www.ijarp.com ISSN 2456-9992 Page: 01-14

International Journal Advanced Research Publications

Review Article

Volume: 02
Issue: 03

AI & GEN-AI IN BANKING: FROM EFFICIENCY PLAY TO STRATEGIC TRANSFORMATION

***Venugopal Madarapu**

Assistant Professor, Vaagdevi Degree & PG College, Hanamkonda, Telangana.

Article Received: 27 January 2026, Article Revised: 15 February 2026, Published on: 07 March 2026

***Corresponding Author: Venugopal Madarapu**

Assistant Professor, Vaagdevi Degree & PG College, Hanamkonda, Telangana.

DOI: https://doi-doi.org/101555/ijarp.5888

ABSTRACT

The integration of Artificial Intelligence (AI) and Generative AI (Gen-AI) in banking has rapidly evolved from experimental pilots to a strategic imperative. This article examines the transformative journey of AI in financial services, tracing its progression from back-office efficiency tools to a fundamental driver of competitive strategy and business model reinvention. Drawing on industry data, institutional case studies, and recent technological developments through early 2026, we analyze how banks are moving beyond isolated use cases toward enterprise-wide AI integration. The research identifies six high-value application domains, examines the emergence of agentic AI and multi-agent systems, and evaluates the governance frameworks necessary for responsible deployment. We argue that the banking industry stands at an inflection point: the strategic question is no longer whether to adopt AI, but how to architect organizational capabilities, governance structures, and partnership models that enable sustainable competitive advantage in an AI-first financial landscape. The article concludes with a strategic framework for institutions navigating this transition and offers projections for the next phase of industry transformation.

KEYWORDS: Generative AI, Banking Strategy, Digital Transformation, Agentic AI, AI Governance, Financial Services Innovation

INTRODUCTION

In May 2025, a panel of financial leaders at The Asian Banker Summit declared an end to the era of AI buzzwords. "Everything that you have learned in the last 10 years," one speaker noted, "the rules are being rewritten dramatically. AI collapses everything we do at the back end and hands the front end over to the customer". This sentiment captures the profound shift

underway in global banking—a transition from viewing artificial intelligence as a tactical efficiency tool to recognizing it as the central engine of strategic transformation. The numbers tell a compelling story. According to the 2025 EY-Parthenon Generative AI in Banking survey, 77% of banks have now actively launched or soft-launched GenAI applications, compared to just 61% in 2023. More significantly, 61% of respondents already report substantial impacts from their deployments, and 89% expect major transformative benefits within the next two years. This represents a decisive move from proof-of-concept to production, from experimentation to execution. Yet the path to value creation remains uneven. McKinsey estimates that while generative AI could unlock $200 billion to $340 billion in annual value for the banking sector, and potentially improve profitability by 14-24%, realizing these gains requires more than technological deployment. It demands fundamental rethinking of operating models, governance structures, and strategic positioning. As Gartner's 2025 Hype Cycle for AI in Banking observes, generative AI has entered the "Trough of Disillusionment," shifting attention to foundational enablers: AI-ready data, engineering practices, and robust governance frameworks.

This article examines the strategic implications of AI and Gen-AI adoption in banking. It addresses three central questions: First, how is the application of AI in banking evolving from isolated efficiency plays toward integrated strategic transformation? Second, what are the high-value domains where AI is reshaping competitive dynamics? Third, what organizational and governance capabilities determine success in this new landscape? The argument proceeds in five parts. Section 2 traces the evolution of AI in banking from rule-based systems to generative and agentic architectures. Section 3 analyzes key application domains and their strategic implications. Section 4 examines implementation challenges and governance imperatives. Section 5 presents a strategic framework for institutions navigating this transition. Section 6 concludes with implications for future research and industry practice.

Literature review

This research article was developed drawing on industry reports, academic publications, and conference proceedings available as of early 2026. The article synthesizes findings from multiple sources to present a comprehensive analysis of AI's strategic role in contemporary banking.

The Evolution of AI in Banking (From Automation to Augmentation to Agency)

Understanding the current inflection point requires appreciating the historical trajectory of AI in financial services. The evolution can be understood as three distinct eras, each building upon but not replacing the last.

1) The First Era: Rule-Based Automation (1980s-2010s)

The first wave of AI in banking was characterized by rule-based systems and early machine learning applications. Expert systems codified human knowledge into decision trees for credit scoring and fraud detection. These systems offered consistency and scale but were inherently limited by their dependence on explicitly programmed rules. They excelled at structured tasks with clear parameters but could not adapt to novel situations or process unstructured information.

