Logistic and Operation Management Research, Vol. 5 No.2 (2026) https://doi.org/10.31098/lomr.v5i2.4859

Check for updates

Research Paper

AFMF-SFBM-A09: The Alaali Behavioral Sustainability Forecasting Artificial Intelligence Model (A-BSF-AI): Development and Comparative Evaluation Through Synthetic Simulations

Hasan Mohamed Husain Alaali*
DFAS Research Center, Bahrain

Received: June 22, 2026	Revised : September 08, 2026	Accepted : October 22, 2026	Online : November 20, 2026

Abstract

This study develops the Alaali Behavioral Sustainability Forecasting Artificial Intelligence Model (A-BSF-AI; AFMF-SFBM-A09) as an adaptive extension of the Alaali Behavioral Sustainability Forecasting Model (A-BSF) within the Alaali Financial Models Meta-Framework (AFMF). It evaluates whether sequential coefficient updating improves forecasting accuracy relative to fixed and alternative adaptive estimation methods under controlled synthetic conditions. The experiment comprises 30 independently seeded simulation runs, each containing 15,000 observations across 10 temporal periods, producing 450,000 observations. The data-generating process incorporates behavioral, environmental, social, and governance (ESG), market, cyclical, and time-dependent influences. Five forecasting specifications are compared: restricted ordinary least squares (OLS), correctly specified OLS, rolling OLS, recursive least squares (RLS), and adaptive clipped normalized least mean squares (NLMS). Chronological partitioning allocates 70% of observations to training, 15% to validation, and 15% to testing. Forecasting performance is assessed using root mean squared error (RMSE) and paired statistical comparisons. Adaptive NLMS achieves a mean testing RMSE of 0.120378, compared with 0.138282 for restricted OLS, representing a 12.94% improvement ($p < 0.001$). It also outperforms the tested RLS specification ($p < 0.001$). However, correctly specified OLS achieves the lowest mean RMSE (0.119822), while rolling OLS performs comparably to adaptive NLMS ($p = 0.118$). These findings demonstrate a conditional forecasting advantage for sequential adaptation rather than universal superiority. The study additionally proposes the Behavioral Intelligence Coefficient (BIC), Ethical Responsibility Index (ERI), and Intelligent Maturity Index (IMI) as conceptual indicators requiring separate operational validation. Its contribution lies in developing the A-BSF-AI architecture and providing a reproducible comparative computational assessment of adaptive behavioral sustainability forecasting within AFMF and Dynamic Financial Applied Meta-Science (DFAS).

Keywords: *Behavioral Sustainability, Adaptive Forecasting, Artificial Intelligence, A-BSF-AI, AFMF, DFAS, Synthetic Simulation.*

INTRODUCTION

Research Background

Financial sustainability concerns an organization's capacity to maintain financial viability and operational continuity while responding to changing managerial behavior, market conditions, and environmental uncertainty. These challenges are particularly relevant to operational planning and supply chain management, where variations in demand, resource availability, supplier performance, and sustainability requirements can affect financial forecasts and managerial decisions. Forecasting systems that rely on historical relationships may experience declining predictive accuracy when the underlying relationships evolve over time.

Adaptive forecasting offers a computational approach to addressing such changes through sequential parameter updating. Adaptive control theory establishes the principles of feedback-based adjustment under changing conditions (Åström & Wittenmark, 1995), while system

Copyright Holder:
© Alaali. (2026)
Corresponding author's email: hasan.mohd.alaali@gmail.com
This Article is Licensed Under:

dynamics and organizational cybernetics emphasize feedback, regulation, and temporal interactions within complex organizations (Forrester, 1961; Beer, 1972). These foundations support the examination of forecasting systems capable of updating estimated relationships as new observations become available.

However, adaptive forecasting should not be equated automatically with advanced artificial intelligence or reinforcement learning. Normalized least mean squares (NLMS) provides a sequential gradient-based estimation mechanism, whereas reinforcement learning involves the optimization of decision policies through environmental interaction (Sutton & Barto, 2018). Distinguishing these computational mechanisms is essential when evaluating the actual capabilities of proposed intelligent forecasting systems.

Within the Alaali Financial Models Meta-Framework (AFMF), the Alaali Self-Funded Behavioral Model (A-SFBM; AFMF-SFBM-F01) establishes a conceptual relationship between behavioral discipline and self-funded financial efficiency (Alaali, 2025a). Subsequent developments introduce behavioral efficiency measurement, predictive assessment, cost recovery, and self-regulation (Alaali, 2025b, 2025c, 2025d, 2025e). The Alaali Expected Behavioral Efficiency Index with Adaptive Intelligence (A-EBEI-AI; AFMF-SFBM-A06), Alaali Behavioral Sustainability Index (A-BSI; AFMF-SFBM-A07), and Alaali Behavioral Sustainability Forecasting Model (A-BSF; AFMF-SFBM-A08) extend this sequence toward adaptive intelligence and sustainability forecasting (Alaali, 2025f, 2025g, 2025h).

The present study develops the Alaali Behavioral Sustainability Forecasting Artificial Intelligence Model (A-BSF-AI; AFMF-SFBM-A09) and evaluates its implemented adaptive forecasting mechanism through comparative synthetic simulations. Its practical relevance concerns the potential application of sequential forecasting adjustments to behavioral, ESG, and market-related uncertainty in organizational planning. Such applications remain prospective because the experiment does not use observed organizational or supply chain data.

Research Problem and Literature Gap

A central forecasting challenge arises when relationships estimated from historical observations change over time. In financial and operational environments, evolving behavioral conditions, ESG requirements, market disturbances, and temporal effects may reduce the reliability of fixed-coefficient forecasting models. This problem is relevant to sustainability-related planning because inaccurate forecasts can affect resource allocation, operational coordination, and managerial risk assessment.

Adaptive estimation provides several established approaches for responding to changing relationships. Rolling OLS periodically re-estimates coefficients using recent observations, recursive least squares (RLS) updates estimates sequentially, and NLMS adjusts coefficients through normalized prediction-error feedback. These methods differ in their updating mechanisms, computational characteristics, and sensitivity to model specification.

An important methodological issue is whether an apparent forecasting advantage arises from sequential adaptation or from deficiencies in the comparison model. A fixed forecasting model that omits a relevant time-dependent predictor may perform worse than an adaptive model even when the adaptive procedure does not possess an inherent advantage over appropriately specified alternatives. Consequently, comparisons against restricted fixed estimation alone are insufficient to establish the incremental value of adaptive forecasting.

Adaptive control theory provides the foundation for sequential adjustment (Åström & Wittenmark, 1995), while managerial control theory emphasizes monitoring, feedback, and organizational accountability (Simons, 1995; Merchant & Van der Stede, 2017). Sustainability accounting further establishes the relevance of environmental and social considerations to

organizational performance (Elkington, 1997; Gray & Bebbington, 2001). Nevertheless, these theoretical foundations do not independently establish the comparative predictive performance of the proposed A-BSF-AI mechanism.

Within AFMF, the Alaali Self-Regulation Multiplier (A-SRM), A-BSI, and A-BSF develop conceptual approaches to behavioral regulation, sustainability assessment, and forecasting (Alaali, 2025e, 2025g, 2025h). The specific computational gap addressed by the present study concerns the absence of a documented comparative evaluation of sequential coefficient updating within A-BSF-AI against both fixed and established adaptive estimation alternatives.

Accordingly, the experiment compares adaptive NLMS with restricted OLS, correctly specified OLS, rolling OLS, and RLS. The synthetic data-generating process includes a known time–ESG interaction that is deliberately excluded from the restricted predictor specification but incorporated into correctly specified OLS. This design makes it possible to distinguish the performance consequences of omitted-variable specification from the benefits of sequential adaptation.

The study therefore addresses a bounded methodological gap: determining whether the implemented A-BSF-AI forecasting procedure provides measurable predictive improvements under controlled time-dependent synthetic conditions and whether those improvements persist against stronger forecasting benchmarks. It does not claim that NLMS is a newly invented estimation algorithm or that the broader AFMF architecture has been empirically validated.

Research Question and Objectives

The study investigates the comparative forecasting performance of sequential adaptive coefficient updating under controlled synthetic conditions, with particular attention to changing behavioral and ESG-related relationships.

