

unknown author

ansh_dubey_revised

 Research Project

Document Details

Submission ID

trn:oid:::27535:140241722

Submission Date

May 24, 2026, 7:47 PM GMT+5:30

Download Date

May 24, 2026, 7:49 PM GMT+5:30

File Name

ansh_dubey_revised.docx

File Size

64.8 KB

39 Pages

10,923 Words

62,105 Characters

2% Overall Similarity

The combined total of all matches, including overlapping sources, for each database.

Filtered from the Report

- Quoted Text

Match Groups

-  **21 Not Cited or Quoted 2%**
Matches with neither in-text citation nor quotation marks
-  **0 Missing Quotations 0%**
Matches that are still very similar to source material
-  **0 Missing Citation 0%**
Matches that have quotation marks, but no in-text citation
-  **0 Cited and Quoted 0%**
Matches with in-text citation present, but no quotation marks

Top Sources

- 1%  Internet sources
- 0%  Publications
- 1%  Submitted works (Student Papers)

Integrity Flags

0 Integrity Flags for Review

Our system's algorithms look deeply at a document for any inconsistencies that would set it apart from a normal submission. If we notice something strange, we flag it for you to review.

A Flag is not necessarily an indicator of a problem. However, we'd recommend you focus your attention there for further review.

Match Groups

- 21 Not Cited or Quoted 2%**
Matches with neither in-text citation nor quotation marks
- 0 Missing Quotations 0%**
Matches that are still very similar to source material
- 0 Missing Citation 0%**
Matches that have quotation marks, but no in-text citation
- 0 Cited and Quoted 0%**
Matches with in-text citation present, but no quotation marks

Top Sources

- 1% Internet sources
- 0% Publications
- 1% Submitted works (Student Papers)

Top Sources

The sources with the highest number of matches within the submission. Overlapping sources will not be displayed.

1	Publication	Richard Gwashy Young. "AI and ML Governance in Financial Services - Balancing I...	<1%
2	Internet	www.classace.io	<1%
3	Student papers	University of Cambridge International Examinations on 2016-05-20	<1%
4	Internet	digilib.unimed.ac.id	<1%
5	Student papers	INTI Universal Holdings SDM BHD on 2026-04-22	<1%
6	Student papers	Jackson State University on 2025-04-16	<1%
7	Internet	mfuir.mfu.ac.th	<1%
8	Internet	rps.wku.edu.et	<1%
9	Student papers	ESoft Metro Campus, Sri Lanka on 2025-02-22	<1%
10	Publication	Raul V. Rodriguez. "Artificial Intelligence for Business - An Implementation Guide ...	<1%

11	Publication	Ton Duc Thang University	<1%
12	Student papers	University of Essex on 2026-05-22	<1%
13	Student papers	University of Southern Queensland on 2018-01-03	<1%
14	Student papers	Westford School of Management on 2026-05-21	<1%
15	Internet	ir.uitm.edu.my	<1%
16	Internet	www.theseus.fi	<1%
17	Student papers	London School of Economics and Political Science on 2008-07-05	<1%

EXECUTIVE SUMMARY

At its core, banking is a business of knowledge. It has always been about collecting information, analyzing patterns, and making decisions based on what you've learned. What sets artificial intelligence apart from all other technologies used in banking is that it offers a unique approach to the business of knowledge. AI isn't just faster processing or better algorithms in the traditional sense. It learns through the analysis of data it receives, it evolves through experience, and it sees connections where no one else can even hope to see any. It analyzes huge volumes of data that put all previous banking technologies to shame.

The topic of this Major Research Project concerned this transformation within the worldwide banking and financial services industry. For this study, only secondary sources confirmed by reliable means were used, including peer-reviewed articles, research reports published by consulting companies, international organizations, and public institutional disclosures. This study reviewed six major applications for artificial intelligence, which included fraud detection, credit assessment, customer relationship management using automation, algorithmic trading, enterprise risk management, and regulatory technology.

In terms of market dynamics, the data revealed a definite trend. In 2023, the banking AI sector was valued at 19.9 billion US dollars, before growing to 26.2 billion in 2024, thus seeing growth of about 32 percent over one year. Projections made by Uptech Research in 2024 set the value at 315.50 billion dollars by 2033. Meanwhile, analysts from McKinsey and Company suggested in 2023 that the use of AI would help unlock value in the range of 200 billion to 340 billion dollars annually for the worldwide banking industry, or between nine and fifteen percent of its total operating earnings.

Three banking institutions, namely JPMorgan Chase, Bank of America, and HDFC Bank, were the subject of this investigation. The difference between the early adopters of artificial intelligence solutions and their rivals lay not only in the investments but

also in their mindset towards the technology. Thus, JPMorgan Chase invested as much as 17 billion dollars into technology development in 2024 and utilized more than 450 artificial intelligence programs, which provided the company with the added value of about 1.5 billion dollars each year. Erica, the virtual assistant of Bank of America, received a million inquiries each day, which would have required 3,000 people to process, for over 42 million digital users of the company. It is worth noting that the case of HDFC Bank proves that such innovation does not belong only to Western banking institutions.

However, not everything could be considered rosy. Only about 29% of financial organizations indicated that AI technology provided cost reductions in practice. Less than 40% of AI initiatives generated the expected return on investment. Implementation challenges that typically took up to fourteen months resulted mostly from an insufficient number of experts combining both solid knowledge in artificial intelligence and profound experience in the financial industry. The report concludes with seven key recommendations for those who take AI seriously.



CHAPTER 1: INTRODUCTION

1.1 Background of the Study

There is a very good reason why banking was one of the first sectors to seriously adopt artificial intelligence. Almost more than any other industry, banks depend upon information. Any decision about credit, any warning about fraud, any advice on investing, any compliance reporting depends on the ability to quickly process huge amounts of data. Up until around the middle of the last century, this was done using computer systems which functioned on rules. Such systems were robust for the task within their programming capabilities. One thing that such systems lacked, however, was the capacity to adapt and to learn. They could not respond to patterns which had not been predicted ahead of time. And thus artificial intelligence transformed the business completely.

The history of algorithms in financial decision-making is actually much longer than one would think. The analysts at Fair, Isaac and Company had started creating rudimentary scoring models, which quantified the risks associated with a borrower based on his/her attributes, as early as the 1950s. The algorithms used back then would seem extremely primitive by today's standards; nevertheless, they represented an important first step in concept formation. More sophisticated rule-based expert systems were introduced through the 1980s and 1990s in areas such as lending and compliance activities. Starting from the early 2000s, statistical methods including neural nets and decision trees were used in such tasks as fraud detection and portfolio management. However, a common flaw of all those techniques is the need for humans to preselect the variables as well as establish rules of decision-making.

The three developments that were introduced during the 2010s were cloud technology, digital banking systems, and advancements in deep learning architecture. The advent of cloud computing allowed institutions to process large amounts of data using cloud-based infrastructure without having to purchase hardware for data processing. Digital banking services, online payment systems, and open banking ecosystems were providing much more transactional data than ever before. The developments in deep

learning architecture provided the ability to identify patterns in large chunks of data even when they were unknown beforehand.

