

Atkinson, David (2026) AI, Judgement and Responsibility: Preserving Decision Authority in the Future Workforce. Other. House of Commons Business and Trade Committee Inquiry: Artificial Intelligence, Business and the Future of the Workforce, London.

Downloaded from: <https://ray.yorksja.ac.uk/id/eprint/15567/>

The version presented here may differ from the published version or version of record. If you intend to cite from the work you are advised to consult the publisher's version:
<https://committees.parliament.uk/writtenevidence/164623/pdf/>

Research at York St John (RaY) is an institutional repository. It supports the principles of open access by making the research outputs of the University available in digital form. Copyright of the items stored in RaY reside with the authors and/or other copyright owners. Users may access full text items free of charge, and may download a copy for private study or non-commercial research. For further reuse terms, see licence terms governing individual outputs. [Institutional Repositories Policy Statement](#)

RaY

Research at the University of York St John

For more information please contact RaY at
ray@yorksja.ac.uk

AI, Judgement and Responsibility: Preserving Decision Authority in the Future Workforce

Written Evidence Submission to the House of Commons Business and Trade
Committee Inquiry: Artificial Intelligence, Business and the Future of the
Workforce

Date: 02 April 2026

Submitted by: Dr David Atkinson, BEng (Hons), CDipAF, PhD, FHEA

Position: Lecturer, York St John University

Contact: d.atkinson@yorksja.ac.uk

Table of Contents

Executive Summary	3
1. Scope of this Submission	4
2. Basis of Evidence.....	4
3. Core Argument.....	4
4. The Future of AI (Q1).....	5
4.1 From Automation to Interpretation	5
4.2 Compression of Decision Space	5
4.3 Three Planning Scenarios	6
5. Impacts on Work and Workers (Q4).....	6
5.1 Redistribution of Responsibility Within Roles	6
5.2 Hidden Risk Transfer	7
5.3 Productivity and its Limits.....	7
5.4 Job Quality and Professional Judgement	7

5.5 Inequality Implications	8
6. Skills, Education and Transitions (Q5)	8
6.1 The Missing Skill: Judgement Under Constraint.....	8
6.2 Reframing AI Literacy	8
6.3 Implications for Education Systems	9
6.4 Reskilling and Upskilling	9
6.5 Job Redesign	9
7. Government Strategy, Regulation and Rights (Q6)	10
7.1 From System Transparency to Decision Transparency	10
7.2 The Limits of “Human-in-the-Loop”	10
7.3 Point-of-Decision Accountability	11
7.4 Policy Levers	11
7.5 Workforce Protections	11
8. Recommendations.....	11
9. Conclusion	12

Executive Summary

This submission addresses Q1, Q4, Q5 and Q6 of the Call for Evidence on Artificial Intelligence, Business and the Future of the Workforce. It argues that AI is not only automating tasks but increasingly shaping how decisions are framed, which options are considered viable, and how outcomes are determined.

The primary risk is therefore not simply job displacement, but a compression of judgement: workers and managers remain accountable for decisions while operating within AI-structured environments that limit meaningful choice. This creates a growing misalignment between decision authority and decision exposure, with implications for job quality, professional judgement, and organisational risk.

AI may increase productivity, but these gains may be offset where decision-making becomes overly narrow, errors propagate quickly, and workers lose the capacity to challenge outputs. This risks forms of “brittle efficiency”—systems that perform well under routine conditions but fail under uncertainty.

Policy should therefore shift from a primary focus on system performance to decision accountability in practice. In particular, Government should: ensure visibility of how AI-supported decisions are structured; align responsibility with genuine decision authority; strengthen education and training focused on judgement under uncertainty; support SMEs in responsible AI adoption; and redefine “human-in-the-loop” standards to require a real capacity to question and override AI outputs.

The effectiveness of AI adoption in the UK economy will depend not only on capability, but on preserving clear, accountable human judgement where decisions become consequential in AI-mediated workplaces.

