Human Technology Institute

UTS

AI ROI: A practical framework for assessing the return on AI investments

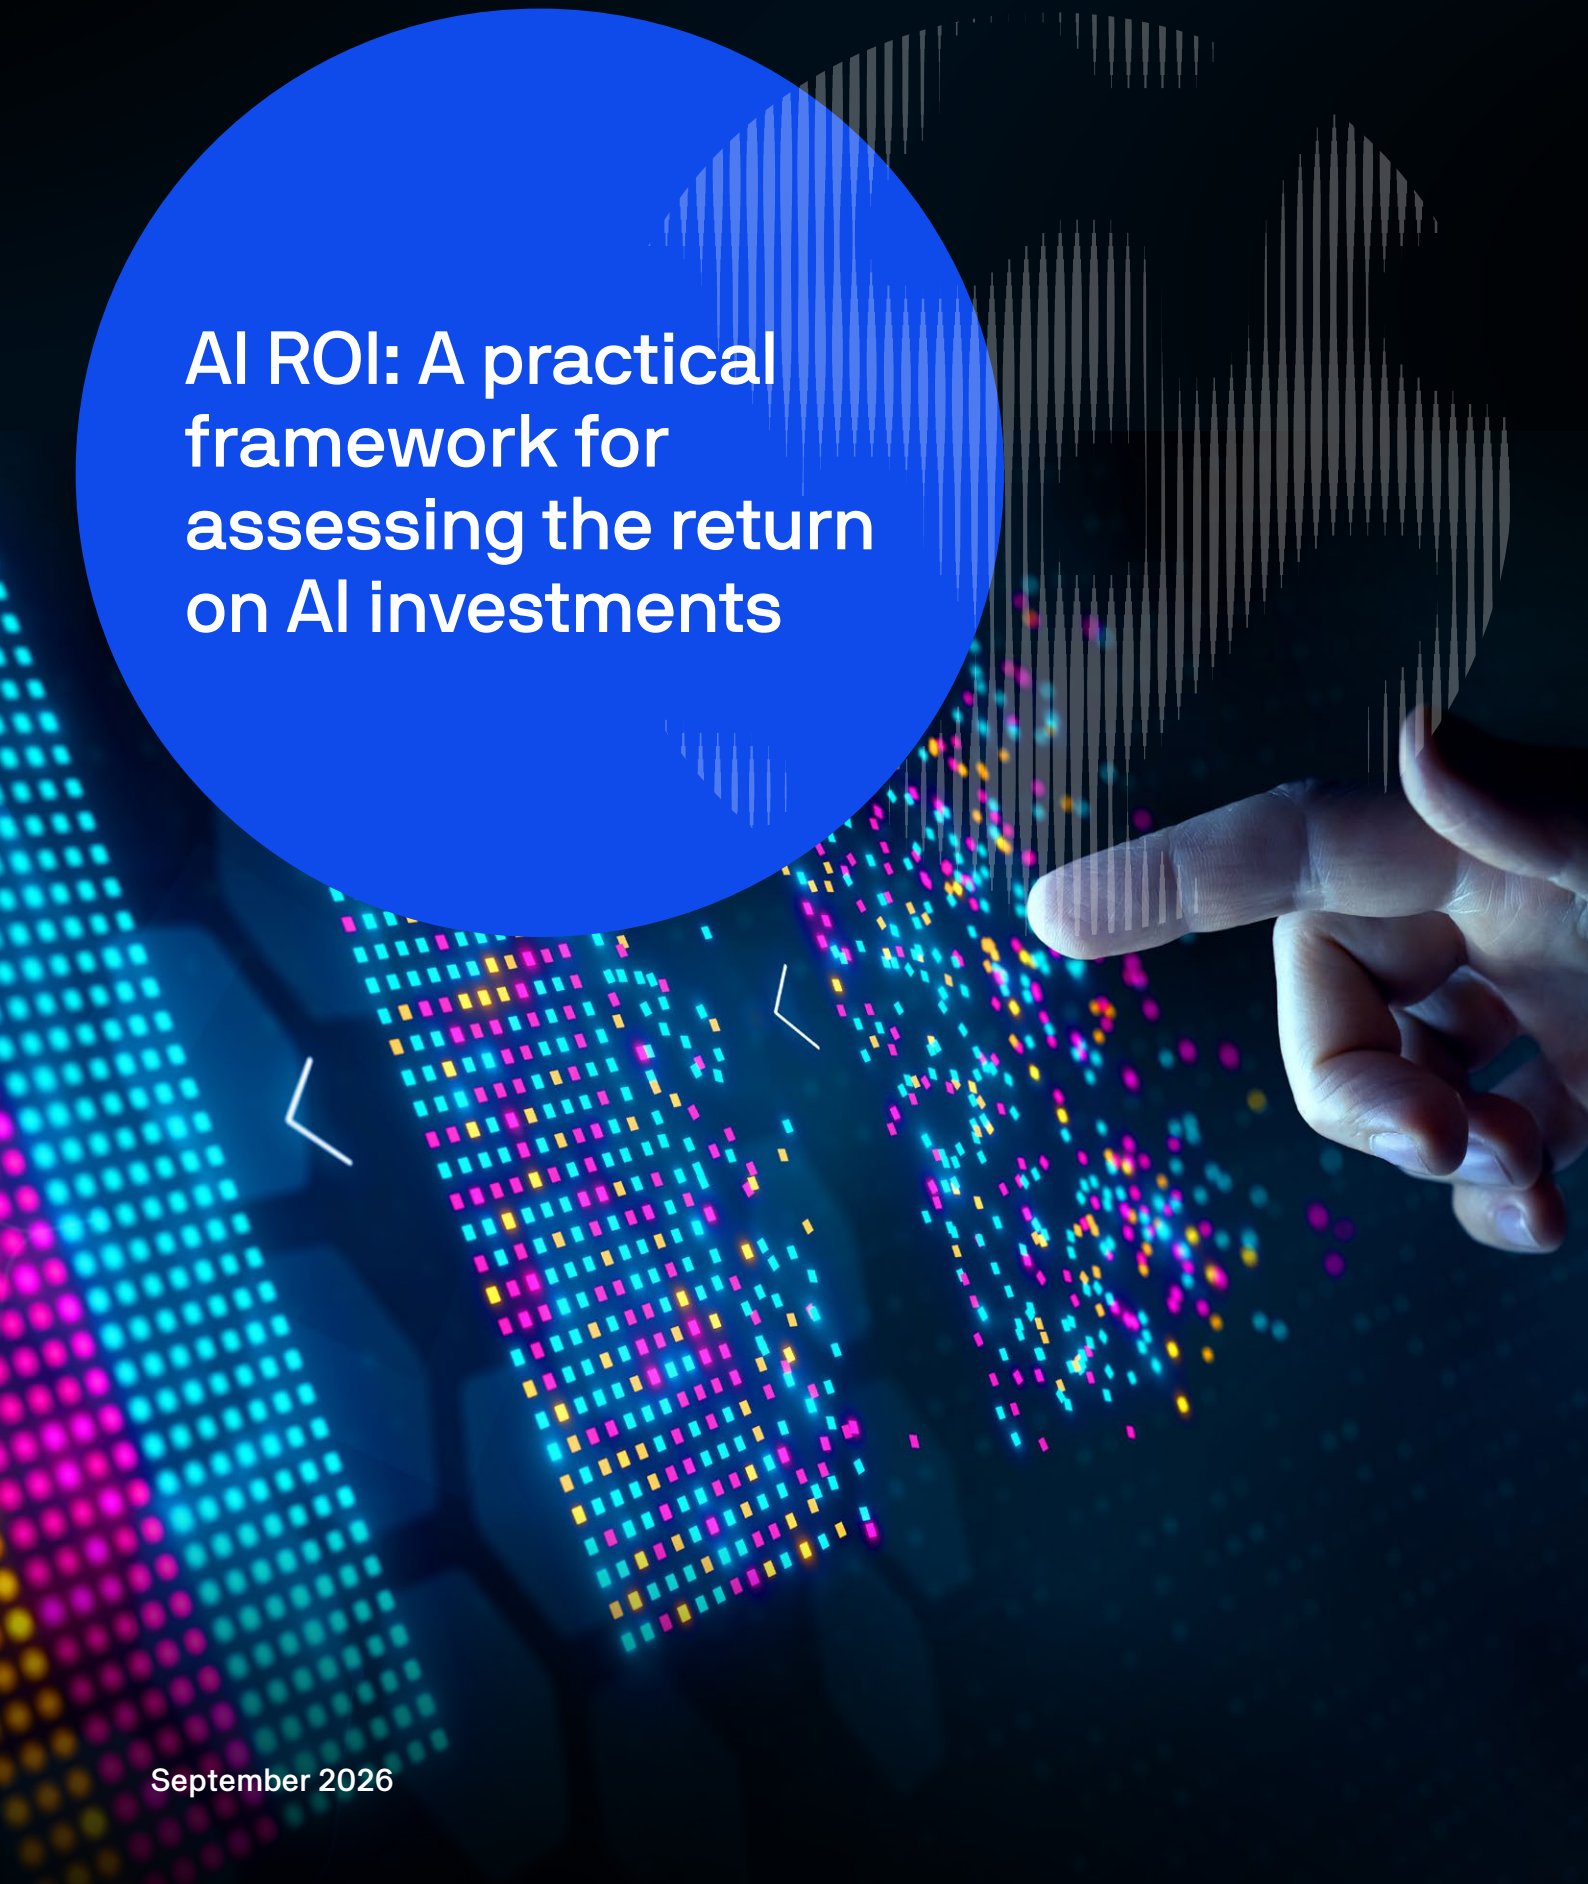

September 2026

The UTS Human Technology Institute (HTI) is an impact-oriented institute building human values into new technologies. Bringing together policy, legal and technical experts, HTI provides independent expert advice, policy development, capability building, and data science solutions to support government, industry and civil society.

