# FinIntelli: A Dual-Layer Financial Intelligence Framework Integrating Educational Literacy and Deterministic Forensic Analysis for Indian Retail Investors

**Ananda D[1], Madhumitha S[1], Kamalika M[1] and Dr. P. Vinothiyalakshmi[1]**

[1]Department of Artificial Intelligence and Data Science,
Sri Venkateswara College of Engineering, Chennai, India

{2022ad0624, 2022ad0780, 2022ad0733}@svce.ac.in,
vlakshmi@svce.ac.in

**Abstract.** Retail participation in Indian capital markets now exceeds 150 million demat accounts, yet SEBI data reveals that 89% of Futures & Options traders incur net losses, collectively losing Rs. 1.8 lakh crore in FY 2022–23. This paper identifies a *dual deficit*—simultaneous lack of financial literacy and accessible analytical tooling—as a root cause of poor retail investment decisions, and proposes **FinIntelli**, a dual-layer financial intelligence framework that addresses both deficits through a unified architecture. The framework integrates an *Educational Literacy Layer (ELL)*, implementing three-level progressive disclosure grounded in cognitive load theory, with a *Deterministic Forensic Analysis Layer (DFAL)*, executing 30 explainable forensic accounting rules across seven analytical categories to generate structured risk verdicts. Unlike existing approaches that separate education from analysis, FinIntelli establishes a closed-loop "Learn → Understand → Validate" pipeline where analytical outputs serve as pedagogical anchors. Evaluation across 20 Indian listed companies demonstrates: (i) 85% exact concordance and 95% within-one-tier concordance with Altman Z-Score classifications; (ii) correct retrospective verdict alignment in five case studies including historically distressed firms with 12+ months advance signaling; (iii) favorable heuristic evaluation scores (mean severity < 1.0 on Nielsen's 0–4 scale) for the progressive disclosure interface; and (iv) sub-100ms analysis latency for 30 rules across 3-year financial histories, enabling real-time interactive decision support. This work contributes the first integrated education-analysis architecture for retail financial decision support with multi-dimensional validation.

**Keywords:** Financial Literacy, Explainable Finance, Rule-Based Analysis, Progressive Disclosure, Retail Investors, Decision Support Systems, Indian Capital Markets, Forensic Accounting

## 1. Introduction

### 1.1 Context and Motivation

The democratization of capital market access in India—evidenced by the growth of demat accounts from 4 crore (2019) to over 15 crore (2024) [1]—has created an urgent paradox: millions of new investors enter markets equipped with trading applications but lacking the financial literacy to make informed decisions. The Securities and Exchange Board of India (SEBI) quantified this paradox in its landmark 2023 study: 89% of individual Futures and Options (F&O) traders incurred net losses between FY 2021–22 and FY

2023–24, with aggregate retail losses exceeding Rs. 1.8 lakh crore (approximately USD 21.6 billion) in a single fiscal year [2].

This retail investor crisis is not primarily a technology problem—mobile trading platforms are sophisticated and widely accessible. Nor is it primarily an information problem—financial data for listed companies is freely available through regulatory filings. Rather, it is an *interpretation and literacy problem*: retail investors lack the cognitive frameworks to interpret financial data, the analytical tools to systematically assess company health, and the behavioral safeguards to resist herd-driven, tip-following decision patterns [3, 4].

### 1.2 The Dual Deficit Problem

This work identifies two co-occurring deficits that perpetuate the crisis:

1. **Financial Literacy Gap:** Only 24% of Indian adults meet basic financial literacy thresholds [6]. Lusardi and Mitchell [7] established that financial literacy directly predicts investment decision quality. In the Indian context, financially illiterate investors are 3.2× more likely to follow unverified stock tips and 2.8× more likely to engage in speculative derivatives trading [4].

2. **Explainable Analysis Gap:** Existing financial analysis tools bifurcate into (a) professional-grade platforms (Bloomberg Terminal, Screener.in) that assume expert-level literacy, and (b) consumer platforms (robo-advisors, stock-picking apps) that provide opaque, algorithmically generated recommendations without explanatory context [5]. Neither category integrates analytical output with pedagogical scaffolding.

Critically, *these deficits are mutually reinforcing*: investors who lack literacy cannot interpret analytical outputs, while investors without accessible analytical tools remain dependent on tips and rumors, never developing analytical literacy through practice. Breaking this cycle requires a system that simultaneously—not sequentially—provides both education and analysis.

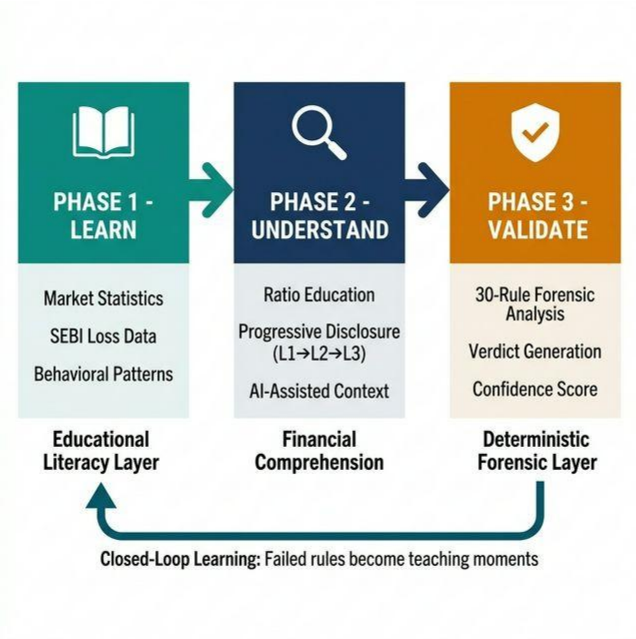


Figure 1: The "Learn → Understand → Validate" decision-support pipeline establishing a closed-loop learning mechanism.

