

Article

Neuro-Symbolic Frameworks for Corporate Leverage and Debt Maturity: Evidence from Econometric and Machine Learning Models

Omar Shawkey ¹, Taha Mohamed Gaber ¹, Esmail Mohamed ¹, Ahmed Hassanein ^{2,3,*}  and Yara Ibrahim ¹ 

¹ Investment and Finance Department, Faculty of International Business and Humanities, Egypt–Japan University of Science and Technology, Alexandria 21934, Egypt; omar.abdelsamie@ejust.edu.eg (O.S.); taha.hamed@ejust.edu.eg (T.M.G.); esmail.abdalgani@ejust.edu.eg (E.M.); yara.magdy@ejust.edu.eg (Y.I.)

² Accounting and MIS Department, College of Business Administration, Gulf University for Science and Technology, Mubarak Al-Abdullah 32093, Kuwait

³ Department of Accounting, Faculty of Commerce, Mansoura University, Mansoura 35516, Egypt

* Correspondence: hassanein.a@gust.edu.kw

Abstract

Forecasting corporate leverage adjustments remains challenging due to persistent financing behavior, firm heterogeneity, and changing macroeconomic conditions. This study investigates whether increasing model complexity improves the forecasting of corporate leverage adjustment by comparing dynamic econometric models, machine learning algorithms, and a neuro-symbolic artificial intelligence framework. The analysis is based on an unbalanced panel of 39,226 firm-year observations from 3001 publicly listed non-financial firms across 18 countries. The empirical analysis compares Fixed Effects and two-step Difference GMM estimators with regularized regression, gradient boosting, artificial neural networks, and a theory-guided neuro-symbolic framework that incorporates economically meaningful financial constraints through a resampling-based approximation of a differentiable rule-based penalty. Model performance is evaluated using out-of-sample forecasting accuracy measured by the Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), and the coefficient of determination (R^2). The results indicate that corporate leverage exhibits substantial persistence, with estimated adjustment speeds of approximately 29–36% annually. Machine learning algorithms do not improve forecasting accuracy relative to benchmark dynamic econometric models when evaluated out of sample, while incorporating symbolic financial constraints provides only limited additional predictive benefits. These findings suggest that leverage persistence dominates model complexity and that parsimonious dynamic econometric models remain highly effective for forecasting corporate leverage adjustment. The study contributes to the growing literature on explainable artificial intelligence in corporate finance by providing a comprehensive comparison of dynamic econometric, machine learning, and neuro-symbolic approaches within a unified forecasting framework for emerging economies in the MENA region.



Academic Editor: Thanasis Stengos

Received: 29 June 2026

Revised: 10 August 2026

Accepted: 13 August 2026

Published: 17 August 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article

distributed under the terms and

conditions of the [Creative Commons](https://creativecommons.org/licenses/by/4.0/)

[Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

Keywords: capital structure; leverage adjustment; debt maturity; machine learning; neuro-symbolic AI; corporate finance; emerging markets; financial forecasting

JEL Classification: G32; G17; C33; C53; C45; O33

1. Introduction

Corporate leverage adjustment is a dynamic process that reflects interactions among firm characteristics, financing conditions, and macroeconomic environments. A substantial body of corporate finance research has developed theoretical frameworks to explain firms' financing decisions, including the trade-off theory, pecking order theory, and market timing theory (Myers & Majluf, 1984; Frank & Goyal, 2009; Graham & Leary, 2011). Empirical evidence consistently shows that firms adjust gradually toward target capital structures because adjustment is subject to transaction costs, financing frictions, and market imperfections. Consequently, firms exhibit partial adjustment dynamics and persistent deviations from target leverage ratios rather than continuously rebalancing their capital structures (Flannery & Rangan, 2006; Lemmon et al., 2008; Faulkender et al., 2012).

Despite these theoretical advances, accurately forecasting corporate leverage adjustment remains challenging. Leverage exhibits strong persistence, state dependence, and considerable heterogeneity across firms and institutional environments. These characteristics are particularly evident in emerging markets, where bank-centered financial systems, concentrated ownership structures, and volatile macroeconomic conditions influence financing behavior. Under these conditions, firms frequently follow economically motivated financial policies, such as maintaining leverage within sustainable ranges and preserving adequate solvency buffers, rather than continuously adjusting toward the theoretical optimum capital structure (Frank & Goyal, 2009; Graham & Leary, 2011).

These characteristics raise an important challenge for predictive modeling. Traditional econometric models effectively capture dynamic adjustment processes and leverage persistence but generally rely on linear specifications that may overlook complex nonlinear relationships among firm characteristics and macro-financial conditions (Flannery & Rangan, 2006; Wooldridge, 2010). In contrast, machine learning algorithms provide greater flexibility for modeling nonlinear interactions and high-dimensional relationships but often sacrifice interpretability and rarely incorporate established financial theory into the learning process (Gu et al., 2020; Athey & Imbens, 2019; Mullainathan & Spiess, 2017). More recently, neuro-symbolic artificial intelligence has emerged as a promising paradigm that integrates data-driven learning with symbolic financial knowledge through explicit financial constraints, thereby improving model transparency while preserving predictive flexibility (Hitzler et al., 2022). However, empirical evidence remains limited on whether increasingly sophisticated predictive frameworks yield meaningful improvements in forecasting corporate leverage adjustments beyond well-established dynamic econometric models, particularly in emerging markets.

Accordingly, the objective of this study is to evaluate whether increasing model complexity, from dynamic econometric models to machine learning algorithms and neuro-symbolic artificial intelligence, improves the forecasting of corporate leverage adjustment. Using an unbalanced panel of 39,226 firm-year observations from 3001 non-financial publicly listed firms across 18 emerging economies in the Middle East and North Africa region during the period 2010 to 2025, the study compares dynamic panel estimators, regularized regression, tree-based machine learning algorithms, neural networks, and a neuro-symbolic framework within a unified forecasting environment. The proposed neuro-symbolic framework incorporates economically motivated financial rules, including leverage bands and solvency guardrails, via differentiable penalty functions to assess whether embedding financial knowledge improves predictive performance while preserving model interpretability (Hitzler et al., 2022).

The empirical findings provide three important insights. First, corporate leverage exhibits substantial persistence, with estimated adjustment speeds ranging from approximately 30% to 53%, confirming that past financing decisions remain the primary determinant of future leverage (Flannery & Rangan, 2006; Faulkender et al., 2012). Second, machine learning algorithms do not yield statistically significant improvements in out-of-sample forecasting performance compared with conventional dynamic econometric models. Third, incorporating symbolic financial constraints into the predictive framework does not improve forecasting accuracy and, under the implemented calibration, slightly reduces predictive performance. Collectively, these findings suggest that the highly persistent nature of corporate leverage adjustment limits the incremental predictive gains obtained from increasingly complex artificial intelligence models.

This study contributes to the literature in four ways. First, it provides one of the first comprehensive empirical comparisons of dynamic econometric models, machine learning algorithms, and neuro-symbolic artificial intelligence for forecasting corporate leverage adjustment in emerging markets using a common dataset, identical forecasting horizons, and a consistent validation framework. Second, it evaluates whether embedding economically grounded financial knowledge within machine learning models improves forecasting performance beyond conventional data-driven approaches. Third, the findings demonstrate that persistence in leverage adjustment dominates model complexity, indicating that increasingly sophisticated artificial intelligence techniques do not necessarily produce superior forecasts. Finally, the study contributes to the emerging literature on explainable artificial intelligence in finance by demonstrating that although neuro-symbolic models improve model transparency by incorporating economically meaningful financial constraints, these advantages do not necessarily translate into superior predictive accuracy under standard model calibration. Overall, the findings reinforce the importance of dynamic econometric modeling, persistence, and model parsimony in forecasting corporate leverage adjustment.

The remainder of the paper is organized as follows. Section 2 reviews the relevant literature on capital structure dynamics and forecasting methods. Section 3 describes the data and modeling framework. Section 4 presents the empirical results and forecast evaluation. Section 5 discusses the implications of the findings for leverage forecasting and neuro-symbolic financial modeling. Section 6 concludes.

2. Theoretical Background and Literature Review

2.1. Theoretical Foundations of Dynamic Leverage Adjustment

Theories of capital structure seek to explain how firms determine their financing policies under conditions of uncertainty, market imperfections, and financing frictions. The trade-off theory proposes that firms choose their capital structures by balancing the tax advantages of debt financing against the expected costs of financial distress and agency conflicts, implying the existence of an optimal leverage ratio toward which firms gradually adjust (Kraus & Litzenberger, 1973). In contrast, the pecking order theory argues that information asymmetry between managers and external investors leads firms to follow a financing hierarchy, preferring internally generated funds, followed by debt, while issuing equity only when alternative financing sources are exhausted (Myers & Majluf, 1984). The market timing theory offers a complementary perspective by suggesting that observed capital structures reflect managers' attempts to exploit temporary market mispricing rather than maintain a stable target leverage ratio (Baker & Wurgler, 2002). Rather than representing competing explanations, these theories capture distinct mechanisms by which firms make financing decisions under varying economic and institutional conditions, collectively contributing to the observed heterogeneity in corporate leverage.

Subsequent research has extended these static theories by recognizing that adjustment toward target leverage is neither immediate nor costless. Instead, firms rebalance their capital structures gradually because of transaction costs, financing frictions, and institutional constraints. Using a partial adjustment framework, [Flannery and Rangan \(2006\)](#) provide evidence that firms systematically converge toward target leverage over time rather than instantaneously restoring optimal capital structures. [Lemmon et al. \(2008\)](#) further show that a substantial proportion of the variation in leverage reflects persistent firm-specific characteristics, highlighting the importance of dynamic adjustment rather than static optimization. Similarly, [Faulkender et al. \(2012\)](#) demonstrate that financing decisions are strongly influenced by internally generated cash flows, implying that leverage adjustment depends on firms' financial flexibility and available financing capacity.

More recent theoretical developments emphasize that leverage adjustment is inherently state-dependent. Rather than following a constant adjustment speed, firms modify their financing behavior in response to changing macroeconomic conditions, financing constraints, and institutional environments. [Dang et al. \(2012\)](#) report asymmetric adjustment behavior, showing that firms above their target leverage reduce debt more slowly than underleveraged firms increase borrowing. Likewise, [DeMarzo and He \(2021\)](#) demonstrate that limited commitment and refinancing frictions generate persistent deviations from target leverage, implying that financing decisions evolve within economically feasible adjustment ranges rather than around a single optimal point. Recent evidence further suggests that macro-financial conditions substantially influence adjustment dynamics, particularly during periods of monetary tightening and financial uncertainty, reinforcing the importance of modeling leverage as a dynamic and context-dependent process rather than a static financial outcome.

The growing emphasis on dynamic leverage adjustment has shifted the focus of capital structure research from identifying the determinants of leverage toward understanding how firms modify their financing decisions over time. Consequently, persistence, state dependence, and adjustment costs have become central features of contemporary capital structure theory, providing the conceptual foundation for forecasting future leverage rather than merely explaining observed capital structures. This perspective is particularly relevant for emerging markets, where financing frictions, institutional heterogeneity, and macroeconomic volatility amplify adjustment costs and contribute to greater persistence in corporate leverage dynamics.

2.2. Macroeconomic and Institutional Determinants of Dynamic Leverage Adjustment

Although capital structure theories primarily emphasize firm-specific determinants of financing decisions, an increasing body of research demonstrates that macroeconomic conditions and institutional environments also shape leverage adjustment. Firms do not adjust their capital structures in isolation. Instead, financing decisions evolve in response to changes in monetary policy, credit availability, economic uncertainty, and institutional quality, all of which affect the costs and benefits of external financing. Consequently, the speed and direction of leverage adjustment vary across economic regimes rather than remaining constant over time.

Macroeconomic conditions influence leverage adjustment by altering borrowing costs, financial constraints, and firms' access to capital markets. Periods of accommodative monetary policy and abundant liquidity generally reduce financing costs, allowing firms to move more rapidly toward their target capital structures. Conversely, higher interest rates, tighter credit conditions, and heightened economic uncertainty increase adjustment costs and slow the rebalancing process. [Alter and Elekdag \(2020\)](#) show that global financial conditions significantly affect corporate leverage growth in emerging markets, particularly

where domestic financial systems are highly dependent on external capital flows. Similarly, [Ottonello and Winberry \(2020\)](#) demonstrate that firm heterogeneity plays a central role in transmitting monetary policy through financing and investment decisions, implying that firms respond differently to identical macroeconomic shocks according to their financial positions.

Recent studies further suggest that leverage adjustment is regime-dependent rather than uniform across economic environments. Firms accelerate or delay capital structure adjustments depending on macroeconomic conditions, financial flexibility, and financing constraints. Evidence from emerging economies indicates that the speed of leverage adjustment differs substantially between periods of economic expansion and contraction, with financially constrained firms adjusting more slowly due to limited access to external finance. These findings highlight that macroeconomic variables are not merely control variables but fundamental drivers of dynamic capital structure behavior.

Institutional characteristics further influence firms' financing decisions by affecting investor protection, financial market development, and access to external capital. Compared with developed economies, firms operating in emerging markets typically face higher information asymmetry, weaker legal enforcement, less developed capital markets, and greater dependence on bank financing. These institutional characteristics increase financing frictions and contribute to greater persistence in leverage ratios. Recent evidence confirms that institutional quality significantly influences firms' investment and financing behavior, reinforcing the importance of considering country-level characteristics when analyzing capital structure dynamics across emerging economies.

Collectively, these findings demonstrate that corporate leverage adjustment reflects the interaction of firm-specific characteristics, macroeconomic conditions, and institutional environments. This interaction generates substantial heterogeneity in financing behavior across firms and countries, suggesting that accurately forecasting leverage adjustment requires modeling dynamic relationships rather than relying solely on static financial characteristics. The presence of nonlinear responses to changing economic conditions also motivates the use of more flexible predictive methods that capture complex interactions while remaining consistent with established corporate finance theory.

2.3. Machine Learning and Neuro-Symbolic Artificial Intelligence in Corporate Finance

The increasing availability of large financial datasets and advances in computational methods have accelerated the adoption of machine learning techniques in corporate finance. Unlike conventional econometric models, which generally require predefined functional forms and rely on linear assumptions, machine learning algorithms can identify complex nonlinear relationships, high-order interactions, and hidden patterns directly from data. These capabilities have led to growing applications of machine learning in financial forecasting, including bankruptcy prediction, credit risk assessment, asset pricing, earnings forecasting, and corporate financial decision-making ([Gu et al., 2020](#); [Mullainathan & Spiess, 2017](#)).

Despite these advantages, evidence regarding the superiority of machine learning over traditional econometric approaches remains mixed. Although flexible algorithms often improve predictive accuracy in highly nonlinear settings, their performance depends heavily on data quality, sample size, feature engineering, and model calibration. Furthermore, many machine learning models operate as "black-box" systems, making it difficult to interpret the economic rationale underlying their predictions or to verify whether the generated forecasts are consistent with established financial theory. These limitations have become increasingly important in finance, where transparency, explainability, and theoretical consistency are essential for both academic research and practical decision-making. Recent reviews of explainable artificial intelligence emphasize that improving interpretabil-

ity while maintaining predictive performance remains one of the principal challenges in financial applications of artificial intelligence.

To address these limitations, recent research has increasingly focused on explainable and theory-informed artificial intelligence. Rather than relying exclusively on statistical learning, these approaches incorporate domain knowledge into the learning process to improve transparency and model reliability. Neuro-symbolic artificial intelligence is one of the most promising developments in this direction, combining the pattern-recognition capabilities of neural networks with symbolic reasoning grounded in explicit knowledge and logical constraints (Hitzler et al., 2022). Recent surveys indicate that neuro-symbolic systems offer greater interpretability, improved reasoning capability, and stronger consistency with domain knowledge than purely data-driven models, although their application in finance remains limited.

The integration of financial theory into machine learning is particularly relevant for corporate finance because well-established economic principles constrain financing decisions. Firms generally operate within economically feasible ranges of leverage, respond to financing frictions, and adjust their capital structures gradually rather than arbitrarily. Embedding these theoretical principles within predictive models may improve model transparency while reducing economically implausible predictions. However, empirical evidence evaluating whether theory-guided machine learning actually outperforms conventional econometric models in forecasting corporate leverage adjustment remains scarce.

This review reveals three important research gaps. First, the capital structure literature has predominantly focused on explaining the determinants and speed of leverage adjustment rather than evaluating the comparative forecasting performance of competing modeling approaches. Second, although machine learning has become increasingly popular in financial prediction, relatively few studies directly compare dynamic econometric models, machine learning algorithms, and neuro-symbolic artificial intelligence within a unified forecasting framework using identical datasets and validation procedures. Third, evidence from emerging markets remains limited despite their distinctive institutional characteristics, financing constraints, and macroeconomic environments. Consequently, it remains unclear whether increasing model complexity generates meaningful improvements in forecasting corporate leverage adjustment or whether conventional dynamic econometric models continue to provide comparable or superior predictive performance. Addressing these gaps constitutes the primary objective of the present study.

2.4. Research Gap and Hypothesis Development

The preceding review demonstrates that substantial progress has been made in understanding corporate capital structure decisions. Classical theories, including the trade-off, pecking order, market timing, and dynamic capital structure theories, explain how firms balance debt and equity financing while gradually adjusting toward their target leverage due to adjustment costs and financing frictions (Frank & Goyal, 2009; Flannery & Rangan, 2006). However, no single theory fully explains financing decisions across different institutional and economic settings. A recent systematic review confirms that the predictions of existing capital structure theories remain complementary rather than conclusive, leaving Myers' (1984) fundamental question regarding firms' financing decisions unresolved (Missaoui & Brahmi, 2025).

