



AI GOVERNANCE AND FINANCIAL RISK: DISCLOSURE BASED AI READINESS INDEX FOR INDIAN BANKS

Snehitha Bommidi¹, Dr. Golli Mohan²

¹Research Scholar, Department of Business Administration, KL Business School, KL University,
Koneru Lakshmaiah Education Foundation.

²Research Guide, Assistant Professor, Department of Business Administration, KL Business School, KL University,
Koneru Lakshmaiah Education Foundation,

Article DOI: <https://doi.org/10.36713/epra31842>

DOI No: 10.36713/epra31842

ABSTRACT

Indian banks have been moving rapidly towards the use of artificial intelligence, but there is still no consensus on the best way to assess how well they govern it. In this paper, a disclosure-based AI-Readiness Index is developed and tested for Indian banks, employing sixteen indicators divided into four categories: governance and oversight, risk management, capability and adoption, and transparency and ethics. Each indicator is assigned a score from 0 to 2 according to the actual information that the bank provides in its FY2025–26 annual report. Fifteen listed banks were assessed in this manner – five being public-sector and ten private-sector – resulting in an overall score out of 100. HDFC Bank achieves the highest score at 90.6 and Yes Bank comes second at 84.4; in both cases these are the only two banks that explicitly link their AI policies to the RBI's Framework for Responsible and Ethical Enablement of AI (FREE-AI), which was issued in August 2025. SBI is next at 81.3. Canara Bank is at the lowest level with a score of 25.0.

The index is statistically robust (with a Cronbach's alpha of 0.87) and the ranking remains almost unchanged when different weighting methods are applied – whether equal weights, entropy weights or TOPSIS all show very similar results. A k-means clustering method divides the banks into three broad groups that correspond to the categories of adopters as described by the Diffusion of Innovation theory: Leaders, Followers and Laggards. On average private banks score higher than public-sector banks (61.3 compared to 48.8), although with only five public-sector banks in the sample this difference does not reach statistical significance. The correlations between the index and financial-risk indicators – NPAs, ROA, CRAR – are weak and inconclusive, and the paper states this clearly rather than presenting the figures in a more favourable light. What the study provides is a practical, regulator-linked method of measuring AI governance readiness, directly aligned with the structure of FREE-AI, thus addressing a gap that is not covered by any of the studies reviewed here: an index of this kind, designed for Indian banks, tested against real risk outcomes after the FREE-AI framework had been introduced.

1. INTRODUCTION

Today, credit decisions, fraud monitoring, customer service and internal risk assessments routinely make use of some kind of machine learning or automation, with these systems often integrated so deeply into everyday operations that it's easy to fail to notice just how much the situation has evolved. However, there has been no real method for measuring how banks are overseeing this change — for example, who is responsible, what kind of supervision is in place, whether the models are examined before going live, and how much of this is actually put down in writing. While industry surveys offer a lot of information regarding AI adoption figures, the topic of governance is one that is less

developed and more recent, and almost no one has systematically tried to measure it among Indian banks specifically.

When the Reserve Bank of India released its Framework for Responsible and Ethical Enablement of AI (FREE-AI) in August 2025, it included seven guiding principles, six pillars (Infrastructure, Policy, Capacity, Governance, Protection, Assurance) and twenty-six recommendations for regulated entities. Indian banks have, for the first time, something definite to measure themselves. This paper asks three questions: how ready are Indian banks to govern AI, based on their own disclosures, and does readiness differ by ownership type? Does the scoring method hold up statistically across different weighting choices? And does disclosed AI-readiness relate to financial risk and performance (NPAs, ROA, CRAR, cost-to-income)?



2. THEORETICAL FRAMEWORK

Theory	Core citation	Role in this study
Resource-Based View	Barney (1991)	Motivates H2/H3 and the Capability pillar as a proxy for the underlying resource
Institutional Theory	DiMaggio & Powell (1983)	Explains why disclosure converges; frames FREE-AI as regulatory pressure
Signalling Theory	Spence (1973)	Justifies treating disclosure as an imperfect proxy; underlies the 0/1/2 scoring rule
Agency Theory	Jensen & Meckling (1976)	Frames each indicator as a monitoring device; separates existence from quality
Diffusion of Innovation	Rogers (1962/2003)	Supports the Leader/Follower/Laggard grouping in Section 5.7

3. LITERATURE REVIEW

3.1 A field growing rapidly, with little attention to governance

Bouladasse, El Ganich and Yahyaoui (2026) analyse 891 Scopus-listed AI-in-banking papers (2014–2024) and find the field growing rapidly (277 papers in 2024 alone), concentrated on credit scoring, fraud detection and fintech. Governance remains comparatively underdeveloped as a theme — directly motivating this paper’s focus.

