AI in Real Assets

Underwriting, Lease Abstraction, and Asset-Level Monitoring

Dr. Leigh Coney

Founder & Principal Consultant, WorkWise Solutions

Q2 2026

workwisesolutions.org

leigh@workwisesolutions.org

Abstract

Real assets are the most heavily documented holdings in any private portfolio, and the workflows that support them still run on manual retyping. This paper examines the three places where AI is earning its keep in real estate and infrastructure investing: underwriting, lease abstraction, and asset-level monitoring. Adoption is nearly universal and strikingly shallow. JLL's 2025 survey of more than 1,500 senior decision-makers found 88 percent of investors piloting AI, while 5 percent of organisations report achieving all of their programme goals. The information problems that AI addresses, meanwhile, are getting more expensive. $875 billion of US commercial and multifamily mortgages mature in 2026. Office CMBS delinquency set an all-time high of 12.34 percent in January 2026. US appraisals diverged from subsequent sale prices by more than those of any other region in MSCI's annual comparison, and 76 percent of infrastructure valuation survey respondents report selling assets above their latest carrying value. This paper argues that underwriting, abstraction, and monitoring amount to one data decision rather than three procurement decisions. Every cash flow in a real asset portfolio is defined by a contract, and the file built from those contracts at acquisition should never stop being read. The paper sets out a four-layer design (contract, operating, signal, judgment) for making that happen, and the behavioural guardrails that keep machine reading from displacing owner judgment.

Keywords

AI in real estate, AI in infrastructure, commercial real estate underwriting, lease abstraction, rent roll normalisation, asset management, portfolio monitoring, appraisal accuracy, valuation lag, building performance standards, concession agreements, family offices, automation complacency

Contents

1. Introduction: The Most Documented Asset Class 3
2. The 2026 Backdrop: Repriced Assets, Stale Numbers 3
3. Adoption Without Production 4
4. Underwriting: Reading Everything Once 5
5. Lease Abstraction: Boring Work, Priced Consequences 6
6. Asset-Level Monitoring: Where the Hold Period Leaks 7
7. The Persistent Asset File 8
8. Keeping Judgment in the Loop 9
9. Conclusion and Recommendations 10
References 10
About the Author 11
About WorkWise Solutions 11

1. Introduction: The Most Documented Asset Class

Real assets run on paper. A single office acquisition arrives as dozens or hundreds of leases and amendments, estoppels, subordination and non-disturbance agreements, a loan package, three years of operating statements, a rent roll formatted however the seller's property management system happens to export it, an appraisal, environmental and engineering reports, and a tax history. An infrastructure stake adds concession agreements, power purchase agreements, operations and maintenance contracts, and regulatory determinations that run longer than all of them. Every number an investment committee sees, from net operating income to debt service coverage to weighted average lease term, is an extraction from those documents, performed by hand, usually twice: once by the seller's analysts and once by the buyer's.

That makes real assets an unusually clean match for the current generation of AI, which performs most dependably when work is grounded in defined documents and output can be checked against a source. It also explains the speed of the land grab. In July 2023, fewer than 5 percent of real estate organisations had started AI pilots; by late 2025, 88 percent of investors had (JLL, 2025).

What the adoption numbers hide is where the effort concentrates. Pilots cluster at the point of acquisition, where work is deadline-driven, budgets sit with fee earners, and capability is easy to demonstrate. The hold period, which runs seven to ten years for value-add real estate and several decades for core infrastructure, is where the same documents go to be forgotten.

Most of what has been published on AI in real assets is written by people selling software, and it reads accordingly: use cases, accuracy claims, customer logos. This paper is written for the people who own the outcome instead. That means general partners of real estate and infrastructure funds, and the family offices that allocate to both alongside their operating businesses. It takes no position on any vendor. It offers evidence on where machine reading is dependable and where it is not, and a design that connects acquisition to hold, so that the most documented asset class can finally read its own paperwork.

