



BAYES
BUSINESS SCHOOL
CITY ST GEORGE'S
UNIVERSITY OF LONDON

AI and the Future of Business

WHITE PAPER SERIES
NO.1 (2026)

[Always learning]

bayes.citystgeorges.ac.uk



Contents

Foreword, Professor André Spicer	4
Preface, Professor Feng Li	5
Introduction	7
What Today's AI Can and Cannot Do	9
The Transformative Power of AI	11
Our Research Agenda	13
AI, Work and Organisational Transformation	
AI for Innovation, Creativity and Entrepreneurship	
AI, Markets and Financial Systems	
AI-Enhanced Actuarial Science and Insurance	
AI in Supply Chain Management and Healthcare Operations	
Responsible, Ethical and Sustainable AI	
Research Institutes, Centres and Groups at Bayes Business School	20
Education	21
Conclusion	22
Working with Us	23

Foreword

Professor André Spicer
Executive Dean



The City of London has been at the forefront of industrial, financial and professional innovation for centuries. That tradition continues today. The financial institutions, insurance companies and professional services firms around us continue to shape the global economy, while the technology and creative firms around Shoreditch are helping to build the companies and industries of the future.

Bayes Business School sits within this long history of innovation. Over our sixty-year history, we have focused on developing professionals and leaders for the City of London and beyond. We have also produced research and insight that inform the businesses, institutions and policymakers with whom we work in London, nationally and around the world.

As a leading global business school, Bayes continues to produce work of the highest scientific merit. We are home to nearly 180 faculty members, many of whom are internationally recognised scholars in their fields. We also house more than a dozen research centres, through which our researchers work with governments, businesses and non-profit organisations to address practical problems with rigorous evidence and analysis.

Over the years, our research has helped to advance the actuarial profession, support central banks in regulating the financial sector, improve the financing of the shipping industry, inform the allocation of hospital beds, identify ways to debias decision-making, help businesses rethink their models and support professionals as they move into leadership roles. These examples represent only a small part of the work undertaken across Bayes.

At Bayes, we believe in curiosity. We support a wide ecology of curiosity-driven research, from misogyny at work, the workings of the art market and copula functions in mathematical models, to forecasting techniques, business jargon, protest movements and curiosity itself.

We also believe that research should create change. This requires us to connect ideas, disciplines and expertise in ways that help address some of the largest challenges facing business and society.

That is why we have organised our work around four major research themes: AI and the Future of Business; the Future of Finance and Insurance; Tackling Societal Challenges; and Adapting to a Transforming World.

Under each of these themes, we bring together experts working on some of the most pressing questions of the day. Our researchers are examining the implications of an ageing society, how to decarbonise shipping, how to harness AI to support creativity and how traditional financial institutions and regulators can respond to innovation in FinTech.

This white paper series is an invitation to engage with our work. In the coming years, we look forward to working with partners to apply high-quality research insight to the issues affecting their organisations, sectors and communities. In doing so, we hope to provide ideas and evidence that help the City of London and the wider economy and society it serves continue to innovate and thrive in the decades ahead.

Preface

Professor Feng Li

Associate Dean for Research and Innovation



For much of the post-Cold War period, business strategy was shaped by a set of assumptions that appeared broadly stable. Markets were expected to deepen, supply chains to remain open, technological progress to diffuse through existing institutions and major disruptions to be episodic rather than systemic. Competition could be intense, but the overall direction of travel seemed relatively clear.

Many of these assumptions no longer hold. Geopolitical rivalry now permeates trade, technology, capital, energy and industrial policy. Economic openness can no longer be treated as a neutral background condition. Institutional resilience can no longer be assumed. At the same time, technological change is advancing in ways that do not simply create new opportunities, but unsettle established assumptions about governance, expertise, trust, accountability and control. Leaders are not just managing change within stable systems. They are increasingly operating amid fundamental shifts in the systems themselves.