3.2 What the non-AI governance literature already shows

Ocansey and Antwi (2026) find governance quality — not mere existence — is associated with lower risk in financial institutions, and call for more AI-specific, emerging-market research beyond cross-sectional designs — exactly what this paper attempts.

3.3 The nearest measurement precedents

Domi, Godolja and Sevrani (2026) build an AI Maturity Index for Albanian banks via survey and recommend moving from perceptions to institutional documents — the approach this paper takes. Abdo et al. (2026) find smart-governance dimensions inversely related to financial risk in Iraqi banks, though caution against causal claims from two years of sector data. Pluskota et al. (2026) apply a similarly evidence-cited 0–4 maturity rubric to global banks, though for ESG rather than financial-risk outcomes.

3.6 The impact that this has left

3.4 The two Indian studies this paper must be honest about

Two Indian studies are close enough to this paper’s topic to require direct discussion. Rao, Pathak and Ibrahim (2025) build a voluntary AI-disclosure index for twelve Nifty Bank constituents (2019–2023) and find disclosure positively predicts customer deposits, more strongly for public-sector banks. Rao, Srivastava and Mejia-Amaya (2024) use a dynamic panel on the same twelve banks and find AI integration positively associated with Return on Equity, with disclosure quality varying widely across banks. Neither paper separates governance structure from general AI-activity disclosure, neither tests NPA or CRAR, and neither could anchor to FREE-AI, which did not yet exist. This paper does not claim to be first, but to measure governance structure specifically, against risk rather than performance, after FREE-AI redefined what governance should look like in India.

3.5 Measurement methods worth borrowing

Ogundele and Nzama’s (2025) binary 1/0 disclosure coding produces a counter-intuitive sign on credit-risk disclosure, motivating this paper’s graded 0/1/2 scale instead. Kaur et al. (2024) and Albuquerque et al. (2022) respectively supply the size-control variable and the key validity caution used here — that bank narrative can be managed rather than neutral, hence this paper’s consistent framing of “disclosed readiness” rather than “verified capability.” Rajan (2026) and Aljunaid et al. (2025) shaped the Risk and Transparency indicator definitions but are not direct empirical precedents.

Dimension	This study	Rao, Pathak & Ibrahim (2025)	Rao, Srivastava & Mejia-Amaya (2024)	Domi et al. (2026, Albania)
AI construct measured	16-indicator governance-centred index, 4 pillars	AI disclosure index	AI adoption + content analysis	AI Maturity Index (survey)
Data source	Annual reports (governance, risk, MD&A, BRSR)	Annual reports	Annual reports + content analysis	85 professionals' survey responses
Outcome tested	Gross/Net N-+PA, ROA, CRAR, Cost-to-Income	Customer deposits	Return on equity	Perceived benefits, scaling intention
Regulatory anchor	RBI FREE-AI (2025), explicit indicator mapping	None stated	None stated	None (TOE + institutional theory)
Aggregation / robustness	Equal, entropy, TOPSIS weights; leave-one-out; k-means	Random-effects panel, Hausman test	Dynamic panel	Equal-weight check, bootstrap, PLS-SEM



In simple terms, none of the papers mentioned in this review—the eleven main papers together with the two discovered during the additional search for Indian studies—includes all of the following features at the same time: a governance-focused multi-pillar index, a sample from India, risk outcomes instead of deposits or ROE, an explicit FREE-AI anchor, and this sort of multi-method robustness checking; that is exactly what this paper aims to provide.

4. METHODOLOGY

4.1 Sample

The sample consists of fifteen of India's listed scheduled commercial banks: five from the public sector (the State Bank of India, Punjab National Bank, the Bank of Baroda, Canara Bank, and the Union Bank of India) and ten from the private sector (HDFC Bank, ICICI Bank, Axis Bank, Kotak Mahindra Bank, IndusInd Bank, Yes Bank, IDFC First Bank, Federal Bank, Bandhan Bank, and RBL Bank). They were selected on the grounds of convenience—because they are listed and have a publicly available English-language annual report for the financial year 2025–26—and since together they include both large public banks, large private banks, and smaller or banks that are going through a turnaround. This is not a random or stratified sample of India's approximately 140 scheduled commercial banks (the category also encompasses small finance banks, regional rural banks, and foreign banks operating in India), and the index should not be interpreted as reflecting the sector as a whole (more information on this is given in Section 7).