2. The 2026 Backdrop: Repriced Assets, Stale Numbers

Start with the debt calendar, because nothing else in real assets moves on a schedule. Seventeen percent of outstanding US commercial and multifamily mortgages, $875 billion, matures during 2026. That is 9 percent below the $957 billion that came due in 2025, which says the maturity wave has crested without saying anything about how it resolves (MBA, 2026). In Deloitte's 2026 outlook survey of more than 850 senior executives at real estate owner and investor organisations, over half report property loan maturities in the coming year, and 21 percent expect to pay them off in full (Deloitte, 2025).

How the rest resolves is visible in the classifications. Trepp's headline CMBS delinquency rate stood at 7.54 percent in April 2026, with offices at 11.69 percent after January's all-time high of 12.34 percent. Add loans past maturity but still paying interest, the category politely labelled performing matured balloon, and the rate rises 152 basis points to 9.06 percent (Trepp, 2026). The New York Fed documented the same behaviour on bank balance sheets: lenders extended impaired mortgages to avoid recognising losses, crowding out new origination by around 5 percent and building a maturity wall equal to 27 percent of bank regulatory capital by the end of 2023 (Crosignani & Prazad, 2024). A lender's forbearance tells an equity owner about the lender's balance sheet, not about the asset.

Capital, meanwhile, is rotating back in. Deloitte's sentiment index recovered to 65 from a trough of 44 in 2023, and 83 percent of surveyed executives expect revenue growth this year. Alternative lenders

supplied 24 percent of US commercial real estate credit last year against a ten-year average of 14 percent, and $585 billion of dry powder awaited deployment as of August 2025 (Deloitte, 2025). Infrastructure fundraising set a record near $200 billion in 2025, more than 58 percent above the prior year and past the 2022 high of $180 billion (McKinsey, 2026). In McKinsey's survey of roughly 300 limited partners, 51 percent plan to raise infrastructure allocations over the next three years, the highest of any private asset class, ahead of buyout at 35 percent and real estate at 30 percent, with family offices, insurers, and sovereign wealth funds showing the largest one-year jump in intent (McKinsey, 2026). Preqin forecasts private infrastructure assets under management nearly doubling from $1.27 trillion at the end of 2023 to $2.35 trillion by 2029, with private real estate growing from $1.61 trillion to $2.66 trillion over the same window (Preqin, 2024). Family offices already sit on real exposure: UBS's survey of 317 family offices puts real estate at 11 percent of the average strategic allocation, and at 18 percent for US families (UBS, 2025).

Here is the tension this paper sits inside. Money is re-entering the asset class faster than anyone has fixed the information plumbing that embarrassed it between 2022 and 2024. Appraisals still update quarterly or annually. Operator reports still arrive as PDFs, 30 to 45 days after month end. Rent rolls are still rebuilt by hand at every transaction. The 2026 vintage is being underwritten, in most shops, with the same reading capacity that priced the 2021 vintage.

3. Adoption Without Production

The survey evidence describes an industry that has bought the thesis and not yet built the muscle. JLL's 2025 Global Real Estate Technology Survey, covering more than 1,500 senior investor and occupier decision-makers across 16 markets, found 88 percent of investors, owners, and landlords piloting AI, alongside 92 percent of corporate occupiers. The same research identified 56 distinct use cases, with the average organisation running five pilots at once. Against that breadth stands one depth statistic: 5 percent report achieving all of their AI programme goals, and 47 percent have met two or three (JLL, 2025).

Figure 1. AI adoption against AI results in real estate organisations. Eighty-eight percent of investors and 92 percent of occupiers have started AI pilots; 5 percent report achieving all of their AI programme goals. Source: JLL 2025 Global Real Estate Technology Survey (1,500+ senior decision-makers, 16 markets).

Budgets confirm the commitment is real. Eighty-seven percent of investors say their real estate technology budgets grew because of AI, and strategic advice on technology now tops their spending priorities (JLL, 2025). What stands between spend and results is rarely the model. JLL attributes the low maturity to data quality, infrastructure, and change management, and reports that more than 60 percent of investors remain unprepared for at-scale adoption on all three counts. Deloitte's survey reads the

same way from the other side: 19 percent of executives describe their organisations as still early in their AI work, and 27 percent report implementation challenges spanning technical issues, missing expertise, and internal resistance (Deloitte, 2025).

