Governance-Driven Financial Intelligence: A Multi-Modal Framework for Predictive Throttling, Governance-Aware Execution, and Evidence-Driven Learning

Vidhyalakshmi Amarnath,
amarnathvidhyalaksh@cityuniversity.edu

Taejin Kim,
kimtaejin@cityu.edu

Brian Maeng
maengjooyol@cityu.edu

City University of Seattle
Seattle, Washington, 98121, USA

Abstract

Financial forecasting systems frequently emphasize predictive accuracy while offering limited support for decision governance, execution control, and adaptive risk management. This study introduces Governance-Driven Financial Intelligence (GDFI), a multi-modal decision-support architecture that bridges the prediction-to-decision gap by integrating Temporal Fusion Transformer–based directional-edge forecasting, predictive throttling, governance-aware execution, and evidence-driven adaptive learning. The framework uses 168 engineered features derived from technical indicators, macroeconomic variables, news sentiment, graph-relational signals, and point-in-time fundamental information. Rather than executing all forecast-generated signals, GDFI selectively approves or suppresses forecast opportunities using quality, utility, directional-edge, exposure, and risk-aware governance criteria. The framework was evaluated in a sequential paper-trading environment without live-capital deployment using 480 forecast opportunities and 200 horizon-matured evaluations, comprising 58 approved and 142 suppressed candidates. Governance-approved opportunities achieved a 51.7% hit rate versus 38.0% for counterfactual suppressed candidates, producing a governance lift of 13.7 percentage points. Statistical validation showed stronger mean realized PnL among approved opportunities using a two-sided permutation test ($p = 0.001$; Cohen's $d = 0.747$). Predictive throttling reduced maximum drawdown by approximately 98.9% relative to an unrestricted equal-weight benchmark, although this reduction was accompanied by lower absolute return, indicating a risk–return tradeoff. Finally, the Learning Governor deferred adaptive updates because only 36 realized learning outcomes were available versus the pre-specified minimum requirement of 100. The findings indicate that governance-aware execution and evidence-driven adaptation can improve decision quality, risk control, and accountability, while the Learning Governor provides a transferable evidence-sufficiency mechanism for regulating adaptive updates in AI-enabled decision systems.

Keywords: Financial Intelligence, Decision Governance, Predictive Throttling, Trade Execution, Multi-Modal Analytics, Adaptive Learning.

ISSN: 2473-4901
v12 n6752

Governance-Driven Financial Intelligence: A Multi-Modal Framework for Predictive Throttling, Governance-Aware Execution, and Evidence-Driven Learning

Vidhyalakshmi Amarnath, Taejin Kim and Brian Maeng

1. INTRODUCTION

Financial intelligence systems increasingly use artificial intelligence (AI), machine learning (ML), and multi-modal analytics to support forecasting, investment analysis, and risk management (Goodfellow et al., 2016; Cao, 2022; Kumbure et al., 2022; Vuković et al., 2025). However, prediction accuracy alone does not ensure effective financial decisions. Financial decision systems must determine whether model-generated opportunities should be authorized, resized, delayed, or suppressed under utility, risk, confidence, and operational constraints. Executing every model-generated signal may increase volatility and drawdown risk while reducing decision consistency.

From an information systems perspective, predictive models create value when embedded within accountable decision processes rather than treated as automatically authorized actions (Akter et al., 2019; Provost & Fawcett, 2013). This is particularly important in financial environments, where AI-enabled decisions raise concerns about transparency, accountability, risk control, and trustworthy system behavior (Guidotti et al., 2018; Rudin, 2019; Díaz-Rodríguez et al., 2023).

Despite advances in multi-modal forecasting, analytics-based decision support, responsible AI governance, and adaptive learning, these streams remain only partially integrated within financial execution workflows. Limited research connects forecast generation with opportunity-level execution authorization, predictive throttling, counterfactual evaluation of suppressed candidates, and evidence-controlled adaptation. GDFI addresses this gap through a governed financial decision lifecycle linking prediction, execution authorization, outcome evaluation, and adaptation.