Authors: Nicholas Davis, Gaby Carney, Andrea Jachs and Brad Daffy.

Acknowledgement of Country

UTS acknowledges the Gadigal people of the Eora Nation, the Boorooberongal people of the Dharug Nation, the Bidiagal people and the Gamaygal people upon whose ancestral lands our university stands. We would also like to pay respect to the Elders both past and present, acknowledging them as the traditional custodians of knowledge for these lands.

Contents

The challenge of measuring AI ROI 2

A structured framework for AI ROI 3

Step 1: Focus on assessing ROI for a specific, targeted use case, not a general capability or system 4

Step 2: Approach ROI with the implementation phase in mind 5

Step 3: Take a broad and realistic view of the benefits 6

Step 4: Ensure all relevant costs are understood and modelled 9

Bringing the framework together 10

Conclusion 11

Appendix: Questions for Evaluating AI ROI 12

The challenge of measuring AI ROI

Most organisations are failing to measure the ROI from their artificial intelligence investments effectively. This is not because they lack the will or the skills, but because they are trying to apply frameworks designed for deterministic software to probabilistic, adaptive systems.

Traditional technology ROI assumes that benefits can be narrowly defined, that they will scale linearly, that performance is stable, and that risks are well-known. The most recent, flexible and powerful forms of AI break every one of these assumptions. Generative AI hallucinates, agentic AI requires new governance forms, and staff perceptions of benefit don't always survive rigorous measurement: a 2025 trial found that experienced developers using AI coding tools believed they were 20% faster, while controlled measurement showed they were 19% slower when using the tools.[i]

Meanwhile, the pressure to demonstrate returns is intensifying. For example, in 2025, 61% of business leaders said they faced increasing pressure to show returns on AI investments compared to earlier.[ii]

Boards and executives rightly want to know how much value, if any, they are generating from AI, but they lack the tools and methodologies to measure it credibly. This article provides a practical framework to narrow that gap and equips boards and senior leaders to oversee its application.

A structured framework for AI ROI

Boards and executives increasingly recognise that defining AI value is only the first challenge and that realising it over time is the harder test.

Effective AI ROI assessment requires more than a single calculation or business case. It demands both a clear structure for defining and modelling value, and disciplined oversight of whether that value is realised as AI is implemented and scaled. The framework that follows reflects this distinction.

Framework for assessing AI ROI

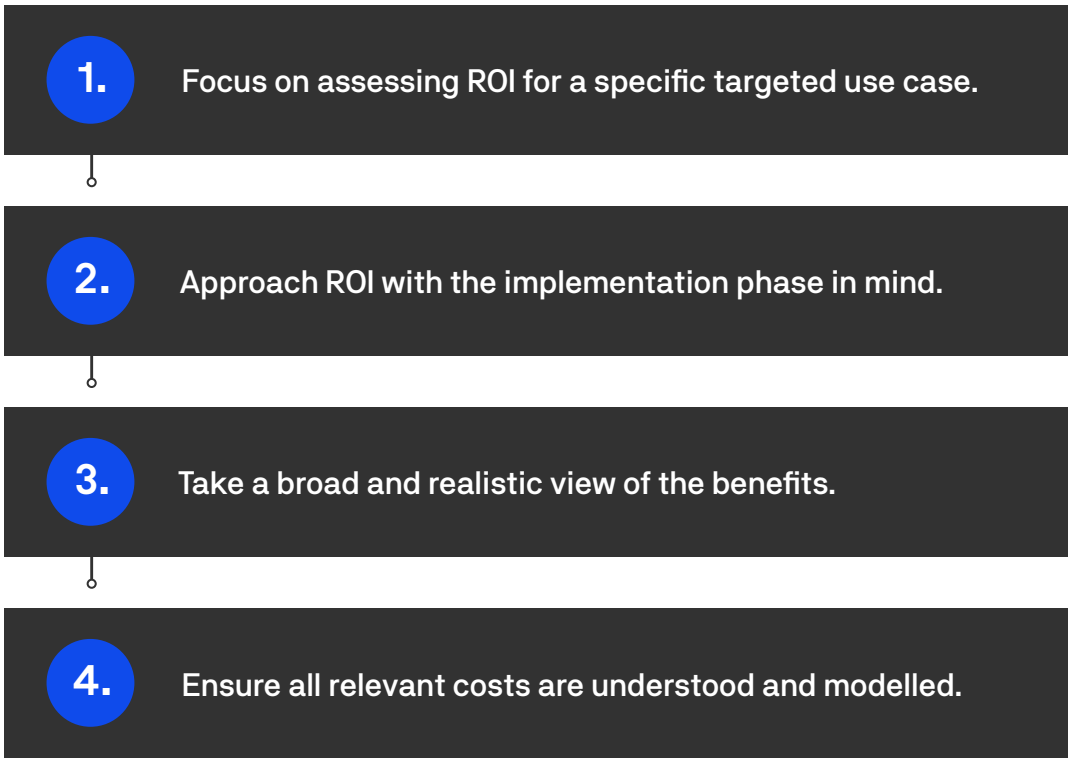

Step 1: Focus on assessing ROI for a specific, targeted use case, not a general capability or system

