

EVALUATING AI ROI

Separating operational leverage from technical hype



Executive Summary

C-suites are currently operating under immense pressure to deploy Artificial Intelligence (AI), yet most treat these investments as necessary sunk costs or vanity projects rather than measurable commercial assets. To move beyond the hype cycle, leadership must apply actuarial precision to tech procurement—mapping structural risks and ensuring every deployment delivers a risk-adjusted return. To drive genuine operational leverage, AI must be structurally designed to improve the EBITDA margin rather than simply inflating the tech stack.

Briefing

The AI Vanity Trap vs. Actuarial Reality

In the current market, AI procurement is often driven by "FOMO" (Fear Of Missing Out) rather than fiscal logic. Boards are authorizing significant capital outlays for Large Language Model (LLM) integrations and generative tools without a clear framework for measuring success. This has led to the "AI Vanity Trap": implementations that look impressive in a demo but add zero measurable value to the bottom line.

From an actuarial perspective, a technology investment is only viable if it improves the business's risk-adjusted yield. In 2026, we must distinguish between "Vanity AI"—superficial chat-bots or siloed pilots—and "Structural AI." Structural AI fundamentally alters the unit economics of the business, automating high-stakes decision-making or removing systemic friction in the middle office. If the technology does not explicitly increase capacity or reduce the cost-to-serve, it is not a commercial strategy; it is an unforced overhead.

Mapping the Risk-Adjusted Return

Evaluating AI ROI requires a cold, objective analysis of the "Technical Hype" versus the "Commercial Asset." Most firms are currently over-investing in the former while ignoring the structural risks of the latter.

Implementation Type	Technical Hype (Vanity implementation)	Commercial Asset (Structural implementation)	Structural Risk Marker
Model Integration	General-purpose LLM wrappers for internal "knowledge management."	Proprietary models trained on private, cleaned data sets for specific underwriting or pricing.	Data leakage and loss of intellectual property to public model providers.
Process Focus	Front-end "help" bots that merely redirect users to existing FAQs.	Middle-office automation of compliance auditing and regulatory reporting.	"Hallucination tax"—the cost of human oversight required to correct AI errors.
Success Metric	User engagement, "cool factor," and marketing optics.	Measurable reduction in Headcount-to-AUM ratio or EBITDA margin expansion.	Regulatory drag if AI-driven decisions cannot be audited or explained.

To achieve a true risk-adjusted return, leaders must account for the "Hallucination Tax"—the necessary cost of a human-in-the-loop to mitigate the systemic risk of automated errors. Without this actuarial buffer, the perceived efficiency gains are often wiped out by the cost of remediation.

The EBITDA Impact: Calculating the AI Efficiency Frontier

To justify strategic AI deployment, we must prove operational leverage through hard numbers. We use the "AI Efficiency Frontier" to determine if a tech investment is diluting or enhancing the margin.

Consider the baseline commercial calculation for a middle-market SME:

- **The Manual Baseline:** A compliance team of 10 people manages 1,000 audits per month. At a fully loaded cost of £60k per person, the cost-per-audit is £600.
- **The AI-Augmented Model:** A firm invests £150,000 per annum in a specialized AI auditing tool. This allows a reduced team of 4 people to manage 2,000 audits per month.
- **The Structural Result:**
 - **Old Total Cost:** £600,000 (Labor).
 - **New Total Cost:** £240,000 (Labor) + £150,000 (Tech) + £50,000 (Oversight/Audit) = £440,000.
 - **Cost-per-Audit:** £220.
- **The ROI:** The business has doubled its capacity while reducing the unit cost by 63%. This is a structural improvement to the EBITDA margin.

If your AI roadmap does not result in a similar downward pressure on unit costs or an exponential increase in capacity without a commensurate increase in headcount, the deployment is likely a vanity exercise.

The Executive Framework for Strategic AI Deployment

True commercial growth requires structural alignment between technology and risk management. I apply three executive heuristics to every AI mandate:

1. **Solve for EBITDA, Not Optics:** Every AI pilot must be tethered to a specific line item on the P&L. If you cannot identify which cost will decrease or which revenue stream will accelerate through capacity expansion, do not authorize the spend.
2. **Map the Data Perimeter First:** AI is a "data-hungry" asset. Before deployment, you must audit your data governance. If your proprietary data is being used to train a third-party model without a clear commercial moat, you are subsidizing your competitors' future growth.
3. **Establish an "Auditability" Score:** In regulated environments, AI is a systemic risk if its outputs are "black boxes." Rate every

implementation on its ability to provide a clean audit trail. If you cannot explain how the AI reached a decision to a regulator, the operational leverage it provides is not worth the corporate risk.

Technology should never be a sunk cost. By applying actuarial rigor to your AI strategy, you transform a hype-driven expense into a resilient, margin-expanding asset.

Mandates & Advisory: Mark Woodruff Consultancy operates exclusively on a B2B basis, providing commercial strategy, fractional board advisory, AI governance, and executive actuarial oversight for financial institutions and scaling SMEs.

To discuss a specific mandate, contact Mark Woodruff Consultancy directly.

Disclaimer: This briefing is for informational and educational purposes only. It does not constitute formal regulatory, legal, or specific corporate financial advice. Professional engagements are subject to a formal letter of agreement.

DISCUSS A MANDATE