2) The Second Era: Predictive Intelligence (2010s-2023)

The proliferation of data and advances in machine learning enabled a second wave focused on prediction and personalization. Banks deployed models that could identify patterns invisible to human analysts—detecting anomalous transactions, predicting customer churn, and optimizing marketing campaigns. This era saw the rise of recommendation engines, dynamic pricing, and risk models that continuously learned from new data. However, these systems remained primarily analytical rather than generative. They could identify what was likely to happen but could not create new content, synthesize disparate information, or engage in natural dialogue. They were tools for human decision-makers rather than autonomous actors in their own right.

3) The Third Era: Generative and Agentic AI (2023-Present)

The introduction of large language models (LLMs) and generative AI marked a qualitative shift. For the first time, banks could deploy systems that understood and generated human language, synthesized unstructured information, and engaged in context-aware dialogue. As Maja Pantic, Chief AI Research Officer at NatWest Group and former generative AI research director at Meta, observes, "The crucial year was not 2022. It was 2017, when the transformer architecture was released". This architectural breakthrough enabled models to understand context and generate coherent, relevant content across domains. The current frontier is "agentic AI"—systems capable of autonomous action across multi-step workflows. Unlike chatbots that respond to queries, agentic AI can independently execute complex processes: onboarding a customer, conducting know-your-customer (KYC) verification, assessing credit

risk, and making underwriting decisions within defined parameters. As one industry executive explains, "Agentic AI is like your own internal team. It does not just respond—it completes multi-step processes".

McKinsey's analysis of multi-agent systems describes a "virtual workplace" where specialized AI agents collaborate under the direction of a coordinating agent. Each agent has domain expertise—one analyzing customer transaction data, another generating investment strategies, a third ensuring regulatory compliance. Together, they function as an integrated team capable of tasks no single model could accomplish alone.

4) The Strategic Significance of the Shift

This evolution matters strategically because it transforms AI's role in the organization. In the first era, AI was a back-office tool for efficiency. In the second, it became a source of analytical insight. In the third, it is becoming a core component of customer experience, product innovation, and competitive differentiation. Consider customer interaction. Traditional digital banking offered structured interfaces: menus, forms, and predefined journeys. Conversational AI adds natural language understanding. But agentic AI offers something fundamentally different: the ability to anticipate customer needs and act autonomously. As one technologist puts it, "Think of AI not just as automation—but as prediction. The next leap is knowing what users want before they ask". This shift has profound implications for competitive strategy. Banks that treat AI as an efficiency tool will achieve cost savings but may find themselves competing on parity with institutions that use AI to reimagine customer relationships, product architectures, and risk management. The strategic question is not whether to adopt AI but what kind of AI adopter to become.

High-Value Application Domains: From Efficiency to Strategic Differentiation

The application of AI in banking spans the entire organization, from customer-facing services to back-office operations. Analysis of current deployments reveals six high-value domains where AI is moving beyond efficiency gains to enable strategic transformation.

A) Customer Experience and Hyper-Personalization

Customer-facing applications represent the most visible and rapidly evolving domain. According to EY's 2025 survey, front-office use cases constitute 43% of GenAI applications that have reached production, compared to 34% for middle office and 23% for back office. This reflects both growing confidence in the technology and recognition of customer experience as a primary competitive battleground. Leading institutions are deploying "hyper-

personalization" engines that combine real-time data with generative capabilities. These systems analyze customer transactions, life events, and market conditions to generate tailored recommendations and proactive advice. One major bank's wealth management platform uses AI to create personalized investment portfolios, adjusting allocations in minutes rather than quarters based on changing market conditions and client circumstances.

The strategic implications extend beyond customer satisfaction. Personalized engagement drives revenue growth: analysis suggests AI-powered relationship management can increase active client coverage per relationship manager by 5-15% and per-customer revenue by 5-10%. More fundamentally, it shifts the bank's role from transaction processor to financial partner—a relationship that deepens customer loyalty and creates switching costs.

B) Risk Management and Credit Underwriting

Risk management has been a traditional strength of AI in banking, but generative capabilities are expanding the frontier. Rather than relying solely on structured historical data, next-generation systems analyze unstructured information—news articles, social media sentiment, supply chain relationships—to assess creditworthiness and detect emerging risks. Synthetic data generation represents a particularly important innovation. By creating realistic but artificial datasets, banks can develop and test risk models without exposing sensitive customer information. Research indicates this approach enables robust predictive modeling while ensuring compliance with data protection regulations.