Accordingly, this study addresses the following research questions:

RQ1: Does governance-aware execution select forecast opportunities with stronger realized outcomes than counterfactual suppressed candidates evaluated under identical horizon-matured outcome rules?

RQ2: Does the Learning Governor enforce pre-specified evidence-sufficiency requirements by deferring model or threshold adaptation when the accumulated realized evidence remains below the authorization threshold?

The study makes four contributions. First, it introduces a governance-driven financial intelligence architecture that separates forecast generation from execution authorization. Second, it operationalizes predictive throttling as a decision-control mechanism that suppresses weaker opportunities before capital allocation. Third, it evaluates governance value through an approved-versus-suppressed comparison under identical ex-post evaluation rules. Fourth, it extends governance beyond execution through a Learning Governor that blocks retraining or threshold modification until sufficient realized evidence has accumulated. Together, these contributions position GDFI as an information systems governance architecture for AI-enabled financial decision support under uncertainty.

2. RELATED WORK

Financial forecasting research has shown that machine learning and deep learning methods can provide measurable value in financial prediction and trading applications. Krauss et al. (2017) reported out-of-sample daily returns exceeding 0.45% before transaction costs for an ensemble-based long-short strategy using S&P 500 constituents. AI-in-finance research has also documented applications across credit scoring, fraud detection, robo-advisory services, digital insurance, and financial inclusion (Cao,

ISSN: 2473-4901
v12 n6752

2022; Vuković et al., 2025). In information systems research, Wamba et al. (2017) linked big data analytics capability to firm performance directly and through dynamic capabilities.

Multi-modal financial intelligence combines technical, sentiment, fundamental, macroeconomic, and relational signals to provide richer market representations than single-source models (Krauss et al., 2017; Kumbure et al., 2022; Vuković et al., 2025; Wamba et al., 2017). However, the literature largely emphasizes prediction quality, analytics capability, or AI adoption rather than whether model-generated opportunities should be approved, suppressed, or deferred before execution.

Decision-support research emphasizes embedding predictive insights within operational decision processes (Akter et al., 2019; Provost & Fawcett, 2013), while responsible-AI research highlights transparency, interpretability, oversight, and governance in high-stakes systems (Guidotti et al., 2018; Rudin, 2019; Díaz-Rodríguez et al., 2023). Financial governance discussions similarly emphasize monitoring, risk control, and responsible AI deployment because of potential effects on market behavior and institutional accountability (Financial Stability Board, 2024; U.S. Department of the Treasury, 2024; Papagiannidis et al., 2025; Choowan & Daovisan, 2026).

Adaptive-learning research examines how systems update behavior as new evidence becomes available (Parisi et al., 2019; Sutton & Barto, 2018). Although adaptation can improve responsiveness under changing conditions, updates based on limited realized evidence may introduce overfitting, instability, or sensitivity to short-term noise. Table 1 synthesizes these literature streams and the integration gap addressed by GDFI.

Literature stream	Primary focus	Limitation addressed by GDFI
Financial forecasting and multi-modal analytics	Predictive accuracy and richer market representations	Forecast outputs are not explicitly governed as execution decisions
Analytics-based decision support	Embedding analytics within operational decision processes	Limited opportunity-level mechanisms for approving, suppressing, or resizing individual forecast opportunities
Responsible AI governance	Transparency, accountability, oversight, and trustworthy system behavior	Governance principles are often conceptual rather than operationalized within financial execution workflows
Adaptive and continual learning	Updating model behavior as new evidence becomes available	Adaptation may occur without an independent evidence-sufficiency authorization layer
GDFI	Forecasting, execution governance, predictive throttling, retained counterfactual outcome evaluation, and controlled adaptation	Integrates the complete prediction-to-decision-and-learning lifecycle

Table 1. Synthesis of Related Literature and the GDFI Integration Gap

3. METHODOLOGY

This section describes GDFI across the decision lifecycle: feature construction, opportunity generation, governance evaluation, predictive throttling, outcome assessment, and evidence-driven adaptive learning.

