OBO (On-Behalf-Of) Governance:

Governing Artificial Intelligence Through Delegated Authority

Author: Neil Hare-Brown for MIFAIR
Submitted to: MIGF
Date: 30 July 2026

This paper proposes a governance model. It is not intended to constitute legal advice or to express legal opinion.

Authors Note

This paper is written from the perspective of a practitioner in cyber risk management, organisational governance and incident response, rather than from that of a practising lawyer. It does not seek to reinterpret or extend the law. Instead, it asks whether long-established legal concepts - particularly agency, delegated authority and acting on behalf of a principal - provide a useful foundation for governing increasingly autonomous AI within organisations.

The legal observations contained in this paper are informed by publicly available legal authorities and consultations. The central contribution of the paper is not a new legal theory, but a proposed governance model intended to complement existing legal and regulatory frameworks.

Abstract

This paper introduces the concept of the **On-Behalf-Of (OBO) State**, a governance relationship in which an AI system acts under explicitly delegated authority from an accountable principal.

Artificial Intelligence is increasingly capable of undertaking actions previously performed exclusively by humans. Existing governance frameworks predominantly treat AI as a software tool requiring oversight.

While this approach remains appropriate for decision-support systems, it becomes increasingly inadequate as organisations deploy autonomous and agentic AI capable of negotiating, authorising, communicating and executing business processes.

The defining governance challenge of Artificial Intelligence is not that machines have become more intelligent. It is that every employee has acquired an unprecedented capability to act on behalf of their organisation. Existing AI governance treats AI as technology.

This paper argues that organisations must instead govern AI as the exercise of delegated organisational authority through a computational proxy. It proposes a complementary governance model founded upon the well-established legal doctrine of delegated authority.

Rather than treating AI solely as technology, organisations should recognise AI systems as **proxies acting on behalf of identifiable principals**. Under this model, authority is delegated - not transferred - and accountability remains with the principal(s) who authorise the AI's actions.

This paper introduces the **OBO Principle** and an associated governance model that is compatible with existing AI standards while addressing a gap concerning delegated organisational authority.

1. Introduction

The adoption of Artificial Intelligence has accelerated rapidly. AI increasingly combines computational capability with functions previously reserved for human expertise.

Initially, AI functioned primarily as a productivity tool. Many organisations still see it exactly as that; a way to improve productivity – in some cases, by orders of magnitude.

Increasingly, however, organisations are deploying AI that is more agentic and:

- negotiates with suppliers,
- writes software,
- authorises transactions,
- communicates with customers,
- makes purchasing decisions,
- schedules resources,
- drives vehicles,
- executes workflows.

This list is by no means complete and is growing by the day.

These AI systems are no longer passive analytical tools. They are no longer software in the traditional sense.

They perform actions and mimic human capabilities, often with remarkable, expert-level capability. As such, they should be considered a hybrid of traditional silicon-powered applications and capabilities previously associated primarily with human expertise.

Existing governance strategy - both that proposed (and, in some cases adopted), for AI and that which applies to controlling human actions - therefore requires reconsideration.

One of the key questions raised in a paper entitled "Artificial intelligence: ethics, governance and regulation" written by Patrick Brione and Devyani Gajjar in 2024 is:

"*Who should be held legally accountable when AI systems violate these legal principles? What should the balance of liability be between the original developer, the organisation employing or disseminating the tools, and the end user(s)?*"

2. The Current Governance Model

The Current Governance Model

ISO/IEC 42001, the NIST AI Risk Management Framework and the EU AI Act concentrate upon:

- governance,
- transparency,
- explainability,
- bias,
- robustness,
- accountability,
- human oversight.

These are essential.

Existing governance frameworks are built upon an implicit assumption:

"AI is a tool."

This assumption becomes increasingly strained as AI becomes agentic.

This paper argues that this characterisation is no longer adequate.

Throughout history, organisations have adopted increasingly powerful tools: the telephone, spreadsheet, internet, search engine and smartphone each improved productivity but remained passive instruments. They extended the capabilities of individuals without fundamentally altering the nature of organisational authority.

Artificial Intelligence is different. AI remains software from a technical perspective, but it can no longer be governed solely as software. It is a cognitive capability that can reason, generate, explain, plan, synthesise, translate, negotiate, code and increasingly execute actions across digital systems.