The pattern inside the pilots is familiar from other private asset classes. Administrative work and customer service went first as the cheap proving grounds, with the largest institutions deploying most aggressively (PwC & ULI, 2025). Priorities for the next 12 to 18 months, per Deloitte's respondents, are tenant relationship management, lease drafting, and portfolio management (Deloitte, 2025). Two of those three are document problems.

A reader of this research series will recognise the diagnosis. The constraint on AI value in investment firms is organisational, not technical, and the firms that treat deployment as workflow redesign rather than software purchase pull away from those that treat it as procurement. Real assets add a specific twist: the industry's data is unusually fragmented across operators, property managers, lenders, and joint venture partners, so the redesign crosses company boundaries. That is harder. It is also why the advantage, once built, persists.

4. Underwriting: Reading Everything Once

Underwriting a building or a concession is the assembly of a cash flow argument from documents under deadline. AI enters that work making two different claims, and they deserve different levels of trust.

The first claim is about reading: extracting terms from leases and loan documents, normalising rent rolls that arrive in whatever format the seller's system exports, tying out trailing-twelve operating statements, reconciling the offering memorandum's numbers to the underlying files. This claim is sound. The work is bounded, document-grounded, and verifiable, exactly the profile where current models perform dependably. The practical effect is that diligence stops being a sampling exercise. A team that once read 30 leases out of 300 and extrapolated now reads all 300 and investigates the exceptions. Sampling was a constraint imposed by hours. Once machines remove the constraint, sampling becomes a choice, and choices about diligence depth get priced eventually: by sellers, by lenders, and by the asset itself.

The second claim is about valuing, and it deserves more scepticism and more precision. The benchmark for machine valuation is the human appraisal, whose error structure is well measured. Cannon and Cole, studying 25 years of NCREIF sales, found appraisals deviating from subsequent transaction prices by more than 12 percent on average, with errors that lag the market in both directions: too low in hot markets, too high in cold ones (Cannon & Cole, 2011). Recent academic work shows machine learning correcting part of that error rather than replacing the appraiser. Deppner and co-authors, using NCREIF properties from 1997 to 2021 and 50 covariates, found that appraisal deviations contain structured variation that gradient-boosted trees capture, raising accuracy and removing the structural bias, with the strongest results in apartments and industrial and weaker ones in office and retail (Deppner, Ahlefeldt-Dehn, Beracha & Schaefers, 2025). The honest summary: models find the systematic component of appraisal error, and miss the idiosyncratic component that makes a specific tower or toll road worth what it is worth. Commercial assets trade thinly and differ structurally, so the comparable density that feeds residential-style automated valuation never arrives.

The operating conclusion follows. Use machine valuation as a challenge function, not a price source. When the model disagrees with the appraisal or with the deal team's number, the deviation becomes an agenda item with a documented resolution. That preserves the proven value, which is bias correction, without handing the decision to a tool that cannot see what makes the asset specific.

Infrastructure underwriting has the same shape with different nouns. The cash flow argument lives in concession terms, availability and performance regimes, indexation mechanics, curtailment provisions, and termination compensation. The documents are fewer and the stakes per clause are higher, which strengthens the case for machine-assisted extraction with human verification of every priced term.

One behavioural note belongs in any underwriting discussion. Deal teams anchor on the seller's pro forma, and the anchor survives contact with evidence unless the process forces a confrontation. An extraction layer that rebuilds the cash flow from source documents, independently of the offering materials, is the cheapest anchor-breaking device the industry has ever had.

5. Lease Abstraction: Boring Work, Priced Consequences

Lease abstraction is the conversion of each lease and its amendments into a structured record: parties, premises, term, renewal and termination options, rent steps, escalations, recovery structures and their caps, co-tenancy conditions, exclusive use clauses, assignment rights, security, and critical dates. It is the least glamorous workflow in real estate. It is also the point where the asset class's information architecture either gets built or does not.

