THE SOCIETY OF PENSION PROFESSIONALS

making pensions work

Governance in the Age of AI:

A Practical Framework for Responsible Leadership

July 2026

Governance in the Age of AI:

A Practical Framework for Responsible Leadership

Foreword

Artificial Intelligence (AI) has moved from an emerging technology to an everyday reality across the pensions industry. This is best illustrated by the Society of Pension Professionals (SPP) annual AI surveys which in 2025 found 87% of pension firms were using AI and in 2026 revealed that 100% of pension firms were doing so. Trustees, advisers, administrators, employers, investment managers and pension scheme members are already using AI in a growing range of activities, from drafting papers and analysing data to supporting investment decisions and interpreting pension information.

The pace of adoption has significantly outstripped the development of dedicated regulation. Rather than creating entirely new legal duties, AI requires trustees to apply existing governance, fiduciary and risk management responsibilities in new ways. The challenge is therefore not whether AI should be used, but how it can be used safely, transparently and with appropriate oversight.

This paper sets out a practical governance framework to help trustees navigate that challenge. It is intended to complement, rather than replace, existing regulatory expectations by identifying the key questions trustees should ask, the controls they should expect to see and the governance arrangements that should underpin responsible AI adoption across the pensions ecosystem.

Jo Fellowes
Chair, SPP Administration Committee

Introduction

AI is already reshaping the pensions landscape. It is increasingly embedded within trustee decision-making, administration platforms, actuarial and investment services, covenant assessment, member communications and the software used across the pensions industry. Members themselves are also beginning to rely on publicly available AI tools to understand their pension benefits and support retirement planning.

Against this backdrop, trustees are not faced with a choice of whether AI exists within their schemes, but how they govern its use. As adoption accelerates, effective AI governance is becoming an integral part of good scheme governance, requiring trustees to understand where AI is being used, the risks it presents, the controls that are in place and the accountability arrangements supporting its deployment.

The UK does not currently have a single, comprehensive statutory framework governing AI. Instead, oversight is delivered through existing legislation together with sector-specific regulators applying established legal and regulatory principles to AI-enabled activities. For pension schemes, this includes the evolving expectations of The Pensions Regulator (TPR), together with the governance requirements contained within the General Code relating to effective systems of governance, risk management, internal controls, data management and cyber resilience.

Accordingly, trustees should view AI through the lens of their existing duties rather than as a separate governance discipline. The same principles that apply to any material operational or outsourcing risk (appropriate oversight, proportionality, accountability, transparency and effective risk management) remain equally applicable where AI is involved.

The SPP believes that AI can deliver significant benefits for schemes, employers and members, provided its use is governed responsibly. This paper therefore proposes a practical, risk-based framework built around five overarching principles: transparency, accountability, proportionality to risk, security by design and meaningful human oversight. The sections that follow consider how those principles can be applied across trustee activities, advisers and administrators, investment management, data governance, contracting, member engagement and scheme oversight.

Risk classification of AI adoption and use

In terms of AI adoption, there are several considerations ranging from model governance considerations with third party vendors and concentration risk, country legislation and AI sovereignty, algorithmic bias and fairness on an ongoing basis, third party validation and audits, costs and other aspects around the subscriptions themselves. Last but not least the creation of a secure operating model and systems architecture with human oversight and control.

Fundamentally, AI use should be assessed by risk level, consistent with the risk-based approach TPR has signalled that it expects trustees and providers to apply:

- **Low risk:** internal drafting, summarisation, administrative support.
- **Medium risk:** member communications, operational processing support.
- **High risk:** outputs influencing member benefits, decisions, or advice.

High-risk uses should be subject to enhanced governance, and in some cases prohibited without robust controls and validation.