AI is not a single category of investment. A general-purpose generative AI assistant rollout (with a user-based cost, broad deployment and diffuse benefits) has a fundamentally different ROI profile from a purpose-built AI tool such as a machine learning model (development-based cost, narrow application, measurable outcomes) or an agentic AI system (token-based cost, elevated risk, transformational potential).

For each use case, organisations should identify the specific business improvement to be achieved and the metrics that will demonstrate its success. A generic goal such as “productivity improvement” is not a useful, measurable objective. Far more useful and specific would be a target such as “Reduce average contract review time from 4 hours to 90 minutes”.

In practice, we see organisations struggle when AI investments are framed too broadly (i.e. “rolling out GenAI” or “building AI capability”) because benefits become diffuse, ownership is unclear, and ROI is difficult to defend with boards. By contrast, organisations that anchor investment decisions to well-defined use cases with specific operational outcomes are far more likely to progress beyond pilots. This use-case focus enables clearer accountability, more credible measurement over time, and sharper decisions about where to scale, pause, or stop AI deployments as evidence emerges.

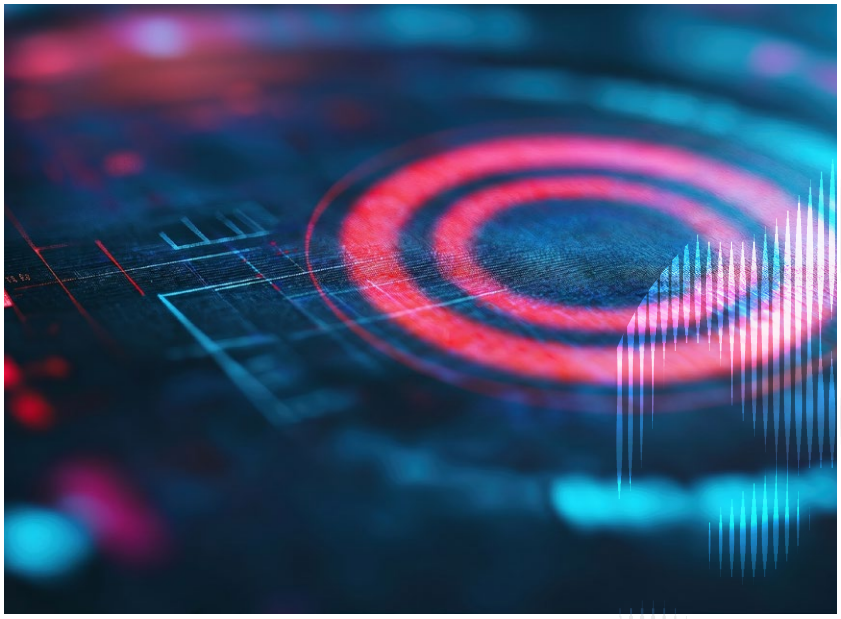

Step 2: Approach ROI with the implementation phase in mind

AI investments typically map to, and often progress through, three different phases. Each has a distinct ROI profile with different benefits, costs and risks rising or falling in importance.

Ideally, investments in the earlier phases will have pre-defined 'stage-gate' measures that must be met before further investment is granted. This helps to manage risks and costs and allow resources to flow to comparatively better value opportunities as appropriate.

- **Experimentation.** The classic experimentation form is a short, bounded pilot focused on learning and answering the question "can and should we do this with AI?". The primary benefits in this phase are insight into capabilities, identification of risks, and validation of key business case assumptions and measurable benefits. Organisations should clearly define the pilot period, the investment thresholds for it, and they should define and baseline clear success measures before the pilot begins.[iii] It is important that pilots come to an end, and that success is evaluated before deciding whether to proceed to integration.

- **Integration.** Once capabilities and risks are well understood, an AI use case is ready to be deployed operationally in order to answer the question "what does it look and feel like if we do this with AI?" The integration phase means working to ensure that the AI solution is practically embedded in organisational workflows and that it is practically used in relevant 'real world' contexts (to withstand the frictions and messiness of daily work). Workflow design and change management are critical investments and areas of focus.[iv] Adoption may be deliberately limited during this phase, and AI investments should pay particular attention to translation efficiency: how well the new solution performs at the boundaries of team handoffs and uneven processes.