The principal research question is:

"Does the implemented A-BSF-AI adaptive forecasting mechanism improve predictive accuracy relative to restricted fixed estimation, correctly specified fixed estimation, and alternative adaptive forecasting methods under repeated synthetic simulation conditions?"

The study addresses four objectives:

1. To specify the conceptual and mathematical architecture of A-BSF-AI as an adaptive extension of A-BSF, while distinguishing proposed components from computationally implemented mechanisms.
2. To construct a reproducible synthetic forecasting experiment comprising 30 independently seeded simulation runs incorporating behavioral, ESG, market, cyclical, and time-dependent influences.
3. To evaluate adaptive NLMS against restricted OLS, correctly specified OLS, rolling OLS, and RLS using chronological forecasting assessment, RMSE, and paired statistical comparisons.
4. To determine the extent to which observed forecasting improvements are attributable to sequential adaptation rather than restricted predictor specification, while identifying the limits of computational validation.

The first objective concerns conceptual development, whereas the second and third objectives concern computational implementation and quantitative evaluation. The fourth objective integrates the comparative findings and establishes the boundaries of their interpretation.

These objectives do not require empirical claims concerning organizational sustainability, ethical autonomy, or operational improvements that are not directly measured by the experiment.

Research Contributions

The study makes three principal contributions

First, it develops A-BSF-AI (AFMF-SFBM-A09) as an adaptive extension of the behavioral sustainability forecasting sequence within AFMF. Building on A-BSF and earlier developments in behavioral efficiency, self-regulation, and adaptive intelligence, the proposed architecture connects sustainability-related forecasting with sequential learning and feedback-based parameter adjustment (Alaali, 2025a, 2025b, 2025c, 2025d, 2025e, 2025f, 2025g, 2025h). Its broader behavioral and governance components remain conceptual rather than computationally validated.

Second, the study provides a reproducible comparative simulation comprising 30 independent runs and 450,000 synthetic observations. The experiment evaluates five forecasting specifications using chronological training, validation, and testing procedures. This design extends beyond comparison with a restricted fixed baseline by incorporating correctly specified OLS, rolling OLS, and RLS.

Third, the study establishes the conditional nature of the implemented adaptive forecasting advantage. Adaptive NLMS achieves a 12.94% average RMSE improvement over restricted OLS and statistically significant superiority over the tested RLS configuration. However, correctly specified OLS achieves the lowest forecasting error, while rolling OLS performs comparably to adaptive NLMS. These findings demonstrate that sequential adaptation can improve performance under restricted model specification without establishing universal superiority over appropriately specified or alternative adaptive estimators.

The practical contribution lies in identifying methodological considerations relevant to organizations evaluating adaptive forecasting systems for financial and sustainability-related planning, including potential logistics and supply chain applications. These considerations concern forecasting design and model selection rather than demonstrated improvements in actual operational performance.

Within Dynamic Financial Applied Meta-Science (DFAS), the study preserves the distinction between conceptual innovation, computational verification, and governance accountability (Alaali, 2025i, 2025j). Its contribution is therefore a developed A-BSF-AI architecture supported by comparative computational evidence, with empirical organizational validation reserved for subsequent research.

LITERATURE REVIEW

Behavioral Sustainability and Managerial Control

Behavioral sustainability concerns the capacity of organizational practices, managerial decisions, and operational behaviors to sustain performance under changing economic and institutional conditions. It encompasses financial discipline, adaptability, and organizational accountability.

Managerial control theory explains how performance measurement, incentives, and strategic boundaries influence organizational behavior (Simons, 1995; Merchant & Van der Stede, 2017). Cybernetics and system dynamics further emphasize feedback, regulation, and temporal interactions in organizational performance (Beer, 1972; Forrester, 1961).

Within the Alaali Financial Models Meta-Framework (AFMF), the Alaali Self-Funded Behavioral Model (A-SFBM; AFMF-SFBM-F01) connects behavioral discipline with self-funded financial efficiency (Alaali, 2025a). Subsequent developments introduce behavioral efficiency and self-funded ROI measures (AFMF-SFBM-A02), the Expected Behavioral Efficiency Index (A-EBEI; AFMF-SFBM-A03), the Behavioral Cost Recovery Ratio (A-BCRR; AFMF-SFBM-A04), and the Self-Regulation Multiplier (A-SRM; AFMF-SFBM-A05) (Alaali, 2025b, 2025c, 2025d, 2025e).

The Alaali Behavioral Sustainability Index (A-BSI; AFMF-SFBM-A07) and Alaali Behavioral Sustainability Forecasting Model (A-BSF; AFMF-SFBM-A08) extend this sequence toward sustainability assessment and forecasting (Alaali, 2025g, 2025h).
A-BSF-AI introduces an adaptive forecasting direction while distinguishing predictive accuracy from actual organizational sustainability. Improved forecasting does not independently establish changes in managerial behavior or operational performance.

Adaptive Forecasting and Learning Systems

Adaptive forecasting updates predictive parameters as new observations become available. Unlike fixed-coefficient models, adaptive estimators can respond to changing relationships, although their effectiveness depends on model specification, learning parameters, and estimation stability.

Adaptive control theory establishes the principles of feedback-based parameter adjustment under uncertainty (Åström & Wittenmark, 1995). System dynamics and cybernetics provide complementary perspectives on temporal interactions and regulatory feedback (Forrester, 1961; Beer, 1972).

Several estimation approaches are relevant to comparative evaluation. Fixed ordinary least squares (OLS) retains coefficients estimated from historical data, whereas rolling OLS periodically re-estimates relationships using recent observations. Recursive least squares (RLS) updates coefficients sequentially, while normalized least mean squares (NLMS) adjusts parameters through normalized prediction-error feedback. Somefun et al. (2024) examine adaptive filtering algorithms, demonstrating that NLMS offers improved convergence and robustness to variations in input signal power, whereas RLS provides rapid convergence through recursive estimation at greater computational complexity. These distinctions support the comparative evaluation of adaptive NLMS and RLS within the proposed A-BSF-AI forecasting mechanism.

These methods differ in their updating mechanisms and sensitivity to changing conditions. Importantly, adaptive estimation does not necessarily outperform a correctly specified fixed model. Reinforcement learning represents a distinct computational approach in which an agent learns to optimize decision policies through interactions involving states, actions, and rewards (Sutton & Barto, 2018). In contrast, sequential gradient-based forecasting updates predictive parameters using observed prediction errors without requiring policy optimization or reward-based decision-making. Accordingly, the NLMS mechanism implemented within A-BSF-AI constitutes online adaptive estimation rather than reinforcement learning.

Within the Alaali Financial Models Meta-Framework (AFMF), the Alaali Expected Behavioral Efficiency Index with Adaptive Intelligence (A-EBEI-AI; AFMF-SFBM-A06) introduces adaptive intelligence into behavioral-financial assessment (Alaali, 2025f). Building on this conceptual foundation, A-BSF-AI extends the AFMF behavioral sustainability sequence through a proposed forecasting architecture incorporating sequential coefficient updating and prediction-error feedback. This development connects the AFMF model series with established online adaptive estimation methods while distinguishing the proposed architecture from the computational mechanisms evaluated in the present study.

Recent research on adaptive learning emphasizes the importance of responding to evolving data distributions and concept drift. Wang et al. (2024) investigated time-segmentation-based data stream learning for concept drift adaptation, providing a contemporary methodological perspective on learning under changing conditions. Consistent with this research direction, the present study evaluates the implemented A-BSF-AI adaptive NLMS mechanism against restricted OLS, correctly specified OLS, rolling OLS, and RLS across 30 independent synthetic simulation runs. The comparison examines forecasting accuracy under time-varying behavioral, ESG, and market

conditions. Advanced reinforcement learning, autonomous decision-making, and explicit concept drift detection remain outside the computational assessment.