4.2 The indicator framework and scoring

There are sixteen indicators distributed among four pillars, each containing four indicators: Governance and Oversight (G1–G4), Risk Management and Controls (R1–R4), Capability and Adoption (C1–C4), and Transparency and Ethics (T1–T4). For each indicator, a score of 0, 1 or 2 is assigned based on the bank's FY2025–26 annual report — specifically the Corporate Governance Report, the Management Discussion & Analysis, the Risk Management section, and the BRSR — by referring to the codebook given in Appendix A. A score of 0 indicates that nothing was found, a score of 1 means that something general or partial was identified — for instance, a general IT policy that does not mention AI, or a committee that is named but whose actual duties are not described — and a score of 2 means that something specific and named has been found — such as a board-approved AI policy or a Chief Data Officer whose responsibilities are outlined. Each of the scores in this study is supported by a page

reference and is recorded in an evidence log, just as Pluskota et al. (2026) did when developing their own maturity rubric.

4.3 Index construction

For each bank $i = 1, \dots, 15$ and each indicator $j = 1, \dots, 16$, let s_{ij} be the raw score and normalise it to $\tilde{s}_{ij} = s_{ij} / 2$ so that it lies in the interval $[0, 1]$. The baseline index, with equal weights, is just the average over all sixteen indicators:

$$I_i^{\wedge EW} = (1/16) \sum_{j=1}^{16} \tilde{s}_{ij} \times 100$$

Two further methods of aggregation serve as a remedy for the fact that equal weighting is, in the end, an arbitrary decision.

Entropy weighting

It is conventional in the field of composite-indicator research (as shown in the OECD's Handbook on Constructing Composite Indicators and in Domi et al.'s 2026 study, which used a similar approach). The principle behind this is that indicators which show a great deal of variation among banks contain more information and therefore should be given greater weight:

$$p_{ij} = \tilde{s}_{ij} / \sum_i \tilde{s}_{ij}; \quad e_j = -(1/\ln n) \sum_i p_{ij} \ln(p_{ij}); \quad w_j^{\wedge ENT} = (1 - e_j) / \sum_k (1 - e_k)$$

There are 15 banks when $n = 15$. The entropy-weighted index is given by $I_i^{\wedge ENT} = \sum_j \tilde{s}_{ij} w_j^{\wedge ENT} \times 100$.

TOPSIS

TOPSIS — which is known as the Technique for Order Preference by Similarity to Ideal Solution — evaluates banks according to their geometric proximity to the best score obtained on each criterion and according to their distance from the worst score. With the same normalised matrix V (and with all the weights being equal, that is, $1/16$ each), the ideal and anti-ideal vectors are given by $v_j^+ = \max_i(v_{ij})$ and $v_j^- = \min_i(v_{ij})$, thus providing each bank with two distances:

$$D_i^+ = \sqrt{\sum_j (v_{ij} - v_j^+)^2}; \quad D_i^- = \sqrt{\sum_j (v_{ij} - v_j^-)^2}; \quad C_i = D_i^- / (D_i^+ + D_i^-)$$

C_i multiplied by 100 is given as a third, independently derived composite score in section 5.6.

Worked example: State Bank of India

It is helpful to look at an example using actual figures. SBI's sixteen scores are G1=2, G2=2, G3=2, G4=2, R1=1, R2=2, R3=1, R4=2, C1=2, C2=2, C3=2, C4=2, T1=1, T2=1, T3=1, T4=1, which amounts to 26 out of the possible 32. The equal-weight index is $(26/32) \times 100 = 81.25$, as shown in Table 4. When broken down by pillar, Governance scores $(2+2+2+2)/8 \times 100 = 100$; Risk scores $(1+2+1+2)/8 \times 100 = 75$; Capability scores $(2+2+2+2)/8 \times 100 = 100$; and Transparency scores $(1+1+1+1)/8 \times 100 = 50$. That final figure is the point worth noting — strong performance in governance and capability, but weaker transparency — a point which is also mentioned in Section 6.

4.4 Hypotheses

Hypothesis	Statement
H1	AI-readiness index scores aren't the same for public and private banks.
H2	Banks with higher AI-readiness scores show lower Gross and Net NPA ratios.
H3	Banks with higher AI-readiness scores tend to have higher Return on Assets.
H4	The patterns in H2 and H3 hold even after factoring in bank size and ownership type.
H5	All sixteen indicators actually work together as a reliable single construct.



4.5 Financial variables and regression design

I pulled numbers for Gross NPA, Net NPA, Return on Assets, CRAR, and Cost-to-Income straight from the FY2025–26 annual reports—always from the Financial Highlights section. If there was a choice between standalone and consolidated figures, I stuck with standalone to keep things consistent across banks. The plan, following Ogundele and Nzama (2025) and Kaur et al. (2024), was to use this regression for H4:

$$\text{Risk}_i = \beta_0 + \beta_1 \cdot \text{AIReadiness}_i + \beta_2 \cdot \ln(\text{Assets})_i + \beta_3 \cdot \text{BankType}_i + \varepsilon_i$$

I ran this separately for Gross NPA, Net NPA, and ROA. The sample is small and not perfectly balanced, so I used robust standard errors. But honestly, with just eleven to fifteen banks and a few predictors, doing a full multiple regression here gives you almost no real statistical power—it ends up looking more precise than it actually is. So instead of reporting those coefficients, most of the results are just straightforward bivariate correlations.

