

ARTIFICIAL INTELLIGENCE ADOPTION AND ITS IMPACT ON FINANCIAL PERFORMANCE AND OPERATIONAL EFFICIENCY: EMPIRICAL EVIDENCE FROM INDIAN BANKING (2014–2023)

GIDEON G

Assistant Professor, Department of MBA, Rao Bahadur Y Mahabaleswarappa Engineering College Ballari, Visvesvaraya Technological University. Email: gideongideon91@gmail.com

Dr. K T GOPI

Professor, Rao Bahadur Y Mahabalewarappa Engineering College Bellari, Visvesvaraya Technological University. Email: gopi.vim@gmail.com

Abstract

This study examines the association between Artificial Intelligence (AI) adoption and the financial performance of Indian commercial banks over a ten-year period (2014–2023), using a panel of 20 banks (10 classified as AI-adopting, 10 as minimally AI-adopting). Because AI-adoption status does not change within any bank over the sample period, a standard fixed-effects estimator cannot identify its coefficient — the treatment is time-invariant and gets absorbed into the bank-specific intercept. The analysis instead uses pooled OLS and random-effects models with standard errors clustered by bank, cross-checked against Welch's t-tests. AI-adopting banks show significantly higher Return on Assets (+0.50 percentage points) and Return on Equity (+5.63 percentage points) and a significantly lower Cost-to-Income Ratio (–11.32 percentage points) than non-adopting banks (all $p < 0.001$). A bootstrapped mediation test does not find a statistically significant indirect effect of AI adoption on profitability through cost efficiency, indicating that the profitability association operates at least partly through channels this dataset does not capture. The design cannot rule out reverse causality — that financially stronger or better-managed banks are simply more likely to adopt AI — and this is discussed as the study's central limitation rather than resolved by the available data. The contribution is a corrected identification strategy for a time-invariant adoption variable in bank-level AI research, together with an explicit account of what a twenty-bank panel can and cannot support.

Keywords: Artificial Intelligence in Banking; Financial Performance; Operational Efficiency; India; Digital Transformation; Regulatory Environment.

1. INTRODUCTION

Artificial Intelligence (AI) has emerged as one of the most transformative technologies in the global financial services industry, fundamentally reshaping the way banking institutions operate, compete, and deliver value to customers. AI technologies—including machine learning, natural language processing, robotic process automation (RPA), predictive analytics, and intelligent chatbots—have enabled banks to automate routine processes, improve decision-making, enhance fraud detection, optimize credit risk assessment, and deliver personalized customer experiences. As digital transformation accelerates worldwide, AI has evolved from a technological innovation into a strategic capability that significantly influences operational efficiency, financial performance, and long-term competitiveness.

The banking industry has experienced rapid digital transformation over the past decade, driven by increasing customer expectations, technological advancements, regulatory

reforms, and growing competition from financial technology (FinTech) firms. In response to these changes, banks across developed and emerging economies have substantially increased investments in AI-driven solutions to improve operational resilience, reduce transaction costs, strengthen risk management, and enhance service quality.

Previous international studies have demonstrated that AI adoption positively influences profitability, operational efficiency, customer satisfaction, and organizational performance. Nevertheless, the magnitude of these benefits varies across countries due to differences in technological infrastructure, regulatory environments, institutional readiness, and digital maturity.

India represents one of the world's fastest-growing digital banking ecosystems. Government initiatives such as **Digital India**, the **Pradhan Mantri Jan Dhan Yojana (PMJDY)**, the **Unified Payments Interface (UPI)**, and continuous regulatory support from the **Reserve Bank of India (RBI)** have accelerated digital financial inclusion and encouraged banks to adopt emerging technologies.

Public and private sector banks have increasingly implemented AI applications in customer relationship management, fraud detection, loan underwriting, anti-money laundering systems, personalized financial advisory services, predictive maintenance, and process automation. These developments have significantly transformed banking operations and strengthened India's position as a leading digital economy.

Despite the growing adoption of AI within Indian banking, empirical evidence regarding its financial and operational impact remains limited. Existing studies have predominantly focused on customer adoption of AI-enabled services, digital banking acceptance, chatbot implementation, or conceptual discussions of AI applications.

Comparatively fewer studies have quantitatively evaluate whether AI adoption contributes to measurable improvements in financial performance indicators such as Return on Assets (ROA), Return on Equity (ROE), and Cost-to-Income Ratio (CIR). Furthermore, limited research has systematically compared the performance of AI-powered banks with that of non-AI banks using longitudinal panel data and advanced econometric techniques. This gap restricts a comprehensive understanding of whether AI investments translate into sustainable financial advantages within the Indian banking sector.

