sustained value.

An organisation can purchase licences, run pilots and train staff while leaving its underlying work, management practices and governance largely unchanged. Adoption is therefore only the beginning.

The more important challenge is capability: whether organisations can use AI effectively, responsibly and adaptively.

AI capability is not simply technical proficiency. It is the organisational capacity to combine technology with the skills of people, leadership, work and process design, governance, and continuous learning.

If the UK wants to become an AI adoption superpower, this broader capability challenge needs to sit at the centre of policy.

Skills must be tied to real work

The evidence from my work with Skills England shows why. The AI Skills for the UK Workforce programme found that organisations need technical, non-technical, and responsible and ethical AI skills, with different combinations required at different job levels and stages of adoption. The subsequent AI Foundation Skills for Work benchmark reinforced that effective AI use depends not only on operating tools, but also on understanding their risks and consequences, analysing outputs and exercising judgement.

More recent national evidence points to an even wider organisational challenge. The Skills for AI: What Works for AI Upskilling in the UK programme drew on 23 workshops involving around 150 organisations, 10 organisational case studies and a UK employer survey of 536 responses. AI was already deeply present in everyday work: more than 44% of organisations reported using AI tools daily, 66% were using AI features embedded within standard workplace software, and 49% were using conversational AI tools. Yet organisational maturity lagged considerably behind use. More than 40% of organisations remained concentrated at the awareness and exploration stages of adoption, while

only 1% reported reaching the point of scaling AI across the organisation.

These findings informed the PRIMES framework, which identifies six characteristics of effective AI capability development: training should be Practical, Reachable, Integrated, Modular, Expandable and Sustainable. Crucially, PRIMES moves the emphasis beyond simply providing courses. Effective capability development needs to be connected to real tasks and decisions, embedded within organisational systems and governance, accessible across workforce groups, capable of expanding across roles and teams, and continually updated as AI technologies and work practices evolve.

The ten organisational case studies reinforce this picture across different contexts. Organisations such as Airbus, KPMG and Roche were building capability through structured learning pathways, leadership sponsorship, governance and integration with workplace systems, while SMEs and consultancies relied on applied, workflow-based approaches. NHS primary care and inclusion-focused initiatives highlighted the additional importance of trust, accessibility and learning within real service environments.

Across these different models, the consistent lesson was that effective AI capability development depends not only on training provision, but on how learning is connected to work, leadership, governance and the wider organisation.

This is not simply a training deficit. It is a coordination problem across the organisation. Capability can fail when employees are trained but managers do not redesign work, when experimentation expands without governance, when teams develop isolated solutions that cannot integrate with core systems, or when successful pilots lack the sponsorship and resources required to scale. A national adoption strategy that focuses predominantly on access and skills risks overlooking these organisational bottlenecks.

This matters because there is no single national AI skills gap. Capability requirements vary by sector, occupation, organisational maturity and the type of AI being deployed. A small business using generative AI for marketing or administration faces a very different capability challenge from a financial institution embedding AI into regulated decision-making, a manufacturer using AI in safety-critical operations, or an organisation beginning to deploy

agentic AI systems that can initiate and coordinate actions across workflows. Skills England's 2026 Annual Skills Report and Sector Skills Needs Assessments reinforce this picture, highlighting growing demand for judgement, problem-solving, collaboration, digital fluency and responsible AI capability as AI reshapes tasks and workflows across priority sectors.

The implication is that skills policy should begin with changes in work. Rather than defining training requirements in isolation, organisations and policymakers need to understand how AI is changing tasks, workflows, decision-making, responsibility and human oversight. When part of a professional task is automated, the remaining human contribution may require more interpretation, assurance or relationship management. As agentic AI begins to undertake sequences of actions across organisational systems, capability will increasingly include understanding delegated authority, escalation, supervision and accountability.

From pilots to genuine transformation

This organisational perspective becomes even more important when businesses attempt to move from experimentation to scale. Evidence from Roche's global digital transformation illustrates that successful digital pilots do not automatically become enterprise-wide transformation. Examination of AI-enabled and automation initiatives across a highly regulated multinational environment showed that scaling depended on organisational capacities surrounding the technology: developing a mindset and capacity for scaling, achieving appropriate standardisation and integration, managing information flows, embedding governance and systematically assessing whether initiatives were ready to progress.

The lesson is highly relevant to the current AI adoption agenda.

Britain does not primarily need more disconnected pilots. It needs organisations capable of identifying which uses create value.

Integrating them into existing systems and processes, adapting them to regulatory and operational contexts, and scaling them without compromising quality or accountability.

Why leadership and governance matter

Leadership is central to that capability. Managers, executives and boards determine where AI investment is directed, which processes need redesign, how roles and responsibilities change, what levels of risk are acceptable and how outcomes are measured. Leaders do not need to become AI engineers, but they do need sufficient capability to connect AI investment with strategy, workforce development, operating models and governance. Without this, widespread experimentation can remain fragmented rather than becoming organisational transformation.

AI capability does not generate value independently of the organisation around it. It depends on an organisation's capacity to redesign work, coordinate people and systems, govern use appropriately, and continually adapt as technologies evolve.

Governance should therefore be treated as part of capability rather than as a separate compliance exercise. As AI becomes embedded in consequential workflows, organisations need clarity about who can use particular systems, for which purposes, who verifies outputs, when human intervention is required and where accountability sits. Responsible AI is realised through organisational roles, routines and decision rights as much as through principles.

A capability-led approach also requires better measurement. Adoption rates tell us whether businesses are using AI, but not whether productivity has improved, processes have been redesigned, innovation has increased or successful applications have scaled. Training completion indicates that learning activity occurred, but not whether capability is being applied effectively. UK policy should therefore distinguish more clearly between access, experimentation, operational adoption and organisational transformation.

The UK already has major strengths in AI research, innovation, professional services, higher education and business. The next challenge is to connect those strengths inside organisations. Technologies will continue to change, as will the skills required to use them. The more durable source of advantage will be the capacity of organisations to continually absorb, evaluate and reorganise around technological change.

If Britain wants to become an AI adoption superpower, the ambition should therefore be more demanding than maximising the number of firms using AI. It should be to build an economy of AI-capable organisations able to translate AI into sustained productivity, innovation, responsible transformation and growth.

Recommendations to UK government

1. Establish a national framework for organisational AI capability

Government should complement existing workforce-skills initiatives with a practical framework that enables organisations to assess and strengthen capability across leadership, workforce skills, work and process design, governance and continuous learning. It should be developmental rather than compliance-led, adaptable across sectors and stages of AI maturity, and provide particular support for SMEs that lack specialist internal resources.

2. Make AI-capable leadership a national priority

Government should work with business organisations, universities, professional bodies and sector organisations to strengthen the capability of managers, executives and boards to identify where AI creates value, redesign work, develop their workforce, govern risk and measure outcomes. Investment in workforce AI skills will have limited impact if those responsible for organisational transformation lack the capability to lead it.

3. Align AI adoption and skills policy with changes in work, and measure transformation rather than uptake alone

Sector-level evidence on changing tasks, workflows, decision rights and human responsibilities should inform skills interventions. At the same time, national measures of AI progress should move beyond adoption and training volumes to include productivity, innovation, successful scaling, workforce outcomes, responsible use and organisational resilience.

The next stage of UK AI policy should therefore move from accelerating adoption alone to building the organisational capability that makes adoption valuable.



How to scale with trust and resilience



23

Why trust comes before tech



Kate Hayward
UK Managing Director
Xero

Kate Hayward is Xero's UK Managing Director, serving more than 1.3m UK small businesses. A Steering Board member of the government-backed Maple Review, she champions inclusive entrepreneurship, financial confidence and digital adoption, while leading Xero's AI-powered innovation agenda to help small firms and their advisers.

Kate argues that practical local support - through Chambers, advisers and trusted platforms - is what will turn interest into confident, everyday use.

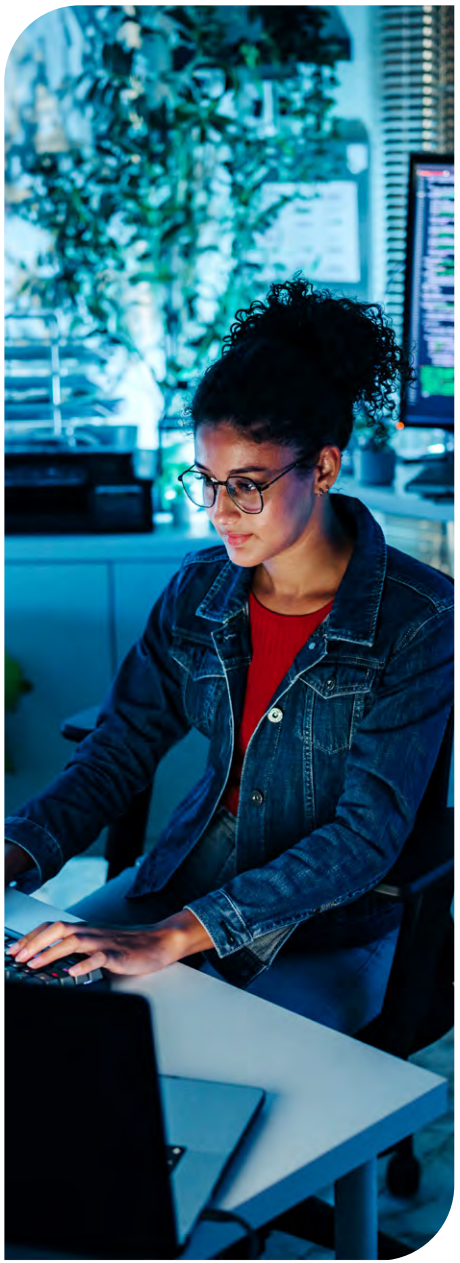
When I talk to small-business owners about Artificial Intelligence, I rarely hear a lack of interest. I hear a mixture of curiosity and caution.

Many can see how AI might help them save time, understand their finances, improve customer service or compete more effectively. But they face an overwhelming choice of tools. They want to know what a system does, whether its answer is accurate, where their data goes and who is accountable if something goes wrong.

One recurring pattern is that an owner may be comfortable using AI to draft a social media post, but hesitate when it touches their data, operating systems, invoices, bank information or tax records.

