KNOWA WHITE PAPER

The Governance Imperative

Why written AI policy is not the same as governed AI use, and what care-architected technology must do to turn one into the other inside the rooms where decisions are made.

Knowa

From policy on paper to governance in the room

AI is already in the trustee meeting. Not because any board chose to deploy it, but because consumer tools are now one prompt away from every laptop in the room. The question is no longer whether AI belongs in trustee work. It is whether the board can show how it was used, and on what.

Most boards govern AI the same way they govern much else: with a written policy. An acceptable-use note, a line in the business plan, a paragraph in an adviser's terms. These matter, and the General Code expects a board to have thought the question through. But a written policy is a statement of intent. It does not, on its own, control what scheme data a tool can see, who is allowed to use it, or whether an answer can be traced back to the document it came from. For most boards, those controls do not yet exist.

This paper draws a line between two kinds of governance. Written governance sets out what the board expects. Operational governance builds those expectations into the tools themselves, so that the policy holds whether or not anyone in the room remembers it on the night. It explains what governance means under trustee duty, why the gap between the written rule and the working control matters now, and how a board can put a workable baseline in place.

> Without governance, AI use does not stop. It simply happens unrecorded, in tools that were never built for fiduciary work.

This is not an argument for caution over progress. It is the opposite. Governance is what lets a board use AI with confidence, evidence its decisions to the sponsor and the Regulator, and serve members better, rather than quietly tolerating ungoverned tools because no one has decided what to allow. The boards that move forward are not the ones with the cleverest tools. They are the ones whose controls let them use any tool at all.

SECTION 01

What governance means in a fiduciary setting

Governance is the system of controls, decision rights, and oversight that determines how AI is used: by whom, for what purposes, with what data, and within what limits. Under trustee duty, that system carries a charge most settings never face. It must protect the confidentiality of scheme and member data across every adviser relationship, preserve the independence of trustee judgement, and keep pace with the board's obligations under the Pensions Acts and the General Code.

A useful distinction runs through the rest of this paper: the difference between governance that is written and governance that is operational.

Written governance sets out what the board expects

Written governance is the layer most boards already have: an acceptable-use policy, a note in adviser terms, a paragraph in the business plan, perhaps a slot for AI on the risk register. It sets expectations and helps evidence that the board has turned its mind to the question. But it relies on every trustee and adviser reading it, remembering it, and applying it correctly under time pressure, in a meeting room, the night before papers are due.

Operational governance puts it into effect

Operational governance builds those expectations into the tools themselves: access scoped to the scheme by design, answers cited back to the source document by default, and a record of what was asked and answered kept whether or not anyone thinks to keep it. It does not depend on memory. The policy is the architecture.

THE PRACTICAL TAKEAWAY

Once a board has a written policy, the next step is to put the controls behind it that make it hold in practice. The written policy describes the standard. Operational governance is the standard, built into the tools the board actually uses.

SECTION 02

Turning policy into practice

The gap between the written rule and the working control matters most at the moment it is hardest to see: when AI use is already widespread but quietly informal. A trustee pastes a consultant's recommendation into a chatbot to sense-check it. An adviser runs a draft paper through Copilot to tighten the language. A scheme secretary asks a generic tool to summarise a Regulator update before the meeting. None of it is governed, and a written policy on its own changes none of it.

Each party around the board needs a different assurance, and a policy document gives none of them evidence. Trustees need to know which uses are permitted and that what they rely on is grounded in the scheme's own record. The sponsor and members need confidence that confidential data is protected. The chair needs to be able to account for how AI was used when the Regulator asks. The point at which these questions arrive in earnest is the point at which a written policy alone stops being enough.

> The boards that move forward are not the ones with the cleverest tools. They are the ones whose controls let them use any tool at all.

The starting point is modest. Make AI use in the scheme visible. Centralise the basic controls over who can access what. And keep a reliable record of how the tools are being used across the board and its advisers. None of this is advanced. It is the foundation that makes everything else possible.

WHERE GOVERNANCE BEGINS

Three foundations: visibility over how AI is used, control over who can access the scheme's data, and a clear record of both. Put these in place first. The rest follows.

