

DISCLAIMER All views or opinions expressed in these essays are the author's own in respect of the essay topic and do not represent the views or opinions of the APLI. The APLI have made the essays available for informational purposes only and the APLI does not make any representation or warranties with respect to the accuracy or applicability of the essays.

Author: Louise Boylan

AI and the Future of Pensions Law: Risks, Governance and Trustee Decision-Making

I. Introduction: Route 66

Route 66 is one of the most famous roads in the world. It carries travellers from their first mile in Chicago to the journey's end on the Pacific coast. Nearly every Irish worker travels a road just like it, from their first payslip to the State pension age of sixty-six, a working lifetime later.¹

A new driver now holds the wheel. Artificial Intelligence ('AI') has arrived in the pensions industry. A recent study found eighty-seven per cent of pension firms using AI in 2025, with complete adoption anticipated this year.² Trustees no longer *decide* whether AI operates within their schemes, only *how* it is governed. A system that fails, fails at scale, in the same confident voice it uses when correct. The Irish courts have already met that voice, in submissions citing authorities that did not exist.³ In a courtroom, such an error costs time. In a pension scheme, it could cost a retirement.

The answer is not to ban the new driver. It is what Susskind calls '*harnessing*', taking the technology's benefits whilst keeping it reined in.⁴ This essay weighs the benefits and limitations of AI in pensions and analyses the legal risks. It recommends an approach for Ireland, examines AI's impact on trustee decision-making, and closes with practical advice on *harnessing* it. Ultimately, where AI informs the trustee's judgment, the law should welcome it. Where it replaces judgment, the law must never hand over the keys.

PART A

I. Benefits and Limitations

Green Lights

Pension administration, largely standardised and repetitive, lends itself readily to AI.⁵ Four benefits are notable. The first pertains to members. AI can draft clearer communications and answer questions put to chatbots at any hour. The second is efficiency, as AI can identify record gaps and process routine work, thereby reducing the cost of running a scheme. The third is protection, as AI can detect scam patterns and flag suspicious transfer requests. The

¹ Social Welfare Consolidation Act 2005, s 2 (as amended)

² Society of Pension Professionals, 'Governance in the Age of AI' (2026) 2.

³ *Guerin v O'Doherty* [2026] IECA 48.

⁴ Richard Susskind, *How to Think About AI* (OUP, 2024).

⁵ PensionsEurope, Answer to the EC's Targeted Consultation on AI (2024) 3.

UK Pensions Regulator ('TPR') has developed a successful tool to identify scam websites.⁶ The fourth is investment advantage, as AI can analyse markets and data at a scale no human could match.

So, why not simply hand the machine the keys?

Speed Bumps

The first limitation is invention. Everyone is familiar with the term 'hallucination', but few realise it is part of the design. As Melham observes, the technology is designed to mimic the intellectual products of humans.⁷ It predicts plausible language rather than testing its truth. As such, *hallucination is a feature, not a failure*. The danger is acute in pensions, where questions encompass historic deeds, trust law, statute and case law. A tool without that context is as fluent as it is wrong.

A further limitation is bias, arising at two levels. At the model level, UNESCO found generative systems display stark prejudice against women.⁸ At the data level, pension records are notoriously error-prone. In 2025, the State ordered 13,000 civil service pensions to be rechecked for systemic deduction errors,⁹ and a model trained on such records learns those errors as the truth.

Third, AI is both a shield and a sword for protection. The same tools that detect scams also sharpen them, making phishing and deepfakes more personalised, faster and scalable.¹⁰ Meanwhile, schemes hold vast member data, from PPS numbers to health records, making them attractive targets. The Capita breach, an action involving 4,000 members, shows how data failures can expose schemes to serious liability.¹¹

Finally, any document or information shared with AI tools risks being placed into the public domain for future searches and may waive legal privilege.¹² Indeed, hundreds of shared conversations with a leading AI chatbot, some containing names and work histories, were recently found indexed by search engines and openly accessible online.¹³

None of these limitations renders AI useless, but each makes it dangerous in proportion to the trust placed in it. The questions are not technical but legal. How far may trustees delegate? What must remain explainable? And, who answers when the machine takes a wrong turn?

II. Legal Risks: Regulation, Data and Governance

AI may perform the task, but the law asks who exercised the judgment. That question is particularly important in pensions because no single framework governs AI within a scheme. Regulation spans, *inter alia*, the EU AI Act, data protection law and pensions governance.