3.1 System Design and Workflow

The proposed GDFI framework transforms heterogeneous market information into governed execution and adaptation decisions while separating prediction generation from execution authorization. Figure 1 illustrates five interacting components: multi-modal intelligence acquisition, directional-edge forecasting, governance evaluation, predictive throttling and execution control, and evidence-driven adaptive learning.

ISSN: 2473-4901
v12 n6752

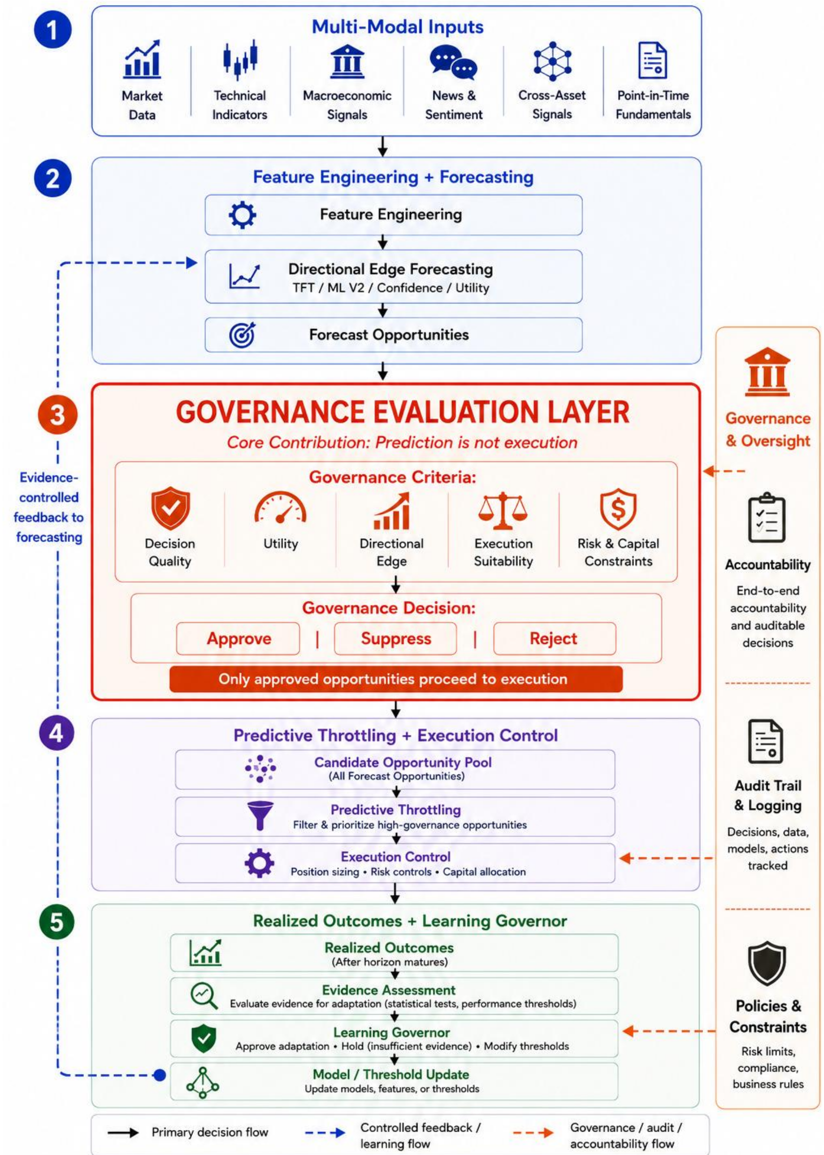

Figure 1. Governance-Driven Financial Intelligence Architecture.

Note: This explanatory architecture figure was created by the author with AI-assisted visualization support and was reviewed, edited, and finalized by the author for accuracy and alignment with the proposed GDFI framework.

The following algorithm summarizes the GDFI decision lifecycle. A forecast opportunity is a model-generated asset-horizon prediction that has not yet been approved or suppressed.

Input: Market observations, engineered feature vector, forecast opportunity, predicted direction, confidence score, expected edge, utility score, trade-quality estimate, and governance thresholds.

Output: Execution decision: APPROVE or SUPPRESS. Adaptation decision: AUTHORIZE or DEFER.