SECTION 03

The evolving risk landscape

Governance has to address a set of risks that are familiar in fiduciary practice, but in a way that enables a board to use AI at scale rather than retreat from it. Four stand out.

> **Data risk and confidentiality.**
>
> Scheme records, member data, and privileged advice must be accessed, used, and shared only within defined boundaries. Consumer AI tools were never designed to honour those boundaries.

> **Accuracy and output integrity.**
>
> A confident summary is not a correct one. The board remains responsible for reviewing and verifying the outputs it relies on, which means the tool must make verification cheap rather than ceremonial.

> **Access and authorisation.**
>
> Boards and advisers collaborate across organisations constantly. Maintaining appropriate access limits while still collaborating, and being able to revoke access cleanly, is a governance problem long before it is a technical one.

> **Compliance and regulatory expectations.**
>
> The regulatory landscape is moving quickly toward transparency, accountability, and documentation. A board that cannot evidence how AI was used is a board that cannot evidence how a decision was made.

None of these risks is new to fiduciary work. What is new is that ungoverned AI amplifies each of them quietly, without anyone deciding that it should. The risk is not that a board chooses badly. It is that the choice is never consciously made at all.

SECTION 04

Who has to be able to answer for it

Using AI well on a trustee board is not one person's job. The responsibility is shared between the people who run the scheme day to day and the people ultimately answerable for its decisions, and each of them needs a different kind of assurance about how AI is being used.

- Trustees need confidence that when they lean on a summary or a search, the output is grounded in the scheme's own record and traceable back to source, so that relying on it is consistent with their duty rather than a shortcut around it.
- The scheme secretary and in-house team need practical control over who can see what, so that access tracks each adviser, project, and relationship cleanly, and can be withdrawn the moment it should be.
- The chair, and professional trustee firms across a portfolio, need a clear, evidenced view of how AI is being used in the scheme's name, because they are the ones who must account for it to the sponsor and the Regulator.

As reliance on AI grows, the ability to bring visibility, control, and a clear record together stops being a governance nicety. It becomes the board's answer when the sponsor, The Pensions Regulator, or a member asks how a decision was actually reached.

> When the Regulator, the sponsor, or a member asks how AI shaped a decision, the board should be able to answer with the record, not a recollection.

SECTION 05

Oversight that works in practice

As a board takes stock of where it stands, a useful starting point is to ask where the gaps sit between what its policy says and what its tools actually enforce. The areas below are the ones boards most often find themselves weighing.

Who can see what, and only what they need

Can access be adjusted cleanly as the board, its advisers, and its projects change, without opening gaps or relying on a favour from whoever set the system up? Each person should see only what their role on the scheme requires, no more, and access should follow them as their involvement begins and ends. A new adviser is added without inheriting the whole history; one who steps away loses access the same day.

Working with advisers without losing the boundary

Most trustee work happens across organisations: the board, the consultant, the actuary, the legal adviser, the administrator. This is exactly where informal sharing, an emailed document, a shared login, quietly defeats the policy. The question is whether collaboration runs through a governed channel where the data stays inside the scheme's boundary and access can be withdrawn when a relationship ends, rather than through copies the board can never call back.

A record that stands up to scrutiny

When the sponsor or The Pensions Regulator asks how AI was used in reaching a decision, can the board answer from the record? Sound governance moves the board from a rough sense of what happened to a clear, retrievable account: what was asked, what the tool drew on, and how that fed the decision. That is the difference between recollection and evidence.

Knowing where AI sits in the scheme's work

A board does not need a usage dashboard. What it needs is the assurance that AI is being used where it should be, on the scheme's own record rather than on documents pasted into a public tool, and that nothing is happening in the dark. Even a simple view of where AI touches the work lets the board confirm the policy is being followed, and surface the quiet, ungoverned uses before they become the norm.

Controls strengthen the practices a board already has

Operational controls do not replace the things a well-run board already does: the risk register, trustee training and knowledge requirements, the work of any sub-committee, the terms it agrees with advisers. They make those practices stronger by tying them to evidence. A board that can see how a tool was actually used governs better than one relying on assurances after the fact.