Extending the classification to enterprise, investment, and member-facing AI

The same three-tier structure can be applied consistently across the breadth of AI uses a scheme is likely to encounter, which fall into four broad categories:

- **Enterprise-level AI adoption and governance:** AI used by the scheme, its sponsoring employer, or its administrators in day-to-day operations - drafting, research, data processing, fraud detection, and workflow automation. Risk depends on the data involved and the extent to which outputs feed into member-affecting decisions without human review.
- **Investment portfolio and external manager AI adoption:** AI used by asset managers and other investment counterparties in the management of scheme assets - including in quantitative strategies, credit and counterparty risk models, and increasingly in agentic systems that initiate trading or rebalancing decisions. Trustees should treat the extent and governance quality of an external manager's AI use as a relevant due diligence consideration alongside other operational due diligence factors, and request evidence of the manager's model risk governance, validation processes, and human oversight arrangements.[1]

1 AI Plan, TPR, May 2026: https://www.thepensionsregulator.gov.uk/en/document-library/corporate-information/ai-plan

- **Data and Large Language Model (LLM) specific governance:** the specific risks introduced by generative and large language model tools, addressed further in Section 5, including data retention, prompt and output storage, and the risk of model training on scheme or member data without consent.
- **Member preferences and external AI use:** members' own use of public AI tools to interpret scheme information, including the risk that members are increasingly likely to encounter unregulated AI tools offering financial guidance or recommendations at the boundary of regulated advice[2], addressed further in Section 9.

Trustee use of AI

Trustee Directors are increasingly using AI to support their preparation for Trustee Board and Committee meetings, with the primary use being a supporting tool to summarise and distil the key information contained in lengthy and complex meeting papers.

It is also being used as a way of enhancing trustee knowledge, and testing external opinions on complex or changed expectations, to mixed effect. AI aims to please, so it looks to return the closest possible answer. If that means a prompt seeks to compare approaches to completing an ORA (own risk assessment) and the closest it finds in the domain being searched is an ORSA, it will respond on that basis. If the incorrect comparison is not identified by the user, it may result in inappropriate or delayed decision making and increased adviser costs.

Whilst not a direct Trustee use, Trustees are increasingly taking part in meetings that are being recorded, transcribed and sometimes minuted by AI, and reading papers which may have been generated using AI, whether as part of the drafting process, or in the peer review process.

> *Trustee Directors are increasingly using AI to support their preparation for Trustee Board and Committee meetings...*

Tools

Public AI tools such as ChatGPT, Copilot (Microsoft) or Gemini (Google), etc. provide accessible tools to explore content on the internet, via prompts entered into the model by the user. Data entered into some AI tools may be retained, reviewed or used to improve services or models, depending on the provider, product terms, settings and contractual arrangements. Trustees should not assume that information entered into a tool remains private unless this is explicitly confirmed. Results will be based on the data available on the internet, and there is a risk that results may not be reliable; referred to as hallucinations. Hallucinations can be mitigated to some extent with effectively written prompts that include requests for source references, which would enable the user to apply some judgement as to whether the sources are credible in the context concerned.

Caution: if a user shares a document with a public AI tool, for distillation, to identify key questions or challenges, or for direct comparison with what is available publicly, the document shared would be made available in the public domain for future searches.

Personal licensed AI tools, or corporate instances of AI, run in a private environment, and ensure user inputs and any data or documents shared, are not made available in the public domain. Although for a private user, the security of that domain may not have been considered, or be as robust as a corporate environment. Those inputs, data or documents may be shared with other users of the personal license (such as other family members) or corporate instance (rest of the organisation). The data residency of the AI model, and location of any backups, will be determined by the agreement in place with the AI provider.

Caution: if a user shares a sensitive document, such as one with member data, or a negotiating position, such as during a valuation and de-risking transaction, or a confidential change project, the information contained may become available to other users of that personal license or corporate instance, which could result in a data breach or revealing other confidential information.

Integrated AI, where AI functionality has been integrated into applications or software, including meeting management solutions, may have other rules around the environment, privacy, residency and security of information and prompts shared.

2 Financial Conduct Authority, Perimeter Report, 26 March 2026: https://www.fca.org.uk/publications/corporate-documents/fca-perimeter-report

Key Controls

Trustees should consider their approach to AI use more generally, the opportunities that AI can present to govern aspects of the scheme and service, and to speed up areas of delivery for the benefit of members, and the new risks (or new causes and consequences of existing risks) as a result of AI. This should include governance of Trustees own use of AI.

Trustees may have policies in place that cover personal device use, email addresses and home working, data and cyber security, adviser management, etc, and they may choose to extend those existing policies and processes to include AI. Alternatively, they may choose to create and adopt a separate AI policy.

Controls and assurance to monitor compliance with the new inclusions will also need to be considered and should be proportionate to the AI risks that the Trustee is exposed to, and their appetite and tolerance for those risks.

