

### Agentic AI: Navigating ROI Challenges and Building a Blueprint for Enterprise Success

Randy McGraw<sup>1</sup>

#### Highlights

- **Agentic AI is a step-change beyond automation—and it fails when “agentic” is treated as branding.** True agents perceive/reason/act/learn across systems (memory, tool-use, orchestration).
- **Enterprises routinely derail the value of agentic AI through *agent-washing*:** underestimating integration/architecture choices, miscalculating ROI, and ignoring governance maturity, regulatory constraints, and vendor/skill requirements.
- **ROI models misfire when they over-index on cost takeout and under-measure value creation.** A cost-only lens misses decision speed/quality, resilience, customer outcomes, risk reduction, and new revenue.
- **Sustainable ROI is fundamentally an execution-and-adoption problem, not a model-performance problem.** Treat ROI as a risk-adjusted measure the net present value of the project depends critically on cross-functional cooperation, capability, and controls.

©2025 Andersen Institute for Finance & Economics. All Rights Reserved. This material is confidential intellectual property of the Andersen Institute for Finance & Economics. The views expressed in this note are those of the authors and do not represent an official position of The Andersen Institute for Finance and Economics or affiliated organizations. By viewing this Andersen Business Insight, you agree that you will not directly or indirectly copy, modify, record, publish, or redistribute this material and the information therein, in whole or in part. No warranty or representation, express or implied, is made by the Andersen Institute or any of its affiliates, nor does Andersen accept any liability with respect to the information and data set forth herein. Distribution hereof does not constitute legal, tax, accounting, investment or other professional advice. The information provided herein is not intended to provide a sufficient basis on which to make an investment decision. Recipients should consult their own advisors, including tax advisors, before making any investment.

<sup>1</sup> Randy McGraw is a Managing Director at Andersen Consulting.

### Introduction

Autonomous systems capable of perceiving, reasoning, and acting toward goals with minimal human intervention—Agentic AI—is rapidly emerging as a transformative force in enterprise technology strategy. Unlike traditional AI, which typically responds to prompts or executes predefined tasks, AI agents possess memory, adapt to feedback, and orchestrate complex workflows across systems. These agents can plan, decide, and execute actions independently, often coordinating with other agents or external tools to achieve business objectives. As enterprise AI accelerates in a digital transformation wrapper, Agentic AI offers unprecedented opportunities to scale operations, enhance decision-making, and unlock new efficiencies.

Yet despite its promise, many enterprises face a sobering reality: the return on investment (ROI) from Agentic AI remains elusive. For example, Gartner predicts that over 40 percent of Agentic AI projects will be canceled by 2027 due to inflated expectations, technical complexity, and unclear business value. But this need not be the case. Real value remains, waiting to be unlocked via Agentic AI.

### Defining Agentic AI: Beyond Automation

Agentic AI encompasses autonomous systems endowed with the ability to perceive, reason, act, and learn. By ingesting diverse data inputs - ranging from structured databases and unstructured text to real-time event streams- these agents build an internal model of their environment. They apply advanced reasoning algorithms to identify patterns, evaluate alternatives, and make decisions aligned with defined objectives. Once a course of action is selected, the agent executes tasks across digital ecosystems- triggering workflows, invoking Application Programming Interfaces (APIs), or coordinating with other sub-agents- to accomplish complex goals with minimal human intervention. Crucially, Agentic AI continuously monitors its own performance, measuring outcomes against targets and incorporating feedback to refine its models and strategies over time.

To place this in context, Agentic AI applications have already delivered concrete impacts across a wide array of industries including:

1. **Retail:** Autonomous agents drive real-time inventory replenishment, dynamic pricing adjustments, and personalized in-store experiences.
2. **Financial Institutions:** Agentic systems are deployed for continuous fraud detection, automated portfolio rebalancing, and regulatory compliance monitoring.
3. **Manufacturing:** Agents orchestrate predictive maintenance schedules, optimize production sequencing, and conduct computer-vision quality inspections with minimal human oversight.
4. **Oil & Gas and Broader Energy Sectors:** Agentic AI autonomously analyzes drilling sensor data, forecasts supply-demand imbalances for grid management, and schedules equipment upkeep to prevent costly downtime.