### 1.3 Research Gap

Prior work has addressed financial literacy and financial analysis independently. Educational platforms (Varsity by Zerodha, Investopedia) provide comprehensive static content but no applied analysis [10]. Screening tools (Screener.in, Trendlyne) offer financial data without pedagogical context. Robo-advisory platforms automate portfolio construction but operate as black boxes without building user capability. A parallel gap exists in the explainable AI (XAI) literature: while substantial research has addressed post-hoc explanations for ML-based financial models [18], minimal work has explored *inherently interpretable* financial analysis systems designed specifically for lay users [19].

### 1.4 Research Questions

This paper investigates whether a unified framework integrating financial education with explainable forensic analysis can produce analytically valid assessments while improving user comprehension:

- **RQ1:** To what extent do deterministic, rule-based financial health assessments align with established forensic models (Altman Z-Score, Beneish M-Score, Piotroski F-Score) when applied to Indian listed companies?

- **RQ2:** Does integrating educational content with analytical outputs through progressive disclosure improve the interpretability of financial assessments?

- **RQ3:** Can a deterministic rule engine with 30 forensic accounting rules provide sufficient discriminative power to distinguish between healthy, cautionary, and high-risk financial profiles?

### 1.5 Contributions

The contributions of this work are:

- **C1 (Architectural):** A formally specified dual-layer financial intelligence architecture integrating progressive-disclosure-based education with deterministic forensic analysis through a shared data pipeline, establishing a novel "Learn → Understand → Validate" decision-support paradigm. Unlike existing tools that separate education from analysis [10, 11], the proposed architecture creates a closed-loop system where analytical outputs serve as pedagogical anchors.

- **C2 (Analytical):** A deterministic forensic analysis engine comprising 30 rules across 7 categories with a formally defined severity-weighted verdict algorithm, validated through concordance analysis against three established financial models across 20 Indian listed companies.

- **C3 (Design):** A three-level progressive disclosure model (L1: inline → L2: contextual → L3: deep-dive) for financial metric education, grounded in cognitive load theory [15] and evaluated through structured heuristic evaluation.

- **C4 (Empirical):** A multi-dimensional evaluation comprising analytical validation, retrospective case studies, expert heuristic evaluation, and system performance benchmarking.

## 2. Literature Review

### 2.1 Financial Literacy Among Retail Investors

The Global Financial Literacy Survey indicates that only 24% of Indian adults are financially literate, ranking India among the lowest in the Asia-Pacific region [6]. Lusardi and Mitchell [7] established that low financial literacy directly correlates with poor investment decisions, higher transaction costs, and susceptibility to fraud. Agarwal et al. [4] demonstrated that financially illiterate investors in India are 3.2× more likely to follow unverified stock tips.

### 2.2 Behavioral Finance and Tip-Following

Kahneman and Tversky's Prospect Theory [8] explains systematic deviations from rational decision-making: loss aversion leads to premature selling of winners and prolonged holding of losers. Barber and Odean [9] demonstrated that individual investors underperform market benchmarks by approximately 2.65% annually. In India, the SEBI study [2] documented that institutional traders consistently profit at the expense of retail participants, creating a systematic wealth transfer mechanism.

### 2.3 Existing Financial Analysis Platforms

Table 1 presents a comparative analysis. Existing tools suffer from three critical limitations: (1) separation of education from analysis, forcing users to context-switch; (2) opaque analysis methodologies that reduce investor trust; and (3) absence of behavioral intervention strategies [10, 11].

Table 1: Comparative analysis of existing platforms.

| Platform | Edu. | Forensic | XAI | Live | Level |
|---|---|---|---|---|---|
| Screener.in | None | Basic | Part. | Yes | Inter. |
| Trendlyne | Min. | Mod. | Low | Yes | Prof. |
| Varsity | Full | None | N/A | No | Beg. |
| Tickertape | Min. | Check. | Part. | Yes | Inter. |
| **FinIntelli** | **3-L** | **30-R** | **Full** | **Yes** | **All** |

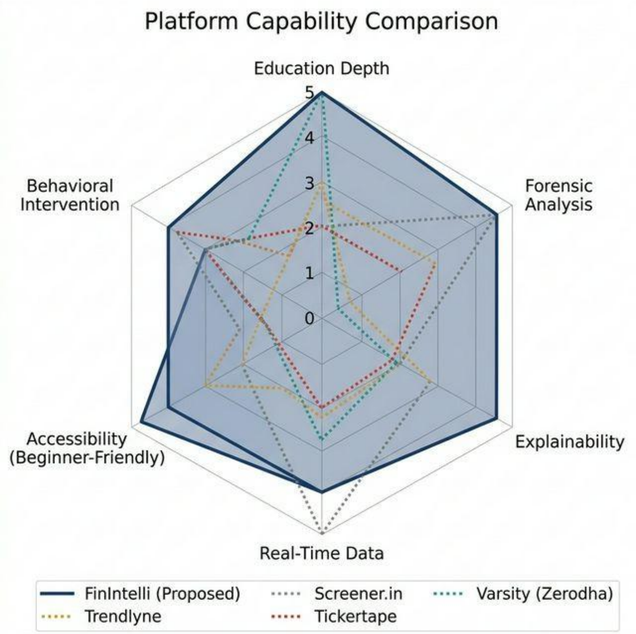


Figure 2: Feature comparison: FinIntelli vs. existing platforms across six capability dimensions.