These figures certainly reflected this trend. In their research report released in 2024, Uptech Research projected the total size of the AI market within the global banking industry to be 19.9 billion US dollars in 2023 and 26.2 billion in 2024, a year-on-year growth rate of 31.7 percent. Estimates forecasted further expansion over the next ten years to hit 315.50 billion dollars by 2033, which is equal to a 31.83 percent compound annual growth rate. For its part, the World Economic Forum estimated in 2023 that the cumulative investment by financial institutions into their AI capabilities had amounted to 35 billion dollars that year, with expectations pointing toward an annual expenditure close to 97 billion dollars per year by 2027. Analysts from McKinsey and Company estimated in 2023 that the yearly value created from artificial intelligence could range from 200 billion to 340 billion dollars for the banking industry alone, which accounted for 9 to 15 percent of the operating income of the industry.

Table 1.1: Annual AI Investment in Global Banking from 2019 to 2024

Year	AI Investment in USD Billion	Year-on-Year Growth	Primary Technology Driver
2019	9.3	Baseline year	Statistical ML in credit and fraud
2020	11.8	26.9 percent	COVID-19 forces shift to digital banking
2021	14.6	23.7 percent	NLP chatbots enter mainstream use
2022	17.2	17.8 percent	Robo-advisory reaches

			commercial scale
2023	19.9	15.7 percent	Generative AI enters enterprise banking
2024	26.2	31.7 percent	Enterprise-wide AI integration begins

Source: Uptech Research, 2024; IDC Worldwide Banking IT Spending Guide, 2024

As is shown in Table 1.1, something needs to be reflected on in this context. The growth in 2024 was faster compared to two preceding years; it was not less fast, but actually faster. This fact correlated well with the first complete year of widespread implementation of generative AI tools, and it was clear evidence of the fact that the technologies themselves generated a new wave of investments rather than just being driven by trends already present in society.

The competitive pressure exerted by fintech companies became another reason to consider AI adoption. The fintech companies, which were born digital and had implemented AI from the very beginning as an integral part of their product, were scaling up at a pace that far exceeded that of the financial service industry overall. As McKinsey & Company reported in 2025, the annual revenue generated by fintech reached approximately 650 billion USD, growing at an impressive 21% yearly compared to 6% in the financial services industry.

1.2 Problem Statement

The key issue regarding AI implementation in the banking industry did not lie in the technology. The problem was much more organisational and strategic. Millions of dollars were being invested into AI, but its payoffs have been much less reliable than expected. The key reason behind such an issue is what this study seeks to uncover.

The legacy infrastructure became a central issue in the challenge for many established organizations. Major banking software used in large banks in the present days is mostly designed in the 1970s and 1980s, with code written using languages and hardware that did not exist when today's AI systems appeared. Customer information was stored in systems that could hardly exchange data between each other. Thus, while one system would store information about mortgages in an organization, another would hold the data on deposits, and yet another would keep the credit card data, without a possibility to create a complete picture of customer activity. This picture creation process itself, which becomes the main prerequisite of successful AI functioning, was associated with a rather lengthy and expensive process of data processing. According to the World Economic Forum report in 2023, lack of high-quality data and issues with the infrastructure were listed as the main impediments of AI implementation by more than half of the financial firms.

The economics behind AI projects was similarly discouraging. In a 2024 study conducted by Boston Consulting Group, just 29 percent of banks had reported that AI was helping reduce costs. Deloitte's Financial AI Adoption Report, published in the same year, made it clear that less than forty percent of banking AI projects achieved the promised return on investment levels. A study carried out by Ernst and Young revealed that 65 percent of banks faced project delays ranging between 14 months, with the key reason for such delays being a lack of individuals combining financial expertise with knowledge of AI technology.

There was also a third issue that did not get much publicity, but was becoming increasingly important from the standpoint of regulation. Many of the most accurate models of artificial intelligence used for credit scoring and fraud detection were neural networks, and their reasoning processes were not understandable in terms of human logic. In almost all cases, credit providers had an obligation to offer an understandable explanation for a negative decision on credit application. A model of artificial intelligence which could only say that the algorithm produced a certain result would be unacceptable according to these regulations. There were ways to make artificial intelligence understandable, but they were not quite developed enough when this study was written, and added costs to implementation.

1.3 Objectives of the Study

The study was anchored on the following specific objectives:

- i. To explore the history and assess the current status of Artificial Intelligence implementations in global banking and finance service firms.
- ii. To conduct an assessment at the domain level of Artificial Intelligence applications in banking in areas such as fraud detection, loan processing, automated customer care, algo-trading, and risk management.
- iii. To assess the effects of Artificial Intelligence in important performance parameters through validated secondary data collected from authentic sources.
- iv. To conduct a thorough assessment of AI strategies and accomplishments of JPMorgan Chase, Bank of America, and HDFC Bank.
- v. To analyze and evaluate the major hurdles and challenges that prevented successful implementation of AI in banks.
- vi. To develop useful and evidence-based strategies for organizations that wish to optimize AI implementations within their business operations.

1.4 Scope of the Study

The study is based on applications of AI in the field of banking and finance, which includes commercial banking, investment banking, universal banking, and other non-banking financial institutions. This analysis is done for a period ranging between 2019 and 2024, as these years were specifically selected to include both periods of pre-generative AI as well as early large language models.

Three organizations were selected as subjects for this case study. JPMorgan Chase was chosen as the most progressive AI user of North American universal banks. Bank of America served to showcase a massive use of AI for consumers. And HDFC Bank was included due to its presence in emerging markets and its capacity to provide a completely different perspective from the other two.

The study relied solely on secondary sources and did not make use of surveys, structured interviews, or experimental designs. AI use in insurance and capital markets, and cryptocurrencies alone were outside the scope of interest, unless the interaction with traditional banking brought relevance to the study.

1.5 Significance of the Study

What made the study distinct from other existing studies is that it synthesized the information gathered from various sources in the literature to generate new knowledge. From a purely analytical perspective, the study incorporated evidence from the literature of journals, consulting firms, institutional statements, and market information in one consolidated format that was not possible individually for each source.

From the viewpoint of professionals, the case studies provided institutional insights that were grounded in actual data verified specifically for each organization, rather than generic knowledge pertaining to the sector as a whole. One such case study was that involving JPMorgan Chase.

From an academic perspective, AI was considered not as one single technology, but as a range of technologies each having its own pathway to adoption, performance level, and organizational impact. This consideration brought a greater degree of depth to the analysis than the simplistic approach that prevailed in the banking literature to date.

And from the perspective of policymakers, the information relating to potential bias risks, concerns with explainability, and issues of regulatory compliance regarding the new AI regulation provided substantial material for deliberation over the question of how to regulate AI within banking without impeding innovation.

CHAPTER 2: LITERATURE REVIEW

2.1 Technology Acceptance Theory in Financial Contexts

It is necessary to take a step back and first consider what motivates organizations to adopt a new technology and why some implementations are successful while others fail. The answer to the first question can be found in what Davis, Bagozzi, and Warshaw called the Technology Acceptance Model in their paper published in 1989. In essence, the authors claimed that a technology will be successfully adopted if users believe that it increases efficiency and that using it does not add excessive complexity to their work. Later scholars have expanded on TAM, incorporating such constructs as trust, social influences, and organizational pressures into the analysis. In the case of financial services, Marikyan and Papagiannidis point out in their 2023 review of the literature in TheoryHub that banking happens to be among the industries with the highest validity evidence supporting frameworks that were built upon the concept introduced by Davis, Bagozzi, and Warshaw because the effect of technology adoption could be readily quantified.