The reason is simple. In real assets, cash flow is a contractual outcome before it is an accounting one. The lease, the concession, the power purchase agreement is where the money legally lives. Underwriting reads those contracts to price the asset. Operations read them to issue invoices and reconcile recoveries. Monitoring reads them to see risk coming. Today those are usually three separate readings by three separate teams, producing three inconsistent files, none of them version-controlled, at three different moments in the deal's life.

The cost of bad abstraction is specific, not abstract. An unflagged early termination option reprices an exit. A missed co-tenancy clause turns one anchor bankruptcy into a corridor of rent reductions, because nine in-line leases reference the anchor by name. A misapplied cap on recoverable expenses leaks money for years, in amounts small enough that nobody investigates and recurring enough to matter at the exit multiple. Critical dates are the sharpest edge: a renewal option that opens and closes unnoticed hands negotiating power to the counterparty as a gift.

Extraction models have made first-pass abstraction nearly free, and the vendor market quotes accuracy in the mid-to-high 90s. Treat those numbers as marketing until reproduced on your own documents. Deloitte's caution is the right one:

> *“Generative AI may summarize standard leases well but may struggle with unique lease terms, though outputs can improve over time with human intervention.” (Deloitte, 2025)*

The terms that price risk are exactly the unique ones. The right design is therefore machine with provenance rather than machine versus human: every extracted field links back to the clause it came from, a person confirms the fields that carry economics (options, co-tenancy, recoveries, termination), and the confirmation is logged. Sampled re-verification then runs continuously, the way a fund audits anything else it relies on.

Done this way once, at acquisition, the abstract stops being a diligence by-product and becomes the first deposit into the asset file the rest of this paper describes. Done as a one-off vendor project, it dies in a folder, and the firm pays to rebuild it at refinancing, at recapitalisation, and at exit.

Infrastructure substitutes its own documents, concession agreements, offtake contracts, and O&M agreements, and the same logic holds with higher stakes per page. A single indexation formula or

change-in-law clause carries more value than any office lease, and deserves the same extraction-plus-confirmation treatment.

6. Asset-Level Monitoring: Where the Hold Period Leaks

An underwriting model is a photograph. The asset starts moving the day the deal closes, and for the next decade the only question that matters is whether the owner sees the movement before the market reprices it or the lender acts on it.

The asset class's standard instruments see slowly. Appraisals are the canonical case. MSCI's annual comparison of appraised values with subsequent sale prices found US valuations diverging by 6.8 percent on a weighted average basis, the widest of any region (MSCI, 2024). The divergence is lag with a direction. Appraised values of US central business district offices had fallen 43 percent from peak by mid-2024, while transaction-based pricing for the same segment stood 51 percent below peak. Translated into performance, 2023 returns were minus 22 percent measured by appraisal and minus 33 percent measured by transactions (MSCI, 2024).

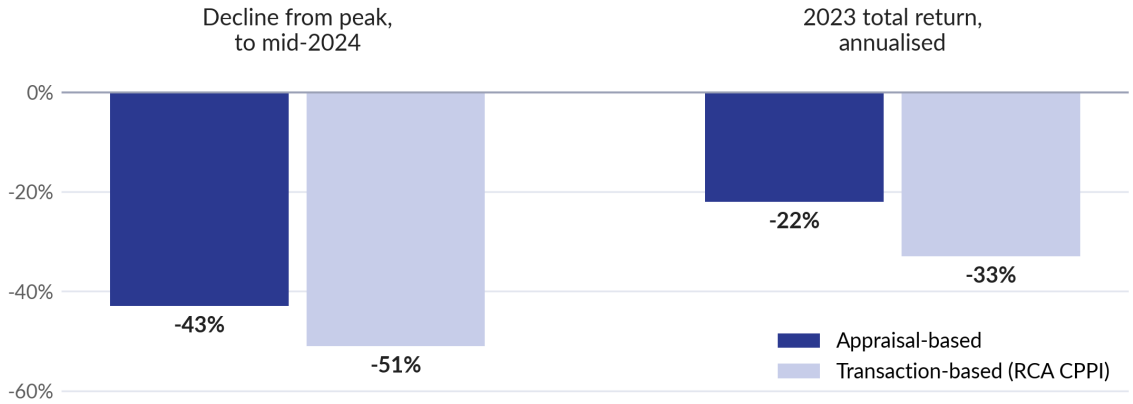

Figure 2. Two measures of the same downturn: US central business district offices. Appraised values fell 43 percent from peak by mid-2024; transaction-based pricing (RCA CPPI) fell 51 percent. Annualised 2023 total returns: minus 22 percent appraisal-based against minus 33 percent transaction-based. Source: MSCI Research (Costello, December 2024).