ESG Uncertainty and Sustainability Forecasting

Sustainability forecasting considers environmental, social, and governance (ESG) conditions alongside behavioral and financial influences. Elkington (1997) establishes the triple-bottom-line perspective, while Gray and Bebbington (2001) emphasize the importance of environmental accounting and organizational accountability. Extending these foundations to operational contexts, Walter et al. (2025) examine the application of artificial intelligence in supply chain demand planning, highlighting the relevance of data-driven forecasting to operational decision-making and resource planning. These perspectives establish the importance of integrating sustainability considerations with forecasting processes in changing organizational environments. However, the reliability of historical forecasting relationships may deteriorate as ESG conditions, market dynamics, and behavioral patterns evolve. Arora et al. (2024) systematically review concept drift detection and adaptation techniques, emphasizing the challenges associated with changing data distributions and the methodological importance of adaptive learning. Their findings provide a foundation for investigating sequential parameter updating when underlying forecasting relationships are unstable. Nevertheless, adaptive estimation does not necessarily guarantee superior predictive performance, particularly when conventional forecasting models adequately capture relevant explanatory relationships.

Accordingly, sustainability-oriented forecasting requires distinguishing between changes in underlying relationships and deficiencies in model specification. While adaptive methods may accommodate evolving patterns through sequential coefficient adjustments, correctly specified fixed models can capture relevant temporal interactions directly. This distinction is particularly important when evaluating forecasting mechanisms intended for behavioral-financial sustainability and potential supply chain applications. It provides the methodological rationale for comparing A-BSF-AI adaptive NLMS against restricted OLS, correctly specified OLS, rolling OLS, and RLS under controlled synthetic conditions.

The simulation incorporates synthetic ESG, behavioral, market, cyclical, and time-dependent variables. A time–ESG interaction permits comparison between restricted forecasting specifications and a correctly specified benchmark.

Because these variables are artificially generated, the experiment evaluates forecasting behavior under controlled assumptions rather than observed organizational sustainability outcomes.

Ethical Governance of Adaptive Financial Systems

Adaptive financial systems raise governance concerns because changing model parameters can alter predictive outputs and their interpretation. Managerial control theory emphasizes accountability, performance evaluation, and strategic boundaries in organizational decision-making (Simons, 1995; Merchant & Van der Stede, 2017).

Within Dynamic Financial Applied Meta-Science (DFAS), the DFAS Future Engine Prototype Protocol (DFAS-FEP) addresses responsible AI use, transparency, authorship integrity, and accountability (Alaali, 2025i, 2025j).

A-BSF-AI proposes three complementary indicators: the Behavioral Intelligence Coefficient (BIC), Ethical Responsibility Index (ERI), and Intelligent Maturity Index (IMI). These address learning-related performance, ethical responsibility, and intelligent-system maturity, respectively. Their operational application requires validated measurement definitions, auditable decision records, and appropriate mathematical safeguards. In particular, BIC requires treatment of numerical instability when changes in behavioral entropy approach zero.

The present experiment does not calculate or validate these indicators. Prediction-error clipping and residual monitoring represent computational controls rather than evidence of ethical compliance or autonomous governance.

Literature Synthesis and Research Positioning

The literature establishes four foundations for A-BSF-AI: behavioral-financial regulation, adaptive estimation, sustainability forecasting, and ethical governance. Managerial control, cybernetics, and system dynamics explain organizational feedback and accountability (Beer, 1972; Forrester, 1961; Simons, 1995; Merchant & Van der Stede, 2017). Adaptive control provides the methodological basis for sequential parameter adjustment (Åström & Wittenmark, 1995), while reinforcement learning represents a separate approach to policy optimization (Sutton & Barto, 2018). Sustainability accounting supports the inclusion of environmental and social considerations (Elkington, 1997; Gray & Bebbington, 2001).

Within AFMF, the developmental sequence progresses from behavioral efficiency and self-regulation toward adaptive intelligence, sustainability assessment, and forecasting (Alaali, 2025a, 2025b, 2025c, 2025d, 2025e, 2025f, 2025g, 2025h). A-BSF-AI extends this sequence through a proposed adaptive forecasting architecture informed by DFAS governance principles (Alaali, 2025i, 2025j).

The study distinguishes conceptual specification, computational implementation, and empirical validation. Its computational contribution consists of 30 independent synthetic simulations comparing adaptive NLMS with four forecasting alternatives. The comparative design evaluates whether sequential adaptation improves accuracy beyond restricted fixed estimation and whether its performance remains competitive with correctly specified and alternative adaptive models. The results support a conditional forecasting advantage without establishing universal superiority or validating the broader conceptual architecture. The study therefore contributes to AFMF through the development of A-BSF-AI and a comparative computational assessment, while reserving claims concerning organizational sustainability and ethical governance for subsequent empirical research.

RESEARCH METHODOLOGY

Research Design

This study adopts a model-development and comparative computational simulation design to examine the adaptive forecasting component of the Alaali Behavioral Sustainability Forecasting Artificial Intelligence Model (A-BSF-AI; AFMF-SFBM-A09). It extends the sustainability forecasting direction established by A-BSF (AFMF-SFBM-A08) and the adaptive intelligence principles introduced through A-EBEI-AI (AFMF-SFBM-A06) (Alaali, 2025f, 2025h).

The research combines conceptual mathematical specification with controlled synthetic forecasting experiments. The conceptual component develops the A-BSF-AI architecture, incorporating adaptive learning, behavioral volatility constraints, and three proposed indicators: the Behavioral Intelligence Coefficient (BIC), Ethical Responsibility Index (ERI), and Intelligent Maturity Index (IMI). The computational component evaluates sequential coefficient updating against fixed and alternative adaptive forecasting methods, drawing on adaptive control theory and dynamic feedback principles (Åström & Wittenmark, 1995; Forrester, 1961).

The experiment comprises 30 independently seeded simulation runs, each containing 15,000 synthetic observations distributed across 10 consecutive periods, with 1,500 observations per period. This produces 450,000 observations across the experiment. The data-generating process incorporates behavioral, environmental, social, and governance (ESG), market, cyclical, and temporal influences, including a time–ESG interaction.

Within each run, observations are partitioned chronologically into nonoverlapping training (70%; n = 10,500), validation (15%; n = 2,250), and testing (15%; n = 2,250) subsets. Training observations support model estimation and initialization, validation observations support applicable hyperparameter selection, and testing observations provide the final comparative assessment.

Five forecasting specifications are evaluated: restricted ordinary least squares (OLS), correctly specified OLS incorporating the time–ESG interaction, rolling OLS, recursive least squares (RLS), and adaptive normalized least mean squares (NLMS). Identical observations are used across models within each run to support paired comparisons.

Root mean squared error (RMSE) serves as the principal forecasting accuracy criterion. Mean RMSE, standard deviations, confidence intervals, and paired statistical tests across the 30 runs assess the consistency of performance differences.

The study addresses conceptual model development and computational implementation rather than external empirical validation. Because all observations are synthetic, the findings remain conditional on the specified data-generating assumptions and do not establish demonstrated improvements in organizational sustainability.

A-BSF-AI Model Architecture and Core Equation

A-BSF-AI extends A-BSF by introducing learning-dependent forecasting adjustments, performance feedback, and behavioral volatility constraints (Alaali, 2025h). Its conceptual equation is:

$$\text{A-BSF-AI}(t+1) = \text{A-BSF}(t) + \eta \times [\text{Reward}(t) - \text{Penalty}(t)] - \lambda \times \sigma_{\text{behavioral}}$$

Where:

- A-BSF-AI(t+1): Updated behavioral sustainability forecasting state.
- A-BSF(t): Previous behavioral sustainability forecasting state.
- η: Learning-rate coefficient.
- Reward(t): Positive sustainability-related learning signal.
- Penalty(t): Negative adjustment associated with behavioral inefficiency, forecasting deterioration, or ethical constraints.
- λ: Behavioral volatility penalty coefficient.
- σ_behavioral: Measure of behavioral volatility.

The equation connects forecasting adjustments with net learning feedback and behavioral instability, consistent with adaptive control principles (Åström & Wittenmark, 1995). However, it remains a conceptual specification because the reward, penalty, and volatility components require operational definitions and independent validation.

Implemented Adaptive Forecasting Mechanism

The computational experiment implements sequential coefficient updating rather than directly calculating the conceptual reward–penalty equation.
The linear forecasting specification is:

$$\hat{y}(t) = w(t)' \times x(t)$$

Where $\hat{y}(t)$ represents the predicted outcome, w(t) is the current coefficient vector, and x(t) contains the intercept and selected synthetic predictors.