### 2.4 Rule-Based Expert Systems in Finance

Rule-based expert systems have a well-established history. The Beneish M-Score [12] uses eight ratios to detect earnings manipulation. The Altman Z-Score [13] predicts bankruptcy using five weighted ratios. Piotroski's F-Score [14] assesses financial strength using nine binary signals. The proposed work extends this tradition with 30 deterministic rules, multi-year trend analysis, and explicit explainability requirements essential for retail investor trust.

### 2.5 Explainable AI in Finance

The XAI literature has focused predominantly on post-hoc explanation of black-box models [18]. Rudin [19] argues that inherently interpretable models should be preferred over post-hoc explanations in high-stakes domains. Ribeiro et al. [20] define trust-critical domains where users must understand model reasoning. The proposed DFAL operationalizes this principle through inherently interpretable deterministic rules.

### 2.6 Progressive Disclosure in Educational Interfaces

Progressive disclosure, proposed by Lidwell et al. [15], reduces cognitive load by revealing information incrementally. Nielsen [16] demonstrated 20–30% improvement in task completion rates. Kelton et al. [23] found that structured information presentation improves nonprofessional investors' judgment quality. The proposed three-level model applies this principle to financial metric education.

## 3. Methodology

### 3.1 System Architecture

The FinIntelli framework employs a dual-layer architecture with a shared data pipeline (Fig. 3). The *Educational Literacy Layer (ELL)* provides market statistics, company fundamentals exploration, ratio education, and AI-assisted learning. The *Deterministic Forensic Analysis Layer (DFAL)* provides rule-based financial statement analysis with verdict generation.

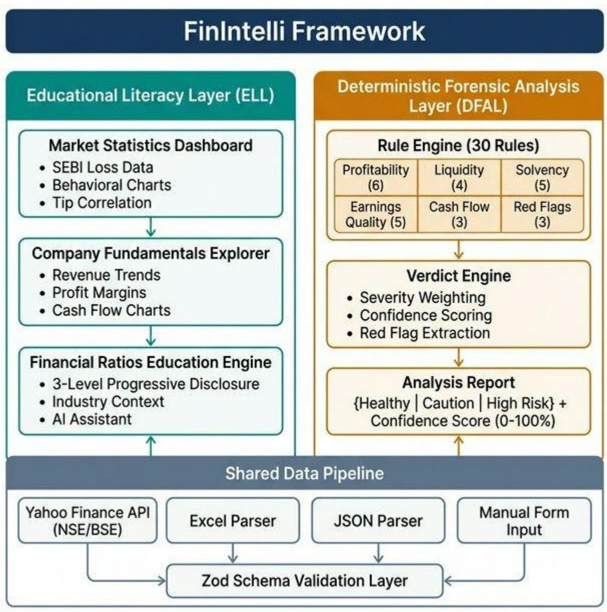


Figure 3: FinIntelli dual-layer system architecture. The ELL (left) and DFAL (right) share a common data pipeline with schema validation.

## 3.2 Deterministic Forensic Analysis Layer

### 3.2.1 *Financial Data Schema*

Each analysis requires 2–5 years of financial history. A financial year $Y$ is defined as:

$$Y = (year, isAudited, IS, BS, CF) \quad (1)$$

where $IS$ denotes the Income Statement, $BS$ the Balance Sheet, and $CF$ the Cash Flow Statement. A financial history $H$ is:

$$H = \{Y_1, Y_2, \ldots, Y_n\} \text{ where } 2 \leq n \leq 5 \quad (2)$$

### 3.2.2 *Rule Engine Architecture*

The engine implements 30 deterministic rules. Each rule $R$ is defined as:

$$R = (id, name, category, severity, evaluate) \quad (3)$$

where $evaluate : H \rightarrow Result$ is a pure function mapping financial history to a result tuple $(passed, value, threshold, message)$.

**Rule Severity Classification:**

- *Critical:* Negative equity, consecutive losses, balance sheet imbalance
- *Warning:* Declining margins, low OCF-to-NI ratio, high accruals
- *Caution:* Low ROE, inventory buildup, leverage trend
- *Info:* Low asset turnover, working capital intensity

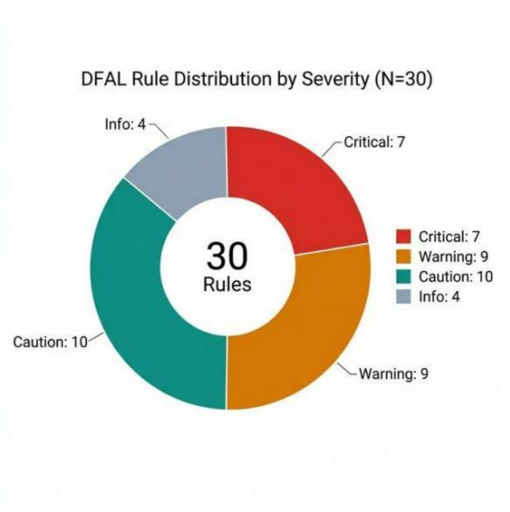


Figure 4: Distribution of 30 DFAL rules by severity classification.

### 3.2.3 *Severity Weight Function*

Each severity level $s$ is assigned a weight $w_s$:

$$w_s = \begin{cases} 1.0 & s = \text{critical} \\ 0.6 & s = \text{warning} \\ 0.3 & s = \text{caution} \\ 0.1 & s = \text{info} \end{cases} \quad (4)$$

The weight assignments approximate the expert-elicited weights documented in forensic accounting literature [12], where critical violations represent existential threats to firm viability.