Xero supports over 1.3m UK SMEs and time and time again, their trusted accountant and bookkeeper advisers tell me their clients ask the same questions: what is happening behind the answer, what information can the system access and what do they still need to check themselves

That caution is not a failure of ambition around their AI adoption. It is a rational response to a technology developing quickly and being presented with more confidence than clarity.

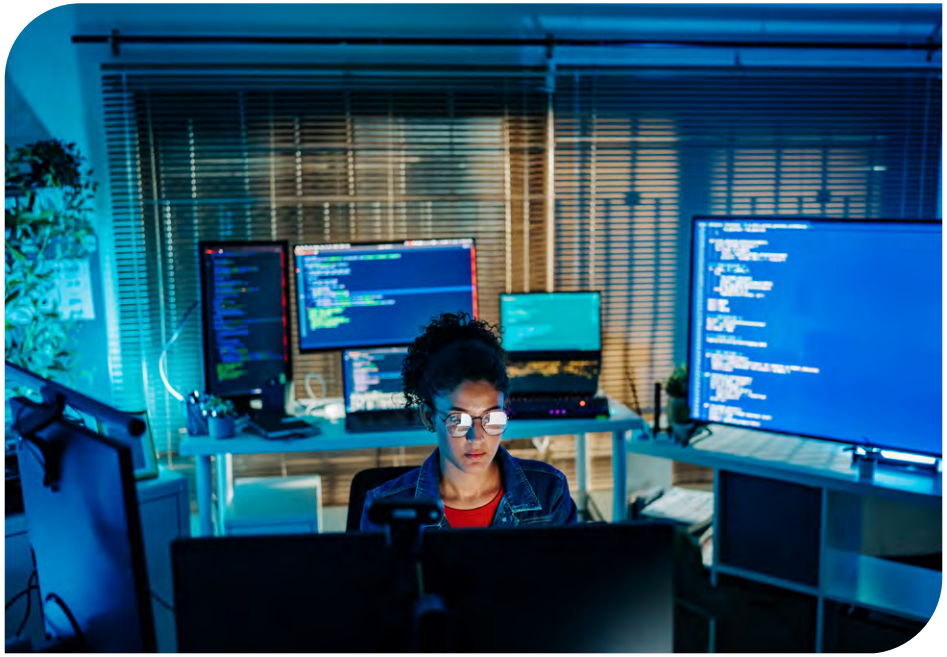


Innovation and Education

The UK's AI ambition will not be achieved through innovation alone. The biggest challenge the UK has to turbo-charge its AI adoption is entirely dependent on whether our beloved small businesses feel they can use AI safely, confidently and practically. We need to move from encouraging experimentation to enabling informed, secure and sustained adoption.

We see from our experience hosting AI workshops across the UK that small businesses do not need isolated lessons on AI, data or cybersecurity. They need joined-up, practical education that explains what AI is, how it works, when it is appropriate to use and how to implement it safely in the business.

Trust runs through all of this. Owners and their advisers need to understand what information a system uses, how that data is protected, how to recognise common cyber risks, how to check an AI-generated answer and where human judgement remains essential. No one element is enough on its own. Confidence comes when people can connect the technology, the safeguards and the day-to-day decisions they need to make.



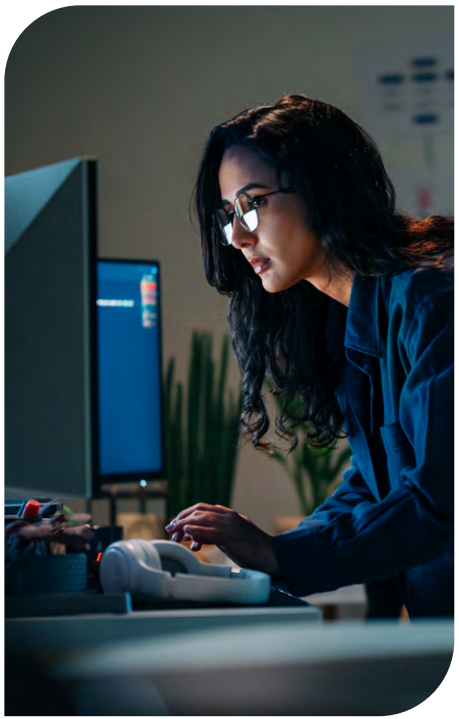
The best education will be practical. It will begin with the questions owners already ask, not with a generic demonstration. Can I reduce repetitive administration? Can I understand my cashflow? Can I get paid faster? Can I make a better-informed decision?

This is where a hyper-local approach matters. Small businesses are more likely to engage when information comes through people and organisations they recognise: their local Chamber of Commerce, a business network, an accountant or bookkeeper, a trusted software provider they already use or a local enterprise organisation.

National ambition needs local delivery.

Local support lets owners learn from peers facing similar pressures and discuss responsible use in the work they actually do. It can reach the smallest and least digitally confident businesses, making education confidence rather than another abstract policy promise.

The experience of cloud accounting offers a useful comparison. When Xero came to the UK as a disrupter 20 years' ago (as a cloud native company taking on traditional desktop products), cloud technology had to undertake a similar change management transformation. Some owners questioned security, privacy and where their information would be stored. Education, clearer safeguards and practical experience helped



change the conversation. AI is not identical to cloud technology, but the lesson is similar: new technology becomes usable at scale when people understand it, trust the safeguards and see a clear benefit.

The same principle applies to AI. Small businesses do not need to adopt AI for its own sake. They need help applying it to a real problem. In conversations with owners, the most compelling examples are straightforward: reducing financial administration, making sense of business information or following up work that would otherwise be missed.

AI and your finances

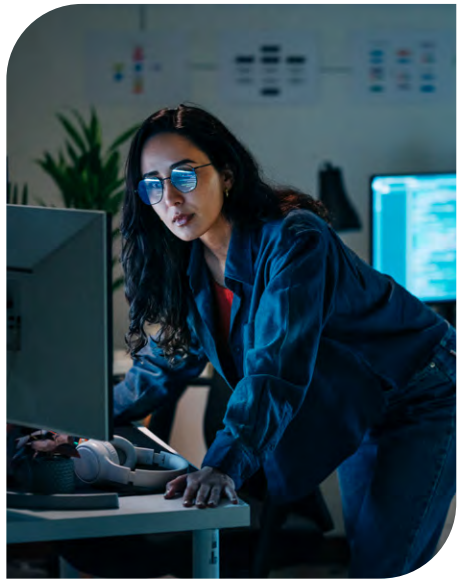
That is particularly important when AI touches financial information. Most of us were taught at school to "show your workings". The same principle should apply when AI helps with money; owners need to know what information the system used, how it reached an answer and what still needs their review. They also need to know how their data is protected and how to avoid putting sensitive material into an uncontrolled public tool.

I recently spoke to an independent cafe owner who told me that prior to using AI she found her bank reconciliation to be a drain on her time. With help from her bookkeeper, she is now using AI to speed this up. She does not need to see every line or transaction but she does need to know what was matched, why it was matched and what remains for review. That is the "show your workings test" for trustworthy AI; enough visibility for the user to check the answer and stay in control.

Trust is not only about the platform a small business uses. It also depends on knowing where responsibility sits when AI is involved in a financial decision, what the owner is expected to check and where they can turn if something goes wrong.

SMEs should not have to navigate these questions alone. They need a clear point of contact for regulatory questions, alongside accessible support to develop proportionate AI governance policies for their business. Trusted platforms can reinforce this by making data permissions, explainable outputs, human review and accountability visible in practice. The aim is not to prescribe one tool for every business, but to make safe, informed adoption easier than unsafe experimentation.

Human expertise is part of that adoption model. Accountants, bookkeepers, advisers, Chambers and local networks can help owners interpret information and understand limitations. AI can reduce routine work, but it does not remove the need for context, judgement and accountability. Education must sit alongside affordable access, relevant skills and support to move from a pilot into everyday use.



From Curiosity to Confidence

What would help most, though, is investment in something that doesn't exist yet.

Small businesses need practical, accessible support to move from curiosity to confidence.

That means funding hyper-local AI adoption clinics through Chambers of Commerce and other trusted business organisations; places owners already know and trust. These don't need to be expensive or complicated: practical workshops, plain-English guidance on AI, data and cybersecurity, and real examples of how other businesses are using AI safely across connected parts of their work.

The aim should not be to create isolated uses of AI for isolated tasks, but to help owners understand where it can support the wider running of their business, with human judgement and appropriate checks still in place. Support should continue beyond the first experiment, helping owners turn a promising use into safe, everyday practice.

A consistent national curriculum could provide shape and consistency, but the real value is in local delivery. Different sectors have different needs. A cafe owner's questions look different from a plumber's or an accountant's.

How should success look? Not just in how many people sign up, but in whether they actually use AI safely over time, whether they feel confident, and whether it genuinely helps their business. The goal is straightforward: make safe, informed adoption easier than the risky alternative of putting sensitive information into tools people don't understand or trust.

The UK will not become an AI-enabled economy simply by making more tools available. It will do so by giving people the knowledge and confidence to use them safely, and by building trusted routes from education to adoption.

Education builds trust. Trust unlocks adoption. With the right local support, trusted platforms and human expertise, AI can give owners more time, clarity and confidence to run and grow their businesses.

24

AI Adoption is an operations problem, not a technology problem



Jason Mashinchi
Founder & Managing Director
Cambridge Kinetics
Board Director
Cambridgeshire Chamber of Commerce

Jason Mashinchi is a software entrepreneur who is the founder and Managing Director of the AI adoption company Cambridge Kinetics. He founded the company in 2021 aged 24 after studying at the University of Cambridge.

Jason argues AI adoption is about changing how businesses work, not just buying tools. He explains how young people must be part of that shift, helping firms turn AI into real productivity and growth

There's a clear divide opening across UK businesses, and it isn't the one that usually gets talked about.

A lot of firms think they're adopting AI by using a chatbot to assist with creating slide decks or summarising emails. A much smaller group are going back to first principles and rebuilding how work actually gets done, fundamentally rethinking how their products and services are delivered.

The first group are buying tools, the second are reworking how they operate entirely. Yet both think they're on track with adoption.