1. Construct the multi-modal prediction-time feature vector.
2. Generate a forecast opportunity and estimate its direction, confidence, edge, utility, and trade quality.
3. Evaluate the opportunity against utility, directional-edge, trade-quality, exposure, and risk-control requirements.
4. If hard governance requirements are violated, assign SUPPRESS and record the reason.

 ISSN: 2473-4901
v12 n6752

5. Otherwise, assign APPROVE when final authorization criteria are satisfied; if not, assign SUPPRESS and retain the opportunity for approved-versus-suppressed outcome evaluation.
6. After horizon maturity, compute realized outcomes and update the realized evidence memory.
7. Evaluate the Learning Governor's evidence requirements and assign AUTHORIZE when they are satisfied; otherwise, assign DEFER and block adaptation.

Appendix A provides the detailed governance authorization logic, including hard vetoes, threshold checks, and final approval rules.

3.2 Intelligence Acquisition and Feature Construction

GDFI constructs a multi-modal feature set from technical indicators, macroeconomic variables, news sentiment, graph-relational signals, point-in-time fundamentals, and additional market features. These variables convert raw market observations into structured prediction-time inputs representing price movement, volatility, volume behavior, macroeconomic conditions, investor sentiment, and cross-asset relationships for forecasting and governance evaluation.

Missing-data treatment was feature-specific. FinBERT-scored news was aggregated by ticker and normalized publication date into daily sentiment and rolling news-event features; missing news-derived variables were set to zero after the ticker-date merge. Event-derived variables used neutral zero values or predefined distance sentinels when unavailable. SEC fundamentals were merged by filing-availability date and carried forward within ticker. The forecasting target was normalized by ticker using a group-specific normalizer.

3.3 TFT-Based Directional Edge Forecasting

Candidate opportunities were generated using a separately trained Temporal Fusion Transformer with a 48-step encoder, five-step prediction horizon, hidden size 64, one recurrent layer, four attention heads, dropout of 0.20, batch size 64, learning rate 0.0001, Quantile Loss, and seven output quantiles. Training followed stable-market, stress-market, and mixed-market curriculum phases, with early stopping and checkpoint selection based on minimum validation loss. Within each phase, observations through *max_time_idx − 5* were used for fitting and the terminal five steps for validation. The best Phase-3 checkpoint remained fixed. The final model was fitted through April 16, 2026, and then frozen; operational signal generation began May 1, 2026, and May 5, 2026, was the earliest evaluated forecast date. During the reported evaluation, the checkpoint was used only for inference, with no retraining, fine-tuning, hyperparameter optimization, or parameter updating.

The TFT estimates directional edge rather than exact price. Each output contains an expected direction, confidence estimate, and supporting evidence from the engineered feature space. The target, target_edge_return_5, represented signed five-step excess return beyond a hybrid neutral threshold defined as the greater of 0.5% and one-half of the 20-period realized volatility scaled by the square root of five.

The forecasting module generates opportunities rather than execution instructions; governance determines authorization because a favorable directional signal may still fail utility, risk, evidence, or execution-suitability requirements.

3.4 Governance Evaluation and Predictive Throttling

The Governance Evaluation Layer assesses each forecast opportunity using forecast quality, expected utility, directional-edge strength, risk exposure, and execution suitability. Stronger opportunities may proceed toward approval, whereas weaker opportunities are suppressed before authorization. Predictive throttling operationalizes this decision by approving opportunities that satisfy governance criteria and retaining suppressed candidates for counterfactual evaluation.

 ISSN: 2473-4901
v12 n6752

3.4.1 Governance Parameter Selection

The frozen production policy reflected prior threshold-tuning audit history and was distinct from Learning Governor authorization. Table 2 reports the governance parameters active at the 36-realized-outcome evaluation snapshot.