Developing and deploying AI agents typically entails several layers of technical and commercial considerations. Choosing the right agentic project and then its architecture requires careful consideration of an organization's unique context. To be successful, enterprise teams must:

1. Define exactly what they seek to accomplish – which specific processes are involved, and the key performance indicators (KPIs) or business levers that need to move and how so.
2. Evaluate their risk tolerance and governance maturity to determine how much autonomy agents can safely wield.
3. Incorporate regulatory requirements and industry-specific constraints that shape permissible data usage, decision boundaries, and auditability standards.
4. Review integration complexity with legacy systems, anticipated scale, and performance service-level agreements that inform whether a centralized "conductor" model or a decentralized mesh of collaborating agents is most appropriate.
5. Review vendor qualifications and experience, as well as which internal skills will be needed to co-deploy and maintain these intelligent systems.
6. Agree on and calculate financial considerations—namely, the total cost of ownership versus projected benefits, the "Return on Investment (ROI) calculation."

### The ROI Challenge: How to Get it Right

In this review, we focus on Item 6 in the list above—the ROI calculation. One hurdle slowing the overall adoption of Agentic AI, and perhaps over-indexing on false negatives in project review (i.e., redlights that could be greenlights) can be described as internal struggles to conceptualize and accurately quantify ROI.

Traditionally, enterprises assess ROI on AI projects almost exclusively through a cost-cutting lens. Finance and operations teams set targets for headcount reduction, cycle-time compression, and unit-cost savings. They establish baseline metrics—full time equivalent (FTE) hours spent on manual tasks, error rates in processing, and per-transaction labor costs—and then overlay projected efficiency gains from automation. Success is judged by monthly dashboards tracking hours saved, error avoidance, and compute-versus-labor cost ratios. Pilots are spun up in controlled environments and only scaled if they meet or exceed predefined cost-savings thresholds.

### Common Pitfalls in Traditional ROI Models

While this methodology is not inaccurate and serves as a necessary baseline, it is far from complete. The cost-centric focus alone is too narrow and fails to acknowledge critical value elements of the equation:

1. It overlooks strategic or intangible benefits such as faster decision-making, improved customer satisfaction, and new revenue streams.
2. It discourages investment in critical foundations (clean data pipelines, robust governance) that don't deliver immediate labor savings.

3. It incentivizes "checkbox" projects that automate low-value tasks rather than tackle complex, high-impact workflows.
4. It fails to account for downstream effects-like reduced compliance risk or enhanced brand reputation-that accrue over time.
5. It fails to quantify actual improvements in output and overfocuses on input-cost reduction. Reducing input costs is simple to measure with standard management accounting methods, whereas measuring output improvements carries uncertainty and requires upper- and lower-bound analysis due to AI's evolving nature and improving performance with larger datasets.

Moreover, to understand why conventional ROI assessments often misfire for Agentic AI initiatives in particular, consider these five recurring pitfalls identified by Andersen Consulting in its reviews of Agentic AI projects:

1. **Superficial "Agent Washing."**  
Vendors and technology providers often mislabel simple automation or chatbot solutions as "Agentic AI," leading to inflated expectations, misaligned costs that show up later, and wasted spend.
2. **Data Readiness and Infrastructure Gaps.**  
Legacy systems and brittle pipelines starve AI agents of reliable, real-time data, crippling their autonomy and adaptability.
3. **Short-Term ROI Metrics.**  
Narrow financial KPIs mask broader value levers-agility, resilience, and decision quality-that Agentic AI can unlock.
4. **Misuse of Hyperscaler or Cloud Credits.**  
Relying on promotional cloud or hyperscaler credits during pilots or in the case modelling can artificially lower compute costs and distort ROI calculations. While these credits improve cash flow in the short term, they don't represent sustainable cost structures and should be excluded from core ROI assessments.
5. **Negative Impact of Internal "Fear-Uncertainty-Doubt" (FUD)**  
FUD acts as an endogenous "adoption risk" that depresses the cross-functional buy-in and operational activity needed to realize benefits, so the ROI is not merely uncertain but systematically biased downward because the organization withholds the complementary effort that converts technical capability into measured value.

Let's wrap all of this into a real example, the case of a well-known multinational logistics firm where Agentic AI was deployed to autonomously analyze weather patterns, port congestion data, and fuel prices to reroute shipments in real time. Previously, such decisions at the enterprise relied on manual coordination and static rules, and several dedicated staff, and often resulted in delays, suboptimal routes, and cost increases. Despite going through the sequence above as described in Section 1, the initial case for deploying an Agentic AI fix was not sufficiently motivational. However, by including the following potential benefits in its ROI calculations, the enterprise expanded its framework from simple cost reduction to account for the significant value delivered via **decision quality metrics**:

1. Reduction in error rates or missed deadlines.
2. Impacts of improved forecast accuracy,
3. Financial impact of better decisions (e.g., avoided penalties, optimized resource allocation).