Infrastructure runs the same physics with the sign flipped. In EDHEC's survey of 79 institutional investors and valuation providers, 76 percent report selling assets above the latest net asset value, typically by 6 to 20 percent, while more than 60 percent rely primarily on management forecasts with limited independent challenge, and only a third adjust valuations during market stress (Whittaker, 2025). Conservative marks feel safe and flatter realised returns. They are still wrong, and an allocator steering by them is steering by a number that is not a price.

For monitoring, the direction of the error matters less than what the error proves: official numbers move after the asset does. Anyone who waits for the appraisal, the quarterly report, or the lender's notice is reading history. The signals that move first sit lower down, at the asset, in data the owner already receives or has the right to receive.

In a commercial building, early looks like this: arrears ageing and payment latency; tenant sales reports where the lease requires them; an anchor's credit event read against every co-tenancy clause that names it; renewal windows opening across the rent roll; expense lines drifting against both budget and the recovery caps that decide who absorbs the drift; capex burn against the engineering report; and energy intensity against a regulatory cap. That last item now carries a price. New York's Local Law 97 has set

emissions caps on most buildings over 25,000 square feet since 2024, with annual compliance reports due from May 2025 and penalties of $268 per ton of CO2-equivalent above the cap; the caps tighten in 2030 (NYC Department of Buildings, 2026). A building's energy meter has become a covenant.

On the debt side of the same asset: coverage headroom and its trajectory, cash sweep triggers, extension conditions and the tests attached to them, and the reporting deadlines in the loan agreement, where lateness is one of the oldest distress signals in lending. Trepp's performing matured balloon category, 152 basis points of stress sitting outside the headline delinquency rate, is what deferred recognition looks like at industry scale (Trepp, 2026).

In infrastructure: availability against the penalty bands in the concession, production against P50 estimates, O&M cost drift against the contract's escalation formula, counterparty credit on offtake, and the calendar of indexation resets and regulatory determinations.

None of this data is exotic. Most of it already arrives, monthly, as PDFs from property managers and operators, and goes unread at portfolio scale because reading it took analyst hours nobody budgeted. Human vigilance also decays in exactly this setting: signals are rare, portfolios are wide, and the person assigned to asset 23's management report has more vivid demands on their attention. The machinery that fixed extraction at acquisition fixes reading during the hold, at lower marginal cost, because operator reports repeat in nearly identical form every month.

7. The Persistent Asset File

The thread running through the three workflows is a single design decision: whether the knowledge assembled at acquisition persists. This section describes the version of yes. Four layers, each with a clear owner and a checkable output.

The contract layer is built at close. Every cash-flow-bearing document, leases and amendments, loan agreements, concessions, offtake and O&M contracts, joint venture agreements, is extracted into structured, versioned records with field-level provenance back to the clause. Amendments trigger re-extraction, and the amendment history is itself a tracked signal. The lease abstracts from diligence are not summarised into the investment memo and discarded; they are the first deposit into the file.

The operating layer maps what arrives during the hold onto what was promised at close. Operator and property manager packages are ingested when they land, and their numbers post against the contract's defined terms rather than a generic chart of accounts: recoveries as the lease defines them, availability as the concession formula computes it. The underwriting model lives here too, as the working baseline rather than an archive. Every actual posts against the assumption that justified the purchase price, and the variance between them becomes the portfolio's most informative time series.

The signal layer turns the file into attention. Three families. Deterministic contract mathematics: critical dates approaching, option windows opening, covenant headroom levels and trajectories, basket and cap usage. Anomaly detection within and across assets: margin drift against peers, arrears latency, energy intensity against the cap path, reporting that arrives late. External overlays mapped to the portfolio: tenant and counterparty credit events read against co-tenancy and offtake exposure, sector rent and capital markets moves, rate resets, and regulatory schedules.