To address these gaps, this study investigates the relationship between AI adoption and financial performance using panel data from twenty Indian commercial banks — ten classified as AI-adopting and ten as minimally AI-adopting — over 2014–2023. Because adoption status is fixed for each bank across the sample window, the analysis uses pooled OLS and random-effects estimation with bank-clustered standard errors, rather than the fixed-effects framework that panel studies typically default to, and tests whether operational efficiency (proxied by the Cost-to-Income Ratio) mediates the relationship between AI adoption and profitability using a bootstrapped mediation model.

The study contributes to the existing literature in three ways. First, it applies an identification strategy suited to a time-invariant adoption variable — a design issue that

appears to have been overlooked in prior India-focused work using panel data. Second, it tests, rather than assumes, whether operational efficiency mediates the AI–performance relationship, and reports the result even where it does not confirm the hypothesised pathway. Third, it is explicit about what a twenty-bank, cross-sectional-adoption panel can support empirically — group comparisons and association, not causal identification — and treats that as a limitation to be stated rather than a caveat to be minimised.

2. LITERATURE REVIEW

Artificial Intelligence (AI) adoption in banking has been studied extensively at the international level, with less causal, India-specific evidence. International studies generally report that AI adoption is associated with higher profitability, greater operational efficiency, and stronger customer engagement, but the mechanisms and magnitude vary enough across markets that findings should not be assumed to transfer directly to the Indian context.

Bansal, Pandey, Goel, and Sharma (2024) examined the barriers Indian banks face in adopting AI-driven chatbots, using an interpretive structural modelling approach, and found that organisational readiness and staff capability — not just the availability of the technology — were the binding constraints on adoption. This is consistent with the pattern in the present dataset, where AI-adopting and non-adopting banks differ systematically on more than one dimension (Section 5), making adoption itself something that likely correlates with unobserved bank characteristics rather than being randomly assigned.

Deepthi, Gupta, Rai, and Arora (2022), writing in *Vision*, assessed the broader dynamics of AI-driven technologies across Indian banks and reported that large private banks have moved fastest on AI adoption, aided by stronger funding and technology infrastructure, while public-sector and smaller banks lag — a pattern that matches the composition of the AI-adopting group in this study's sample (SBI, HDFC, ICICI, and other large private and public banks).

Agarwal, Swami, and Malhotra (2022) reviewed AI adoption drivers in the post-COVID-19 environment and argued that adoption decisions are shaped as much by organisational and leadership factors as by cost pressure or competitive threat — a finding with direct implications for the endogeneity concern raised in Section 4: if adoption is driven by leadership quality or organisational capability that also affects performance, an AI–performance association could partly reflect that underlying capability rather than AI itself.

Sheth, Jain, Roy, and Chakraborty (2022), in the *International Journal of Bank Marketing*, examined AI-driven banking services in emerging markets and found that AI applications in customer-facing functions are associated with improved personalisation and service delivery, though their study — like most of this literature — relies on cross-sectional or perception-based data rather than longitudinal financial-performance panels.


On the financial-performance side specifically, Bandyopadhyay (2022) used a dynamic panel data analysis of public-sector banks in India to study the relationship between capital and financial performance, illustrating both the econometric approach (dynamic panel GMM) that a fully causal AI-adoption study would need, and the extent to which India-specific bank-performance research has so far focused on capital adequacy and risk rather than on technology adoption as the explanatory variable of interest. This is the gap the present study addresses, while being explicit that its own design — static rather than dynamic, and lacking a valid instrument for AI adoption — falls short of establishing causality.

These findings are consistent with three established theoretical lenses: the Resource-Based View (Barney, 1991), under which AI functions as a firm-specific capability that is costly to imitate and can therefore support a performance advantage; Diffusion of Innovation Theory (Rogers, 2003), which explains why adoption spreads unevenly across banks with different resources and risk appetites; and Institutional Theory (DiMaggio & Powell, 1983), which accounts for the role of regulatory bodies such as the RBI in shaping which technologies banks adopt and how quickly. Two tensions remain unresolved in this literature and motivate the analysis that follows: whether AI's profitability effect operates directly or is mediated by cost efficiency, and whether the association reflects AI's causal effect or the tendency of stronger banks to adopt AI in the first place.

3. THEORETICAL FRAMEWORK

The conceptual framework for this study builds on the Resource-Based View (RBV) and Diffusion of Innovation Theory. AI adoption is treated as a strategic resource that enhances competitiveness by improving cost efficiency and profitability. The model proposes that AI adoption impacts bank performance both directly and indirectly, through mediating variables such as operational cost reduction and efficiency improvements (CIR, NPA). This framework also integrates Institutional Theory, acknowledging that regulatory support and organizational culture moderate adoption outcomes.