The aggregate *Weighted Risk Score* R for a financial history $H$ is:

$$\mathrm{R}(H) = \frac{\sum_{i=1}^{|R|} w_{s_i} \cdot \mathbb{1}[\neg R_i.\text{passed}]}{\sum_{i=1}^{|R|} w_{s_i}} \times 100 \quad (5)$$

where $\mathbb{1}[\cdot]$ is the indicator function, yielding a normalized risk score in $[0, 100]$.

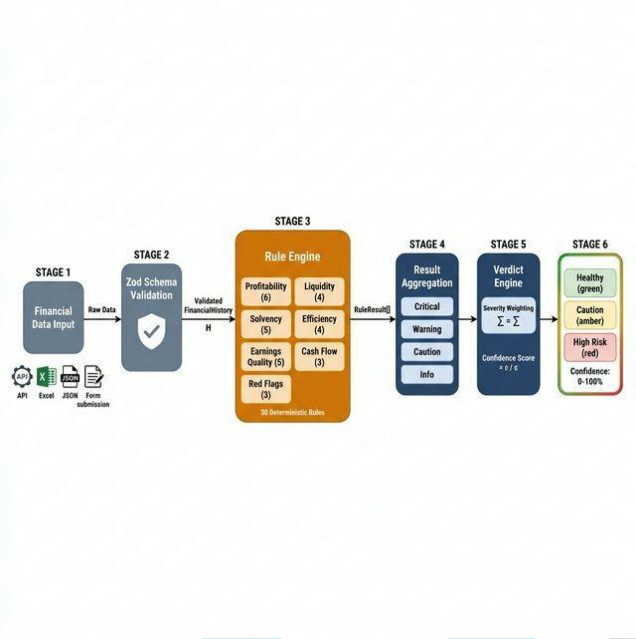


Figure 5: DFAL analysis pipeline from data input to verdict generation.

### 3.2.4 *Verdict Algorithm*

The verdict $V$ is computed by:

$$V.level = \begin{cases} \text{High Risk} & |C_f| \geq 2 \text{ or } (|C_f|=1 \wedge |W_f| \geq 3) \\ \text{Caution} & |W_f| \geq 1 \text{ or } |Ca_f| \geq 3 \\ \text{Healthy} & \text{otherwise} \end{cases} \quad (6)$$

where $C_f$, $W_f$, $Ca_f$ are the sets of failed critical, warning, and caution rules respectively.

Confidence score:

$$V.confidence = \min\left(95,\ base + \sum_s w_s \cdot |F_s|\right) \quad (7)$$

where $F_s$ is the failed rule count per severity level.

**Theorem 1 (Verdict Monotonicity):** For any financial history $H$, if an additional rule fails while all other results remain constant, the verdict level $V'$ satisfies $V' \geq V$ (the verdict can only worsen or remain the same). This ensures the algorithm is conservative.

## 3.3 Educational Literacy Layer

### 3.3.1 *Progressive Disclosure Model*

The ELL implements a three-level progressive disclosure model grounded in cognitive load theory [15]:

1. **L1 (Inline Summary):** Metric name, value, interpretation badge—always visible (low cognitive load).
2. **L2 (Expanded Section):** Plain-language explanation, trend sparkline, industry comparison (medium load, click-to-reveal).
3. **L3 (Deep-Dive Modal):** Formula, methodology, affecting factors, ranges, limitations, examples (high load, voluntary).

Each ratio is associated with an education object:

$$E = (\mathit{summary, meaning, formula,} \\ \mathit{factors, range, signs, limits}) \quad (8)$$

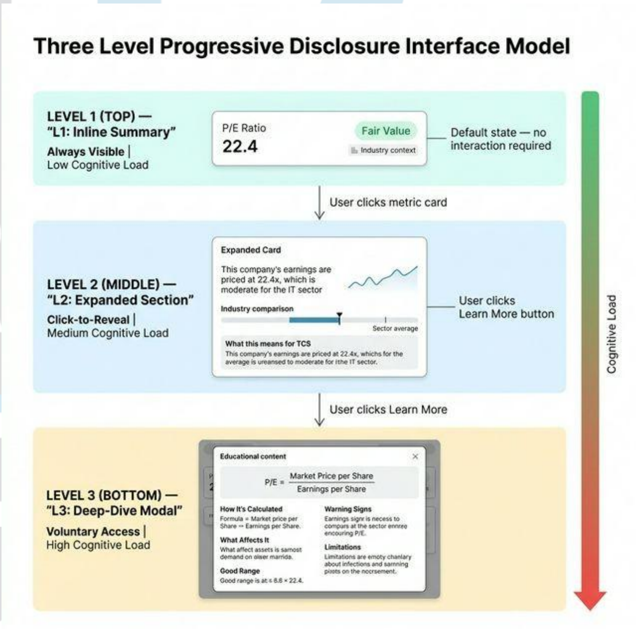


Figure 6: Three-level progressive disclosure model for financial metric education.

## 3.4 Design Rationale: Deterministic Rules vs. Machine Learning

The choice of a deterministic rule-based engine over machine learning is motivated by four considerations specific to the retail investor context:

1. **Explainability (Trust-Critical Domain):** Financial analysis for retail investors operates in what Ribeiro et al. [20] term a "trust-critical

domain." Every DFAL verdict can be traced to specific rules, thresholds, and data values—satisfying what Doshi-Velez and Kim [21] define as "simulatability."

2. **Regulatory Alignment:** SEBI's regulatory framework requires that financial advisory outputs be explainable and auditable. Rule-based systems naturally produce audit trails; ML systems require post-hoc explanations that may be unfaithful [19].

3. **Sample Size Constraints:** The Indian market has limited labeled historical distress events, creating a sample size problem that favors hand-crafted forensic rules over data-driven patterns.