- **Scaling.** Once a use case is well-integrated into workflows, it can be extended across the organisation to a wider set of users and systems, eventually becoming "this is now the way we do things with AI". This is where benefits compound and risks are managed or reduced, and more traditional ROI measures such as revenue impact, cost reduction and customer outcomes become directly applicable. When scaling, adoption becomes a critical factor, as does system performance. Small changes in these factors can have large impacts at scale.[v]

From a governance perspective, failure to align ROI expectations with the different implementation stages is one of the most common causes of AI investment underperformance. Boards are frequently presented with 'pilot-stage results' as evidence of future value, only to see benefits diminish as operational complexity, integration costs and adoption challenges emerge. What we are seeing is that treating ROI as phase-specific enables more credible stage-gate decisions, reduces sunk-cost bias, and supports stronger accountability as AI investments move closer to business-as-usual.

Step 3: Take a broad and realistic view of the benefits

Whichever phase an AI use case is in, it's important to take a broad view of the potential benefits. This helps the organisation to gain a clearer understanding of the range of benefits that AI investments achieve over time.

Organisations often focus on the direct benefits that AI investments can generate, such as time savings, reduced customer churn, and improved quality. It is also important to account for the indirect benefits that the organisation experiences as a result of their investment.

Indirect benefits can include:

- **Better data foundations:** the data uplift required to implement one AI system or use case is likely to have broader benefits for the organisation's data governance[vi];
- **Improved AI literacy:** the benefits of better practical understanding of the technology, the risks and how to manage them, across the workforce;
- **Improved risk management capability:** the process of identifying and controlling for AI related risks should deliver stronger control environments more generally[vii];
- **Better employee experience:** supporting staff with AI tools can improve employee engagement and enable staff and teams to deliver within reduced timeframes[viii];
- **Enhanced customer experience:** supporting customers with AI tools can improve access, navigation, consistency, understanding, and trust; and
- **Improved business experience:** supporting businesses with AI tools to enable complex applications can expedite time to a new business opening and contributing to economic growth.

From a board and executive perspective, a narrow focus on immediate financial benefits can mask whether AI investments are genuinely strengthening the organisation. We see more successful AI portfolios where indirect benefits such as data foundations, risk management capability and workforce readiness are explicitly recognised as sources of strategic value. These benefits often determine whether future AI investments can scale with confidence, speed and trust, or instead require repeated remediation and rework.

Testing realisation assumptions

For AI investments at the integration and scale phases, three key factors determine whether the direct and indirect AI benefits are actually realised. When the investment on an AI use case is material or significant, boards and senior leaders should treat these as critical assumptions that underpin the AI business case, and track them with the same discipline applied to any major investment:

1. **Adoption. What proportion of the target user base is actively using the capability?**

   Adoption is a critical factor for AI use cases that are employee or customer focused. Where this varies widely it can significantly affect benefit realisation. If pilots have shown real benefits and well-managed risks, leaders should be consciously measuring and incentivising uptake of the AI use cases. The most powerful driver of adoption is not executive endorsement alone, but rather whether direct managers actively support, model, and make time for AI tool use. If middle management is not enabled and incentivised (i.e. not included in transparent KPI targets), adoption will stall regardless of top-down commitment.[ix]

---

2. **Translation efficiency. How well does the targeted business improvement translate into broader organisational value?**

   The data show that even if individual use of an AI tool produces measurable benefits, this may not translate into higher productivity at the team or organisational levels. Sources of translation loss can include technical friction related to poorly-designed workflows, poor AI literacy, legacy system constraint or data quality at point of input. It can also comprise cultural factors, such as teams in which AI use is stigmatised as "cutting corners"; where time saved is immediately absorbed by new and additional tasks; or where informal norms discourage reliance on AI outputs. Management should be aware of both kinds of friction, and identify and estimate them. This includes identifying corresponding opportunities to improve the translation from individual to organisational benefits.

---

3. **Accuracy and performance. How well does the AI use case perform, and can we rely on it to deliver the quality and consistency required?**

   Both generative and agentic AI systems have extraordinary capabilities, but they also carry risks that may impact their accuracy and performance. Large Language Models can hallucinate, and agentic AI systems can drift from their goals and can be brittle when carrying out complex tasks.

What we hear from boards and executives is that adoption, translation efficiency and performance are frequently the weakest links between approved AI investment and realised returns. We routinely see pilots demonstrating promise, only for benefits to erode as solutions encounter operational complexity, uneven adoption and trust concerns at scale. Treating these factors as critical realisation assumptions, rather than implementation side-issues, strengthens accountability and enables more credible oversight of AI investments over time.

Step 4: Ensure all relevant costs are understood and modelled

Standard ROI models naturally incorporate direct costs related to technology acquisition, testing and integration; and often extend to data quality uplift and risk controls. In addition, for AI investments it is important to recognise key vendor terms that can impact costs such as contract duration, business resilience requirements and pricing changes, to ensure an accurate assessment of AI use case cost structure over time.

