Review and Analysis: *The GenAI Divide — State of AI in Business 2025*

My assessment: **this is a very useful business-strategy report, but the headline "95% of organizations are getting zero return" is considerably stronger than the underlying methodology can comfortably support.** The report is strongest when it explains *why* AI pilots stall and what successful organizations do differently. It is weaker when its directional interview findings are presented as broad market statistics.

The bigger insight is not "95% of AI fails." The bigger insight is:

> **Enterprise AI has an adoption problem, a workflow problem, a learning problem, and an ROI-measurement problem—not primarily a model-performance problem.**

That distinction matters. 2026 MIT v0.1_State_of_AI_in_Bu…

---

1. What the report actually studied

The report is based on:

- 300+ publicly disclosed AI initiatives
- Structured interviews with representatives from **52 organizations**
- Survey responses from **153 senior leaders**
- Research conducted from **January through June 2025**
- A six-month post-pilot window for measuring ROI in portions of the analysis

Importantly, the report itself labels the work **"Preliminary Findings."** It also acknowledges selection bias, varying definitions of success, self-reported implementation outcomes, uneven sample sizes and the difficulty of isolating AI ROI from other operational changes. 2026 MIT v0.1_State_of_AI_in_Bu…

That means this is best treated as a **strong directional enterprise study**, not as definitive statistical evidence that exactly 95% of all corporate AI projects fail.

---

2. The central finding is powerful—but frequently misinterpreted

The headline finding is:

$30–$40 billion invested in enterprise GenAI, yet only about 5% of integrated AI pilots show meaningful P&L impact.

The report calls the gap between experimentation and measurable business transformation the **“GenAI Divide.”** 2026 MIT v0.1_State_of_AI_in_Bu…

But there are really **three different numbers** buried inside the narrative:

Stage	General-purpose LLMs	Embedded/task-specific GenAI
Investigated	80%	60%
Piloted	50%	20%
Successfully implemented	40%	5%

The chart on page 6 is one of the most important in the entire report.

It demonstrates something more nuanced than “AI doesn't work.”

Consumer AI works reasonably well.

ChatGPT, Copilot and similar tools have relatively high use because employees can immediately apply them to:

- drafting
- summarizing
- research
- basic analysis
- brainstorming

The problem occurs when companies attempt to move GenAI into **persistent operational workflows**.

That is where the implementation rate collapses.

2026 MIT v0.1_State_of_AI_in_Bu…

---

3. The report's strongest insight: adoption ≠ transformation

This may be the most important sentence executives should take from the report:

Enterprises do not have an AI experimentation shortage. They have an AI conversion shortage.

Organizations are launching pilots.

Employees are using AI.

Executives have budgets.

Vendors are everywhere.

But relatively few experiments become systems producing measurable recurring business value.

The report finds that more than **80% of organizations explored or piloted general-purpose GenAI**, while nearly **40% reported deployment**.

Yet industry-level structural transformation remained limited.

That distinction between **activity and value creation** is enormously important.

A company can have:

> 10,000 Copilot licenses
> 5 AI committees
> 14 pilots
> 3 innovation labs
> and still have almost no measurable improvement in EBITDA.

That is essentially the GenAI Divide.

---

4. The most revealing section may actually be “Shadow AI”

Page 8 contains one of the report's most strategically important findings:

40% of companies had purchased official LLM subscriptions.

But employees at **more than 90% of surveyed companies** reported regularly using personal AI tools for work.

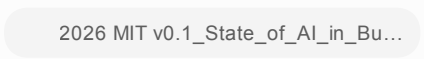

This is much more significant than simply demonstrating widespread ChatGPT use.

It suggests employees are effectively conducting their own AI adoption experiments.

They are independently determining:

- which work should be automated;

- which tasks benefit from AI;
- where AI saves time;
- where AI is unreliable;
- where workflow integration is necessary.

In other words:

The workforce may know where the AI ROI opportunities are before senior leadership does.

That challenges the conventional centralized-AI-transformation model.

---

5. This creates a fascinating organizational contradiction

The report essentially finds:

Employees resist enterprise AI tools while enthusiastically adopting AI.

Those two things sound contradictory until you examine *what* workers are rejecting.

The leading barrier on page 11 is described as:

“Unwillingness to adopt new tools.”

But those same workers frequently use ChatGPT and other tools privately.

Therefore, the evidence does **not necessarily demonstrate generalized employee resistance to AI.**

It could instead indicate:

> **Employees resist badly designed AI implementations.**

That is an extremely important distinction.

Workers may not be saying:

> “I don't want AI.”

They may be saying:

> “I don't want another corporate system that makes my job harder.”

That changes the diagnosis completely.