3.1 Conceptual Framework Diagram:



3.2 Control Variables

Bank Size – Measured as the logarithm of total assets, this controls for the scale effect in bank performance. Larger banks often have economies of scale, more diversified portfolios, and greater capacity to invest in technology, which can influence ROA, ROE, and CIR independently of AI adoption.

Capital Adequacy Ratio (CAR) – Expressed as a percentage, CAR measures a bank's capital relative to its risk-weighted assets. It ensures financial stability and resilience against credit risk. A higher CAR can improve investor confidence and support higher profitability, potentially confounding the effect of AI adoption on performance.

3.3 Moderating Factors

Institutional and Regulatory Environment – This includes the influence of central bank policies (e.g., RBI guidelines), financial regulations, and compliance requirements. A supportive regulatory framework can accelerate AI adoption by providing standards, incentives, and risk management guidelines. Conversely, restrictive or unclear regulations can slow adoption and limit performance impact.

3.4 Research Objectives:

Objective 1: To study the impact of AI on the performance of the banking and financial industry

Objective 2: To measure the reduction of operational costs due to AI implementation

Objective 3: To study the financial performance of AI-powered banking institutions versus non-AI banks in India

3.5 Research Hypotheses

- **H1:** AI adoption positively influences ROA.
- **H2:** AI adoption positively influences ROE.
- **H3:** AI adoption negatively influences CIR (improves efficiency).
- **H4:** Operational efficiency mediates the relationship between AI adoption and bank performance.
- **H5:** Bank size and CAR act as control variables, moderating the effect of AI adoption.
- **H6:** Institutional and regulatory factors strengthen the relationship between AI adoption and financial outcomes.

This visual illustration emphasizes the causal path: AI adoption drives operational efficiency, which in turn enhances financial performance, with moderating and control variables shaping the intensity of outcomes.

This framework provides the basis for econometric estimation and Structural Equation Modeling (SEM), ensuring both direct and mediated effects are captured.

4. METHODOLOGY

4.1 Sampling Design: The sample includes 20 Indian commercial banks, divided into 10 AI-intensive (e.g., SBI, HDFC, ICICI) and 10 with minimal AI adoption. Classification was conducted using content analysis of disclosures in annual reports and regulatory filings.

4.2 Data Sources and Variables

The dataset spans FY 2014–2023. Data sources include bank annual reports, RBI bulletins, investor presentations, and secondary portals such as Screener and Moneycontrol.

4.3 Variable Definition Table

Variable	Definition	Proxy/Unit	Source
ROA	Return on Assets	%	Annual Reports
ROE	Return on Equity	%	Annual Reports
CIR	Cost-to-Income Ratio	%	Annual Reports
AI Score	Composite AI index (0–10)	Score	Content analysis
AI Dummy	Binary adoption	1 or 0	Classification
Size	Bank size	Log Assets	Annual Reports
NPA	Non-Performing Asset ratio	%	Annual Reports
CAR	Capital Adequacy Ratio	%	Annual Reports

Explanation: The table defines the variables specified in the research framework. The estimation sample (Objective4_Performance_Comparison_WithFormulas.xlsx, $n = 200$) contains ROA, ROE, Net Profit Margin, CIR, NPA, and a binary AI-adoption flag (AI Dummy) for each bank-year.

Bank Size, CAR, and a continuous 0–10 AI Adoption Score are part of the conceptual framework but were not present in the analysed dataset and are not included in the regressions reported in Section 5; they are listed here as variables for future data collection (Section 6.2).

4.4 Preliminary Tests

- Shapiro-Wilk Test (real output, $n = 200$): ROA — $W = 0.977$, $p = 0.002$; ROE — $W = 0.975$, $p = 0.001$; CIR — $W = 0.960$, $p < 0.001$. All three variables depart significantly from normality, so inference relies on cluster-robust standard errors and Welch's t-test (Section 5.2), which do not assume normally distributed residuals, rather than on tests that do.
- Variance Check: No implausible outliers were found on visual inspection of ROA, ROE, and CIR; skewness is moderate (0.35, 0.05, and 0.36 respectively) and does not indicate the presence of extreme values driving the results.

Explanation: These checks confirm the data are usable for pooled and random-effects panel regression, but the non-normality result is a reason to rely on robust/clustered standard errors throughout rather than an incidental detail.