At a minimum, Trustee controls and assurance should ensure:

- No sensitive scheme data is input into unapproved tools
- Outputs are independently reviewed
- Tool usage aligns with scheme and employer policies
- AI is not used for decision-making without human validation

Trustees should consider whether a more formal AI governance operating model is also needed, for example covering:

- who owns AI governance at trustee level;
- whether there should be an AI register / inventory of use cases;
- how new use cases are approved;
- what goes to the board vs a subcommittee vs management;
- how often AI use is reviewed;
- what management information trustees should receive;
- how AI risks are recorded in the risk register / ORA

Trustees should also be mindful of future challenge and have a clear record and audit trail of material AI use, including retention of relevant records and documentation of human validation.

Trustee training or workshops on effective and secure use of AI will also be important, and recognising AI is rapidly evolving, this is unlikely to be a one and done exercise.

Adviser and administrator use of AI

AI is increasingly likely to be used by advisers, administrators and other service providers across scheme administration, actuarial work, investment support, covenant analysis, legal services, member communications and operational processing.

TPR recognises that AI can support communications, guidance and administration through personalised member support, targeted communications, automation of routine tasks, faster processing and improved fraud detection.

Trustees and scheme managers should understand where AI is being used by, or on behalf of, the scheme and whether it influences scheme outputs, member communications, calculations, case prioritisation, risk scoring, advice or member experience. This will include direct use by the adviser/administrator but also perhaps use by fourth parties such as any of their providers and sub-processors.

Although activities may be delegated, trustees remain accountable for decisions and outcomes, and TPR expects them to understand where and how AI is used across their scheme.

Possible questions that trustees might like to consider asking include but should not be limited to:

- What AI is used in delivering services to the scheme?
- What data is used, and where is it processed?
- Is any personal or confidential data used to train models?
- Which outputs are reviewed by a human?
- Are any significant decisions solely automated?
- What testing and validation has been done?
- What incidents have occurred?
- What sub-processors or model providers are involved?
- What contractual rights do trustees have to notification/audit?

Trustees should engage proactively with service providers to identify where AI and automated decision-making (*"ADM"*) are being used and require evidence of appropriate controls. Drawing on TPR's expectations, trustees and scheme managers should:

- establish clear governance, accountability and oversight for AI use across administrators, advisers and other providers;
- ensure AI systems are appropriately tested, assured and monitored before implementation and throughout their use;
- identify AI-related risks, maintain proportionate controls and review them as AI evolves;
- maintain a clear data strategy, with high-quality scheme and member data underpinning AI-supported processes;
- comply with data protection requirements, including those relating to ADM and the use of personal data in AI systems;
- understand how AI systems use, process and protect scheme data, including cyber security implications; and
- remain alert to AI-enabled scams and fraud, with appropriate preventative and response measures.

Trustees should also consider AI-enabled fraud and operational resilience. TPR warns that AI may increase the sophistication of cyber-attacks, scams, clone websites, phishing and impersonation fraud. Service providers should have controls addressing risks such as deepfake impersonation, voice cloning, fraudulent member or trustee instructions, malicious documents and AI-generated phishing, and explain how those controls operate.

Automated decision-making

Particular care is needed where a service provider's AI involves ADM or profiling under UK data protection law. Trustees are data controllers for scheme personal data and therefore need appropriate oversight of processor activity, including whether AI or automation involves ADM or profiling. Where providers act as processors, trustees remain responsible for ensuring personal data is processed lawfully and with appropriate controls.

ADM involves a significant decision made solely by automated means without meaningful human involvement. A significant decision is one that produces a legal effect or a similarly significant effect on an individual.

Examples in a pension scheme context include automated decisions on entitlement to benefit options, benefit calculations without meaningful human review, personalised investment or retirement communications that materially influence member choices, automated engagement or case prioritisation, and stopping or suspending a pension because of suspected fraud.

Trustees should not assume that nominal human oversight avoids the ADM regime. Meaningful human involvement requires a reviewer who can genuinely influence the outcome, has authority to change it, understands the system's logic, outputs, limitations and risks, and considers the relevant facts for each individual decision. Spot-checking or retrospective quality assurance is unlikely to be sufficient where decisions are otherwise made automatically. This reflects the ICO's draft approach to meaningful human involvement.