4.6 Reliability and validity

To check if all sixteen indicators hold together as one construct (that's H5), I calculated Cronbach's alpha the standard way. I also compared the three different weighting methods using Spearman rank correlations (see Section 5.6 for details). For sensitivity, I used a leave-one-indicator-out test: basically, I recalculated the index sixteen times, each time leaving out a different indicator, and tracked how the rankings changed—the same way Domi et al. (2026) did it (covered in Section 5.8).

Statistic	Value
N	15
Mean	57.08
Std. Dev. (population)	18.31
Minimum	25.00 (Canara Bank)
Maximum	90.63 (HDFC Bank)
Cronbach's α (16 items)	0.871

An alpha score of more than 0.87 far exceeds the threshold value of 0.7, suggesting that the 16 items function as a single assessment rather than independent measurements, despite the mix of indicators drawn from four pillars. Thus, H5 is supported. 5.2 Full ranking

4.7 A typology: k-means clustering

Instead of stopping at a ranked list, I also grouped the banks using k-means clustering ($k=3$) based on their four pillar scores. This follows the Leader/Follower/Laggard categories used by Domi, Godolja, and Sevrani (2026), and it lines up with the adopter categories from Diffusion of Innovation theory (see Section 2.5). The clusters are ranked by their centroid averages and labeled accordingly—you'll find those results in Section 5.7.

4.8 Endogeneity

The primary concern is reverse causality, as better performing and less risky banks could be both inclined to disclose more information. A single cross-sectional observation cannot identify this possibility, and an instrumental-variable analysis was not attempted, as no plausible instruments (predictors strongly related to the disclosure variables but not directly associated with risk) were identified. The causal nature of the linkages reported in this article is deliberately understated, as emphasized in Section 7.

5. RESULTS

5.1 Descriptive statistics

The equal-weight index ranges from 25 (Canara Bank) to 90.6 (HDFC Bank) across the 15 banks analyzed (Table 3), with the mean value being 57.1 and the population standard deviation being 18.3.

Gov/Risk/Cap/Trans are the four pillar scores, each ranging from 0 to 100. GNPA/NNPA/ROA/CRAR are the standalone financial indicators for FY2025–26; “n/a” means that the number was not found in any of the pages scanned

Rank	Bank	Type	Index	Gov	Risk	Cap	Trans	GNPA%	NNPA%	ROA%	CRAR%
1	HDFC Bank	Private	90.63	87.5	100.0	87.5	87.5	1.15	0.38	1.94	19.70
2	Yes Bank	Private	84.38	87.5	87.5	75.0	87.5	1.30	0.20	0.80	15.80
3	State Bank of India	Public	81.25	100.0	75.0	100.0	50.0	1.49	0.39	1.07	15.40
4	Axis Bank	Private	68.75	75.0	75.0	75.0	50.0	1.23	0.37	1.46	16.42
4	Bandhan Bank	Private	68.75	100.0	62.5	37.5	75.0	3.27	0.97	0.64	17.30
6	Punjab National Bank	Public	59.38	87.5	62.5	62.5	25.0	2.95	0.29	0.89	17.74
6	IndusInd Bank	Private	59.38	75.0	75.0	50.0	37.5	3.43	1.00	n/a	17.48
8	ICICI Bank	Private	56.25	87.5	75.0	25.0	37.5	1.44	0.35	n/a	17.18
9	Kotak Mahindra Bank	Private	53.13	75.0	37.5	25.0	75.0	1.23	0.31	2.06	22.40
10	RBL Bank	Private	50.00	50.0	75.0	50.0	25.0	1.45	0.39	0.53	14.25
11	Union Bank of India	Public	40.63	62.5	37.5	25.0	37.5	2.82	0.48	1.25	18.10



11	IDFC First Bank	Private	40.63	75.0	25.0	50.0	12.5	1.61	0.56	n/a	15.60
11	Federal Bank	Private	40.63	50.0	50.0	12.5	50.0	1.62	0.20	1.15	17.25
14	Bank of Baroda	Public	37.50	62.5	37.5	37.5	12.5	1.89	0.45	1.06	16.31
15	Canara Bank	Public	25.00	37.5	37.5	12.5	12.5	1.84	0.43	1.10	17.04

5.3 Patterns across the four pillars

The first pillar, Governance, demonstrates the highest degree of consistency across the banks analyzed, with the mean score of 73.3 out of 100. Fourteen out of 15 banks report the existence of a technology committee at the board level, which supports the findings of the theoretical institutional analysis in Section 2.2 on the emergence of disclosure practices mandated by regulators well before the introduction of the FREE-AI initiative.