Alongside traditional metrics, these potential outcomes were identified and quantified during a proof of concept (POC) exercise, allowing leadership to see how Agentic AI enhanced strategic responsiveness and operational reliability-benefits that would have been invisible in a purely cost-efficiency model.

Within the agentic framework, the new AI agent system:

1. Dynamically adjusted routes in actionable timeframes,
2. Reduced delivery times by 18 percent,
3. Avoided \$3.2 million in annual penalties from missed service-level agreements, and
4. Enabled the company to create new capabilities with the time savings, such as more high-touch customer-centric activity from the same team.

A more robust ROI model was the difference between a false negative and a successful deployment of an Agentic AI application.

### Global AI Investment Trends and Strategic Imperatives

Agentic AI and its supporting cloud infrastructure are no longer confined to regional growth stories—they are shaping economic strategies on every continent. From the hyperscaler expansions in Asia and North America to state-backed initiatives in Europe and Africa, governments and enterprises alike are recalibrating their digital transformation roadmaps around the promise (and risks) of agentic AI:

- **North America** remains the epicenter of AI R&D, with hyper-scalers like AWS, Google, and Microsoft doubling down on multi-agent frameworks, observability tooling, and governance infrastructure. The U.S. Inflation Reduction Act and CHIPS Act also indirectly fuel AI competitiveness through semiconductor investments.
- **Europe** is balancing sovereign AI ambitions with stricter regulatory oversight. The EU's AI Act and General Data Protection Regulation (GDPR)-driven frameworks have encouraged local investments in trusted AI and cloud infrastructure, particularly in France, Germany, and the Nordics.
- **Asia-Pacific** is seeing both scale and speed, with China and India investing billions in sovereign models, whilst Southeast Asia, Japan, and South Korea focus on AI-ready data centers and 5G integration.
- **Latin America** is emerging as a fast follower, with Brazil and Mexico investing in AI cloud zones and digital skills programs, though political and economic volatility complicates timelines.
- **EMEA** is betting on leapfrogging. The UAE and Saudi Arabia are investing in sovereign AI stacks and hyperscaler partnerships, while African nations like Kenya, Nigeria, and South Africa are positioning AI to solve infrastructural and developmental challenges. Said **Omar Sultan AI**

**Olama, UAE Minister of State for AI:** *"We are not consumers of AI; we aim to be creators of sovereign AI systems for our future economy."*

We can summarize the projected value of AI investment by region as follows:

## AI Market Outlook by Region

Continent	Hyperscaler Activity	Key National Initiatives	Market Outlook: Infrastructure Market (2030, \$Billions)
North America	Mature, densest region for AI R&D and infrastructure	CHIPS Act, NIST AI Safety Institute	\$150
Europe	Expansion with sovereign cloud zones	EU AI Act, Gaia X\$75B AI infrastructure market	\$75
Asia Pacific	Rapid hyperscaler deployment, sovereign AI in China, India	Digital India, China's Next Gen AI Plan, Japan's AI Strategy	\$200+
Latin America	Moderate growth; Brazil&Mexico driving infrastructure	Brazil's AI Strategy, Mexico AI policy drafts	\$25
Middle East	Sovereign clouse plus AI city projects (e.g., NEOM)	UAE National AI Strategy 2031, Saudi Vision 2030	\$40
Africa	Early stage, infrastructure plus AI for development	Smart Africa Alliance, Kenya's Digital Master Plan	\$10

Created with Datawrapper

### Common Pitfalls in Global Agentic AI Adoption

Despite the optimism and investment acceleration, recurring themes of caution are emerging across regions that require attention as an Agentic AI strategy develops for an enterprise:

1. **Fragmented Data Ecosystems**

Data silos, incompatible standards, and uneven interoperability across borders hinder the ability of Agentic AI to access holistic, real-time insights. Multinationals struggle to unify customer, supply chain, and regulatory data across jurisdictions, slowing automation and cross-functional decision-making.
2. **Regulatory Complexity**

The global AI regulatory landscape is fragmented, with Europe leading strict oversight while emerging markets often tend to adopt lighter-touch or less clear policies. This mismatch increases compliance costs for enterprises and can result in risk-averse deployments, stifling innovation in highly regulated sectors like finance and healthcare.
3. **Environmental & Infrastructure Constraints**