The same techniques allow banks to stress-test portfolios against scenarios that have no historical precedent—a capability increasingly valuable in a volatile world. Multi-agent systems are transforming credit underwriting. An AI coordinator can decompose the underwriting process into subtasks: verifying income documentation, assessing collateral value, analyzing industry risk, and generating loan documentation. Each subtask is handled by a specialized agent, with results integrated into a comprehensive risk assessment. This approach not only accelerates decisions but improves consistency and auditability.

C) Software Development and Engineering Productivity

Perhaps the most significant near-term productivity gains are occurring in software development. Banks are applying AI-powered coding tools to accelerate development, modernize legacy systems, and reduce technical debt. Automated software engineering platforms can generate code, write tests, refactor existing applications, and produce documentation—tasks that previously consumed substantial developer time. The strategic importance extends beyond efficiency. Legacy modernization has been a persistent challenge for banking, with core systems dating back decades constraining innovation. AI-assisted

refactoring enables banks to incrementally modernize these systems, extracting business logic while preserving stability. As one industry analysis notes, "explainable refactoring becomes the connective tissue enabling this hybrid state"—AI-generated code transformations accompanied by audit trails and traceability to original logic.

D) Knowledge Management and Employee Empowerment

Banks accumulate vast repositories of knowledge—product documentation, regulatory guidance, customer interaction histories, research reports—that are difficult to access and synthesize. GenAI systems are transforming this landscape by creating unified knowledge layers that enable natural language querying across disparate sources. Morgan Stanley's deployment of AI tools for financial advisors illustrates the potential. Approximately 15,000 advisors now use AI to generate meeting summaries, identify next-best actions, and access relevant research. Rather than searching multiple systems, advisors can ask questions in natural language and receive synthesized answers grounded in institutional knowledge. This represents a shift from "find" to "synthesize." Traditional knowledge management required users to locate relevant documents and extract insights themselves. AI-powered systems do the synthesis, presenting not just information but analysis and recommendations. The result is faster decisions, more consistent advice, and reduced training time for new employees.

E) Compliance and Regulatory Reporting

Regulatory compliance has long been a cost center for banking—essential but non-differentiating. AI is beginning to change this equation by automating compliance processes while improving their effectiveness. Natural language processing can analyze regulatory updates and map them to internal policies. Generative systems can draft regulatory filings and compliance documentation for human review. The automation of anti-money laundering (AML) monitoring illustrates the potential. Traditional systems generate numerous false positives requiring manual investigation.

AI-powered systems can analyze transaction patterns with greater sophistication, reducing false positives while improving detection of actual suspicious activity. Some institutions are deploying agentic systems that not only flag anomalies but initiate investigation workflows and generate suspicious activity reports.

F) Fraud Detection and Cybersecurity

Fraud detection has been an AI success story for decades, but generative capabilities are enabling new approaches. Anomaly detection models can identify subtle patterns indicative of fraud, while generative systems can simulate attack scenarios to test defenses. As financial crime becomes more sophisticated, AI-powered defenses become correspondingly critical.

The strategic importance of this domain will only grow. As banking becomes increasingly digital and real-time, the window for detecting and preventing fraud shrinks. Institutions that can protect customers effectively gain competitive advantage through reduced losses and enhanced trust.

Implementation Challenges and Governance Imperatives

The potential of AI in banking is substantial, but so are the challenges of realizing it. Industry data reveals significant implementation hurdles: only 16% of use cases reach full deployment, and 40% of implemented use cases fail to meet expectations. Understanding these challenges is essential for developing effective strategies.

- **Data Foundations and Infrastructure**

Data quality remains the primary constraint on AI performance. Legacy banking architectures typically create data silos that fragment customer information and complicate model training. McKinsey reports that 40% of enterprises have more than 50 data silos, leading to model accuracy degradation of 20-30%. The path to resolution requires investment in modern data infrastructure. Real-time data pipelines, event-driven architectures, and domain-driven data models enable the data readiness that AI performance demands. Banks that modernize data foundations first will see the greatest returns from AI investments. Synthetic data offers a complementary approach. By generating realistic artificial data, banks can develop and test models without accessing sensitive customer information. This enables innovation while maintaining privacy and compliance. However, synthetic data must be carefully validated to ensure it accurately represents real-world patterns and does not introduce bias.