The thing that distinguished the advent of AI in banks from other technological advancements in banking was the notion of usefulness itself. When it came to technology such as ATMs, internet banking, or electronic payment systems, usefulness could be considered a matter of convenience, speed, and accessibility. In the case of AI usefulness, it had more to do with the aspect of enhancing accuracy, minimizing risks, and offering something that had not existed before at all. This was validated by Gautam through an empirical study done in 2023 where he analyzed the effect of AI on risk management and fraud detection in banks and found that usefulness perceptions in AI were directly related to improved operational efficiency and results.

An article on the scientometrics of artificial intelligence in finance was recently published in Humanities and Social Sciences Communications that examined development of scholarly thinking on AI in financial services over a period of five decades through published scientific literature. There has been an evident trend from expert systems to statistical learning to deep learning and generative AI algorithms. It

is also worth pointing out that each time, AI technology evolved in terms of what it was able to do within the field of finance rather than becoming better at doing the things which were already being done. It is the present phase which stands out as special as it tackles the cognitive work of writing documents.

2.2 Machine Learning in Fraud Identification

Fraud prevention was one of the early and most commercially successful applications of AI in the banking sector. The nature of binary outcomes, vast amounts of data from past experiences, high losses incurred by banks in either direction, and the adversarial relationship between banks and criminals who would continue to adapt and change methods made fraud prevention a prime field for testing the potential of machine learning technology. A study on AI-assisted fraud detection published in the Journal of Big Data by Springer Nature in 2024 revealed an accelerated trend of publication starting from 2022. Three main factors contributed to this trend: improved deep learning architectures for handling sequence data, more access to actual payments data via open banking standards, and increasing regulatory scrutiny over fraud prevention practices in banks.

Analysis presented on arXiv in the year 2025 highlighted the evolution in year-on-year improvements in algorithmic solutions for fraud detection in the finance industry. Long Short-Term Memory (LSTM) networks had seen the greatest progression in their adoption rate, especially in the years 2022 to 2024. This was quite logical given that fraud often occurred in unusual sequences of transactions, not necessarily just individual occurrences, and LSTM networks were capable of understanding long-term dependencies in sequences of data. In terms of practical applications, ensemble techniques combining the predictions of several different algorithms proved superior in their ability to generalize to new patterns of fraud.

The practical ramifications of this difference were empirically recorded in a study carried out by Sabareesh and published in the Journal of Innovative Engineering Research in 2024. In a comparison of AI-based approaches in fraud detection algorithms, especially XGBoost and deep neural network algorithms compared with

traditional rule-based algorithms, it was revealed that AI techniques performed better in terms of accuracy, false-positive rate, and speed of processing. According to Caspian One in 2025, the use of AI-based algorithms for detecting fraud in banks has resulted in accuracy rates of over 90 percent, compared with traditional rule-based algorithms' 65 to 70 percent accuracy rates.

2.3 Creditworthiness Evaluation Through AI Models

Credit scoring was at the point of confluence of financial profits, equality in society, and legal obligations, making this topic of research the most popular among all those related to the application of artificial intelligence in banking studies. The key disadvantage of logistic regression, which was a traditional method for conducting credit scoring, was based on the idea that all dependencies between independent and dependent variables had to be linear. However, real-life creditworthiness was represented by complicated non-linear functions of dozens of variables.

Amarnadh and Moparthi conducted an extensive review article on the use of AI models in credit risk assessment in 2023, which was published in the IDT Journal. The primary takeaway from their study was that ensembling models were more effective than single classifiers, in terms of both prediction accuracy and stability. They further pointed out the conflict between model accuracy and interpretability, in that while the highly accurate deep learning models generated results via very complicated processes, thus making it impossible to fulfill the requirements of regulatory explainability.

Teodorescu and Obreja Brasoveanu provided empirical proof of the existence of the accuracy problem in a comparative research conducted in 2025 through an 80 to 20% data split, five-fold cross validation applied to a real banking dataset. The outcome showed a high accuracy range of 0.89 to 0.91 of Area Under the Receiver Operating Characteristic Curve scores obtained from artificial intelligence ensemble models, compared to 0.74 for logistic regression models. In 2022, Uddin et al., in their article published in the International Journal of Finance & Economics, analyzed the application of random forests models in improving accuracy and interpretability in the context of credit decision-making for microenterprises.

2.4 Conversational AI and Customer Engagement

The application of natural language processing in customer interactions within banks received considerable focus from both academia and practitioners during the research period. According to a report issued by Forrester Research in 2023, banking and finance emerged as one of the top industries in terms of adoption of conversational AI due to inherent characteristics making them fit for use of this technology. The nature of customer interaction in the banking industry was characterised by frequent repetition of questions, round-the-clock availability requirement, cost pressures per engagement, and large number of engagements that could be resolved without any human intervention.

Customer adoption of AI-based interactions in banking operations was another area of academic research which always found that trust was an important mediating variable in such interactions. Technical proficiency alone could not be sufficient in influencing customer adoption behavior. It was necessary that customers should have perceived the AI-based platform to be truthful in communicating its technical limitations, cautious in handling customers' private information, and able to hand over the communication seamlessly to human representatives in case of a problem beyond the capacity of the AI system.

2.5 Automated Trading and Wealth Advisory Systems

In the sphere of trading and portfolio management, AI technology had been applied much more extensively than anywhere else in finance. By the mid-2020s, automated trading platforms utilizing artificial intelligence software were responsible for a dominant share of transaction volume on the biggest equity and bond markets. Manual human discretionary trading was gradually restricted to situations where subjective analysis of qualitative factors - managerial style, geopolitical risks, new market environments – had an edge over computer algorithms.

The LOXM equity trading platform created by JPMorgan Chase provided a clear example of how reinforcement learning could be used in practical trading execution.

Having been trained using a massive amount of transaction data numbering in the billions, the LOXM algorithm was able to optimize both minimizing market impact and maximizing speed of execution, which were dual objectives achieved by human traders through heuristic processes but pursued by AI through learning policies from transaction data. McKinsey and Company reported a productivity improvement rate of 20 to 60 percent along with a credit decisioning cycle time reduction of about 30 percent using multiagent AI in credit underwriting processes in 2024.

2.6 Compliance Automation and RegTech

The use of artificial intelligence for compliance purposes became one of the fastest-growing application areas in banking due to the joint effect of two factors: rising costs associated with compliance and increasing computing power needed to conduct regulatory compliance. Global banks invested from \$1 billion to \$5 billion per year into compliance operations after the financial crisis of 2008. Artificial intelligence could not only help save money but also provide better results than alternative solutions.

Regulatory regimes themselves were getting into the act when it came to the regulation of AI. With the adoption of the AI Act in July 2024, the European Union classified financial services as a "high risk" form of AI applications, requiring AI systems involved in credit decisions, fraud detection, and insurance underwriting to demonstrate certain levels of transparency, documentation, and human involvement in their decision-making processes. The Bank of England's 2024 survey of AI use by UK financial services organizations contained new questions concerning generative AI, which demonstrated just how quickly AI technologies were gaining regulatory relevance. An example in point is Citibank's deployment of generative AI to analyze an entire 1,089-page report reviewing the implications of new US capital regulations.

