a blueprint for better business

A framework for human-centred AI

A guide for business

July 2026



© Blueprint for Better Business 2026

Introduction

We are living through a genuine paradigm shift - AI is reshaping how businesses operate, how people work, and how society functions, at a rapidly increasing pace and scale. Although the full implications are still unknown, and the regulatory context is still evolving, it is already clear that decisions being made in the next months and years will establish norms, embed assumptions and create new economic and societal dependencies that will shape the range of possible futures.

The views of society on AI are running ahead of business or regulation. In society the conversation is about trust, fairness, what AI means for jobs, and what it means to be human. The dominant conversation on AI in business focuses on growth, efficiency, productivity, risk and compliance. These immediate priorities are no doubt important but can draw attention from the deceptively simple but more fundamental question business leaders ought to be asking: what is the role of our business in building a better future, and is the way we are thinking about the deployment of AI consistent with that?

The gravitational pull of the current economic system is toward short-termism and extraction. An example of this would be the development of social media which began with a positive intention to create connection and community but was drawn by its economic model to addict users and increase outrage in order to maximize engagement, leading to isolation and lower societal trust.

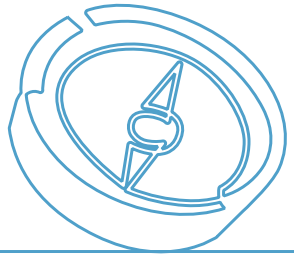
What can we do to appeal to strong leadership in business to help chart a different, more purposeful course?

We believe that businesses positioned to not only survive but thrive in a less predictable world will be those that create genuine value for the people and communities they serve, building the kind of trust with their stakeholders that increases their adaptability, relevance and resilience.

In this guide we introduce a series of tools to encourage broader conversations inside businesses, inviting different perspectives and a more holistic exploration of the potential that widespread AI implementation might bring.

In the race to effectively deploy rapidly evolving technology in a competitive market, most businesses are starting from the wrong place. Understandably, they are asking: how do we adopt AI safely? How do we manage the risks? How do we avoid being left behind?

These are good questions. But they are not the right starting point. The right starting point is: what is the role of our business in building a better future, and is the way we are thinking about the deployment of AI consistent with that?



How to use this guide

This guide is designed to help Boards and leadership teams ask better questions about AI - not just about risk and compliance, but about purpose, people, and the kind of organisation you want to become. It is to help your leadership begin a different kind of conversation: one that looks across the business as a whole and across multiple time horizons, grounds AI adoption in purpose rather than just possibility, and builds the shared understanding needed for genuine oversight.

The guide brings together **four interconnected tools**: **opening questions** to spark Board-level conversation, a **framework** to assess your organisation's current approach, **five principles** that describe what purpose-led AI looks like in practice, and a **workshop outline** to go deeper as a Board or for use in leadership teams across the business. This page shows how the components fit together and how to use them - you can work through them in sequence or return to each as your thinking develops.

1. Questions for Boards

A set of high-level questions to guide Board discussion on AI - helping leaders consider long-term resilience, impacts on people and the planet, and governance and accountability.

How to use it

Use these questions to guide open Board discussions on AI - highlighting where deeper focus is needed, not solving everything at once.

2. The Framework

A guiding framework to help catalyse deeper discussion at different levels in the organisation - that supports leaders to consider how choices shape people, relationships with stakeholders and wider society..

How to use it

Use the framework to assess where your organisation sits today and where you aspire to be. Read bottom-to-top within each column.

3. The Five Principles

A set of aspirational outcomes that describe what a purpose-led, human-centred organisation looks like in practice when designing and deploying AI in service of people, society and the planet.

How to use it

Use the Principles to challenge and guide decision-making - helping your organisation picture what good looks like and continuously shape how AI is developed and used over time.

4. Workshop outline

An illustrative workshop outline designed to help use the tools to guide a discussion with Boards or leadership teams.

How to use it

Use this as a guide to structure a focused discussion - working through key prompts to reflect, challenge assumptions and identify practical implications for strategy, governance and next steps.

© Blueprint for Better Business 2026

Why Boards should engage

AI is changing how organisations operate - how decisions are made, how people work, and how value is created and distributed. For most organisations, that reshaping is happening fast, in multiple places at once, and largely without a coherent view of the whole.