⁶ TPR, 'Pension Scam Fighters Stand United as AI Takes the Fight to the Fraudsters' (9 April 2025).

⁷ Tom Melham in Zou (ed), Cambridge Handbook of Generative AI and the Law (CUP, 2025).

⁸ UNESCO, Challenging Systematic Prejudices (2024).

⁹ O'Donovan, 'Misunderstanding and Inconsistency behind Pension Errors' (RTÉ, September 2025).

¹⁰ International AI Safety Report (2025).

¹¹ Spurgeon & Ors v Capita PLC [2026] EWHC 241.

¹² Munir v Secretary of State for the Home Department [2026] UKUT 81.

¹³ Hays, "Some People's Chats with Claude-AI Found Publicly Available Online" (BBC, July 2026).

Additionally, responsibility is divided. Administrators operate the systems, employers supply data and may shape their use, and external providers design tools that schemes often cannot inspect. Each may incur liability within its own role, but trustees retain overarching accountability, whilst members ultimately bear the consequences when the chain fails.

This accountability follows from the trustee's role. An Irish occupational scheme is constituted as a trust. Its assets are held separately from the employer, whilst trustees administer the scheme, invest its funds and pay benefits.¹⁴ Their function is not merely administrative; the deed also confers discretion requiring judgment. An ill-health pension may turn on whether medical evidence establishes permanent incapacity. Where a member dies in service, trustees decide which dependants receive a lump-sum benefit. Neither question is mechanical. Thus, AI may inform the decision, but it cannot replace the trustee.

i. The EU AI Act

The AI Act ranks systems by risk.¹⁵ Article 5 prohibits a narrow class of practices outright. Annex III lists the high-risk systems that require conformity assessment and assigns responsibilities to deployers under Article 26. Depending on who controls the system's use, the deployer may be the trustee board, the administrator, or another stakeholder.

Most current pension uses fall outside Annex III. PensionsEurope has concluded that the current use of AI in pensions is not high-risk.¹⁶ For instance, the use of chatbots under Article 50 generally requires only disclosure that users are interacting with a machine. However, classification follows function, and a tool informing discretionary benefit decisions could fall within Annex III.

Despite this, it is important to note the AI Act was designed for harms that announce themselves, yet pension harm does not. A hidden error can compound silently for years and emerge only at retirement, when the remedy arrives decades too late.

Moreover, enforcement deepens the gap. Nine Irish authorities enforce the Act, including the Financial Services and Pensions Ombudsman, but not the Pensions Authority.¹⁷ The body hearing complaints holds AI Act powers, whilst the scheme's supervisor holds none.

ii. Data Protection

Data protection is where day-to-day risk accumulates. As Pollard noted long before generative AI, data protection was already a trustee duty.¹⁸ Trustees will generally act as controllers, while ill-health cases engage special-category health data under Article 9 GDPR. Article 22(1) also protects members against significant decisions made solely by automated processing. That line is crossed once human involvement becomes ceremonial.¹⁹

¹⁴ Buggy, Finucane and Tighe, *Irish Pensions Law and Practice* (TRH, 2006).

¹⁵ Regulation (EU) 2024/1689.

¹⁶ n5.

¹⁷ SI No 366 of 2025.

¹⁸ David Pollard, *Freshfields Corporate Pensions Law* (Bloomsbury, 2015) para 11.6.1.

¹⁹ Regulation (EU) 2016/679.

Other stakeholders are not immune. Administrators act as processors or joint controllers, whilst employers remain responsible for the workforce data supplied to the scheme. Therefore, AI lengthens the chain of accountability rather than transferring responsibility to one participant.

Confidentiality creates a further risk. In *Munir v Secretary of State for the Home Department*, a representative uploaded privileged material to an open-source AI tool, and the court held privilege waived.²⁰ Trustees, employers and administrators routinely exchange privileged advice on funding, benefits and disputes, and one careless prompt in an unapproved system may hand it to the other side.²¹

iii. Pensions Governance

The third framework is pensions governance. Newman's taxonomy identifies three recurring pension mistakes: document, administration and decision errors.²² AI does not invent a fourth category; it magnifies all three. For instance, a hallucinated deed provision is a document error, a flawed output paid to 10,000 members is an administrative error, and an unexamined recommendation is a decision error. Existing governance duties still apply, but AI increases the speed, scale and opacity of failures.