Aspect	Governance	Risk	Capability	Transparency
Governance	1.00	0.49	0.62	0.54
Risk	0.49	1.00	0.67	0.56
Capability	0.62	0.67	1.00	0.35
Transparency	0.54	0.56	0.35	1.00

Each of the six pairs of pillars is positively correlated at the medium level (between 0.35 and 0.67), suggesting a general tendency for similarity between different domains without indicating the presence of measurement duplication (none of the correlations exceed the threshold of 0.70 typically associated with measurement equivalence). The highest similarity is observed between Capability and Risk (0.67), which is consistent with the RBV theory's proposition regarding the concurrent development of a bank's ability to design and implement AI systems and the corresponding validation and governance infrastructure needed to operate them safely. The association of Transparency with other domains is the weakest, with the maximum value observed for its correlation with Risk at 0.56. This result aligns with the qualitative observations made during the development of Table 2:

Pillar	Public mean (n=5)	Private mean (n=10)	Difference
Governance	70.0	76.25	+6.25
Risk	50.0	66.25	+16.25
Capability	47.5	48.75	+1.25
Transparency	27.5	53.75	+26.25

Make up most of the gap, while Governance and Capability are similar between ownership structures. This aligns with the institutional-theory expectation that coercive pressure has the strongest effect on governance, a domain historically more uniformly guided by RBI expectations, and that the more discretionary areas (explicit ethics statement, a detailed description of risk controls) reflect ownership-related differences in disclosure culture or resourcing. Running a separate Mann-Whitney test per pillar (rather than the overall index) highlights

In contrast, the Transparency/Ethics pillar is the most varied, with a mean score of 43.3 and a large standard deviation of 24.1 (Table 4). Only three banks (Yes, HDFC, and Bandhan) exceed the threshold of 75, reporting specific responsible/ethical-AI statements. At the same time, the five banks (Canara, Baroda, PNB, IDFC First, and Union Bank) score equal to or lower than 25. The associations between the pillars are presented in Table 4.

several banks demonstrate strong disclosure of capabilities and governance arrangements but do not report specific ethical-AI commitments.

5.4 H1 — Public vs private

Private banks (n = 10, mean 61.25, SD 16.94) have higher average scores than public-sector banks (n = 5, mean 48.75, SD 21.94). A Mann-Whitney U test, the appropriate choice given the small sample and highly skewed distribution, gives U = 35.5, p = 0.218; a Welch's t-test gives t = 1.12, p = 0.303. Neither reaches conventional significance, and H1 is not supported — as it is noted in the reporting, this is an inconclusive result given only five public banks and not a finding that there is no difference. Table 5 shows the mean differences by pillar.

this: Transparency comes nearest to significance (U = 39.5, p = 0.082), followed by Risk (n = 50__, n = 51__), while Governance (n = 52__, n = 53__) and Capability (n = 54__, n = 55__) show virtually no separation between ownership structures. None of the four pillars reach conventional 0.05 significance — again due to the small public n — but Transparency is the pillar in which a larger sample would plausibly reveal a real difference.

5.5 H2/H3/H4 — Index against financial-risk indicators

Table 6 shows Pearson correlations between the equal-weight index and the five financial-risk variables.



Variable	Pearson r	N (pairs)	Direction consistent with H2/H3?
Gross NPA ratio	-0.166	15	Yes (weak)
Net NPA ratio	-0.017	15	Negligible
Return on Assets	+0.135	11	Yes (weak)
CRAR	+0.058	15	Negligible
Cost-to-Income ratio	-0.227	10	Yes (weak)

All correlations are weak, and with n peaking at 15 (and dropping as low as 10 where a figure couldn't be located), none would reach conventional significance at the 0.05 level. H2, H3 and H4 are not supported, and the multiple regression planned for H4 isn't run with reported coefficients for the same reason — three or four predictors against eleven to fifteen data points is not reliable. However, three of the five variables — Gross NPA,

ROA, Cost-to-Income — do point in the direction H2/H3 would predict, and can be noted as a pattern to follow with a larger sample. With up to five separate correlation tests on a sample this small, the chance of one coming up significant by chance alone is real — and this paper does not correct for them, largely because none came close enough to being significant.

Variable	Banks with data (of 15)	Banks missing
Gross NPA ratio	15	0
Net NPA ratio	15	0
Return on Assets	11	4 (ICICI, IndusInd, IDFC First, and one other)
CRAR	15	0
Cost-to-Income ratio	10	5
Provision Coverage Ratio	12	3

Table 6b provides the exact breakdown of each financial variable available, as the sample size varies substantially by variable and affects how much weight any single coefficient should have.