This is the context in which Bayes Business School has identified and articulated four Strategic Research Themes: Tackling Societal Challenges, AI and the Future of Business, The Future of Finance and Insurance and Adapting to a Transforming World. These are not separate agendas assembled for convenience. They are different ways of approaching the same underlying challenges around how organisations create value, manage risk and sustain legitimacy when the conditions that once underpinned strategy and operations are becoming unreliable.

The purpose of this white paper series is not to offer slogans, reassurance or fashionable language dressed up as insight. It is to bring serious research to bear on the practical and institutional questions now confronting business and society. This means looking beyond surface

disruption to the deeper shifts underneath it, questioning inherited frameworks where necessary and engaging with firms, policymakers and other partners who recognise that older playbooks are losing their grip. The unifying conviction behind this series is that rigour matters most when the world becomes harder to read.

Artificial intelligence (AI) is the most visible expression of these wider shifts. It attracts both inflated claims and shallow dismissal, neither of which is particularly useful. Its importance does not lie in spectacle, nor in the familiar cycle of hype and backlash that accompanies each new wave of technology. It lies in the growing ability to embed machine intelligence into the operating fabric of organisations, in generating content, supporting decisions, shaping workflows and coordinating activity. In a growing number of settings, it also takes action with limited human intervention.

Firms are not simply collections of tasks and processes. They are structured around roles, routines, incentives, professional judgement, lines of accountability and culture that have evolved over time. Once intelligence begins to migrate into technical systems, the issue becomes a question of organisational architecture. Which decisions remain human; which are delegated; how is oversight exercised; how is expertise reconfigured and trust maintained; and who carries responsibility when outcomes are consequential but causality is opaque?

These questions go to the heart of management and governance. The central challenge is not model capability.

In many domains, the more pressing constraint is whether organisations and institutions are able to deploy AI responsibly, effectively and at scale.

The barriers are managerial, social and institutional as much as technical, including legacy systems, path dependence, weak governance, unclear incentives, skills gaps, concentration risk, cyber vulnerability, regulatory uncertainty and fragile public trust. The history of technology suggests that transformation is rarely driven by capability in isolation. It depends on whether institutions are able to comprehend, absorb, direct and legitimise what new technology makes possible.

Bayes Business School is well placed to contribute to this agenda. We examine AI through the lenses of strategy, operations, organisation, finance, insurance, work, innovation, governance and public value. We are concerned with how intelligence is operationalised, how organisations are redesigned around it, how benefits and risks are distributed and how institutions respond without mistaking motion for progress. The purpose of this white paper is to clarify the organisational and institutional

choices that will determine whether AI becomes a source of real capability, resilience and public value. The gap between what AI can do and what organisations can responsibly govern will be the defining strategic terrain of the next decade.



Introduction

AI is moving into the operating fabric of organisations. For most firms, the question is no longer whether to adopt it but how to use it well.



AI is moving into the operating fabric of organisations. For most firms, the question is no longer whether to adopt it but how to use it well. The technology is advancing quickly. The real constraint lies in the managerial and institutional work needed to deploy AI responsibly and at scale, in ways that organisations and society can absorb and govern.


This paper is a positioning statement. It sets out how Bayes Business School reads the challenges, where our research is focused and where we think the most productive ground lies for collaboration. Several questions run through this white paper.

Where does AI create real capability and where does visible adoption conceal shallow change? How should firms govern decisions taken by systems whose reasoning is often opaque and whose failures may be correlated across the economy? Which constraints on adoption matter most, the technical limits of the systems or the managerial and institutional conditions that determine whether a technology is actually absorbed? And how should leaders respond to a technology that is reshaping organisations, business models and markets at the same time?

These questions cut across familiar boundaries. Bayes Business School is well placed to address them. Our AI research spans the school's three faculties of Management, Finance and Actuarial Science and Insurance and is organised around six interrelated research areas that reflect both our strengths and the issues now facing firms, regulators and policymakers.

What Today's AI Can and Cannot Do

The most significant shift is that AI, especially in combination, is moving from decision support towards delegated intelligence.