The adaptive update equation is:

$$w(t+1) = w(t) + \alpha \times clip[y(t) - \hat{y}(t), -0.6, 0.6] \times x(t) / [1 + x(t)'x(t)]$$

Where:

- $w(t+1)$: Updated coefficient vector.
- α: Learning-rate parameter controlling adjustment magnitude.
- $y(t) - \hat{y}(t)$: Prediction error calculated after observing the outcome.
- $clip[\cdot, -0.6, 0.6]$: Restriction of the error contribution to the interval $[-0.6, +0.6]$.
- $x(t)$: Current predictor vector.
- $1 + x(t)'x(t)$: Normalization term moderating update magnitude.

The procedure generates each prediction before observing its corresponding outcome, calculates the prediction error, and updates the coefficients for the subsequent observation. Error clipping limits unusually large adjustments, while normalization reduces sensitivity to predictor magnitude. The learning rate is selected independently within each simulation run using validation RMSE from the candidate values 0.0005, 0.001, 0.002, 0.004, 0.008, 0.012, 0.02, 0.04, and 0.08. The selected value is 0.02 in all 30 runs. The implemented procedure is evaluated against restricted OLS, correctly specified OLS, rolling OLS, and RLS. These comparisons distinguish the potential benefits of sequential adaptation from those associated with predictor specification and alternative updating mechanisms. Sequential coefficient updating can accommodate changing relationships but does not guarantee convergence, optimality, or stability under all data-generating conditions.

Relationship Between the Conceptual and Computational Models

The conceptual and computational equations represent distinct levels of A-BSF-AI development. The conceptual formulation connects sustainability forecasting with rewards, penalties, and behavioral volatility, whereas the implemented procedure updates linear forecasting coefficients using observed prediction errors. Although η and α both regulate adjustment magnitude, they are not mathematically identical. Likewise, prediction-error feedback does not constitute a complete implementation of the conceptual reward–penalty mechanism.

The repeated simulations evaluate the forecasting component against fixed and alternative adaptive estimators. They do not implement the conceptual behavioral volatility penalty, autonomous ethical constraints, deep reinforcement learning, federated learning, or multi-agent coordination. Accordingly, the computational findings concern the predictive properties of the implemented adaptive forecasting procedure rather than validation of the complete proposed A-BSF-AI architecture.

Behavioral Intelligence Coefficient (BIC), Ethical Responsibility Index (ERI), and Intelligent Maturity Index (IMI)

The conceptual architecture of A-BSF-AI introduces three complementary indicators: the Behavioral Intelligence Coefficient (BIC), Ethical Responsibility Index (ERI), and Intelligent Maturity Index (IMI). These constructs extend the AFMF behavioral sustainability sequence by addressing learning performance, ethical accountability, and intelligent-system maturity within the broader governance principles of Dynamic Financial Applied Meta-Science (DFAS) (Alaali, 2025f, 2025g, 2025h, 2025i, 2025j). The indicators remain proposed mathematical constructs. Although the computational experiment evaluates forecasting performance across 30 independent simulation runs, it does not operationalize or validate BIC, ERI, or IMI.

Behavioral Intelligence Coefficient (BIC)

BIC is proposed to assess the relationship between forecasting improvement, behavioral durability, and changes in behavioral uncertainty.

$$BIC = \Delta(Accuracy \times Durability) / \Delta Entropy$$

Where:

- Accuracy: Predictive performance of the forecasting system.
- Durability: Capacity to maintain behavioral-financial performance over time.
- Entropy: Measure of behavioral uncertainty or disorder.
- Δ: Change between defined evaluation periods.

The numerator represents changes in combined accuracy and durability, while the denominator captures changes in behavioral entropy. Operational implementation requires standardized definitions, reproducible entropy measurement, and explicit treatment of positive and negative entropy changes. The ratio is undefined when ΔEntropy equals zero and may become numerically unstable as the denominator approaches zero.

The simulation evaluates forecasting accuracy and its variation across independent runs but does not measure behavioral durability or entropy. Accordingly, no operational BIC score is calculated.

Ethical Responsibility Index (ERI)

ERI is proposed to assess ethical responsibility by relating the frequency of ethical violations to their consequences.

$$ERI = 1 - (Ethical\ Violations / Total\ Decisions) \times Decision\ Impact$$

Where:

- Ethical Violations: Number of decisions violating defined ethical requirements.
- Total Decisions: Number of decisions evaluated.
- Decision Impact: Assessed severity or consequence of ethical violations.

The formulation combines violation frequency with an impact adjustment, reflecting the accountability principles associated with DFAS-FEP (Alaali, 2025i).

Operational implementation requires explicit ethical criteria, auditable decision records, reliable violation classifications, and a justified impact scale. Without normalization, the impact adjustment may produce values outside the zero-to-one range. The ratio is also undefined when no decisions are recorded.

The simulation does not generate ethical decision records or evaluate compliance. Prediction-error clipping and residual monitoring are numerical controls rather than validated ethical safeguards. Consequently, no operational ERI score is calculated.

Intelligent Maturity Index (IMI)

IMI is proposed to characterize intelligent-system maturity through autonomy, learning efficiency, ethical alignment, and operational complexity.

$$IMI = (Autonomy\ Level \times Learning\ Efficiency \times Ethical\ Alignment) / System\ Complexity$$

Where:

- Autonomy Level: Degree of independent execution of defined system functions.
- Learning Efficiency: Effectiveness of learning relative to specified performance criteria.
- Ethical Alignment: Conformity with established ethical requirements.
- System Complexity: Operational or computational complexity of the system.

The formulation combines learning, autonomy, and ethical alignment while adjusting for complexity, consistent with DFAS governance principles (Alaali, 2025i, 2025j). Empirical application requires validated measurement scales, task-specific autonomy definitions, independently assessed ethical alignment, and a strictly positive complexity measure. Universal maturity thresholds cannot be established without calibration.

Although the implemented procedure updates forecasting coefficients sequentially, the experiment does not independently assess autonomy, ethical alignment, or system complexity as maturity indicators. Therefore, no operational IMI score is calculated.

Methodological Status of the Three Indicators

BIC, ERI, and IMI represent distinct conceptual dimensions of A-BSF-AI: learning-related performance and uncertainty, ethical responsibility, and intelligent-system maturity. Their mathematical specification does not establish measurement validity, statistical independence, or causal relationships. Each requires separate operational definitions, calibration, and empirical testing. The 30-run computational assessment strengthens the evaluation of forecasting accuracy through repeated simulations and comparisons with four alternative estimators. However, statistical evidence of predictive improvement does not establish behavioral durability, ethical responsibility, or intelligent-system maturity. Accordingly, the indicators remain conceptual contributions, while computational conclusions are restricted to the implemented forecasting mechanism.

Synthetic Data Generation and Simulation Parameters

The computational experiment comprises 30 independently seeded simulation runs, each containing 15,000 synthetic observations across 10 consecutive periods, with 1,500 observations per period. The combined experiment generates 450,000 observations.

The outcome-generating equation is:

$$y(t) = 0.32 \times ESG(t) + 0.25 \times Behavior(t) - 0.18 \times Market(t) + 0.16 \times \sin(2\pi t/3000) + 0.12 \times Time(t) \times ESG(t) + \varepsilon(t)$$

Where:

- y(t): Simulated sustainability-related outcome.
- ESG(t): Synthetic ESG condition.
- Behavior(t): Synthetic behavioral condition.
- Market(t): Synthetic market shock.
- sin(2πt/3000): Cyclical component.
- Time(t): Normalized observation index ranging from 0 to 1.
- ε(t): Gaussian disturbance with standard deviation 0.12.

ESG, behavioral, and market variables are independently generated from standard normal distributions. The normalized time variable is calculated as t/(N − 1), where t ranges from 0 to N −

1 and N = 15,000.

The coefficients represent experimental assumptions rather than empirical estimates. The time–ESG interaction introduces a changing relationship between ESG conditions and the simulated outcome.

Table 1. Synthetic Simulation Design and Parameters

Parameter	Specification
Observations per run	15,000
Independent simulation runs	30
Total generated observations	450,000
Simulated periods per run	10
Observations per period	1,500
Random seeds	20261105–20261134
ESG coefficient	0.32
Behavioral coefficient	0.25
Market coefficient	−0.18
Cyclical coefficient	0.16
Time–ESG interaction coefficient	0.12
Gaussian noise standard deviation	0.12
Training observations	10,500 (70%)
Validation observations	2,250 (15%)
Testing observations	2,250 (15%)

Two predictor specifications are evaluated. The restricted specification includes an intercept, ESG, behavior, market conditions, and the cyclical component. The correctly specified OLS benchmark additionally incorporates the time–ESG interaction. This distinction permits comparison between adaptive forecasting, restricted fixed estimation, and a benchmark containing the known interaction. Rolling OLS and RLS provide additional comparisons with conventional adaptive estimation procedures.