A note on cross-study comparison, for transparency: an independently created disclosure index covering the same fifteen banks and fiscal year indicates a markedly different finding for this section — a strong, statistically significant positive correlation between AI-readiness and Cost-to-Income (reported as a Spearman $n = 56$, $n = 57$). This paper's own Cost-to-Income correlation is weak and negative (Pearson $r = -0.227$, $n = 58$), and the two studies diverge significantly in individual scores (e.g., HDFC Bank and Federal Bank score substantially differently across the two indices). Rather than selecting one result over the other without having re-verified both from scratch, this paper reports the figures it calculated and flags the

disagreement explicitly, as additional support for the single-scorer reliability limitation raised in Section 7: two independent readings of similar disclosures, by different scorers, can diverge substantially, and Cohen's kappa testing is thus listed as a priority next step.

5.6 Robustness: three ways of weighting the same data

In addition to the plain equal-weight index, two other methods were run: entropy weighting, which de-emphasizes indicators that don't vary much across banks, and TOPSIS, which ranks banks by the distance from the best and worst observed profiles (Section 4.3). All three agree closely: Spearman $n = 59$ between equal and entropy weighting, and 0.98 between equal weighting and TOPSIS, both $U = 35.50$.

The ranking isn't an artefact of picking equal weights.

Bank	Equal-weight	Entropy-weight	TOPSIS (closeness ×100)
HDFC Bank	90.63	80.31	77.84
Yes Bank	84.38	85.89	72.57
State Bank of India	81.25	75.58	69.10
Axis Bank	68.75	53.93	57.52
Bandhan Bank	68.75	58.38	57.52
Punjab National Bank	59.38	34.22	47.36
IndusInd Bank	59.38	43.58	50.00
ICICI Bank	56.25	37.66	45.55
Kotak Mahindra Bank	53.13	35.26	42.92
RBL Bank	50.00	34.91	43.30
Union Bank of India	40.63	19.94	32.94
IDFC First Bank	40.63	24.42	35.01
Federal Bank	40.63	19.78	32.21
Bank of Baroda	37.50	20.32	33.33
Canara Bank	25.00	3.62	17.91



HDFC and Yes Bank remain as the top two, and Canara Bank as last, regardless of methodology. What's interesting are the changes: entropy weighting discounts indicators almost every bank has (the board IT committee, G1, thirteen of fifteen banks scoring 2) and emphasizes those that actually distinguish banks (a named AI policy, as an example). Punjab National Bank falls most heavily under this methodology, going from an rank-6, 59.4 score down to 34.2 — a larger drop in ranking proportionally than any other bank. Its equal-weight score, in other words, is heavily

reliant on indicators most banks already have, and not on what sets it apart. Yes Bank, however, barely moves (actually increases slightly, to 85.9), as its score was already largely based on more differentiated disclosure to begin with.

5.7 Sorting banks into tiers: k-means

Running k-means ($U = 35.51$) on the four pillar scores gives three groups, labelled using the DOI-style language from Sections 2.5 and 4.7.

Tier	Banks (n)	Centroid — Gov	Centroid — Risk	Centroid — Cap	Centroid — Trans
Leaders	SBI, HDFC Bank, Yes Bank (3)	91.7	87.5	87.5	75.0
Followers	PNB, ICICI, Axis, IndusInd, Bandhan, RBL (6)	79.2	70.8	50.0	41.7
Laggards	BoB, Canara, Union, Kotak, IDFC First, Federal (6)	60.4	37.5	27.1	33.3

What most distinguishes leaders from followers is Capability (87.5 vs 50.0 — a 37.5-point gap), and what distinguishes leaders from laggards is risk (87.5 vs 37.5, a 50-point difference). The dividing line between tiers is not the governance structure, which is reasonably comparable across all three groups (as per the isomorphism point from Section 5.3), but demonstrated capability and a described risk process. One case worth noting directly is Kotak Mahindra Bank, which falls into the Laggards tier despite having a detailed, named enterprise AI framework (built around “build, control and secure” principles, per the evidence log). Its Risk (37.5) and Capability (25.0) scores drag it down as no described model-validation process or dedicated AI team appeared in the pages reviewed. It demonstrates the signalling-theory point from Section 2.3: naming a framework but not following up with described operational detail does not score it as highly as one that does.

5.8 How much does any one indicator matter?

To determine if the ranking relies heavily on any single indicator, the equal-weight index was rerun sixteen times, each time omitting one indicator, and each resulting ranking compared back to the full one using Spearman correlation. The range across all sixteen reruns is 0.980 (dropping C3, the dedicated AI/analytics team indicator) to 1.000 (dropping G1, the near-universal committee indicator), averaging 0.992. No single indicator, dropped individually, has much effect — which is reassuring, as a second scorer could conceivably disagree with one or two of this study's judgment calls without affecting the overall picture meaningfully.