6. The report inadvertently makes a strong human-readiness argument

The authors ultimately frame the principal issue as a **technology learning gap**:

> systems cannot remember, adapt or learn sufficiently.

That is certainly part of the problem.

But the report's own evidence suggests another gap exists simultaneously:

The Human–Workflow Gap

Look at the reported implementation barriers:

- unwillingness to adopt tools;
- output-quality concerns;
- poor user experience;
- insufficient executive sponsorship;
- difficult change management;
- poor integration with workflows.

Those aren't exclusively model problems.

They are an interaction between:

Technology × people × workflow × leadership × organizational design.

That distinction is important because simply adding persistent memory to an AI agent will not solve poor leadership, employee distrust, badly designed workflows or weak incentives.

The report occasionally comes close to implying:

> Memory + agents = enterprise ROI.

Its own evidence suggests reality is considerably more complicated.

2026 MIT v0.1_State_of_AI_in_Bu…

---

7. The 90/10 result deserves considerably more attention

Page 13 asks users whether they would assign work to AI or a junior colleague.

For quick tasks:

70% AI / 30% human

For complex, multi-week work:

10% AI / 90% human

2026 MIT v0.1_State_of_AI_in_Bu…

That is an extraordinary result.

It suggests AI adoption has already separated into two distinct economic categories.

Transactional work

AI is becoming highly competitive.

Examples:

- email
- summaries
- simple analysis
- document preparation
- routine research

Contextual work

Humans retain enormous advantage.

Examples:

- client management
- ambiguous projects
- long-running assignments
- organizational judgment
- exception management

That means the near-term workplace may evolve less toward:

Human → AI replacement

and more toward:

Human → AI operator → AI supervisor → exception manager.

---

8. The report makes a convincing argument that ROI is hiding in the back office

One of its strongest financial observations appears on pages 20–21.

Organizations disproportionately invest in visible applications such as sales and marketing.

Yet some of the largest measurable returns appeared in:

- operations
- procurement
- finance
- document processing
- customer service
- outsourced services

Reported examples include:

- **$2–$10 million annual BPO savings**
- **30% reduction in external agency spending**
- **$1 million annual outsourced-risk-management savings**
- **40% faster lead qualification**
- **10% improvement in customer retention**

2026 MIT v0.1_State_of_AI_in_Bu…

This exposes another possible reason companies report weak AI ROI:

They may be deploying AI where executives can see it rather than where economics are strongest.

Sales AI gives executives an easy story.

An accounts-payable process improvement may generate more profit but receive less attention because the value is operational rather than glamorous.

---

9. One of the best observations: AI ROI doesn't necessarily mean layoffs

The report finds that successful deployments generally were **not producing broad-based workforce reductions**.

Instead, many companies captured value through reductions in:

- BPO spending
- agencies
- outside consultants
- outsourced risk management
- external processing

2026 MIT v0.1_State_of_AI_in_Bu…

This is strategically significant.

The first wave of AI labor substitution may therefore occur outside the organization before it occurs inside it.

Think of the sequence as:

Consultants → agencies → contractors → BPO → open positions → internal jobs

rather than immediately:

employees → layoffs.

That is a materially different workforce scenario.

---

10. “Buy vs. Build” is another important finding—but needs caution

The report says externally partnered implementations reached deployment approximately:

67% of the time

versus:

33% for internally built solutions.

2026 MIT v0.1_State_of_AI_in_Bu…

That produces the memorable conclusion:

> External partnerships are roughly twice as successful.

But the researchers correctly insert an important caveat: **correlation does not prove causation.**

Companies buying externally could differ systematically from companies building internally.

They might have:

- clearer use cases;
- stronger procurement;
- better executive sponsorship;
- more mature processes;
- higher willingness to change;
- more disciplined ROI requirements.

So the finding should not become:

“MIT proved buying AI works twice as well.”

A more defensible interpretation is:

Organizations in this sample using strategic vendor partnerships reported materially higher deployment rates than organizations relying on internal development.

That wording will stand up much better under scrutiny.

---

11. The report's proposed success model is essentially this

Across the report, successful implementations exhibit six characteristics:

1. Start narrow.

Don't automate an entire company.

Automate one expensive workflow.

2. Solve a real operational problem.

Not a hypothetical “AI use case.”

3. Integrate into existing workflows.

Employees should not have to build an entirely new work routine around the technology.

4. Allow feedback and adaptation.

The system should improve as users interact with it.

5. Let frontline managers identify use cases.

Not every AI opportunity should originate in an innovation lab.

6. Measure economics.

Successful buyers evaluate business outcomes rather than AI benchmark performance.

That is probably the most useful executive playbook in the report.

---

12. There are several significant inconsistencies in the report

These are important if you intend to quote it publicly.