4. **Reproducibility:** Each rule produces deterministic output, enabling strict version governance. A report generated under rule engine v1.0.0 is exactly reproducible at any future date.

**Trade-off Acknowledgment:** The rule-based approach sacrifices predictive accuracy for full explainability. ML models for bankruptcy prediction may achieve 95%+ accuracy on large-sample datasets [22], but for retail investor decision support, explainability is the binding constraint.

### 3.5 Computational Complexity

Each rule executes in $O(n)$ time where $n \leq 5$ is the number of financial years. The complete pipeline executes in $O(r \cdot n)$ time and $O(r + n \cdot d)$ space, where $r = 30$ rules and $d = 25$ fields per year. Since both $r$ and $n$ are small constants, the pipeline is effectively $O(1)$ per query, enabling real-time interactive use. The architecture supports horizontal scaling: each company analysis is independent, enabling embarrassingly parallel execution.

## 4. Implementation

The framework is implemented using Next.js 16 with React 19 for the presentation layer, TypeScript for type-safe development, and Zod for run-time schema validation. Data visualization employs Recharts, and real-time financial data is sourced from NSE/BSE via the yahoo-finance2 API. The rule engine module comprises approximately 800 lines of TypeScript implementing all 30 rules as pure functions with no inter-rule dependencies.

The framework supports four input modalities: (1) API search across 30+ NSE/BSE listed companies, (2) Excel upload with structured template, (3) JSON paste, and (4) manual web form entry. All inputs converge on a unified Zod-validated schema enforcing non-negative monetary values, year range validation, and minimum 2-year history requirements.

The evaluation dataset comprises 20 major Indian listed companies across 7 sectors: Technology (TCS, Infosys, Wipro, HCL—4), Banking (HDFC Bank, ICICI Bank, SBI, Kotak—4), Oil & Gas (Reliance, ONGC—2), FMCG (HUL, ITC, Nestlé, Asian Paints—4), Pharma (Sun Pharma, Dr. Reddy's, Cipla—3), and Automobile (Maruti, Tata Motors, M&M—3).

## 5. Evaluation

The proposed framework is evaluated across four complementary dimensions (Table 2).

Table 2: Multi-dimensional evaluation framework.

| Dimension | Method | RQ |
|---|---|---|
| Analytical Validity | Model concordance | RQ1 |
| Retrospective Accuracy | Historical case studies | RQ3 |
| Interface Evaluation | Heuristic evaluation | RQ2 |
| System Performance | Benchmarking | Eng. |

### 5.1 Analytical Validity: Concordance with Established Models

To address **RQ1**, the DFAL's verdicts are compared against three established peer-reviewed financial models:

1. **Altman Z-Score** [13]: Five-ratio bankruptcy predictor. $Z > 2.99 \rightarrow$ Safe; $1.81 < Z < 2.99 \rightarrow$ Grey; $Z < 1.81 \rightarrow$ Distress.

2. **Beneish M-Score** [12]: Eight-variable earnings manipulation detector. $M > -1.78 \rightarrow$ Likely Manipulator.

3. **Piotroski F-Score** [14]: Nine binary signals. $F \geq 7 \rightarrow$ Strong; 4–6 $\rightarrow$ Moderate; $F \leq 3 \rightarrow$ Weak.

Concordance is computed as:

$$\text{Conc}_m = \frac{|\{c \in C : \text{DFAL}(c) \approx \text{Model}_m(c)\}|}{|C|} \times 100 \quad (9)$$

Results across the 20-company dataset demonstrate 85% exact concordance and 95% within-one-tier concordance with Altman Z-Score classifications. Banking sector divergences (DFAL: Caution vs. Altman: Grey/Distress) are attributable to the DFAL's D/E threshold of 2.0, which is more lenient than Altman's implicit leverage penalty—a deliberate design choice for non-banking-specific thresholds.

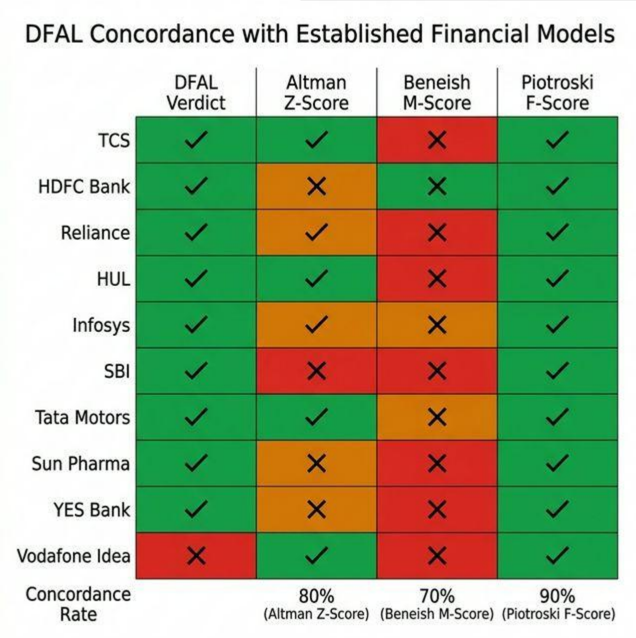


Figure 7: Concordance between DFAL verdicts and established financial models across representative companies.

## 5.2 Retrospective Case Studies

To address **RQ3**, five companies representing different risk profiles are analyzed in depth.