120kV/ 375mAs
1.50s/4mm/1.0x4

M:213.75mm(181.30)
4791: 4: 7*
24.00mm
(24.0mm)
9.5D
(256,216)

120kV/ 375mAs
1.50s/4mm/1.0x4

M:213.75mm(181.30)
4791: 4: 8*
28.00mm
(28.0mm)
9.5D
(256,216)

WL=40
WW=150

WL=40
WW=150

120kV/ 375mAs
1.50s/4mm/1.0x4

M:213.75mm(181.30)
4791: 4:23*
40.00mm
(40.0mm)
9.5D
(256,216)

120kV/ 375mAs
1.50s/4mm/1.0x4

M:213.75mm(181.30)
4791: 4: 9*
44.00mm
(44.0mm)
9.5D
(256,216)

Serious strategy starts with clarity about what AI can and cannot do. Too much current commentary slides between sweeping claims about machine intelligence and narrow observations about particular tools. The capabilities, risks and governance requirements differ sharply.

Artificial General Intelligence (AGI), in the sense of systems with human-level cognitive abilities across tasks, does not exist yet. Whether it ever will and on what timescale, remains contested. What does exist is a range of working AI that, for practical purposes, can be grouped into three broad families.

Analytical AI applies rules and statistical models to structured tasks such as classification, prediction and optimisation. It is already pervasive in fraud detection, credit scoring, logistics and medical imaging. In many narrowly defined tasks it outperforms human experts.

Generative AI, built on large language and multimodal models, produces new outputs in text, code, images and audio. It is used for drafting, summarisation, coding assistance and customer interaction. Its capabilities are broad but uneven across domains and it is prone to fabrication (sometimes described as ‘hallucination’).

Agentic AI combines these capabilities with planning and action. It can set goals, sequence steps and execute tasks with limited human supervision. This is the frontier where questions of oversight, liability and control are least settled.

The most significant shift is that AI, especially in combination, is moving from decision support towards delegated intelligence. The system no longer simply informs human decisions. In a growing number of cases, it takes the decision and acts on it. The managerial and governance implications are still unresolved.

What AI is not is equally important. It may display operational autonomy in bounded domains, but its behaviour remains shaped by human-defined objectives, training data and institutional context.

Biases in the data reappear in outputs. Optimisation targets set by developers become consequential decisions about who benefits and who bears the cost. In insurance pricing, for example, AI can enable highly granular risk assessment that improves accuracy while straining long-standing norms around fairness and mutualisation. Even where sensitive attributes are formally excluded, models can reconstruct them indirectly. AI is a social and institutional technology as much as a technical one.

The Transformative Power of AI

The history of consequential technologies suggests that their primary effect is rarely straightforward labour substitution.



Public debate about AI has become fixated on job displacement, especially the claim that white-collar knowledge work is now exposed to automation in the way manual and routine work was in earlier waves. The concern is real, but it frames the question too narrowly.

The history of consequential technologies suggests that their primary effect is rarely straightforward labour substitution. They reorganise work, firms and markets around new capabilities, often in ways that only become clear long after the initial disruption. Radiology is a telling case. AI systems can now match or exceed human performance on some image-recognition tasks and for years radiologists were presented as obvious candidates for replacement. This has not happened. Demand for radiologists has continued to increase. Image reading is only one part of the role. Automation may replace or improve a limited set of tasks, but radiologists still carry broader responsibilities involving interpretation, judgement, communication and oversight. The effect is not the disappearance of the profession, but a change in the work it does.

The more important effects of AI will come from this kind of reshaping across firms and sectors. Business models will change. Workflows will be redesigned. The boundaries between firms, platforms and infrastructures will shift as intelligence becomes embedded in how the economy operates. So too will the distribution of value.