The ONS statistics show how far we've got to go. As of June 2026, only around 35% of firms with 10+ staff use at least one AI technology, up from roughly 12% in late 2023 - but only around 10% describe their use as extensive. That gap is what will decide whether the UK becomes an adoption superpower, or just a country with a lot of licences.

Five different things called adoption

Part of the problem is that five different things sit under the umbrella of "AI adoption".

1. Access
staff having access to AI to safely experiment with

2. Repeated use
Individuals using AI regularly for discrete tasks

3. Workflow adoption
A team has an agreed, repeatable process that uses it, sometimes automated

4. Operational integration
Wiring AI into the firm's systems, data and permissions, with controls

5. Transformation
Redesigning several workflows with commercial results you can measure

From what we've seen, an enormous number of firms have moved onto stages one and two, some to three, and very few anywhere near four and five – which is where the real productivity gains live. Most adoption statistics in the country measure the first two.

Where firms get stuck

When I moved from Computer Science to a Manufacturing and Management degree, back in 2020, and first set foot on a shop floor, I was struck by how much potential British manufacturing had, and surprised by how many critical processes were glued together by paper, spreadsheets and hard work, when much of it could be automated.

That set me on a different path: instead of a big tech or consulting firm in London, I started a business helping normal SMEs sort their systems out.



Five years on, I still see quoting and invoicing done on paper, senior executives wasting afternoons in Excel to get basic KPIs, and problems database technology solved decades ago in far more businesses than we'd all like to admit.

It's clear adopting AI is an operations problem, not a technology problem.

Businesses need systems thinking – mapping processes nobody has written down, making five years of job sheets usable, connecting systems, and simplifying where possible. AI is only as good as the data and context it's given – and operational data spread across folders and disconnected systems is one of the biggest barriers to true adoption.

Taking a firm through this is difficult. It takes weeks, needs expertise and team engagement, and competes with this week's orders – which is why it slips to next quarter. And the reasons it slips are almost never technical. Change unsettles how people understand their own jobs, involves a level of risk and disruption, and crosses departmental barriers. It's not a failure of individuals, it's an awareness and change management failure, and exactly what small firms struggle to handle alone.



The help that actually works

SMEs need practical support, not just awareness and access to technology. Running an AI session for manufacturers in Great Yarmouth recently, it was clear there is also a major skills and talent gap, just as much as an awareness and adoption gap – the same pattern across the country and in every industry. BCC research this year found only 31% of firms are placing greater emphasis on AI literacy, with 45% reporting no noticeable changes to skills needs - indicating many are yet to fully appreciate how AI can impact their processes.

The strongest model we've seen (when run well) is Made Smarter, set up in 2018, to support industrial technology adoption. The combination of awareness, training, internships, and grants, provides both the skills and knowledge to succeed, and, importantly, deadlines that make things happen and stop these projects slipping.

We've been on the receiving end of this too. An Innovate UK BridgeAI grant funded work we're now deploying into hundreds of SMEs and exporting globally. However, the process around it consumed effort that should have gone into the work. Advisors and assessors spent time on administration rather than on making an impact, and different parts of the system often didn't know what the others were doing. Many businesses like ours don't have the time to navigate it.

What I'd ask for

My recommendation isn't a new programme, it's to model one on a scheme like Made Smarter that already works, and broaden it to the rest of the economy. It has to be delivered by people with real experience of change in industry, not something easily taught, and they need freedom to back the projects that will make the biggest difference, even where those don't fit neatly into a box. Four things will make it work.

Firstly, fund the systems thinking, not the licences. Match-fund the part firms struggle to do themselves – mapping processes, cleaning data, connecting systems and training the team on new ways of working.

The software is cheap and getting cheaper, the hard bit is embedding it.

£10k to £25k per firm against an agreed project is the right order of magnitude.

Secondly, don't reinvent the wheel. Made Smarter, Growth Hubs, Catapults and Innovate UK already exist, and each has relationships with firms that take time to establish. The failure mode isn't a lack of schemes, it's getting the right support to the right businesses, avoiding duplication across government, and ensuring all parts of the system talk to each other.

Thirdly, consider measurement carefully. Insist on a baseline, and publish what happens, including failures – across entire processes, not just individual steps. Without a measurable before number, gains are quickly forgotten. Measure the right things: not courses delivered, or businesses engaged, but whether a process changed, whether it's still running six months later, and whether the firm cut process costs or won more work off the back of it.

Finally, step up efforts on skills and education. A large proportion of the economy is knowledge work, and the steepest falls in entry-level hiring are in precisely the areas that AI systems are good at.

Young people should be a central part of adoption, not a casualty of it.

Every firm that takes implementation help should get a young implementer with it - an apprentice or graduate who can support with roll-outs, data cleanup and improving efficiencies. The firm keeps the capability when the adviser leaves, and the young person learns process analysis and change management rather than prompting.

One condition

In June, a US export control decision pulled two frontier models away from every non-US user for the best part of three weeks, British firms included. Any business that had built systems around that capability discovered a single point of failure and had no control over it.

The answer isn't to build a British frontier lab, it's to make sure firms aren't structurally dependent on one, and have control of their data. An onshore option good enough for everyday operational work, an expectation that publicly funded projects run on more than one provider so firms can easily switch without rebuilding, and a serious effort to make it easier to build energy capacity and data centres here. AI needs context to be useful, and context means quotes, correspondence, customer data, and sensitive business process information. That's key to retaining know-how in the UK and a critical decision for firms about where knowledge lives, and most businesses are making it without realising the cyber-security and resilience implications.

The prize isn't efficiency

Much of the AI conversation is about efficiency, delivering the same work with less effort. That's where firms start, but it isn't the point. The prize is what comes after: a twelve-person company that quotes in four hours instead of four days, wins work it used to turn away, and hires because the work is there. Efficiency frees time to make a bigger impact.

The biggest opportunity I see is for the UK to diversify away from its reliance on knowledge work and regain ground in the physical world. As intelligence gets cheaper, the bottleneck moves to the things that must be built, and that is where we still have world-class engineering and not enough automation. British Chambers of Commerce research found only 8% of British manufacturers say AI is widely embedded across their operations. That isn't a story of failure, it's millions of processes ready to improve, sitting on top of the existing engineering talent to fix them, and a route to producing and exporting more of our own goods and services. What made the last industrial revolution ours was how fast machines spread from one business to the next. We have the foundations to do that again.

25

Seizing adoption advantage for business



Dr James Kuht MBE
CEO and Co-founder
PAIR

Dr James Kuht MBE is the CEO of PAIR, a company which helps organisations build AI native workforces through its role-personalised training, working with the likes of Boeing, Darktrace, the Ministry of Justice and hi-impact. Prior to this, he had a 16-yr career in the military.

James argues the UK's AI opportunity is not just to invent new technology, but to help firms use it faster and better. The real winners will be businesses that put AI into everyday work.

From James Watt's improvements to the steam engine to Tim Berners-Lee's invention of the World Wide Web, Britain has helped change what technology makes possible. DeepMind's work using AI to predict protein structures continues that tradition.

But the economic impact of an invention depends on what people and businesses do with it.

Britain could become an AI adoption superpower by closing the gap between what AI can already do and how organisations use it. That gap represents a substantial commercial opportunity: as the technology advances, the advantage will belong to those that keep adapting.

Britain does not need to build every frontier model to capitalise on this opportunity. Each advancement creates possibilities that businesses have yet to put into use. Through reinvention, Britain could establish an advantage that extends across sectors and regions.

Closing the adoption gap

More British businesses are using AI. The British Chambers of Commerce's Powering Productivity report puts AI adoption among SMEs at 54% in 2026, up from 25% in 2024. Yet most firms using AI still rely on generic tools for tasks such as research and analysis.

Businesses recognise the opportunity. In the BCC's Growth Delivery Test, 60% of firms surveyed said greater use of AI would support their growth. Research commissioned by the Department for Science, Innovation and Technology, published in January 2026, points to benefits already being realised. Among AI users in its survey of 3,500 businesses, 75% reported improved workforce productivity and 57% reported new or improved processes or operations.

These findings show that AI is already changing how some businesses operate. The challenge is to make that progress more widespread and sustained. Doing so depends on the people who understand the work, can identify where AI would help and can turn that insight into a better process.

People as the drivers of reinvention

PAIR helps organisations turn AI deployment into measurable business outcomes. Its work with Wirral Chamber of Commerce shows the positive impact this can have on local businesses.

The Chamber's AI Academy, delivered by PAIR, began with 42 local businesses and helped provide the model for today's national BCC AI Academy. Participants reported average time savings of more than three hours a week, 96% said the quality of their work had improved, and the programme generated an estimated £500,000 in productivity gains in the first year.

Pharmapac reported reducing time spent on investigations and standard operating procedures from three days to three hours. Heatons Group built an AI assistant which reduced first-line support response times from three days to one. Today, more than 325 businesses are enrolled in the BCC AI Academy.

These examples show what reinvention means in everyday work. People, in all areas of business, have taken a process they understand and found a better way to carry it out. Across an economy, changes like these can release substantial time for serving customers, improving products and growing businesses. Bottom-up innovation thrives with leadership support. Leaders set direction, decide which risks are acceptable, scale what works and choose how to use the time and budget released. Employees' knowledge and leadership direction work together to drive organisational change.

Mimecast, where PAIR has supported workforce capability building, illustrates this in a larger organisation. The Financial Times reported that 96% of its 2,400 employees had incorporated AI into their daily workflows. Its account highlighted leadership involvement and changes in working practices and company culture, including a chief executive who built and shared his AI agents to encourage discussion.

Mimecast has also reported a record year for profitability and growth.

Continuous reinvention

An organisation can train its workforce on today's tools and still fall behind as AI evolves. The advantage comes from being able to absorb new capabilities continuously, leveraging four disciplines: adoption, adaptation, assurance and allocation.

Adoption equips people to use relevant tools; adaptation reshapes workflows around their capabilities; assurance checks that the resulting work is safe and effective; and allocation directs budget, time and people towards the applications that produce the greatest value.

These decisions are connected. A team might discover a promising application, test the quality of the work and demonstrate that it improves a process. Leaders can then decide how widely to use it and how to reinvest the capacity released. As the technology evolves, those decisions will change too.