- **Governance, Risk, and Compliance**

The regulatory environment for AI in banking is evolving rapidly. Supervisors increasingly demand explainability—the ability to understand and articulate how AI systems reach decisions. This poses challenges for complex models where decision logic is not transparent. Leading institutions are responding with structured governance frameworks. According to EY's survey, 75% of banks have formal AI governance committees, and 60% grant decision-making authority to executive leadership teams. Centralized governance enables consistent standards while allowing flexibility for business-unit implementation. Key governance elements include:

❖ **Model risk management** adapted for generative AI, including validation protocols and ongoing monitoring

❖ **Explainability requirements** calibrated to use case criticality

- ❖ **Bias detection and mitigation** frameworks that assess models for disparate impact
- ❖ **Audit trails** enabling reconstruction of AI-influenced decisions
- ❖ **Human oversight** protocols defining when and how AI recommendations require human review

The concept of "governed intelligence" captures the integration of AI capabilities with the controls required in regulated industries. This includes observability, traceability, and explainability as design principles rather than afterthoughts.

- **Talent and Organizational Capabilities**

The talent challenge in AI is acute. Estimates suggest China alone will face a shortfall of 5 million AI professionals by 2030, with fewer than 15% possessing the combination of technical and financial expertise needed for banking applications. Addressing this challenge requires multiple strategies. Some institutions will build internal capabilities in priority domains. Others will partner with technology providers and fintech companies to access specialized expertise. Most will need both approaches, combining internal development with ecosystem partnerships. Organizational structure matters as well. Traditional "tech vs. business" silos impede AI integration. Leading institutions are creating fusion teams that combine domain expertise with technical capability, enabling AI development that is grounded in business reality. These teams require new career paths, compensation models, and ways of working.

- **Cost Management and Financial Sustainability**

AI investment carries significant costs, from infrastructure to talent to model development and operation. The energy requirements are substantial: training large models requires computing power equivalent to multiple nuclear power plants. Inference costs—the expense of running models in production—can also accumulate rapidly. Banks are responding with pragmatic approaches. Rather than applying generative AI universally, they are strategically selecting use cases where the technology's capabilities justify its costs. For simple queries, traditional retrieval-based methods may suffice; generative AI is reserved for complex reasoning tasks. FinOps for AI—financial management practices adapted for AI workloads—is emerging as a distinct discipline. Investment funding has also evolved. In 2023, 27% of GenAI initiatives were funded from corporate strategy budgets; by 2025, this had dropped to 6%, with IT and technology budgets now the primary source (65%). This shift reflects maturation: AI is becoming part of normal technology investment rather than a special strategic initiative.

Strategic Frameworks for AI Transformation

Given the complexity and stakes of AI adoption, banks need coherent strategic frameworks to guide their decisions. This section synthesizes industry insights into a structured approach encompassing strategic positioning, use case prioritization, and implementation pathways.

1) Strategic Positioning: Builder, Innovator, or Adopter

McKinsey's analysis identifies three distinct strategic postures for AI in banking:

- **Builders** develop proprietary AI capabilities that differentiate core business processes. They invest in custom models, unique datasets, and specialized talent to create competitive advantage. ING's development of intelligent assessment systems exemplifies this approach. Building requires substantial investment but offers the greatest potential for differentiation.
- **Innovators** focus on transforming specific domains through AI. They may not build foundational models but develop distinctive applications that reshape customer experiences or operational processes. Morgan Stanley's AI-powered advisor tools represent this category: built on underlying platforms but uniquely configured for wealth management.
- **Adopters** leverage commercially available AI tools to improve efficiency. They implement SaaS-based solutions for customer service, document processing, or compliance monitoring. This approach offers rapid deployment and lower investment but limited differentiation as competitors can access similar tools.

The choice among these postures depends on institutional resources, strategic aspirations, and competitive context. Regional banks may find adoption the most practical path, while global institutions may need builder capabilities to maintain competitive positioning. Critically, the choice is not permanent: institutions can migrate between postures as capabilities develop.

2) Use Case Prioritization: From Experimentation to Value

With dozens of potential AI applications, prioritization becomes essential. Analysis of successful implementations reveals common principles:

➢ **Start with operational pain points, not technology.** The most effective use cases address real business problems—processes that are slow, error-prone, or resource-intensive. Technology should be the solution, not the starting point.

➢ **Align with business outcomes.** Success metrics should connect to business performance: time saved, revenue generated, risk reduced. Model performance metrics (accuracy, precision) are intermediate indicators, not ultimate measures of success.