In line with the integration phase of the AI use case, organisations should also directly include a view of indirect costs, which can be split into two categories:

1. **Operational indirect costs** recur with each deployment. These include specific system training, anticipated productivity dip during adoption, process redesign, and other change management investments. Change management is consistently the largest determinant of whether AI investments deliver projected ROI. These costs may decline over time for subsequent AI system deployments as employees and managers become accustomed to new ways of working. However, they should always be explicitly recognised and estimated by use case.

2. **Structural indirect costs** are often essential for reliable and long-term benefit realisation. These are more general, organisational-level assets. They tend to be front-loaded costs that create enduring capability and can be attributed across a portfolio of investments. They include AI governance structures, risk management systems, AI policies and practices, data infrastructure, and regulatory compliance capabilities.

How to treat risk in AI ROI

For AI investments, risk does not simply affect costs. It directly affects whether projected benefits can be relied upon, sustained, and scaled. In practice, we see otherwise promising use cases constrained by trust-related issues such as data integrity concerns, model performance variability, insufficient human oversight, or regulatory uncertainty. These risks rarely eliminate value altogether, but they materially reduce confidence in benefit realisation.

As a result, risk should be treated as an input into valuation, not just an item to be mitigated. Organisations should explicitly assess how trust, governance and control maturity influence the certainty of projected returns, and reflect this through confidence-weighted or scenario-based ROI estimates. Stronger controls, explainability and assurance mechanisms may increase upfront costs, but they frequently improve long-term returns by reducing volatility, rework and deployment delays particularly as AI use cases scale.

From our experience we see AI ROI assessments consistently underestimate total cost—because indirect and structural costs are treated as secondary or absorbed “somewhere else” in the organisation. As AI use cases progress beyond pilots, costs associated with integration, change, governance, risk management and regulatory readiness frequently exceed initial expectations and materially affect realised returns. Organisations that explicitly model these costs early and revisit them as AI scales are far better positioned to make defensible investment decisions and avoid value erosion through unplanned remediation.

> “The key things when assessing AI ROI are to be specific about the use case and expected benefits, test the realisation assumptions as you go, and be disciplined in reassessing the investment if the benefits are not being achieved.”
>
> — Gaby Carney, Senior Fellow, Strategic AI, Human Technology Institute, University of Technology Sydney

Bringing the framework together

This framework reflects what differentiates AI initiatives that deliver sustained value from those that stall after early promise. By combining expert research into emerging AI systems with insight gained from real-world implementation, it provides boards and executives with a structured way to govern AI investments across their full lifecycle, from initial use-case definition through to scale.

What we have seen in practice is that organisations that apply this discipline are better able to set credible expectations, challenge assumptions as evidence emerges, and make informed decisions about where to accelerate, redesign, or stop investment. At an executive level, this means treating AI not as a collection of isolated experiments, but as a portfolio of value-creating assets—overseen with the same rigour applied to any strategic investment including clear accountability, active stakeholder engagement, and an ongoing focus on trust, adoption, and realised outcomes.

A minimum viable AI ROI approach

While every board may want management to provide a comprehensive assessment of the ROI on AI investments, particularly for large investments, not every organisation will be in a position to implement the framework immediately.

For mid-market firms, or organisations early in their AI maturity, three disciplines will capture the majority of the value:

1. **Define specific business improvements** for each AI use case before deployment, set measurable target metrics that your organisation can practically track, and make sure the board and management agree on these as central to establishing the 'AI business case'.

2. **Measure use and adoption measures.** Track which teams, individuals or processes are using or relying on the capability and how consistently. If use is below expectations, diagnose why before investing further.

3. **Review costs and benefits regularly.** Benefits should be measured in line with the improvements identified through the original business case. Adjust assumptions based on what you learn. Be willing to stop use if the benefits are not being realised or costs seem to climb above the perceived benefits.

Creating specific targets, measuring real-world use, and honest review are the foundation on which the more comprehensive framework can be built over time.

At a board and executive level, the absence of perfect information should not result in the absence of discipline. We frequently see stronger outcomes where organisations adopt a minimum viable approach to AI ROI early, using simple but consistent measures to create transparency and accountability. These disciplines form a practical control mechanism that can scale as AI investments become more material and complex.

Conclusion

The challenge of measuring and realising AI ROI is real, but organisations that succeed approach it as a governance and execution discipline, not a one-off calculation. Those that assess AI investments at the use-case level, align expectations to implementation stage, rigorously test realisation assumptions, and manage AI as a portfolio of strategic assets are consistently better positioned to deliver sustained, defensible returns.