All five forecasting methods use identical observations within each simulation run, supporting paired performance comparisons. The resulting evidence remains conditional on the synthetic data-generating process and does not establish external empirical validity.

Baseline and Adaptive Forecasting Procedures

Five forecasting specifications are evaluated using identical generated observations within each simulation run: restricted fixed ordinary least squares (OLS), correctly specified OLS, rolling OLS, recursive least squares (RLS), and A-BSF-AI adaptive clipped normalized least mean squares (NLMS).

Fixed OLS Baselines

The restricted OLS model is estimated using 10,500 training observations:

$$\hat{y}_OLS(t) = w_OLS' \times x(t)$$

Where w_OLS represents the training-estimated coefficient vector and x(t) contains the restricted predictor set. The correctly specified OLS benchmark additionally incorporates the Time(t) × ESG(t) interaction included in the synthetic data-generating process.

Both OLS specifications retain fixed coefficients during validation and testing. Their comparison distinguishes the effects of predictor specification from those of sequential coefficient adaptation.

Adaptive and Recursive Forecasting Procedures

The A-BSF-AI computational implementation applies the normalized gradient update defined in Section 3.2.1. Its coefficients are initialized using restricted training-based OLS estimates and updated sequentially after each observed outcome.

Three adaptive estimation approaches are evaluated:

- A-BSF-AI adaptive NLMS: Sequential normalized gradient updating with prediction-error clipping at ±0.6.
- Rolling OLS: Re-estimation using the latest 1,500 available observations, with coefficients refreshed every 250 observations.
- Recursive least squares (RLS): Sequential coefficient updating with a forgetting factor of 0.995.

Rolling OLS uses only observations available before each refitting point. RLS begins testing from the adaptive NLMS model's post-validation coefficient estimates and applies recursive updates. These procedures represent alternative mechanisms for parameter adaptation. Their comparison does not imply equivalent initialization or hyperparameter optimization, and none constitutes reinforcement learning.

Learning-Rate Selection

Nine candidate learning rates are evaluated for adaptive NLMS:

$$\alpha = 0.0005, 0.001, 0.002, 0.004, 0.008, 0.012, 0.02, 0.04, \text{and } 0.08$$

The learning rate minimizing validation RMSE is selected independently for each simulation run. The selected value is 0.02 in all 30 runs. The corresponding post-validation adaptive coefficients are carried forward into testing. Each test prediction is generated before its outcome is observed, and subsequent coefficient updates use only information available at that point. Testing observations are excluded from learning-rate selection.

Table 2. Forecasting Procedures and Evaluation Specifications

Model	Estimation or updating
Restricted OLS	Fixed training-estimated coefficients; excludes time–ESG interaction
Correctly specified OLS	Fixed training-estimated coefficients; includes time–ESG interaction
Rolling OLS	1,500-observation window; refitting every 250 observations
Recursive least squares	Sequential updating; forgetting factor 0.995
A-BSF-AI adaptive NLMS	Sequential normalized updates; validation-selected learning rate; ±0.6 error clipping

All five models are evaluated using the same 2,250 testing observations per run, producing 67,500 testing observations across the 30 independent simulations.

Computational Prediction Bound

An absolute prediction-residual threshold of 2.5 is used to monitor unusually large adaptive forecasting errors. Residual exceedances are recorded separately from the ±0.6

prediction-error clipping mechanism applied during coefficient updating.
These numerical controls do not constitute ethical compliance measures, operational governance safeguards, or validation of the Ethical Responsibility Index (ERI).

Model Evaluation and Validation Metrics

Forecasting accuracy is evaluated using root mean squared error (RMSE):

$$RMSE = \sqrt{[\Sigma(y(i) - \hat{y}(i))^2 / n]}$$

Where $y(i)$ represents the simulated outcome, $\hat{y}(i)$ the corresponding forecast, and n the number of evaluated observations.

Relative forecasting improvement is calculated as:

$$Improvement\ (\%) = [(RMSE_comparator - RMSE_adaptive) / RMSE_comparator] \times 100$$

Positive values indicate lower forecasting error for adaptive NLMS, while negative values indicate inferior performance relative to the comparator. Across the 30 independent runs, forecasting performance is summarized using mean RMSE, standard deviation, standard error, and 95% confidence intervals. Two-sided paired t-tests assess run-level RMSE differences between adaptive NLMS and each comparator. The number of runs in which adaptive NLMS achieves lower RMSE is also recorded. These comparisons evaluate performance consistency under the specified synthetic conditions rather than universal forecasting superiority.

Chronological Evaluation Structure

Each simulation run is partitioned chronologically into training (10,500 observations), validation (2,250 observations), and testing (2,250 observations). Training establishes initial model coefficients, validation supports adaptive learning-rate selection, and testing provides the final comparative evaluation.

Adaptive forecasts are generated before observing their corresponding outcomes. After each observation becomes available, the relevant coefficients are updated for subsequent predictions.

The resulting procedure constitutes one-step-ahead online forecasting rather than a fixed-coefficient, multi-period holdout assessment. Chronological partitioning prevents future observations from entering earlier predictions, while permitting sequential learning from previously observed testing outcomes.

Numerical Verification of Forecasting Performance

The computational outputs include run-level RMSE values, learning-rate validation results, coefficient trajectories, descriptive statistics, and observation-level testing predictions for the first simulation run.

Forecasting accuracy is verified using matched actual and predicted observations, independently of spreadsheet summary formulas. Training R^2, prediction-error clipping frequency, and residual-threshold exceedances are also recorded.

The 30 independent simulations permit assessment of stochastic performance variation under the same data-generating structure. However, repeated sampling does not establish robustness across alternative structural assumptions, disturbance distributions, or observed organizational environments.

Reproducibility and Simulation Documentation

The study provides supplementary computational documentation comprising executable simulation code, parameter specifications, generated observations, forecasting outputs, and comparative evaluation results.

The experiment includes 30 independently seeded runs, each containing 15,000 observations. The documented procedures specify the data-generating equation, random seeds, chronological partitions, predictor specifications, estimation methods, learning-rate candidates, and evaluation metrics.

The executable simulation implements random-number generation, predictor construction, fixed OLS estimation, rolling OLS, recursive estimation, adaptive NLMS updating, and RMSE calculations.

Run-level performance results and observation-level predictions support independent numerical verification. The accompanying spreadsheet organizes simulation outputs and comparative statistics, while the executable code provides the computational basis for reproducing the experiment.

The documentation distinguishes implemented forecasting procedures from the conceptual BIC, ERI, and IMI indicators and other proposed A-BSF-AI capabilities.

Computational reproducibility supports verification of the reported synthetic results but does not establish external empirical validity, ethical governance effectiveness, or organizational sustainability improvements.

FINDINGS AND DISCUSSION

Simulation Dataset and Experimental Configuration

The computational experiment comprises 30 independently seeded simulation runs, each containing 15,000 synthetic observations across 10 consecutive periods, producing 450,000 observations in total. Each run allocates 10,500 observations to training (70%), 2,250 to validation (15%), and 2,250 to testing (15%).

Five forecasting models are evaluated: restricted ordinary least squares (OLS), correctly specified OLS, rolling OLS, recursive least squares (RLS), and A-BSF-AI adaptive normalized least mean squares (NLMS). Identical observations are used across models within each run to support paired statistical comparisons.

The synthetic outcome incorporates environmental, social, and governance (ESG) conditions, behavioral factors, market shocks, cyclical variation, and a time–ESG interaction. The correctly specified OLS benchmark includes this interaction, whereas the remaining models employ the restricted predictor specification.

The adaptive learning rate is selected independently for each run using validation RMSE. A value of 0.02 is selected in all 30 runs. Testing follows a one-step-ahead forecasting procedure, with adaptive coefficients updated only after observing the corresponding outcome.