- **Consider data readiness.** Use cases requiring new data sources or significant data cleanup face higher implementation risk. Prioritizing areas with available, high-quality data accelerates time to value.
- **Balance quick wins with strategic bets.** Some investments should target near-term efficiency gains; others should build capabilities for long-term transformation. The optimal portfolio combines both, with quick wins funding and informing strategic initiatives.
- **Involve control functions early.** Risk, compliance, and legal teams should participate in use case development, not merely review completed work. Early engagement prevents later surprises and accelerates approval.

3) Implementation Pathways: From Pilot to Scale

Moving from successful pilot to enterprise scale requires deliberate planning. Key considerations include:

- **Technical architecture:** Pilots may use experimental tools and temporary infrastructure. Scaling requires integration with enterprise systems, security controls, and operational processes. Architecture decisions made early constrain later options.
- **Change management:** AI implementation changes how people work. Effective programs invest in training, communication, and support to help employees adopt new tools. Resistance often reflects inadequate change management rather than technology limitations.
- **Vendor strategy:** The AI vendor landscape is crowded and evolving. Banks must decide which capabilities to build internally, which to source from strategic partners, and which to procure as commoditized services. These decisions shape both cost structure and strategic flexibility.
- **Continuous learning:** AI systems require ongoing monitoring and refinement. Model performance degrades over time as patterns change; feedback loops must capture this degradation and trigger updates. Implementation is not a one-time event but an ongoing process.

4) The Ecosystem Imperative

For most banks, particularly regional institutions, ecosystem partnerships are essential. The complexity and pace of AI development exceed any single organization's capacity. Effective partnerships enable access to specialized capabilities without requiring internal development of everything. Key partnership domains include:

- **Cloud providers** offering AI platforms and infrastructure

- **Model developers** providing foundation models and specialized capabilities
- **System integrators** with implementation expertise
- **Fintech companies** offering domain-specific applications
- **Academic institutions** providing research access and talent pipelines

The strategic challenge is managing these partnerships while maintaining control of proprietary data and distinctive capabilities. "Sovereignty-by-design" architectures enable this balance, ensuring that sensitive data and models remain within institutional control even when leveraging external platforms.

The Future: Toward Autonomous Finance

As AI capabilities continue to evolve, the banking industry faces a future that looks fundamentally different from its present. This section explores emerging trends and their strategic implications.

❖ **The Rise of Agentic AI**

The next frontier is agentic AI—systems capable of autonomous action across complex workflows. Unlike current applications that respond to queries or execute predefined tasks, agentic AI can initiate action, coordinate across functions, and adapt to changing circumstances. Consider mortgage origination. Today's AI might assist with document collection or preliminary credit assessment. Agentic AI could manage the entire process: verifying income, ordering appraisals, assessing title, underwriting the loan, and generating closing documents—all while keeping customers informed and escalating exceptions for human review. This shift has profound implications for operating models. If AI agents can execute end-to-end processes, organizational structures designed around human-performed tasks will require fundamental rethinking. Jobs will shift from execution to oversight, from following procedures to managing exceptions.

❖ **The Personal AI Assistant**

The vision of a personal AI financial assistant is moving toward reality. Unlike today's fragmented tools, future assistants will have comprehensive visibility into customer finances, understand goals and constraints, and act autonomously within defined boundaries. Such assistants could negotiate with service providers, optimize savings across accounts, identify financial opportunities, and alert customers to risks. They would operate continuously, not just when customers initiate interaction. The bank-customer relationship would shift from episodic transactions to continuous partnership. Privacy implications are substantial.

Realizing this vision requires customers to share extensive financial data—raising questions about data sovereignty, security, and trust. Maja Pantic envisions a future where "your biometric data, your financial profile, everything stays on your phone. It does not go anywhere else". Device-local AI, rather than cloud-based processing, could enable personalization while preserving privacy.

❖ **Quantum AI and Long-Term Horizons**

Looking further ahead, quantum AI represents a potential paradigm shift. By exploiting quantum mechanical properties, quantum systems could solve problems intractable for classical computers—portfolio optimization at unprecedented scale, risk simulation with greater accuracy, fraud detection with higher sensitivity. However, practical quantum AI remains distant. Gartner estimates mainstream adoption is more than a decade away, with current systems too error-prone and expensive for enterprise use. Banks should monitor developments and support academic research but maintain primary focus on near-term AI opportunities.