The judgment layer is deliberately human. Signals route to named owners with severity tiers and response deadlines. Investment committees review exceptions and trajectories instead of reporting packets. Management tracks two numbers that barely exist in the industry today: time from asset event

to firm awareness, and time from awareness to action.

A vignette shows the design earning its keep. A retail anchor misses two consecutive sales-report deadlines and its arrears age from 15 to 45 days. The contract layer knows that nine in-line leases name that anchor in co-tenancy clauses with rent relief triggers, and that two of the nine are inside their renewal windows. No covenant has tripped and no appraisal will move for two quarters. As a pattern, this is a watch-list entry with months of option value attached: re-lease, renegotiate, restructure the anchor relationship, or prepare the lender conversation before the lender starts it. The infrastructure version fits in one sentence: availability drifting toward the penalty band while O&M overtime climbs and an indexation reset approaches is a priced problem long before it is a reported one.

Note what this design is not. It is not a product category, and nothing here requires custom research models. The extraction capability is the one deal teams already license; the operating data already flows; the missing piece is the decision that the file persists, with owners, instead of dissolving into folders at close. That is an organisational decision, which is why it is rarer than the technology.

8. Keeping Judgment in the Loop

Machine reading at portfolio scale introduces failure modes of its own, and they are behavioural before they are technical.

Automation complacency arrives first. When extraction is usually right, people stop checking, and the error that matters will sit in the small share nobody re-verified. The countermeasures are structural rather than exhortative: sampled re-verification of extracted fields as a standing process, blind review of randomly selected asset files, and periodic audits of dismissed alerts. An earlier paper in this series, Combating Automation Complacency, treats these mechanisms in detail.

Alert fatigue is second. A monitoring system that cries weekly gets dismantled informally within a year. Thresholds need tuning by severity, duplicates need suppression, and alert volume itself belongs on the governance dashboard, reviewed as seriously as the alerts.

Anchoring works in both directions. The same model output that breaks a deal team's anchor on the seller's pro forma becomes an anchor itself when an asset manager adopts the machine's value without argument. The discipline is symmetry: deviations between model, appraisal, and house view get documented resolutions, whichever direction they run.

Confidentiality is a design requirement, not a compliance afterthought. Rent rolls, tenant payment histories, and loan terms are commercially sensitive and personally identifying. Deloitte's survey found nearly half of organisations interested in synthetic data precisely because this material cannot be sent into public model training (Deloitte, 2025). Architecture that keeps portfolio data out of shared models entirely is the cleaner answer.

Last, say less than the system does. Limited partners now ask about AI capability as part of allocation diligence, and they are alert to overstatement. A firm that can show time-to-awareness metrics and a living asset file needs no adjectives.

9. Conclusion and Recommendations

Real assets entered 2026 with capital flowing back, a maturity calendar forcing decisions, official numbers that move after the fact, and an adoption boom that is wide and thin. Underneath all four sits the same structural fact: this is an asset class whose cash flows are written down in contracts, serviced by monthly documents, and read, today, by almost no one at portfolio scale.

For managers and allocators who want the advantage, the working agenda is as follows:

1. Build the contract layer at the next acquisition rather than waiting for a platform decision, then backfill the existing book in order of equity at risk and nearness of debt maturity.
2. Keep the underwriting model alive as the monitoring baseline. Report actuals against acquisition assumptions quarterly, and treat the variance series as the portfolio's primary early-warning instrument.
3. Treat lease and concession abstraction as infrastructure, with field-level provenance and human confirmation of every term that carries economics. Reproduce any vendor's accuracy claim on your own documents before relying on it.
4. Move monitoring cadence from the calendar to arrival. The trigger for analysis is a document landing, an amendment, or an external event touching the portfolio, never the quarter end alone.
5. Make critical dates machine-owned and responses human-owned. No renewal window, extension test, or regulatory deadline should be discovered after it closes.
6. Hold concentration views as continuously refreshed facts: tenant and counterparty exposure, sector, lender, rate reset profile, and regulatory cap exposure across the book.
7. Install complacency countermeasures on day one: sampled re-verification, blind file reviews, dismissed-alert audits, and alert volume on the governance dashboard.
8. Tell limited partners what was built, in operational terms. Demonstrated monitoring capability is scarcer than demonstrated underwriting capability, and allocators can tell the difference.