These risks engage the governance and risk-management requirements of IORP II, as transposed into the Pensions Act 1990. Any AI tool that uses member data falls within it and must be identified, assessed, and controlled.²³

Employers may also be exposed. In *Phillips v National Grid Gas*,²⁴ an entitlement to an incapacity pension depended on the employer's opinion. The High Court held that forming the opinion without proper process and reasonable care was a breach. Consequently, an employer outsourcing that opinion to an algorithm invites the same finding.

The risk extends across the supply chain. As noted, administrators often purchase systems from external providers, which may in turn rely on models that neither the administrators nor the trustees can inspect. Donald called the modern fund a 'virtual' institution for its outsourcing.²⁵ AI deepens that dependence. Although liability may arise at several points in the chain, trustees remain answerable to members for the scheme's governance. Many schemes will use the same system, so a single defect may be replicated, turning an isolated failure into a systemic one.

III. Should Ireland Follow the UK?

²⁰ n12.

²¹ n13.

²² Newman, *A Practitioner's Guide to Correcting Mistakes in Pension Schemes* (6th edn, Bloomsbury 2022) para 1.1.

²³ Directive 2016/2341.

²⁴ [2025] EWHC 693.

²⁵ Scott Donald, *The Pension Fund as a 'Virtual Institution'* in Agnew, *Pensions: Law, Policy and Practice* (Hart, 2020) 111.

In May 2026, TPR published an AI Plan built on governance, supply chains and scam protection. The overarching message is that ‘trust is the most valuable asset in our system’.²⁶ It emphasises that AI should support, not replace, human decision-making. By contrast, Ireland appears to be building machinery rather than guidance, based on two main indicators.

Firstly, the Regulation of Artificial Intelligence Bill 2026 contains over one-hundred sections, yet no provision tells a trustee whether their administrator’s chatbot may draft a benefit statement. That is no criticism of a Bill necessarily *general by design*. Its generality is precisely why sector-specific guidance is needed. Who else will translate broad obligations into the realities of pension administration?

Secondly, the Irish Pensions Authority is drafting an AI policy, but only for ‘internal’ use.²⁷ The schemes it supervises remain *without* guidance. Arguably, European law does not fully capture the sector’s particular risks, while national enforcement gives no express role to the regulator. Where, then, are schemes to find practical direction?

The EU AI Act and GDPR impose obligations, so concerns about duplication are understandable. Yet restrictions are not instructions, and delay is no answer when widespread adoption is already anticipated.²⁸ As Susskind observes, patients want health, not neurosurgeons, and litigants want justice, not judges.²⁹ Likewise, members want pensions paid correctly, not a particular process. Guidance would help trustees deliver that without surrendering responsibilities.

PART B

The Trustee and the Black Box

AI changes the tools available to trustees, but not the identity of the legal decision-maker. The central question is whether a trustee who relies on an AI output has still exercised the judgment required by trust law. That question has three dimensions: whether reliance is consistent with fiduciary duty, how far tasks may be delegated, and when reliance becomes an unlawful surrender of discretion. Trust law, whose problems ‘change with the times’ but whose principles adapt, can answer all three.³⁰

Fiduciary duties are neither modified nor diluted by AI and extend to every decision resting on its output. Trustees must act in members’ best financial interests, to the prudent person of business standard.³¹ In *Edge v Pensions Ombudsman*, it was held that discretionary decisions stand only where trustees consider the right questions and weigh the relevant facts.³² The courts went further in *Pitt v Holt*, holding that sufficiently serious ‘inadequate deliberation’ is a breach of duty.³³ Similarly, in *Greene v Coady*, Charleton J emphasised that trustees must

²⁶ TPR, AI Plan (20 May 2026)

²⁷ Dáil Debate 14 April 2026, Written Answer 1859.

²⁸ n2.

²⁹ Richard Susskind, *The End of Lawyers?* (OUP, 2008) 28.

³⁰ Lord Nicholls, ‘Trustees and Their Broader Community’ (1995).

³¹ *Cowan v Scargill* [1985], *Bartlett v Barclays Bank* [1980].

³² [2000] Ch 602

³³ [2013] UKSC 26.

inform themselves and deliberate.³⁴ Therefore, the black box creates no escape from judgment. Trustees may use AI to inform a decision, but they must still understand, test, and deliberate.