Business leaders are often recognising their business models are under threat whilst watching competitors move at speed to adopt AI with a clear focus on maximising productivity or efficiency. All businesses need to be profitable to survive, but focusing on short term financial gains alone can create blind spots for the long-term implications of consequential decisions.

The result is that the most important questions go unasked:

- What does our use of AI and the choices we are making say about what kind of business we are, and what we are becoming?
- Are we deploying AI in ways that are consistent with our purpose and values — or simply in ways that are convenient or competitive?
- What are the cumulative effects on our people, our customers, the trust we have built with our stakeholders, and our long-term resilience?

These are not technical questions. They are questions about purpose, culture, and the future of the organisation. They belong to the Board.

The opportunity: a more resilient organisation

Engaging with AI at Board level is not only about managing risk. It is an opportunity to build a more resilient, more trusted, and more purposeful organisation.

Used well, AI can genuinely augment human capability, extend access to services, improve decision-making, free people for more meaningful work and help us to solve some of the most pressing social and environmental challenges we are facing as a society.

The question is whether your organisation is making the choices about design, deployment, governance, and culture that will produce those outcomes rather than their opposites.

The organisations most likely to thrive in an AI-enabled future are not those that adopted AI fastest. They are those that adopted it most thoughtfully: with a clear sense of what they were for, genuine attention to the people affected, and the foresight to shape the technology around their values rather than letting the technology shape them.



© Blueprint for Better Business 2026

1. Questions for Boards - enabling the right conversations

The questions which follow are a starting point and are deliberately broad. Some will be more pressing for your business than others. The goal is not to answer them all in one session, but to identify where the most important conversations need to happen and to ensure they do.

Purpose and long-term resilience

- How can our business be a force for good, and what is the role of AI in enabling or accelerating the good that we can contribute?
- What kind of business do we want to be in one, five and ten years' time, and are the AI choices we are making today helping us become that business, or pulling us in a different direction?
- Are we thinking seriously about the long-term effects of our AI strategy on our people, our sector, and wider society?
- Are we thinking carefully about who gains from our AI adoption and who bears the costs? Is the value created for us by AI flowing to people, our communities and the planet, and not only to our shareholders?

People and the planet

- Are we open, internally and externally, about where and how AI is involved in decisions and do we maintain clear human accountability for AI-generated decisions and outputs that affect people's lives?
- Are our employees genuinely involved in shaping how AI is used in their work and are we investing in the skills and understanding they need to work well alongside AI, including critical thinking and the judgement to know when not to use it?
- Are we satisfied that AI is not being used to exploit vulnerability and manipulate behaviour?
- Are we taking responsibility for the human cost of the AI systems and supply chains we rely on?
- Are we seeking to not only minimise the negative impact of the demands of AI on the planet but to actively use technology to solve some of our biggest sustainability challenges?

Governance and accountability

- Can we see how and where AI is used clearly enough to be able to provide meaningful oversight and have clear accountability lines been established for when something goes wrong?
- Are questions of fairness, dignity and human impact shaping our AI choices from the start, or are we trying to manage problems once they are visible?
- Do we conduct real impact assessments before deploying AI, particularly in contexts that affect people significantly?
- Do we maintain trusted processes for raising concerns about AI use without fear of consequence?
- Do we advocate for conditions that allow all organisations to adopt AI in a human centred way and engage in the policy conversation?

2. The Framework

An organisation is a product of peoples' thinking and doing. Ways of thinking and their subsequent impact on others can enable or limit what any organisation can achieve for itself and society, including through new technology. Emerging technology, including AI, offers profound and transformative potential for people, business and society as a whole. This framework aims to steer the design and implementation of technology in ways that promote the dignity and value of people and the common good.

Foundational understanding (mind-set)

Each person is a someone, not a something (Dignity and value of people)	Delivering value by serving society (The common good)
How are we thinking about deploying AI as a tool to **enable people to realize their full potential**, retaining their agency and autonomy, and not reducing people to passive recipients of what AI can deliver?	How are we considering **human flourishing, fairness and the good of society** as a whole in designing new products and services and the way we interact with all of our stakeholders, and not simply financial return?