Recent evidence suggests that governance quality, sustainability practices, and debt maturity also shape corporate financing decisions. Board characteristics influence both leverage and debt maturity decisions (Hamida et al., 2025), corporate governance moderates the relationship between leverage and firm value (Tulcanaza-Prieto et al., 2024), ESG performance reduces leverage manipulation by improving transparency and reducing information asymmetry (Mo et al., 2024), and green credit policies affect corporate risk-taking through debt maturity mismatch (Wang et al., 2025). Furthermore, advances in machine learning have introduced flexible techniques capable of capturing nonlinear relationships that conventional econometric models may overlook, although these methods often lack economic interpretability.

Despite these advances, several research gaps remain. First, most studies focus on explaining the determinants of leverage rather than comparing the forecasting performance of alternative modeling approaches. Second, few studies evaluate dynamic econometric models, machine learning algorithms, and neuro-symbolic artificial intelligence within a unified framework using identical data and validation procedures. Third, governance, ESG, debt maturity, and macroeconomic factors have rarely been incorporated simultaneously into integrated predictive models. Finally, evidence from emerging MENA economies remains relatively limited despite their unique institutional and financial characteristics.

To address these gaps, this study compares dynamic econometric models, machine learning algorithms, and a theory-guided neuro-symbolic artificial intelligence framework using a common sample of publicly listed non-financial firms from emerging MENA economies. Unlike previous studies, the proposed framework integrates firm-specific characteristics, macroeconomic conditions, and symbolic financial constraints within a unified predictive architecture. The study therefore evaluates whether increasing model complexity and embedding financial knowledge improve out-of-sample forecasting performance while maintaining consistency with established corporate finance theory.

Based on the theoretical and empirical literature, the following hypotheses are proposed.

- H1.** *Corporate leverage exhibits significant dynamic adjustment toward its target level.*
- H2.** *Firm-specific characteristics, corporate governance, ESG performance, debt maturity, and macroeconomic conditions significantly influence corporate leverage adjustment.*
- H3.** *Machine learning models provide superior out-of-sample forecasting performance compared with traditional dynamic econometric models.*
- H4.** *Incorporating symbolic financial knowledge into machine learning models improves forecasting performance relative to conventional machine learning models.*

Figure 1 presents the conceptual framework that links firm-specific characteristics, governance, ESG, debt maturity, and macroeconomic conditions to corporate leverage adjustment, and compares the forecasting performance of dynamic econometric models, conventional machine learning algorithms, and neuro-symbolic artificial intelligence.




Figure 1. Research Model.

3. Methodology

The study applies a hybrid methodology combining panel econometrics, statistical learning, and neural-symbolic modeling to analyze neuro-symbolic capital structure. It integrates causal inference with predictive modeling to balance interpretability and non-linear learning (Athey & Imbens, 2019; Mullainathan & Spiess, 2017). Using a design science approach, the framework is treated as a computational artifact that embeds financial theory into a predictive system, following stages of problem identification, design, demonstration, and evaluation (Hevner et al., 2004; Peffers et al., 2007) as presented in Figure 2. This method addresses the limits of linear panel models and enhances the ability to capture complex financial dynamics (Dang et al., 2012; Flannery & Rangan, 2006; Athey & Imbens, 2019).

The model integrates firm-level and macroeconomic data, combining econometric estimation with statistical learning methods such as regularized regression and boosting, and embedding financial constraints as differentiable penalties within the learning objective (Wooldridge, 2010; Hitzler et al., 2022). The framework is applied to panel data from MENA firms to generate predictions of leverage adjustment and debt maturity under varying macro-financial conditions. Model evaluation includes both econometric diagnostics and out-of-sample validation, with robustness checks across alternative specifications (Jungherr et al., 2024). This dual evaluation framework ensures statistical reliability and predictive performance, while demonstrating how theory-based constraints can be integrated with data-driven models in financial forecasting contexts (Hevner et al., 2004; Peffers et al., 2007).




Figure 2. Design Science Research Methodology.

To enhance methodological transparency and facilitate reproducibility, Algorithm 1 summarizes the complete computational workflow employed in this study. The proposed framework integrates traditional econometric approaches with machine learning techniques and a neuro-symbolic learning framework within a rolling-window forecasting design. The procedure begins with data acquisition and preprocessing, followed by variable construction and econometric estimation, and concludes with the implementation of benchmark machine learning models. Symbolic financial constraints are subsequently incorporated to refine neural network predictions, and the predictive performance of all competing models is evaluated using multiple forecasting accuracy measures. This algorithm provides a concise overview of the analytical pipeline implemented in the study and serves as a roadmap for the empirical methodology in Algorithm 1, which is described in detail in Appendix A. The procedure integrates dynamic panel econometric models, machine learning algorithms, and symbolic financial constraints within a rolling-window forecasting framework to generate one-step-ahead leverage predictions.

Algorithm 1. Computational workflow of the proposed neuro-symbolic forecasting framework.

BEGIN

1. Data Acquisition and Integrity Verification

- 1.1 Load raw firm-level panel from Refinitiv extract
- 1.2 Remove non-firm placeholder rows (no ticker/ISIN/financials populated)
- 1.3 Set panel key to WORLDScope IDENTIFIER (FIRM_ID),
- 1.4 VERIFY: assert zero duplicate (FIRM_ID, YEAR) pairs

2. Variable Construction

- 2.1 Construct leverage measures: Debt/Assets, Market Leverage, Net Debt/EBIT
- 2.2 Construct profitability (ROA, ROE), size (Log Assets), liquidity (Cash/Assets, Current Ratio), tangibility, growth (Market-to-Book, Sales Growth), maturity proxy (ST Debt Share), and interest coverage
- 2.3 Construct one-period-ahead targets: $\text{delta_Lev}_{(i,t+1)}$, $\text{delta_Mat}_{(i,t+1)}$
- 2.4 Construct one-period lagged macroeconomic variables (Inflation, Volatility)

3. Missing Data and Outlier Treatment

- 3.1 Impute missing ratio values hierarchically:
 - (a) linear interpolation within firm (ordered by year)
 - (b) nation-year median for any residual gap
 - (c) overall column median (fallback, rarely invoked)
- 3.2 Winsorize each constructed ratio at the 1st/99th percentile
- 3.3 Standardize (z-score) Net Debt/EBIT, Debt/Assets, and Market Leverage within nation-year cells (for structural factor construction only)
- 3.4 Extract first principal component of the three standardized leverage proxies $\rightarrow \text{Lev}_{(i,t)}$ [structural leverage factor]

4. Descriptive and Correlation Diagnostics

- 4.1 Compute summary statistics (mean, median, SD, P25, P75) for all study variables
- 4.2 Compute pairwise Pearson correlations with p -values
- 4.3 Compute Variance Inflation Factors to rule out multicollinearity

5. Panel Econometric Estimation

- 5.1 Two-way Fixed Effects (dynamic partial-adjustment specification):
 - (a) demean data by firm, by year, and add back the grand mean (exact two-way within-transformation)
 - (b) estimate by OLS on the demeaned data
 - (c) compute firm-clustered (sandwich) standard errors
 - (d) repeat for firm-only FE and firm + year FE specifications
- 5.2 Two-Step Difference GMM (Arellano-Bond):
 - (a) first-difference the model to eliminate the firm fixed effect
 - (b) instrument the differenced lagged dependent variable; instrument other regressors with their own first differences
 - (c) estimate one-step GMM, form the optimal weighting matrix from one-step residuals, re-estimate (two-step)
 - (d) compute two-step standard errors WITH the N-scaling correction validated by Monte Carlo simulation
 - (e) repeat with a restricted lag-2-only instrument set
 - (f) compute AR(1), AR(2) serial-correlation tests and the Sargan/Hansen overidentification test
- 5.3 Compute Speed of Adjustment = $(1 - \text{lagged-leverage coefficient}) \times 100$ for each specification (Firm FE, Firm + Year FE, full GMM, restricted GMM)

Algorithm 1. *Cont.***6. Symbolic Rule Layer Construction****6.1** Define five differentiable rule-violation functions $g_r(x_{it})$:

- (a) Leverage Band: deviation from nation-year median beyond $\kappa \times$ rolling firm-level volatility
- (b) Coverage Guardrail: interest coverage below its 10th percentile
- (c) Maturity Ladder: ST debt shares outside the nation-year IQR band
- (d) Refinancing-under-Stress: low ST debt shares when lagged regional volatility exceeds its 90th percentile
- (e) Asset-Liability Matching: tangibility/maturity mismatch beyond its in-sample mean $\pm$ one SD

6.2 Smooth each violation with a softplus relaxation:

$\phi_r = (1/\beta) \times \log(1 + \exp(\beta \times g_r))$, $\beta = 10$
 [replaces the non-differentiable hinge $\max(0, g_r)$]

6.3 Mask any rule to $\phi_r = 0$ (and flag “missing”) wherever a required input, or ϕ_r itself, is non-finite**6.4** Aggregate: $\Phi_{(i,t)} = \text{sum over } r \text{ of } [w_r \times \phi_r]$, equal weights $w_r = 1$ **6.5** Define Rule Adherence Rate per observation:

$\text{RAR}_{(i,t)} = (\text{count of evaluable rules with } \phi_r \leq \ln(2)/\beta) / (\text{count of evaluable rules})$
 [threshold $\ln(2)/\beta$ is the value of ϕ_r exactly AT the true rule boundary $g_r = 0$]

7. Rolling-Window Walk-Forward Forecasting Framework

FOR each fold $k = 0, 1, 2, \dots$ DO

7.1 Training window: $\text{YEAR} \leq (2016 + k)$ **7.2** Validation year: $(2017 + k)$ **7.3** Test year: $(2018 + k)$ [stop once test year > 2024 , since 2025 is the panel’s final year and has no subsequent period to build a one-period-ahead target against]**7.4** Standardize predictors using TRAINING-FOLD statistics only (scaler fit on train, applied to validation/test—no leakage)**8. Benchmark and Neuro-Symbolic Model Training (within each fold)****8.1** Train LASSO ($\alpha = 0.001$)**8.2** Train Ridge ($\alpha = 1.0$)**8.3** Train Gradient Boosting Machine (200 trees, depth 3, $\text{lr} = 0.05$)**8.4** Train feedforward Neural Network (2 hidden layers [64, 32], ReLU, Adam optimizer $\text{lr} = 1 \times 10^{-3}$, $\text{L2} = 1 \times 10^{-4}$, batch = 256, early stopping on a 15% validation split, patience 15 epochs)**8.5** Train Neuro-Symbolic variant: resample the training fold with probability weights $w_i = 1/(1 + \lambda \times \Phi_i)$, then fit the same Neural Network architecture on the resampled data [proxy for a joint gradient penalty—see Verification Note C]

Algorithm 1. *Cont.***9. Forecast Generation and Symbolic Projection**

- 9.1 Generate one-period-ahead $\Delta \text{Lev}_{\text{hat}}$, $\Delta \text{Mat}_{\text{hat}}$ for each benchmark model on the held-out test fold
- 9.2 Project next-period state: $\text{Debt}/\text{Assets}_{\text{hat}} = \text{Debt}/\text{Assets} + \Delta \text{Lev}_{\text{hat}}$, $\text{ST_Share}_{\text{hat}} = \text{ST_Share} + \Delta \text{Mat}_{\text{hat}}$
- 9.3 Re-evaluate the Leverage Band and Maturity Ladder rules against the PROJECTED state (not the observed state) to obtain a prediction-dependent Φ and RAR for this fold/model/ λ [this step was found to be missing in an earlier version—see Verification Note D]

10. Model Evaluation

- 10.1 Compute RMSE, MAE, R-squared for every model on pooled out-of-fold predictions
- 10.2 Compute Diebold-Mariano test (Newey-West HAC variance) for each benchmark model against the Neural Network
- 10.3 Sweep λ over $\{0, 0.2, 0.4, 0.6, 0.8, 1.0\}$; record MSE and prediction-dependent RAR at each value
- 10.4 Store per-fold results

11. NSD Composite Index Construction and Validation

- 11.1 Combine five globally standardized components: $\text{Lev}_{(i,t)}$ [3.1], $\text{Mat}_{(i,t)} = 1 - \text{ST_Share}_{(i,t)}$, $\Delta \text{Lev}_{\text{hat}_{(i,t+1)}}$, $\Delta \text{Mat}_{\text{hat}_{(i,t+1)}}$ [walk-forward, no look-ahead], $\Phi_{(i,t)}$
- 11.2 NSD_equal: unweighted sum of the four positive components minus standardized Φ
- 11.3 NSD_pca: first principal component across the same five components (robustness weighting)
- 11.4 Standardize the aggregated index within nation-year cells
- 11.5 Redundancy check: correlate NSD against standalone Debt/Assets and Market Leverage
- 11.6 Incremental explanatory power: nested two-way FE models predicting one-period-ahead ROA, with and without NSD
- 11.7 Variance decomposition: sequential R-squared as structural, adjustment, and rule-adherence blocks are added

12. Ablation Analysis

- 12.1 FOR each of the five symbolic rules DO
 - remove the rule, recompute Φ and RAR on the remaining four, refit the model at $\lambda = 1$, record RAR/MSE impact
 END FOR

13. Model Comparison and Reporting

- 13.1 Compare benchmark forecasting models
- 13.2 Compare Linear vs. Neural vs. Neuro-Symbolic RAR/MSE,
- 13.3 Report econometric estimates, forecast performance, symbolic diagnostics, and NSD construct validation

END

3.1. Variable Specification

The model examines how firms adjust leverage and debt maturity using firm-level financial data and macroeconomic conditions across countries. Variables are selected based on capital structure and debt maturity theories, as well as econometric criteria such as identification, stability, and comparability. The variables are shown in Tables 1 and 2. Firm-level variables capture leverage, liquidity, profitability, growth opportunities, and collateral. These reflect the main drivers of financing decisions. Table 1 provides their definitions and theoretical rationale.

Table 1. Firm-Level Variables Definition.

Category	Variable	Definition	Reference
Capital Structure	Net Debt/EBITDA	(Total Debt – Cash)/EBITDA	(Frank & Goyal, 2009)
	Debt/Assets	Total Debt/Total Assets	(Rajan & Zingales, 1995)
	Market Leverage	Debt/(Debt + Market Capitalization)	(Baker & Wurgler, 2002)
Maturity Structure	Average Debt Maturity	Weighted average maturity of debt obligations	(Barclay & Smith, 1995)
	% Debt Due $\leq$ 3 Years	Short-term debt share	(Almeida et al., 2011)
Coverage & Liquidity	Interest Coverage	EBIT/Interest Expense	(Graham & Leary, 2011)
	Cash/Assets	Cash & Equivalents/Total Assets	(Opler et al., 1999)
	Current Ratio	Current Assets/Current Liabilities	(Graham & Leary, 2011)
Profitability & Size	ROA	Net Income/Total Assets	(Myers & Majluf, 1984)
	ROE	Net Income/Equity	(Myers & Majluf, 1984)
	Log (Assets)	Natural logarithm of total assets	(Rajan & Zingales, 1995)
Growth & Tangibility	Market-to-Book	Market Value/Book Value	(Baker & Wurgler, 2002)
	Sales Growth	Year-over-year revenue growth	(Myers & Majluf, 1984)
	Tangibility	Net Fixed Assets/Total Assets	(Frank & Goyal, 2009)

Table 2. Macroeconomic Variables.

Variable	Definition	Rationale
Credit Spread	Lending–deposit spread	Captures credit market frictions and financing cost premia
Term Spread	Lending rate–deposit rate	Reflects yield curve conditions and monetary stance
Policy Rate	Central bank lending rate	Monetary policy tightness affecting borrowing cost
Inflation	Consumer price inflation	Macroeconomic stability and real debt burden
Regional Volatility Index	Regional stock volatility index	Financial uncertainty and risk aversion proxy

Firm-level variables are winsorized at the 1st and 99th percentiles to reduce the influence of outliers. To ensure cross-country comparability, variables are standardized within country-year cells using z-score normalization. Multicollinearity is assessed using variance inflation factors, and highly correlated leverage measures are combined using principal component analysis to improve stability. Stationarity is evaluated using panel unit root tests, with first-difference transformations applied in robustness checks where necessary. Missing data are handled using interpolation for macro variables and exclusion thresholds for firms with insufficient time-series observations. Macroeconomic variables are included to capture financing conditions and are measured at the country-year level. These variables account for cost-of-capital effects, refinancing incentives, and macro-financial volatility, which are particularly relevant in emerging markets. Standard errors are clustered to address cross-sectional dependence, and lagged macro variables are used to mitigate simultaneity bias.

3.2. Sample Description

The empirical analysis is based on an unbalanced panel comprising 39,226 firm-year observations from 3001 non-financial publicly listed firms across 18 Middle East and North African (MENA) countries over the period 2010–2025. The sample provides extensive cross-sectional and temporal variation, making it well-suited for examining corporate financing decisions for developing predictive models using conventional econometric techniques, machine learning algorithms, and the proposed neuro-symbolic framework.

The initial sample was extracted from Refinitiv Workspace and included all publicly listed non-financial firms with available financial statements, market, and macroeconomic data. Financial institutions, including banks, insurance companies, and other regulated financial firms, were excluded because their capital structures and financing decisions are subject to industry-specific regulatory requirements that differ substantially from those of non-financial firms. The dataset was subsequently screened for duplicate observations, yielding a final panel with no duplicate records or duplicate firm-year observations.

Financial variables were obtained directly from Refinitiv Workspace, which provides standardized and harmonized accounting and market information across reporting jurisdictions. Consequently, no additional currency conversion or accounting-standard harmonization procedures were required. Missing values were addressed using a hierarchical imputation procedure that combined firm-level temporal imputation with country-level median substitution, while continuous financial variables were winsorized at the 1st and 99th percentiles to mitigate the influence of extreme observations. Following the data-cleaning procedures, the final dataset contained only a negligible number of residual missing values (80 observations, representing less than 0.01% of the sample), thereby ensuring a high-quality panel suitable for econometric estimation and machine learning applications.