For boards and senior leaders, the value of this framework lies less in precision forecasting than in decision quality. It provides a structured basis to set clear objectives, challenge assumptions as evidence emerges, intervene early when value is at risk, and make confident decisions to scale, reshape or stop AI investments. In an environment of rapid technological change, regulatory scrutiny and workforce impact, this level of discipline is increasingly essential.

Our research and delivery experience indicate that organisations achieving stronger AI outcomes consistently invest in four foundations that improve returns across an AI portfolio:

1. **Strategy alignment.** A clearly articulated AI strategy that links use cases to business objectives, sets risk appetite, and enables prioritisation. For boards, this strategy becomes a standing item for oversight and annual review.

2. **AI governance structures.** Fit-for-purpose, cross-functional governance that integrates risk, technology, legal, workforce and business accountability, supported by standardised risk assessment and clear policies on human oversight and acceptable use.

3. **Workforce engagement.** Genuine engagement with employees in use-case design, transparent transition and reskilling pathways, and active measurement of employee experience alongside adoption and productivity.

4. **Monitoring and measurement infrastructure.** Practical mechanisms to track adoption, performance, quality and value realisation over time that is strong enough to support oversight and assurance, but proportionate to organisational context and respectful of worker rights.

Taken together, these capabilities distinguish organisations that can scale AI with confidence from those that remain trapped in cycles of pilots, rework and unrealised promise. As AI becomes increasingly embedded in core business processes, boards and executives will need this level of integrated insight across Value, Risk and Trust to govern AI investments effectively.

"Having a structured, shared and flexible framework for AI ROI allows boards and management to align on and deliver to the promise of investments with greater confidence, speed and safety."

— Professor Nicholas Davis, Co-Director, Human Technology Institute, University of Technology Sydney

Appendix: Questions for Evaluating AI ROI

Questions for management to ask

Before investing:

- What specific business improvement will this use case achieve, and how will we measure it?
- What stage of implementation is appropriate, and what criteria should we have for transitioning from one stage to another?
- Have we identified all direct and indirect costs, including integration and workflow redesign costs, change management investments, and workforce consultation?
- What are our realisation assumptions (adoption rate, translation efficiency, literacy level), and what evidence supports them?
- How does this investment fit within our broader AI portfolio?
- What is the impact to our key operational risks from implementing this solution?

After investing:

- What specific improvements have we measured against the target metric?
- What is the adoption rate, and what is driving or limiting it, including at the middle management level (where cultural impact can be critically important)?
- What new indirect or unanticipated costs have emerged?
- Are our risk controls operating as intended?
- How does realised ROI compare to projected ROI, and have we updated our realisation assumptions?
- What cross-cutting benefits has this investment built or contributed to (data, literacy, governance) that reduces the cost of the next one?

Questions for the board/directors to ask

Before investing:

- Is the investment aligned with our AI strategy and risk appetite?

- Has management quantified the realisation assumptions and stress-tested them? What reporting will we receive, at what frequency, and using what metrics?

- Who is accountable for the kill decision if gate criteria are not met?

- Have regulatory and workforce consultation obligations been addressed?

After investing:

- How do we ensure reporting is balanced and not selectively positive?

- How is the investment tracking against projected ROI?

- What corrective action is management taking where returns are below expectations?

- What portfolio-level value has this investment created for future deployments?

- Are we confident management would tell us if this investment should be stopped?

Putting the AI ROI framework into practice – a simple, hypothetical example

Step 1: Define a specific use case and business outcome

MarketGenius, a mid-market marketing firm with 150 staff, is evaluating a Microsoft Copilot deployment. Rather than framing the investment as a general "GenAI rollout", management identifies a specific business friction point: the time spent on routine proposal drafting.

Internal estimates indicate that proposal drafting accounts for approximately 12,000 staff hours per year (around 2.5 hours per week for 120 staff). Management sets a clear and measurable target: a 15% reduction in time spent on routine proposal drafting. At a blended staff cost of $120 per hour, this implies a potential gross benefit of $216,000 per annum.

Step 2: Align ROI expectations to the implementation phase

Given that many staff are already using generative AI tools informally, management proposes moving directly to broad deployment rather than running a prolonged pilot. However, the executive team explicitly recognises that early benefits will depend less on the technology itself and more on adoption, workflow integration, and quality controls.

As a result, expectations for the first year are deliberately conservative, with success defined as evidence of sustained use and measurable time savings rather than full financial return.

Step 3: Take a broad and realistic view of the benefits

In addition to the direct time-saving benefit, management identifies a set of indirect and enabling benefits expected from the deployment. These include:

- Improved AI literacy across proposal and account teams;
- Standardisation of proposal content and quality, reducing rework;
- Clearer proposal workflows and approval points, improving transparency; and
- Foundational governance and usage controls that can be reused for future AI use cases.

While these benefits are not all immediately quantifiable, the board explicitly acknowledges that they reduce the cost and risk of subsequent AI deployments and therefore form part of the overall value case.