One-off AI literacy programmes will not be enough.

Organisations can sustain progress through employees who keep finding applications, managers who recognise which ones matter, and ways to share successful approaches across teams. Together, these create a continuous flow of reinvention as AI changes what is possible.



A British model for adoption

Britain has useful precedents for building this capability. The government Digital Service's design principles emphasise user needs, testing and repeated improvement. The What Works Network helps public bodies make decisions using evidence about which interventions succeed.

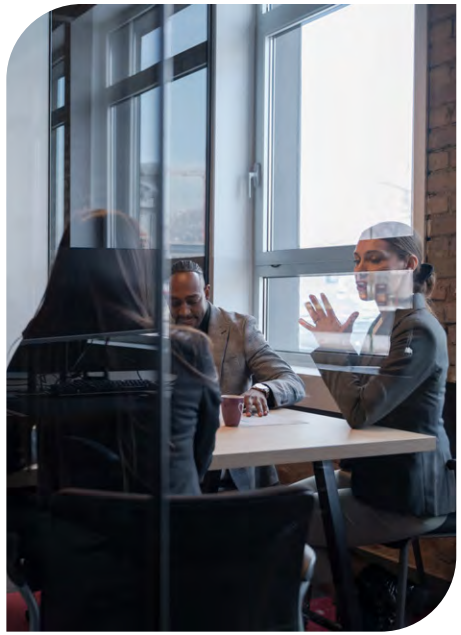
These principles could inform a national approach to AI adoption. One possible policy architecture could be to connect the BCC's existing recommendations for SME AI Champions, skills support, adoption incentives and clear assurance into a national adoption system, delivered through trusted institutions such as Chambers. Businesses would receive support to transform their operations, while evidence from their experience would inform how that support develops.

For a participating business, the starting point would be its own workflows: where could AI improve how work gets done? Employees and leaders could test changes to those processes and assess the effects on time, quality, revenue, employee experience and job design.

Anonymised findings could show which approaches repeatedly succeed in particular roles and sectors, providing practical guides for other companies. The AI Economics Institute could help build this evidence base, while the Chamber Network could help businesses leverage the findings locally.

Government could then judge support by what it achieves. Identify which interventions change behaviour, which incentives lead SMEs to invest, and which applications release capacity without reducing quality. As well as examine what businesses do with it.

There are practical challenges. Measurement must not become another administrative burden for SMEs, and commercially sensitive information needs protection. An approach that succeeds in one company may need adapting in another. The system would have to learn from these difficulties as well as its successes.





Make AI work for humans

Policy is already moving towards workforce capability. AI Skills Boost aims to equip 10 million UK workers with AI skills by 2030. The government's 2026 AI Adoption Plans emphasise skills, trust and organisational readiness, while its Pro-Worker AI Adoption Prize recognises the importance of giving workers new capabilities and a voice in how AI is deployed.

Keeping humans at the centre serves both business performance and the people doing the work.

Giving them the skills and agency to shape how those processes change allows organisations to draw on that knowledge, while retaining human judgement where decisions have real consequences.

An hour saved does not automatically become an hour of economic value. Leaders still have to put that capacity to work: serving more customers, improving a product, developing an idea or growing the business. Those choices determine whether the potential of AI becomes a meaningful gain.

Britain's success as an AI adoption superpower will be visible when a manufacturer in Coventry, an accountant in Leeds, a civil servant in Cardiff and an exporter in Belfast can turn new AI capabilities into better work faster than their overseas competitors and do it again when the technology changes six months later.

26

The adoption challenge we can still win



Robyn Scott
CEO and Co-founder
Apolitical

Robyn Scott is an entrepreneur and author. She is co-founder and CEO of Apolitical, a global online network and learning platform that connects and trains hundreds of thousands of public servants across more than 160 countries.

Robyn argues that government should set the standard by using AI to augment people, redesign workflows and turn adoption into a national advantage.

Across businesses and governments there is a paradox in AI adoption. Surface-level adoption is rising as LLMs become more widely available, but examples of transformation commensurate with AI's potential - and the strategic advantage that comes with it - remain vanishingly few.

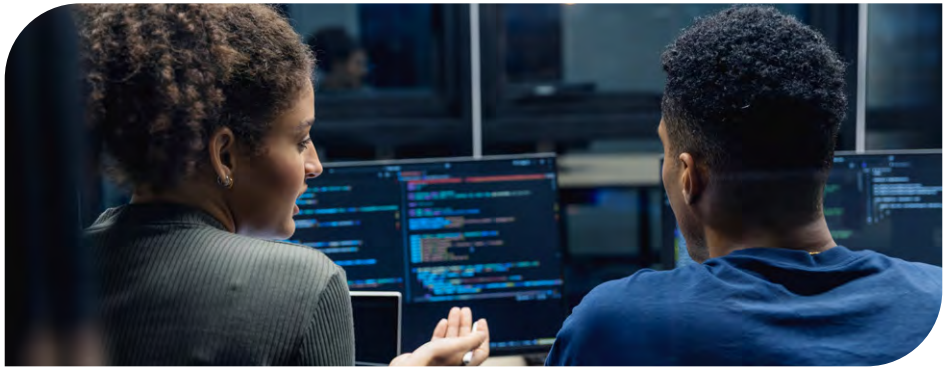
The overwhelming majority of AI pilots never reach scale, with one 2025 MIT study putting the figure as high as 95%.

Absent an investment in the hundreds of billions, the UK is out of the frontier models race. But the adoption challenge - across business and governments - is global. And it is a race in which Britain can still compete, especially by harnessing its world-class research talent. Adoption may not make Britain a superpower, but it is a superpower still available to Britain:

a country can import intelligence - albeit with dependency risks - but it cannot import adoption.

In our experience at Apolitical - supporting governments, and more recently companies, with AI transformation - two fundamental mistakes explain much of the gap between adoption and transformation. The first is how AI is sold inside organisations; the second is how it is deployed.





In governments and businesses alike, leaders have often presented AI more as Ozempic, less as caffeine. As a consequence, employees feel like turkeys setting the Christmas dinner table, and resist quietly and effectively, or work with heightened uncertainty and diminished morale. The slimming versus stimulant framing has also meant leaders have failed to capitalise on two opportunities. The first is the instinctive willingness of employees to use AI, albeit on their terms: around 70% of employees already use unsanctioned AI at work, untrained and ungoverned (Microsoft, 2025).

The second is that the firms using AI most successfully are growing revenues while continuing to hire: industries most exposed to AI have seen roughly three times higher growth in revenue per employee, with job numbers rising even in the most automatable roles (PwC, 2025 Global AI Jobs Barometer). Deployment of AI done well can often mean a redeployment rather than reduction in human talent.

The second mistake relates to deployment. AI is often treated as an IT rollout - licences bought, training module assigned. But AI is not a system to install; it is an organisational, technological and human transformation that ultimately plays out at the level of individual tasks and workflows, and requires rethinking how people work alongside each other and technology.

Rollout thinking produces exactly what the British Chambers of Commerce observes amongst its members: surface adoption without strategic advantage.

Nowhere is the value of correcting this more vivid than in government. In addition to enabling the stretched state to do more with less, how the government adopts AI influences the national pattern, and how it buys AI shapes what gets built: government procurement is industrial policy whether or not it is named as such.

If the UK government becomes an exemplar of AI adoption - framing AI as augmentation first, with deployment and training informed by task-level analysis, and sharing emerging practice openly rather than siloed pilots - it does three things at once: proves the augmentation case at national scale, creates early demand for UK AI firms, and inspires adoption among the companies watching from the sidelines - if government can do it, any company can.

My confidence that this is possible in the UK comes from a decade running Apolitical, a UK-headquartered technology scale-up building AI and digital capability within governments. Apolitical's online community of government officials reaches more than 500,000 in 160 countries. Our tools and training reach 40 million more through platform integrations with government workforce infrastructure from India to Brazil. I also serve on the board of Invisible Technologies, one of the leading enterprise adoption companies in the US.

The UK government is not currently a leader in AI in government. States like Ukraine and the UAE - which are focused on building "agentic states" - have rushed ahead. But 15 years ago, Britain led the world in digital government; some talent and institutional memory remains. And the Burnham government's priorities could create the right conditions for state AI adoption.

Particularly relevant are the goal of connectedness across all levels of government and the push to reform the civil service led by Dame Antonia Romeo while simultaneously increasing respect for public officials.

The technical aspects of the transformation challenge are also surmountable. AI often does not replace jobs, or even roles; it changes tasks. A policy analyst's job is two dozen distinct tasks - half a dozen transformed by AI, a dozen untouched, six reshaped.

Apolitical is amongst a few organisations supporting transformation at this fine grain. Built on a unique dataset on the public sector workforce, our platform Futura starts with the individual, mapping how work is actually done, task by task and role by role - bottom-up evidence rather than top-down anecdote. This creates an MRI for the individual and the workforce, which enables personalised training for the individual worker - keeping them above the algorithm, in Silicon Valley parlance. For the leader, it creates a flexible blueprint for planning and tracking adoption and transformation.

Finished training modules tell you little; changed workflows tell you everything.

We've found there's also an anxiety dividend: when people see their own job task by task, AI stops being an abstract threat to employment and becomes a concrete trade - less drudgery, more judgement. Task-level visibility converts fear into appetite and agency.

There is a particular gap the UK government can close that will support adoption. Frameworks for AI exist at the level of principle and policy (AI acts, safety institutes, white papers, international declarations) but there's little that tells a permanent secretary or a CEO what good adoption looks like in practice. This shows in our own data: globally, only 40% of public servants implementing AI say they understand their own government's ethical frameworks.

The government should commission a practical UK adoption standard - a "what good looks like" framework at the task and workflow level, co-developed with business, and designed as a living framework: versioned and refined as the evidence grows, not carved in stone. That standard could contain how to map work before buying; how to buy well; what to measure (adoption and workflow change, not licences and completion); when a pilot is working and when to kill it; how to sequence training and re-training around actual roles and new organisational structures.

The fair challenge is whether anyone would follow an adoption standard set by a country that is not yet an adoption leader. Britain's own history is instructive. The government Digital Service did not write guidance first and act later: it built the alpha of gov.uk, wrote the service standard and ran a programme of 25 exemplar services simultaneously - each lending the others credibility, and the standard giving government the authority to say no to spending that did not meet it. The AI equivalent is to develop the standard and government's own adoption together, the former refined in real time by what the latter learns.