The “95%” claim is rhetorically stronger than the underlying evidence.

The executive summary states:

> “95% of organizations are getting zero return.”

But elsewhere the report describes **5% of task-specific enterprise AI implementations reaching successful implementation**.

Those are not precisely the same thing.

“No successful workflow implementation” does not automatically equal **zero economic return across an organization**.

Industry counts are inconsistent.

The executive summary refers to:

> **2 of 8 major sectors**

Elsewhere the report states:

> **seven of nine sectors**

Yet the industry table appears to contain eight categories.

That should make anyone using the statistic publicly cautious.

AI investment numbers are inconsistent.

Page 9's takeaway says approximately:

> **50% of GenAI budgets go to sales and marketing.**

The accompanying paragraph says executives allocated:

> **approximately 70%.**

The report also admits this wasn't actual spending data—it asked executives how they would distribute a hypothetical $100.

So this should not be described as:

> “70% of corporate AI spending goes to sales and marketing.”

That would overstate what was measured.

2026 MIT v0.1_State_of_AI_in_Bu…

---

13. The biggest methodological weakness is the definition of success

Page 7 defines a successfully implemented task-specific GenAI tool as one where:

> users or executives remarked that it caused marked and sustained productivity and/or P&L impact.

That is not the same as independently measured ROI.

Later the appendix says measurable KPIs were considered, but the study still relies heavily on interview and survey evidence.

Therefore, the report mixes:

perceived value

with

measured operational outcomes

and sometimes with

P&L ROI.

Those three concepts should be separated much more carefully.

---

14. The Agentic AI conclusion is plausible—but more speculative than the earlier findings

The report argues that persistent memory, agentic systems, MCP, A2A and NANDA could solve much of the learning gap.

That is a reasonable technological thesis.

But the evidence presented earlier in the study does not conclusively prove:

lack of AI memory → 95% failure

or

agentic AI → successful enterprise ROI.

The report establishes correlation between workflow fit, adaptability and successful implementation much better than it establishes that agentic architecture itself causes ROI.

The **"Agentic Web"** section should therefore be read largely as a forward-looking thesis rather than as a finding demonstrated by the research.

---

15. The deeper economic interpretation

I would reduce the entire report to one equation:

AI ROI = Capability × Workflow Fit × Adoption × Learning × Economic Value

If any factor approaches zero, the return collapses.

A brilliant model with poor workflow integration:

≈ zero ROI

A perfectly integrated platform nobody trusts:

≈ zero ROI

High employee adoption of trivial use cases:

≈ little P&L impact

A strong technology deployed against a low-value process:

≈ low ROI

That is why counting licenses or pilots gives executives such a misleading picture of AI transformation.

---

16. The real “gap” may actually be larger than the report describes

The report names the problem the **GenAI Divide**.

But its evidence suggests at least five separate divides:

Technology Divide
Can the system actually perform the task?

Workflow Divide
Can it operate inside real business processes?

Human Divide
Will employees trust and use it?

Leadership Divide
Do managers redesign work around the capability?

Economic Divide
Can anyone demonstrate that adoption created measurable value?

The industry has concentrated enormous energy on the first divide.

The report suggests the remaining four may increasingly determine ROI.

---

17. What leadership should do Monday morning

Instead of asking:

> “Where can we use AI?”

Leadership should ask:

> **“Where is expensive work currently occurring, who performs it, what prevents that workflow from being more productive, and how will we know whether AI changed the economics?”**

Then follow a disciplined sequence:

Measure current cost, productivity and readiness.

Segment technology problems from workflow, skill, trust and organizational problems.

Select narrow, economically meaningful workflows.

Partner or build based on measurable business requirements.

Deploy alongside the people who actually perform the work.

Measure again against the original baseline.

Only then can executives distinguish **AI utilization from AI adoption—and AI adoption from AI ROI.**

Overall assessment

I would rate the report:

Dimension	Assessment
Executive relevance	**Excellent**
Framing of the problem	**Excellent**
Workflow/adoption insight	**Very strong**
ROI insight	**Strong but directional**
Statistical rigor	**Moderate**
Causal evidence	**Limited**
Workforce analysis	**Strong and underdeveloped**
Agentic-AI forecast	**Interesting but speculative**

The strongest part of this report is not the sensational 95%.

It is the evidence showing **why billions of dollars of experimentation can coexist with very little economic transformation.**

The report reveals an emerging reality:

> **The next stage of enterprise AI competition will not be won simply by the company with the most sophisticated model. It will be won by organizations that can turn AI capability into repeatable human workflows, sustained adoption and measurable economic value.**

That is a much stronger—and more defensible—conclusion than simply saying **“95% of AI projects fail.”** 2026 MIT v0.1_State_of_AI_in_Bu…

Sources