To avoid overstating value, management stress-tests two critical assumptions before finalising the business case:

1. **Adoption:** Based on staff surveys and structured interviews, the CTO recommends a conservative first-year adoption rate of 50%, reflecting a mix of early adopters, sceptics and teams with limited immediate relevance.
2. **Translation efficiency**: The COO estimates that only 70% of individual time savings will translate into organisational benefit due to workflow constraints, quality assurance steps and approval requirements.

Applying these assumptions reduces the first-year realised benefit to approximately $75,600 ($216,000 × 50% × 70%).

Step 4: Model full costs, including indirect and structural costs

The firm models both direct and indirect costs:

1. **Direct costs:** Copilot licensing at $45 per user per month, equating to $81,000 per year at full deployment.
2. **Operational indirect cos**ts: An additional $25,000 is budgeted for training, workflow redesign and change management support.

This results in a total estimated first-year cost of $106,000.

Bringing it together: board-level oversight and portfolio thinking

On this basis, the investment shows a modestly negative first-year ROI. However, the board and executive team agree that the indirect benefits, particularly AI literacy, workflow clarity and governance foundations, justify proceeding, provided assumptions are actively monitored.

For year two, management revises assumptions, projecting adoption to rise to 65% and translation efficiency to improve to 80% as training costs fall and workflows stabilise. Improved adoption, workflow maturity and reduced rework allow a portion of previously indirect benefits to be realised as measurable time savings. Incorporating these effects increases adjusted benefits to approximately $134,800.

Before approving further scale, the board requests explicit assurance that:

- Quality controls and data handling practices are operating as designed;
- Workforce impacts remain consistent with expectations; and
- Adoption, performance variability and risk indicators are reported alongside financial outcomes.

This ensures that the AI investment is governed not as a one-off technology deployment, but as a value-creating asset subject to ongoing oversight, trust controls and evidence-based decision-making.

Endnotes

i J. Becker, N. Rush, E. Barnes, D. Rein, “Measuring the Impact of Early-2025 AI on Experienced Open Source Developer Productivity” (METR, 2025).

ii S. Lazzaro, ‘The Big AI New Year’s Resolution for Businesses in 2026’ *Fortune* (15 December 2025) <https://fortune.com/2025/12/15/aritficial-intelligence-return-on-investment-aiq/>; Kyndryl, Readiness Report: *The Readiness Tipping Point* (Report, 2025) <https://www.kyndryl.com/content/dam/kyndrylprogram/doc/en/2025/readiness-report.pdf>.

iii For low-risk, off-the-shelf deployments, this stage may be brief or skipped entirely. For higher-risk or novel deployments (particularly where workforce consultation obligations or formal risk assessment processes apply) experimentation may appropriately run longer. Be careful not to simply assume that pilot success will translate into broader benefits without additional costs.

iv Benefits should be tracked against the specific business improvement identified, alongside indirect benefits such as AI literacy and improved team capability. The single largest risk at this stage is integration failure: where the AI does not connect effectively with existing systems, data pipelines, or security architecture.

v Here, the portfolio view becomes critical: scaling one AI use case should reduce the cost experimentation and integration costs of subsequent ones through shared infrastructure.

vi Across all phases, data foundations and AI literacy are often the highest-value indirect benefits because they transfer to compound across future investments. Every AI use case that requires a data quality uplift improves the organisation’s readiness for the next one.

vii Risk management capability delivers value through stronger control environments, reduced remediation costs, and better auditability. Like data and literacy, these benefits transfer across use cases.

viii Every deployment that builds staff capability and confidence with AI reduces the cost and risk of the next one subsequent deployments. Employee experience matters more than its indirect nature suggests, particularly in talent-intensive industries.

ix Note also that higher AI literacy correlates directly with both higher adoption and more effective use. For generative AI systems, prompt quality is directly related to output accuracy and usefulness. A 2026 survey found that significant positive ROI from AI investments nearly doubled from 21% to 42% among organisations with mature data literacy programs. Datacamp, *2026 State of Data and AI Literacy Report* (Report, 2026) <https://www.datacamp.com/business/2026-data-ai-literacy-report>.

Authors

Professor Nicholas Davis
Industry Professor, Emerging Technology and
Human Technology Institute Co-Director

nicholas.davis@uts.edu.au

Gaby Carney
Senior Fellow, Strategic AI, Human Technology Institute
gaby.carney@uts.edu.au

Andrea Jachs
Partner, AI Risk Lead, KPMG Australia

Brad Daffy
Partner, Powered Data & AI, KPMG Australia

Human Technology
Institute

For more information

Human Technology Institute
hti@uts.edu.au

University of Technology Sydney
PO Box 123
Broadway NSW 2007

uts.edu.au

UTS CRICOS 00099F

Stock photographs: Adobe Stock Images (www.stock.adobe.com) and Pexels (www.pexels.com).