Diffusion of a general-purpose technology is not a single process. Infrastructure, applications and the organisational adaptation around them each move at different speeds and the slowest phase sets the pace. Earlier digital general-purpose technologies unfolded over long periods and many of their most consequential uses were not apparent at the outset. AI's delivery layer is already partly in place, but its compute and power substrate is not and the institutional work of absorbing it has barely begun. The pattern will be uneven and partly unexpected, shaped as much by organisational choices, institutional capacity and culture as by the technology itself.

For leaders and policymakers, the strategic question is not how to use AI to do existing work more cheaply. It is how to redesign roles and workflows and build organisations, markets and institutions that are more adaptive, resilient and legitimate.

This means looking beyond immediate applications to the governance frameworks, skills systems and institutional arrangements that determine how benefits are distributed and risks contained.

Our Research Agenda

Bayes Business School sits in the City of London, with direct exposure to the institutions, markets and policy networks in which many of these questions are now playing out. The six research areas below reflect where our expertise is strongest and where business leaders, regulators and policymakers are confronting the hardest questions.



AI, Work and Organisational Transformation

AI is reshaping the architecture of work. As intelligence increasingly migrates into technical systems, authority is redistributed inside firms. Decisions that used to be embedded in professional roles are unbundled. Workflows that relied on tacit coordination are reconstituted around systems. Accountability, which has long travelled with role and rank, becomes harder to locate.

Our research examines the conditions under which AI is genuinely absorbed into organisations at scale. We study how human-AI collaboration works in practice, where algorithmic trust is warranted and where it is misplaced and how complementarities between people and machines are designed and sustained. We place AI within the longer history of digital transformation, which over three decades has already reshaped job design, leadership, professional identity and organisational form.

These organisational questions sit alongside more difficult institutional challenges. When intelligence is delegated, service delivery, regulation and public value are all reconfigured and the existing models of organising

and governing become strained. Firms and the leaders within them now face a sustained period of structural change for which few are well prepared.

A consistent finding runs through this work. The gains from AI depend on the quality of the organisational work done alongside the technical deployment. Firms that treat AI as a drop-in replacement typically see limited returns. Firms that redesign roles, processes, incentives and oversight around it capture more value and are better placed to avoid the governance failures that accompanied earlier waves of automation.

Research in this theme is led by

Professor Nicolas Bacon
Dr Aneesh Banerjee
Dr Eugenia Cacciatori
Professor Itziar Castelló-Molina
Dr Daisy Chung
Professor Laura Empson
Professor Gianvito Lanzolla
Professor Feng Li
Dr Yusuf Oc
Dr Nettra Pan
Dr Jafar Sabbah
Dr Simone Santoni

IN FOCUS:

Digital transformation and delegated intelligence

One of the most consequential developments in digital transformation is delegated intelligence. Organisations are assigning bounded areas of judgement and action to digital systems. The implications are organisational and institutional. Authority shifts. Professional expertise is reworked around oversight and intervention. Accountability becomes more difficult when outcomes are shaped by models, data, external platforms and design choices outside the firm.

AI is not just another wave of automation. How far can organisations delegate intelligence before they weaken the very capacity needed to govern it? Research led by Professor Feng Li investigates what separates delegation that creates sustained value from delegation that accumulates hidden risk. It builds on earlier Bayes work on digital transformation and platform competition, funded by UK Research Councils and global technology firms and feeds into current debates on algorithmic accountability and AI governance.

AI for Innovation, Creativity and Entrepreneurship

The dominant framing of AI in the creative industries is defensive, with AI seen as a threat to livelihoods, intellectual property and the distinctive value of human creative work. The concern is legitimate, but it narrows the perspective. The more consequential issue is how AI reshapes the processes through which novelty, insight and commercial value are produced and how firms, individuals and institutions position themselves within these changing processes.

Earlier general-purpose technologies rarely destroyed creative fields. Photography did not end painting. It redirected it. Digital tools did not end graphic design. They enlarged it and changed who could participate. The pattern is redistribution and reorganisation, not disappearance. What is distinctive about AI, especially generative AI, is the speed and scope of the change and the extent to which the creative processes themselves are being reshaped. Idea generation, iteration and refinement are now routinely mediated by co-creative tools.

