How Leaders Should Evaluate AI Vendors (Without Getting Sold AI Theater)

Most organizations don't fail at artificial intelligence because the technology doesn't work.

They fail because the *decision-making* around AI is broken long before any tool is deployed.

Executives feel pressure to "do something with AI." Vendors promise transformation. Internal teams want momentum. And before long, six-figure contracts are signed for tools no one fully understands — or can clearly explain.

This is how AI theater happens.

Most organizations don't fail at artificial intelligence because the technology doesn't work.

They fail because the *decision-making* around AI is broken long before any tool is deployed.

Executives feel pressure to "do something with AI." Vendors promise transformation. Internal teams want momentum. And before long, six-figure contracts are signed for tools no one fully understands — or can clearly explain.

This is how AI theater happens.

The AI Vendor Problem Isn't Technical — It's Strategic

AI vendors are very good at demos. They are less good at answering the questions leaders actually need to ask:

- What decision does this improve?
- What risk does this introduce?
- What happens if we *don't* do this?
- How will this look to regulators, customers, or the public in two years?

When these questions aren't asked early, organizations confuse *activity* with *strategy*.

Why Most AI Pitches Sound Impressive (and Why That's a Problem)

Most AI pitches rely on three familiar tactics:

1. **Abstract benefits**
   "Efficiency," "insight," "personalization," "automation" — all real, all vague.
2. **Borrowed authority**
   References to large models, research papers, or brand-name clients without context.
3. **Fear-based urgency**
   "Everyone else is adopting this."

“You’ll fall behind.”
“This is inevitable.”

None of these prove the tool is right for *your* organization.

A Better Framework: Evaluate Decisions, Not Tools

Leaders should evaluate AI proposals the same way they evaluate major strategic decisions — by slowing down the *judgment*, not the innovation.

Here’s a practical framework I recommend.

1. Start With the Decision, Not the Technology

Ask one simple question:

What specific decision will this AI system improve?

If the answer is unclear, abstract, or constantly shifting, that’s a red flag.

AI that doesn’t map to a real decision usually ends up:

- Underused
- Misused
- Or quietly abandoned

2. Identify the New Risks You’re Accepting

Every AI system introduces *new* risks, even when it reduces others.

These may include:

- Reputational risk
- Bias or fairness concerns
- Explainability problems
- Dependency on opaque vendors
- Governance and compliance exposure

If a vendor can’t clearly articulate these risks — or dismisses them — leadership should be cautious.

3. Separate Capability From Accountability

Many AI tools *can* do impressive things.

The harder question is:

Who is accountable when the system is wrong?

If accountability is diffuse (“the model,” “the data,” “the system”), the organization is absorbing more risk than it realizes.

4. Pressure-Test the “Do Nothing” Option

A strong strategy always compares action to inaction.

Ask:

- What happens if we delay six months?
- What knowledge would we gain by waiting?
- What is the *real* cost of not deploying this now?

Urgency that collapses under these questions usually isn’t strategic urgency — it’s sales urgency.

5. Ask How This Will Be Explained Internally and Publicly

If leadership can’t clearly explain:

- Why this AI exists
- What it does
- What it does *not* do
- How it’s governed

…then trust will erode long before value appears.

AI decisions are communication decisions.

The Role of Independent Judgment

Many organizations rely on:

- Internal teams who are too close to the technology
- Vendors who benefit from approval
- Committees that optimize for consensus, not clarity

What’s often missing is **independent strategic judgment** — someone whose role is to ask uncomfortable questions *before* money, credibility, or trust is spent.

This is not about slowing innovation.
It’s about making decisions leaders will still stand behind later.

Final Thought

AI is not a shortcut around leadership responsibility.

It amplifies it.

The organizations that succeed with AI aren’t the ones who adopt the most tools — they’re the ones who make the clearest decisions.