Hyperscaler data centers require immense amounts of energy and water, straining local utilities and raising sustainability concerns. Climate volatility—from droughts to extreme weather—further threatens continuity of operations, underscoring the need for greener infrastructure and adaptive resilience planning. Any strategic plan for Agentic AI deployment needs to carefully look at power cost and accessibility carefully.
4. **Talent and Capability Gaps**

While demand for AI talent surges, there is a pronounced shortage of expertise in areas like agent orchestration, observability, and governance. This scarcity affects both advanced economies and developing nations, slowing deployment timelines and increasing reliance on a small pool of specialists.
5. **Cultural & Organizational Resistance**

AI adoption is as much about mindset as it is about technology. Hierarchical cultures and entrenched legacy systems often resist decentralization, while low levels of digital trust in some regions amplify skepticism, delaying enterprise-wide transformation.
6. **Misinformation & Disinformation**

The global spread of deepfakes, synthetic media, and election-related disinformation destabilizes trust in institutions and media. Agentic AI, when deployed in sensitive sectors like governance or healthcare, must navigate this volatile information landscape without inadvertently amplifying false narratives.
7. **AI-Powered Scams & Fraud**

Fraudulent activity leveraging AI—such as voice cloning, impersonation, and investment scams—has surged across financial services and e-commerce globally. This not only undermines consumer confidence but also creates significant regulatory and reputational risks for organizations deploying agentic systems.
8. **Algorithmic Opacity & Accountability Gaps**

The "black box" nature of many agentic systems leaves critical decisions unexplained, eroding trust among regulators, enterprises, and citizens. Without robust observability and auditing frameworks, accountability lapses can become systemic risks, particularly in safety-critical and regulated industries.
9. **Uncertain Human Impact**

While the Government of Singapore took a global first move in January 2026 by codifying a Model AI Governance Framework for Agentic AI (MGF for Agentic AI, Version 1.0), setting

out a legal framework for the responsible deployment of autonomous agents via (i) assessing and bounding risks upfront (e.g., limit autonomy/tool/data access), (ii) making humans meaningfully accountable (clear responsibility + effective oversight), implementing technical controls and processes (testing, guardrails, monitoring), and (iii) enabling end-user responsibility (information/training so users can oversee agents appropriately), much of the FUD mentioned earlier remains unaddressed.

### Blueprint for ROI Success with Agentic AI

At Andersen Consulting, we work with clients to establish a rigorously validated, institutionally aligned perspective on every Data & AI initiative before it advances. Drawing on our experience navigating Agentic AI's unique pitfalls and the intertwined technical and organizational complexities, as noted throughout, we recommend the following heuristic to guide strategic decision-making and ensure accurate (sustainable) ROI calculations:

1. Clarify Use Cases and Business Objectives.  
Focus on high-impact workflows (e.g., onboarding, procurement, or customer engagement) where agents can deliver measurable outcomes.
2. Redesign ROI Frameworks.  
Expand beyond cost savings to include decision speed and quality, improved KPIs and business levers, operational agility, risk mitigation, and employee and customer experience.
3. Invest in Data and Governance.  
Build an agile, clean data infrastructure at all times and start now. Implement metadata systems, observability tools, and governance agents to ensure transparency and trust.
4. Start with Human-in-the-Loop Systems.  
Deploy agents in controlled environments with oversight. Gradually scale autonomy as confidence and maturity grow and as the impact of Agentic AI in the organization can be fully absorbed.
5. Leverage Global and Local Insights.  
Ask Andersen Consulting to provide both broad, global best practices and nuanced, local market perspectives, enhancing your learning from case studies and de-risking strategic decisions.
6. Eliminate FUD in the Execution Path.  
Treat FUD as an endogenous "adoption risk" that depresses cross-functional buy-in and execution, and mitigate it with credible leadership commitments (redeployment/reskilling), co-design with impacted teams, and success metrics framed as capacity/quality gains—thereby raising the probability of implementation and stabilizing ROI.

### A Necessary Word on Fear-Uncertainty-Doubt in the Enterprise

Eliminating FUD within internal teams may be the most difficult of all Agentic AI critical success factors as it has dependencies on both organizational culture and leadership, and it emerges in unexpected ways at unexpected times.