Although the UK ADM regime is more permissive than previously, safeguards remain. Trustees should establish with service providers whether ADM is used and, where it is used for a significant decision, ensure affected members or beneficiaries:

- are informed that ADM is being used;
- have an opportunity to make representations;
- have access to meaningful human intervention; and
- can challenge the automated decision.

The rules are stricter where special category data, such as health data, is involved. In practice, ADM is unlikely to be appropriate for ill-health or incapacity pension decisions, given the use of health data, the need for medical judgment and the legal and financial significance of the outcome.

Trustees should therefore require service providers to identify any AI-enabled process involving, or potentially involving, ADM or profiling, and explain:

- the purpose of the ADM;
- whether profiling is involved;
- whether decisions are solely automated or include meaningful human involvement;
- whether decisions have legal or similarly significant effects;
- the lawful basis relied upon;
- whether special category data is processed;
- what information is provided to affected individuals;
- how members can challenge decisions or request human review;
- how accuracy, bias and fairness are assessed; and
- how the process is documented and monitored.

Where a service provider proposes to use, or is already using, ADM, trustees should ask whether a DPIA has been carried out. If so, they should request a copy. If not, the provider should explain why, and trustees should consider whether those reasons are consistent with the nature, scale and potential impact of the processing.

Data governance

Key requirements include:

- Clear rules on what data can be entered into AI systems
- Restrictions on personal, financial, or member-level data use
- Understanding whether prompts or outputs are stored or used for model training
- Awareness of cross-border data transfer risks

LLM-specific governance considerations

LLM and generative AI tools introduce data governance questions that go beyond traditional data protection compliance, and trustees should expect providers to be able to answer the following:

- **Training data provenance:** whether scheme or member data submitted to a tool is used, even temporarily, to fine-tune or improve the underlying model, as distinct from being processed solely to generate the requested output.
- **Output provenance and verification:** LLM outputs can be factually incorrect or fabricated (so-called hallucination) with no inherent indication of confidence or source; trustees should require that material outputs are independently verifiable against an authoritative source before being relied upon.
- **Automated decision-making safeguards:** Article 22 of the UK GDPR gives individuals the right not to be subject to decisions based solely on automated processing, including profiling, which produce legal or similarly significant effects[3]. The Data (Use and Access) Act 2025 relaxes some of the previous restrictions on automated decision-making, but only where organisations implement safeguards allowing individuals to make representations, obtain meaningful human intervention, and challenge decisions made by solely automated means[4]. Any AI use that influences member benefit calculations or decisions should be assessed against this standard specifically.
- **Cross-border transfer and jurisdiction of processing:** many AI tools route data through infrastructure located outside the UK and EEA; trustees should understand not just where a provider is headquartered but where the underlying model inference and any data retention physically occurs.

Incident and breach management

AI-related risks should be incorporated into existing governance and breach frameworks, including:

- Error detection and escalation routes
- Reporting mechanisms to trustees
- Integration with existing breach reporting expectations under TPR
- Response plans for AI-enabled fraud or misinformation

Practical escalation framework

Trustees should ensure the following are clearly defined and tested, not merely documented:

- **Detection:** how an AI-related error or anomaly is first identified - whether through automated monitoring, user report, member complaint, or third-party notification - and the maximum acceptable time to detection for high-risk use cases as classified under Section 2.
- **Severity classification and escalation:** a defined threshold for when an AI-related incident must be escalated to trustees directly, rather than handled at provider or administrator level, calibrated to the risk classification of the AI use involved.
- **Regulatory reporting alignment:** AI-related breaches should be assessed against TPR's existing notifiable events and breach reporting framework rather than treated as a separate category requiring new reporting infrastructure; trustees should confirm with advisers that an AI-related error affecting member benefits would be correctly identified as reportable under existing duties.
- **AI-enabled fraud and misinformation response:** given TPR's specific warning on the risk of deepfake impersonation and member reliance on generative AI for pension decisions[5], response plans should include a pre-agreed protocol for verifying member identity where deepfake or voice-cloning fraud is suspected, and a communication plan for alerting members rapidly if a scheme-specific scam pattern is identified.

3 UK General Data Protection Regulation, Article 22 (right not to be subject to automated decision-making producing legal or similarly significant effects).