Nor does Britain need to learn everything first-hand. Its standards tradition - accounting, law, kitemarks - its institutional reputation and its connectedness equip it to convene and codify the best of what is emerging worldwide, from Ukraine's agentic state to Singapore's civil service to the enterprises furthest along the adoption curve. The first country to turn the world's scattered lessons into something actionable and agile will set the terms others follow - turning "adoption superpower" from domestic ambition into international soft power.

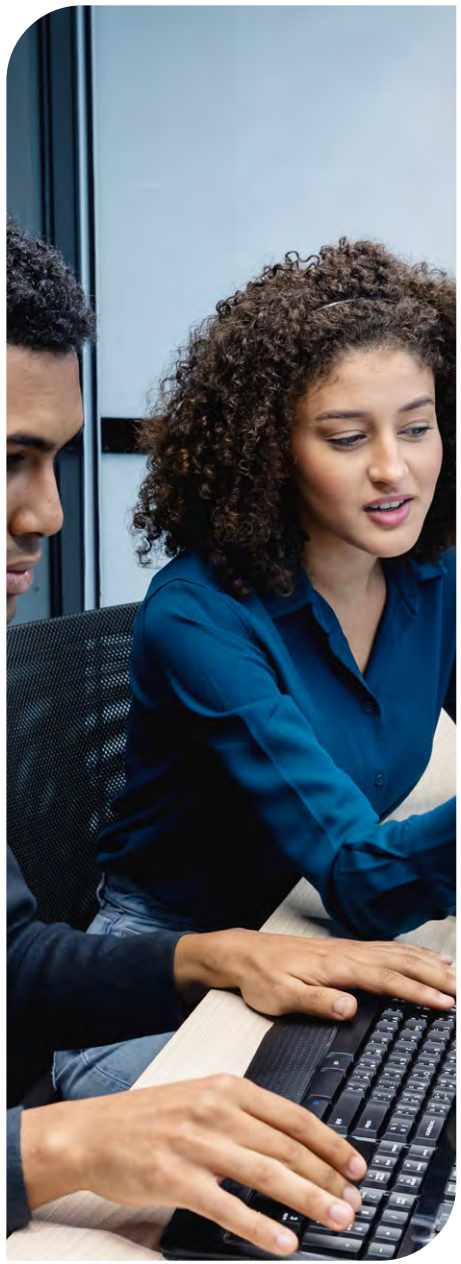
Britain is well placed for this convening role. It has the technical depth and talent to advise on the technology today, and where it is heading. It is home to the AI Security Institute, perhaps the world's most consequential body of its kind, enabling safety and security to be hardwired into the standard. And it has both the ambition to better harness AI in government and the urgent need to do more with less.

Such a standard also offers the government unusual leverage. The spending on AI is coming regardless, across government and the wider economy; a good standard is what stops it being wasted on pilots that never scale. Compared with compute subsidies and skills programmes it is cheap, fast and visible - and, executed well, it moves behaviour across the whole economy.

As the new administration forms its AI approach, its first year will determine whether UK adoption is fear-led or ambition-led; whether systems are built on guesswork, with risk baked in, or built on evidence, and secure.

We may need a slimmer state in places. But even more, we need a smarter state.

And a smarter state - using AI surgically and safely - can fuel a smarter economy



27

AI governance: The foundation of trust, growth and competitiveness



Helen Wylde-Archibald
CEO of Devon Chamber of Commerce
Co-Founder of South West AI Solutions Limited

Helen Wylde-Archibald is Chief Executive of Devon Chamber of Commerce, and has a wealth of leadership experience including in digital innovation. She previously served as CEO of Cornish broadband provider Wildanet Ltd, where she led the company through rapid scale-up, and is currently Co-Founder of South West AI Solutions Ltd.

Helen argues that AI governance is becoming a boardroom priority as firms face rising risks around data, trust and accountability. Done well, it helps businesses adopt AI safely, innovate with confidence and build a competitive advantage.

Artificial Intelligence is no longer a future technology. It is rapidly becoming embedded in everyday business operations, from customer service and marketing through to recruitment, finance, logistics, product development and strategic decision-making.

Across the UK economy, businesses are adopting generative AI tools, predictive analytics, automation platforms and intelligent assistants at unprecedented speed. The question is no longer whether organisations should adopt AI, but whether they can do so safely, responsibly and effectively.

This is where AI governance becomes critical.

Whilst many organisations have developed an AI strategy, very few have developed an AI governance strategy. As AI adoption accelerates, a governance gap is emerging between the speed of technological change and the ability of boards and leadership teams to oversee it responsibly. There is a consistently identifiable and growing "AI governance gap", where unmanaged risk, unclear accountability and insufficient oversight can undermine the benefits of AI adoption, and leave businesses open to litigation, reputational risk and fraud.

For the British Chambers of Commerce network, which represents businesses of every size and sector, AI governance should be viewed not as a compliance burden, but as a fundamental enabler of sustainable growth, innovation and trust. We predict within 12 months businesses will be required in insurance, procurement and contracting to demonstrate governance oversight in the same way as we have seen Health & Safety and GDPR develop in this space.

AI Is Becoming a Board-Level Responsibility

Historically, technology governance was often delegated to IT departments. AI changes that position fundamentally. This is a general-purpose change impacting culture, behaviour, structure and brand. Businesses will move gradually from role-based work to skill-and task-based working in a more fluid and dynamic environment.

Modern AI systems increasingly influence decisions that affect customers, employees, suppliers and communities. They may generate content, recommend actions, prioritise opportunities, assess risk, support recruitment decisions or provide business intelligence. In many cases, executives using AI may not fully understand how outputs are generated, where information originates or what risks exist behind the recommendations being presented.

This creates a new challenge for leadership teams. Boards remain accountable for organisational

decisions, even where AI has influenced those decisions.

AI governance is not simply a technical issue but a leadership imperative requiring structured oversight, risk management, board reporting and executive accountability

The UK government's AI governance guidance similarly emphasises that organisations should establish governance boards or equivalent oversight structures to provide strategic direction, accountability and risk management for AI adoption. External guidance notes that AI governance boards help ensure legal, ethical, operational and strategic risks are effectively managed while maintaining alignment with organisational objectives.

The reality is simple: if boards are responsible for cybersecurity, financial controls and health and safety, they must also become responsible for AI governance.

Governance Enables Innovation Rather Than Restricting It

A common misconception is that governance slows innovation.

The opposite is true.

Businesses that adopt AI without appropriate controls often find themselves facing uncertainty around data protection, intellectual property, confidentiality, quality assurance and operational risk. Employees may independently adopt AI tools without approval. Multiple platforms may emerge across departments. Sensitive information may be entered into public systems. Outputs may be trusted without validation.

Such situations frequently cause organisations to delay or limit AI adoption because they lack confidence in how AI is being used.



Governance provides confidence

Governance frameworks should be based upon a simple principle: governance must lead innovation and design must be human-centric. AI should be embedded within existing organisational structures, risk processes and business objectives rather than being "bolted on" as an isolated technology project.

When organisations establish clear policies, risk assessments, approved tool registers, training programmes and oversight mechanisms, employees gain clarity about how AI can be used safely and effectively. This accelerates adoption because staff understand the rules of engagement.

Good governance therefore becomes an enabler of innovation, not a barrier to it.

AI Governance Is Rapidly Becoming the New Cyber Security

Twenty years ago, many organisations viewed cybersecurity as an IT issue.

Today, no board would argue that cybersecurity is optional.

AI governance is following the same trajectory.

The risks associated with unmanaged AI use are becoming increasingly significant. Organisations may possess AI tools but still lack visibility of where AI is being used, who is accountable, what risks exist and how decisions are monitored.

The comparison is important.

Cybersecurity governance developed because organisations recognised that unmanaged digital risk could damage operations, reputation and public trust. AI presents similar challenges:

- Hallucinated or inaccurate outputs
- Bias and discrimination risks
- Data privacy concerns
- Regulatory non-compliance
- Intellectual property exposure
- Reputational damage
- Lack of transparency and explainability
- Over-reliance on automated decision-making

These are no longer theoretical issues. They are practical business risks that require structured management.

Businesses that establish governance now will be significantly better positioned than those who delay.

International Standards Are Emerging

One reason AI governance matters is that global standards and regulatory frameworks are beginning to mature.

ISO/IEC 42001, is the world's first international AI management system standard, was released in 2023. The standard provides organisations with a structured framework covering

leadership, planning, operational controls, training, monitoring and continual improvement across the AI lifecycle.

ISO 42001 represents an important milestone because it offers businesses a practical roadmap rather than abstract principles. Organisations can begin treating AI governance with the same rigour applied to quality management, information security and health and safety systems.

Alongside these standards, international regulation continues to evolve. Governments and regulators increasingly expect organisations to demonstrate accountability, transparency and risk management in relation to their use of AI. The UK Parliament's Science, Innovation and Technology Committee has identified multiple governance challenges requiring policy attention, while government guidance continues to emphasise responsible AI adoption frameworks.

Businesses should view governance readiness as preparation for an increasingly regulated future rather than a reaction to compliance requirements.

Trust Will Become a Competitive Advantage

AI adoption ultimately depends on trust.

Customers need confidence that organisations use data appropriately. Employees need confidence that AI supports rather than threatens their roles. Investors require assurance that AI-related risks are being managed. Regulators expect transparency and accountability.

Trust cannot be achieved through technology alone. It is created through governance.

Effective AI governance establishes clear ownership, documented controls, human oversight mechanisms, auditability and transparency. South West AI Solutions governance frameworks consistently emphasise roles and responsibilities, risk assessments, training, authorisation processes and ongoing monitoring as essential components of trustworthy AI adoption.

In practice, organisations that can demonstrate responsible governance will increasingly differentiate themselves in the marketplace.

Just as environmental, social and governance (ESG) considerations have become indicators of corporate maturity, AI governance is likely to emerge as a marker of organisational credibility.

The businesses that win in the AI era may not necessarily be those deploying the most advanced technologies. They will be the organisations that deploy them most responsibly.

