

Key Findings Summary

- 1. Technology is not the hardest part.** 77% of the hardest challenges were invisible and intangible costs: change management, data quality, and process redesign. 61% of successful projects included at least one prior failure, whose costs never appear in the final ROI.
- 2. Timeline variance is organizational, not technical.** Similar use cases took weeks at one company and years at another. The difference was executive sponsorship, existing organizational processes, and end user willingness.
- 3. Escalation-based models were associated with better results.** Escalation-based models (AI handles 80%+ autonomously, humans review exceptions) delivered 71% median productivity gains versus 30% for approval models. This may, in part, reflect different types of tasks addressed.
- 4. Executive sponsorship is about actions, not approval.** Effective sponsors clear blockers weekly, bridge business and technical teams, and tie AI adoption to corporate OKRs. Most critically, they create a culture that gives permission to fail.
- 5. Staff functions are the most frequent source of resistance, but some parts may become enablers after buy-in.** Legal, HR, Risk, and Compliance were the most frequent source of resistance at 35%, ahead of internal end-users at 23%.
- 6. Headcount reduction is common but not inevitable.** Headcount reduction was the largest outcome in 45% of the deployments, but alternatives (hiring avoided, redeployment, no reduction) accounted for 55%. Broader labor market data suggests entry level roles in AI exposed occupations are already declining.
- 7. Revenue from AI is real, but still rare, and follows three patterns.** Personalization that converts, speed that wins deals, and internal tools repackaged as products. A small subset of cases also shows AI enabling work that was previously impossible.
- 8. Agentic AI works, but most firms have not used it, yet.** Agentic implementations showed 71% median productivity gains versus 40% for high-automation but represented only 20% of cases. Agentic AI isn't a new UI; it's a redefinition of the role of humans and machines in the workflow.
- 9. Messy data is not a blocker if you design around it.** LLMs fixed many of the data problems they were supposed to struggle with. Store everything, connect it, and let the models do the cleaning.
- 10. Security enables more than it blocks.** Security was not a project killer in any of the cases we studied. Requirements that were initially barriers later enabled projects to handle sensitive data.
- 11. Model choice is a commodity for many use cases.** For 42% of implementations, model choice was fully interchangeable. Companies don't always need the best available AI models. The durable advantage is in the orchestration layer, not the foundation model.