4 Data (Use and Access) Act 2025: https://www.legislation.gov.uk/ukpga/2025/18/contents

5 Financial Conduct Authority, Perimeter Report, 26 March 2026: https://www.fca.org.uk/publications/corporate-documents/fca-perimeter-report

Covenant considerations

Monitoring the strength and resilience of the employer covenant is a core component of effective governance, with TPR's 2026 Annual Funding Statement highlighting that *"[t]he employer covenant supporting schemes remains integral when assessing the level of supportable risk in a scheme's journey plan"*.

The pace of AI development and adoption is already reshaping businesses models across multiple sectors, with direct implications for employer covenant strength. For some, particularly in the services sector, there is a significant opportunity from efficiency gains. With the technology improving at an exponential rate, these impacts will continue to spread across more sectors.

However, the economic benefits are unlikely to be evenly distributed – there will inevitably be winners and losers. Employers that are slower to adapt may find it increasingly difficult to compete on price and/or capability with rivals further along the curve. These effects might first emerge as reduced cash generation but longer-term could mean more uncertain prospects.

AI is also rapidly transforming the cyber risk faced by employers. The Five Eyes cyber security agencies have issued a call to action to business leaders, stating: *"Frontier AI models are anticipated to exceed current industry expectations…The timeline is not years, it is months. In this environment, cyber resilience is integral to advancing business continuity, market confidence, and long-term value"*[6]. In a covenant context, a significant cyber event can cause business interruption, reduced revenues and reputational damage.

The importance of understanding such risks will be further heightened in run-on and potential surplus release scenarios, where the longevity and resilience of the employer becomes paramount.

Trustees will want to ensure their covenant monitoring processes are adapted to capture these emerging opportunities and threats. This may include:

- discussions with management on the impacts of AI adoption on their industry.
- assessing the employer's relative positioning versus competitors.
- benchmarking revenue growth and profitability trends.
- understanding the readiness of the employer (and its key customers and suppliers) to defend against frontier AI-enabled cyber threats.

Ultimately AI does not just change how schemes are run, it changes the risk profile of the employer supporting them.

Contracting

As pension scheme advisers and administrators increasingly incorporate AI into their service delivery, trustees should begin considering whether existing contractual arrangements remain adequate. Many current appointments were negotiated before the widespread deployment of generative AI and machine learning tools and so are unlikely to contain provisions specifically addressing their use. Given trustees' continuing responsibility for scheme governance, contractual documentation should evolve to provide greater transparency, oversight and accountability where AI forms part of the services provided. Set out below are areas that trustees should ideally seek to incorporate into contracts with their advisers and administrators.

Disclosure of AI use

Trustees should consider requiring service providers to disclose whether, and to what extent, AI is used in delivering contracted services. This should extend beyond generative AI tools to include machine learning systems, automated decision-making processes and AI-enabled workflow tools where they have a material role in producing advice, calculations, member communications or governance reporting. Disclosure provisions should identify:

- the categories of AI technologies used;
- the functions for which AI is deployed in provision of the services;
- whether outputs are reviewed by appropriately qualified personnel before being relied upon or communicated; and
- any material limitations or risks associated with the technology.

Such transparency enables trustees to understand where AI is influencing service delivery and to determine whether additional governance measures are appropriate.

The pace of AI development and adoption is already reshaping businesses models across multiple sectors, with direct implications for employer covenant strength.

6 The AI shift in cyber risk: why leaders must act now, 22 June 2026: https://www.ncsc.gov.uk/news/the-ai-shift-in-cyber-risk-why-leaders-must-act-now

Governance and monitoring requirements

Contracts should also require providers to maintain appropriate internal governance arrangements for AI systems used in connection with the pension scheme services. Although responsibility for the provider's internal governance remains with that provider, trustees have a legitimate interest in ensuring that appropriate controls exist where those systems may affect scheme administration or decision-making. Appropriate contractual obligations may include requirements for providers to:

- maintain documented governance frameworks for AI systems;
- undertake risk assessments before deploying AI into live services;
- implement appropriate human oversight for material outputs;
- maintain policies addressing accuracy, bias, security and model performance;
- provide periodic reporting on AI governance where requested; and
- cooperate with trustee governance reviews.

This enables trustees to demonstrate appropriate oversight of outsourced activities without assuming responsibility for the provider's operational management.