4.5 Econometric Model:

Performance indicators (ROA, ROE, CIR) are regressed on AI-adoption status and available control variables (NPA). AI-adoption status does not vary within a bank over 2014–2023, so a fixed-effects (within) estimator cannot identify its coefficient — attempting to fit one throws an `AbsorbingEffectError`, since the entire AI-adoption effect is absorbed into the bank fixed effects.

Pooled OLS and random-effects models are used instead, with standard errors clustered by bank to account for within-bank serial correlation; the two estimators are cross-checked against each other rather than compared via a Hausman test, since Hausman's logic (comparing an FE and RE estimate of the same coefficient) does not apply when FE cannot estimate that coefficient at all.

4.6 Diagnostic Tests Summary

4.6.1 Why a Hausman Test Is Not Reported

Purpose (standard use): The Hausman test ordinarily compares Fixed Effects (FE) and Random Effects (RE) estimates of the same coefficient to decide which is more appropriate, on the logic that both are consistent under the null of no correlation between the effects and the regressors, but only RE is efficient.

Result in this dataset: AI-adoption status is fixed for every bank across all ten years, so the within-bank (FE) transformation removes all variation in the variable of interest and FE cannot estimate its coefficient (confirmed empirically: `linearmodels.PanelOLS` raises an `AbsorbingEffectError` when entity effects are included). A Hausman comparison requires both estimators to produce an estimate of the same coefficient, which is not possible here.

Interpretation: This is reported as a specification decision, not a hidden limitation — pooled OLS and random-effects estimates are used instead (Section 4.5), and their close agreement (coefficients match to three decimal places, Section 5.3) is offered as the relevant robustness check in place of a Hausman test.

4.6.2 Variance Inflation Factor (VIF)

Purpose: To detect multicollinearity among independent variables.

Real output (pooled model, AI Dummy and NPA): $VIF(\text{AI Dummy}) = 3.80$; $VIF(\text{NPA}) = 3.80$.

These values are below the conventional threshold of 10 (and below the more conservative threshold of 5), so multicollinearity does not invalidate the coefficient estimates.

They are, however, well above 1, reflecting genuine correlation between AI adoption and asset quality in this sample — AI-adopting banks also have substantially lower NPA ratios (2.44% vs. 7.02%, Section 5.1) — which is discussed as a substantive finding relevant to the endogeneity question, not just a diagnostic footnote.

Interpretation: Coefficients are stable, but the AI–NPA correlation means NPA cannot be treated as a clean, independent control; see Section 5.4 for the implication.

4.6.3 Breusch–Pagan Test

Purpose: To test for heteroskedasticity in the residuals of the regression model.

Real output (pooled OLS residuals, $n = 200$):

ROA model: $\chi^2 = 10.89$, $p = 0.004$ (heteroskedasticity present). ROE model: $\chi^2 = 0.34$, $p = 0.845$ (no evidence of heteroskedasticity). CIR model: $\chi^2 = 26.90$, $p < 0.001$ (heteroskedasticity present).

Interpretation: Heteroskedasticity is present in the ROA and CIR models, which is why all coefficient tables in Section 5.3 report standard errors clustered by bank rather than ordinary OLS standard errors — clustering is robust to this kind of heteroskedasticity as well as to within-bank serial correlation.

Overall Interpretation

The diagnostics collectively indicate:


- Identification: AI-adoption status is time-invariant per bank, so pooled OLS / random effects — not fixed effects — is the applicable estimator, and a Hausman test does not apply.
- VIF: No destabilising multicollinearity ($VIF = 3.80$), though AI adoption and NPA are meaningfully correlated, which matters for interpretation, not just estimation.
- Breusch–Pagan: Heteroskedasticity is present in two of three models, addressed throughout via bank-clustered standard errors.

Conclusion: The regression results reported in Section 5.3 come from an estimator that is actually identifiable given this data's structure, with standard errors appropriate to the heteroskedasticity and clustering the diagnostics reveal.

Test	Result
VIF (max)	3.80 (< 10 threshold)
Breusch-Pagan (ROA)	$\chi^2=10.89$, $p=0.004$ (heteroskedastic)
Breusch-Pagan (CIR)	$\chi^2=26.90$, $p<0.001$ (heteroskedastic)

Explanation: This diagnostic sequence replaces the fixed-effects/Hausman framework used in earlier drafts, which could not have produced valid estimates given the time-invariant treatment variable.

The table below reports the real VIF and Breusch–Pagan values against conventional thresholds; a Hausman row is omitted for the reasons given in Section 4.6.1.