1. Scope of this Submission

This submission addresses the following questions from the Call for Evidence:

- Q1: The future of AI
- Q4: Impacts on work and workers
- Q5: Skills, education and transitions
- Q6: Government strategy, regulation and rights

These areas are selected to focus on how AI is reshaping decision-making, responsibility, and workforce dynamics.

2. Basis of Evidence

This submission draws on:

- Professional experience in high-stakes decision environments, including prior service as a Royal Air Force Communications Engineering Officer working with UK and NATO operational systems;
- Current academic work in business and management education, focusing on decision-making under uncertainty and the integration of AI into teaching and assessment;
- Ongoing research into decision-making in AI-mediated environments, including the development of a framework examining how responsibility and judgement are exercised under conditions of constraint (Material Entropic Dynamics); and
- Early-stage applied work using AI-supported scenario exercises and decision simulations in educational and professional contexts.

Together, these inform a practice-oriented perspective on how AI is reshaping decision-making, responsibility, and workforce dynamics.

3. Core Argument

Artificial Intelligence is often framed in terms of automation, productivity, and efficiency. While these are important considerations, they do not fully capture the nature of the transformation underway. AI is increasingly not just performing tasks, but shaping how decisions are framed, how relevance is determined, and how options are presented to human decision-makers.

This shift has significant implications. The primary risk is not simply that tasks will be automated, but that decision-making environments will become increasingly structured by AI systems in ways that narrow perceived alternatives before a human decision is made. This can lead to a compression of judgement: individuals remain formally responsible for decisions, but operate within pre-structured environments that limit meaningful choice. This perspective is therefore relevant to current UK policy discussions on AI governance and workforce impacts, including recent Government and Parliamentary inquiries into AI adoption and regulation.

The central policy challenge is therefore not only how to support AI adoption, but how to preserve accountable human judgement at the point where decisions become consequential and irreversible.

This submission argues the following:

- AI is evolving from a tool of execution to a system that structures interpretation and decision contexts.
- This creates a growing misalignment between decision authority (who can act) and decision exposure (who bears consequences).
- Productivity gains may therefore coincide with a degradation of responsibility, job quality, and professional judgement.
- Policy, skills development, and regulation must focus on maintaining visibility of alternatives and clarity of responsibility at the point of decision.

4. The Future of AI (Q1)

In response to Q1, this submission argues that over the next decade, AI is likely to evolve along three interrelated trajectories:

4.1 From Automation to Interpretation

AI systems are increasingly capable of:

- summarising complex environments,
- identifying patterns of relevance,
- proposing courses of action,
- structuring decision workflows.

This represents a shift from task execution to interpretive mediation—AI is shaping how situations are understood before decisions are made.

4.2 Compression of Decision Space

As AI systems filter, rank, and prioritise information, they can:

- reduce cognitive load,
- but also reduce the range of options that are considered viable.

This creates a form of decision compression, where alternatives are excluded or deprioritised before human judgement is exercised. In practice, this can be recognised where AI systems reduce the number of options presented, prioritise a single recommended course of action, or make alternative choices operationally difficult to access or justify.

4.3 Three Planning Scenarios

Policy should consider three plausible scenarios:

4.3.1 Augmented Judgement (Desirable Scenario)

AI expands visibility of options, surfaces uncertainty, and supports deliberation. Human decision-makers retain genuine authority and responsibility.

4.3.2 Compressed Judgement (Likely Risk Scenario)

AI systems pre-structure decisions, narrowing options. Human actors remain nominal decision-makers but operate within constrained frameworks.

4.3.3 Delegated Decision-Making (Failure Scenario)

Responsibility becomes obscured as decisions are effectively made by systems, while accountability remains assigned to individuals or organisations.

5. Impacts on Work and Workers (Q4)

In response to Q4, this submission argues that current debates often focus on which jobs will be automated or augmented. However, a more fundamental transformation is occurring within roles themselves.

5.1 Redistribution of Responsibility Within Roles

AI systems are increasingly:

- shaping decisions upstream,
- while workers remain responsible for outcomes downstream.

This creates a growing misalignment between:

- decision authority (who determines the outcome), and
- decision exposure (who bears the consequences).