The final sample spans 16 years (2010–2025) and includes firms from 19 MENA countries, although the distribution of observations varies considerably across countries because of differences in stock market size, listing activity, and data availability. The panel is unbalanced, with firms contributing between 1 and 16 annual observations, an average of 13.07 observations per firm, and a median of 16 observations, reflecting firm entry, exit, delisting, and differences in historical data availability over the study period.

Table 3 [Panels A and B] summarizes the geographical composition and overall characteristics of the sample. Israel and Türkiye together account for approximately 42.6% of all firm-year observations, followed by Saudi Arabia, Egypt, Jordan, Kuwait, and the United Arab Emirates. The broad geographical coverage and large sample size provide a comprehensive representation of listed non-financial firms across the MENA region and establish a robust empirical setting for evaluating the performance of econometric, machine learning, and neuro-symbolic approaches under diverse institutional and macroeconomic environments.

3.3. Construction of Neuro-Symbolic Debt (NSD)

The Neuro-Symbolic Debt (NSD) construct is a theory-based index that measures corporate debt positioning by combining current balance sheet conditions with dynamic adjustment behavior. It is grounded in trade-off and partial adjustment theories, which view leverage as moving toward target levels over time (Flannery & Rangan, 2006; Frank & Goyal, 2009). NSD is designed as a descriptive and predictive measure of how closely firms align with financially sound debt practices, rather than a causal latent variable (Rajan & Zingales, 1995). To avoid circularity, NSD has two components. The structural component captures the firm's current debt position using factors like profitability, size, asset tangibility, and growth opportunities (Frank & Goyal, 2009; Rajan & Zingales, 1995). The adjustment component reflects forward-looking changes generated by a neural-symbolic model that

embeds financial rules into predictions (Hitzler et al., 2022). Separating these components helps reduce endogeneity and ensures more reliable modeling of debt dynamics (Flannery & Rangan, 2006).

Table 3. Panel A: Sample Distribution by Country. Panel B: Sample Characteristics.

Panel A			
Country	Firms	Firm-Year Observations	Percentage (%)
Israel	661	8388	21.38
Türkiye	657	8329	21.23
Saudi Arabia	386	4061	10.35
Egypt	244	3386	8.63
Jordan	188	2913	7.43
Kuwait	139	2149	5.48
United Arab Emirates	166	2067	5.27
Oman	108	1586	4.04
Iraq	97	1171	2.99
Morocco	73	1092	2.78
Tunisia	71	1042	2.66
Qatar	53	759	1.94
Bahrain	42	639	1.63
Malta	46	601	1.53
Palestine	37	525	1.34
Syria	21	319	0.81
Lebanon	11	167	0.43
Sudan	1	16	0.04
Total	3001	39,226	100.00

Panel B	
Characteristic	Value
Firm-year observations	39,226
Number of firms	3001
Number of countries	18
Sample period	2010–2025
Average observations per firm	13.07
Median observations per firm	16.0
Minimum observations per firm	1
Maximum observations per firm	16
Duplicate firm-year observations	0

3.3.1. Structural Debt Component

The structural component of the Neuro-Symbolic Debt (NSD) index captures the firm's contemporaneous debt position using multiple complementary measures of leverage and debt structure. Consistent with empirical capital structure research, leverage is represented by a composite factor rather than a single accounting ratio, recognizing that alternative leverage measures capture different dimensions of corporate financing decisions (Frank & Goyal, 2009; Rajan & Zingales, 1995). Let $Lev_{i,t}$ denote the first principal component extracted from the nation-year standardized values of $NetDebt/EBIT_{i,t}$, $Debt/Assets_{i,t}$, and $MarketLeverage_{i,t}$:

$$Lev_{i,t} = PC_1 \left\{ z, (NetDebt/EBIT_{i,t}), z(Debt/Assets_{i,t}), z(MarketLeverage_{i,t}) \right\}. \quad (1)$$

Principal component analysis reduces multicollinearity among the underlying leverage measures while preserving the dominant variation in firms' financing structures, thereby providing a parsimonious representation of debt intensity (Jolliffe, 2002). Standardization within nation-year cells further removes systematic cross-country differences arising from institutional and macroeconomic environments prior to component extraction, thereby improving comparability across firms operating in different financial systems (Rajan & Zingales, 1995).

Debt maturity is represented by the long-term debt share, defined as the complement of the short-term debt ratio:

$$Mat_{i,t} = 1 - STDebtShare_{i,t}.$$

This transformation ensures that higher values correspond to longer average debt maturity and lower refinancing exposure, consistent with established theories of debt maturity and financial flexibility (Barclay & Smith, 1995). Together, $Lev_{i,t}$ and $Mat_{i,t}$ describe the firm's contemporaneous debt configuration without incorporating information on future financing adjustments, thereby maintaining a clear distinction between the structural and dynamic components of the NSD index.

3.3.2. Dynamic Adjustment Component

The dynamic component of the Neuro-Symbolic Debt (NSD) index captures firms' expected adjustments toward future debt positions through one-period-ahead predictions of leverage and debt maturity. Let $\widehat{\Delta Lev}_{i,t+1}$ and $\widehat{\Delta Mat}_{i,t+1}$ denote the predicted changes in leverage and maturity generated by the neuro-symbolic learning architecture:

$$\widehat{\Delta Lev}_{i,t+1}, \widehat{\Delta Mat}_{i,t+1}.$$

The predictive model is estimated using a genuine walk-forward validation protocol that preserves the chronological ordering of the panel data. For each rolling estimation window, the neural network is trained exclusively on observations available up to the corresponding training cutoff year and subsequently generates predictions only for the immediately following test year. This forecasting framework prevents look-ahead bias by ensuring that future information is never incorporated into model estimation, thereby replicating the information set available to decision makers at each point in time.

The predictive architecture uses lagged firm-specific financial characteristics, together with macroeconomic variables, to estimate future debt adjustments, consistent with dynamic capital structure models in which firms gradually move toward target leverage due to adjustment costs and financing frictions (Flannery & Rangan, 2006; Lemmon et al., 2008). Incorporating constraint-aware neural-symbolic learning further enables economically meaningful predictions by integrating data-driven forecasting with structured financial reasoning, consistent with emerging research on interpretable neural-symbolic systems (Hitzler et al., 2022). By separating contemporaneous debt measures from forward-looking adjustment predictions, the dynamic component minimizes mechanical overlap with the structural component and reduces potential circularity in constructing the NSD index (Flannery & Rangan, 2006).

3.3.3. Symbolic Rule Layer

The symbolic rule layer incorporates economically grounded financial constraints into the predictive framework by penalizing debt positions and predicted adjustments that deviate from established principles of capital structure, solvency, and debt maturity management (Barclay & Smith, 1995; Myers & Majluf, 1984; Opler et al., 1999). Each

symbolic rule r is represented by a differentiable violation function $g_r(X_{i,t}, \hat{y}_{i,t})$, where $X_{i,t}$ denotes the observed firm characteristics and $\hat{y}_{i,t}$ represents the predicted debt adjustment. Rather than applying a non-differentiable hinge penalty, rule violations are transformed using a smooth softplus relaxation:

$$\phi_r(X_{i,t}, \hat{y}_{i,t}) = \frac{1}{\beta} \log(1 + \exp(\beta g_r(X_{i,t}, \hat{y}_{i,t}))), \quad (2)$$

where β controls the smoothness of the approximation. This formulation preserves differentiability throughout the optimization process while converging to the conventional hinge penalty as β increases, making it suitable for gradient-based neural network training (Hitzler et al., 2022).

The symbolic layer comprises five complementary financial rules governing leverage dispersion, interest coverage, debt maturity structure, refinancing exposure under periods of elevated market stress, and asset-liability matching. These rules are calibrated using the empirical distribution of the estimation sample, rolling firm-level volatility, nation-year distributional thresholds, and in-sample benchmark values to ensure that constraint evaluation reflects prevailing financial conditions rather than arbitrary cut-off values. Missing or non-finite inputs are excluded from rule evaluation to prevent invalid observations from contributing artificial penalties.

The aggregate symbolic penalty is computed as the equally weighted sum of the individual rule violations:

$$\Phi_{i,t} = \sum_{r=1}^R w_r \phi_r(X_{i,t}, \hat{y}_{i,t}), \quad (3)$$

where w_r denotes the weight assigned to rule r . Rule adherence is additionally summarized using the Rule Adherence Rate (RAR), which measures the proportion of applicable symbolic rules satisfied by each observation. Together, the penalty term and adherence measure quantify the extent to which observed debt positions and predicted financing adjustments remain consistent with economically grounded financing principles, thereby embedding structured financial reasoning within the learning architecture (Hitzler et al., 2022).

3.3.4. Composite Index Construction

The Neuro-Symbolic Debt (NSD) index integrates contemporaneous debt positioning, predicted financing adjustments, and adherence to symbolic rules into a unified measure of corporate debt management. To ensure comparability across heterogeneous variables, each of the five underlying components—structural leverage ($Lev_{i,t}$), debt maturity ($Mat_{i,t}$), predicted leverage adjustment ($\widehat{\Delta Lev_{i,t+1}}$), predicted maturity adjustment ($\widehat{\Delta Mat_{i,t+1}}$), and the aggregate symbolic penalty ($\Phi_{i,t}$)—is first variance-normalized using standardized scores, thereby preventing any single component from dominating the composite index because of differences in measurement scale (Jolliffe, 2002).

The normalized components are then aggregated using two complementary weighting approaches. The primary specification constructs the NSD index as an equally weighted composite of the four positive components, offset by the standardized symbolic penalty, following the common practice of assigning equal weights when no theoretical or empirical justification exists for differential weighting (Frank & Goyal, 2009). As a robustness assessment, an alternative specification is obtained by applying principal component analysis to the same standardized variables, allowing the underlying covariance structure of the data to determine the relative contribution of each component (Jolliffe, 2002). The principal component is oriented so that higher index values consistently represent stronger debt positioning.

Following aggregation, the composite index is standardized within nation-year cells to remove systematic cross-country differences arising from macroeconomic and institutional conditions, thereby improving international comparability (Rajan & Zingales, 1995). By combining structural debt characteristics, forward-looking financing adjustments, and symbolic rule adherence only after each component has been independently constructed and standardized, the NSD index avoids mechanical circularity while providing a comprehensive measure of firms' debt positioning that extends beyond conventional leverage ratios.

3.3.5. Construct Validation

Construct validation is conducted to assess whether the Neuro-Symbolic Debt (NSD) index captures information beyond conventional leverage measures and provides a meaningful representation of firms' debt positioning. Consistent with empirical corporate finance practice, validation focuses on the incremental informational content of the composite construct rather than treating NSD as a latent causal variable (Flannery & Rangan, 2006; Lemmon et al., 2008). Three complementary validation procedures are employed.

First, correlation diagnostics are used to examine the degree of redundancy between the NSD index and conventional leverage proxies, including Debt/Assets and Market Leverage. Moderate rather than near-perfect correlations indicate that the composite index captures additional dimensions of debt management beyond static leverage measures (Frank & Goyal, 2009). Second, incremental explanatory power is evaluated by comparing nested regression models with and without the NSD index, following established model comparison procedures in the corporate finance literature (Flannery & Rangan, 2006). Improvements in model fit and predictive performance are used to assess whether the composite index contributes explanatory information beyond traditional debt measures.

Finally, variance decomposition is performed to evaluate the independent contribution of the structural, dynamic adjustment, and symbolic rule components to the overall construct. This analysis confirms that each component contributes unique variation after accounting for the information contained in the remaining components, thereby reducing redundancy within the composite index. By integrating contemporaneous debt characteristics, forward-looking financing adjustments, and theory-driven rule adherence within a unified framework, the NSD index operationalizes debt positioning as a multidimensional construct consistent with hybrid neural-symbolic learning and established theories of corporate capital structure (Frank & Goyal, 2009; Hitzler et al., 2022).

3.4. Symbolic Logic Layer ($\mathfrak{R}$)

The symbolic logic layer incorporates economically grounded financing constraints within the neural learning framework by encoding established principles of capital structure, financial flexibility, debt maturity management, and refinancing risk into a set of differentiable symbolic rules. Rather than assuming unconstrained optimization, the framework embeds bounded rationality by allowing predictive outputs to be evaluated against empirically calibrated financial constraints derived from observed financing behavior (Myers & Majluf, 1984; Opler et al., 1999; Barclay & Smith, 1995; Flannery & Rangan, 2006). This hybrid formulation enables the model to combine data-driven learning with domain knowledge while preserving the flexibility required to capture heterogeneous financing decisions across firms (Frank & Goyal, 2009; Hitzler et al., 2022).

Five symbolic rules are incorporated into the learning architecture, representing complementary dimensions of corporate debt management. The leverage band rule constrains debt positions relative to firm-specific risk and prevailing market conditions. The coverage guardrail limits excessive debt accumulation when debt-servicing capacity deteriorates.

The maturity ladder rule penalizes abnormal concentrations of short-term debt that increase refinancing exposure. The refinancing-under-stress rule incorporates macro-financial conditions by increasing penalties during periods of elevated market uncertainty. Finally, the asset-liability matching rule evaluates the consistency between asset tangibility and debt maturity structure, reflecting established duration-matching principles in corporate finance (Barclay & Smith, 1995; Myers & Majluf, 1984). Collectively, these rules provide an economically interpretable mechanism for evaluating whether predicted financing decisions remain consistent with theoretically grounded financing practices.

3.4.1. Empirical Calibration of Rule Parameters

To avoid arbitrary parameter selection, all symbolic rules are calibrated using empirical characteristics of the estimation sample rather than predetermined constants. The leverage band is defined relative to the nation-year median Debt/Assets ratio and firm-specific leverage volatility, allowing acceptable leverage ranges to vary across firms and macroeconomic environments. Coverage thresholds are determined from the empirical distribution of interest coverage ratios, while maturity constraints are calibrated using the interquartile range of short-term debt shares within each nation-year to account for institutional differences in debt markets (Barclay & Smith, 1995; Frank & Goyal, 2009). The refinancing stress rule is activated only when lagged volatility exceeds empirically determined stress thresholds, whereas the asset-liability matching rule evaluates deviations from the observed relationship between asset tangibility and debt maturity using sample-based benchmark statistics. This empirical calibration strategy anchors all symbolic constraints in observed financing behavior, improving economic realism while avoiding reliance on subjective threshold selection (Flannery & Rangan, 2006; Lemmon et al., 2008).

3.4.2. Differentiable Penalty Formulation

Each symbolic rule r is represented by a violation function $g_r(X_{i,t}, \hat{y}_{i,t})$ that measures departures from the corresponding financial constraint. To maintain compatibility with gradient-based optimization, violations are transformed using a smooth softplus relaxation rather than a non-differentiable hinge function:

$$\phi_r(X_{i,t}, \hat{y}_{i,t}) = \frac{1}{\beta} \log(1 + \exp(\beta g_r(X_{i,t}, \hat{y}_{i,t}))), \quad (4)$$

where β controls the smoothness of the approximation. The softplus formulation remains differentiable over the entire parameter space while closely approximating the conventional hinge penalty for sufficiently large values of β , thereby supporting efficient optimization within the neural learning framework (Hitzler et al., 2022). Individual rule violations are aggregated using equal weights to form the overall symbolic penalty, whereas observations containing non-finite rule evaluations are excluded from the corresponding rule calculation to prevent invalid numerical values from influencing optimization. In addition to the aggregate penalty, a Rule Adherence Rate (RAR) is computed as the proportion of applicable symbolic rules satisfied by each observation, providing an interpretable measure of compliance with the embedded financial principles.

3.4.3. Sensitivity and Stability Analysis

The robustness of the symbolic logic layer is evaluated using a structured sensitivity analysis designed to assess the contribution of individual symbolic rules to predictive performance. An ablation framework sequentially removes each rule from the hybrid architecture and compares predictive accuracy, symbolic penalty behavior, and rule adherence across alternative specifications. This procedure follows established validation practices

in empirical finance and machine learning by evaluating the incremental contribution of each modeling component rather than relying on a single specification (Flannery & Rangan, 2006; Lemmon et al., 2008). Additional robustness analyses examine alternative parameter settings for the symbolic constraints to verify that model performance remains stable across reasonable calibration choices. Together, these procedures ensure that improvements in predictive performance are attributable to economically meaningful constraints rather than excessive regularization or overfitting (Hastie et al., 2009).

3.4.4. Epistemic Role of Symbolic Constraints

The symbolic logic layer is intended to encode economically interpretable prior knowledge rather than impose deterministic financing rules. Consequently, symbolic violations are interpreted as indicators of financing stress, refinancing pressure, or temporary deviations from target debt policies rather than evidence of irrational corporate behavior (Flannery & Rangan, 2006; Frank & Goyal, 2009). The use of empirically calibrated thresholds, differentiable penalties, and robustness analyses allows the framework to integrate financial theory within a flexible predictive architecture while preserving statistical learning capability. By combining domain knowledge with data-driven prediction, the symbolic layer enhances interpretability and theoretical consistency without sacrificing the adaptability of the neural learning process, thereby providing a transparent implementation of hybrid reasoning in corporate finance (Hitzler et al., 2022).

3.5. Neural Learning Layer ($\mathcal{N}$)

The neural learning layer models the nonlinear and dynamic relationships between firm-specific financial characteristics and future debt adjustments using a supervised feedforward neural network. Rather than estimating structural causal effects, the network is designed as a predictive model that learns complex interactions between firm fundamentals and macro-financial conditions while preserving the temporal ordering of the panel data (Gu et al., 2020; Hastie et al., 2009). All explanatory variables are lagged by one period to prevent simultaneity bias and ensure that predictions are based exclusively on information available at the forecasting date, consistent with dynamic capital structure models (Flannery & Rangan, 2006). Let $x_{i,t} \in \mathbb{R}^p$ denote the standardized feature vector containing lagged firm-level and macroeconomic variables. Hidden representations are generated through successive nonlinear transformations:

$$h_{i,t}^{(l)} = \sigma(W^{(l)}h_{i,t}^{(l-1)} + b^{(l)}), \quad (5)$$

where $h_{i,t}^{(0)} = x_{i,t}$, $W^{(l)}$ and $b^{(l)}$ denote the trainable weights and biases of layer l , and $\sigma(\cdot)$ represents the rectified linear unit (ReLU) activation function (Goodfellow et al., 2016). The network jointly predicts one-period-ahead changes in leverage and debt maturity through a multitask output layer:

$$\widehat{\Delta Lev}_{i,t+1} = f_{Lev}(h_{i,t}^{(L)}), \widehat{\Delta Mat}_{i,t+1} = f_{Mat}(h_{i,t}^{(L)}), \quad (6)$$

thereby capturing the interdependence between leverage and maturity decisions within a unified predictive framework (Barclay & Smith, 1995).