4.7 Software and Tools Used

- 4.7.1 Python (pandas, statsmodels, linearmodels) — Panel Regression and Diagnostics
- Data cleaning, pooled OLS and random-effects estimation, VIF, Shapiro-Wilk, and Breusch–Pagan tests were run in Python (pandas, statsmodels, linearmodels, scipy). Standard errors are clustered by bank throughout.

4.7.2 Regression-Based Mediation (replaces the Panel VAR / EViews subsection)

- A Panel VAR / Granger-causality analysis of AI adoption's dynamic effect on performance is not identifiable with this dataset for the same reason a fixed-effects model is not: AI-adoption status never changes within a bank, so there is no within-bank timing of adoption for a VAR to exploit. Rather than retain a Panel VAR analysis that could not have been estimated from this data, the dynamic question is set aside as a direction for future research requiring banks that adopted AI partway through the sample window (Section 6.3).

4.7.3 Mediation Analysis (replaces the SmartPLS / SEM subsection)

- The data contain single observed indicators for AI adoption and operational efficiency (CIR), not the multi-item scales a latent-variable SEM measurement model requires, so no CFA, factor loadings, AVE, or composite reliability could be estimated — and the RMSEA/CFI/GFI figures in earlier drafts did not come from an estimated model. In their place, a regression-based (Baron–Kenney style) mediation test with bootstrapped standard errors was run on the observed variables: AI adoption → CIR (path a), then CIR and AI adoption → ROA/ROE (paths b and c'). Results are reported in Section 4.9.

4.7.4 Microsoft Excel — Data Compilation

- Excel was used to compile and structure the source data (ROA, ROE, Net Profit Margin, CIR, NPA, and AI-adoption status for 20 banks × 10 years) before analysis in Python.

4.8 Dynamic Effects: Not Estimated

As noted in Section 4.7.2, Panel VAR / Granger-causality estimates of AI adoption's effect over time require within-bank variation in adoption timing, which this dataset does not have (every bank's AI-adoption status is fixed across 2014–2023). No Granger-causality or lag-structure results are reported. Section 6.3 outlines what a future dataset would need to support this analysis.

4.9 Mediation Analysis (real output, bootstrapped, 2,000 resamples clustered by bank)

- Path a (AI adoption → CIR): -12.32 , $p < 0.001$ — AI-adopting banks have significantly lower CIR.
- Path b (CIR → ROA, controlling for AI adoption): ≈ 0.00 , not significant. Path b (CIR → ROE, controlling for AI adoption): -0.066 , not significant.
- Indirect effect ($a \times b$): ROA — 0.003 , 95% bootstrap CI ($-0.137, 0.128$), includes zero. ROE — 0.808 , 95% bootstrap CI ($-0.059, 1.606$), includes zero.

Explanation: The indirect (mediated) effect of AI adoption on ROA and ROE through CIR is not statistically distinguishable from zero at conventional levels, even though AI adoption significantly affects both CIR and profitability on their own.

In Baron–Kenny terms, this pattern (significant a-path, non-significant $a \times b$ indirect effect) does not support H4 as originally stated — operational efficiency, at least as proxied by CIR alone, does not statistically explain the link between AI adoption and profitability in this sample.

This is reported as a genuine finding rather than smoothed into a confirmed hypothesis: the direct association between AI adoption and ROA/ROE (Section 5.3) is likely carried by channels this single-mediator model does not capture — for example revenue effects (cross-selling, fee income) that CIR does not measure.

5. RESULTS AND DISCUSSION

5.1 Descriptive Statistics

Metric	AI Mean	Non-AI Mean
ROA (%)	1.23	0.90
ROE (%)	15.07	9.80
CIR (%)	43.00	55.31

Explanation: Real sample means ($n = 200$ bank-years) show AI-adopting banks outperforming non-adopting banks on profitability and cost efficiency. AI-adopting banks

average NPA of 2.44% versus 7.02% for non-adopting banks — a gap comparable in relative size to the CIR gap, and one reason NPA does not survive as an independent control once AI adoption is included (Section 4.6.2).

These are simple group means; the inferential tests below (Sections 5.2–5.3) establish whether the differences are statistically robust, and Section 5.4 discusses why they should not be read as proof of a causal AI effect.

5.2 Group Comparison (Welch's t-test, unequal variances — appropriate given the non-normality found in Section 4.4)

Metric	t-Statistic	p-Value	Result
ROA	6.49	7.4×10^{-10}	Significant
ROE	12.74	1.6×10^{-27}	Significant
CIR	-17.51	5.9×10^{-40}	Significant

Explanation: Differences in ROA, ROE, and CIR between AI-adopting and non-adopting banks are statistically significant at any conventional threshold. These t-statistics confirm the group difference is real and large relative to within-group variation; they do not, on their own, establish that AI adoption caused the difference — that question is addressed in Section 5.4 alongside the regression results.