Audit and assurance rights

Traditional audit clauses often focus on financial controls or operational processes but may not extend adequately to AI governance. Trustees should therefore consider ensuring that contractual audit provisions encompass AI-related controls. Depending upon the nature of the services, contracts could require providers to make available evidence of:

- AI governance policies;
- relevant independent assurance reports;
- internal control testing;
- cybersecurity controls relating to AI systems;
- validation and quality assurance processes; and
- remediation of significant AI-related incidents.

Trustees are unlikely to require access to proprietary algorithms or commercially confidential model architecture. Instead, the emphasis should be on obtaining sufficient assurance that appropriate governance, controls and oversight arrangements are operating effectively.

Data protection and confidentiality

While existing contracts will generally require compliance with UK data protection legislation and the maintenance of confidentiality, AI introduces new and additional risks. Contracts should be updated to address AI-assisted processing of personal data rather than focusing solely on compliance with data protection legislation. Trustees should therefore consider requiring providers to give specific contractual commitments regarding the use of AI in processing pension scheme data, including confirmation that personal data and confidential scheme information will not be inputted into or used to train public or shared AI models without the trustee's express authorisation.

Contracts should also require providers to maintain appropriate technical and organisational security measures for AI-enabled processing, ensure that confidentiality obligations apply equally to information processed using AI systems, and implement appropriate controls over any AI-related sub-processors or international data transfers.

Notification of new or expanded AI use

Many service providers are adopting AI incrementally, meaning that services initially delivered entirely by human personnel may become increasingly automated over time. Without contractual notification requirements, trustees may remain unaware that significant aspects of service delivery have changed. Notification provisions might require providers to inform trustees of:

- the introduction of new AI technologies affecting scheme services;
- material changes to existing AI deployment;
- significant changes in the role AI plays in producing advice or administrative outputs;
- AI-related incidents that could affect scheme members or trustee decision-making; and
- changes in sub-contracting arrangements involving AI providers.

For particularly significant changes, trustees may wish to reserve the right to discuss the proposed deployment, request additional information regarding governance arrangements or, in exceptional cases, withhold consent where the proposed use would materially alter the agreed service model.

Scheme members and external use of AI

AI is also changing how members engage with their pensions - with AI tools increasingly being used by members to summarise scheme communications, interpret benefit structures, support complaints and for retirement planning.

Emerging Risks

While AI has the potential to improve member engagement and understanding, it also introduces a number of new risks.

- **Inaccurate outputs.** AI tools may generate inaccurate, inconsistent or misleading information – for example, through misinterpreting or oversimplifying scheme provisions, relying on incomplete, outdated or irrelevant sources (e.g. information relating to another pension scheme), or by producing *"hallucinated"* content.
- **False confidence.** The confident and authoritative language used by many AI tools may create a sense of accuracy and certainty, even in relation to outputs that are fundamentally incorrect or unsupported.
- **Unregulated *"advice"*.** While most AI tools sit outside the perimeter of regulated financial advice, their conversational and seemingly tailored output can blur the distinction between information and a personalised recommendation, meaning members may treat them as personalised advice and rely on them accordingly (despite their limitations).
- **Data privacy and fraud.** Member use of public AI tools may create data privacy and security risks, particularly where benefit statements or other personal financial data is inputted into such tools. This may result in member's personal data being exposed or misused and could leave members more vulnerable to fraud and scams.
- **Accessibility and Digital Exclusion.** While AI can enhance engagement, trustees must consider whether AI processes could create challenges or even unfair outcomes for certain groups. Trustees will need to consider how these processes might impact vulnerable members, or could amplify existing inequalities. There is a distinct risk of digital exclusion; therefore, it will be important to ensure there remain appropriate, easily accessible routes to human support alongside any AI-driven tools. For further information, please see the SPP's paper, *"Pensions in a Digital World: Embedding Inclusion"* due to be published in the summer of 2026[7].

These risks present a particular challenge for trustees. Where members are independently using AI tools, trustees cannot control how members use those tools, or the tools themselves (or even tell exactly how they are being used). However, poor outcomes resulting from misinformation or inappropriate guidance as a result of AI usage could still result in member complaints or undermined member confidence in the scheme more generally.