The implemented architecture consists of two fully connected hidden layers with 64 and 32 neurons, respectively, using ReLU activation functions and linear output units. Network parameters are estimated using the Adam optimizer with adaptive learning rates, while model complexity is controlled through L2 weight regularization and early stopping based on validation performance to reduce overfitting (Goodfellow et al., 2016; Hastie et al., 2009). Hyperparameter selection is performed via a predefined search over alternative

network architectures, regularization strengths, and learning rates, with the final specification selected based on validation performance within the rolling training framework. Additional stability analyses across multiple random initialization seeds are conducted to verify that predictive performance is not driven by a particular network initialization.

Model estimation follows an expanding rolling-window cross-validation procedure that preserves the panel's chronological structure. For each rolling fold, the network is trained using observations available up to the corresponding training period, validated on the immediately subsequent year, and evaluated on the following out-of-sample test year. The training window then expands sequentially through time until the final forecasting period, ensuring that every prediction is generated without access to future information. Firm-year observations occurring before the first forecast year do not receive adjustment predictions because no valid out-of-sample forecast can be produced without violating the walk-forward design. This estimation strategy eliminates look-ahead bias, closely mimics real-time forecasting, and provides a rigorous evaluation of the model's out-of-sample predictive performance (Gu et al., 2020).

Within the hybrid neural-symbolic framework, the neural learning layer captures nonlinear predictive relationships that may not be adequately represented by conventional econometric specifications, while the symbolic logic layer constrains these predictions using economically interpretable financial principles. This complementary interaction enables the model to balance predictive flexibility with theoretical consistency, thereby improving interpretability without restricting the learning algorithm's capacity to identify complex financing dynamics.

3.6. Fusion and Training

The hybrid neuro-symbolic framework integrates data-driven prediction with economically grounded financial constraints through a unified training objective that simultaneously optimizes predictive accuracy and theoretical consistency (Hitzler et al., 2022; Goodfellow et al., 2016). Within this architecture, the neural learning layer estimates one-period-ahead adjustments in leverage and debt maturity, while the symbolic logic layer evaluates these predictions against empirically calibrated financing rules. Rather than treating the symbolic rules as deterministic restrictions, the framework incorporates them as differentiable regularization terms that guide learning while allowing the model to accommodate heterogeneous financing decisions and changing macro-financial conditions.

The overall optimization objective combines the predictive loss generated by the neural network with the aggregate symbolic penalty according to

$$\mathcal{L}_{total} = \mathcal{L}_{pred} + \lambda \sum_{r=1}^R w_r \phi_r(X_{i,t}, \hat{y}_{i,t}), \quad (7)$$

where $\mathcal{L}_{pred}$ denotes the prediction error associated with the joint estimation of leverage and maturity adjustments, $\phi_r(\cdot)$ represents the differentiable violation of symbolic rule r , w_r is the corresponding rule weight, and λ controls the overall strength of symbolic regularization (Hitzler et al., 2022). The predictive component is estimated using a multitask regression objective, whereas the symbolic component penalizes deviations from economically grounded financing principles through the softplus formulation. The equal weighting of symbolic rules reflects the absence of a theoretical basis for assigning differential importance across the embedded financial constraints.

Model estimation follows the same expanding rolling-window framework adopted for the neural learning layer to preserve the chronological ordering of the panel data and eliminate look-ahead bias (Gu et al., 2020). Within each rolling fold, model hyperparameters are selected using a validation subset drawn exclusively from the corresponding training

period, while the final evaluation is performed on the subsequent out-of-sample test year. Hyperparameter tuning considers alternative network architectures, learning rates, and regularization strengths, with the final specification selected according to validation performance. Early stopping, L2 regularization, and repeated estimation across multiple random initialization seeds further improve model stability and reduce the likelihood of overfitting (Goodfellow et al., 2016; Hastie et al., 2009).

For computational tractability within the available implementation environment, the joint objective above was operationalized via the resampling-based procedure summarized in Algorithm 1 rather than by directly minimizing L_{total} end-to-end. Specifically, within each training fold, observations were resampled with probability weights $w_i = 1/(1 + \lambda\Phi_i)$, so that observations associated with larger symbolic-rule violations were down-weighted before the neural network architecture was estimated on the reweighted sample. This resampling scheme is a first-order proxy for the differentiable joint-penalty formulation: it shifts the effective training distribution toward theory-consistent observations in a manner analogous to penalizing violations directly within the loss function, without requiring a custom differentiable training loop within the available computational environment. All neuro-symbolic results reported in this study and the associated figures were generated using this resampling-based estimator. The differentiable softplus penalty formulation 2 and the joint loss function above specify the theoretical objective that the resampling procedure approximates; implementations with access to full automatic-differentiation frameworks could instead estimate this joint objective directly, which we flag as a direction for future refinement of the framework.

To evaluate the contribution of the symbolic component, the hybrid framework is compared with alternative specifications via a structured ablation analysis in which symbolic constraints are sequentially removed or modified while maintaining the same neural architecture. This procedure isolates the incremental contribution of economically motivated constraints to predictive performance and model stability, consistent with validation practices in both machine learning and empirical finance (Flannery & Rangan, 2006; Hitzler et al., 2022). By jointly optimizing prediction accuracy and symbolic consistency within a rolling out-of-sample framework, the proposed architecture provides a transparent integration of statistical learning and financial theory while preserving both predictive flexibility and economic interpretability.

3.7. Econometric Framework

To evaluate the dynamic behavior of corporate debt decisions and assess the incremental informational content of the Neuro-Symbolic Debt (NSD) index, the study employs complementary panel econometric estimators to account for unobserved heterogeneity, leverage persistence, and potential endogeneity. Consistent with the study's predictive objective, econometric estimation is used to quantify dynamic debt adjustment and to benchmark the proposed neuro-symbolic framework against established panel data approaches, rather than to establish causal relationships (Angrist & Pischke, 2009; Wooldridge, 2010).

3.7.1. Two-Way Fixed Effects

The baseline specification is estimated using a two-way fixed-effects model that controls for both firm-specific heterogeneity and common time effects. The empirical specification is given by

$$\text{Debt} / \text{Assets}_{i,t} = \rho \text{Debt} / \text{Assets}_{i,t-1} + \beta' X_{i,t} + \mu_i + \lambda_t + \varepsilon_{i,t}, \quad (8)$$

where $X_{i,t}$ denotes the vector of firm-specific explanatory variables, μ_i represents firm fixed effects, λ_t captures year fixed effects, and $\varepsilon_{i,t}$ is the idiosyncratic error term. Including the

lagged dependent variable allows the model to capture persistence in corporate leverage, consistent with partial adjustment theories of capital structure (Flannery & Rangan, 2006; Lemmon et al., 2008).

Estimation is performed using the exact within transformation, whereby both firm-specific and year-specific means are removed from each variable prior to estimation:

$$\tilde{y}_{i,t} = y_{i,t} - \bar{y}_i - \bar{y}_t + \bar{y}, \quad (9)$$

thereby eliminating time-invariant firm characteristics and common macroeconomic shocks before ordinary least squares estimation. To account for serial dependence within firms, statistical inference is based on firm-level heteroskedasticity-consistent cluster-robust covariance estimators (Arellano, 1987; Wooldridge, 2010). For comparison, one-way fixed-effects specifications that include only firm effects are also estimated.

3.7.2. Difference GMM

To address the dynamic nature of leverage adjustment and the potential endogeneity of the lagged dependent variable, the study additionally estimates the Arellano–Bond Difference Generalized Method of Moments (Difference GMM) estimator (Arellano & Bond, 1991). First differencing removes unobserved firm-specific effects, yielding

$$\Delta Debt / Assets_{i,t} = \rho \Delta Debt / Assets_{i,t-1} + \Delta X'_{i,t} \beta + \Delta \varepsilon_{i,t}. \quad (10)$$

The differenced lagged dependent variable is instrumented using appropriately lagged levels of leverage, while the remaining explanatory variables enter the instrument matrix according to their assumed predetermined properties. Instrument selection follows a parsimonious strategy to limit instrument proliferation and preserve estimator efficiency in finite samples (Roodman, 2009).

Parameter estimates are obtained using the two-step efficient GMM estimator,

$$\hat{\beta} = (X'ZWZ'X)^{-1}X'ZWZ'y, \quad (11)$$

where Z denotes the instrument matrix and W is the optimal weighting matrix estimated from first-step residuals. Model validity is evaluated using the Arellano–Bond serial-correlation tests for first- and second-order autocorrelation in differenced residuals together with the Sargan/Hansen over-identification tests for instrument validity (Arellano & Bond, 1991). These diagnostics provide evidence on the appropriateness of the instrument set and the consistency of the dynamic panel estimator.

3.8. Machine Learning Benchmarks

To evaluate whether the proposed neuro-symbolic framework provides predictive improvements beyond conventional statistical and machine learning approaches, a comprehensive benchmarking framework is implemented using linear, regularized, ensemble, and neural network models. All benchmark models are estimated using the same expanding rolling-window validation procedure, thereby ensuring that performance comparisons are based on identical training, validation, and out-of-sample test periods. This common evaluation framework eliminates differences arising from alternative data partitions and provides a consistent basis for comparing predictive performance across competing methods (Gu et al., 2020).

The benchmark models include Least Absolute Shrinkage and Selection Operator (LASSO), Ridge Regression, Gradient Boosting Regression, and a conventional feedforward neural network without symbolic constraints. LASSO and Ridge regression are included to

address multicollinearity and overfitting via coefficient regularization and are widely used linear predictive methods in high-dimensional financial applications (Hastie et al., 2009; Tibshirani, 1996). Gradient Boosting is an ensemble learning approach that sequentially combines weak learners to improve predictive accuracy and capture complex interactions among explanatory variables (Friedman, 2001). The unconstrained neural network uses the same architecture as the neuro-symbolic model but excludes the symbolic logic layer, thereby isolating the incremental contribution of theory-guided constraint integration.

All benchmark models predict one-period-ahead changes in leverage using an identical set of standardized explanatory variables comprising firm-level financial characteristics and lagged macroeconomic variables. To prevent information leakage, feature standardization is performed separately within each rolling training fold, and the resulting transformation is subsequently applied to the corresponding validation and test observations. Consequently, no information from future periods enters model estimation or feature scaling, preserving the integrity of the out-of-sample forecasting exercise (Gu et al., 2020).

In addition, sensitivity analyses are conducted to examine whether predictive performance is affected by alternative definitions of market-based explanatory variables that share accounting components with leverage measures. These analyses evaluate the robustness of model performance under alternative feature constructions and ensure that forecasting accuracy is not mechanically driven by overlapping accounting identities. By estimating all benchmark models under identical forecasting protocols and data partitions, the benchmarking framework provides a rigorous assessment of the incremental predictive value contributed by nonlinear learning and symbolic financial reasoning beyond conventional econometric and machine learning approaches.

3.9. Model Evaluation

The predictive performance of the proposed neuro-symbolic framework is evaluated using complementary measures of forecasting accuracy, statistical significance, and model robustness. Consistent with best practices in financial prediction, model evaluation emphasizes genuine out-of-sample performance obtained under the expanding rolling-window validation framework, ensuring that every prediction is generated exclusively from information available at the corresponding forecasting date (Gu et al., 2020). This evaluation strategy minimizes look-ahead bias and provides a realistic assessment of predictive performance under sequential forecasting conditions.

Forecast accuracy is assessed using the Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), and the coefficient of determination (R^2), which jointly evaluate the magnitude of prediction errors and the proportion of variation explained by each model (Hastie et al., 2009; Wooldridge, 2010). RMSE places greater emphasis on larger forecasting errors, whereas MAE provides a scale-consistent measure of average prediction error that is less sensitive to extreme observations. The coefficient of determination is reported to facilitate comparison of overall predictive fit across econometric, machine learning, and neuro-symbolic models.

To evaluate whether observed differences in forecasting performance are statistically significant, competing models are compared using the Diebold–Mariano test based on forecast error differentials (Diebold & Mariano, 1995). The test evaluates the null hypothesis that two forecasting models possess equal predictive accuracy while accounting for serial dependence in the sequence of forecast errors through heteroskedasticity- and autocorrelation-consistent variance estimation. This procedure provides a formal statistical assessment of whether improvements in predictive performance reflect genuine forecasting gains rather than sampling variability.

The contribution of the symbolic reasoning layer is further examined through a structured ablation analysis in which each symbolic rule is sequentially removed while maintaining an identical neural network architecture and estimation procedure. Following the removal of each rule, the aggregate symbolic penalty and Rule Adherence Rate (RAR) are recomputed, and the hybrid model is re-estimated under the same regularization settings to evaluate changes in predictive performance. This approach isolates the incremental contribution of each symbolic constraint to model accuracy, stability, and theoretical consistency, providing a transparent assessment of the role of domain knowledge within the neuro-symbolic learning framework (Hitzler et al., 2022). Together, these complementary evaluation procedures provide a comprehensive assessment of predictive performance, statistical robustness, and the contribution of economically grounded symbolic reasoning to financial forecasting.

3.10. Ablation Analysis

To evaluate the contribution of symbolic reasoning to predictive performance, the study employs a structured ablation analysis in which each symbolic rule is sequentially removed while maintaining an identical neural network architecture, training protocol, and regularization settings. This approach isolates the incremental contribution of each symbolic constraint by ensuring that all remaining components of the hybrid framework remain unchanged across specifications (Hitzler et al., 2022). The ablation design follows established validation practices in interpretable machine learning, where individual model components are systematically excluded to assess their influence on predictive accuracy, model stability, and theoretical consistency.

For each ablation experiment, the symbolic penalty and Rule Adherence Rate (RAR) are recalculated after excluding the selected rule, and the hybrid model is subsequently re-estimated using the same expanding rolling-window cross-validation procedure.

The strength of symbolic regularization and all neural network hyperparameters are held constant throughout the analysis to ensure that observed differences in performance are attributable solely to the removal of the corresponding symbolic constraint. This controlled experimental design enables direct comparison across alternative model specifications while eliminating the influence of changes in model complexity or parameter tuning.

Model performance is evaluated by comparing prediction errors and symbolic consistency across the alternative specifications. Changes in forecasting accuracy provide evidence on the predictive contribution of individual symbolic rules, whereas variations in the aggregate symbolic penalty and Rule Adherence Rate indicate the extent to which each rule contributes to maintaining economically consistent financing predictions. This complementary evaluation distinguishes the predictive role of individual constraints from their contribution to theoretical coherence within the hybrid learning framework.

The ablation analysis also provides insight into the relative importance of the symbolic rules and identifies potential concentration of the symbolic penalty among specific financial constraints. By systematically examining the impact of each rule on both predictive performance and symbolic consistency, the analysis enhances the transparency and interpretability of the proposed neuro-symbolic framework while demonstrating the incremental value of integrating economically grounded financial knowledge into machine learning models (Hitzler et al., 2022).

4. Results

4.1. Preliminary Data Diagnostics

Table 4 presents the descriptive statistics for all variables included in the empirical analysis. The sample exhibits considerable cross-sectional variation in firms' capital structures and financial characteristics. The mean debt-to-assets ratio is 0.2118, indicating that debt finances approximately one-fifth of total assets on average, while market leverage displays greater dispersion than book leverage. Net debt-to-EBIT shows substantial variability, reflecting heterogeneous debt-servicing capacities across firms. The macroeconomic variables also exhibit considerable variation throughout the sample period, particularly inflation, which is consistent with the diverse economic conditions across the MENA countries included in the analysis. Overall, the descriptive statistics suggest sufficient variation in both firm-specific and macroeconomic variables to support the subsequent econometric and machine learning analyses.

Table 4. Summary Statistics.

Variable	Mean	Median	SD	P25	P75
Debt-to-Assets	0.2118	0.1460	0.2220	0.0151	0.3381
Market Leverage	0.4056	0.3508	0.3543	0.0323	0.7517
Net Debt-to-EBIT	2.9001	0.8262	22.8794	−1.0811	4.5747
Short-Term Debt Share	0.5246	0.5020	0.3459	0.2164	0.8774
Cash-to-Assets	0.1449	0.0662	0.2000	0.0192	0.1790
Return on Assets (ROA)	0.0129	0.0392	0.2407	0.0061	0.0857
Market-to-Book	0.9341	0.5076	1.6346	0.2416	0.8920
Tangibility	0.2599	0.1602	0.2712	0.0233	0.4320
Log of Total Assets	20.2424	20.1232	2.7010	18.3347	21.9188
Lagged Inflation	18.8642	8.8916	20.0636	7.7751	16.3325
Lagged Market Volatility	23.4967	23.9123	3.6185	22.3364	24.6390

Notes: This table reports the descriptive statistics for the variables used in the empirical analyses. Mean and Median represent measures of central tendency, SD denotes the standard deviation, and P25 and P75 correspond to the 25th and 75th percentiles, respectively.