Our research examines how AI augments creative processes, how sectors such as media, fashion, retail and professional services are being reshaped by AI-enabled design and customisation and how entrepreneurship itself changes as the capital and technical thresholds for building new products and services fall. The internet and mobile platforms enabled a generation of capital-light start-ups, but these still required significant technical capability. AI is lowering the threshold further. In the coming decade, entrepreneurship may depend increasingly on the ability to identify worthwhile problems and frame them imaginatively, rather than on the technical capacity to build every part of the solution from scratch.

Research in this theme is led by

Dr Aneesh Banerjee
Dr Eugenia Cacciatori
Professor Clive Holtham
Dr Sara Jones
Professor Feng Li
Professor Neil Maiden
Dr Jafar Sabbah

IN FOCUS:

INSIGHTS and the co-creative AI toolkit

Centre for Creativity enabled by AI's (CebAI) portfolio of co-creative AI tools is built on a specific premise: that generative AI is most useful in creative work when it supports human thinking rather than substitutes for it. The portfolio includes Business Sparks for strategic ideation, Design Sparks for service design and Sport Sparks, developed with the English Institute of Sport.

INSIGHTS, the latest addition, combines generative AI with structured creative-thinking methods to help users identify opportunity spaces that large language models used on their own tend to miss. Partners include Miele and the Centre for Management Consulting Excellence.

This work now continues through the Institute for Creativity and AI, which examines the future of creative work in the presence of AI and develops new forms of co-creative AI for professional contexts, in collaboration with research partners across the UK and internationally.

AI, Markets and Financial Systems

Finance has been reshaped by decades of technological change. AI is now entering financial systems at every layer, from trading and portfolio construction to customer service, credit decisions, compliance and regulatory surveillance. The central question is whether its adoption strengthens or undermines the stability, inclusiveness and legitimacy of the financial system.

Our work builds on a substantial body of research at Bayes on the digitalisation of finance, including central bank digital currencies, stablecoins and decentralised finance and on the role of digital infrastructure in expanding access to liquidity for small and medium-sized enterprises. AI raises a set of significant questions. How can it be used to improve financial decision-making, market efficiency and regulatory oversight in ways that are genuinely reliable? What new risks does it introduce, including opacity, vulnerability to adversarial attack, embedded bias, correlated failure and reliance on a small number of dominant model providers? And what governance, disclosure and supervisory frameworks will be needed as AI becomes structurally embedded in financial infrastructure?

Concentration risk deserves particular attention. If a small number of foundation-model providers become the backbone of credit scoring, fraud detection, compliance monitoring and algorithmic trading across the financial system, the result will be a new form of correlated exposure. Existing regulatory frameworks are not well designed for this kind of dependency. It is becoming an important issue for central banks, regulators and firms alike.

Research in this theme is led by

Professor Pawel Bilinski
Professor Barbara Casu
Professor Michel Driessen
Professor Vasso Ioannidou
Professor Feng Li
Dr Guangyu Li
Professor Scott Moeller
Professor Iqbal Owadally
Dr Chuanping Sun
Dr Andrew Yim
Dr Junli Zhao
Dr Rui Zhu

IN FOCUS:

AI in mergers and acquisitions

The Mergers & Acquisitions Research Centre (MARC), founded at Bayes in 2008 by Professor Scott Moeller and now directed by Professor Michel Driessen, is a leading centre dedicated to both the research and practice of M&A. It is a resource used by corporations, regulators, professional services firms and universities for access to current research and practical ideas. Its sponsors have included Ardian, BP, KPMG, SS&C Intralinks, ICAEW, Mergermarket, White & Case and WTW.

Its 2023 book *The Deal Paradox*, by Anna Faelten, Michel Driessen and Scott Moeller, examines how AI, big data and machine learning are changing deal sourcing, due diligence, valuation and post-merger integration. MARC also runs public-facing events such as its annual Outlook for M&A and annual academic research conference, which continue to draw senior practitioners and market participants.