Table 3. OLS Coefficient Estimation and Model Fit Across 30 Simulation Runs

Predictor	True Coefficient	Restricted OLS	Correctly Specified OLS
Intercept	0.000000	0.000006	0.000013
ESG	0.320000	0.361863	0.319745
Behavioral Factor	0.250000	0.249410	0.249416
Market Shock	−0.180000	−0.180167	−0.180241
Cyclical Component	0.160000	0.159791	0.159822
Time–ESG Interaction	0.120000	Excluded	0.120454
Training R^2	—	0.940681	0.943018

Note. Estimated coefficients and R^2 represent averages across 30 independently seeded simulation runs, using 10,500 training observations per run. The restricted OLS specification excludes the time–ESG interaction, whereas the correctly specified OLS specification includes it.

The correctly specified OLS model recovers the underlying coefficients more accurately, particularly for ESG (0.319745 versus the true coefficient of 0.320000). Excluding the time–ESG interaction increases the estimated ESG coefficient to 0.361863. The correctly specified model also achieves a higher average training R^2 (0.943018 versus 0.940681), demonstrating the importance of incorporating relevant temporal interactions into the forecasting specification.

Table 4. Learning Rate Sensitivity Analysis Across 30 Simulation Runs

Learning Rate (α)	Mean Validation RMSE
0.0005	0.128279
0.0010	0.127009
0.0020	0.125134
0.0040	0.122999
0.0080	0.121376
0.0120	0.120856
0.0200	0.120603
0.0400	0.120838
0.0800	0.121805

Note. Mean validation RMSE values are calculated across 30 independently seeded simulation runs. The learning rate of 0.02 achieved the lowest mean validation RMSE and was selected in all 30 runs. The learning rate sensitivity analysis demonstrates that validation RMSE decreases as the learning rate increases from 0.0005 to 0.02, reaching its minimum of 0.120603. Further increases to 0.04 and 0.08 produce higher forecasting errors. The consistent selection of 0.02 across all 30 runs indicates that this learning rate provides the best validation performance among the nine tested candidates under the specified synthetic conditions.

Comparative Forecasting Performance

Table 5 presents testing RMSE across the 30 independent simulation runs.

Table 5. Forecasting Performance Across 30 Simulation Runs

Forecasting model	Mean RMSE	Standard deviation	95% confidence interval
Restricted OLS	0.138282	0.002005	0.137534–0.139031
Correctly specified OLS	0.119822	0.001610	0.119221–0.120424
Rolling OLS	0.120312	0.001643	0.119699–0.120926
RLS	0.121456	0.001543	0.120880–0.122032
A-BSF-AI adaptive NLMS	0.120378	0.001565	0.119794–0.120963

A-BSF-AI adaptive NLMS reduces mean testing RMSE from 0.138282 under restricted OLS to 0.120378, corresponding to an average relative improvement of 12.94%. The adaptive procedure outperforms restricted OLS in all 30 simulation runs.

However, correctly specified OLS achieves the lowest mean RMSE (0.119822), followed by rolling OLS (0.120312). Adaptive NLMS achieves lower forecasting error than RLS (0.121456) but slightly higher error than correctly specified and rolling OLS.

These results indicate that sequential coefficient updating improves forecasting accuracy relative to a restricted fixed specification, although incorporating the relevant time–ESG interaction directly into a fixed model produces superior performance under the simulated conditions.
Paired statistical comparisons assess the consistency and significance of these differences across independent runs.

Table 6. Paired Statistical Comparisons Against A-BSF-AI Adaptive NLMS

Comparator	Mean RMSE difference	Mean improvement (%)	p-value	Adaptive wins
Restricted OLS	+0.017904	+12.94%	<0.001	30/30
Correctly specified OLS	−0.000556	−0.46%	<0.001	0/30
Rolling OLS	−0.000066	−0.06%	0.118	13/30
RLS	+0.001077	+0.89%	<0.001	30/30

Note: Mean RMSE difference is calculated as comparator RMSE minus adaptive NLMS RMSE. Positive values indicate an adaptive forecasting advantage. Mean improvement percentages represent the average of run-level relative improvements.

The paired comparisons demonstrate statistically significant improvements over restricted OLS and the tested RLS specification ($p < 0.001$). Adaptive NLMS achieves lower RMSE than both models in all 30 runs.

Conversely, correctly specified OLS significantly outperforms adaptive NLMS ($p < 0.001$). The difference between rolling OLS and adaptive NLMS is not statistically significant ($p = 0.118$), indicating comparable performance under the specified experimental conditions.

The RLS comparison requires qualification because RLS is initialized using the adaptive NLMS model's post-validation coefficients. Consequently, its results represent the tested recursive configuration rather than an independently optimized RLS benchmark.

Overall, the evidence supports a conditional forecasting advantage for A-BSF-AI adaptive NLMS, particularly when the fixed forecasting specification omits a relevant time-dependent interaction.

Adaptive Forecasting Behaviour and Performance Interpretation

The adaptive NLMS procedure updates its coefficient vector sequentially using prediction-error feedback. Each prediction is generated before observing the corresponding outcome, and the subsequent error informs the next coefficient update.

The simulation incorporates a time–ESG interaction, introducing gradual variation in the relationship between ESG conditions and the simulated sustainability-related outcome. Because adaptive NLMS uses the restricted predictor set, sequential updating provides a mechanism for partially accommodating this changing relationship without explicitly incorporating the interaction term.

The improvement over restricted OLS is consistent across all 30 runs, suggesting that coefficient adaptation reduces forecasting error associated with the restricted fixed specification. Nevertheless, the superior performance of correctly specified OLS demonstrates that direct inclusion of the relevant interaction can be more effective than sequential adjustment of an incomplete predictor specification.

Rolling OLS achieves forecasting accuracy comparable to adaptive NLMS, indicating that periodic re-estimation using recent observations can provide similar predictive benefits. The absence of a statistically significant difference between these methods limits claims concerning the

incremental advantage of the implemented NLMS mechanism.

Adaptive NLMS also outperforms the tested RLS configuration. However, this result remains conditional on the specified forgetting factor, initialization procedure, and synthetic data-generating process.

Table 7. Period-wise Forecasting Performance Across 30 Simulation Runs

Forecasting Model	Period 9 RMSE	Period 10 RMSE
Restricted OLS	0.135085	0.139831
Correctly Specified OLS	0.119990	0.119719
Rolling OLS	0.120383	0.120260
RLS	0.123240	0.120532
A-BSF-AI Adaptive NLMS	0.120537	0.120279

Note. Period-wise RMSE values represent averages across 30 independently seeded simulation runs. The testing partition covers the final 2,250 observations, comprising 750 observations from Period 9 and 1,500 observations from Period 10.

The period-wise results indicate that restricted OLS forecasting error increases from 0.135085 in Period 9 to 0.139831 in Period 10. Correctly specified OLS maintains the lowest RMSE in both periods, while adaptive NLMS and rolling OLS exhibit comparable forecasting accuracy. RLS shows a reduction in forecasting error between the two periods. These findings reinforce the importance of predictor specification and suggest that adaptive estimation can partially accommodate changing relationships without consistently outperforming correctly specified fixed estimation.

The computational monitoring procedures include prediction-error clipping at ±0.6 and an absolute residual threshold of 2.5. Across the testing observations, no clipping events or residual-threshold exceedances are recorded. These findings describe the numerical behaviour of the implemented algorithm under the simulated conditions and do not establish general stability guarantees.

Table 8. Residual Diagnostics and Forecasting Stability Across 30 Simulation Runs

Diagnostic Indicator	A-BSF-AI Adaptive NLMS
Independent simulation runs	30
Total testing observations	67,500
Mean testing RMSE	0.120378
Standard deviation of run-level RMSE	0.001565
95% confidence interval for mean RMSE	0.119794–0.120963
Prediction-error clipping threshold	±0.6
Testing clipping events	0
Clipping event rate	0.00%
Absolute residual monitoring threshold	2.5
Residual-threshold exceedances	0

Note. Diagnostic results are based on 30 independently seeded simulation runs, each containing 2,250 testing observations. Clipping events and residual-threshold exceedances refer exclusively to the testing partition.

Overall, the results demonstrate that adaptive parameter updating can improve forecasting accuracy when relationships vary over time. However, its effectiveness depends on predictor

specification, estimation procedures, and experimental conditions. The evidence supports a conditional computational advantage for A-BSF-AI adaptive NLMS rather than universal forecasting superiority.