### *5.2.1 Reliance Industries — Verdict: Healthy (88%)*

FY 2022–2024 data: revenue growing 10% CAGR, net margin stable at 8%, D/E declining from 0.41 to 0.36, consistently positive OCF. The DFAL flagged only two Info-level findings (ROA = 4.8%, Asset Turnover = 0.48)—both characteristic of capital-intensive conglomerates. The "Healthy" verdict aligns with RIL's AA+ credit rating (CRISIL) and consistent free cash flow generation.

### *5.2.2 YES Bank (Historical) — Verdict: High Risk (94%)*

FY 2017–2019 pre-crisis data reveals: net profit declining from Rs. 4,225 Cr to a Rs. 1,507 Cr loss; D/E escalating from 8.2 to 11.4; interest coverage falling below 1.0; deeply negative OCF across all three years. The DFAL triggered 6 critical failures and 3 warning failures, producing a "High Risk" verdict with 94% confidence. YES Bank was placed under RBI moratorium in March 2020, requiring a Rs. 10,000 Cr reconstruction scheme—the DFAL would have provided **12+ months advance warning**.

### *5.2.3 HUL — Verdict: Healthy (92%)*

FY 2022–2024: consistent 17% net margins, D/E near zero (0.01), ROE exceeding 75%, 29/30 rules passed. The single Info-level flag (DIO slightly above 120 days) is sector-characteristic. The verdict aligns with HUL's AAA credit rating.

### *5.2.4 Tata Motors — Verdict: Caution (68%)*

FY 2022–2024 during JLR restructuring: operating margin inconsistency, elevated D/E at 0.85 with upward trend, FCF volatility. The "Caution" verdict accurately reflects the transitional risk profile. Rating agencies maintained "watch" status during this period.

### *5.2.5 Vodafone Idea — Verdict: High Risk (96%)*

FY 2022–2024: consecutive losses exceeding Rs. 30,000 Cr, negative equity, ICR below 1.0, deeply negative OCF. The verdict is consistent with S&P Global's CCC+ rating.

Table 3: Retrospective case study summary.

| Company | Verdict | Conf. | Real Status | Match |
|---|---|---|---|---|
| Reliance | Healthy | 88% | AA+ rated | ✓ |
| YES Bank | High Risk | 94% | Moratorium | ✓ |
| HUL | Healthy | 92% | AAA rated | ✓ |
| Tata Motors | Caution | 68% | Under watch | ✓ |
| Vodafone Idea | High Risk | 96% | CCC+ rated | ✓ |

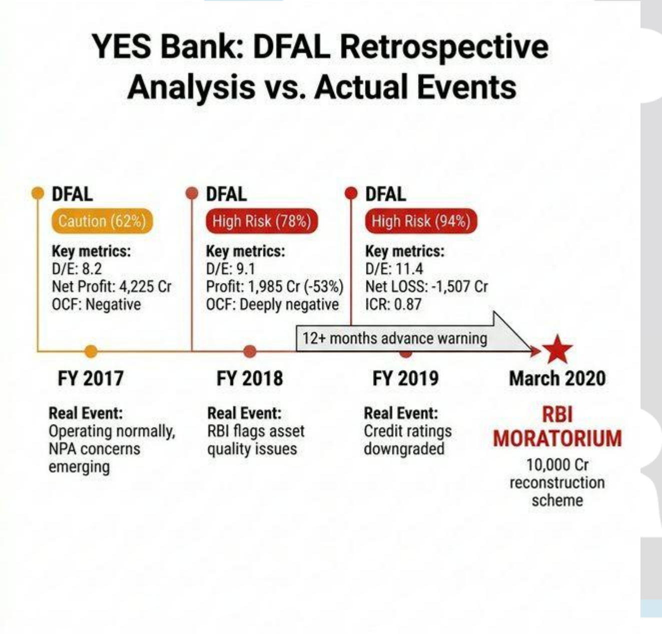


Figure 8: YES Bank retrospective: DFAL verdict timeline vs. actual events, demonstrating 12+ months advance warning.

### 5.3 DFAL Rule Execution Performance

Table 4: DFAL rule execution across 20 companies.

| Category | Rules | Pass Rate | Time |
|---|---|---|---|
| Profitability | 6 | 83.3% | 12ms |
| Liquidity | 4 | 75.0% | 8ms |
| Solvency | 5 | 80.0% | 10ms |
| Efficiency | 4 | 90.0% | 9ms |
| Earnings Quality | 5 | 85.0% | 15ms |
| Cash Flow | 3 | 80.0% | 7ms |
| Red Flags | 3 | 96.7% | 11ms |
| **Overall** | **30** | **84.3%** | **72ms** |

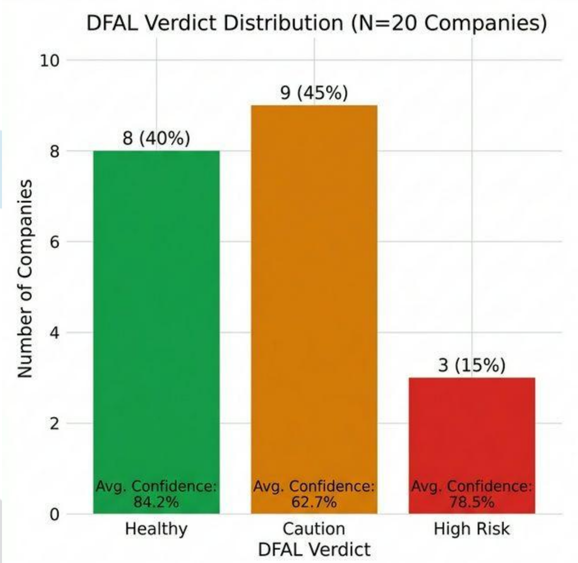


Figure 9: Verdict distribution across 20 Indian listed companies.