6. LIMITATIONS AND ENDOGENEITY

A number of things limit what can be claimed here, and they're worth stating plainly rather than burying.

- Disclosure isn't capability. A high score means a bank wrote something down, not that an outside party verified it's actually practised. Albuquerque et al. (2022) showed bank narrative can be managed rather than neutral, and that caution applies just as much here — it's why this paper keeps calling the index a measure

of “disclosed readiness,” not “verified capability,” rather than letting the distinction blur.

- One scorer, no second opinion. All fifteen banks were scored by a single researcher. Without a second independent scorer, Cohen's kappa — the check Domi et al. (2026) use, and that this paper's own literature review found missing from almost every comparable disclosure-coding study — couldn't be run. That's a real gap, not a minor caveat, and it should be closed before this index gets used to support any firm conclusion.

- Not every bank got the same depth of read. Five banks — SBI, PNB, Bank of Baroda, Canara, Federal — were scored through close, page-by-page reading with full context checked for every indicator. The other ten went through a faster, structured search-and-verify process across the full text of each report, necessarily quicker given there were over 4,000 combined pages across fifteen reports to get through. Both approaches produced cited, page-referenced scores, but the depth of cross-checking wasn't perfectly even, and a handful of borderline calls — documented in the evidence log — would benefit from a second reviewer's eyes.

- Small and not randomly chosen. Fifteen banks, picked by convenience, out of roughly 140 scheduled commercial banks in India. Neither the index nor the k-means tiers should be read as describing the sector as a whole — just this particular set of fifteen.

- Not enough statistical power. At $n = 15$ (and as low as 10 for some variables), none of the tests in Section 5 can reliably detect anything short of a fairly large effect. The non-significant results for H1 through H4 mean “inconclusive,” not “no effect,” and that's why H4's regression wasn't estimated with reported coefficients (Sections 4.5, 5.5).

- A single year can't settle cause and effect. This design can't tell whether better AI governance lowers risk, whether already low-risk banks simply disclose more, or whether both come from something else entirely — bank size, management quality, general disclosure habits. No instrument was available, and even a multi-year panel extension would still only have fifteen cross-sectional units, too few for reliable fixed-effects estimation.



- Timing works against finding much. FY2025–26 reports are the first full-year reports published after FREE-AI came out in August 2025. Banks have had very little time to actually change anything in response, which probably understates whatever effect FREE-AI eventually has on disclosed practice in later years.

- Some financial figures are just missing. ROA couldn't be located for three banks (ICICI, IndusInd, IDFC First), and Cost-to-Income for five. These are left blank rather than estimated or guessed, which shrinks the effective sample for the correlation tests (Table 6) further, but avoids putting made-up numbers into the analysis.

- A second independent index of the same banks disagrees with this one in places. As noted in Section 5.5, an independently produced disclosure index covering the same fifteen banks and fiscal year reports different individual scores for at least two banks (HDFC Bank, Federal Bank) and a differently signed Cost-to-Income correlation. This paper has not attempted to adjudicate which reading is more accurate — doing so properly would require re-reading both sets of source pages side by side, which is exactly the second-scorer exercise recommended as the top priority in Section 8. Readers should treat both studies' bank-specific scores as provisional until that reconciliation happens.

8. CONCLUSION AND FUTURE RESEARCH

This paper builds a sixteen-indicator, four-pillar AI-Readiness Index for Indian banks, scored from FY2025–26 annual reports and tied directly to the RBI's FREE-AI framework. Across fifteen listed banks, HDFC and Yes Bank come out on top, both with explicit AI governance policies that name FREE-AI directly; Canara Bank sits at the bottom. The index holds together statistically (alpha of 0.87) and barely moves across three different aggregation methods — equal weighting, entropy weighting, TOPSIS — all correlating above 0.95, and survives a sixteen-way sensitivity test averaging 0.99. A k-means grouping sorts the sample into Leader, Follower and Laggard tiers that line up with how Diffusion of Innovation theory describes adopters. Private banks disclose more governance structure on average than public-sector banks, mostly concentrated in Transparency and Risk, though the gap doesn't reach statistical significance given the sample size. The correlations with actual financial-risk indicators come out weak and inconclusive — reported here as an honest null result rather than stretched into something it isn't, which is more than can be said for parts of the literature this paper reviewed.

Five things would make an obvious next step. A second independent scorer, even on a subsample, would finally allow a real Cohen's kappa test — probably the single biggest gap in this version of the study. Extending the financial data to a three-year panel (FY2023–24 through FY2025–26) would let someone check whether readiness changes actually come before risk changes, which would chip away at the endogeneity problem even if it couldn't fully resolve it. A bigger, more representative sample — small finance banks, regional rural banks included — would fix both the power problem and the generalisability problem, and would make the H4 regression actually estimable instead of just specified. As FREE-AI implementation matures over the next few

reporting cycles, redoing this index in two or three years would allow a genuine before-and-after comparison this single-year study can't offer, and would let the k-means tiers be tested for whether banks actually move between them — a direct test of the diffusion dynamic DOI theory predicts but this snapshot can't see. And pairing this paper's governance-disclosure approach with Rao, Pathak and Ibrahim's (2025) customer-response data might finally explain the odd mismatch noted in Section 6.3, between which banks disclose the most and which banks' disclosure actually moves customers.