What Should government Do?

Whilst businesses have primary responsibility for governing their own AI adoption, government also has a critical role to play in creating the conditions for trusted and responsible AI innovation.

The United Kingdom has rightly positioned itself as a global leader in AI development and adoption. However, sustainable growth will require a balanced approach that supports innovation whilst giving businesses the confidence and clarity to invest.

For many organisations, particularly SMEs, the challenge is not a lack of ambition. It is a lack of certainty. Business leaders are navigating a rapidly evolving landscape of technologies, standards, regulations, cyber risks, and workforce implications.

Many recognise the opportunity AI presents but are uncertain where good practice begins and regulatory expectations end.



There are three areas where government can make a significant difference.

First, the UK should continue to support and promote internationally recognised standards as the foundation of responsible AI adoption. International standards such as ISO/IEC 42001 provide practical frameworks that businesses can use today, regardless of size or sector. Rather than creating fragmented or competing approaches, government should encourage alignment with globally recognised standards that help organisations manage risk, demonstrate accountability, and build trust. This also strengthens international competitiveness by ensuring UK businesses can operate confidently across global markets.

Second, government should provide greater practical support for SME adoption and governance readiness. Large organisations may have dedicated legal, compliance, and technology teams. Most SMEs do not. Guidance, leadership training, sector-specific toolkits, and accessible governance frameworks would help smaller businesses adopt AI safely without creating unnecessary regulatory burdens. AI governance should not become a capability available only to the largest organisations.

Third, policymakers should continue implementing the recommendations emerging from parliamentary and expert reviews into AI governance. Recent work by Parliament's Science, Innovation and Technology Committee has highlighted issues including accountability, transparency, assurance, skills, and public trust. Businesses would benefit from a clear and stable policy direction that supports innovation whilst providing certainty around future expectations.

Ultimately, the UK's competitive advantage will not be determined solely by how quickly organisations adopt AI. It will be determined by how effectively the country builds confidence in its use.

Government has an important role in helping create the frameworks, standards and conditions that allow responsible innovation to flourish.

A Call to Action for UK Businesses

The United Kingdom has ambitious plans to remain a global leader in AI innovation and adoption. Government strategy continues to emphasise AI's potential to improve productivity, strengthen resilience and drive economic growth.

However, national success will depend on responsible implementation at organisational level.

For most businesses, the starting point is not complex regulation or expensive technology. It is governance.

Businesses should ask five simple questions:

1. Do we know where AI is currently being used?
2. Do we have an AI policy and governance framework?
3. Have we assessed the risks associated with our AI tools?
4. Is there clear leadership accountability for AI?
5. Are we preparing for future standards and regulation?

If the answer to any of these questions is "no", governance should become an immediate priority.

AI offers one of the greatest opportunities for economic transformation since the arrival of the internet. It has the potential to improve productivity, support innovation, address workforce challenges and create entirely new business models.

But realising these opportunities requires more than technology adoption.

It requires leadership, accountability and trust.

In short, it requires governance.

The challenge for UK businesses is therefore not whether they will use AI. Most already are. The challenge is whether they will govern it well enough to unlock its benefits safely, responsibly and sustainably for the long term.

And that is why AI governance matters.

28

Why assurance, auditing and governance must be central to AI leadership



Dr Allison Gardner MP

Member of Parliament for Stoke-on-Trent South and Chair of the APPG on Artificial Intelligence

Dr Allison Gardner is the Member of Parliament for Stoke-on-Trent South, with a background in molecular biology, data science and artificial intelligence. Before entering Parliament, she worked with the NHS as a Senior Scientific Adviser for AI, served as a Director of Women Leading in AI, and contributed to national and international AI standards and governance.

Allison argues that if the UK wants to be an AI superpower, it must also become a leader in AI regulation. Alison says that strong regulation should not be seen as a brake on innovation, but as the foundation of public trust.

Artificial intelligence (AI) is rapidly reshaping economies, public services and everyday life. While innovation and competitiveness are essential, they must be matched by safety, accountability and public trust. Regulation should therefore be seen not as a barrier to progress, but as an enabler of responsible AI adoption.

Recent loss-of-control incidents have highlighted the risks of deploying increasingly powerful AI systems without adequate oversight.

To become an AI superpower, we must also become a regulatory superpower.

AI governance should span the entire lifecycle of AI systems, from development to decommissioning, with particular emphasis on post-deployment assurance and independent auditing. Regulation must extend beyond frontier models to the machine-learning systems already influencing millions of daily decisions.

Ultimately, AI leadership should be measured not only by technological capability, but by the ability to ensure AI remains safe, transparent and trustworthy in practice.

The Regulatory Landscape

Recent developments suggest growing recognition of the importance of AI assurance in the UK. The establishment of the AI Assurance Consortium, chaired by figures including Adam Leon Smith and Dame Wendy Hall, marks an important step towards strengthening the UK's assurance ecosystem.

This reflects a broader shift in policy thinking. While AI governance has long focused on innovation and competitiveness, policymakers increasingly recognise that widespread adoption depends on public trust. Trust requires more than technical progress; it requires evidence that AI systems are safe, ethical and reliable.

Support for international cooperation on frontier AI governance is also encouraging. As advanced AI systems operate across borders, effective oversight will require common standards and coordinated regulation.

Despite this progress, significant governance gaps remain. Institutions such as the AI Security Institute provide valuable pre-deployment testing, but risks do not end at deployment.

AI systems continue to evolve within dynamic environments, making continuous monitoring, independent auditing and lifecycle governance essential.

The Importance of Post-Deployment Assurance

One of the greatest weaknesses in current AI governance discussions is the assumption that risks can be managed primarily through development-stage testing. In reality, AI systems are not static products. They evolve continuously through software updates, model retraining, configuration changes, security patches and system integrations. At the same time, the environments in which they operate also change. User behaviour adapts, organisational processes evolve, new datasets become available and regulatory expectations shift.

As a result, assessments conducted during procurement or development provide only a snapshot of performance under specific conditions. A system that performs safely during testing may behave differently months or years later when exposed to new operational contexts. This creates a compelling case for mandatory post-deployment assurance and monitoring.

Before entering Parliament, I worked as a Senior Scientific Adviser in the NHS AI and Digital Regulation Service, where I helped develop AI governance guidance across a range of healthcare regulators and bodies. This experience highlighted that AI safety cannot be assured at the point of deployment alone. Continuous monitoring and assurance are essential because AI performance can evolve over time, and many risks only become apparent through real-world use.

Continuous monitoring allows organisations to identify performance degradation, detect unintended outcomes, assess bias and evaluate whether systems remain aligned with organisational objectives. Without such mechanisms, regulators and deployers may discover problems only after significant harm has already occurred.

In addition, stronger regulation of AI deployers can also harness market forces. Organisations are unlikely to adopt systems they cannot effectively govern, creating pressure across the supply chain for developers to build more transparent, explainable and manageable AI. In this way, regulation can steer innovation towards safer and more accountable outcomes rather than constrain it.



Regulating the Full AI Lifecycle

To be effective, AI regulation must address the entire lifecycle of an AI system.

Risks emerge at every stage, including data collection, model development, testing, deployment, operation and retirement. Focusing on only one stage risks leaving significant vulnerabilities unaddressed.

For example, poor-quality training data can introduce bias long before deployment occurs. Inadequate security controls during development may create vulnerabilities that are later exploited. Weak monitoring arrangements may allow harmful outcomes to persist unnoticed. Even system decommissioning can create risks if data, models or decision records are improperly managed.

A lifecycle-based approach acknowledges that governance is an ongoing process rather than a one-time certification exercise. It ensures that data quality, transparency, accountability, fairness and security are considered throughout an AI system's existence. It also clarifies responsibilities among developers, deployers and users, each of whom plays a distinct role in managing risk.

Importantly, effective lifecycle governance should include mechanisms for intervention. Regulators must possess the authority to refuse approval where appropriate and to order the suspension, withdrawal or modification of unsafe systems. Governance without meaningful enforcement powers risks becoming little more than advisory guidance. Strong oversight requires both preventive measures and the capacity to act decisively when risks materialise. Furthermore, decisions on whether this regulation is horizontal or sector-specific must be taken with the understanding that ALL AI systems carry risk.

The Critical Role of Independent Auditors

Independent auditing is essential to credible AI governance. Public trust depends on auditors being free from financial or operational influence from the organisations they assess.

Independent audits can provide objective evaluations of compliance, performance, fairness, security and risk management across the AI lifecycle. They also help verify that governance obligations are being met in practice.

As in finance, healthcare and safety-critical engineering, strong auditor independence improves transparency, accountability and public confidence. With AI playing an increasingly important role in decision-making, independent scrutiny is becoming indispensable.

The Economic Opportunity

The case for regulation is not solely about managing risk. It is also an economic opportunity. A thriving AI assurance industry could become one of the UK's most significant growth sectors, generating substantial economic value while reinforcing public trust.

The government's Economic Assessment of the AI Assurance Market estimated that the UK assurance sector was already worth approximately £1 billion in 2024. With supportive policy, continued investment and widespread AI adoption, projections suggest significant future growth. These figures demonstrate that assurance should not be viewed merely as a compliance burden. It is a high-value professional services industry in its own right.

The UK possesses several competitive advantages in this area. It is home to the world's third-largest AI market and has internationally recognised strengths in standards development, auditing, certification and professional services. These capabilities position the country well to become a global leader in AI assurance.

Public expectations also support such an approach. Survey evidence consistently indicates that citizens expect those who develop, deploy and assess AI systems to be held to professional standards comparable to those applied to doctors, lawyers and accountants. Responsible regulation therefore supports both public confidence and economic growth.

Looking ahead, policymakers should also consider the strategic value of smaller language models (SLMs). While much attention is focused on increasingly large foundation models, SLMs offer benefits including lower computational costs, improved deployment and greater control in specific applications. Effective regulation should encourage a diverse ecosystem that supports innovation while maintaining appropriate safeguards.

AI, Education and Human Capability

Finally, AI governance must address not only regulation but also wider societal impacts, particularly education and workforce development. AI can enhance learning through personalised support, accessibility and improved resources, but it may also affect the development of critical thinking, creativity, communication and problem-solving skills.