Parameter	Value
Alpha threshold	0.0002
Utility-Q threshold	0.001
Directional-edge threshold	0.00005
ML V2 quality threshold	0.85
Fallback ML threshold	0.66
Allocation range	1.0%–5.0%
Maximum approved opportunities per day	3
Maximum gross exposure	3.1381%
Minimum realized outcomes for adaptation	100
Minimum new outcomes after an approved update	20

Table 2. Governance Parameters at the Frozen Evaluation Snapshot

3.5 Evidence-Driven Adaptive Learning

GDFI includes an Evidence-Driven Learning Governor that regulates model adaptation. Because unrestricted retraining under limited evidence may introduce overfitting, instability, or sensitivity to short-term noise, adaptation is treated as a governed rather than automatic process.

Before retraining or threshold modification, the Learning Governor evaluates whether sufficient realized outcomes have accumulated; otherwise, adaptation is deferred. The minimum threshold was set to 100 as a conservative, pre-specified governance control to prevent adaptation based on small or noisy samples and was not tuned to improve reported performance. Following an authorized update, at least 20 additional realized outcomes were required before another modification could be considered.

3.6 Evaluation Protocol

The study used sequential paper trading without live-capital deployment. Governance decisions used information available at each forecast timestamp, and realized outcomes were incorporated only after horizon maturity. Suppressed candidates were retained for counterfactual evaluation under the same horizon-specific rules as approved opportunities. At the frozen cutoff, 200 of 480 opportunities had matured: 58 were approved and 142 suppressed. Separately, 36 closed outcomes were available to the Learning Governor.

Portfolio comparisons used gross realized paper-trading returns without transaction-cost, slippage, or market-impact deductions. The equal-weight benchmark represented unrestricted participation in eligible mature opportunities, whereas GDFI applied selective governance and controlled allocation. Because gross exposure was not matched, the comparison should be interpreted as descriptive rather than as an estimate of deployable investment performance.

3.7 Threats to Validity and Leakage Controls

Several controls were used to reduce forward-looking bias and protect temporal validity. Governance decisions were based only on information available at the corresponding prediction timestamp, realized outcomes were incorporated only after the applicable forecast horizon matured, and suppressed candidates were retained for counterfactual evaluation without using their future outcomes in the original authorization decision. SEC fundamental variables were aligned using filing-availability dates, the best Phase-3 TFT checkpoint remained fixed throughout the reported governance evaluation, and

ISSN: 2473-4901
v12 n6752

adaptive model or threshold changes were blocked unless the Learning Governor's pre-specified evidence requirements were satisfied.

The study is limited by the ten-equity universe, 200 horizon-matured evaluations, 58 approved opportunities, a non-exposure-matched equal-weight benchmark, and the use of gross paper-trading returns. The modest approved sample and smaller horizon-specific subsets limit statistical power, particularly for hit-rate and between-horizon comparisons. Accordingly, the findings indicate governance selectivity and risk-control behavior within the frozen evaluation snapshot rather than causal superiority, universal profitability, or deployable investment performance.

4. DATA COLLECTION

Historical price and volume data for the ten-equity universe and benchmark instruments were collected from Yahoo Finance through the yfinance Python library. Benchmark and market-context instruments included SPY, QQQ, IWM, XLK, XLF, XLE, and the VIX index as a market-stress proxy. Macroeconomic variables were collected from the Federal Reserve Economic Data (FRED) database using the fredapi client. FRED series were ordered by observation date, forward-filled, and backward-merged onto each market date using the latest dated observation. Financial news and sentiment inputs were derived from the Nasdaq/FNSPID dataset and Alpha Vantage news-sentiment data and transformed using ProsusAI/FinBERT where applicable. Point-in-time fundamental indicators were collected from SEC EDGAR companyfacts filings and merged using filing availability dates to reduce forward-looking bias. Additional event and graph-relational features were engineered internally from earnings-calendar information, price and volume shocks, benchmark relationships, and cross-asset structure. These sources were processed through a time-aligned feature-engineering pipeline to construct the final modeling dataset.

The final modeling dataset contained 20,700 observations spanning March 7, 2018, through June 1, 2026, across ten major equities. Table 3 summarizes its principal characteristics.

Metric	Value
Observation Rows	20,700
Feature Columns	168
Start Date	March 7, 2018
End Date	June 1, 2026
Number of Tickers	10

Table 3. Dataset characteristics used to evaluate the Governance-Driven Financial Intelligence framework.