Potential considerations for trustees

While trustees cannot control members' use of AI there are steps trustees can consider taking to help mitigate some of these risks:

- **Improve existing scheme content / communications** – reviewing and improving existing scheme communications and website content may be a particularly impactful way to mitigate some of these risks. Firstly, improved content may reduce the number of members looking to AI in the first place for information about their benefits. Secondly, as AI outputs are only as good as the underlying inputs – the more clear, comprehensive and easily findable scheme content can be, the more likely that AI tools base their output on the correct underlying inputs (rather than hallucinating or pulling information from other schemes). For some schemes, this may involve considering whether to make scheme content public for the first time, to help AI tools to extract the information trustees would like members to see.
- **Education and guidance** – consider providing education or guidance to members around how to use AI tools more safely (e.g. verifying AI outputs and being mindful of inputting any confidential or personal information).
- **Risk warnings** – consider whether to include AI caveats or *"health warnings"* in member communications on the risks of using AI to provide pensions information or retirement planning as well as signposting members to any other resources or support available.
- **Enhance in-scheme support and resources** – consider improving the support, resources and tools provided by the scheme to reduce the likelihood of members looking to AI for answers. For example, could well-deployed targeted support help steer members away from using AI for retirement planning purposes?

7 "Pensions in a Digital World: Embedding Inclusion", SPP, 2026: https://the-spp.co.uk/document-category/guides-and-reports/

- **Scheme specific AI tools** – consider whether introducing scheme-specific AI tools, such a member-facing chatbots or virtual assistants might help minimise members' use of public AI tools. Unlike public AI tools, these solutions could be trained on scheme-specific information, operated within agreed guardrails and subject to ongoing monitoring and oversight.

Implementation, accountability and oversight

Trustees' fiduciary duties are not altered, suspended or diluted by the use of AI. They continue to apply in full force whether decisions are made directly by trustees, informed by AI-enabled tools, or delegated to third parties using AI within their service delivery. The core obligations to act prudently, in members' best interests, and with the skill and care of an ordinary prudent person of business extend equally to the decision to rely on, or accept, AI-generated outputs.

This principle is reinforced by TPR's recently published 2026 AI Plan. This makes clear that trustees and scheme managers *"remain accountable for decisions and outcomes even when they delegate activities to providers or advisers,"*[8] and are expected to understand where and how AI is being used across the scheme and its service providers. TPR also expects trustees to assure themselves that administrators, advisers and other service providers have appropriate governance arrangements in place, supported by rigorous testing, assurance and ongoing monitoring both at implementation and on an ongoing basis.

This regulatory direction is understandably consistent with wider supervisory messaging. In May 2026, the TPR Chief Executive emphasised that trustees, administrators and scheme managers should *"act now"* by strengthening governance, improving data quality, understanding where AI is being used, and protecting members from AI-driven fraud. This reflects the fact that AI governance is now a supervisory priority across UK financial services, with pensions uniquely exposed given the scale of assets under management and membership coverage.

> *Rather than being treated as a standalone governance framework, AI should be embedded within existing trustee systems and processes.*

In parallel, the DWP is developing further guidance on the scope of trustees' fiduciary duties. This work signals a more expansive approach to fiduciary decision-making, including consideration of systemic risks and broader member outcomes. While climate change is cited as a leading example of a systemic risk, AI should be understood in the same context: it is a cross-cutting technological and operational factor affecting the pensions ecosystem as a whole, including scheme administration, investment management, advisory services and member behaviour.

Taken together, these developments reinforce a central point: trustees do not require new legislation to conclude that AI governance falls squarely within the scope of their existing fiduciary duties. AI is already embedded across the value chain, and its governance is therefore an inherent part of acting prudently in members' interests.

Against this backdrop, clear accountability for AI use must be established and documented within existing governance structures. This should include clarity on who is responsible for AI-related governance within the scheme, the respective roles of trustees, advisers, administrators and other service providers, ownership of AI-related risks and controls, and defined escalation routes for errors, incidents, breaches or concerns arising from AI use. It should also include mechanisms for challenge, review and correction of AI-generated outputs where they are relied upon in scheme decision-making or communications.

Rather than being treated as a standalone governance framework, AI should be embedded within existing trustee systems and processes. This includes integration into risk registers and risk management frameworks, internal controls and assurance processes, service provider oversight arrangements, cyber security and data governance frameworks, and business continuity and incident response planning.

Effective oversight also requires appropriate information flow from service providers. Trustees should receive proportionate and risk-based reporting on AI usage within scheme services. This may include disclosure of material AI use cases affecting scheme operations, changes to AI-enabled systems or processes, identified incidents or control failures, results of assurance or validation activity, and emerging regulatory or operational risks.