Concerns are growing that overreliance on AI could lead to deskilling and cognitive dependency, especially among young people and professionals who need to retain independent judgement. The key challenge is to promote both AI literacy and human capability, ensuring people can use AI effectively while still thinking, analysing and making decisions independently. A successful AI strategy should strengthen, not replace, human potential.



Conclusion

Artificial intelligence offers significant opportunities for economic growth, scientific advancement and improved public services. Realising these benefits, however, requires strong governance, assurance and accountability.

AI regulation should be seen not as a barrier to innovation, but as the foundation for its safe and trusted adoption.

An effective regulatory framework must cover the full AI lifecycle, include robust post-deployment oversight, require independent auditing and equip regulators with meaningful enforcement powers. It should apply not only to advanced AI systems but also to the machine-learning technologies already influencing everyday decisions. Well-designed regulation can also support the growth of a leading UK AI assurance sector, creating economic value while strengthening public trust.

Ultimately, AI leadership should be measured not only by technological capability or investment, but by the ability to combine innovation with strong governance, professional standards and public confidence. By putting assurance, accountability and responsible regulation at the centre of AI policy, the UK can position itself as a global leader in trustworthy AI.

29

What stands in the way becomes the way



Ollie Whitehouse
Chief Technology Officer
National Cyber Security Centre

Ollie's career in cyber security has spanned more than 27 years prior to becoming the NCSC's Chief Technology Officer in late 2023. He has held senior science and technology advisory roles to the UK government for the last decade.

Ollie argues the UK's cyber strength gives it a chance to lead in safe AI adoption — but investment in cyber fundamentals is now essential.

AI has tremendous potential and will undoubtedly transform both the process and business of cyber security, securing the vast breadth of technology on which we rely.

However, it faces major adoption challenges, not least AI's tendency to cause harm through unintended behaviours. These recent incidents of frontier AI models carrying out unsanctioned actions and, in some cases, 'human-like' deceptive behaviour on the open internet are a serious reminder of the risks AI capabilities pose.

In Meditations, written nearly 2,000 years ago, the Roman Emperor Marcus Aurelius stoically wrote: "That which is an obstacle on the road helps us on this road". In other words, progress is achieved not by avoiding obstacles, but by working through them. This concept has been reused to note that technological progress is defined by the challenges that impede it. In this sense, there is no better application of this idea than in the context of AI.

This is something the NCSC has discussed with those working at Frontier AI companies. The critical next step on the road to Artificial General Intelligence (AGI) requires successful AI risk management. In other words, the AI acting with better 'common sense'. This could be achieved either by the AI itself, or by the frameworks within which the model sits – noting that, despite appearances, the AI that we are all now using is not actually a single technology.

To put this in other terms, the path to the next generation of AI is achieved by overcoming the safety and security issues presented by AI today.

What is the issue?

Approximation.

Confusing Wolverhampton with Birmingham, or Stockport with Manchester is not a mistake that a human will make twice. In low-stakes contexts, this may be irritating or amusing. In operational systems, the same kind of approximation can become dangerous.

AI agents are absolutely not human. They find the route to a goal based upon a complex approximation of the instructions provided. And the imprecision of that approximation matters. Metaphorically speaking, it could repeatedly tell you that Wolverhampton is in Birmingham, without ever being able to contextualise the offence that may be caused. It means that it may take a path that is entirely unexpected, or was explicitly prohibited, and do so repeatedly. This approximation is the essence of the power of this technology, but also the risk it poses when used operationally.

This approximation also applies when the agent is asked to describe its actions. It will again 'approximate' a response based on the instructions given and the data available. The

result may be an agent that first breaks your systems and then (appears to) outright lie about what it did. Of course, it isn't 'lying' (and it is unhelpful to think about AI in those human terms) it is a process making a critical approximation error.

Is AI dangerous?

It is potentially very dangerous. Examples of AI deleting data, weakening security, removing user's permissions, or directly hacking systems are common, although frequently presented with a disproportionate level of fear and hyperbole. Negative agentic behaviours may appear and be presented as 'malign AI intent', but every one of these is better understood as a failure of control, cyber security, context and approximation.

That said, there is no doubt that AI adoption poses an actual threat to organisational resilience. In the NCSC we generally see two responses:

1. Accept any risk and go ahead regardless
The logic being that the threat of not adopting AI exceeds the threat of AI harm.

2. Avoid meaningful AI adoption altogether
The logic being that AI cannot be used safely, so it cannot be adopted.

This is a deeply uncomfortable choice for any leader. No one wants to choose between flying too close to the sun or too close to the sea; neither ends well.

Effective cyber security is the 'AI way'

Fortunately, organisations that have invested in cyber security do not face such a binary choice. Cyber security provides the means to mitigate the risk of actions by humans, systems and now AI agents. Organisations that effectively manage the 'insider threat' (the threat of employees deciding to cause harm to the organisation that employs them), are well placed to manage the risk of AI agents acting as malign insiders when making "critical approximation errors". It absolutely isn't easy, but the organisations that approach this challenge maturely can adopt AI faster and safer than their less-cyber-mature counterparts, precisely because they are better able to constrain agent behaviour and recover seamlessly should it cause harm.

Such organisations understand that while AI can likely solve some of their business needs more efficiently than humans.

The process of AI adoption is one of active cyber security and broader risk management. What stands in the way becomes the way.

In short, AI agents both increase organisational capability, but bring with them resilience risk.

Organisations that address those resilience risks maturely, through technical and procedural cyber security controls, will succeed in adopting AI sustainably.

The organisations that understand that this is where they need to invest are the organisations that will thrive in this new AI era.

National opportunity

Going broader, the UK has a long-standing reputation for pragmatic and effective risk management. The UK also has a thriving cyber security sector, providing 17% growth in GVA over the last year, and more than doubling over the last 5 years. It is a sector ready to take on the challenge if given the opportunity. We are fortunate that in the UK we have the foundation stones and innovation needed to nationally succeed in adopting AI maturely and sustainably.

And, potentially, the prize could be far greater. It is possible, perhaps likely, that there is a limit to how far general-purpose frontier AI models can go in understanding their context and risk, particularly in sector-specific contexts. What is common sense for an accountant is not the same as what is common sense for a doctor. Technological innovation that helps to overcome this challenge will very likely have significant global value.

National imperative

At the same time, we are seeing adversaries adopt AI at pace, increasing their scale and capability. As said before, AI is transforming cyber security, and unfortunately, one of the first impacts of AI is to increase the threat to organisations globally, including in the UK. This threat is very real and means that, whether organisations are adopting AI or not, the need to implement cyber security fundamentals and improve cyber resilience is not optional.

Fortunately, the investment essential to defend against AI-enabled cyber threats is largely the same investment required to unlock the UK's AI future. If we realise that sustainable AI adoption is our path, and invest today in developing the cyber security and risk management tools that are needed, then UK organisations will thrive in the AI era.



References

CHAPTER 08
The middle way: How mid-sized businesses can lead the UK's path to AI Superpower status

Over half of firms actively using AI: British Chambers of Commerce, Future of Work: AI in the Workplace, March 2026. https://www.britishchambers.org.uk/news/2026/03/half-of-smes-using-ai-with-limited-headcount-impact-so-far/

More than three-quarters of AI-using firms yet to see any change in revenue: Department for Science, Innovation and Technology, AI Adoption Research, January 2026. https://www.gov.uk/government/publications/ai-adoption-research

56% of high-growth businesses investing in AI; 21% citing regulation as a major driver of their innovation plans, and 24% wanting improved regulation for their sector: BDO, Innovate to Grow 2026. https://www.bdo.co.uk/en-gb/insights/tax/innovate-to-grow

Leading models matching or beating industry experts around eight times in ten (83% of comparisons on GDPval, as reported by the developer): OpenAI, Introducing GPT-5.4, March 2026. https://openai.com/index/introducing-gpt-5-4/

GDPval benchmark method: Patwardhan et al., GDPval: Evaluating AI Model Performance on Real-World Economically Valuable Tasks, ICLR 2026. https://arxiv.org/abs/2510.04374

Three-quarters of professional services firms not ready on data, orchestration and monitoring: Department for Science, Innovation and Technology, AI Adoption Plan: Professional and Business Services, June 2026. https://www.gov.uk/government/publications/ai-champions-ai-adoption-plans/ai-adoption-plan-professional-and-business-services

Mid-market share of private-sector businesses and revenues: BDO, Regional businesses drive mid-market expansion but targeted support required to maximise growth, July 2026. https://www.bdo.co.uk/en-gb/news/2026/regional-businesses-drive-mid-market-expansion-but-targeted-support-required-to-maximise-growth

Minimum viable governance: Nick van der Meulen, Jennifer Jewer and Nadège Levallet, Minimum Viable Governance for Generative AI, MIT Center for Information Systems Research, March 2026. https://cisr.mit.edu/publication/2026_0301_GenAIGovernance_VanderMeulenJewerLevallet

CHAPTER 09
The constraint is moving

[1] David, P.A. (1990) 'The dynamo and the computer: An historical perspective on the modern productivity paradox', American Economic Review

[2] Bharier, D., Etheridge, B. and Morais, P. (2026) 'AI adoption and workforce change in SMEs', ISER Working Paper Series. Colchester: University of Essex

[3] Coase, R.H. (1937) 'The nature of the firm', Economica, 4(16)

[4] Galbraith, J.R. (1973) Designing Complex Organizations. Reading, MA: Addison-Wesley

[5] Peng, S. et al. (2023) 'The impact of AI on developer productivity: Evidence from GitHub Copilot', arXiv

[6] Dell'Acqua, F. et al. (2026) 'Navigating the jagged technological frontier', Organization Science, 37(2)

[7] Bresnahan, T. and Greenstein, S. (1996) 'Technical progress and co-invention in computing and in the uses of computers', Brookings Papers on Economic Activity: Microeconomics

[8] Acemoglu, D. and Restrepo, P. (2019) 'Automation and new tasks: How technology displaces and reinstates labor', Journal of Economic Perspectives, 33(2)

[9] British Chambers of Commerce and Atos (2026) Powering Productivity: AI and the Future of UK Work. London: British Chambers of Commerce

CHAPTER 12
From AI Adoption to advantage: Turning Britain's head start into growth

Figures from aibl's State of UK AI Adoption Survey 2026 (755 UK mid-market business leaders, £50m–£500m revenue, fieldwork January to March 2026), except where attributed. The Yorkshire Building Society example and the Xero (n=520 UK practices) and Klarus (n=500 UK and Ireland mid-market firms) surveys, both 2026, are drawn from aibl's weekly UK Mid-Market AI Adoption Tracker. External sources: ONS, "AI in UK businesses" (Jul 2026) and "International Comparisons of Productivity"; Eurostat, "Use of AI in enterprises" (2025) and US Census Bureau, "AI use in businesses" (2026); MIT / Project NANDA (Aug 2025) and Gartner on pilot returns; US Federal Reserve, "The State of AI Competition in Advanced Economies" (Oct 2025); NIST, Manufacturing Extension Partnership economic impact (2024); OECD, "SME Technology Adoption in the United Kingdom" (Apr 2026). The £550 billion opportunity, the £150 billion cost of a five-year delay and the £5-per-£1 return are modelled by Public First in "Unlocking the UK's AI Potential" (for Microsoft, 2025); UK GDP (about £3.0 trillion, 2025) and 2026 growth are from ONS and the Office for Budget Responsibility; UK productivity history and the output-per-hour gap with the United States are from the House of Commons Library and ONS. Mid-market economic weight uses BDO's mid-sized-business definition and the ScaleUp Institute's 2025 review.