The distinction is profound. Unlike previous technologies, AI can dramatically amplify the effective capability of almost any individual. An inexperienced employee can produce software, draft legal correspondence, analyse financial statements, interpret technical standards or generate strategic recommendations that would previously have required years of specialist training. Equally, an individual acting recklessly or dishonestly can use the same capability to create convincing misinformation, fraudulent communications, malicious software or deceptive content.

The technology itself is indifferent. Its effect depends entirely upon the authority under which it is used. Consequently, every employee and contractor with access to AI, whether through an approved organisational platform or through unauthorised "shadow AI", possesses a capability that can materially influence customers, suppliers, regulators and colleagues.

Every AI-assisted output that is relied upon, communicated or used in organisational activity has the potential to represent the organisation.

Historically, organisations have exercised control over representation by controlling expertise. Specialist knowledge was concentrated within qualified professionals, subject matter experts and experienced managers. AI fundamentally alters this balance by making sophisticated reasoning and content generation available across the entire workforce.

The effect of this revolution in capability can be likened to increasing your workforce by hiring subject matter experts to be available 'on-demand' to every member of existing staff and expecting every member of staff – with little to no guidance – to be able to use that capability responsibly.

In the scenario described here, just like in the past and regardless of each member of staff's personal sensibilities, they remain representatives of their organisation and existing law describes the liabilities that flow from that reality.

So, the question therefore is no longer, *"How should employees use AI responsibly?"*

The question becomes, **"How should an organisation govern thousands of AI-assisted acts of representation undertaken on its behalf every day?"**

Existing governance frameworks concentrate on the risks posed by AI systems.

This paper proposes that organisations must also govern the relationship between the organisation, the individual and the AI itself.

The central challenge is therefore not simply technological governance. It is the governance of delegated organisational authority.

3. Existing Legal Doctrine

A January 2026 consultation issued by the UK Jurisdiction Taskforce examined how existing English private law is likely to apply to harms arising from AI. Its draft Legal Statement argues that established common-law principles can address many AI-related harms without requiring an entirely new liability regime.

However, while this paper answers the question of who is legally liable for a range of harms that may arise out of the use of AI, what this paper seeks to answer is "**How should organisations govern AI so that delegated authority, operational responsibility and accountability are clearly established before harm occurs?**"

So, when it comes to whether existing legal doctrine will help us to define the governance strategy for AI, the answer may well already exist.

Most legal systems recognise the doctrine of **agency**. Agency concerns relationships in which one party acts on behalf of another.

Key other concepts include:

Principal: The individual or organisation possessing legal authority.

Agent (or Proxy): A person authorised to exercise delegated authority on behalf of the principal.

Delegated Authority: Authority expressly or implicitly granted by the principal.

Fiduciary Responsibility: The obligation to act within the authority granted.

Revocation: The authority may be withdrawn at any time.

Accountability: The principal remains legally responsible for acts undertaken within the delegated authority.

These concepts have existed for centuries and should be equally useful when defining strategy for the governance of AI.

4. The Missing Governance Concept

When an employee instructs an AI to:

- negotiate a contract,
- approve expenditure,
- communicate with customers,
- modify production systems,

what is the legal relationship between the employee and the AI?

Current standards do not answer this question explicitly.

5. AI As Proxy

This paper proposes that organisational AI should be governed using these same principles.

Under this model, an AI possesses:

- no legal personality,
- no independent organisational authority,
- no formal corporate mandate of its own,
- no fiduciary duty.

Instead:

The AI operates solely as a computational proxy exercising authority delegated by an identifiable principal within an approved Authority Profile.

Authority therefore originates from the principal - not from the AI.

6. Two Distinct Modes of AI Operation

Not every interaction with Artificial Intelligence represents the exercise of organisational authority.

There is an important distinction between AI being used as an **individual cognitive capability** and AI acting as an **organisational proxy**. This paper therefore distinguishes between two governance modes.

Capability State

In the Capability State, AI operates as an advanced cognitive assistant to an individual.

Typical activities include:

- research,
- drafting,
- brainstorming,
- summarisation,
- translation,
- coding assistance,
- learning and professional development.