Evolving thinking which shapes habitual behaviour

Other people matter (Solidarity)	Freedom with responsibility (Subsidiarity)	Building trust and trusted relationships (Reciprocity)	Valuing diversity and building bridges (Plurality)	Stewardship, values and resources (Sustainability)
How are we actively seeking to design or deploy AI systems to enhance livelihoods and **serve the broadest community?**	How are we creating the conditions to enable people **to develop wise judgement** and to **be accountable** for both their own decisions and those made by AI?	How are we seeking to enhance **social cohesion** and build **trusted relationships** through the use of AI?	How are we actively seeking **other perspectives**, inviting, considering and responding to **scrutiny and challenge**, internally and externally?	How are we seeking to use our influence and collaborating with others for the **benefit of all,** including **future generations** and **the planet?**
How are we valuing our employees and seeking to use AI to enhance their capabilities and work and better serve our customers without bias?	How are we ensuring that people retain appropriate accountability for decisions made by AI?	How are we intending to create mutual value with key stakeholders through implementing AI?	How are we consulting and encouraging others to participate in decisions on AI deployment?	How are we seeking to understand and minimize any negative impacts of our organisation's use of AI on people and the planet?
Are we only considering the financial case for using AI and seeking to automate as much as possible?	Are we simply seeking to automate decisions, without considering accountability?	Are we using our position to exploit AI for our self-interest alone?	Are we imposing our own views on AI deployment on others?	Are we using AI to maximize our profits leaving government to regulate for any negative impacts?

How to read the framework

The framework sets out questions for an organisation to consider at three levels of maturity – going from the bottom of the arrows and working up:

Human-centred mindset:
Considers the broader implications of how AI is used and how it might promote human flourishing, the good of society as a whole and the impact on the planet

Risk mitigation mindset:
Seeks to minimize the negative consequences from the use of AI, particularly on key stakeholders

Efficiency mindset:
AI is used with little or no regard for human and environmental consequences

© Blueprint for Better Business 2026

2. The Framework

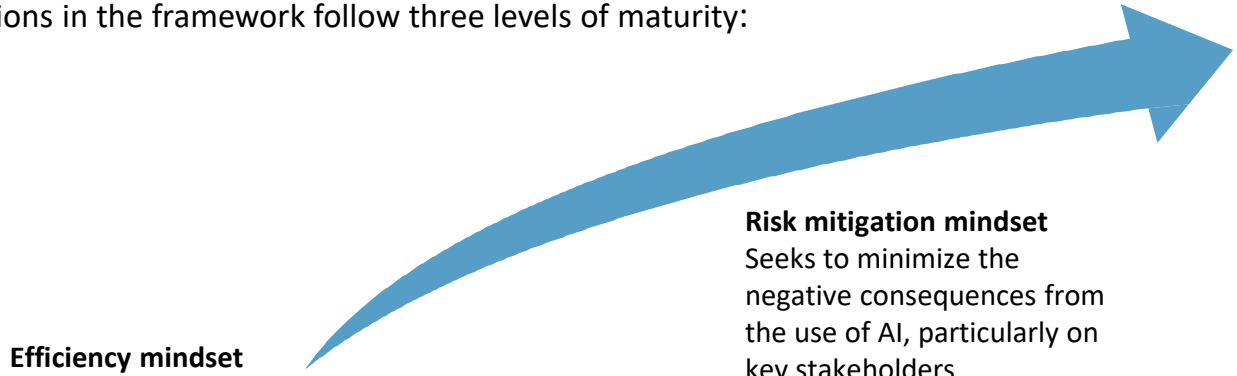
Guiding mindset – maturity levels

The framework is intended to help catalyse broader conversations, internally and externally, that explore the longer-term impact of AI on society, what that might mean for the future resilience of their business and inform a "guiding mindset" that can help organisations navigate towards a better future.

It is not a compliance framework, nor does it aim to address all the detailed considerations that need to be taken into account when implementing AI. The questions included are not intended to be an exhaustive list, but a way of starting discussions within businesses and across sectors. Those discussions should also involve voices from civil society, investors, government and regulators. It will take intentional effort from across society to steer a path through disruption towards a future that delivers long-term wellbeing for all people and the planet.

In having these discussions, it is of course important for organisations to recognise that many of the most significant impacts will come from their own buying power - the AI tools and platforms they purchase, the mandates they create for suppliers and how AI is integrated into their own customer-facing products. In the absence of global regulation on the development of AI, organisations can have a meaningful influence on how it is developed and deployed through what they choose to buy. Those organisations that spend time on these discussions now will be more likely to be trusted by stakeholders, more able to adapt to rapid change, and will be more resilient in an AI-enabled future.