For example:

- A manager may rely on AI-generated recommendations for hiring or resource allocation;
- A frontline worker may follow AI-guided processes in healthcare, logistics, or customer service;
- In both cases, individuals remain accountable for outcomes, even where meaningful alternatives were not fully visible.

This dynamic is observable across multiple contexts. For example, in recruitment, AI systems may pre-filter candidates and rank suitability, leaving hiring managers formally responsible for final decisions made within a narrowed field of options. In healthcare, triage or prioritisation systems may shape treatment pathways before clinician judgement is exercised. In financial services, credit scoring systems may determine admissible lending decisions, while responsibility for outcomes remains with individuals or institutions. In each case, decision environments are structured upstream, while accountability remains downstream.

5.2 Hidden Risk Transfer

This misalignment can result in:

- transfer of risk from systems to individuals,
- increased legal and professional exposure for workers,
- reduced ability to contest or override AI-influenced decisions.

Over time, this may:

- erode trust in organisational systems,
- discourage initiative and professional judgement,
- create defensive or compliance-driven behaviour rather than responsible decision-making.

5.3 Productivity and its Limits

AI can deliver productivity gains through:

- faster processing,
- improved coordination,
- reduced duplication of effort.

However, productivity gains may be offset if:

- decisions are made on overly narrow grounds,
- errors propagate quickly through automated systems,
- workers lose the capacity to identify when AI outputs are inappropriate.

There is therefore a risk of “brittle efficiency”: systems that perform well under expected conditions but fail under uncertainty. This may be observed where systems perform efficiently under expected conditions but degrade rapidly when confronted with novel, ambiguous, or unforeseen situations.

5.4 Job Quality and Professional Judgement

Job quality is affected not only by pay and security, but by:

- the ability to exercise judgement,

- clarity of responsibility,
- the capacity to influence outcomes.

Where AI systems:

- reduce discretion,
- obscure reasoning,
- or limit alternatives,

workers may experience:

- reduced autonomy,
- diminished professional identity,
- increased stress due to responsibility without control.

5.5 Inequality Implications

The effects of AI are unlikely to be evenly distributed:

- Lower-paid and routine roles may experience increased monitoring and constraint;
- Professional roles may retain more discretion but face heightened accountability;
- SMEs may adopt AI tools without sufficient governance structures, increasing risk exposure.

This could widen inequalities not only in income, but in control over work and exposure to risk.

6. Skills, Education and Transitions (Q5)

In response to Q5, this submission argues that the dominant policy focus on technical skills—such as coding, data analysis, and AI literacy—is necessary but insufficient.

6.1 The Missing Skill: Judgement Under Constraint

Workers increasingly need the ability to:

- recognise when alternatives remain available,
- question how AI outputs have been generated,
- identify when decision spaces have been prematurely narrowed,
- exercise discretion in the presence of uncertainty.

These are not purely technical skills. They are forms of judgement.

6.2 Reframing AI Literacy

AI literacy should include:

- understanding how AI systems prioritise and filter information,
- recognising limitations, biases, and uncertainties,
- knowing when to rely on AI and when to challenge it.

This requires moving beyond tool usage towards critical engagement with decision environments. This aligns with emerging policy discussions on critical digital skills and AI literacy, but places greater emphasis on the capacity to interrogate and resist inappropriate system outputs.

6.3 Implications for Education Systems

Current education and training systems are not fully adapted to this shift. In particular:

- Curricula often separate technical knowledge from decision-making practice;
- Assessment systems may incentivise correct answers rather than responsible judgement;
- Professional education may not adequately address AI-mediated decision contexts.

There is a need to integrate:

- scenario-based learning,
- decision-making under uncertainty,
- interdisciplinary approaches combining technical, ethical, and practical reasoning.

6.4 Reskilling and Upskilling

Effective models should:

- be embedded within workplace contexts,
- focus on real decision scenarios rather than abstract training,
- support workers in understanding how AI affects their specific roles.

For SMEs in particular:

- accessible, modular training frameworks are needed,
- alongside guidance on governance and responsible AI use.