AI-Enhanced Actuarial Science and Insurance

Insurance is, by its nature, a data-intensive industry built on risk modelling, variable pricing and the management of long-tail uncertainty. It is a natural terrain for AI and the sector is being reshaped rapidly by hybrid actuarial-AI models and new data sources including text, images and telematics. The potential gains in pricing accuracy, claims handling, fraud detection and operational efficiency can be substantial but so are the risks, which cut across fairness, accountability, explainability, privacy and the long-standing institutional norm of risk mutualisation.

Our research at Bayes sits at the interface of actuarial science, machine learning and regulation. We examine how hybrid models can improve pricing, reserving and mortality modelling while remaining transparent, interpretable and auditable. We investigate how natural language processing and reinforcement learning can support investment portfolio construction for pension funds and individual savers. This includes features that standard models often handle poorly, such as sustainability preferences, illiquid assets like housing, health and long-term care costs and tax. We also examine where large language models and synthetic

data add real operational value in insurance, from claims triage and document automation to regulatory reporting and what controls, including human oversight, audit trails and bias safeguards, are required for responsible deployment.

A connected strand examines fairness and explainability in insurance AI, including how firms can demonstrate compliance through modelling, validation, documentation and ongoing monitoring aligned with model risk management and regulatory expectations. These questions are now live for regulators developing supervisory frameworks for AI in insurance and for insurers themselves, for whom unexplainable models carry material business and reputational risk.

Research in this theme is led by

Professor Vali Asimit
Professor David Blake
Dr Cormac Bryce
Dr Russell Gerrard
Professor Steven Haberman
Dr Simone Krummaker
Professor Les Mayhew
Professor Pietro Millossovich
Professor Iqbal Owadally
Professor Andreas Tsanakas
Dr Rui Zhu

IN FOCUS:

AI and disruptive technologies in actuarial work

Research led by Dr Simone Krummaker examines how AI and other disruptive technologies are reshaping actuarial work itself, from pricing and reserving to claims handling, fraud detection and regulatory reporting. The work asks how far the profession's methods, skills and institutional practices need to change as machine learning and generative AI move from tools used by actuaries to systems that increasingly carry out actuarial tasks directly.

The underlying question is central for insurers and regulators. What does responsible deployment of AI look like inside a profession whose authority rests on demonstrable rigour, reproducibility and professional accountability? This work speaks directly to current debates within the actuarial profession and insurance regulation about how AI can be used without weakening auditability, professional judgement or trust.

AI in Supply Chain Management and Healthcare Operations

Supply chains and healthcare operations are natural candidates for AI-driven efficiency gains. Both involve large volumes of data, complex interdependencies and scheduling problems that strain human cognition. The central question is how far AI can move operational decision-making from reactive to anticipatory and what organisational and governance conditions are required to make the shift safely.

Our research examines AI in procurement, logistics and supply chain planning, where its ability to handle high-dimensional data and combinatorial complexity exceeds what any human team can do unaided. It looks at how AI can improve governance in buyer-supplier relationships and across supply networks, particularly in combination with the Internet of Things and distributed ledger technologies. Traceability and transparency requirements are now rising from customers and from legislation on modern slavery, due diligence and environmental disclosure. It also examines applications in healthcare operations, including the dynamic scheduling of elective procedures,

where AI meets ethical, regulatory and clinical considerations that rarely arise in commercial supply chains.

A growing strand of this work sits at the intersection of AI with geopolitical and climate risk. As post-Cold War assumptions about open, integrated supply chains have broken down and as climate impacts have become material to operational planning, firms face a set of risks in which political, environmental and operational factors are correlated in ways that traditional models did not capture. AI can help firms assess, monitor and respond to these risks, from construction firms designing for flood resilience to extractive industries managing the interplay of drought and political instability in their operating regions. The methodological challenge is to integrate these risks into decision models without creating false precision. The governance challenge is deciding how much discretion to delegate to the resulting systems.