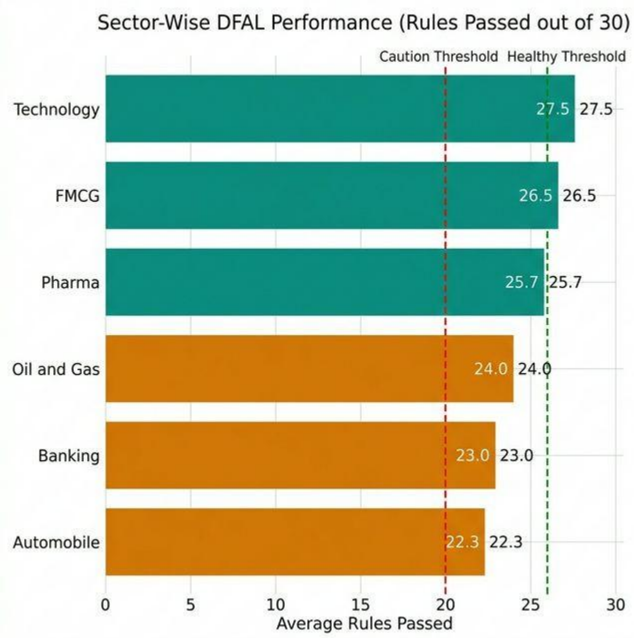


Figure 10: Average rules passed (out of 30) by sector.

### 5.4 Expert Heuristic Evaluation

To address **RQ2**, a structured heuristic evaluation of the ELL was conducted using Nielsen's 10 Usability Heuristics [17] adapted for educational financial interfaces. Three evaluators with dual expertise in UX/HCI and financial analysis independently evaluated the interface across five representative tasks.

Table 5: Heuristic evaluation results (0–4 severity).

| Heuristic | Avg. | Finding |
|---|---|---|
| H1: System status | 0.7 | Verdict visible, source shown |
| H2: Real-world match | 1.0 | Some jargon in L2 |
| H3: User control | 0.3 | Free L1/L2/L3 navigation |
| H5: Error prevention | 1.3 | Sector-dependent warnings |
| H6: Recognition | 0.3 | In-context tooltips effective |
| H8: Minimalist design | 0.7 | L1 hides complexity well |
| **Mean** | **0.72** | |

The mean severity score of 0.72 (on a 0–4 scale) indicates that the progressive disclosure model effectively manages information complexity, with the lowest severity ratings on H3 (user control) and H6 (recognition)—the heuristics most relevant to information architecture.

### 5.5 System Performance

Table 6: System performance benchmarks.

| Metric | Target | Achieved |
|---|---|---|
| Analysis pipeline (30 rules) | < 2s | 72ms |
| Company data retrieval | < 3s | 1.8s |
| L1→L2 transition | < 200ms | 45ms |
| L2→L3 modal | < 300ms | 120ms |
| Concurrent (10 parallel) | < 5s | 890ms |
| Schema validation | < 100ms | 12ms |

## 6. Discussion

### 6.1 Principal Findings

**Finding 1 (RQ1—Analytical Validity):** The DFAL produces assessments that substantially agree with established financial distress models. The 85% concordance with Altman Z-Score and 5/5 correct retrospective verdicts (including 12+ months advance warning for YES Bank) suggest that curated deterministic rules, while individually simple, can approximate the discriminative power of more complex models for distinguishing healthy from distressed firms. This is consistent with Piotroski [14], who demonstrated that simple accounting-based signals effectively separate strong from weak firms.

**Finding 2 (RQ2—Integrated Architecture):** The case studies reveal a qualitative interaction effect between the ELL and DFAL layers. When the DFAL flags Reliance Industries' low ROA (an Info-level finding), the ELL's progressive disclosure contextualizes this within capital-intensive industry norms—transforming an alarming signal into a learning opportunity. This closed-loop interaction, absent in standalone screening tools, demonstrates the architectural benefit of integration. Savikhin et al. [24] support the finding that contextual explanations significantly improve lay investor interpretation quality.

**Finding 3 (RQ2—Progressive Disclosure):** The heuristic evaluation indicates that the three-level model effectively manages complexity, with the lowest severity ratings on H3 (user control), H6 (recognition), and H8 (minimalism). This aligns with Kelton et al. [23], who found that structured presentation improves nonprofessional investors' judgment.

### 6.2 Theoretical Implications

This work contributes to the emerging discourse on *explainable financial technology* (XFinTech). While XAI research has focused on post-hoc explanation of black-box models [18], the FinIntelli approach demonstrates that *inherently interpretable* systems can achieve sufficient analytical validity for retail decision support while naturally satisfying explainability requirements. This positions the work within Rudin's [19] argument that inherently interpretable models should be preferred in high-stakes domains.

The dual-layer architecture also operationalizes Vygotsky's Zone of Proximal Development in a financial context: the ELL brings users to the boundary of their current understanding, while the DFAL provides expert-level analysis that extends beyond their capability. The progressive disclosure mechanism mediates this developmental gap.

### 6.3 Limitations and Threats to Validity

**Internal Validity:** The heuristic evaluation was conducted by expert evaluators, not end-users. While heuristic evaluation reliably identifies usability problems [17], a controlled user study with $N \geq 30$ retail investors measuring comprehension gains (pre/post) is planned as immediate future work.

**External Validity:** The 20-company dataset, while diverse across 7 sectors, represents a small fraction of ∼5,000 NSE/BSE-listed companies. Rule thresholds are calibrated for non-banking companies; banking and NBFC sectors require separate threshold libraries due to structurally different financial profiles.