During the Capability State, AI supports the user's thinking but does not exercise organisational authority or formally represent the organisation.

OBO State

The On-Behalf-Of (OBO) State begins when AI is authorised to perform, assist with or materially influence an activity that represents the organisation or exercises delegated organisational authority.

Typical activities include:

- communicating with customers,
- negotiating contracts,
- approving expenditure,
- authorising transactions,
- generating regulatory submissions,
- publishing organisational statements,
- making operational decisions,
- controlling technical systems.

Once the OBO State is entered, the AI ceases to operate merely as an individual productivity capability and instead functions as an organisational proxy.

The governance requirements described throughout the remainder of this paper apply specifically to AI operating within the OBO State.

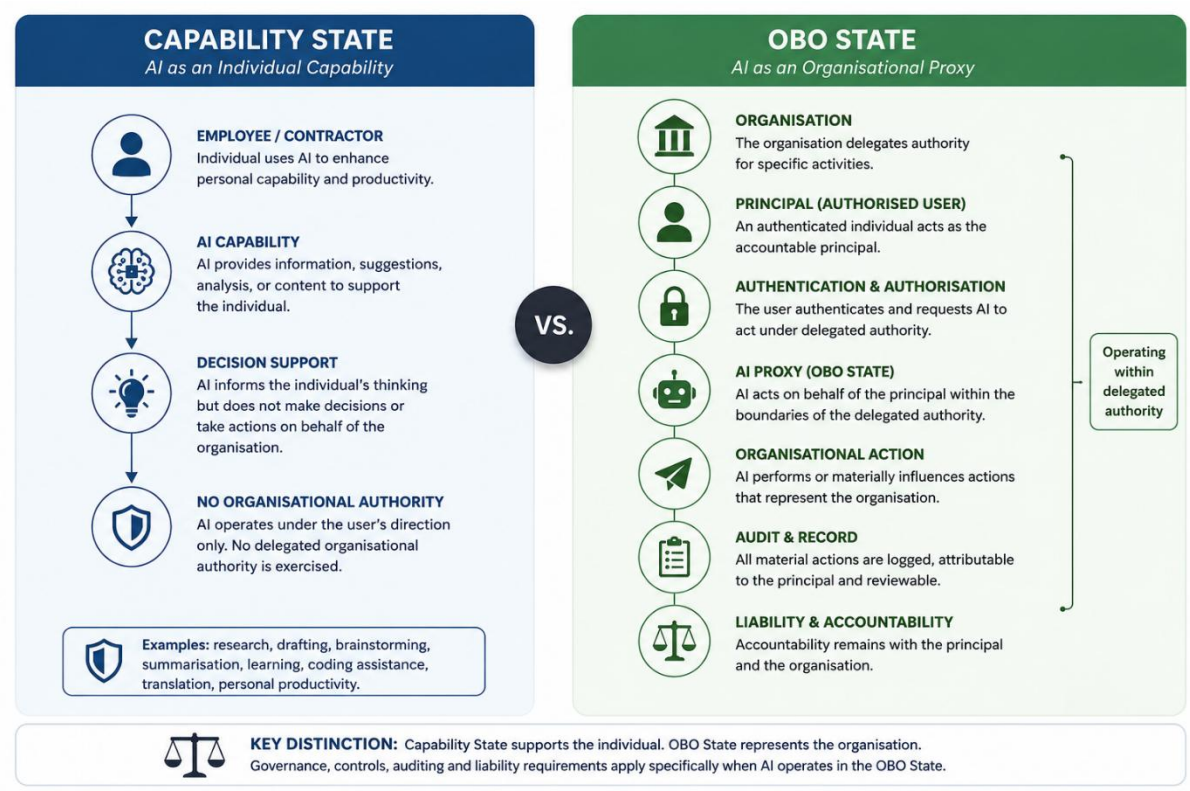

7. The OBO State

This paper introduces the concept of the **On-Behalf-Of (OBO) State**.

The OBO State is entered whenever an authenticated user authorises an approved AI system to undertake an activity that exercises delegated organisational authority or formally represents the organisation. It may be activated automatically through organisational workflow, by explicit user selection, or by system policy based on the nature of the task being undertaken.