Table 2.1: Consolidated Summary of Literature Reviewed

Author and Year	Research Focus	Principal Finding	Publication
-----------------	----------------	-------------------	-------------

Davis et al., 1989	Technology Acceptance Model foundations	Perceived usefulness and ease of use are the primary adoption drivers	Management Science
Amarnadh and Moparthy, 2023	AI methods for credit risk evaluation	Ensemble models outperform individual classifiers consistently	IDT Journal
Gautam, 2023	AI impact on risk and fraud in banking	AI adoption correlates with efficiency and risk-adjusted returns	AI and IoT Review
Sabareesh, 2024	AI-driven fraud detection evaluation	AI achieves above 90 percent accuracy versus 65 to 70 percent for rule-based	JIER
Uddin et al., 2022	Random forest for micro-enterprise credit	Improves accuracy and interpretability in small business lending	Int'l J. Finance and Econ.
Nature, 2025	Scientometric review of AI in finance	Structural shift from rule-based systems to generative AI across banking	Humanities and Soc. Sci. Comms.
ArXiv, 2025	Deep learning trends in fraud detection	LSTM shows strongest growth; ensemble methods dominate production systems	arXiv preprint
McKinsey and Company, 2023	Generative AI potential in banking	200 to 340 billion USD annual value potential; 20 to 60 percent productivity gains	McKinsey Reports
Forrester Research, 2023	Conversational AI in customer service	AI handles 70 to 80 percent of routine banking queries autonomously	Forrester
Teodorescu and Obreja, 2025	AI versus logistic regression in credit scoring	AI ensemble AUC-ROC of 0.89 to 0.91 versus 0.74 for logistic regression	Banking Research
Springer Nature, 2024	AI fraud detection systematic review	Research acceleration from 2022; deep learning is the preferred approach	Journal of Big Data

Source: Compiled from academic and industry publications, 2022 to 2025



CHAPTER 3: RESEARCH METHODOLOGY

3.1 Research Design

The research design used for the study was a descriptive approach based on secondary data analysis. This was not accidental. In this case, the purpose of the research did not lie in exploring cause-and-effect relationships but in describing and analyzing an already existing situation. Hence, the use of a descriptive approach was justified. From the perspective of philosophy, this study was conducted from the point of view of positivism, according to which knowledge about the economic environment could be achieved through observations and inferences from measurable data.

The decision to use secondary data exclusively had an analytical rationale behind it rather than being a convenience. There was no method of collecting primary data that would come close to the geographical scope, range of organizations involved, and quantity of information that could be gleaned from the analysis of secondary data sources. The performance indicators, amounts invested in technology, and outcomes achieved by the organizations under consideration - JPMorgan Chase, Bank of America, and HDFC Bank – could only be found in their annual reports and other public domain research.

Evidence triangulation was used throughout the methodological process. Evidence gained from scholarly publications, reports of consultants, market studies, and self-revelation by institutions was compared across all steps. In cases where multiple independent sources revealed similar results, it was considered as conclusive evidence. But when there was a conflict of information, such as an overly optimistic prediction from the industry and a slightly conservative viewpoint from the academic community, both views were considered in the analysis without discarding either of them.

3.2 Secondary Data Collection Strategy

The process of collecting data was based on four predetermined criteria used uniformly throughout all sources used. First among these is the reliability of the source

organisation itself. Second is the recency of the article, which is prioritised because of the rapid evolution of AI technology, and which was heavily preferred in 2020 and onwards. The third criterion is relevance to the topic of artificial intelligence in banking. The fourth criterion for choosing a source is its verifiability, via an available public link to the source article.

Database research was conducted through Google Scholar, Springer Nature, SAGE Journals Online, Elsevier ScienceDirect, and arXiv. The keywords used included combinations of “artificial intelligence banking”, “machine learning financial services”, “generative AI banking”, “credit scoring deep learning”, “fraud detection neural network”, and “banking digital transformation AI”. Only studies published after January 2019 were considered except in cases where the study had theoretical significance and value regardless of its age.

The following industry publication databases belonging to McKinsey & Company, Deloitte Insights, Boston Consulting Group, World Economic Forum, Gartner, International Data Corporation, and Uptech Research have been searched directly for any relevant reports that address the topics related to the application, investment made, and results achieved through AI within financial services. In addition, annual reports, investor presentations, technology strategy presentations, and verified journalist analysis from reputable publications like the Wall Street Journal, Bloomberg, Financial Times, and fintech magazines have been thoroughly reviewed.

3.3 Source Portfolio and Authentication

Table 3.1 presents the categorised portfolio of the primary secondary data sources drawn upon in this study.

Table 3.1: Portfolio of Authenticated Secondary Data Sources

Category	Organisation or Publication	Nature of Data Provided	Year
----------	-----------------------------	-------------------------	------

Global Consulting	McKinsey and Company	AI value potential, adoption surveys, generative AI analyses	2023 to 2025
Global Consulting	Boston Consulting Group	ROI data, efficiency gains, implementation benchmarks	2024
Global Consulting	Deloitte Insights	AI adoption rates, project ROI achievement, implementation surveys	2023 to 2024
Global Consulting	Ernst and Young	Implementation delays, talent shortage statistics	2024
International Organisation	World Economic Forum	Total AI spending projections, financial inclusion analysis	2023
Market Intelligence	Gartner Inc.	AI adoption scale projections, technology maturity cycle data	2023
Market Intelligence	IDC	Banking AI investment data, sector rankings	2024
Market Intelligence	Uptech Research	AI market size in banking, CAGR projections to 2033	2024
Academic Journals	Springer Nature and Elsevier	Peer-reviewed empirical studies on AI models in banking	2022 to 2025
Institutional	JPMorgan Chase	AI platforms deployed, technology budget, performance outcomes	2023 to 2024
Institutional	Bank of America	Erica virtual assistant metrics, digital user statistics	2023
Industry Analysis	Emerj AI Research	JPMorgan AI case study, industry use case analysis	2024 to 2026
Industry Analysis	Caspian One	Fraud detection accuracy benchmarks, AI adoption survey data	2025

Source: Compiled by Researcher, 2025

3.4 Analytical Framework

This analysis was conducted using an innovative three-level approach designed specifically for this research to categorise information and to make sure that all aspects

of AI implementation in banking were considered. Firstly, it considered the macro level by analysing the market scale and investment volumes, and AI adoption indicators in the industry as a whole. Secondly, it considered the mid-level aspect by analysing the AI performance and adoption aspects within six areas of applications outlined in the research questions. Thirdly, it considered the micro level of analysis by examining how particular banks managed AI.

Research from academia and the markets proved valuable at the first and second levels because of the strength of their quantitative approach and their broad institutional reach. At the third level, where specific program data was necessary, case studies were conducted using information from corporate reports, corroborated journalistic reporting, and specialized AI research. Cross-triangulation of data between the three levels – where macro-level adoption trends were verified through the use of domain-level performance statistics and ultimately confirmed by institutional-level case studies – formed the bedrock of the study's strongest conclusions.

3.5 Limitations and Mitigation Strategies

Four potential issues associated with the research design should be acknowledged. First, there was source optimism bias. Consultancies that stood to profit from providing AI advisory services might have been structured in such a way as to promote positive results. However, the possibility for this type of bias was controlled through cross-referencing the findings of consulting sources with those found in academic literature.

The second limitation was temporal lag. Given that AI technology was evolving rapidly in the banking industry, information published in 2024 could already be partially outdated at the time of writing this paper. This issue was handled by emphasizing on the direction of the results instead of regarding specific numbers as eternal truths. The geographic distribution of available data constituted the third limitation. While there was plenty of information available about banking practices in North America and Europe, there was considerably less available data regarding banks in Africa, Southeast Asia, and Latin America. Disclosure of information about banking

organizations was the fourth limitation since financial organizations were prone to promoting success and hiding failures.