Research in this theme is led by

Dr Oben Ceryan
Professor Feng Li
Dr Florian Lucker
Professor Nikos Nomikos
Dr Dimitris Paraskevopoulos
Professor ManMohan S Sodhi
Dr Byung-Gak Son
Dr Chuanping Sun

IN FOCUS:

Reducing waiting times for elective surgery

A recent project in the Bayes operations group applied AI and advanced analytics to one of the NHS's most persistent operational problems: waiting times for elective surgery. Conventional scheduling models often treat procedures as broadly independent. In practice, theatre performance varies significantly with the experience of particular surgeon-anaesthetist pairings and with how that variability propagates across a full day's operating list.

This work quantifies that variance and develops new algorithms for constructing elective lists that keep performance within acceptable bounds. The first application is predictive: estimating the variability of a list before surgery and acting on it, for example by placing the most variable list at the end of the day. The second is structural: building lists whose variance is deliberately constrained, so that excessive under- or over-utilisation of theatres is minimised. The intended result is shorter waiting times and, indirectly, better patient outcomes through earlier treatment.

Responsible, Ethical and Sustainable AI

AI governance has moved from the margins of adoption to the centre of it. Firms that deploy AI without a credible account of how they are managing fairness, transparency, accountability and environmental impact will find their licence to operate contested by regulators, employees, customers and the public. Our research examines how this account can be built and it informs how firms, regulators and other actors are approaching the task.

The range of concerns is wide, from the appropriation of creative work into training data and the generation of deepfake content, to the reproduction of social biases in decisions about healthcare, insurance and employment, to the substantial and rising energy and water demands of large-scale AI infrastructure. The benefits of AI accrue unevenly, the costs often fall on those least able to absorb them and the institutional arrangements for redress remain underdeveloped.

Our work develops frameworks for ensuring AI systems are fair, transparent and aligned with organisational values and public expectations and examines what it takes to make these frameworks operational instead of symbolic. We investigate how AI is reshaping the regulation of economies and

societies, drawing on earlier general-purpose technologies for context without overstating the parallels. We ask how the benefits of AI can be distributed more fairly and how the risks can be prevented from settling disproportionately on already disadvantaged groups. This theme also cuts across the others in this white paper. Responsible deployment is a condition of effective deployment in work, finance, insurance and operations alike.

Research in this theme is led by

Dr Eugenia Cacciatori
Professor Itziar Castelló-Molina
Professor Feng Li
Professor André Spicer

With contributions from colleagues across the school and the wider City St George's, University of London

IN FOCUS:

Algorithmic accountability in digital transformation

Research led by Professor Itziar Castelló-Molina examines how organisations can design credible forms of algorithmic accountability in digital transformation. Drawing on work in business ethics, sustainability and platform governance, it asks how firms, regulators and other actors can move beyond declared principles to mechanisms that those affected by algorithmic decisions can actually rely on. It is particularly relevant where firms are under pressure to translate broad AI principles into governance arrangements that hold up under regulatory, legal and public scrutiny.

The issue is now live as AI moves from experimental use to operational deployment. Accountability cannot rest on abstract commitments. It requires governance arrangements that are visible, enforceable and capable of commanding trust among employees, customers, regulators and wider society. This work speaks directly to current debates on AI governance in the UK, Europe and beyond.

Research Institutes, Centres and Groups

Several Bayes research centres and groups contribute directly to this work, including:

- Institute for Creativity and AI (ICAI)
- Centre for Creativity enabled by AI (CebAI)
- Centre for Healthcare Innovation Research (CHIR)
- Mergers & Acquisitions Research Centre (MARC)
- Centre for Banking Research (CBR)
- Pensions Institute (PI)

BAYES BUSINESS SCHOOL

Education

Generative AI has been a routine part of how recent cohorts of students have learned and worked throughout their university education. The first of these cohorts has now entered the workplace. This shifts the debate beyond whether students should be allowed to use large language models in assignments.