Theoretical and Managerial Implications

The findings contribute to the Alaali Financial Models Meta-Framework (AFMF) by providing a comparative computational assessment of the adaptive forecasting component proposed within A-BSF-AI (AFMF-SFBM-A09).

The principal theoretical implication concerns the distinction between adaptive learning and predictor specification. Sequential coefficient updating improves forecasting performance relative to restricted fixed estimation, but correctly specified OLS achieves superior accuracy. This demonstrates that adaptation and model specification represent complementary rather than interchangeable forecasting considerations.

The comparable performance of rolling OLS and adaptive NLMS further indicates that forecasting improvements may arise through different parameter-updating mechanisms. Accordingly, the computational evidence supports the relevance of adaptive estimation without establishing the unique superiority of the A-BSF-AI implementation.

From a managerial perspective, the results suggest that organizations evaluating adaptive forecasting systems should compare them against appropriately specified fixed models and established updating procedures. Forecasting accuracy, model transparency, parameter stability, and computational requirements should be considered before adopting an adaptive approach.

However, these implications remain methodological because the experiment uses synthetic observations rather than actual organizational data. The results do not demonstrate improvements in financial sustainability, managerial behaviour, ESG performance, or operational decision-making.

The proposed Behavioral Intelligence Coefficient (BIC), Ethical Responsibility Index (ERI), and Intelligent Maturity Index (IMI) extend the conceptual scope of A-BSF-AI toward learning effectiveness, ethical accountability, and intelligent-system maturity. Nevertheless, these indicators are not operationalized or validated by the present simulations.

Within Dynamic Financial Applied Meta-Science (DFAS), the study maintains a distinction between conceptual model development, computational evidence, and governance accountability. Predictive accuracy alone does not establish ethical responsibility, autonomous intelligence, or institutional sustainability.

The findings therefore support A-BSF-AI as a developing adaptive forecasting model with demonstrated conditional computational performance. Further empirical testing, alternative structural simulations, and independent validation of its proposed indicators are required before broader organizational or governance claims can be established.

Potential Applications in Logistics and Supply Chain Forecasting

The findings of this study provide a methodological basis for exploring adaptive forecasting applications in logistics and supply chain management, particularly where demand patterns, supplier performance, operational costs, and sustainability-related conditions change over time. Walter et al. (2025) demonstrate the growing relevance of data-driven forecasting in supply chain demand planning, while Arora et al. (2024) highlight the importance of adapting predictive models to evolving data distributions. Within this context, A-BSF-AI offers a candidate adaptive estimation mechanism whose practical relevance requires empirical validation.

Demand Forecasting and Inventory Planning

Changes in customer demand, seasonal fluctuations, and market conditions can weaken forecasting relationships estimated from historical observations. Walter et al. (2025) identify demand planning as an important application of data-driven forecasting in supply chain management. The sequential coefficient-updating mechanism of A-BSF-AI could potentially support periodic demand forecast revisions, inventory replenishment planning, and production scheduling. However, the present simulation does not establish improvements in actual inventory performance or demand forecasting accuracy.

Supplier Performance and Lead-Time Forecasting

Variations in supplier reliability, procurement conditions, and transportation disruptions may alter relationships between operational predictors and delivery outcomes. Concept drift research emphasizes the importance of recognizing and adapting to changing predictive relationships (Arora et al., 2024). Accordingly, A-BSF-AI could be investigated for updating supplier lead-time estimates and procurement forecasts as new operational observations become available. Such applications would require supplier-specific datasets and comparisons with established forecasting alternatives.

Sustainability-Oriented Supply Chain Forecasting

Sustainability considerations introduce additional forecasting challenges involving resource consumption, environmental costs, supplier compliance, and ESG-related operational conditions. The time-varying ESG interaction examined in the present simulation provides a controlled illustration of how changing sustainability-related relationships can influence predictive accuracy. This suggests a potential application of adaptive estimation to sustainability-related operational indicators, although the study does not directly evaluate environmental performance, supplier sustainability, or organizational resilience.

Forecasting Model Selection and Operational Decision-Making

The comparative simulation results provide a qualification for managerial adoption. A-BSF-AI reduced mean RMSE by 12.94% relative to restricted OLS and outperformed the tested RLS configuration. However, correctly specified OLS achieved the lowest forecasting error, while rolling OLS demonstrated statistically comparable performance to A-BSF-AI. These findings indicate that adaptive updating should not be treated as inherently superior to conventional forecasting. Instead, model selection should consider explanatory-variable specification, temporal interactions, updating frequency, and out-of-sample predictive performance. This interpretation is consistent with the broader methodological emphasis on selecting appropriate adaptation strategies under changing data conditions (Arora et al., 2024).

Implementation Requirements and Future Empirical Validation

Practical application would require clearly defined operational forecasting targets, chronological training and testing procedures, representative historical datasets, and benchmark comparisons. Potential evaluation settings include product-level demand forecasting, inventory requirements, supplier delivery times, procurement expenditures, and sustainability-related operating costs. Following the research opportunities identified by Walter et al. (2025), future studies could examine whether adaptive forecasting produces measurable improvements in supply chain decisions rather than predictive accuracy alone. Evaluation should also consider computational costs, forecast stability, and the operational consequences of prediction errors.

Overall, A-BSF-AI provides a computationally evaluated adaptive estimation mechanism with potential relevance to logistics and supply chain forecasting. Its demonstrated advantage over restricted OLS does not establish general superiority over correctly specified or alternative adaptive models. Consequently, its contribution to operational decision-making remains a testable proposition requiring validation with observed supply chain data.

CONCLUSION

Principal Findings

This study developed the conceptual architecture of the Alaali Behavioral Sustainability Forecasting Artificial Intelligence Model (A-BSF-AI; AFMF-SFBM-A09) and evaluated its implemented adaptive forecasting mechanism through repeated synthetic simulations.

The experiment comprised 30 independently seeded runs, each containing 15,000 observations, producing 450,000 observations in total. Five forecasting specifications were compared: restricted ordinary least squares (OLS), correctly specified OLS, rolling OLS, recursive least squares (RLS), and adaptive clipped normalized least mean squares (NLMS).

A-BSF-AI adaptive NLMS achieved a mean testing RMSE of 0.120378, compared with 0.138282 for restricted OLS, representing a 12.94% improvement. The adaptive procedure outperformed restricted OLS and the tested RLS specification in all 30 runs, with statistically significant differences ($p < 0.001$).

However, correctly specified OLS achieved the lowest mean RMSE (0.119822) and significantly outperformed adaptive NLMS ($p < 0.001$). Rolling OLS achieved comparable forecasting performance (RMSE = 0.120312), with no statistically significant difference from adaptive NLMS ($p = 0.118$).

These findings demonstrate that sequential coefficient updating can improve forecasting accuracy relative to a restricted fixed specification under changing synthetic relationships. Nevertheless, incorporating relevant predictors directly into a correctly specified model may produce superior results. The evidence therefore supports a conditional computational advantage rather than universal superiority of adaptive NLMS.

The findings remain specific to the implemented forecasting procedure and the documented synthetic conditions. They do not establish improvements in actual organizational sustainability or validate the complete proposed A-BSF-AI architecture.

Theoretical Contribution to AFMF and DFAS

A-BSF-AI extends the behavioral sustainability forecasting sequence within the Alaali Financial Models Meta-Framework (AFMF) by connecting the forecasting structure of A-BSF with sequential adaptive estimation (Alaali, 2025h).

The proposed architecture incorporates learning-dependent forecasting adjustments, behavioral volatility constraints, and three conceptual indicators: the Behavioral Intelligence Coefficient (BIC), Ethical Responsibility Index (ERI), and Intelligent Maturity Index (IMI). These constructs establish directions for investigating learning effectiveness, ethical accountability, and intelligent-system maturity, although they remain unvalidated.

The computational contribution consists of a reproducible comparative assessment across 30 independent simulations and five forecasting specifications. The results distinguish the benefits of sequential adaptation from those associated with predictor specification and alternative updating methods. They also demonstrate that adaptive estimation does not necessarily outperform correctly specified fixed models or conventional rolling estimation.