Table 5 reports the Pearson correlation matrix. Most pairwise correlations are low to moderate in magnitude, indicating that the explanatory variables capture distinct dimensions of firms' financial characteristics. As expected, debt-to-assets is positively associated with market leverage, while cash holdings and short-term debt share exhibit negative correlations with leverage. Firm size is positively correlated with market leverage but negatively associated with market-to-book ratio, whereas profitability displays a negative relationship with leverage measures. Although several correlations are statistically significant, none indicate excessive linear dependence among the regressors. This conclusion is further supported by the variance inflation factor (VIF) diagnostics, with all VIF values ranging from 1.00 to 2.89, well below the commonly accepted threshold of 5. These results confirm that multicollinearity is not a concern and that the explanatory variables are suitable for multivariate estimation.

Table 5. Correlation Matrix.

Variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	VIF
(1) Debt-to-Assets	1.000 ***											2.180
(2) Market Leverage	0.630 ***	1.000 ***										2.890
(3) Net Debt-to-EBIT	0.120 ***	0.170 ***	1.000 ***									1.049
(4) Short-Term Debt Share	−0.170 ***	−0.100 ***	−0.020 ***	1.000 ***								1.106
(5) Cash-to-Assets	−0.220 ***	−0.330 ***	−0.140 ***	−0.020 ***	1.000 ***							1.241
(6) ROA	−0.060 ***	0.090 ***	0.010	−0.020 ***	−0.200 ***	1.000 ***						1.203
(7) Market-to-Book	0.090 ***	−0.270 ***	−0.030 ***	−0.010 **	0.170 ***	−0.310 ***	1.000 ***					1.337
(8) Tangibility	0.220 ***	0.110 ***	−0.020 ***	−0.140 ***	−0.230 ***	0.070 ***	−0.000	1.000 ***				1.136
(9) Log Assets	0.090 ***	0.530 ***	0.140 ***	−0.170 ***	−0.240 ***	0.300 ***	−0.380 ***	−0.020 ***	1.000 ***			1.845
(10) Inflation (Lagged)	−0.010 **	−0.000	−0.010 *	0.020 ***	0.000	0.020 ***	−0.000	0.010 **	−0.000	1.000 ***		1.003
(11) Market Volatility (Lagged)	0.010 **	0.020 ***	−0.020 ***	−0.030 ***	0.030 ***	−0.010	−0.010	−0.010 **	0.030 ***	−0.040 ***	1.000 ***	1.006

Notes: This table reports the Pearson pairwise correlation coefficients and variance inflation factor (VIF) diagnostics for the variables included in the empirical analyses. Statistical significance is denoted as follows: *** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$.

4.2. Dynamic Capital Structure Estimation

Table 6 reports the dynamic fixed-effects estimation results for both firm-fixed-effects and firm-plus-year-fixed-effects specifications. The coefficient on lagged debt-to-assets is positive and highly significant across both models, indicating substantial persistence in firms' capital structure and supporting the partial adjustment hypothesis. Firm profitability is negatively associated with leverage, whereas asset tangibility is positively associated with it. The results further show that inflation has a small but statistically significant negative association with leverage, while market volatility is not statistically significant. The estimated speed of adjustment ranges from 35.21% to 36.42%, suggesting that firms adjust approximately one-third of the gap between actual and target leverage within one year.

Table 6. Dynamic Fixed Effects Estimation.

Panel A. Firm Fixed Effects	Coefficient	Clustered SE	t-Statistic	p-Value
Lagged Debt-to-Assets	0.6358 ***	0.0107	59.1910	<0.001
Return on Assets (ROA)	−0.0910 ***	0.0107	−8.5345	<0.001
Log of Total Assets	0.0044 ***	0.0013	3.4240	0.0006
Market-to-Book	0.0073 ***	0.0016	4.6508	<0.001
Tangibility	0.0796 ***	0.0094	8.4868	<0.001
Lagged Inflation	−0.0001 ***	0.0000	−3.4271	0.0006
Lagged Market Volatility	−0.0001	0.0002	−0.5268	0.5984
Panel B. Firm and Year Fixed Effects	Coefficient	Clustered SE	t-Statistic	p-Value
Lagged Debt-to-Assets	0.6479 ***	0.0106	61.2333	<0.001
Lagged Return on Assets (ROA)	−0.0164 **	0.0083	−1.9911	0.0465
Lagged Log of Total Assets	0.0022	0.0014	1.5581	0.1192
Lagged Market-to-Book	−0.0029 ***	0.0010	−2.9150	0.0036
Lagged Tangibility	0.0129 *	0.0067	1.9325	0.0533
Lagged Inflation	−0.0001 **	0.0000	−1.9892	0.0467
Lagged Market Volatility	0.0001	0.0001	0.7855	0.4322
Panel C. Model Summary	Within R²	Observations	Firms	Speed of Adjustment (%)
Firm Fixed Effects	0.4755	36,206	2998	36.42
Firm and Year Fixed Effects	0.4528	36,187	2998	35.21

Notes: The speed of adjustment is calculated as $1 - \rho$, where ρ is the coefficient on the lagged debt-to-assets ratio. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Table 7 presents the results of the two-step Difference GMM estimator, which accounts for the dynamic nature of capital structure and potential endogeneity. The lagged leverage coefficient remains positive and highly significant, confirming the persistence of leverage decisions after controlling for endogeneity. Profitability has a positive and statistically significant effect, while market-to-book ratio and tangibility are negatively associated with leverage. Firm size is not statistically significant in the dynamic specification. The diagnostic tests indicate no evidence of second-order serial correlation, and the Sargan test supports the validity of the instrument set.

Table 7. Difference GMM Results.

Variable	Coefficient	Standard Error	z-Statistic	p-Value
Lagged Debt-to-Assets	0.6959 ***	0.0398	17.4817	<0.001
Lagged Return on Assets (ROA)	0.0547 ***	0.0156	3.5073	0.0005
Lagged Log of Total Assets	−0.0057	0.0043	−1.3070	0.1912
Lagged Market-to-Book	−0.0200 ***	0.0035	−5.7119	<0.001
Lagged Tangibility	−0.0881 ***	0.0162	−5.4439	<0.001

Notes: *** denote statistical significance at the 1%.

Table 8 reports the restricted-instrument Difference GMM specification that uses only second-order lags as instruments. The estimated coefficient on lagged leverage is consistent with the full instrument specification, indicating that the results are robust to alternative instrument sets. Market-to-book ratio remains negatively associated with leverage, whereas firm size continues to be statistically insignificant. Overall, the restricted specification yields results qualitatively similar to those from the full instrument model.

Table 8. Restricted Instrument GMM.

Variable	Coefficient	Standard Error	z-Statistic	p-Value
Lagged Debt-to-Assets	0.7083 ***	0.0461	15.3570	<0.001
Lagged Log of Total Assets	0.0024	0.0053	0.4450	0.6563
Lagged Market-to-Book	−0.0227 ***	0.0039	−5.7665	<0.001

Notes: *** denote statistical significance at the 1%.

Figure 3 illustrates the estimated speed of adjustment across the alternative econometric specifications, while Table 9 summarizes the corresponding adjustment coefficients. The fixed effects model produces the highest adjustment speed, whereas both Difference GMM specifications indicate slightly slower adjustment. Nevertheless, all estimators consistently suggest that firms gradually converge toward their target capital structure, with annual adjustment speeds ranging from approximately 29% to 35%.


**Figure 3.** Comparison of Adjustment Speeds.

Table 9. Comparison of Adjustment Speeds Results

Model	Lagged Debt-to-Assets Coefficient (ρ)	Speed of Adjustment (%)
Firm and Year Fixed Effects	0.6479	35.21
Difference GMM (Full Instrument Set)	0.6959	30.41
Restricted Instrument Difference GMM (Lag 2)	0.7083	29.17

Table 9 summarizes the estimated persistence coefficients and the corresponding speed of adjustment across the alternative econometric specifications. The fixed effects model produces the highest adjustment speed of 35.21%, whereas the Difference GMM estimators yield slightly lower adjustment speeds of 30.41% and 29.17% for the full and restricted instrument specifications, respectively. Despite these differences, all estimators provide consistent evidence that firms adjust their capital structures gradually rather than instantaneously, with approximately one-third of the deviation from target leverage being corrected within a year. The similarity of the estimated adjustment speeds across alternative estimation approaches supports the robustness of the dynamic capital structure findings.

4.3. Machine Learning Forecasting Performance

Table 10 reports the out-of-sample predictive performance of four machine learning algorithms across two alternative market-to-book specifications. Overall, the Gradient Boosting Machine (GBM) consistently delivers the highest predictive accuracy across both feature sets, achieving the lowest RMSE and MAE, as well as the highest coefficient of determination (R^2). The improvement is more pronounced when the alternative market-to-book measure, which excludes the shared debt component, is employed, suggesting that the revised specification provides a cleaner representation of firms' growth opportunities.

The linear regularization models, LASSO and Ridge, exhibit similar predictive performance, with only marginal differences in prediction errors and explanatory power. In contrast, the neural network produces the weakest results across both specifications, recording the highest prediction errors and the lowest R^2 values. The negative R^2 obtained under the conventional market-to-book specification indicates that the neural network performs worse than a simple mean-based benchmark.

The Diebold–Mariano tests further confirm these findings. Negative and statistically significant DM statistics indicate that the forecasting errors of LASSO, Ridge, and particularly GBM are significantly smaller than those of the neural network benchmark. Across both feature sets, GBM records the largest and most significant DM statistics, providing strong evidence of its superior out-of-sample forecasting performance. Collectively, these results demonstrate that ensemble tree-based methods outperform both linear regularization techniques and neural networks in predicting capital structure dynamics, while the alternative market-to-book specification further enhances predictive accuracy.

To provide a comparable evaluation for debt maturity, Table 10b reports the out-of-sample forecasting performance of the same four benchmark algorithms (LASSO, Ridge, Gradient Boosting, and the feedforward Neural Network) for one-period-ahead changes in debt maturity, $\Delta \text{Mat}_{(i,t+1)} = \text{Mat}_{(i,t+1)} - \text{Mat}_{(i,t)}$, where $\text{Mat}_{(i,t)} = 1 - \text{ST_Share}_{(i,t)}$ is the long-term debt share. Models are estimated using the identical expanding rolling-window walk-forward protocol, feature set, hyperparameters, and evaluation metrics (RMSE, MAE, R^2 , and the Diebold–Mariano test against the Neural Network benchmark).

Table 10. a. Out-of-Sample Forecast Performance; b. Out-of-Sample Forecast Performance—Debt Maturity.

(a)					
Panel A. Conventional Market-to-Book Ratio (Tobin's Q with Shared Total Debt Component)	RMSE	MAE	R ²	DM Statistic vs. Neural Network	DM p-Value
LASSO	0.0934	0.0507	0.0514	−3.5903	0.0003
Ridge	0.0934	0.0513	0.0501	−3.6399	0.0003
Gradient Boosting Machine (GBM)	0.0920	0.0494	0.0788	−7.7717	<0.001
Neural Network	0.0964	0.0544	−0.0101	—	—
Panel B. Alternative Market-to-Book Ratio (Market Capitalization/Book Equity)	RMSE	MAE	R ²	DM Statistic vs. Neural Network	DM p-Value
LASSO	0.0933	0.0506	0.0527	−2.4949	0.0126
Ridge	0.0934	0.0513	0.0515	−2.5380	0.0111
Gradient Boosting Machine (GBM)	0.0916	0.0493	0.0877	−8.3725	<0.001
Neural Network	0.0954	0.0547	0.0099	—	—
(b)					
Panel A. Conventional Market-to-Book Ratio (Tobin's Q with Shared Total Debt Component)	RMSE	MAE	R ²	DM Statistic vs. Neural Network	DM p-Value
LASSO	0.1874	0.0988	−0.0021	−9.9405	<0.001
Ridge	0.1874	0.0993	−0.0025	−9.9191	<0.001
Gradient Boosting Machine (GBM)	0.1883	0.1015	−0.0115	−6.1266	<0.001
Neural Network	0.1898	0.1061	−0.0281	—	—
Panel B. Alternative Market-to-Book Ratio (Market Capitalization/Book Equity)	RMSE	MAE	R ²	DM Statistic vs. Neural Network	DM p-Value
LASSO	0.1875	0.0989	−0.0030	−9.7029	<0.001
Ridge	0.1876	0.0995	−0.0038	−9.4944	<0.001
Gradient Boosting Machine (GBM)	0.1884	0.1021	−0.0130	−5.6096	<0.001
Neural Network	0.1899	0.1067	−0.0291	—	—

Notes: a. RMSE denotes root mean squared error, MAE denotes mean absolute error, and R² represents the out-of-sample coefficient of determination. The Diebold-Mariano (DM) statistic tests the predictive accuracy of each model relative to the neural network benchmark. Negative DM statistics indicate superior forecasting performance compared with the neural network. Lower values of RMSE and MAE and higher values of R² indicate better predictive performance. b. RMSE denotes root mean squared error, MAE denotes mean absolute error, and R² represents the out-of-sample coefficient of determination for the one-period-ahead change in debt maturity ($\Delta \text{Mat}_{(i,t+1)}$). The Diebold-Mariano (DM) statistic tests the predictive accuracy of each model relative to the neural network benchmark; negative DM statistics indicate significantly smaller forecast errors than the neural network. Feature set, rolling-window folds, and hyperparameters are identical to those underlying Table 10a.

In contrast to leverage adjustment, none of the four models achieve positive out-of-sample explanatory power for debt-maturity adjustment: R² is negative for every model under both market-to-book specifications, ranging from −0.002 (LASSO) to −0.029 (Neural Network), indicating that none of the benchmark models beat a naive mean forecast out-of-sample. Prediction errors are also substantially larger in absolute terms than for leverage: pooled RMSE is approximately 0.187–0.190 for debt maturity compared with approximately 0.092–0.096 for leverage (Table 10), consistent with the greater volatility of the short-term debt share documented in Table 4. The linear regularized models (LASSO and Ridge)

attain the lowest RMSE and MAE and the least negative R^2 , and the Diebold-Mariano tests indicate that both significantly outperform the Neural Network ($p < 0.001$), while Gradient Boosting also significantly outperforms the Neural Network but no longer dominates the linear models as it did for leverage—a reversal of the ranking reported in Table 10. Taken together, these results indicate that, conditional on the firm-level and macroeconomic predictors employed in this study, debt-maturity adjustment is considerably harder to forecast out-of-sample than leverage adjustment, and that increasing model complexity does not compensate for this. We interpret this asymmetry as evidence that debt-maturity decisions are shaped more by supply-side and institutional factors, such as lender-driven covenant structures and debt-market access, that are not captured by the firm-specific and macroeconomic variables used here, whereas leverage exhibits the strong mechanical persistence that dominates its forecasting exercise. Accordingly, we temper the paper's predictive claims by focusing on leverage adjustment, for which evidence of genuine out-of-sample predictability is considerably stronger, and we treat the debt-maturity component of the neuro-symbolic framework primarily as a structural input to the NSD index rather than as a high-accuracy forecasting target in its own right.

4.4. Symbolic Constraint Analysis

Table 11 reports the distribution of the symbolic penalty term used to quantify violations of the predefined financial constraints. The distribution is positively skewed, with a mean of 18.164 and a maximum value of 807.421, indicating that a limited number of observations incur substantial penalties. In contrast, the median (0.286) and interquartile range (0.185–0.501) remain low, suggesting that most observations satisfy the symbolic constraints with only minor deviations. The absence of missing values confirms complete evaluation of the symbolic penalty across the full sample.

Table 11. Distribution of Symbolic Penalty.

Statistic	Value
Mean	18.164
Median	0.2857
1st Quartile	0.1848
3rd Quartile	0.5011
Maximum	807.4213
Missing	0

Table 12 and Figures 4–6 evaluate the sensitivity of the neuro-symbolic framework to the symbolic regularization parameter (λ). Across the examined values of λ , prediction accuracy remains highly stable, with the mean squared error (MSE) fluctuating only between 0.0012 and 0.0014. Similarly, the Rule Adherence Ratio (RAR) varies within a narrow interval from 0.8467 to 0.8497, indicating consistently high compliance with the imposed symbolic constraints.

Table 12. Lambda Sensitivity.

λ	MSE	RAR
0	0.0012	0.8481
0.2	0.0013	0.8497
0.4	0.0014	0.8482
0.6	0.0012	0.8467
0.8	0.0014	0.8481
1	0.0013	0.8493




Figure 4. Predictive Accuracy versus λ .




Figure 5. Rule Adherence versus λ .

Figure 4 illustrates that predictive accuracy is largely unaffected by changes in λ , while Figure 5 demonstrates that rule adherence remains consistently high across all regularization levels. Figure 6 further highlights the trade-off between predictive performance and symbolic consistency, showing that increasing the weight assigned to symbolic constraints does not materially reduce forecasting accuracy. Collectively, these findings indicate that the proposed neuro-symbolic framework is robust to the choice of λ , achieving stable predictive performance while maintaining a high degree of compliance with the embedded financial constraints.




Figure 6. Prediction–Constraint Trade-off.

4.5. Comparative Performance of Alternative Models

Table 13 and Figure 7 compare the predictive performance of the linear, neural, and neuro-symbolic models using mean squared error (MSE) and the Rule Adherence Ratio (RAR). The linear model achieves the lowest prediction error, with an MSE of 0.0007, while maintaining a high level of rule compliance (RAR = 0.8500). The neural model records a slightly higher MSE of 0.0012 and a comparable RAR of 0.8481.

Table 13. Consolidated Model Comparison.

Model	MSE	RAR
Linear	0.0007	0.85
Neural	0.0012	0.8481
Neuro-Symbolic	0.0013	0.8493




Figure 7. Model Comparison

The neuro-symbolic model produces an MSE of 0.0013 and an RAR of 0.8493, indicating that incorporating symbolic constraints results in only a marginal reduction in predictive accuracy while preserving a high degree of adherence to the predefined financial rules. Figure 7 visually confirms that the performance differences among the three models are relatively small, with all approaches achieving similar levels of rule compliance.