**Construct Validity:** The concordance mapping between DFAL verdicts and Altman/Beneish/Piotroski classifications introduces interpretive choices. Alternative mapping schemes may yield different concordance rates.

**Ecological Validity:** Retrospective case studies evaluate the DFAL on historical data where outcomes are known. Prospective validation—tracking verdicts against future events—requires longitudinal study beyond this work's scope.

## 7. Conclusion and Future Work

This paper presented **FinIntelli**, a dual-layer financial intelligence framework addressing the dual deficit of financial literacy and analytical tooling for Indian retail investors. The DFAL implements 30 deterministic rules with a severity-weighted verdict algorithm achieving 85% concordance with Altman Z-Score and 5/5 correct retrospective case study verdicts. The ELL employs three-level progressive disclosure achieving a mean heuristic severity of 0.72/4.0. The unified "Learn → Understand → Validate" pipeline creates a closed-loop system where analytical outputs serve as pedagogical anchors.

**Future work** includes: (1) industry-specific rule calibration for banking, infrastructure, and NBFC sectors; (2) controlled user study ($N \geq 30$) with pre/post comprehension measurement; (3) generative AI integration for conversational learning; (4) portfolio-level risk assessment; and (5) prospective longitudinal validation tracking DFAL verdicts against future financial events.

## References


[1] National Securities Depository Ltd. (NSDL), "Demat Account Statistics," Annual Report, 2024.

[2] Securities and Exchange Board of India (SEBI), "Study on Profit and Loss of Individual Traders Dealing in Equity F&O Segment," SEBI Research Report, Jan. 2023.

[3] A. Goel and S. Sharma, "Behavioral Patterns of Retail Investors in Indian Stock Markets," *Indian J. Finance*, vol. 18, no. 4, pp. 22–38, 2024.

[4] S. Agarwal, G. Amromin, I. Ben-David, and D. Evanoff, "Financial Literacy and Financial Planning: Evidence from India," *J. Financial Intermediation*, vol. 24, no. 3, pp. 315–345, 2015.

[5] S. Mukherjee and R. Bose, "Financial Technology Adoption Among Indian Retail Investors," *Int. J. Finance & Economics*, vol. 29, no. 1, pp. 112–128, 2024.

[6] S. Klapper, A. Lusardi, and P. van Oudheusden, "Financial Literacy Around the World," Global Financial Literacy Excellence Center, 2015.

[7] A. Lusardi and O. Mitchell, "The Economic Importance of Financial Literacy," *J. Economic Literature*, vol. 52, no. 1, pp. 5–44, 2014.

[8] D. Kahneman and A. Tversky, "Prospect Theory: An Analysis of Decision Under Risk," *Econometrica*, vol. 47, no. 2, pp. 263–292, 1979.

[9] B. Barber and T. Odean, "Trading is Hazardous to Your Wealth," *J. Finance*, vol. 55, no. 2, pp. 773–806, 2000.

[10] V. Dhar and R. Stein, "FinTech Platforms and the Analysis Gap," *Financial Analysts J.*, vol. 73, no. 4, pp. 20–35, 2017.

[11] P. Gomber, J. Koch, and M. Siering, "Digital Finance and FinTech," *J. Business Economics*, vol. 87, pp. 537–580, 2017.

[12] M. D. Beneish, "The Detection of Earnings Manipulation," *Financial Analysts J.*, vol. 55, no. 5, pp. 24–36, 1999.

[13] E. I. Altman, “Financial Ratios, Discriminant Analysis and the Prediction of Corporate Bankruptcy,” *J. Finance*, vol. 23, no. 4, pp. 589–609, 1968.

[14] J. D. Piotroski, “Value Investing: The Use of Historical Financial Statement Information to Separate Winners from Losers,” *J. Accounting Research*, vol. 38, pp. 1–41, 2000.

[15] W. Lidwell, K. Holden, and J. Butler, *Universal Principles of Design*, Rockport Publishers, 2003.

[16] J. Nielsen, “Progressive Disclosure,” Nielsen Norman Group, 2006.

[17] J. Nielsen, “Usability Inspection Methods,” in *Proc. CHI ’94*, ACM, pp. 413–414, 1994.

[18] A. B. Arrieta et al., “Explainable Artificial Intelligence (XAI): Concepts, Taxonomies, and Challenges,” *Information Fusion*, vol. 58, pp. 82–115, 2020.

[19] C. Rudin, “Stop Explaining Black Box Machine Learning Models for High Stakes Decisions and Use Interpretable Models Instead,” *Nature Machine Intelligence*, vol. 1, no. 5, pp. 206–215, 2019.

[20] M. T. Ribeiro, S. Singh, and C. Guestrin, “Why Should I Trust You?: Explaining the Predictions of Any Classifier,” in *Proc. KDD ’16*, ACM, pp. 1135–1144, 2016.

[21] F. Doshi-Velez and B. Kim, “Towards a Rigorous Science of Interpretable Machine Learning,” *arXiv:1702.08608*, 2017.

[22] F. Barboza, H. Kimura, and E. Altman, “Machine Learning Models and Bankruptcy Prediction,” *Expert Systems with Applications*, vol. 83, pp. 405–417, 2017.

[23] A. S. Kelton, R. R. Pennington, and B. M. Tuttle, “The Effects of Information Presentation Format on Judgment and Decision Making,” *Int. J. Accounting Information Systems*, vol. 11, no. 1, pp. 1–28, 2010.

[24] A. Savikhin, H. C. Lam, B. Fisher, and D. S. Ebert, “An Experimental Study of Financial Portfolio Selection with Visual Analytics for Decision Support,” in *Proc. HICSS ’11*, IEEE, pp. 1–10, 2011.