CHAPTER 4: ANALYSIS, DISCUSSION AND RECOMMENDATIONS

4.1 The Macro Landscape of AI Adoption in Banking

What was revealed on a macro level in the course of research was the presence of tremendous momentum and considerable diversity in both quality and consequences of deployments. In terms of general tendencies, the situation was clear-cut. From an experimental to a strategic resource, AI has transformed in banking from 2019 until 2024. Yet again, the distribution of its benefits turned out to be rather uneven.

Table 4.1: AI Market Valuation in Banking from 2023 to 2033

Year	Market Value in USD Billion	Status	CAGR
2023	19.9	Recorded	Baseline
2024	26.2	Recorded	31.7 percent
2025	34.5	Projected	31.83 percent
2027	59.8	Projected	31.83 percent
2030	143.6	Projected	31.83 percent
2033	315.5	Projected	31.83 percent

Source: Uptech Research, 2024; World Economic Forum, 2023

An annual growth rate of 31.83 percent per year over the course of ten years was certainly something worthy of close consideration instead of just blind acceptance. This meant that the market would grow over 15 times larger than its size in 2024 by the year 2033. There certainly have been historical instances where technology markets have grown at a similarly high rate within a similar time frame, such as the early commercial Internet, mobile phones, and cloud computing. The common factor that all of these trends had was their wide range of applicability in multiple industries, powerful network effects, and decreasing unit costs as more users adopted the technologies. These were precisely the same characteristics possessed by AI in

banking. The key issue with this sort of projection is that it was contingent on many external factors.

Table 4.2: Geographic Distribution of AI Adoption in Financial Services in 2024

Region	Adoption Rate	Lead Application Areas	Primary Adoption Driver
North America	78 percent	Algorithmic trading, fraud detection, generative AI workflows	Scale, fintech competition, large technology budgets
Asia-Pacific	71 percent	Mobile AI banking, credit scoring, conversational AI	Mobile-first customer base, digital-native population
Western Europe	63 percent	RegTech, compliance automation, conversational AI	EU AI Act compliance pressure
United Kingdom	63 percent	Open banking AI, fraud systems	FCA regulatory sandbox, open banking mandate
India	52 percent	Credit scoring, chatbots, KYC automation	Digital India initiative, UPI transaction data ecosystem
Middle East	45 percent	Digital banking transformation, mobile services	National AI strategies in UAE and Saudi Arabia

Source: McKinsey and Company, 2024; EY Financial Services CTO Survey, 2024; Uptech Research, 2024

As we can see from Table 4.2, the reasons behind the patterns are distinctly structural in nature and not simply varying levels of technological advancement. With an impressive 78 percent usage rate, North America, which is filled with major technology firms, has ample resources, and faces stiff competition from fintech startups, comes as no surprise. The Asia-Pacific region, second place holder with its 71 percent usage rate, demonstrates quite a contrasting reason for such high numbers. Unlike in Europe and North America, many countries in the Asia-Pacific region have

not had the legacy banking system hindering AI implementation there and in addition to this, huge amounts of transactional data provided by such giants as China and India offer a unique environment for AI learning. On the other hand, only 52 percent in India represents another case study of sorts.

Another interesting case study is the situation in the United Kingdom, where the results showed that the regulation design could be used as a mechanism facilitating AI usage. The project of the regulatory sandboxing launched by the FCA made it possible to implement AI technology in practice in trading environment under control. Together with open banking project that implemented the standardization of the data exchange processes, such situation was deemed to be quite fruitful for AI implementation in the UK. Based on the research carried out by the Lloyds' Bank, it became evident that the rate of AI implementation in the UK doubled within a year and reached 63%.

4.2 Domain-Level Analysis of AI Applications

4.2.1 Fraud Prevention through AI

Among all the use cases that were explored, it is perhaps no coincidence that the domain where AI provided the most compelling and consistent evidence of creating demonstrable value for banking was fraud prevention. This was because virtually every feature of the underlying problem domain was ideal for machine learning algorithms. To start with, there was an unambiguous binary target variable. Secondly, large amounts of historical labeled training data were available. Thirdly, the cost of missing frauds and the cost of false positives were each significant in their own right. Finally, the nature of the adversarial relationship meant that the efficacy of rule-based systems became less relevant over time.

Table 4.3: Comparative Effectiveness of Fraud Detection Approaches

Detection Method	Accuracy Rate	False Positive Rate	Processing Latency	Adaptive Capability
Rule-based Systems	65 to 70 percent	30 to 35 percent	Hours	Low, requires manual update

Standard ML Models	85 to 90 percent	10 to 15 percent	Seconds	Moderate, periodic retraining
Deep Learning Models	90 to 96 percent	5 to 10 percent	Milliseconds	High, continuous learning
AI Ensemble Systems	92 to 98 percent	3 to 8 percent	Milliseconds	Very High, multi-model adaptation

Source: BioCatch Industry Report, 2024; Caspian One, 2025; Sabareesh, 2024

The figures shown in Table 4.3 translated into practical implications beyond simple percentage differences within a statistical measure. The difference between having 65% to 70% success for rules-based systems versus 92% to 98% success for AI ensemble models amounted to going from missing about one-third of cases of fraudulent activity to missing fewer than one out of every twelve. For a bank handling millions of transactions on a day-to-day basis, such an improvement would save the institution hundreds of millions of dollars annually by preventing losses from fraud. As far as false positives are concerned, moving from 30% to 35% false positives to 3% to 8% made all the difference in providing customers with a smooth experience. In the former scenario, the institution was blocking two to four legitimate customer transactions for every single fraudulent case that was detected.

In the case of NatWest Bank, it was clearly shown through its recorded outcomes that using AI resulted in such an advantage. Post AI fraud detection systems installation, new account fraud reduced by 90 percent since 2019, along with lower customer inconvenience due to inaccurate transaction blocking. Similarly, a second case related to the implementation of AI can be found in the AI-based KYC automation system used by HSBC Bank. Through computer vision and NLP technology, the bank used these tools to validate documents and compare them against its fraud database, resulting in a 12-day process shortened to less than 24 hours and improving accuracy from 87 percent to 99 percent.

4.2.2 AI-Driven Credit Evaluation

AI had its most systematic and significant impact in credit risk assessment, an area where the benefits of AI over conventional approaches in terms of analysis had been most well demonstrated. Conventional credit assessment involved using logistic regression analysis to predict default probabilities, which implied a linear relationship between various attributes of the borrower and the chance of default. This was helpful for computational purposes but rather simplistic when it came to modeling the actual factors involved in credit scoring. In fact, there were thousands of different variables at work, and their relationship to each other was anything but linear.

Table 4.4: Predictive Performance of Credit Scoring Models

Performance Metric	Logistic Regression	Random Forest	XGBoost	Deep Neural Network
Overall Accuracy	72.4 percent	84.6 percent	87.2 percent	89.1 percent
AUC-ROC Score	0.74	0.86	0.89	0.91
Precision	68.3 percent	81.2 percent	84.7 percent	87.3 percent
Recall	71.8 percent	82.9 percent	85.6 percent	88.4 percent
Processing Speed	Fast	Moderate	Fast	Slow
Interpretability	High	Medium	Medium	Low
Regulatory Acceptance	High	Moderate	Moderate	Low

Source: Adapted from Amarnadh and Moparthy, 2023; Teodorescu and Obreja Brasoveanu, 2025

Table 4.4 explained the reason behind the predominance of XGBoost as the favored algorithm used in practice for credit scoring applications even though it did not provide the highest accuracy scores. By virtue of its very high accuracy score, its fast processing time, and moderate interpretability, it offered an optimal trade-off between accuracy and interpretability in line with regulatory requirements. Deep neural networks provided the highest accuracy scores in each metric but were limited by lack of interpretability, relatively slower processing, and poor acceptability within regulatory frameworks for credit decision making.