5.3 Pooled OLS / Random Effects Regression Results (cluster-robust SE by bank; real output, n = 200)

Variable	ROA Coefficient	ROE Coefficient	CIR Coefficient
AI Adoption	+0.498***	+5.627***	-11.321***
Size (log assets)	N/A — not in dataset	N/A — not in dataset	N/A — not in dataset
NPA	+0.036**	+0.079 (n.s.)	+0.217 (n.s.)
CAR	N/A — not in dataset	N/A — not in dataset	N/A — not in dataset

R-squared: ROA = 0.190, ROE = 0.451, CIR = 0.609. (Notes: *** $p < 0.001$, ** $p < 0.01$, n.s. = not significant at $p < 0.05$. Bank size and CAR are listed for completeness against the framework in Section 3.2 but were not available in the analysed dataset; see Section 6.2.)

The table reports pooled OLS / random-effects coefficients (Section 4.5) for three dependent variables — ROA, ROE, and CIR — for the two explanatory variables actually available in the dataset: AI adoption and NPA.

1) AI Adoption

- ROA: Coefficient +0.498, $p < 0.001$ — AI-adopting banks have significantly higher returns on assets.
- ROE: Coefficient +5.627, $p < 0.001$ — AI-adopting banks achieve substantially higher shareholder returns.
- CIR: Coefficient -11.321, $p < 0.001$ — AI-adopting banks operate at a materially lower cost-to-income ratio.

2) NPA

- ROA: Coefficient +0.036, $p = 0.004$ — significant, but the sign is opposite to what a simple risk story would predict; this likely reflects the NPA–AI-adoption correlation noted in Section 4.6.2 ($VIF \approx 3.8$) rather than a stable, independent NPA effect, and should not be over-interpreted.
- ROE: Coefficient +0.079, not significant ($p = 0.493$).
- CIR: Coefficient +0.217, not significant ($p = 0.418$).

Key Insights

- The largest, most consistent effect is AI adoption on CIR, and it is the effect most plausibly tied to a concrete mechanism (RPA and automated processing reducing operating cost) rather than to unobserved bank quality alone.
- AI adoption's effect on ROA and ROE is significant but, per the mediation test in Section 4.9, not statistically explained by the CIR channel — the profitability association likely reflects additional factors (revenue effects, or unobserved differences between adopting and non-adopting banks) that this model does not isolate.
- NPA loses independent explanatory power once AI adoption is in the model, consistent with AI-adopting banks also being the lower-risk banks in this sample — a pattern this cross-sectional design cannot resolve into cause and effect (Section 5.4).
- Bank size and CAR could not be tested; both are plausible confounders and their omission is a limitation (Section 6.2), not a gap the current results can speak to.


5.4 Segment-Wise Context (from prior literature, not directly tested in this dataset)

- **Lending:** prior studies (e.g., Bansal et al., 2024) associate predictive analytics with improved credit scoring.
- **Payments:** AI-driven chatbots are widely reported to reduce service-delivery cost (Sheth et al., 2022).
- **Insurance:** claim-automation applications of AI are reported elsewhere to shorten turnaround times.


Explanation: These segment-level patterns are drawn from the literature discussed in Section 2, not estimated from this study's bank-level panel, which does not disaggregate by business line. They are included as context for interpreting the aggregate CIR effect (Section 5.3), not as findings of this paper.

5.5 Trend Graphs


• 5.5.1 ROA Trend (2014–2023): AI vs. Non-AI banks




• 5.5.2 ROE Trend (2014–2023): AI vs. Non-AI banks



- 5.5.3 CIR Trend (2014–2023): AI vs. Non-AI banks



- 5.5.4 Operational Cost Trend (2014–2023): AI vs. Non-AI banks



Explanation: The trend charts show the same AI-adopting vs. non-adopting gap documented in Sections 5.1–5.3, widening over the sample period for ROA, ROE, and CIR. Because AI-adoption status does not change within a bank in this sample, these trends cannot distinguish "AI causing a widening gap" from "already-diverging banks

having different adoption propensities" — the same identification limitation that applies to the regression results, restated here because the visual pattern is the one most likely to be over-read as causal.

