

How to Deflates the “AI Bubble” Without Killing Innovation



The AI bubble will not burst because the technology fails. It will burst where delivery discipline is absent.

In Part 1, I argued that bubble pressure is building not in AI models, but in how (according to studies by MIT, Deloitte, McKinsey, etc) organisations are deploying them, capital without discipline, spend without business cases, autonomy without governance.

This risk is not inevitable, but companies that treat AI as a **delivery problem** are already avoiding these traps.

Reframe AI: From Capability to Portfolio

The "AI strategy" concept sounds great but is too general. Start thinking **AI investment portfolio**.

What is the difference?

Small, scoped bets instead of platform transformations. Each initiative is independently funded, with its own success criteria and stop conditions. You are not all-in. You are diversified.

Explicit stop / scale decisions instead of momentum-driven expansion. Every initiative has predefined thresholds: what result triggers more investment, what result triggers a pause, what result triggers a kill. Decisions are made on evidence, not optimism.

Clear business owners instead of cross-functional committees own the P&L impact. A person owns the user experience. Someone owns the risk. Not "the AI team," but a named individual with a stake in the outcome.

This turns the AI (like any other) project from a bet the business made into a managed set of deliverables with a known treatment.

Narrow AI Beats Broad Autonomy

Broad and general systems inflate cost, risk, and testing burden, whilst narrow and specific systems create clarity and accountability:

- **Classification.** Route this email. Flag this transaction. Identify this pattern.
- **Summarisation.** Condense this document. Extract these key points. Highlight exceptions.
- **Anomaly detection.** Does this look normal? Should this be reviewed? Is this pattern breaking?
- **Embedded inside workflows** where humans remain decision owners.

The AI assists. It doesn't replace. It surfaces options. It doesn't execute autonomously.

Why are they different approaches, with different levels of successes?

- Because narrow systems can be **tested, governed, and rolled back** without dismantling entire operations.
- They integrate into existing controls. They don't replace them.
- This addresses governance lag without slowing progress.

You are not relying on new AI governance frameworks but building systems that fit into governance frameworks that already exist.

Governance as a Delivery Feature

Governance is what enables scalability.

Without governance:

- Every initiative is a one-off.
- Every deployment is custom.
- Every risk assessment starts from scratch.

Governance creates **repeatable patterns** that accelerate delivery while reducing risk.

Practical controls that work:

1. **Decision thresholds.** The AI can auto-execute up to £X, auto-suggest up to £Y, and must escalate beyond that. The thresholds are not arbitrary; they are based on materiality and risk appetite.
2. **Human-in-the-loop checkpoints.** At defined stages, a human reviews, approves, or overrides. These aren't bureaucratic gates. They are decision quality checks.
3. **Audit trails by default.** Every decision, every input, every confidence score is logged. Not for compliance theatre. For learning and improvement.

4. **Kill switches defined upfront.** Before deployment, not after failure. What triggers a pause? What triggers a rollback? What triggers an escalation? Everyone knows, and everyone can act.
-

Financial Discipline: AI That Pays or Stops

AI initiatives must **earn continued funding** through metrics that matter, not through storytelling.

What is important to track:

1. **Cost per decision.** What does each AI-assisted action cost? Is it cheaper than the manual alternative? By how much?
2. **Accuracy within tolerance.** You don't need perfection. You need performance within an acceptable range. Define the range. Measure against it. Report deviations.
3. **Exception rates.** How often does the AI escalate? How often does it fail? How often does a human override? Track trends, not just snapshots.
4. **Adoption by role.** Are the intended users using it? If adoption is low, either the problem wasn't real or the solution isn't fit for purpose. Either way, that is a signal.

These metrics prevent sunk-cost escalation and if an initiative isn't performing, you will find out early, and therefore adjust or stop.

The Role of PMO, Data, and Product Thinking

AI fails when it sits between functions, but it succeeds when **delivery ownership is explicit**. And that is where hybrid leadership matters.

1. **PMO discipline** ensures scope control, dependency management, and risk visibility. AI projects aren't exempt from delivery fundamentals. They need them more.
2. **Data discipline** ensures the inputs are fit for purpose. Garbage in, garbage out hasn't been repealed. If the data quality isn't there, no model will save you.
3. **Product thinking** ensures user-centricity and value focus. AI isn't successful when it is technically impressive. It is successful when people use it and outcomes improve.

This is where **AlfaFinTec** sits naturally as a delivery partner who bridges **technical implementation with business strategy**, ensuring AI initiatives are scoped, governed, and measured like the operational systems they are becoming.

Risk-managed experimentation. Outcome-driven scaling. Not hope. **Discipline.**

The Reset Is Coming. Discipline Determines Who Benefits.

The AI bubble will not burst because the technology fails, it will burst where delivery discipline is absent.

- Where capital has been deployed without business cases.
- Where autonomy has been granted without governance.
- Where spend has become political instead of evidence based.

Organisations that treat AI as a **project execution problem**, with clear ownership, constraints, and financial logic, will not just survive the reset, but **they will benefit from it**.

Because what is important is the operational value, and **operational value requires delivery discipline**.

That is not sexy. It is not viral. It won't trend on LinkedIn. But it works.

And when the bubble pressure releases, **working is what matters**.

Wences Alfageme is the founder of AlfaFinTec, an AI transformation consultancy specialising in delivery discipline for financial services. With over 20 years of experience bridging technical implementation and business strategy, he helps organisations deploy AI with the rigour operational systems demand.

#AIDelivery #AIGovernance **#RiskManagement** #DigitalTransformation #ArtificialIntelligence