The frequency and depth of oversight should be calibrated to risk. Higher-risk AI applications, particularly those affecting member outcomes, financial decisions or regulated activities, should be subject to enhanced scrutiny, documented review and appropriate independent assurance.

8 AI Plan, TPR, May 2026: https://www.thepensionsregulator.gov.uk/en/document-library/corporate-information/ai-plan

Trustees and those involved in scheme governance should also maintain an appropriate level of knowledge and understanding of AI-related risks and opportunities. This may require periodic training, updates on regulatory and industry developments, and the incorporation of AI considerations into trustee knowledge and understanding (TKU) frameworks.

Finally, as the foreword to this paper makes clear, AI governance should be treated as a rapidly evolving discipline. In light of the pace of change, trustees should ensure that policies, controls, oversight arrangements and service provider expectations are reviewed regularly to remain aligned with evolving technology, industry practice and regulatory expectations.

The question is no longer whether trustees should engage with AI, but whether they are governing its use with the same discipline, challenge and oversight that they apply to any other material source of risk and opportunity.

Conclusion

AI is no longer an emerging issue for pension schemes. It is becoming part of the fabric of scheme governance. Whether used by trustees, advisers, administrators, investment managers, employers or members themselves, AI is already influencing how schemes are governed, administered and experienced. The question is no longer whether trustees should engage with AI, but whether they are governing its use with the same discipline, challenge and oversight that they apply to any other material source of risk and opportunity.

The good news is that trustees do not need an entirely new governance framework. The principles that have always underpinned good trustee decision-making - prudence, accountability, transparency, proportionality, effective risk management and appropriate challenge – each remain the right principles for the age of AI. What is changing is how those principles are applied. Trustees must understand where AI is being used, ask informed questions of their advisers and service providers, ensure appropriate controls are in place and satisfy themselves that meaningful human accountability remains at the centre of important decisions.

AI also presents a significant opportunity. Used responsibly, it has the potential to improve efficiency, strengthen administration, enhance decision-making, reduce fraud, improve member engagement and ultimately deliver better outcomes for schemes and their beneficiaries. Equally, unmanaged or poorly governed AI has the potential to create operational, legal, cyber, financial and reputational risks. Good governance is therefore not a barrier to innovation, it is what allows innovation to be adopted with confidence.

This paper is intended to provide a practical framework for that governance. It is not a static checklist, but a set of principles and questions that should evolve alongside technology, regulation and market practice. Trustees who embed AI governance within their existing systems of governance today will be better placed to respond to the opportunities and challenges of tomorrow.

The pensions industry has successfully adapted to profound regulatory, economic and technological change over many decades. AI is the next chapter in that evolution. Those schemes that approach it with curiosity, appropriate scepticism and robust governance will be best placed to harness its benefits while continuing to protect the interests of the members they serve.

Acknowledgements

The SPP is grateful to all its members for their contribution towards this publication, particularly Jo Fellowes, Tom Partridge, Sonya Fraser, Hannah Savill, Ian D'Costa and Chandra Gopinathan.

About The Society of Pension Professionals

Founded in 1958 as the Society of Pension Consultants, today SPP is the representative body for a wide range of providers of pensions advice and services to schemes, trustees and employers. These include actuaries, accountants, lawyers, investment managers, administrators, professional trustees, covenant assessors, consultants and pension specialists.

Thousands of individuals and pension funds use the services of one or more of the SPP's members, including the overwhelming majority of the 500 largest UK pension funds.

The SPP seeks to harness the expertise of its around 90 corporate members - who collectively employ over 20,000 pension professionals - to deliver a positive impact for savers, the pensions industry and its stakeholders including policymakers and regulators.

Further information

If you have any queries or require any further information about this discussion paper, please contact the SPP's Director of Policy & PR, Phil Hall phil.hall@the-spp.co.uk or telephone 07392 310264

To find out more about the SPP please visit the SPP web site: https://the-spp.co.uk/

Connect with us on LinkedIn at: https://www.linkedin.com/company/the-society-of-pension-professionals/

Follow us on X (Twitter) at: https://twitter.com/thespp1

We are a Living Wage Employer

Prompt Payment Code

THE SOCIETY OF PENSION PROFESSIONALS

making pensions work