Visibility

Where AI touches the scheme's work, and on what data. The starting point for every other control.

Control over access

Who can see the scheme's records, decided once and applied consistently across every adviser.

Citations to source

Every answer traceable to the document behind it, so the board can check rather than trust.

A defensible record

An account that answers the question before it is asked: how was this decision actually reached.

Each of these is modest on its own. Together they are the difference between a policy that exists and a policy that holds.

SECTION 06

Building the foundation now

AI will keep changing, and what a board is expected to govern will widen with it. The principles will not change, but the need for clarity over what AI may touch, who may use it, and how its use is recorded will only grow. A few things help a board stand on firm ground as that happens.

Start with the controls behind your most important rules.

No board needs everything at once. Put in place the few controls that address the clearest risks and make the policy real: knowing where AI is used, controlling access to scheme data, and keeping a record that holds up.

Choose tools that were built for this work.

A tool that scopes access to the scheme, cites its sources, and keeps confidential papers inside its boundary is governing by design. A general consumer tool leaves the board to govern everything by hand, every time.

Be ready to show your working.

A board that can explain, with the record in front of it, how AI was used in reaching a decision is a board that can reassure members, satisfy the sponsor, and answer the Regulator without scrambling.

Treat governance as what lets you proceed, not what holds you back.

Done well, governance does not slow the board down. It is what lets it use AI at all, with confidence rather than quiet unease.

> The standard a board sets for itself now is the standard it will be judged against later: in the record, and in how well members were served by it.

SECTION 07

Where to start

1. **Take stock of where you stand.** Set your written rules on AI against what your tools actually enforce. Which expectations are backed by a real control, and which rely on everyone simply remembering? The gap between the two is where to begin.

2. **Decide what good looks like, and write it down.** Agree clear answers to the questions the Regulator, the sponsor, or a member could ask about how AI is used in the scheme's work. If the honest answer today is "we'd have to recall," the work is not yet done.

3. **Close the gap with the right tools, in order of risk.** Address the clearest exposures first: scheme data leaving a governed boundary, answers that cannot be traced to source, use that leaves no record. Make sure the board's reach never runs ahead of its control.

Governing AI well is not a separate project bolted onto trustee work. It is part of how a board discharges its duty in a world where the tools are already in the room. A policy without a control is a hope. A control without a policy is an accident. Governance is what binds the two.

A board that puts the controls behind its rules now is one that can use AI without surrendering its judgement, and answer for that use with confidence. Building the baseline now is what makes a board ready for the capabilities, and the questions, that follow.

THIS IS WHAT KNOWA IS BUILDING

Knowa Q runs across the platform as the layer that makes governance real: access scoped to the scheme, a full record of what was asked and answered, citations back to the source document, and a boundary that confidential papers never cross. It surfaces the record. The judgement stays with the trustees and their advisers.

Knowa

REFERENCES & FURTHER READING

Sources and further reading

A short bibliography for readers who want to follow the human-factors and governance references behind this paper.

Human-factors research

1. LaCour, M., Cantú, N. G. & Davis, T. (2019). When calculators lie: a demonstration of uncritical calculator usage among college students. PLOS ONE, 14(10): e0223736.
2. Parasuraman, R. & Manzey, D. H. (2010). Complacency and bias in human use of automation: an attentional integration. Human Factors, 52(3): 381–410.
3. Dratsch, T. et al. (2023). Automation bias in mammography: the impact of AI BI-RADS suggestions on reader performance. Radiology, 307(4): e222176.

Governance and regulation

1. Pensions Act 1995, ss. 35–36; Pensions Act 2004, ss. 247–248 (trustee knowledge and understanding); The Pensions Regulator (2024), General Code of Practice. The General Code is not itself a statement of law, but courts and tribunals must take its relevant terms into account when assessing how legal duties are met.

Background reading

- Bainbridge, L. (1983). Ironies of automation. Automatica, 19(6): 775–779.
- Goddard, K., Roudsari, A. & Wyatt, J. C. (2012). Automation bias: a systematic review of frequency, effect mediators and mitigators. JAMIA, 19(1): 121–127.

Knowa

The AI governance operating system for Boards and advisers · knowa.co