Agentic AI ROI is best treated as a risk-adjusted investment, not a pure model-performance forecast. Conceptually, this might be captured in computing a “risk adjusted NPV” that incorporates factors beyond direct cash flow: In this framework, the risk-adjusted NPV is presented by an equation:

$$\text{Risk-adjusted NPV} = \pi (\text{cooperation, capability, controls}) \times \text{NPV}_{\text{technical}} - \text{Costs}.$$

In practice, the main forecasting error in the consideration of new agentic AI projects is not mis-estimating  $\text{NPV}_{\text{technical}}$ , but ignoring that  $\pi$ —the probability the organization actually implements, adopts, and scales—can fall sharply when teams anticipate displacement, loss of autonomy, or reputational exposure. Under FUD, ROI becomes endogenously “unmeasurable” in the short run because the complementary actions required for value capture (process redesign, data hygiene, exception-handling knowledge, disciplined usage) can be selectively withheld, and the resulting underperformance is then misread as a technology failure rather than an adoption equilibrium.

Eliminating FUD therefore requires making  $\pi$  a first-class design variable with incentive-compatible commitments and metrics. We advise clients to (i) separate “capacity released” from “jobs removed” in early-stage evaluation, using cycle time, quality, backlog reduction, and risk reduction as the primary proof points; (ii) embed credible redeployment pathways and reskilling into the program charter (and budget), so “augmentation” is not a slogan but a contract; and (iii) co-design agent workflows with the roles most affected, converting threatened groups into co-producers of the operational playbook (guardrails, escalation logic, exception taxonomies) that determines whether agents succeed at scale. This reframes agentic AI from a headcount story to a capability story—raising adoption, improving performance, and making ROI calculations stable rather than politically contingent.

The same logic applies to the IT organization, which is inevitably implicated because Agentic systems are integration-heavy (identity, access, data permissions, logging/audit, network controls): IT “blocking” is often rational when downside risk is concentrated on them while upside is diffuse elsewhere, and when automation threatens routine IT tasks. To raise  $\pi$ , we establish an “approved lane” for agentic deployments—vetted models and connectors, auditable workflows, continuous monitoring—paired with joint business/IT ownership (shared objectives and key results (OKRs) for value and risk posture) and explicit funding for platform capacity (observability, evaluation, red-teaming, and control planes). In this governance architecture, cross-functional buy-in becomes measurable and engineered, enabling sustainable ROI rather than one-off pilot results.

In a real-time example of the efficacy of this approach, take the recent agenticification of the IT Service Help Desk function of a major Japanese industrial enterprise. Seeking alleviation from deteriorating service level KPIs and commensurate cost increases as the benchmarked 1:100 ratio of staff to internal customers worsened with growth and service complexity, the company understood the potential impact of an Agentic AI offload of repeating functions to unmanned channels. But in a savvy move, rather than develop the project in a purely cost savings mold that would invariably entail staff cuts, they made the affected IT team the single-threaded project owner, required cross-functional support of the project, and stated up front that the value of KPI improvements and “found time” would be largely distributed to the

team to expand capabilities around more hands-on high-value work and explore new service areas that would increase their value to the enterprise. The project reached greenlight in just 3 months despite expected scope creep in a new area; what emerged from the project became the classic “glass half empty vs half full” scenario inherent in all Agentic AI deployments- harvest the value in terms of pure cost cuts and staff reduction, or use creativity aligned to strategic planning and deploy the value in whole or in part into new capabilities.

Overall, it remains our belief that Agentic AI should be treated as a transformative capability to be measured, governed, and aligned with enterprise goals.

### **Conclusion: Strategic Investment with Eyes Wide Open**

Agentic AI is not just another tech trend-it's a structural evolution in how enterprises operate, and it is the next significant and transformative step in AI for enterprise. But success hinges on strategic clarity, robust infrastructure, and, importantly, new and nuanced ROI models.

The global investment trajectory is promising, and while the overall AI business is accelerating in 2026, it must be matched by strategy, governance, and smart vendor selection. Andersen Consulting professionals have pioneered more than 400 enterprise AI and Data deployments over more than a decade and seek to help clients engage with agentic AI deliberately.

The path forward does not have to be as Gartner predicts. Agentic AI should be treated as a transformative capability to be measured, governed, and aligned with enterprise goals. With the right blueprint, including proper ROI calculations and a comprehensive approach to project initiation and evaluation, Agentic AI can deliver not just automation but an enduring strategic advantage.