Higher education can no longer be organised on the assumption that students will learn, think and work as previous generations did.

The implications reach well beyond assessment. They extend to curriculum, teaching, learning and the wider purpose of business education itself. Universities cannot meet this challenge by teaching the use of new tools alone. As AI becomes embedded in how organisations operate, business education must develop judgement, critical reasoning, domain understanding and a stronger sense of professional responsibility in work increasingly mediated by machine intelligence.

What we teach, how we teach it and the capabilities graduates leave with all require fundamental rethinking. At Bayes, this is already under way across our undergraduate, master's, MBA and executive programmes and through our PhD and Executive PhD pathways. Our research on AI in organisations is shaping the redesign. We collaborate across disciplines, including science and technology, engineering, medicine and law and with partners in industry, government and the third sector.

AI cannot be taught well in isolation and we welcome active engagement from a diverse range of partners.



Conclusion

AI is moving from decision support to delegated intelligence. It is becoming part of the operating fabric of organisations, shaping workflows, coordinating activity and acting within authority boundaries that firms themselves are still in the process of defining. The open question is on what terms and for whose benefit.

The challenges being tackled now are primarily organisational and institutional, not technical. What is at stake is whether firms, institutions and societies can build the conditions under which AI creates durable value, distributes that value fairly and does so without eroding the trust on which modern economies depend. This work is slow, unglamorous and uneven. It is also where the real decisions of the next decade will be taken.

The obstacles are well known. Legacy systems, weak governance, concentration risk, regulatory uncertainty, cyber vulnerability, skills gaps and fragile public confidence will not yield to further technical progress.

This is the environment in which AI will either become a source of broad capability and resilience, or settle into

another uneven technology whose gains accrue narrowly and whose costs fall on those least able to bear them. Which of these outcomes we get is not a matter of forecast. It is a matter of choice, made by a large number of actors, mostly in the next few years.

Meeting this challenge requires serious research on the organisational and institutional work of deploying AI and education that equips graduates to exercise judgement over AI-enabled systems. It requires sustained engagement with firms, regulators and policymakers and collaboration across disciplines. It requires honesty about what AI can do, about where benefits and costs will fall and about how much of the institutional architecture is still to be built.

This is the agenda Bayes is pursuing, collaboratively.



Working with Us

This paper is written to be useful to three audiences: business leaders trying to deploy AI responsibly in organisations whose people, processes and accountabilities were not designed around it; regulators and policymakers developing the supervisory and policy frameworks AI will need in order to become trustworthy at scale; and researchers working on the theoretical and empirical questions this transition now raises.

The harder questions are best worked out with partners. We are particularly interested in working with organisations grappling with the managerial and governance problems of AI deployment; with regulators and policy bodies looking for research partners on the structural questions AI is now posing for their fields; and with research groups inside and outside City St George's whose work complements ours on these questions.

Readers who would like to engage on specific areas of our research are welcome to contact the theme leads named in this paper. Enquiries of a broader nature can be directed to the Associate Dean for Research and Innovation.

The institutional scaffolding for the AI transition is still being built. We would like to work on it with you.



[bayes.citystgeorges.ac.uk/
faculties-and-research/research](https://bayes.citystgeorges.ac.uk/faculties-and-research/research)

Bayes Business School

106 Bunhill Row/33 Finsbury Square

London EC1Y 8TZ/EC2A 1AG

T: +44 (0)20 7040 8600

bayes.citystgeorges.ac.uk



BayesBusinessSchoolOfficial



Bayes-Business-School



@BayesBSchool



BayesBSchool



@BayesBSchool



All the information contained within this white paper was correct at the time of going to print.

City St George's, University of London is an independent member of the University of London which was established by Royal Charter in 1836. It consists of 17 independent member institutions of outstanding global reputation and several prestigious central academic bodies and activities.



UNIVERSITY
OF LONDON