Nevertheless, the inclusion aspect of AI-based credit scoring was worth exploring despite not being prominently featured in the practitioner literature. Conventional credit scoring left out many individuals from accessing loans despite meeting eligibility criteria because they did not have any credit history – young people, new immigrants, people living in rural areas, and entrepreneurs with no credit history. The incorporation of alternative data in AI-based credit scoring systems allowed for more credit-worthy customers to be identified. In the UK, some banks were using AI-based credit scoring to automatically approve loans of up to 100,000 US dollars within a very short time.

4.2.3 Conversational AI in Retail Banking

In fact, by 2024, it can be stated that conversational AI had become the most recognizable usage of artificial intelligence in retail banking. The reason behind such an outcome is connected with the fact that every bank offered its clients at least one type of customer service system utilizing artificial intelligence technology. They varied from basic chatbot assistants for frequently asked questions to more sophisticated virtual assistants capable of performing more complex functions.

The McKinsey analysis of the state of retail banking for 2024 revealed that the share of banking clients who use mobile and digital banking services had risen by 18 percentage points from 2020 to 2023 to reach 57 percent worldwide. The frequency of contact between banking customers and the digital interfaces of their banks rose by 72 percent in the same period. It was this growing engagement on digital platforms that resulted in both the need for AI-based solutions and the availability of training data to make such solutions highly effective.

The best deployments were those where the efficiencies of volume processing using AI were augmented by intelligent routes for escalations to the human agents. According to research carried out by Forrester Research in 2023, up to 70 percent to 80 percent of questions posed to banks could be dealt with through AI without any sacrifice in quality. For more complicated transactions like dispute resolution, mortgage advice, and situations requiring dealing with financial hardship, however, there was still a strong customer preference for human interaction. The issue here was

not whether AI or humans were best but where each fit in the service chain and how to make the switch imperceptible to the customer.

4.2.4 Algorithmic and AI-Driven Trading

Regarding the use of artificial intelligence in trading and investment management, the development was further and more comprehensive in this sector than in any other part of finance during the period considered for this research. Trading systems based on AI constituted a large proportion of the executed volume of transactions in significant stock and bond markets. Discretionary human trading became more prevalent in situations in which factors, like the credibility of management or geopolitical uncertainty, could be used to gain an edge not yet achieved by algorithms.

JPMorgan Chase developed the LOXM equities execution system which is the highest level of development of such systems as they used reinforcement learning relying on training data involving billions of transactions that allowed them to learn execution algorithms depending on the constantly changing microstructure of the present-day stock market environment, with live liquidity, price impact, and volatility accounted for. This system was unique compared to earlier innovations since what it delivered was not an optimization of the process, but rather a system that improves itself all the time. According to McKinsey & Company in 2024, the use of AI-based multiagent systems for preparing credit memos was associated with productivity improvement of 20 to 60 percent alongside with increased decision-making efficiency by 30 percent. As anticipated, in 2026, productivity improvement in front-office operations of investment banks will be at the 27 to 35 percent level.

4.2.5 Enterprise Risk Management and RegTech

The use of artificial intelligence in the domain of enterprise risk management had the capacity to change the practice of risk monitoring from one that was sporadic and reactive to one that was proactive. Conventional systems of risk reporting relied on cycles during which data became stale as soon as it made its way to those making decisions. Artificial intelligence systems working with live data feeds would be capable of detecting risk trends as they occurred – abnormal trading behaviors that

signaled market manipulation, suspicious transactions that suggested credit deterioration.

Research by Kazakh National University published in 2025 based on machine learning models applied to open banking data and surveys conducted among banking experts in the Middle East showed that ensembles, such as XGBoost and Random Forests, greatly surpassed traditional risk assessment techniques in classifying default risks in loans. It is particularly interesting how the dual methodology employed by the researchers shed light on the importance of explainability for practitioners' use of AI risk management solutions: decision-makers needed to know why something worked.

This use-case study shows how effectively and efficiently AI technologies can be applied to regulatory compliance purposes. For example, AI technology could be effectively used to detect and prevent money laundering schemes. The rules-based approach led to excessive rates of false positives in anti-money laundering transaction monitoring systems. As much as 95% - 99% of alerts generated by such systems appeared to be false positives and required additional investigation by the compliance department; however, these attempts yielded no results. The introduction of AI technology in such processes reduced the number of false positives and increased detection rates. This was especially well shown by the AI system developed by JPMorgan Chase – COIN – which could analyze commercial loan contracts, which required an estimated 360,000 hours per year of manual work, within several seconds and without any errors.

4.3 Institutional Case Studies

4.3.1 JPMorgan Chase

The position held by JPMorgan Chase in relation to AI in global banking was one truly unique in nature, characterized not just by the large scale of investments but by the strategic approach employed for developing its AI capacity. The investment of 17 billion US dollars in technology in 2024, over 450 AI use cases being implemented and generating revenues from the applications of AI amounting to 1.5 billion dollars each year, made the institution the best studied example of implementing AI in global

banking. The uniqueness of JPMorgan Chase's strategy for developing AI lies not in any particular program but in the conscious efforts directed towards establishing an infrastructure for collaborative applications of AI technologies.

Table 4.5: AI Programme Benchmarking Across Major Global Banks in 2023 to 2024

Institution	AI Platform or System	Primary Application Area	Documented Outcome
JPMorgan Chase	LLM Suite	Research, drafting, and analytics	Over 200,000 employees onboarded; 1.5 billion USD annual value
JPMorgan Chase	COIN	Contract analysis automation	360,000 manual hours replaced annually
JPMorgan Chase	Coach AI	Wealth management advisory	95 percent faster response; 20 percent gross sales increase
JPMorgan Chase	OmniAI	Real-time fraud detection	250 million USD annual savings; over 1 billion USD loss prevention
Bank of America	Erica	Virtual customer assistant	Over 1 million daily inquiries; 42 million plus active users
HSBC	AI KYC System	Know Your Customer automation	12 days to 24 hours; accuracy from 87 to 99 percent
Goldman Sachs	AI Code Assistant	Software test generation	Manual test generation substantially automated
NatWest	AI Fraud Detection	Transaction fraud prevention	90 percent reduction in new account fraud since 2019
HDFC Bank	EVA and AI Credit Models	Customer service and lending	2.7 million plus conversations; 85 percent autonomous resolution

Source: AI Expert Network, 2024; Emerj, 2026; Bank of America Annual Report, 2023; Uptech Research, 2024

The JPMorgan program was placed within a comparative context in Table 4.5. While the broad nature of the portfolio is important to note, what is also remarkable about JPMorgan's portfolio is its inclusion of both narrow-task oriented software such as COIN, and general productivity platforms such as the LLM Suite. Organizations which used both approaches created greater value than organizations which only used one approach. The depth of JPMorgan's portfolio ranging from productivity tools to fraud prevention, wealth advice, and compliance demonstrates the importance of scale over application when realizing value from AI technology.