Yet AI strains these duties. The first issue is explainability. The duty to weigh considerations is exactly what most models frustrate. A trustee cannot explain what they do not understand, and a court cannot review what cannot be explained. Discretion is vital precisely in cases the model has never seen, and those decisions *must* remain human. As Porter MR highlighted, the trustee is ‘an honourable and upright, as well as an intelligent and capable’ officeholder.³⁵ The model is capable and intelligent. Often, nothing more.

The second is over-reliance. Trustees have always leaned on advisers, but delegation does not extinguish responsibility. In *Scott v National Trust*, it was held that advisers advise, but trustees decide.³⁶ AI complicates that distinction because it is an adviser that cannot be questioned, cross-examined or held to account.

The governing line falls between AI *informing deliberation* and AI *supplanting it*, and an Irish court has drawn it. In *Guerin v O’Doherty*, hallucinated AI citations were submitted to the Courts, which subsequently laid down principles that transpose readily to trustees. AI may be used, but responsibly. Its use must be disclosed, every product independently verified, and responsibility remains with whoever deploys it.³⁷

PART C

Holding the Trust: Advice to Trustees

What should trustees do whilst guidance is awaited? Melham warns AI is ‘notoriously hard to define’, and what cannot be defined resists neat regulation.³⁸ However, trustees need not wait. Existing fiduciary duties can be translated into three disciplines. To harness AI use, there must be transparency, accountability, risk controls, security from the outset, human oversight and a verifiable audit trail.³⁹

1. Plan the Route

The first discipline is putting the scheme’s use of AI in writing. Every scheme should adopt an AI policy that identifies approved tools, permitted tasks, and a named person responsible for oversight. Additionally, usage should be graded by risk, from low-risk internal drafting to decisions affecting benefits, which should be tightly controlled. Any new tool should start on the test track, sandboxed away from live member data until tested, assured, and monitored.

Three rules are imperative. Firstly, no scheme data or privileged material should be entered into an unapproved tool. Secondly, no AI output should be relied upon without independent

³⁴ [2014] IEHC 38.

³⁵ *BoI v Cogry* [1900] IR 219, in Biehler, *Equity and the Law of Trusts* (2016).

³⁶ [1998] 2 All ER 705.

³⁷ n3.

³⁸ n7.

³⁹ n2.

review. Third, no shared model should be trained on member data without express authorisation.

Trustees should know what technology is used, what it receives, and who is accountable for its output.

2. Verify Everything

The second discipline is verification. Nothing AI-generated should reach a member or a minute without checking and sign-off by a named person. An audit trail is crucial, as any challenge will be judged on the output and who reviewed it. In *Turner v Turner*, trustees signed appointment deeds they had never read, and the appointments failed.⁴⁰ A trustee approving unread machine output makes the same mistake but with more sophisticated software.

3. Stop for Every Traveller

The third discipline is preserving the route to human review. Automation may process routine tasks, but it must not trap exceptional cases or determine significant benefits without meaningful intervention. As Warner J observed, pension members are not volunteers, and access to a fair human process is part of what they were promised.⁴¹

Every automated system should allow members who fall outside the model's assumptions to reach a person who can reconsider the outcome. This does not require trustees to become technologists. It requires only what equity has always demanded - prudence, loyalty and deliberation.

Conclusion: Journey's End

As Cardozo wrote, the law, like the traveller, must be ready for the morrow.⁴² The morrow is already here. AI is embedded in pension administration, investment, and member communication, bringing great benefits but concerning limitations.

The law has not been absent. The AI Act, data protection and pensions governance each impose obligations, while trust law continues to demand its irreducible core.⁴³ Yet the framework remains fragmented: trustees bear the accountability, and members face the consequences. Ireland should follow the UK by issuing pensions-specific guidance that makes those responsibilities clear before harm emerges, decades later, at retirement.

Route 66 was decommissioned in 1985, replaced by faster roads. The travellers kept travelling, and their destinations endured. The AI systems used by pension schemes will likewise be replaced many times before today's youngest members retire. The duties protecting those members will remain. Trustees must plan the route, verify every output and preserve a path to human review. Where AI *informs* their judgment, trustees should harness it.

⁴⁰ [1984] Ch 100.

⁴¹ *Mettoy v Evans* [1990] 1 WLR 1587, cited in Biehler.

⁴² Cardozo, *The Growth of the Law* (1924)

⁴³ *Armitage v Nurse* [1998] Ch 521

Where it would *replace* that judgment, they must never hand it the keys. The technology is only ever a vehicle. The trustees, as always, must stay in control of the destination.