The RAR of 0.8493 reported here, and the corresponding $\lambda = 1.0$ row of Table 12, are computed under the observed-state rule evaluation, i.e., symbolic rules are assessed against firms' currently observed debt positions. Table 14 additionally reports RAR under the projected-state evaluation, in which the leverage-band and maturity-ladder rules are re-evaluated against the model's own one-period-ahead forecasts rather than the observed state. Because the projected state incorporates forecast uncertainty, the projected-state RAR for the equivalent full-model, $\lambda = 1.0$ specification is lower (0.7427; Table 14) than its observed-state counterpart reported here. The two figures are therefore not directly comparable: the observed-state RAR (Tables 12, 13 and 15) characterizes how consistent firms' current debt positions are with the symbolic rules, whereas the projected-state RAR (Table 14) characterizes how consistent the model's forecasts are with those rules, which we consider the more stringent and decision-relevant metric. We retain both measures for comparability with the λ -sensitivity analysis and clarify this distinction to avoid the appearance of an internal inconsistency between Tables 12–14.

Table 14. Ablation Study.

Configuration	RAR	MSE	Mean Φ
Full model (all 5 rules)	0.7427	0.0015	18.164
Without: leverage_band	0.8643	0.0021	18.0173
Without: coverage	0.7034	0.0017	0.3304
Without: maturity_ladder	0.6813	0.0012	18.1579
Without: refinancing_stress	0.6969	0.0014	18.082
Without: asset_liability	0.7679	0.0015	18.0685

Table 15. Redundancy Analysis.

NSD Variant	vs.	Correlation
NSD_equal	debt_assets	0.2321
NSD_equal	market_leverage	0.361
NSD_pca	debt_assets	0.0134
NSD_pca	market_leverage	0.0741

4.6. Monetary Regime and Dynamic Adjustment

Figures 8–11 examine how monetary conditions influence leverage adjustment and assess whether firms' financing decisions remain consistent with the symbolic constraints embedded in the proposed framework.

Figure 8 presents the average change in leverage across lagged regional volatility regimes. Leverage adjustment varies systematically with monetary stress. Firms operating in the highest volatility regime exhibit the largest positive adjustment, whereas intermediate stress conditions are associated with a substantial reduction in leverage. These findings indicate that leverage decisions are sensitive to changes in macro-financial conditions and that adjustment behavior is heterogeneous across monetary environments.




Figure 8. Monetary Tightening and Leverage Adjustment.

Figure 9 illustrates the partial adjustment process by relating the leverage gap to subsequent changes in leverage under normal and monetary tightening regimes. In both cases, the fitted relationships are positive, confirming the existence of partial adjustment toward target leverage. However, the flatter slope during monetary tightening suggests a slower adjustment process when financing conditions become more restrictive, indicating that firms face greater frictions in restoring their desired capital structure during periods of financial stress.




Figure 9. Partial Adjustment under Monetary Conditioning.

Figure 10 further compares the distribution of leverage adjustments under normal and stress regimes. Although the median adjustment remains close to zero in both cases, the stress regime exhibits a wider dispersion and a broader interquartile range, reflecting greater heterogeneity in firms' financing responses during adverse monetary conditions. This increased variability is consistent with firms facing different degrees of financing constraints and adjustment costs.




Figure 10. Leverage Adjustment under Stress Regime.

Finally, Figure 11 evaluates leverage dynamics relative to the symbolic constraint bands. Firms with leverage below or within the prescribed bands record positive average adjustments, indicating gradual movement toward the target region. In contrast, firms above the upper symbolic band exhibit a negative average adjustment, reflecting systematic deleveraging. This asymmetric adjustment pattern is consistent with the symbolic financial rules incorporated into the neuro-symbolic framework and provides evidence that the estimated dynamics remain economically interpretable while respecting the predefined financial constraints.




Figure 11. Leverage Dynamics within Symbolic Bands.

4.7. Ablation Analysis Results

Table 14 presents the ablation analysis, which evaluates the contribution of each symbolic rule by sequentially removing one constraint from the full neuro-symbolic framework. The full model, incorporating all five symbolic rules, achieves a Rule Adherence Ratio (RAR) of 0.7427, a mean squared error (MSE) of 0.0015, and an average symbolic penalty (Φ) of 18.164, serving as the benchmark for comparison.

Removing the leverage band rule yields the largest increase in rule adherence (RAR = 0.8643) but also the highest prediction error (MSE = 0.0021), indicating that this constraint imposes the strongest regularization on the learning process. In contrast, excluding the coverage rule substantially reduces the average symbolic penalty to 0.3304 while decreasing rule adherence (RAR = 0.7034), suggesting that this rule is the primary contributor to the overall penalty term.

The exclusion of the maturity ladder rule yields the lowest prediction error (MSE = 0.0012) but also the lowest rule adherence (RAR = 0.6813), implying a trade-off between predictive accuracy and compliance with the symbolic constraints. Similarly, removing the refinancing stress rule slightly improves predictive performance while reducing rule adherence, whereas excluding the asset–liability rule has only a limited effect on both MSE and the symbolic penalty.

4.8. Construct Validation of the NSD Index

The construct validity of the proposed Neuro-Symbolic Diagnostics (NSD) index is evaluated through redundancy analysis and variance decomposition. Together, these analyses assess whether the NSD index captures information beyond conventional leverage measures and whether it provides incremental explanatory power for firms' capital structure dynamics.

Table 15 reports the correlations between the two NSD variants and traditional leverage measures. The equally weighted NSD index exhibits low-to-moderate positive correlations with both debt-to-assets (0.2321) and market leverage (0.3610), indicating that it is related to firms' leverage decisions while remaining empirically distinct from these conventional indicators. In contrast, the principal component-based NSD index shows only weak correlations with debt-to-assets (0.0134) and market leverage (0.0741), suggesting that it captures complementary information rather than merely reflecting existing leverage measures.

Table 16 presents the variance decomposition analysis. The structural variables alone explain a limited proportion of the variation in the dependent variable (within $R^2 = 0.0009$). Incorporating the adjustment variables substantially increases the explanatory power, raising the within R^2 to 0.0028 and contributing an incremental R^2 of 0.0018. The subsequent inclusion of the symbolic rule adherence measure (Φ) further increases the within R^2 to 0.0032, providing an additional incremental R^2 of 0.0004.

Table 16. Variance Decomposition.

Block	Within_ R^2	Incremental_ R^2	N
Structural only (Lev, Mat)	0.0009	—	16,789
+Adjustment (DeltaLevHat, DeltaMatHat)	0.0028	0.0018	16,789
+Rule adherence (Phi)	0.0032	0.0004	16,789

5. Discussion

This study examines the leverage adjustment and the relative effectiveness of econometric, machine learning, and neuro-symbolic approaches in forecasting financing decisions. Overall, the empirical evidence strongly supports the theoretical foundations of dynamic capital structure, confirming that leverage adjustment is a gradual and persistent process shaped by both firm-specific characteristics and macroeconomic conditions. The results also demonstrate that while machine learning algorithms can improve predictive performance, the gains depend on the underlying characteristics of the adjustment process. In contrast, incorporating symbolic financial knowledge improves the theoretical coherence and interpretability of the learning framework but does not generate statistically meaningful

improvements in forecasting accuracy. These findings refine, rather than overturn, existing theories of capital structure by showing that the value of increasingly complex artificial intelligence models depends on the economic structure of the underlying decision process.

5.1. Dynamic Adjustment Remains the Dominant Feature of Corporate Capital Structure

The first objective of this study was to examine whether firms gradually adjust their capital structures toward target leverage. The estimated adjustment speeds of approximately 29% to 35% provide strong support for the proposed hypothesis and confirm that leverage adjustment is inherently dynamic rather than instantaneous. The consistency of these estimates across both the fixed-effects and Difference GMM models suggests that the persistence observed in leverage is not driven by model specification but reflects the underlying financing behaviour of firms.

This evidence is highly consistent with the dynamic trade-off theory, which argues that firms rebalance their capital structures gradually because transaction costs, financial distress costs, and financing frictions make immediate adjustment suboptimal (Flannery & Rangan, 2006). Similarly, Lemmon et al. (2008) argue that leverage largely reflects persistent firm-specific characteristics, while Faulkender et al. (2012) show that financing decisions evolve gradually in response to internal financial capacity. The present findings therefore reinforce the view that corporate financing policies are characterized by persistence and gradual convergence towards target leverage rather than continuous optimization. Rather than challenging existing theory, the results provide additional evidence from emerging MENA economies that dynamic adjustment remains a fundamental feature of corporate financing behaviour.

The significance of profitability, market-to-book ratio, tangibility, firm size, and inflation further supports the second hypothesis that leverage adjustment depends on both firm-specific and macroeconomic determinants. Profitability influences firms' reliance on internal financing, consistent with the pecking order theory (Myers & Majluf, 1984), while the significance of market-to-book ratio reflects the importance of growth opportunities predicted by both the trade-off and market timing theories (Baker & Wurgler, 2002). The significant effect of inflation confirms that financing decisions are conditioned by the macroeconomic environment, supporting recent evidence that monetary conditions and financing frictions influence firms' ability to adjust their capital structures (Alter & Elekdag, 2020; Ottonello & Winberry, 2020). Collectively, these findings demonstrate that leverage adjustment reflects the interaction between internal firm characteristics and external economic conditions, consistent with the conceptual framework developed in this study.

5.2. Model Complexity Does Not Necessarily Improve Forecasting Performance

The comparison of alternative forecasting models provides only partial support for the third hypothesis that machine learning models outperform conventional dynamic econometric approaches. Although the Gradient Boosting Machine consistently produces the lowest prediction errors and the highest out-of-sample explanatory power among all machine learning algorithms, its improvement over the dynamic econometric models is relatively limited. At the same time, the fixed-effects and Difference GMM models continue to generate economically meaningful and theoretically interpretable estimates of leverage adjustment.

This result can be explained by the nature of the leverage adjustment process itself. Corporate capital structure evolves through gradual adjustment toward target leverage, implying that much of the predictable variation is already captured by dynamic persistence and observable firm characteristics. Under these conditions, the potential benefits of increasingly flexible learning algorithms become inherently limited because the underlying

economic process is relatively stable and highly autoregressive. Consequently, sophisticated nonlinear models have fewer opportunities to identify additional predictive structure beyond that already contained in lagged leverage and conventional financial variables.

Nevertheless, GBM consistently outperforms the remaining machine learning algorithms because it offers an effective balance between model flexibility and statistical stability. Sequential boosting enables GBM to capture nonlinear interactions among firm characteristics and macroeconomic variables while controlling prediction variance through ensemble learning. By comparison, the neural network performs less effectively because leverage adjustment does not appear to involve sufficiently complex latent nonlinearities to justify a high-capacity deep learning architecture. Instead, the persistence embedded in firms' financing decisions favours algorithms capable of modelling moderate nonlinear relationships without excessive model complexity. These findings therefore support the argument of [Mullainathan and Spiess \(2017\)](#) and [Gu et al. \(2020\)](#) that the effectiveness of machine learning depends on the characteristics of the economic process rather than on algorithmic sophistication alone. The evidence therefore provides only partial support for the third hypothesis: machine learning can improve forecasting performance, but only selected algorithms, particularly GBM, generate meaningful improvements over conventional econometric models.

5.3. Symbolic Financial Knowledge Improves Interpretability Rather than Prediction

The fourth hypothesis proposed that embedding symbolic financial knowledge within machine learning would improve forecasting performance. The empirical evidence does not support this hypothesis. Across alternative values of the symbolic regularisation parameter, predictive accuracy changes only marginally, while the neuro-symbolic model performs similarly to the conventional neural network. These findings indicate that the symbolic constraints do not introduce sufficient independent predictive information to materially improve out-of-sample forecasts.

The ablation analysis helps explain this outcome. Removing individual symbolic rules occasionally produces slightly lower prediction errors, but these improvements are consistently accompanied by weaker compliance with economically meaningful financial constraints. Similarly, the construct validation analysis demonstrates that the Neuro-Symbolic Diagnostics (NSD) index captures information largely distinct from conventional leverage measures, while contributing only modest incremental explanatory power. Together, these findings suggest that symbolic rules encode complementary financial knowledge rather than additional predictive signals.

This interpretation is fully consistent with the objectives of neuro-symbolic artificial intelligence. [Hitzler et al. \(2022\)](#) argue that symbolic reasoning should primarily enhance explainability, transparency, and consistency with domain knowledge rather than maximise predictive accuracy. Within corporate finance, financing decisions are constrained by leverage capacity, debt servicing ability, refinancing risk, and maturity management. Embedding these principles into the learning process, therefore, improves the model's theoretical credibility by ensuring that predictions remain economically plausible. Consequently, the principal contribution of the proposed neuro-symbolic framework lies not in outperforming conventional machine learning but in integrating established financial theory directly into the prediction process. The rejection of the fourth hypothesis should therefore not be interpreted as evidence against neuro-symbolic learning; rather, it indicates that its value in corporate finance lies primarily in improving model interpretability and theoretical consistency, rather than predictive performance alone.

5.4. Contributions to Capital Structure Research

Taken together, the findings contribute to the capital structure literature in three important ways. First, they provide robust evidence from emerging MENA economies that leverage adjustment remains a gradual and persistent process, reinforcing the predictions of dynamic capital structure theory. Second, they demonstrate that although machine learning can improve forecasting performance, the superiority of artificial intelligence is not universal and depends on matching algorithmic complexity to the underlying economic process. Finally, they show that neuro-symbolic artificial intelligence offers value by embedding financial theory within data-driven prediction, thereby improving interpretability and theoretical coherence without sacrificing predictive performance. These findings bridge the gap between traditional corporate finance theory and modern explainable artificial intelligence and suggest that future predictive models should prioritise the integration of economic reasoning with statistical learning rather than pursuing algorithmic complexity alone.

6. Conclusions

6.1. Summary of Findings

This study developed and evaluated an integrated framework combining dynamic econometric models, machine learning algorithms, and a neuro-symbolic artificial intelligence architecture to analyse corporate leverage adjustment in emerging MENA economies. Using a unified empirical framework, the study examined the determinants of leverage adjustment and compared the predictive performance of conventional econometric models, machine learning methods, and theory-guided neuro-symbolic learning.

The empirical findings provide strong evidence that corporate leverage adjustment is a gradual and persistent process. The estimated adjustment speeds obtained from both fixed-effects and Difference GMM models indicate that firms progressively converge toward their target capital structures, providing strong support for the dynamic capital structure hypothesis (H1). Moreover, profitability, market-to-book ratio, tangibility, firm size, and inflation significantly influence leverage adjustment, supporting the hypothesis that both firm-specific characteristics and macroeconomic conditions shape financing decisions (H2). These findings reinforce the predictions of the trade-off theory (Kraus & Litzenberger, 1973), the pecking order theory (Myers & Majluf, 1984), and the dynamic adjustment framework proposed by Flannery and Rangan (2006).

The comparison of forecasting models provides only partial support for the hypothesis that machine learning outperforms conventional econometric models (H3). Among all machine learning algorithms, the Gradient Boosting Machine consistently achieves the highest predictive performance, demonstrating its ability to capture moderate non-linear relationships in structured financial data. However, the improvement over the dynamic econometric models is relatively modest, indicating that much of the predictable variation in leverage adjustment is already explained by persistence and conventional financial variables.

In contrast, the hypothesis that incorporating symbolic financial knowledge improves predictive performance (H4) is not supported. Although the neuro-symbolic framework maintains economically meaningful financial constraints and produces interpretable predictions, it does not significantly improve forecasting accuracy relative to conventional machine learning models. Instead, the symbolic rules primarily enhance theoretical consistency and explainability rather than predictive performance.

6.2. Theoretical and Practical Implications

The study makes several contributions to the corporate finance literature. First, it extends the evidence supporting dynamic capital structure theory by demonstrating that gradual leverage adjustment remains a dominant characteristic of firms operating in emerging MENA economies. The results further confirm that financing decisions are jointly determined by firm-specific characteristics and macroeconomic conditions, reinforcing the complementary roles of the trade-off, pecking order, and dynamic adjustment theories.

Second, the findings contribute to the growing literature on artificial intelligence in finance by demonstrating that increasing model complexity does not necessarily improve forecasting performance. The superior performance of the Gradient Boosting Machine suggests that algorithms capable of modelling moderate nonlinear relationships while maintaining statistical robustness are better suited to corporate finance applications than more computationally intensive neural networks.

Third, the study advances the emerging field of explainable artificial intelligence by demonstrating that the primary contribution of neuro-symbolic learning lies in embedding financial theory directly into predictive models. Rather than improving predictive accuracy, symbolic reasoning enhances model transparency, theoretical coherence, and consistency with established financial principles, thereby providing a foundation for more explainable financial decision-support systems.

From a practical perspective, the findings provide useful implications for corporate managers, financial institutions, and policymakers. Corporate managers can use the identified determinants of leverage adjustment to better understand financing behaviour under changing macroeconomic conditions. Financial institutions may benefit from the proposed Neuro-Symbolic Diagnostics (NSD) index as a complementary measure of debt discipline that extends beyond traditional leverage ratios. For regulators and policymakers, the relatively rapid adjustment speeds observed across MENA firms suggest that corporate balance sheets respond actively to changes in economic conditions, implying that monetary policy and financing conditions can be transmitted relatively quickly through corporate financing decisions.

6.3. Limitations and Future Research

Despite its contributions, this study has several limitations that provide opportunities for future research.

First, the empirical analysis is limited to publicly listed non-financial firms in emerging MENA economies. Future studies could examine whether the findings generalise to private firms, financial institutions, state-owned enterprises, or firms operating in developed markets.

Second, the symbolic financial rules were specified using predetermined theoretical constraints. Future research could investigate adaptive neuro-symbolic frameworks in which symbolic constraints evolve dynamically according to changing macroeconomic conditions, firm characteristics, or reinforcement learning mechanisms.

Third, although several machine learning algorithms were evaluated, recent developments in artificial intelligence, including transformer-based architectures, graph neural networks, Bayesian neural networks, and large language model-assisted financial forecasting, remain largely unexplored in the context of capital structure adjustment and deserve further investigation.