Within Dynamic Financial Applied Meta-Science (DFAS), the study reinforces the distinction between conceptual development, computational verification, and governance

accountability (Alaali, 2025i, 2025j). Forecasting accuracy alone does not establish ethical autonomy, institutional legitimacy, or organizational sustainability.

The principal contribution is therefore the development of the A-BSF-AI conceptual architecture, supported by reproducible evidence of the conditional forecasting performance of its implemented adaptive mechanism.

LIMITATIONS AND FUTURE RESEARCH

Research and Simulation Limitations

The principal limitation is the exclusive reliance on synthetic observations. The experiment does not incorporate actual organizational financial records, observed behavioral conditions, or externally measured environmental, social, and governance indicators. Its findings therefore remain conditional on the specified data-generating process.

Although the study employs 30 independently seeded runs, these repetitions use the same underlying structural assumptions and Gaussian disturbance distribution. Consequently, they assess stochastic variability rather than robustness across alternative economic environments, noise distributions, or structural-change scenarios.

The simulation deliberately excludes the time–ESG interaction from the restricted forecasting specification. While correctly specified OLS provides an additional benchmark, the adaptive procedure's advantage over restricted OLS partly reflects this model misspecification. Furthermore, the comparisons do not include nonlinear forecasting models or independently optimized configurations of every alternative estimator.

The computational implementation uses sequential normalized gradient updating rather than Deep Q-Learning, federated reinforcement learning, Bayesian optimization, or Long Short-Term Memory (LSTM) mechanisms. These remain proposed architectural extensions.

Similarly, BIC, ERI, and IMI are not operationally calculated or empirically validated. Prediction-error clipping and residual monitoring are numerical controls, not measures of ethical responsibility.

The availability of executable simulation code and numerical outputs supports computational reproducibility under the documented software environment. However, reproducibility does not establish external validity or demonstrate organizational sustainability improvements.

These limitations restrict the findings to a controlled comparative assessment of adaptive forecasting rather than validation of the complete A-BSF-AI architecture.

Future Empirical Validation and Model Extensions

Future research should evaluate A-BSF-AI using observed financial, behavioral, and sustainability-related data across industries and institutional environments. Such studies should determine whether adaptive forecasting improvements persist under actual operational conditions.

Further simulations should examine alternative data-generating processes, non-Gaussian disturbances, varying market-shock intensities, structural breaks, and alternative learning-rate and estimator configurations. Comparisons with nonlinear and time-varying forecasting methods would strengthen assessment of the incremental value of adaptive coefficient updating.

The proposed BIC, ERI, and IMI indicators require operational definitions, measurement calibration, mathematical safeguards, and independent reliability and validity testing. Subsequent development may also implement reinforcement learning, federated learning, and behavioral memory mechanisms through explicitly specified algorithms and independently evaluated computational experiments.

The longer-term objective is to integrate adaptive forecasting, behavioral sustainability measurement, and accountable governance into an empirically testable financial system. This progression would extend the AFMF research sequence while preserving the distinction between conceptual innovation, computational evidence, and external validation (Alaali, 2025f, 2025h).

REFERENCES

Alaali, H. M. H. (2025a). AFMF-SFBM-F01: The Alaali Self-Funded Behavioral Model (A-SFBM): A new financial–behavioral paradigm integrating incentive economics, operational discipline, and self-regenerative efficiency in family enterprises. In *The Alaali Financial Models Meta-Framework (AFMF): Vol. 1. Behavioral finance, sustainability, and ethical intelligence* [Preprint]. SSRN. https://ssrn.com/abstract=5611650

Alaali, H. M. H. (2025b). AFMF-SFBM-A02: Behavioral efficiency and self-funded ROI: Dual metrics under the A-SFBM framework. In *The Alaali Financial Models Meta-Framework (AFMF): Vol. 1. Behavioral finance, sustainability, and ethical intelligence* [Preprint]. SSRN. https://ssrn.com/abstract=5611650

Alaali, H. M. H. (2025c). AFMF-SFBM-A03: A-EBEI: The Expected Behavioral Efficiency Index—A predictive extension of the A-SFBM framework. In *The Alaali Financial Models Meta-Framework (AFMF): Vol. 1. Behavioral finance, sustainability, and ethical intelligence* [Preprint]. SSRN. https://ssrn.com/abstract=5611650

Alaali, H. M. H. (2025d). AFMF-SFBM-A04: Alaali Behavioral Cost Recovery Ratio (A-BCRR). In *The Alaali Financial Models Meta-Framework (AFMF): Vol. 1. Behavioral finance, sustainability, and ethical intelligence* [Preprint]. SSRN. https://ssrn.com/abstract=5611650

Alaali, H. M. H. (2025e). AFMF-SFBM-A05: Alaali Self-Regulation Multiplier (A-SRM). In *The Alaali Financial Models Meta-Framework (AFMF): Vol. 1. Behavioral finance, sustainability, and ethical intelligence* [Preprint]. SSRN. https://ssrn.com/abstract=5611650

Alaali, H. M. H. (2025f). AFMF-SFBM-A06: A-EBEI-AI—Adaptive intelligence model for behavioral forecasting and regulation. In *The Alaali Financial Models Meta-Framework (AFMF): Vol. 1. Behavioral finance, sustainability, and ethical intelligence* [Preprint]. SSRN. https://ssrn.com/abstract=5611650

Alaali, H. M. H. (2025g). AFMF-SFBM-A07: The Alaali Behavioral Sustainability Index (A-BSI). In *The Alaali Financial Models Meta-Framework (AFMF): Vol. 1. Behavioral finance, sustainability, and ethical intelligence* [Preprint]. SSRN. https://ssrn.com/abstract=5611650

Alaali, H. M. H. (2025h). AFMF-SFBM-A08: Alaali Behavioral Sustainability Forecasting Model (A-BSF). In *The Alaali Financial Models Meta-Framework (AFMF): Vol. 1. Behavioral finance, sustainability, and ethical intelligence* [Preprint]. SSRN. https://ssrn.com/abstract=5611650

Alaali, H. M. H. (2025i). *The DFAS-FEP protocol: A global governance standard for responsible AI use and authorship integrity in financial modelling* [Preprint]. SSRN. https://ssrn.com/abstract=5260517

Alaali, H. M. H. (2025j). *The doctrine of Dynamic Financial Applied Science (DFAS): Foundations, frameworks, and the future of financial epistemology* [Preprint]. SSRN. https://ssrn.com/abstract=5311455

Arora, S., Rani, R., & Saxena, N. (2024). A systematic review on detection and adaptation of concept drift in streaming data using machine learning techniques. *WIREs Data Mining and Knowledge Discovery, 14*(4), Article e1536. https://doi.org/10.1002/widm.1536

Åström, K. J., & Wittenmark, B. (1995). *Adaptive control* (2nd ed.). Addison-Wesley.

Beer, S. (1972). *Brain of the firm*. Allen Lane.

Elkington, J. (1997). *Cannibals with forks: The triple bottom line of 21st century business*. Capstone.

Forrester, J. W. (1961). *Industrial dynamics*. MIT Press.

Gray, R., & Bebbington, J. (2001). *Accounting for the environment*. SAGE Publications.

Merchant, K. A., & Van der Stede, W. A. (2017). *Management control systems: Performance measurement, evaluation and incentives* (4th ed.). Pearson.

Simons, R. (1995). *Levers of control: How managers use innovative control systems to drive strategic renewal*. Harvard Business School Press.

Somefun, C. T., Daramola, S. A., & Somefun, T. E. (2024). Advancements and applications of adaptive filters in signal processing. *Journal Européen des Systèmes Automatisés, 57*(5), 1259–1272. https://doi.org/10.18280/jesa.570502

Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction* (2nd ed.). MIT Press.

Walter, A., Ahsan, K., & Rahman, S. (2025). Application of artificial intelligence in demand planning for supply chains: A systematic literature review. *The International Journal of Logistics Management, 36*(3), 672–719. https://doi.org/10.1108/IJLM-02-2024-0120

Wang, K., Lu, J., Liu, A., & Zhang, G. (2024). TS-DM: A time segmentation-based data stream learning method for concept drift adaptation. *IEEE Transactions on Cybernetics, 54*(10), 6000–6011. https://doi.org/10.1109/TCYB.2024.3429459

Supplementary Materials

- Supplementary File A: ABSF_AI_All_Tests.xlsx
- Supplementary File B: ABSF_AI_simulation.py