6.5 Job Redesign

Employers introducing AI should:

- actively redesign roles to preserve meaningful human contribution,
- ensure workers retain the ability to question and override AI outputs,
- align authority with responsibility.

Without this, AI adoption risks creating roles that are:

- operationally efficient but strategically and ethically fragile.

7. Government Strategy, Regulation and Rights (Q6)

In response to Q6, this submission argues that current regulatory discussions focus on:

- bias,
- transparency,
- fairness.

These are essential, but do not fully address the structural effects of AI on decision-making.

7.1 From System Transparency to Decision Transparency

It is not sufficient to understand how an AI system works in isolation. What matters is how decisions are made in practice.

Policy should therefore prioritise:

- decision traceability: how options were generated, filtered, and selected;
- visibility of alternatives considered and excluded; and
- clarity on who had the authority to make or override decisions.

These concerns complement existing debates on automation bias, algorithmic management, and meaningful human control, but extend them by focusing specifically on how decision environments are structured prior to action.

7.2 The Limits of “Human-in-the-Loop”

The concept of a “human-in-the-loop” is often treated as a safeguard. However, this is insufficient if:

- the human is presented with only one viable option,
- time pressures or system design discourage challenge,
- responsibility is assigned without meaningful control.

Regulation should therefore ask:

- Does the human decision-maker have genuine alternatives?
- Is there a real capacity to say no?

Meaningful human oversight requires, at a minimum: sufficient time to review decisions; visibility of alternative options; the authority to challenge or reject system outputs; and organisational support for exercising that judgement.

7.3 Point-of-Decision Accountability

A key principle for governance should be point-of-decision accountability:

- identifying where decisions become consequential,
- ensuring responsibility is aligned with authority at that point,
- preventing situations where individuals bear responsibility without control.

7.4 Policy Levers

Government can support this through:

- Standards and guidance on responsible AI deployment in workplaces;
- Procurement requirements that prioritise decision transparency and human oversight;
- Support for SMEs in implementing governance frameworks;
- Regulatory clarity on liability in AI-mediated decisions.

7.5 Workforce Protections

Employment frameworks may need to evolve to address:

- AI-influenced decision-making in recruitment and management,
- the right to challenge or review AI-assisted decisions,
- protection against undue responsibility for system-driven outcomes.

8. Recommendations

- **Adopt a decision-focused approach to AI policy**
Move beyond capability and performance to consider how AI shapes decision-making environments.
- **Ensure visibility of decision processes**
Require organisations to demonstrate how options are generated, filtered, and selected in AI-supported decisions.
- **Align authority with responsibility**
Avoid situations where workers bear consequences without meaningful control over decisions.
- **Strengthen judgement-focused skills development**
Integrate decision-making under uncertainty into education, training, and professional development.
- **Redefine “human-in-the-loop” standards**
Ensure human oversight includes genuine capacity to question, challenge, and override AI outputs.
- **Support SMEs in responsible AI adoption**
Provide practical frameworks and resources for governance, training, and risk management.

- **Develop regulatory approaches centred on decision accountability**
Focus on where decisions become consequential, not only on how systems operate.

9. Conclusion

AI presents significant opportunities for productivity and innovation. However, its most profound impact may lie in how it reshapes decision-making itself.

If AI systems narrow the range of perceived options before decisions are made, there is a risk that responsibility becomes detached from control. Workers and organisations may remain accountable for outcomes without having exercised meaningful judgement.

The long-term success of AI adoption in the UK economy will therefore depend not only on technical capability, but on maintaining clear, accountable, and responsible decision-making in increasingly AI-mediated environments.

Preserving this capacity is essential for:

- economic resilience,
- workforce wellbeing,
- and public trust in AI-enabled systems.

A related position paper outlining the underlying research framework is available via Zenodo: Atkinson, D. (2026). Material Entropic Dynamics: Why Faster Prediction Cannot Deliver Freedom in the Age of AI. Zenodo. <https://doi.org/10.5281/zenodo.18902944>.