Real assets have always written everything down. For the first time, reading all of it, continuously, costs less than not reading it. The firms that act on that arithmetic will know their buildings and concessions better than the market does, and in an asset class repriced by information lag, that is the whole game.

References

Cannon, S. E., & Cole, R. A. (2011). How Accurate Are Commercial Real Estate Appraisals? Evidence from 25 Years of NCREIF Sales Data. The Journal of Portfolio Management, 37(5), 68-88.

Crosignani, M., & Prazad, S. (2024). Extend-and-Pretend in the U.S. CRE Market. Federal Reserve Bank of New York Staff Reports, no. 1130, October 2024.

Deloitte. (2025). 2026 Commercial Real Estate Outlook. Deloitte Center for Financial Services, 2025.

Deppner, J., Ahlefeldt-Dehn, B., Beracha, E., & Schaefers, W. (2025). Boosting the Accuracy of Commercial Real Estate Appraisals: An Interpretable Machine Learning Approach. The Journal of Real Estate Finance and Economics, 71(2), 314-351.

JLL. (2025). 2025 Global Real Estate Technology Survey. Jones Lang LaSalle, October 2025.

McKinsey. (2026). Global Private Markets Report 2026. McKinsey & Company, February 2026.

Mortgage Bankers Association. (2026). 2025 Commercial Real Estate Survey of Loan Maturity Volumes. MBA, February 2026.

MSCI. (2024). The Risks of Appraisal Uncertainties in Private Real Estate (J. Costello). MSCI Research, December 2024.

NYC Department of Buildings. (2026). Local Law 97: Greenhouse Gas Emissions Reductions. City of New York.

Preqin. (2024). Future of Alternatives 2029. Preqin Ltd, September 2024.

PwC & ULI. (2025). Emerging Trends in Real Estate 2026. PwC and the Urban Land Institute, 2025.

Trepp. (2026). CMBS Delinquency Reports, January to April 2026. Trepp Inc., as reported in Multi-Housing News, February to May 2026.

UBS. (2025). Global Family Office Report 2025. UBS AG, May 2025.

Whittaker, T. (2025). Fair Value or Fair Guess? Inside the Engines of Infrastructure Valuation. EDHEC Infra & Private Assets Research Institute, June 2025.

About the Author

Dr. Leigh Coney is the Founder and Principal Consultant of WorkWise Solutions. He holds a PhD researching how humans interact with emerging technology, and has spent over a decade at the intersection of AI, behavioural science, and organisational design. His work focuses on decision making in high stakes environments, with particular attention to why sophisticated AI systems fail to achieve adoption and how governance and deployment systems can be designed to account for human cognitive and social dynamics.

Dr. Coney advises private equity firms, family offices, private credit funds, real estate and infrastructure funds, independent sponsors, and investment banks on AI strategy, governance design, and psychology informed change management. This paper is part of an ongoing research series on responsible AI adoption in financial services. Previous publications include AI in Private Credit: Underwriting, Covenant Surveillance, and the Monitoring Gap (Q2 2026), The AI Exit Premium: Preparing for Buyer Scrutiny (Q2 2026), Agentic AI Governance in Private Equity (Q1 2026), and Combating Automation Complacency (Q1 2026).

About WorkWise Solutions

WorkWise Solutions builds secure, purpose built AI systems for private equity, venture capital, family offices, private credit funds, real estate and infrastructure investors, and investment banking firms. The firm specialises in zero retention AI architecture that ensures proprietary deal flow and portfolio data never train public models.

WorkWise's approach is grounded in a core insight: most AI implementations fail because of broken workflows and poor adoption strategies, not because of the technology itself. Every engagement integrates behavioural science and organisational psychology into the technical design, so that AI systems become a durable part of how investment teams actually work.

Contact: leigh@workwisesolutions.org | workwisesolutions.org