The questions in the framework follow three levels of maturity:



Efficiency mindset
AI is used with little or no regard for human and environmental consequences

Risk mitigation mindset
Seeks to minimize the negative consequences from the use of AI, particularly on key stakeholders

Human-centred mindset
Considers the broader implications of how AI is used and how it might promote human flourishing, the good of society as a whole and the impact on the planet

2. The Framework

Underlying assumptions

In keeping with Blueprint's approach - challenging the underlying assumptions that shape business and society - we asked ourselves the following questions to help us reflect - not just on the impacts of AI, but on the assumptions that underpin our thinking:

- **When we ask how AI might impact the economy -** are we assuming the economy is simply a mechanism for efficiently delivering GDP? Or do we see it as a web of relationships that enable people to build meaningful lives, exercise agency, and form interdependent bonds with others?
- **When we ask what AI might do to jobs -** are we treating jobs as mere vehicles for wages—replaceable by Universal Basic Income in a future of AI-enabled abundance? Or do we recognize jobs as vital to human dignity, social contribution, and wellbeing, offering people a sense of purpose and belonging?
- **When we ask about the implications of AI for humanity -** are we viewing ourselves as rational, individualistic utility-maximisers? Or as complex beings with emotions, relationships, hopes and fears—driven not only by self-interest but also by concern for others?

We think that to truly realise the positive potential of this extraordinary technology, businesses must find the time and space to consider fundamental questions – ones that encourage a more imaginative, participatory and human-centered response that aims to promote the common good, human dignity and a sustainable relationship with the environment.

For instance:

- how might we use this technology to help us serve the broadest community and build societal cohesion?
- how might we build trusted relationships that we could not have imagined before through the use of this technology?
- how do we support and develop the capacity for critical thinking and wise judgement, and protect human agency and autonomy?

How we frame and answer these questions will fundamentally shape how this technology is implemented, and how the benefits, and costs, will be shared.



3. The Five Principles

Blueprint's Five Principles offer a picture of what a purpose-led organisation might look like if it were to implement the framework. Each section sets out aspirational outcomes, but these should not be seen as exhaustive or a compliance checklist. Instead, these principles should be used to continually challenge and help an organisation picture how they design and deploy AI systems and tools in service of the common good and human flourishing.

Honest and fair with customers and suppliers

- Is transparent when a customer or supplier is interacting with an AI agent
- Is open with suppliers about AI tools used in procurement or evaluation and works with them to ensure AI efficiencies are shared across the value chain
- Rejects the use of algorithms to manipulate or coerce through exploiting human psychology, personal data, or vulnerability
- Recognises that data belongs to the people whose lives it reflects, and gives individuals meaningful ability to understand and contest decisions made about them

A good citizen

- Uses its agency to advance public dialogue about AI to promote conditions that are good for society as a whole, including our collective understanding of truth, social justice and human flourishing
- Engages in broader societal efforts to address bias in AI systems and promotes access to opportunities for under-represented groups
- Takes responsibility for the full human cost of its AI systems, using its purchasing power and supplier relationships to drive up human rights standards across the supply chain and eliminate exploitation

Has a purpose which delivers long-term sustainable performance

- Implements an AI strategy which is designed to reinforce the purpose of the business and its human-centred approach and welcomes public scrutiny
- Maintains robust guardrails and effective boundaries to limit and govern its use of AI, including intervention and escalation capacity
- Measures success in terms of the long term value created for all stakeholders and wider society, alongside financial performance

A responsible and responsive employer

- Uses AI systems wherever possible to enhance capabilities and human autonomy and provide more meaningful work, rather than pursuing cost efficiency as the sole objective
- Supports people to develop the AI literacy and skills they need to participate confidently, supports any transition and rewards people fairly for the productivity gains achieved
- Engages employees in constructive and continuing dialogue about the best use of AI systems in the business
- Does not use AI to monitor or manage employees in ways that undermine their dignity, autonomy or sense of being trusted

A guardian for future generations

- Shares what it learns about AI use and invests in contributing to the development of standards and norms for AI that promote the common good and human flourishing
- Evaluates and considers the long-term impact of its AI business strategy on society and the future of work, particularly for young people
- Actively uses AI to solve the biggest sustainability challenges and minimises the demands of AI on energy and natural resources
- Does not treat the pace of AI adoption as a good in itself - actively asks whether slower, more considered deployment better serves its people, its customers,, the planet and wider society

4. Workshop outline – for use with Boards or leadership teams

Having had an initial discussion looking at the questions for Boards in this pack, this is an outline workshop to facilitate a deeper discussion by Boards or leadership teams using the Framework and Principles. Using an external facilitator could help you to get the most value out of the discussion, building on this outline to develop a bespoke version for your organisation.
If you would like to explore how Blueprint can support these discussions, please get in touch.