5.6 Diagnostic Test Summary

The diagnostics in Section 4.6 support three conclusions. First, AI-adoption status is time-invariant per bank, so pooled/random-effects estimation — not fixed effects — is the correct and, in fact, only identifiable approach for this design. Second, VIF values (3.80) indicate no destabilising multicollinearity, though AI adoption and NPA are correlated enough to matter for interpretation. Third, Breusch–Pagan tests found heteroskedasticity in the ROA and CIR models, addressed throughout via bank-clustered standard errors. Collectively, these diagnostics support the coefficient estimates in Section 5.3 as correctly specified given the data, while the mediation test in Section 4.9 and the endogeneity discussion below (Section 5.7) are the reasons the results are presented as an association rather than a demonstrated causal effect.

5.7 Endogeneity and Robustness

The central threat to a causal reading of these results is reverse causality: banks that are already better-managed, better-capitalised, or lower-risk plausibly both perform better and adopt AI sooner, which would produce the same association observed here even if AI itself had no effect. Two pieces of evidence in this dataset are consistent with that concern — the correlation between AI adoption and NPA ($VIF \approx 3.8$, Section 4.6.2) and the fact that the mediation test (Section 4.9) does not find the hypothesised operational-efficiency channel, suggesting the AI–profitability association is not fully explained by the mechanism the theoretical framework proposes. Addressing this properly would require either an instrument for AI adoption uncorrelated with unobserved bank quality, or panel variation in adoption timing (banks that adopted AI partway through the sample), neither of which is available here. What robustness checks are possible with this data — clustering standard errors by bank, cross-checking pooled OLS against random effects, and reporting the non-normality and heteroskedasticity diagnostics rather than assuming them away — have been applied throughout. The honest summary is that this study establishes a robust association, not a demonstrated causal effect, and Section 6.2 states that limitation directly.

6. CONCLUSION

This study examined the association between Artificial Intelligence (AI) adoption and the financial performance of 20 Indian banks (10 AI-adopting and 10 non-adopting) over 2014–2023, using pooled OLS and random-effects panel regression with bank-clustered standard errors — an estimator chosen because AI-adoption status does not vary within any bank in the sample, which rules out fixed-effects estimation. AI-adopting banks show significantly higher ROA and ROE and a significantly lower CIR than non-adopting banks. A bootstrapped mediation test does not find that this profitability association is statistically explained by operational efficiency (CIR) alone, indicating the relationship likely operates

through additional channels not captured in this dataset. Because adoption status is not randomly assigned and the design has no instrument or adoption-timing variation to exploit, the results should be read as a robust association documented with an identification strategy appropriate to the data, not as evidence of a causal AI effect.

6.1 Policy and Practice Implications (interpreted cautiously, given Section 5.7)

- 1) **Regulatory Support for AI Integration** – The Reserve Bank of India (RBI) should develop standardized AI adoption and disclosure frameworks, enabling benchmarking and facilitating informed investment decisions.
- 2) **Targeted Digital Investment Schemes** – Public sector banks, which often lag in AI maturity, should receive targeted funding or incentives to deploy AI in high-impact areas such as fraud detection, credit scoring, and process automation.
- 3) **Workforce Reskilling Programs** – Industry associations and training institutes should implement large-scale reskilling initiatives to ensure bank employees can effectively work alongside AI systems.
- 4) **Encouraging AI for Financial Inclusion** – Policy measures should promote AI deployment in rural banking operations, microfinance, and cooperative banks to expand access to credit and financial services.

6.2 Limitations of the study

- **AI Adoption Measurement** – The dataset classifies banks with a binary AI-adopter flag; the 0–10 composite AI Adoption Score described in the original framework (Section 3, Table 4.3) was not available and is not what the regressions in Section 5 use. A validated composite score, with a documented coding protocol and inter-rater reliability (Cronbach's alpha / Cohen's kappa), remains a task for future data collection.
- **Sample Size and Time-Invariant Treatment** – 20 banks / 200 bank-years support the regressions and t-tests reported here but are underpowered for SEM with latent constructs or dynamic panel GMM. More fundamentally, no bank in the sample changes AI-adoption status across 2014–2023, which is why fixed-effects estimation is not identifiable (Section 4.5) and why the dynamic (Panel VAR) analysis planned in earlier drafts could not be estimated (Section 4.8).
- **Missing Controls** – Bank size (log assets) and Capital Adequacy Ratio are part of the theoretical framework (Section 3.2) but were not present in the analysed dataset; only NPA was available as a control. Both are plausible confounders of the AI–performance relationship and their omission should be treated as a real gap, not a minor one.
- **Endogeneity** – Reverse causality (financially stronger or better-managed banks self-selecting into AI adoption) cannot be ruled out with this cross-section; Section 5.7 discusses this directly rather than assuming the causal direction implied by the original hypotheses. Addressing it would require an instrument for AI adoption or panel variation in adoption timing, neither of which is available here.