9. REFERENCES

1. Abdo, S. M. S., Al-Jubouri, M. I. M., Ali, S. A. G., & Saeed, Y. J. (2026). *The role of smart governance in mitigating financial risks in the era of artificial intelligence: Evidence from Iraqi banks*. *F1000Research*, 15 211.
<https://doi.org/10.12688/f1000research.175117.2>
2. Albuquerque, F., Stoltzemburg, V., & Cariano, A. (2022). *Identifying impression management strategies using intelligence: A research applied to the banking sector in Portugal*. *Polish Journal of Management Studies* 26(1).
<https://doi.org/10.17512/pjms.2022.26.1.01>
3. Aljunaid, S. K., Almheiri, S. J., Dawood, H., & Khan M. A. (2025). *Secure and transparent banking: AI-driven federated learning model for financial fraud detection*. *Journal of Risk and Financial Management*, 18(4) 179.
<https://doi.org/10.3390/jrfm18040179>
4. Barney, J. (1991). *Firm. Sustained competitive advantage*. *Journal of Management*, 17(1) 99–120.
5. Bouladasse, H., El Ganich, S., & Yahyaoui T. (2026). *Artificial intelligence in banking: A analysis of research trends, intellectual structure and performance-related themes*. *Journal of Risk and Financial Management*, 19(8) 581.
<https://doi.org/10.3390/jrfm19080581>
6. DiMaggio, P. J., & Powell W. W. (1983). *The iron cage revisited: isomorphism and collective rationality in organizational fields*. *American Sociological Review*, 48(2) 147–160.
7. Domi, L., Godolja, M., & Sevrani, K. (2026). *Artificial intelligence. Organizational readiness in banking: Evidence from Albania*. *Administrative Sciences*, 16(8) 385.
<https://doi.org/10.3390/admsci16080385>
8. Jensen, M. C., & Meckling W. H. (1976). *Theory of the firm: behavior, agency costs and ownership structure*. *Journal of Financial Economics*, 3(4) 305–360.
9. Kaur, A., Joshi, M., Singh, G., & Sharma, S. (2024). *The impact of reputation on cost of debt: A panel data analysis of Indian listed firms*. *Journal of Risk and Financial Management*, 17(8) 367.
<https://doi.org/10.3390/jrfm17080367>
10. Ocansey E. O. N. D., Antwi, M. O. (2026). *Corporate. Financial risk management in financial institutions: A systematic literature review and future research agenda*. *Journal of Economics, Finance and Accounting Studies*, 8(9) 1–15.
<https://doi.org/10.32996/jefas.2026.8.9.1>
11. Ogundele, O. S., & Nzama, L. (2025). *Risk management. Financial performance: Analysing credit and liquidity risk management and disclosures by Nigerian banks*. *Journal of Risk and Financial Management*, 18(4) 198.
<https://doi.org/10.3390/jrfm18040198>



12. Pluskota, P., Slupinska, K., Wawrzyniak, A., & Wasikowska, B. (2026). The application of intelligence (AI) in the implementation of ESG-oriented sustainable development strategies in the banking sector: A case study. *Sustainability*, 18(2) 732. <https://doi.org/10.3390/su18020732>
13. Rajan, A. (2026). Artificial intelligence, in decision-making: Opportunities, risks and challenges. *International Scientific Journal of Engineering and Management* 5(9). <https://doi.org/10.55041/ISJEM08595>
14. Rao, K. P. V., Pathak, D., & Ibrahim F. (2025). Does voluntary AI disclosure influence customer behaviour?
15. Panel evidence from Indian banks. *Banks and Bank Systems*, 20(4), 76–86. [https://doi.org/10.21511/bbs.20\(4\).2025.07](https://doi.org/10.21511/bbs.20(4).2025.07)
16. Rao, P., Srivastava, N., & Mejia-Amaya, A. F- 2024. Effect of artificial intelligence on the financial performance of Indian banking sector. *Journal of Infrastructure, Policy and Development*, 8(15), 9511. <https://doi.org/10.24294/jipd9511>
17. Reserve Bank of India. (2023). Master Direction – Information Technology Governance, Risk, Controls and Assurance Practices.
18. Reserve Bank of India. (2025). Framework for Responsible and Ethical Enablement of Artificial Intelligence (FREE-AI).