Finally, future studies could extend the proposed framework by incorporating additional dimensions of corporate financial behaviour, including environmental, social, and governance (ESG) performance, corporate governance quality, financial distress, and sustainability-related financing decisions. Comparative analyses across different institu-

tional settings would also improve understanding of whether the observed performance of neuro-symbolic learning is specific to emerging MENA economies or reflects more general characteristics of corporate financing behaviour.

6.4. Concluding Remarks

This study demonstrates that corporate leverage adjustment in emerging MENA economies remains fundamentally a dynamic and economically structured process. Conventional dynamic econometric models continue to provide robust and theoretically consistent estimates, while machine learning algorithms offer incremental improvements in forecasting performance when appropriately matched to the characteristics of the underlying data. Among the evaluated algorithms, the Gradient Boosting Machine provides the strongest predictive performance by balancing nonlinear modelling capability with statistical robustness.

Although the proposed neuro-symbolic framework does not improve forecasting accuracy, it represents an important methodological contribution by integrating financial theory directly into the learning process. The findings therefore suggest that the future of artificial intelligence in corporate finance lies not in replacing established financial theory with increasingly complex algorithms, but in developing intelligent systems that combine the predictive capabilities of machine learning with the transparency, interpretability, and theoretical consistency of domain-specific financial knowledge. Such integration provides a promising direction for the development of explainable and theory-driven decision-support systems in corporate finance.

Author Contributions: Conceptualization, O.S., E.M. and A.H.; methodology, O.S., E.M., T.M.G. and Y.I.; software, O.S. and E.M.; validation, T.M.G., Y.I. and A.H.; formal analysis, O.S., E.M. and T.M.G.; investigation, O.S., E.M. and Y.I.; resources, A.H.; data curation, O.S., E.M. and T.M.G.; writing—original draft preparation, O.S., E.M., T.M.G. and Y.I.; writing—review and editing, O.S., E.M., T.M.G., Y.I. and A.H.; visualization, O.S. and Y.I.; supervision, A.H.; project administration, A.H. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: This study is based exclusively on secondary data obtained from commercial and publicly available databases and does not involve human participants, personal data, or animal subjects. Therefore, ethical approval and informed consent were not required.

Informed Consent Statement: Not applicable.

Data Availability Statement: The firm-level financial data used in this study were obtained from the Refinitiv database under an institutional subscription. Macroeconomic variables were collected from publicly available international sources, including the World Bank's World Development Indicators and the International Monetary Fund, as described in the manuscript. Owing to licensing restrictions imposed by Refinitiv, the firm-level data cannot be made publicly available. Researchers with access to the Refinitiv database can replicate the dataset using the sample selection criteria and variable definitions provided in the manuscript. The processed dataset and programming code used in the analysis are available from the corresponding author upon reasonable request, subject to the licensing terms of the original data providers.

Acknowledgments: We would like to thank the Gulf University for Science and Technology (GUST), Kuwait, for facilitating the completion of this research project. During the preparation of this manuscript, the authors utilized Grammarly and ChatGPT (GPT-5.5) for proofreading purposes and in creating Figures 1 and 2. The authors have thoroughly reviewed and edited the output and assume full responsibility for the content of this publication.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

NSD	Neuro-Symbolic Debt
NSCS	Neuro-Symbolic Capital Structure
FE	Fixed Effects
GMM	Generalized Method of Moments
ML	Machine Learning
AI	Artificial Intelligence
NN	Neural Network
GBM	Gradient Boosting Machine
LASSO	Least Absolute Shrinkage and Selection Operator
RMSE	Root Mean Squared Error
MAE	Mean Absolute Error
RAR	Rule Adherence Rate
ROA	Return on Assets
ROE	Return on Equity
EBITDA	Earnings Before Interest, Taxes, Depreciation, and Amortization
EBIT	Earnings Before Interest and Taxes
PCA	Principal Component Analysis
VIF	Variance Inflation Factor
MENA	Middle East and North Africa
MTB	Market-to-Book Ratio
SOA	Speed of Adjustment
CC BY	Creative Commons Attribution License
DM Test	Diebold–Mariano Test
ST Debt	Short-Term Debt
R ²	Coefficient of Determination

Appendix A

Appendix A.1 Computational Environment

The empirical analyses were implemented entirely using Python 3.12 within a Linux-based computational environment configured to ensure deterministic execution and complete computational reproducibility. The analytical workflow was developed as a fully automated and modular pipeline in which all stages of the empirical analysis, including data preprocessing, financial variable construction, panel econometric estimation, machine learning forecasting, neuro-symbolic optimization, statistical diagnostics, figure generation, and final document assembly, were executed programmatically without manual intervention.

The computational framework follows a sequential execution architecture in which the output generated by each analytical module serves as the validated input for the subsequent stage. This design ensures consistency throughout the empirical analysis and eliminates discrepancies that may arise from manual processing or independent execution of individual procedures. Consequently, all descriptive statistics, econometric estimations, machine learning forecasts, neuro-symbolic diagnostics, robustness analyses, tables, and figures reported in the manuscript originate from a single integrated computational workflow.

All numerical computations were performed using IEEE-754 double-precision floating-point arithmetic. Machine learning algorithms were initialized using fixed random seeds to ensure deterministic estimation and reproducible forecasting results. The primary analyses employed Seed = 42, while robustness analyses were repeated using Seeds {0, 1, 2, 3, 4} to verify the stability of the reported forecasting performance and neuro-symbolic optimization results. Under identical inputs, repeated execution of the complete analytical

pipeline produces identical parameter estimates, predictions, diagnostic statistics, and evaluation metrics.

The computational environment operated within a Linux container using CPU-based execution only, without GPU acceleration. Internet access was unavailable during execution, preventing the installation of additional statistical libraries. As a result, external panel econometric packages, including statsmodels and linearmodels, were not available. The two-way fixed effects estimator and the two-step Arellano-Bond Difference GMM estimator were therefore implemented directly using matrix algebra based on NumPy and pandas, following the estimation procedures. Machine learning estimation, statistical inference, and visualization were performed using specialized Python scientific libraries.

To ensure platform independence and facilitate complete replication, all file paths were defined using relative directory structures. Every script expects its required input files within the working directory and automatically exports standardized outputs that serve as inputs for subsequent analytical stages.

Table A1. Computational Environment.

Component	Detail
Programming language	Python 3.12
Operating environment	Linux container
Hardware	CPU-only execution
GPU acceleration	Not used
Numerical precision	IEEE-754 double precision (64-bit floating point)
Internet access	Disabled during execution
Panel econometric implementation	Custom NumPy and pandas matrix algebra implementations
External econometric libraries	statsmodels and linearmodels unavailable
Random initialization	Primary seed = 42; robustness analysis using Seeds {0, 1, 2, 3, 4}
Computational workflow	Fully automated sequential pipeline
File management	Relative directory structure
Output formats	Excel, CSV, JSON, PNG,
Reproducibility	Deterministic end-to-end computational pipeline

Appendix A.2 Primary Computational Libraries

The computational framework was implemented using a collection of open-source Python scientific computing libraries that collectively support data engineering, panel econometric estimation, machine learning forecasting, statistical inference, visualization, and automated document generation. Each library performs a dedicated function within the computational workflow, ensuring a modular implementation while maintaining complete consistency across all analytical stages.

Data management and numerical computation were performed using pandas and NumPy, which provide efficient manipulation of longitudinal panel datasets, as well as the matrix algebra required for variable construction, panel transformations, and econometric estimation. Because external panel econometric libraries were unavailable within the execution environment, all fixed effects and Difference GMM estimators were implemented directly using NumPy-based matrix operations and pandas data structures, following the algorithms.

Machine learning models were implemented using scikit-learn, which provided the computational framework for Artificial Neural Networks via the MLPRegressor estimator, as well as for LASSO, Ridge Regression, Gradient Boosting Regression, principal compo-

nent analysis (PCA), feature standardization, and hyperparameter optimization. Statistical inference and model evaluation relied on SciPy, which was used to compute Pearson correlation coefficients and associated significance levels, to generate probability distributions required for GMM diagnostic tests, and to perform the Diebold-Mariano predictive accuracy test.

All graphical outputs reported in the manuscript were generated using Matplotlib 3.10.5. Figures were exported as publication-quality PNG files using a standardized blue color palette and image resolutions of 150–200 dpi to ensure consistency throughout the manuscript. Following completion of the statistical analyses, all tables and figures were automatically assembled into the final Microsoft Word manuscript using the docx package, executed via Node.js, thereby eliminating manual formatting and ensuring complete consistency between computational outputs and the reported results.

The modular organization of the computational framework allows each analytical component to be executed independently while remaining compatible with the remaining stages of the replication pipeline. Consequently, modifications to individual modules, such as alternative forecasting algorithms or additional robustness analyses, can be incorporated without altering the overall computational architecture.

Table A2. Primary Computational Libraries.

Library	Primary Function	Application
pandas	Data manipulation	Panel construction, variable engineering, data preprocessing
NumPy	Numerical computing and matrix algebra	Fixed effects estimation, Difference GMM implementation, numerical optimization
scikit-learn	Machine learning	Artificial Neural Networks (MLPRegressor), LASSO, Ridge Regression, Gradient Boosting, StandardScaler, PCA
SciPy	Statistical inference	Pearson correlations, significance tests, probability distributions, GMM diagnostics, Diebold-Mariano test
Matplotlib	Data visualization	Publication-quality figures (PNG, 150 to 200 dpi)
Node.js (docx package)	Automated document generation	Assembly of the final Microsoft Word manuscript

Appendix A.3 Data Engineering and Preprocessing Pipeline

The data engineering pipeline was developed as a deterministic multi-stage process to transform the raw Refinitiv financial dataset into a clean and internally consistent panel suitable for econometric estimation, machine learning forecasting, and neuro-symbolic optimization. Every preprocessing operation was executed programmatically using a predefined sequence of validation and transformation procedures, thereby eliminating manual intervention and ensuring complete reproducibility of the analytical dataset.

Raw observations were imported directly from the Refinitiv source file and organized into a longitudinal panel using the WORLDScope IDENTIFIER (FIRM_ID) as the unique firm identifier. During the initial validation stage, the computational pipeline identified and removed 16 placeholder observations that contained no valid ticker symbol, ISIN, or financial statement information. The pipeline subsequently verified the uniqueness of every FIRM_ID-Year observation.

Following panel validation, all accounting variables, financial ratios, lagged variables, rolling statistics, and model-specific predictors were constructed automatically using

firm-level chronological ordering. Missing values were handled using a hierarchical imputation procedure designed to preserve firm-specific information whenever possible. The procedure first applied linear interpolation within individual firms. Remaining missing observations were then replaced using the corresponding nation-year median. Any values still missing after these steps were replaced using the overall global median. This hierarchical approach maximized data completeness while preserving both cross-sectional and temporal consistency.

Continuous ratio variables were subsequently winsorized at the 1st and 99th percentiles to reduce the influence of extreme observations without removing valid firm-year records. Winsorization was performed after ratio construction and prior to model estimation to ensure that both econometric estimators and machine learning algorithms were not disproportionately influenced by outlying values.

To accommodate cross-country heterogeneity, variables used in the structural component of the neuro-symbolic framework were standardized within each nation-year cell using z-score normalization. This standardization was applied exclusively to variables entering the structural principal component analysis and did not affect the remaining estimation variables. Principal Component Analysis (PCA) was then used to construct the composite leverage indicator employed within the neuro-symbolic diagnostic framework.

The final processed dataset constituted the common analytical input for every econometric model, forecasting algorithm, robustness analysis, and neuro-symbolic optimization procedure. Consequently, all empirical results reported in the manuscript were generated from an identical validated dataset using a single deterministic preprocessing pipeline.

Table A3. Data Engineering Configuration.

Setting	Configuration
Firm identifier	WORLDSCOPE IDENTIFIER (FIRM_ID), not ticker symbol
Panel structure	Firm-Year longitudinal panel
Placeholder observations removed	16 observations containing no valid ticker, ISIN, or financial information
Duplicate validation	FIRM_ID-Year uniqueness verified at two stages of the pipeline
Missing-value imputation	Firm-level linear interpolation → nation-year median → global median
Variable construction	Financial ratios, lagged variables, rolling statistics, and model predictors
Outlier treatment	Winsorization at the 1st and 99th percentiles
Cross-country standardization	Nation-year z-score normalization for structural PCA variables only
Principal Component Analysis	Composite leverage factor construction
Data integrity	Automatic execution stops if duplicate firm-year observations are detected
Analytical output	Fully validated analysis panel used throughout the replication workflow

Appendix A.4 Econometric Estimation Configuration

The econometric estimation framework was implemented as a hierarchical sequence of panel estimators designed to account for unobserved heterogeneity, dynamic adjustment, and potential endogeneity within the capital structure model. All estimations were performed using the same validated analytical dataset generated during the preprocessing stage, thereby ensuring complete comparability across estimation methods and robustness analyses.

The baseline specification employed a two-way fixed-effects estimator implemented with the exact within transformation. Firm-specific and year-specific means were removed from each variable prior to estimation, thereby eliminating both time-invariant firm heterogeneity and common temporal effects. Parameter estimation was subsequently performed

using ordinary least squares on the transformed data, while statistical inference relied on firm-clustered sandwich covariance estimators to account for heteroskedasticity and within-firm serial correlation.

To address the dynamic nature of leverage adjustment and the potential endogeneity of the lagged dependent variable, the study employed the two-step Arellano-Bond Difference Generalized Method of Moments (GMM) estimator. Endogeneity was controlled through internally generated instruments constructed from lagged values of the dependent variable and differenced regressors. Instrument proliferation was limited by using collapsed instrument matrices, thereby improving estimator stability and reducing finite-sample bias.

Two Difference GMM specifications were estimated. The full specification employed collapsed lags two through four of the dependent variables, together with the differenced values of the remaining endogenous regressors, yielding a total of seven instruments for five estimated parameters. A restricted specification was subsequently estimated using only the second lag of the dependent variable, producing an exactly identified model with three estimated parameters. The restricted specification served as an additional robustness analysis to evaluate the sensitivity of the results to instrument selection.

Model adequacy was evaluated using the standard diagnostic procedures for dynamic panel estimation. Serial correlation in the differenced residuals was assessed using the Arellano-Bond AR(1) and AR(2) tests, while instrument validity was evaluated using the Sargan and Hansen over-identification tests. These diagnostics were computed automatically following each GMM estimation and incorporated into the reported empirical results.

Because the computational environment did not permit the installation of specialized econometric libraries, all fixed-effects and Difference GMM estimators were implemented directly using NumPy and pandas-based matrix algebra, following the algorithms. The two-step covariance estimator was not adjusted using the Windmeijer finite-sample correction, and this limitation should be considered when interpreting the reported standard errors.

Table A4. Econometric Configuration.

Setting	Configuration
Fixed effects estimator	Two-way exact within transformation (firm and year demeaning)
Estimation method	Ordinary Least Squares on transformed data
Covariance estimator	Firm-clustered sandwich covariance estimator
Dynamic estimator	Two-step Arellano-Bond Difference GMM
Instrument construction (full specification)	Collapsed lags 2 to 4 of the dependent variable (3 columns) together with differenced endogenous regressors (4 columns), producing 7 instruments for 5 parameters
Instrument construction (restricted specification)	Lag 2 of the dependent variable only, producing an exactly identified model with 3 parameters
Instrument reduction	Collapsed instrument matrix
Diagnostic tests	Arellano-Bond AR(1), Arellano-Bond AR(2), Sargan test, Hansen J-test
Computational implementation	Custom NumPy and pandas matrix algebra implementation
Small-sample correction	Windmeijer correction not applied
Analytical output	Fixed effects estimates, Difference GMM estimates, diagnostic statistics, and robustness analyses

Appendix A.5 Machine Learning Forecasting Configuration

The machine learning framework was developed to evaluate the out-of-sample predictive performance of alternative forecasting algorithms relative to the econometric benchmark models. All machine learning models were estimated using the same validated analytical dataset generated during the preprocessing stage and trained within a common rolling-window forecasting framework to ensure a fair comparison across competing algorithms.

The primary forecasting model was an Artificial Neural Network (ANN) implemented using the MLPRegressor algorithm in scikit-learn. The network consisted of two hidden layers with 64 and 32 neurons, respectively, using Rectified Linear Unit (ReLU) activation functions, and a linear output layer for continuous prediction. Model parameters were optimized using the Adam optimizer with a learning rate of 0.001, L2 regularization of 0.0001, and a mini-batch size of 256 observations.

To reduce overfitting and improve model generalization, the neural network employed an early stopping procedure based on a validation sample comprising 15% of the training observations. Training terminated automatically if validation performance did not improve by at least 0.0001 for 15 consecutive epochs, subject to a maximum of 200 training epochs. These settings ensured stable convergence while limiting unnecessary model complexity.

Three additional machine learning algorithms were estimated for comparative evaluation. LASSO regression was estimated using a regularization parameter of 0.001, while Ridge regression employed a penalty parameter of 1.0. Gradient Boosting Regression was implemented with 200 decision trees, a maximum tree depth of 3, and a learning rate of 0.05. Collectively, these models represent linear, regularized, ensemble, and non-linear forecasting approaches that provide a comprehensive benchmark for evaluating predictive performance.

Forecast evaluation followed an expanding rolling-window design to preserve the temporal ordering of financial information and eliminate look-ahead bias. For each forecasting iteration, models were trained using observations through year 2016 + k , validated using observations from year 2017 + k , and evaluated on the subsequent test year 2018 + k , where $k = 0, 1, \dots, 6$. The testing period was capped at 2024, ensuring that each prediction was generated exclusively from information available prior to the forecast horizon.