Time needed – at least 60-90 minutes:

Welcome, context & aims (5 mins)	This is not a technical session, but a leadership-level conversation about purpose, ethics, and organisational choices. exploring not just what we can do with AI, but how our approach to AI affects our existing purpose and commitments made to stakeholders and ultimately how it might help us become a better, more sustainable business.
Big picture: AI as a paradigm shift (10 mins)	Reflection (in pairs or small groups): • Where do we feel tension between speed, competitive pressure, and careful, human-centred deployment? • What is the outcome we are trying to move towards—and how can AI help?
Challenging underlying assumptions (10 mins)	Explore 3 questions: • Economy: are we assuming AI's role is GDP efficiency—or strengthening the web of human relationships? • Jobs: are jobs only wage mechanisms—or sources of dignity, contribution and belonging? • Humanity: are people rational utility-maximisers—or complex relational beings? Ask participants to choose the assumption that most influences decision-making in this organisation today. What would change if we adopted the alternative assumption?
An organisational assessment (5 mins)	Walk through the Framework's three levels - Where does our organisation currently sit? Where do we aspire to be? • Efficiency mindset • Risk mitigation mindset • Human centred mindset
Main Discussion (30-40 mins)	Break into groups – more detailed discussions referring to the Framework and the aspirational outcomes envisaged in the Principles: • What tensions (e.g., productivity vs. dignity, value extraction vs. proposition improvement) do we need to consciously manage? • What organisational capabilities do we need to steward AI wisely—judgement, critical thinking, ethical literacy? • How do we activate these capabilities by embedding this approach into processes and behaviours?
Implications for purpose, strategy, governance (10 mins)	• How is our approach to AI consistent with our corporate purpose and values? • How is our approach to AI in the interests of our stakeholders and have we ascertained their views? • How is our approach to AI making our organisation more sustainable as a business and a better business in the long term? • Does the Board have sufficient understanding of AI to provide meaningful oversight, and if not, what development is needed? • What governance mechanisms and controls (committees, review processes, principles) might be needed? • Where do we need more data or insight before deciding?
Closing reflection & next steps (5 mins)	What is the one commitment or question we will take forward from this session? Next steps?

Acknowledgements and references

This guide was created following a series of roundtables with businesses, academics and others and with input from a number of Chairs, CEOs and other senior business leaders. We would like to thank them all for their input and for taking the time to review earlier drafts.

Blueprint's original Framework and Principles on which the versions for AI are drawn can be seen here.

This guide builds on an earlier paper: Purpose for PLCs - Time for Boards to focus. which explores questions for a Board to enable it to fulfil its role as collective trustee of the corporate purpose of the business, such that it behaves - and is seen to behave - in a way that respects and cultivates all the relationships on which the long-term successful pursuit of that purpose depends, and through which the company makes its profits.

A Blueprint for Better Business is an independent charity whose purpose is to create a better society through better business. Our work is about stimulating and energising a different way of thinking and behaving in business. Our work is a public good and we encourage anyone to use it widely. This booklet and all of the information and resources on our website is free and publicly available – all we ask is that you reference Blueprint if you use specific tools or workshops.

SUPPORT OUR WORK

Our work is funded by donations, including from large corporates who see the value of our independence and catalytic role. **The window for shaping how AI is deployed – not just constrained – is open now. It will not remain open indefinitely. We are seeking funding to enable us to continue this important work during this critical time.**

If you would like to make a donation you can do so here.
If you would like to discuss how we might work together, donate via bank transfer and/or would like a receipt for your donation please contact us.

For more information about Blueprint contact:
Sarah@blueprintforbusiness.org or
Soulla@blueprintforbusiness.org

Blueprint Trust is a registered charity no: 1159150
Registered in England & Wales as a Company Limited by Guarantee no: 9006403

© Blueprint for Better Business 2026