The implementation of the LLM Suite was arguably one of the most valuable experiences with AI scalability at any multinational organization in the world. Introduced over the summer of 2024, the provider-agnostic platform included several large language models by different vendors and linked to in-house databases. It was implemented for more than 200,000 employees within eight months through a hands-on approach that prioritized practical implementation over theoretical training. Users reported an estimated efficiency increase between 30 and 40 percent in knowledge-based tasks. Engineers working with AI-assisted coding tools saw a productivity increase ranging from 10 to 20 percent. The overall annual value attributed to LLM Suite applications summed up to roughly 1.5 billion US dollars per year, including faster research, faster document creation, improved risk analysis, and client interaction preparations.

JPMorgan used the Coach AI advisory system in their asset and wealth management unit. This AI application proved to be successful as it brought in revenue for the bank through growth opportunities instead of cutting costs. By offering real-time access to market analysis, synthesized research, and an individualized investment recommendation structure via natural language processing, Coach AI enabled advisors to respond to clients' queries almost 95 percent faster when there were issues in the market. This allowed for 20 percent more gross sales in asset and wealth management between 2023 and 2024, thus proving that AI in banking could bring profits, contrary to popular belief.

Concerning fraud, OmniAI detected suspicious activity by analyzing about 10 trillion US dollars in daily transactional data using advanced analytics such as anomaly detection and behavioral pattern analysis. Accuracy in fraud detection was 98 percent, resulting in over 1 billion US dollars being protected against losses, as well as an estimated savings of 250 million US dollars per year due to fewer investigations.

4.3.2 Bank of America

The strategy used by Bank of America regarding artificial intelligence relied on Erica, which was perhaps the most obvious element. Launched in 2018 as an entry-level solution for transaction-based assistance, Erica had, by 2023, developed into a comprehensive advisory solution capable of managing over one million inquiries each day from over 42 million digital users. The amount of work it could manage daily equated to the work of over 3,000 employees. Hiring a staff of 3,000 employees would cost over a billion dollars annually in salaries.

Erica's development journey had many analytical merits due to the amount by which its progress was driven through the process of implementation and not just engineering. Its ability to comprehend queries improved from 85 percent upon launch to 95 percent in two years of active use. The progress was driven by the training signal provided by the millions of customer interactions that cannot be achieved within laboratory conditions. Her functional scope grew commensurately from checking balances and reviewing transaction histories to offering insights into finances, spending, credit scores, and product offerings.

The sector analysis conducted by Deloitte in 2023 revealed that organizations that used generative AI were able to automate about 41 percent of their back office processes as against 25 percent using the previous technology, robotic process automation. Similarly, the Bank of America also witnessed similar results in all its processes relating to fraud monitoring, compliance checks, and transaction processing. It had established an AI governance structure that took into account responsibility of AI with regular bias audits for credit-related algorithms, transparency of AI-driven decision-making for individuals, and human supervision of high impact decisions.

4.3.3 HDFC Bank

The perspective offered by HDFC Bank was something that the studies on JPMorgan Chase and Bank of America failed to provide. Being the biggest private sector bank of India, HDFC faced the challenge of functioning in a landscape comprising a huge customer base with many individuals who did not have any credit records, a wide variety of languages among its customer base of hundreds of millions of people, RBI regulations, and an online payments framework built around the infrastructure of Unified Payments Interface.

EVA, the virtual assistant of the bank, which was introduced back in 2017, provided HDFC Bank with seven years of practical experience at the time of conducting this study. With EVA having already handled over 2.7 million conversations with customers and an autonomy rate of over 85 percent, it was proven that conversational artificial intelligence could provide performance results in India, similar to that of artificial intelligence systems used in Western countries. It is especially relevant due to the high costs of running a call centre in the market.

The AI-based credit scoring system of HDFC bank helped solve one of the biggest problems that Indian banking has faced since long: that of giving formal credit facilities to borrowers who did not have enough formal credit history to be scored using traditional models. The bank used information from GST filings, utility bill payments, and even transactions through UPI to gauge the risk profile of the borrowers and price their loans accordingly. The process time for approving retail loan applications was reduced from several days to just a few hours or even minutes in cases of pre-approval by the bank.

With regard to the RBI's regulatory framework for the use of AI in banking, HDFC Bank was tasked with ensuring that they documented their model validation efforts as well as having controls for ensuring human involvement in the decision-making process for loans supported by AI technologies. This approach helped them prove that they could innovate responsibly through AI technology without compromising regulatory requirements.

4.4 Synthesis of Findings

By combining macro analysis, domain-level evidence, and institutional case studies, seven principal findings were identified.

Table 4.6: Consolidated Research Findings Matrix

No.	Finding Theme	Evidentiary Basis	Strategic Implication
F1	Critical Mass Achieved	78 percent of organisations using AI in at least one function by 2024; banking among top five investor sectors	AI adoption is structural, not a passing trend
F2	Persistent ROI Gap	Only 29 percent report cost savings from BCG 2024; 38 percent meet ROI targets from Deloitte 2024	Execution quality matters as much as technology selection
F3	Fraud Detection Leadership	AI accuracy of 92 to 98 percent; 9.6 billion USD projected annual savings from Caspian One 2025	Highest consistent ROI application in banking AI
F4	Talent as Binding Constraint	Domain specialists 80 percent faster in implementation from Goldman Sachs 2024; 14-month average delay from EY 2024	AI talent strategy is an existential priority
F5	Enterprise Versus Point-Solution Divergence	JPMorgan Chase and Bank of America enterprise-wide architecture consistently outperforms siloed deployments	Holistic strategy multiplies the value of individual tools
F6	Regulation as Adoption Driver	EU AI Act and AML and KYC requirements are driving RegTech investment according to Uptech Research 2024	Compliance pressure creates a clear AI investment case
F7	Generative AI as Inflection Point	58 percent of institutions attribute revenue growth to AI from McKinsey 2024; generative AI reshapes employee workflows	Gen AI shifts banking AI from narrow to general-purpose capability

Source: Compiled from authenticated secondary sources, 2023 to 2025

F2 discovery initially seemed to be at odds with the optimistic path suggested by discoveries F1, F3, and F7. The solution to the seeming contradiction lay in understanding the bimodal nature of AI penetration in the banking sector. Certain

companies like JPMorgan Chase, Bank of America, HSBC, and a few others have invested in strategic integration of AI and are generating significant profits from the investment. However, there are many more firms that have invested in piecemeal projects in AI and lack the required data, governance structure, and capabilities that can enable their transition to AI success. The statistics on 29 percent of respondents experiencing benefits in terms of cost savings can largely be attributed to this latter group.

F7 marked a turning point in the findings of the research in terms of structure. Before the development of the large language model technology in 2023, AI in banking could only be described as a set of applications, each operating independently in its own defined scope. However, the advent of generative AI provided a means for a general application of the technology: the same tool used for reviewing legal documents could also be utilised for communication purposes, regulatory assessment, information management, and memo writing. In 2024, McKinsey & Company revealed that 58 percent of financial firms directly credited AI investment for their increased revenue, with generative AI gaining prominence in this area.

4.5 Strategic Recommendations

Seven recommendations were derived from the evidence base assembled in this research. Each addressed one or more of the barriers identified in the findings and was grounded in specific documented institutional experience rather than in abstract principle.