Model selection was based on a systematic hyperparameter optimization procedure. Candidate neural network architectures included hidden layer structures of (32), (64, 32), and (128, 64) neurons. The search additionally evaluated L2 regularization values of 10^{-5} , 10^{-4} , and 10^{-3} , learning rates of 10^{-4} , 10^{-3} , and 10^{-2} , and neuro-symbolic weighting parameters ranging from 0.0 to 1.0 in increments of 0.2. The optimal configuration was selected on basis to out-of-sample validation performance prior to final model estimation.

To ensure computational reproducibility, all stochastic learning algorithms were initialized using a primary random seed of 42. Forecasting performance was subsequently re-evaluated using random seeds 0, 1, 2, 3, and 4 to verify that the reported predictive results were not sensitive to initialization. The resulting forecasting models therefore provide reproducible and robust estimates of predictive performance across multiple machine learning paradigms.

Table A5. Machine Learning Configuration.

Setting	Configuration
Primary forecasting model	Artificial Neural Network (MLPRegressor)
Hidden layer architecture	Two hidden layers (64, 32)
Activation function	ReLU

Table A5. Cont.

Setting	Configuration
Output layer	Linear
Optimization algorithm	Adam
Learning rate	0.001
L2 regularization	0.0001
Batch size	256
Early stopping	15% validation split
Patience	15 epochs
Convergence tolerance	0.0001
Maximum epochs	200
LASSO penalty (α)	0.001
Ridge penalty (α)	1.0
Gradient Boosting	200 trees, maximum depth = 3, learning rate = 0.05
Forecast evaluation	Expanding rolling-window validation
Training period	Years $\leq 2016 + k$
Validation period	Year = 2017 + k
Testing period	Year = 2018 + k (maximum year 2024)
Hyperparameter search	Hidden layers {(32), (64,32), (128,64)}, L2 $\{10^{-5}, 10^{-4}, 10^{-3}\}$, learning rate $\{10^{-4}, 10^{-3}, 10^{-2}\}$, λ {0.0, 0.2, 0.4, 0.6, 0.8, 1.0}
Random initialization	Primary seed = 42; robustness analysis using seeds {0, 1, 2, 3, 4}

Appendix A.6 Neuro-Symbolic Optimization Configuration

The neuro-symbolic optimization framework was developed to integrate data-driven prediction with economically interpretable financial rules. Rather than relying solely on statistical learning, the framework incorporates domain-specific constraints through a differentiable penalty function that guides model optimization while preserving flexibility in the learning process. This hybrid approach enables the forecasting model to improve predictive accuracy while remaining consistent with established financial principles governing leverage adjustment, debt capacity, and refinancing behavior.

To facilitate gradient-based optimization, all symbolic rules were transformed into differentiable penalty functions using the Softplus approximation. A smoothing parameter of $\beta = 10$ was applied uniformly across all symbolic constraints to provide a close approximation to the corresponding piecewise penalty functions while ensuring numerical stability during optimization.

The leverage stability rule was defined using a dynamic tolerance band centered on the nation-year median leverage. The allowable deviation from this benchmark was determined by the firm's rolling five-year standard deviation, multiplied by a scaling parameter $\kappa = 1.0$. Firms whose predicted leverage exceeded the upper or lower limits of this adaptive interval incurred a symbolic penalty proportional to the magnitude of the violation.

The interest coverage rule identified firms exhibiting elevated financial distress using the 10th percentile of the in-sample interest coverage distribution as the minimum acceptable threshold. Predictions for firms below this threshold incurred an additional symbolic penalty, thereby discouraging leverage recommendations inconsistent with firms' debt-servicing capacity.

Debt maturity structure was evaluated using the firm's short-term debt share relative to the corresponding nation-year interquartile range. The admissible maturity interval was defined using an interquartile range multiplier of 1.0, allowing moderate variation across countries while penalizing predictions that implied unusually concentrated debt maturity profiles.

Refinancing risk was incorporated through a regional market volatility indicator. Periods in which lagged regional volatility exceeded the 90th percentile of its historical distribution were classified as high-refinancing-stress regimes. During these periods, symbolic penalties increased for leverage recommendations that implied additional refinancing pressure under adverse market conditions.

The overall symbolic objective function combined the five rule-specific penalties with equal weights, assigning each rule a weight of 1.0. The resulting composite penalty was added to the prediction loss function, thereby balancing predictive accuracy and economic consistency throughout model estimation. Optimization terminated when the relative improvement in the objective function fell below a convergence tolerance of $\varepsilon = 10^{-6}$, ensuring stable parameter estimation while avoiding unnecessary computational iterations.

The neuro-symbolic framework therefore provides an interpretable extension of conventional machine learning by embedding theoretically motivated financial constraints directly within the optimization process. This approach allows the model to generate forecasts that are both statistically accurate and economically consistent with established principles of corporate financing behavior.

Table A6. Neuro-Symbolic Optimization Configuration.

Setting	Configuration
Penalty relaxation	Softplus approximation
Softplus smoothing parameter (β)	10
Leverage band parameter (κ)	1.0
Leverage reference	Nation-year median leverage
Dynamic tolerance	Firm rolling five-year standard deviation
Interest coverage threshold	10th percentile of the in-sample distribution
Debt maturity rule	Nation-year interquartile range of short-term debt share
Interquartile range multiplier	1.0
Refinancing stress threshold	90th percentile of lagged regional volatility
Rule weighting	Equal weights ($w = 1.0$) for all five symbolic rules
Convergence tolerance	$\varepsilon = 10^{-6}$
Optimization objective	Prediction loss plus symbolic penalty
Computational output	Rule penalties, composite symbolic loss, optimized forecasts

Appendix A.7 Forecast Evaluation and Computational Validation

The predictive performance of all econometric and machine learning models was evaluated using a consistent set of statistical measures of accuracy and validation procedures. These evaluation metrics quantify forecasting performance, assess statistical differences between competing models, and examine the extent to which the neuro-symbolic

framework satisfies the embedded financial rules. All evaluation statistics were computed automatically from the out-of-sample forecasts generated by the rolling-window validation framework.

Forecast accuracy was assessed using three conventional regression performance measures: Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), and the coefficient of determination (R^2). RMSE measures the average magnitude of prediction errors while assigning greater weight to larger deviations because of the squared-error loss function. MAE provides a complementary measure based on the absolute prediction error and is less sensitive to extreme observations. The coefficient of determination evaluates the proportion of variation in the observed dependent variable explained by each forecasting model.

To determine whether differences in predictive performance between competing forecasting models were statistically significant, the study employed the Diebold-Mariano (DM) test of predictive accuracy. The DM statistic was calculated using squared forecast-error loss differentials, together with a Newey-West heteroskedasticity- and autocorrelation-consistent covariance estimator to account for serial dependence in the forecast errors. Statistical significance was evaluated using a two-sided test based on the asymptotic normal distribution.

The neuro-symbolic framework was additionally evaluated using the Rule Adherence Rate (RAR), which measures the proportion of model predictions satisfying the symbolic financial constraints incorporated into the optimization procedure. The RAR provides a direct assessment of the degree to which the forecasting model produces economically consistent predictions while maintaining predictive accuracy.

The computational pipeline generated all evaluation statistics automatically after model estimation and forecasting. This automated implementation ensured that every reported performance measure was computed using identical test observations and evaluation procedures, thereby maintaining consistency across econometric models, machine learning algorithms, and neuro-symbolic forecasting approaches.

Table A7. Forecast Evaluation Metrics.

Metric	Purpose
Root Mean Squared Error (RMSE)	Measures average squared prediction error
Mean Absolute Error (MAE)	Measures average absolute prediction error
Coefficient of Determination (R^2)	Measures explanatory and predictive performance
Diebold-Mariano Test	Compares predictive accuracy between competing forecasting models using Newey-West adjusted variance
Rule Adherence Rate (RAR)	Measures compliance with the symbolic financial constraints embedded in the neuro-symbolic framework
Evaluation sample	Out-of-sample forecasts generated from the expanding rolling-window validation procedure

Appendix A.8 Computational Reproducibility Protocol

The complete empirical analysis was implemented as a deterministic computational workflow in which each stage of the analytical process was executed sequentially using dedicated scripts. Every script receives standardized input files, performs a predefined analytical task, and generates validated outputs that serve as inputs for subsequent stages of the pipeline. This modular architecture ensures complete consistency across all reported empirical results while eliminating manual intervention during data processing, estimation, and document preparation.

The computational workflow begins with reconstruction of the analytical panel from the raw Refinitiv dataset. The preprocessing script validates firm identifiers, removes placeholder observations, constructs all financial variables and ratios, imputes missing values according to the predefined hierarchical procedure, applies winsorization to ratio variables, and generates the principal component analysis leverage factor required for the neuro-symbolic framework. The resulting validated analytical dataset serves as the common input for all subsequent econometric and machine-learning analyses.

Following data preparation, dedicated scripts generate the descriptive statistics, correlation analyses, econometric estimates, robustness analyses, machine-learning forecasting results, neuro-symbolic diagnostics, and validation statistics reported throughout the manuscript. Additional scripts produce the publication-quality figures directly from the computational outputs generated during model estimation. Finally, all tables and figures are automatically assembled into the final Microsoft Word manuscript using a Node.js document-generation script.

All scripts were executed sequentially within the same computational environment. Every script expects its required input files within the current working directory and automatically exports standardized output files that are consumed by subsequent stages of the workflow. Because the entire analytical process is executed programmatically, every numerical result reported in the manuscript, including descriptive statistics, parameter estimates, diagnostic tests, forecasting results, robustness analyses, tables, and figures, is generated directly from executable source code rather than manually copied or edited.

The computational workflow also incorporates multiple integrity checks designed to ensure reproducibility and data consistency. Duplicate FIRM_ID-Year observations are verified at two independent stages of the preprocessing pipeline, and execution terminates automatically if duplicate records are detected. Intermediate datasets generated throughout the workflow are retained to facilitate verification of every analytical stage and to provide a transparent audit trail from the original data source to the final reported results. Consequently, independent execution of the complete workflow using the same input files produces identical empirical findings.

Table A8. Computational Reproducibility Checklist.

Step	Script	Purpose
1	prepare_analysis_dataset.py	Builds the validated Analysis_Panel.xlsx from the raw Refinitiv data, constructs variables, imputes missing values, applies winsorization, builds the PCA leverage factor, and performs integrity checks.
2	build_tables_4_5_7_8_9.py	Generates Tables 4, 5 and 7–9, including descriptive statistics, correlation analysis, Difference GMM estimations, and speed-of-adjustment comparison.
3	build_remaining_tables.py	Generates Tables 1–3 and 6, performs the symbolic-rule ablation study, and exports fixed-effects metadata required for subsequent analyses.
4	build_figure9.py	Produces the datasets required for Figure 9, illustrating partial adjustment under alternative monetary conditions.
5	results_analysis_v2.py	Generates Tables 10–13 and exports the datasets required for Figures 8, 10 and 11, including machine learning forecasting results and neuro-symbolic diagnostics.
6	nsd_composite_index.py	Constructs the NSD Equal and NSD PCA indices, performs walk-forward prediction, and evaluates construct validity through redundancy analysis, incremental explanatory power, and variance decomposition.
7	build_all_figures.py	Generates publication-quality Figures 3–11 from the outputs produced by the preceding scripts.

References

- Almeida, H., Campello, M., & Weisbach, M. S. (2011). Corporate financial and investment policies when future financing is not frictionless. *Journal of Corporate Finance*, 17(3), 675–693. [CrossRef]
- Alter, A., & Elekdag, S. (2020). Emerging market corporate leverage and global financial conditions. *Journal of Corporate Finance*, 62, 101590. [CrossRef]
- Angrist, J. D., & Pischke, J.-S. (2009). *Mostly harmless econometrics: An empiricist's companion*. Princeton University Press. [CrossRef]
- Arellano, M. (1987). Computing robust standard errors for within-groups estimators. *Oxford Bulletin of Economics and Statistics*, 49(4), 431–434. [CrossRef]
- Arellano, M., & Bond, S. (1991). Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations. *The Review of Economic Studies*, 58(2), 277–297. [CrossRef]
- Athey, S., & Imbens, G. W. (2019). Machine learning methods that economists should know about. *Annual Review of Economics*, 11(1), 685–725. [CrossRef]
- Baker, M., & Wurgler, J. (2002). Market timing and capital structure. *Journal of Finance*, 57(1), 1–32. [CrossRef]
- Barclay, M. J., & Smith, C. W. (1995). The maturity structure of corporate debt. *The Journal of Finance*, 50(2), 609–631. [CrossRef]
- Dang, V. A., Kim, M., & Shin, Y. (2012). Asymmetric capital structure adjustments: New evidence from dynamic panel threshold models. *Journal of Empirical Finance*, 19(4), 465–482. [CrossRef]
- DeMarzo, P. M., & He, Z. (2021). Leverage dynamics without commitment. *Journal of Finance*, 76(3), 1195–1250. [CrossRef]
- Diebold, F. X., & Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business and Economic Statistics*, 13(3), 253–263. [CrossRef]
- Faulkender, M., Flannery, M. J., Hankins, K. W., & Smith, J. M. (2012). Cash flows and leverage adjustments. *Journal of Financial Economics*, 103(3), 632–646. [CrossRef]
- Flannery, M. J., & Rangan, K. P. (2006). Partial adjustment toward target capital structures. *Journal of Financial Economics*, 79(3), 469–506. [CrossRef]
- Frank, M. Z., & Goyal, V. K. (2009). Capital structure decisions: Which factors are reliably important? *Financial Management*, 38(1), 1–37. [CrossRef]
- Friedman, J. H. (2001). Greedy function approximation: A gradient boosting machine. *The Annals of Statistics*, 29(5), 1189–1232. [CrossRef]
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT Press. Available online: <https://www.deeplearningbook.org> (accessed on 28 December 2025).
- Graham, J. R., & Leary, M. T. (2011). A review of empirical capital structure research and directions for the future. *Annual Review of Financial Economics*, 3, 309–345. [CrossRef]
- Gu, S., Kelly, B., & Xiu, D. (2020). Empirical asset pricing via machine learning. *The Review of Financial Studies*, 33(5), 2223–2273. [CrossRef]
- Hamida, A., Colot, O., & Kechad, R. (2025). Do board characteristics influence leverage and debt maturity? Empirical evidence from a transitional economy. *Journal of Risk and Financial Management*, 18(8), 418. [CrossRef]
- Hastie, T., Tibshirani, R., & Friedman, J. (2009). *The elements of statistical learning* (Springer series in statistics). Springer. [CrossRef]
- Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly*, 28(1), 75–105. [CrossRef]
- Hitzler, P., Eberhart, A., Ebrahimi, M., Sarker, M. K., & Zhou, L. (2022). Neuro-symbolic approaches in artificial intelligence. *National Science Review*, 9(6), nwac035. [CrossRef] [PubMed]
- Jolliffe, I. (2002). *Principal component analysis* (2nd ed.). Springer. [CrossRef]
- Jungherr, J., Meier, M., Reinelt, T., & Schott, I. (2024). *Corporate debt maturity matters for monetary policy* (International Finance Discussion Paper No. 1402). Board of Governors of the Federal Reserve System. [CrossRef]
- Kraus, A., & Litzenberger, R. H. (1973). A state-preference model of optimal financial leverage. *The Journal of Finance*, 28(4), 911–922. [CrossRef]
- Lemmon, M. L., Roberts, M. R., & Zender, J. F. (2008). Back to the beginning: Persistence and the cross-section of corporate capital structure. *Journal of Finance*, 63(4), 1575–1608. [CrossRef]
- Missaoui, I., & Brahmi, M. (2025). Decision to determine the optimal firm's capital structure: A systematic literature review. In *Converging economic policy, corporate strategy, and technology for emerging economies* (pp. 1–20). IGI Global. [CrossRef]
- Mo, Y., Wei, F., & Huang, Y. (2024). Does fulfilling ESG responsibilities curb corporate leverage manipulation? Evidence from Chinese-listed companies. *Sustainability*, 16(13), 5543. [CrossRef]
- Mullainathan, S., & Spiess, J. (2017). Machine learning: An applied econometric approach. *Journal of Economic Perspectives*, 31(2), 87–106. [CrossRef]
- Myers, S. C. (1984). The capital structure puzzle. *The Journal of Finance*, 39(3), 574–592. [CrossRef]

- Myers, S. C., & Majluf, N. S. (1984). Corporate financing and investment decisions when firms have information that investors do not have. *Journal of Financial Economics*, 13(2), 187–221. [[CrossRef](#)]
- Opler, T., Pinkowitz, L., Stulz, R., & Williamson, R. (1999). The determinants and implications of corporate cash holdings. *Journal of Financial Economics*, 52(1), 3–46. [[CrossRef](#)]
- Ottonello, P., & Winberry, T. (2020). Financial heterogeneity and the investment channel of monetary policy. *Econometrica*, 88(6), 2473–2502. [[CrossRef](#)]
- Peffer, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007). A design science research methodology for information systems research. *Journal of Management Information Systems*, 24(3), 45–77. [[CrossRef](#)]
- Rajan, R. G., & Zingales, L. (1995). What do we know about capital structure? Some evidence from international data. *The Journal of Finance*, 50(5), 1421–1460. [[CrossRef](#)]
- Roodman, D. (2009). A note on the theme of too many instruments. *Oxford Bulletin of Economics and Statistics*, 71(1), 135–158. [[CrossRef](#)]
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1), 267–288. [[CrossRef](#)]
- Tulcanaza-Prieto, A. B., Lee, Y., & Anzules-Falcones, W. (2024). The moderating role of corporate governance in the relationship between leverage and firm value: Evidence from the Korean market. *Risks*, 12(1), 11. [[CrossRef](#)]
- Wang, Z., Bian, B., & Wang, J. (2025). Influence of green credit policy on corporate risk-taking: The mediating effect of debt maturity mismatch and the moderating effect of executive compensation. *Sustainability*, 17(7), 2862. [[CrossRef](#)]
- Wooldridge, J. M. (2010). *Econometric analysis of cross section and panel data* (2nd ed.). MIT Press.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.