- **Measurement Model** – No multi-item indicators are available for AI adoption or operational efficiency, so a latent-variable SEM with measurement-model diagnostics (loadings, AVE, CR, HTMT) could not be estimated; Section 4.9 substitutes an observed-variable mediation test, which is a narrower but honestly estimable alternative.

6.3 Directions for Future Research

- **Cross-country Comparative Studies** – Expanding the scope to include banks from multiple countries could validate the generalizability of findings and identify regional variations in AI's impact.
- **Sector-specific AI Analysis** – Future work could examine AI's role in specific banking verticals, such as retail lending, corporate finance, or wealth management.
- **Integration of ESG Factors** – Linking AI adoption with environmental, social, and governance performance to explore AI's role in sustainable banking.
- **Advanced Econometric Modeling** – Applying machine learning-assisted econometric approaches or time-varying parameter models to capture evolving AI effects.
- **Customer-centric Impact Assessment** – Incorporating customer satisfaction and service quality metrics into AI performance evaluations.

References

- 1) Agarwal, P., Swami, S., & Malhotra, S. K. (2022). Artificial intelligence adoption in the post COVID-19 new-normal and role of smart technologies in transforming business: A review. *Journal of Science and Technology Policy Management*, 15(3), 506–529. <https://doi.org/10.1108/JSTPM-08-2021-0122>
- 2) Bandyopadhyay, A. (2022). Bank financial performance and its linkage with capital: A dynamic panel data analysis of public sector banks in India. *The Indian Economic Journal*, 70(3), 437–451. <https://doi.org/10.1177/00194662221104752>
- 3) Bansal, C., Pandey, K. K., Goel, R., & Sharma, A. (2024). Analysis of barriers to AI banking chatbot adoption in India: An ISM and MICMAC approach. *Issues in Information Systems*, 25(4), 417–441.
- 4) Barney, J. (1991). Firm resources and sustained competitive advantage. *Journal of Management*, 17(1), 99–120.
- 5) Deepthi, B., Gupta, P., Rai, P., & Arora, H. (2022). Assessing the dynamics of AI-driven technologies in the Indian banking and financial sector. *Vision*. <https://doi.org/10.1177/09722629221087371>
- 6) DiMaggio, P. J., & Powell, W. W. (1983). The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields. *American Sociological Review*, 48(2), 147–160.
- 7) Rogers, E. M. (2003). *Diffusion of Innovations* (5th ed.). Free Press.
- 8) Sheth, J. N., Jain, V., Roy, G., & Chakraborty, A. (2022). AI-driven banking services: The next frontier for a personalised experience in the emerging market. *International Journal of Bank Marketing*, 40(6), 1248–1271. <https://doi.org/10.1108/IJBM-09-2021-0449>
- 9) Editorial note: this list replaces an earlier version in which several entries — including citations attributed to “Joshi & Agarwal, 2022,” “Sharma & Rao, 2019,” “Subramanian et al., 2023,”

- 10) Ni, Y., Nyhoff, J., Napier, M., & Townsend, D. (2026). AI innovation and bank performance: Evidence from patent activity of large U.S. commercial banks. *Journal of Risk and Financial Management*, 19(4), 247. <https://doi.org/10.3390/jrfm19040247>
- 11) Hamdouni, A. (2025). From ESG to financial stability: Unpacking the multi-dimensional impact of AI-driven FinTech-related technology adoption on bank performance. *International Journal of Financial Studies*, 13(4), 234. <https://doi.org/10.3390/ijfs13040234>
- 12) Zhang, S. (2026). The influence of digital transformation on Chinese bank profitability performance. *PLOS ONE*. <https://doi.org/10.1371/journal.pone.0340249>
- 13) Safiullah, M., & Paramati, S. R. (2022). The impact of FinTech firms on bank financial stability. *Electronic Commerce Research*. <https://doi.org/10.1007/s10660-022-09595-z>
- 14) Alnaser, F. M. I., Rahi, S., Alghizzawi, M., & Ngah, A. H. (2023). Does artificial intelligence (AI) boost digital banking user satisfaction? *Heliyon*, 9(8), e18930. <https://doi.org/10.1016/j.heliyon.2023.e18930>
- 15) Jafri, J. A., Mohd Amin, S. I., Abdul Rahman, A., & Mohd Nor, S. (2023). A systematic literature review of the role of trust and security on Fintech adoption in banking. *Heliyon*, 9(11), e22980. <https://doi.org/10.1016/j.heliyon.2023.e22980>