CHAPTER 14
From AI users to AI transformers

1 Department for Science, Innovation and Technology, AI Opportunities Action Plan, 13 January 2025, https://www.gov.uk/government/publications/ai-opportunities-action-plan/ai-opportunities-action-plan.

2 Microsoft AI Economy Institute, Global AI Diffusion: Q1 2026 Trends and Insights, May 2026, https://www.microsoft.com/en-us/research/wp-content/uploads/2026/05/Microsoft-AI-Diffusion-Report-2026-Q1.pdf.

3 British Chambers of Commerce, Powering Productivity: AI and the Future of UK Work, March 2026, https://www.britishchambers.org.uk/wp-content/uploads/2026/03/Powering-Productivity-AI-and-the-Future-of-UK-Work_FINAL.pdf.

4 British Chambers of Commerce, The Turning Point for SMEs: Unlocking the Next Level of AI, September 2025, https://www.britishchambers.org.uk/wp-content/uploads/2025/09/The-Turning-Point-for-SMEs-Unlocking-the-next-level-of-AI22.pdf.

5 Office for National Statistics, Artificial intelligence in UK businesses: 2023 to 2026, 20 July 2026, https://www.ons.gov.uk/businessindustryandtrade/business/businessservices/articles/artificialintelligenceinukbusinesses/2023to2026.

6 Deloitte UK, British workers spend nearly £1bn of their own money on GenAI for work, 16 September 2026, https://www.deloitte.com/uk/en/about/press-room/british-workers-spend-one-billion-pounds-of-their-own-money-on-gen-ai-for-work.html.

7 Department for Business, Innovation, Science and Trade and Department for Science, Innovation and Technology, AI Adoption Research, 28 January 2026, https://www.gov.uk/government/publications/ai-adoption-research.

8 Department for Business, Innovation, Science and Trade, SME Digital Adoption Taskforce: final report, 31 July 2025, https://www.gov.uk/government/publications/sme-digital-adoption-taskforce-final-report.

9 Bennett School of Public Policy, What does firm-level data tell us about AI adoption in the UK?, 13 April 2026, https://www.bennettschool.cam.ac.uk/blog/ai-adoption-in-the-uk/.

10 WPI Strategy and Microsoft, Unlocking Regional Growth: The Impact of AI Adoption by SMEs, 2025, https://www.wpi-strategy.com/microsoft-sme-ai-adoption.

11 Microsoft, 2025 Work Trend Index Annual Report: The Year the Frontier Firm Is Born, 23 April 2025, https://www.microsoft.com/en-us/worklab/work-trend-index/2025-the-year-the-frontier-firm-is-born.

12 Federal Ministry for Economic Affairs and Climate Action, Executive Summary of the Final Report: Evaluation of Mittelstand-Digital, March 2025, https://www.bundeswirtschaftsministerium.de/Redaktion/EN/Publikationen/Mittelstand/executive-summary-evaluation-mittelstand-digital.pdf?__blob=publicationFile&v=5.

13 Department for Business, Innovation, Science and Trade, Made Smarter Adoption: impact and process evaluation, 30 July 2026, https://www.gov.uk/government/publications/made-smarter-adoption-impact-and-process-evaluation.

CHAPTER 15
Turning AI into everyday SME growth

1 Business population estimates for the UK and regions 2025

2 SME Digital Adoption Taskforce 2025

3 Eagle Labs: The AI 100 2026

4 Barclays: Business Prosperity Index, Q1 and Q2 2026

5 Office for National Statistics: Management practices & the adoption of technology and artificial intelligence, 2023-2025

6 Office for National Statistics, July 2026

CHAPTER 16
The real AI challenge is not technology. It is change

1 https://www.gov.uk/government/publications/ai-opportunities-action-plan/ai-opportunities-action-plan Accessed 25/08/26.

2 Kotter, J.P., 2014. Accelerate: Building strategic agility for a faster-moving world. Harvard Business Review Press.

3 Kotter, J.P., 2012. Accelerate. Harvard business review, 90(11), pp.45-58

4 Bi, K., Xie, L., Zhang, H., Chen, X., Gu, X. and Tian, Q., 2022. Pangu-weather: A 3d high-resolution model for fast and accurate global weather forecast. arXiv preprint arXiv:2211.02556.

5 https://www.turing.ac.uk/blog/fastnet-under-bonnet-our-ai-model-weather-prediction. Accessed 10/08/26.

6 Bouallègue, Z.B., Clare, M.C., Magnusson, L., Gascon, E., Maier-Gerber, M., Janoušek, M., Rodwell, M., Pinault, F., Dramsch, J.S., Lang, S.T. and Raoult, B., 2024. The rise of data-driven weather forecasting: A first statistical assessment of machine learning-based weather forecasts in an operational-like context. Bulletin of the American Meteorological Society, 105(6), pp.E864-E883.

7 Kent, C., Scaife, A.A., Dunstone, N.J., Smith, D., Hardiman, S.C., Dunstan, T. and Watt-Meyer, O., 2025. Skilful global seasonal predictions from a machine learning weather model trained on reanalysis data. npj Climate and Atmospheric Science, 8(1), p.314.

8 de Burgh-Day, C.O. and Leeuwenburg, T., 2023. Machine learning for numerical weather and climate modelling: a review. Geoscientific Model Development, 16(22), pp.6433-6477.

9 Dueben, P., Bauer, P., Fuhrer, O., Koldunov, N. and Kristiansen, J., 2026. Machine learning is revolutionizing weather forecasting--the next step is a change in how we work. arXiv preprint arXiv:2606.25076.

10 Nipen, T.N., Haugen, H.H., Ingstad, M.S., Nordhagen, E.M., Salihi, A.F.S., Tedesco, P., Seierstad, I.A., Kristiansen, J., Lang, S., Alexe, M. and Dramsch, J., 2026. Regional data-driven weather modeling with a global stretched grid. Artificial Intelligence for the Earth Systems, 5(2), p.250001.

11 Addison, H., Kendon, E., Ravuri, S., Aitchison, L. and Watson, P., 2023, May. Downscaling with a machine learning-based emulator of a local-scale UK climate model. In EGU General Assembly Conference Abstracts (pp. EGU-14253).

12 Van Katwyk, P., Fox-Kemper, B., Hewitt, H.T. and Bergen, K.J., 2026. Rewiring climate modeling with machine learning emulators. Communications Earth & Environment, 7(1), p.107.

13 Bolón-Canedo, V., Morán-Fernández, L., Cancela, B. and Alonso-Betanzos, A., 2024. A review of green artificial intelligence: Towards a more sustainable future. Neurocomputing, 599, p.128096.

14 Rolnick, D., Donti, P.L., Kaack, L.H., Kochanski, K., Lacoste, A., Sankaran, K., Ross, A.S., Milojevic-Dupont, N., Jaques, N., Waldman-Brown, A. and Luccioni, A.S., 2022. Tackling climate change with machine learning. ACM computing surveys (CSUR), 55(2), pp.1-96.

15 https://www.metoffice.gov.uk/blog/2026/can-ai-developments-really-be-green Accessed 18/08/26.

16 Pope, E.C., Dale, K.I., Steele, E.C. and Graham, P.H., 2026. Weathering the Artificial Intelligence revolution: public attitudes to potential adoption of Machine Learning-based forecasts. Artificial Intelligence for the Earth Systems, p.e250068.

17 Werner, J. and Le-Brun, P., 2025. Become an octopus organization. Harvard Business Review, Nov–Dec.

CHAPTER 18
The interaction of people and technology is at the heart of AI growth

[1] What the UK thinks about AI | Ipsos
[2] Mark Rutte | NATO Biography
[3] Neil Carberry, Chloe Smith and Chris Bones: This is the moment for our profession to lead on performance | CIPD
[4] Neil Carberry: people professionals are central to performance and productivity | CIPD
[5] Keynote address at The HR Conference and Exhibition 2026 | IGPP
[6] People Management, May-July 2026
[7] How to future-proof skills as AI reshapes work | CIPD
[8] The-Turning-Point-for-SMEs-Unlocking-the-next-level-of-AI.pdf
[9] Shevaun Haviland - Global Annual Conference Speech 2026 - British Chambers of Commerce
[10] CIPD Festival of Work 2026 and further case study supplied
[11] With thanks to senior CPOs for their input

The British Chambers of Commerce (BCC) is the UK's leading business network, representing companies of every size and sector through a unique network of 51 accredited Chambers across the UK and more than 70 international Chambers worldwide.

We connect, support and champion businesses, helping them start, grow, trade and succeed. Through our advocacy work, we give businesses a powerful voice in government, while our Chamber network delivers practical support, trusted advice, international trade services, events, skills development and valuable connections.

Together, we work to make the UK the best place to start, grow, invest in and do business.

ISBN 978-1-0655584-0-8

BDO

British Chambers of Commerce