Recommendation 1: Shift AI from Project Portfolio to Enterprise Infrastructure

There is clear distinction between organizations that derived value from AI implementation on a wide scale compared to those that have failed in doing so through the evidence. One clear distinction between such organizations is whether AI was viewed as a capability infrastructure for the organization as a whole, characterized by centralized data, platforms, and governance, or as individual projects implemented by different business units without coordination. It was found by McKinsey and Company

that more than 50 percent of financial institutions with considerable resources had centralized their AI infrastructure. Centralization of AI infrastructure has been linked with positive performance impact.

Recommendation 2: Resolve Data Infrastructure Before Deploying Advanced Models

The ability of the models to perform depended on the availability and quality of the available data. Organisations that invested in building complex architectures on top of poor-quality, fragmented and inconsistent data always delivered lower results compared to organisations that built simpler models on top of good quality, integrated data environments. The importance of investing in data infrastructure transformation to include integrated customer data environments, data definitions standardisation, cross-system data integration and data quality improvements had to be made at this stage as a prerequisite to deploying AI capabilities.

Recommendation 3: Prioritise Domain-Specialised AI Talent

AI skills in general were essential, but they were not enough for an efficient application of AI within the banking industry. The study carried out by Goldman Sachs in 2024 proved that an AI specialist with deep knowledge about the banking industry would be able to complete his or her work almost twice as fast as a generalist who had the same technical skills, but who lacked financial expertise. The 14 months' implementation time, which was mentioned in 2024 by Ernst and Young, could be mainly explained by the lack of people combining machine learning skills with banking knowledge.

Recommendation 4: Sequence AI Deployment from Proven to Complex Applications

Based on evidence collected, the following applications proved themselves capable of delivering consistent value add regardless of the context: fraud detection, process automation in documents handling, and automation of KYC and AML compliance were all characterised by advanced technological capabilities, reliable performance

indicators, robust alignment with regulations, and ease of replication. These should be the focus areas for an organisation when embarking on its early AI journey, as they will help gain traction, establish necessary skills, and secure trust from stakeholders that allows further commitment down the line. Such a strategy helped organisations accumulate knowledge about AI and experience that made investments into advisory and credit assessment possible.

Recommendation 5: Implant Explainability and Fairness from the Start

Responsible AI could not have been a compliance process that would need to be addressed after deploying AI solutions via audits. Responsible AI had to be an architectural principle that would have to be considered and designed into AI solutions right from the outset. Tools like SHAP values and LIME ought to have formed part of credit scoring software right from their design process and not just bolt-on features. Testing biases among protected classes had to form a mandatory feature of all models used by AI to determine customer eligibility for financial products.

Recommendation 6: Engage Proactively with Emerging Regulatory Frameworks

However, the regulatory framework for AI use in banks was developing quickly in various banking jurisdictions. The labeling of credit scoring, fraud detection, and insurance underwriting as high-risk AI activities by the EU AI Act created demands for transparency, accountability, and documentation that many existing applications had failed to meet. Rather than viewing such demands as hindrances that were to be minimized, it would behoove institutions to consider them quality assurances that would enhance the reliability of AI systems and protect customers.

Recommendation 7: Build AI Literacy Programmes Across All Organisational Levels

With the launch of the JPMorgan Chase LLM Suite, 200,000 employees went from having no active use of AI to using AI daily within eight months. This shows that organisational AI literacy is possible if there is true alignment between commitment

at an executive level, availability of appropriate tools, and practical training. For the programme to be effective, it would require different levels of literacy depending on the roles within the organisation. The operations staff would need to have enough literacy so they could utilise AI outputs while being able to understand its limitations. The middle management would need to be literate in such a way that they could evaluate AI investment cases and manage AI-augmented teams.

4.6 Limitations of the Study

Four limitations of this research design require honest acknowledgement to allow readers to calibrate appropriately the weight they place on the findings and recommendations.

First came source optimism bias. It is possible that consulting firms profiting from advisory fees on AI projects would have structural biases towards representing positive AI implementation results. Given that only 29% of organizations cited cost reductions from their initiatives as being significant, according to BCG data, this was an important exception to the rule. However, it is prudent for the reader to recognize industry bias when reading the material.

The second bias was temporal lag. The speed at which AI was developing in the banking sector meant that literature written in 2024 could have been outdated by the time this study was conducted. To address this issue, the study relied more on directional findings as opposed to using numerical information that could easily become outdated. Thirdly, institutional secrecy was another problem. Financial institutions tended to share their success stories while hiding their shortcomings. Thus, analysis based on case studies was likely to paint a more rosy picture of the implementation process as opposed to internal evaluations of an institution's processes. Lastly, geographic skewness was another challenge since there was considerably less information available on banking firms from Africa, Southeast Asia, and Latin America compared to North America and Europe.

CHAPTER 5: CONCLUSION

The present study was undertaken with the intention of studying the impact of AI on the industry of banking and financial services, and it was indeed found that the transformation being undergone by this sector was very real, and very structural. In each and every one of the six areas studied, AI technology has consistently provided tangible improvements as compared to the previous generation technologies. Better detection rates, better decision-making, and greater efficiencies were just some of the outcomes obtained through the use of AI technology.

The statistics around the market highlighted the magnitude of what was happening. A sector that had grown from 19.9 billion US dollars to a forecast of 315.50 billion US dollars in 2033 at a compound annual growth rate of 31.83 percent clearly showed how far removed the technology was from a mere concept. Major financial institutions had made billions annually from AI technologies. Yet, the most important type of AI capability to be unveiled during the research period, generative AI and LLMs, had not fully been embraced by businesses as at the time of this writing.

The ROI gap finding required special attention in any discussion of conclusions since it disrupted an appealing storyline of AI's inherent benefits and its automatic worth for companies that implemented them. This story was contradicted by the facts. Not even 30 percent of financial organizations found that AI helped them achieve some cost-saving results. Less than 40 percent of AI initiatives achieved their ROI goals. This did not look like a minor problem. In fact, it highlighted a certain systemic issue with implementation that was not linked to AI's capabilities at all but rather to the quality of company infrastructure, management skills, and employee expertise.

For the three companies mentioned above, we got the other side of the story as well. We learned how possible it was for the companies like JPMorgan Chase, Bank of America, and HDFC Bank to realize possibilities when they view artificial intelligence as an enterprise-wide strategic initiative instead of a series of technology experiments. The realization of \$1.5 billion of AI value per year by JPMorgan, use of Erica

technology in Bank of America to deal with a million customers daily, and using credit scoring technology to extend loans to customers that cannot be considered otherwise.

This diversity in outcomes raised concerns about implications extending well beyond individual banks' approaches to the adoption of AI. In the face of unevenly distributed gains for the application of artificial intelligence in finance in developed countries, the possibility arose that such gains could exacerbate imbalances in terms of the accessibility of financial services, as well as competition therein. The HDFC Bank case indicated that this scenario was not preordained but rather the product of certain choices. Artificial intelligence, carefully tailored to local conditions, could make a positive difference even in developing countries.

Indeed, as of the time of this research, next-generation AI-driven systems able to achieve multiple objectives on their own already existed in prototype form. They possessed an even greater transformative potential in that they would enable process automation and full-fledged end-to-end loan origination and comprehensive personal financial planning. In addition, such systems could be used for automated compliance monitoring. According to McKinsey's Global Banking Annual Review in 2025, AI had a thirty percent chance of transforming banking operations and the financial practices of consumers in a matter of ten years' time. Whether this transformation would work for the best of everyone – the bank clients, workers, local communities, and institutions – depended on the decision-making process taking place right here and now.