Once activated:

- the organisation remains the principal from which authority originates,
- the authenticated individual acts as the authorised user responsible for invoking and supervising that authority,
- the AI operates as a computational proxy,
- every material action is attributable to the authenticated user and the organisational principal,
- the authority exercised by the AI remains bounded by the user's Authority Profile and organisational policy,
- accountability for actions undertaken during the OBO State survives termination of the session.

While the two states, both Capability and OBO describe different modes of AI engagement, organisations may consider that requiring compliance with only the OBO state in their governance programmes would be simpler and less ambiguous.

8. Authority Profile

Every approved AI system operating within the OBO State shall be subject to an Authority Profile defining, as applicable:

- financial authority,
- contractual authority,
- procurement authority,
- data-access and disclosure authority,
- communication and publication authority,
- operational and technical authority,
- approval thresholds,
- prohibited activities,
- escalation requirements,
- duration, expiry and revocation conditions.

This extends to AI the same principles that well-governed organisations already apply when defining delegated authority for employees, officers and contractors.

9. OBO Roles

Roles for operating the OBO Governance model include:

- **Organisational Principal:** The legal entity on whose behalf authority is exercised.
- **Authorised User:** The authenticated employee or contractor who invokes and supervises the AI.
- **AI Proxy:** The AI system exercising delegated authority within the applicable Authority Profile.

10. Benefits

The OBO Governance model provides:

- **Clear accountability**: Responsibility never becomes ambiguous.
- **Familiar governance**: Boards already understand delegated authority.
- **Better auditability**: Every AI action maps directly to a principal.
- **Better insurance**: Insurers can assess authority rather than attempting to define AI liability independently.
- **Better regulation**: Existing legal principles require minimal adaptation.

11. Relationship to Existing Standards

OBO Governance complements:

- ISO/IEC 42001,
- ISO/IEC 23894,
- NIST AI RMF,
- COBIT,
- COSO,
- Corporate Governance Codes.

It introduces an additional governance layer which we term **Authority Governance** rather than replacing existing risk governance.

12. Future Development

The paper proposes further work on:

- AI Authority Registers,
- Proxy Certificates,
- Delegated Authority Profiles,
- OBO Session Management,
- AI Authority Audit Trails,
- AI Revocation Mechanisms.

13. Conclusion

Artificial Intelligence represents a significant change not merely because it produces information but because it increasingly performs actions.

AI does not democratise knowledge; it democratises capability.

Governance should therefore evolve beyond treating AI solely as software or tools.

The long-established legal doctrines of agency, proxy and delegated authority provide a mature and widely understood foundation upon which AI governance can develop.

Recognising AI as a **proxy acting on behalf of an accountable principal** preserves existing principles of responsibility whilst providing organisations with a practical framework for governing increasingly autonomous AI systems.

The question is no longer whether organisations should govern AI. It is whether they can continue to govern AI effectively without explicitly governing the authority it exercises on their behalf.

Just as organisations define the authority of every employee, officer and contractor, they must now define the authority of every AI operating on their behalf.

Limitations

This paper proposes a governance framework rather than a legal doctrine. It is intended to stimulate discussion among governance professionals, regulators, standards bodies, legal practitioners and technologists. The author welcomes legal scrutiny of the proposed OBO model, particularly regarding its interaction with agency law, corporate governance and emerging AI regulation.

References

(Key policy, legal and governance sources informing this paper)

Department for Science, Innovation and Technology (2024). *Implementing the UK's AI Regulatory Principles: Initial Guidance for Regulators*.

UK Jurisdiction Taskforce (2026). *Consultation on the Legal Statement on Liability for AI Harms under the Private Law of England and Wales*.

European Union (2024). *Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence*.

ISO/IEC 42001:2023. *Artificial Intelligence — Management System*.

ISO/IEC 23894:2023. *Artificial Intelligence — Guidance on Risk Management*.

ISO/IEC 38507:2022. *Governance Implications of the Use of Artificial Intelligence by Organizations*.

NIST (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*.

OECD (2024). *Recommendation of the Council on Artificial Intelligence*.

UK Government (2023). *A Pro-Innovation Approach to AI Regulation*.

Brione, P., & Gajjar, D. (2024). *Artificial Intelligence: Ethics, Governance and Regulation*. UK Parliament, POST.