❖ **The Strategic Imperative**

The trajectory is clear: AI will become increasingly central to banking operations, customer relationships, and competitive positioning. Institutions that treat AI as an efficiency play—deploying isolated tools without strategic integration—risk falling behind competitors that use AI to transform their businesses. Success requires more than technology investment. It demands strategic clarity about where and how AI creates value, governance frameworks that enable responsible innovation, organizational capabilities that integrate AI into daily work, and partnership strategies that provide access to specialized capabilities. The institutions that lead will be those that embrace "governed intelligence": the ability to scale AI with transparency, traceability, and regulatory confidence. This is not merely a technology posture but a trust posture—determining which organizations innovate safely, move faster with assurance, and navigate regulatory scrutiny without losing momentum.

CONCLUSION

This article has examined the evolution of AI and generative AI in banking from tactical efficiency tools to strategic transformation drivers. The evidence suggests a fundamental shift: AI is no longer a supporting capability but is becoming the primary interface, intelligence layer, and optimization engine across banking operations. Six high-value domains are driving this transformation: customer experience and hyper-personalization, risk management and credit underwriting, software development and engineering productivity,

knowledge management and employee empowerment, compliance and regulatory reporting, and fraud detection and cybersecurity. Across these domains, AI is moving beyond automation toward augmentation and ultimately agency—systems capable of autonomous action within defined parameters. Realizing this potential requires navigating significant challenges: data infrastructure limitations, governance and regulatory requirements, talent shortages, and cost management. Successful institutions are responding with structured frameworks that align AI investment with strategic priorities, govern deployment with appropriate controls, and build organizational capabilities for sustained innovation.

Looking forward, the rise of agentic AI, personal AI assistants, and ultimately quantum AI will continue reshaping the competitive landscape. The strategic question for banking leaders is not whether to adopt AI but what kind of AI institution to become: builder, innovator, or adopter. The answer will shape competitive positioning for years to come. In an industry built on trust, the institutions that thrive will be those that combine AI's capabilities with the governance, transparency, and human judgment that trust requires. The future of banking is not fully automated or fully human but a thoughtful integration of both—powered by AI, guided by purpose, and accountable to the customers and communities it serves.

REFERENCES

1. Shabu, K., Prabhash, N., & Mageswari, S. S. (2025). Enhancing banking efficiency through synthetic data generation and generative AI applications. In *Advances in economics, business and management research/Advances in Economics, Business and Management Research* (pp. 23–41). https://doi.org/10.2991/978-94-6463-896-7_3
2. Foo Boon Ping. *The end of buzzwords - Financial leaders define the roadmap for generative AI and LLMs in banking*. The Asian Banker. https://live.theasianbanker.com/videocategory/the-end-of-buzzwords-financial-leaders-define-the-roadmap-for-generative-ai-and-llms-in-banking
3. Guntumadugu, V. K. (2025). Generative AI in Financial Services: a Strategic framework for digital transformation. *International Journal of Scientific Research in Computer Science Engineering and Information Technology*, *11*(1), 1553–1559. https://doi.org/10.32628/cseit2 51112145
4. Kumar, Y. (2025). The Dual-Faceted Integration of Generative AI in Banking: Balancing innovation and governance. *Journal of International Crisis and Risk Communication Research*, 354–365. https://doi.org/10.63278/jicrcr.vi.3284

5. Kher, P., Gupta, S., & Trull, Z. (2025, September 30). *How AI in banking can result in major transformative benefits*. https://www.ey.com/en_us/insights/banking-capital-markets/ai-in-banking-ey-parthenon-genai-survey-insights
6. Hettiarachchi, I. (2025). The Rise Of Generative Ai Agents In Finance: Operational Disruption And Strategic Evolution. *Zenodo (CERN European Organization for Nuclear Research)*. https://doi.org/10.5281/zenodo.15285625
7. Sindhu, J. K. (2025, July 9). *Hype Cycle for Artificial Intelligence in Banking, 2025* (ID G00831999). Gartner. https://www.gartner.com/doc/reprints?id=1-2LY4W948&ct=250922 &st=sb
8. Jana. (2025, June 18). *Four GenAI use cases delivering value in banking today*. NayaOne. https://nayaone.com/blog/four-genai-use-cases-delivering-value-in-banking-today/
9. Blakey, D., & Blakey, D. (2026, February 2). 2026: The year BFSI shifts from modernisation ambition to governed intelligence. *Electronic Payments International*. https://www.electronicpaymentsinternational.com/comment/2026-bfsi-shifts-modernisation-ambition-governed-intelligence/?utm_campaign=recommended-articles-pi