The framework used 168 engineered features spanning technical, macroeconomic, sentiment, graph-relational, point-in-time fundamental, and other market information. As shown in Table 4, technical and price-based indicators accounted for 68 variables (40.5%), followed by 34 other engineered, 24 news and sentiment, 16 macroeconomic, 14 graph-relational, and 6 fundamental features. The larger technical share reflects price, return, volatility, volume, and benchmark-relative transformations across multiple lookback windows, while TFT variable selection learned their relative contributions. Remaining fields comprised identifiers, target labels, and market-regime variables.

 ISSN: 2473-4901
v12 n6752

Feature Category	Count	Percentage
Price / Technical	68	40.5%
Other Engineered Features	34	20.2%
News / Sentiment	24	14.3%
Macroeconomic	16	9.5%
Graph / Relational	14	8.3%
Fundamental	6	3.6%
Identifier / Time Index	3	1.8%
Target / Decision Label	2	1.2%
Regime / Stress	1	0.6%
Total	**168**	**100.0%**

Table 4. Distribution of engineered features across the multi-modal intelligence architecture.

A complete variable-level dictionary for the reported 168-variable frozen TFT modeling schema, including each field's domain, modeling role, definition, and construction source, is provided in Appendix C (Table C1).

Encoder variable-selection outputs from the frozen Phase-3 TFT identified dollar_index_proxy as having the highest encoder variable-selection weight among named engineered covariates; realized_vol_10, graph_avg_corr_20, peer_mean_rsi_14, rolling_beta_spy_20, news_count_recent_5d, and news_sentiment_z_20 were also influential. These model-specific weights are interpretive rather than causal.

As a supplementary redundancy diagnostic, within-ticker Pearson correlations were evaluated across 153 numeric prediction-time covariates after excluding identifiers, targets, realized outcomes, prediction errors, governance outputs, and decision labels. Among 11,628 evaluated pairs, 73 (0.63%) had $|r| \geq 0.90$, primarily among OHLC prices, benchmark measures, sentiment transformations, volume-shock indicators, and graph-spillover features. The strongest correlations are reported in Appendix Table B1.

5. DATA ANALYSIS

The analysis evaluated five dimensions: governance selectivity, decision quality, portfolio-level risk control, forecast-horizon performance, and adaptive-learning oversight.

First, governance selectivity was analyzed by tracking forecast opportunities through outcome maturity and authorization. At the frozen reporting cutoff, 200 of 480 opportunities had reached their required T+4 or T+5 horizons and formed the mature evaluation population; 58 were approved and 142 were suppressed.

Second, decision quality compared approved and suppressed opportunities after horizon maturity under identical horizon-specific outcome rules. Measures included approved and suppressed hit rates, governance lift, mean realized-PnL difference, permutation-test significance, and Cohen's d. Hit rate was the percentage of mature opportunities with positive realized directional PnL, and governance lift was the approved-minus-suppressed hit-rate difference. Wilson 95% confidence intervals were calculated for each hit rate, and the Newcombe method estimated the interval for the difference between independent proportions. Mean realized-PnL differences were assessed using a two-sided test with 5,000 permutations of the pooled observations while preserving the approved-group sample size. Cohen's d standardized the difference using the pooled standard deviation and was reported with a noncentral-t confidence interval.

ISSN: 2473-4901
v12 n6752

Third, portfolio-level risk control compared GDFI with an unrestricted equal-weight benchmark. GDFI applied selective authorization and controlled allocation, and gross exposure was not matched. Measures included total return, mean return, volatility, Sharpe ratio, and maximum drawdown based on gross realized paper-trading returns without modeled transaction costs, slippage, or market impact. The comparison therefore assessed risk–participation consequences rather than exposure-matched investment superiority or deployable net performance.

Fourth, forecast-horizon performance was analyzed using hit rate, mean PnL, Sharpe ratio, and maximum drawdown. Because the executed T+4 and T+5 samples contained only 32 and 26 opportunities, respectively, the horizon comparison was treated as a descriptive diagnostic rather than as a formal statistical test of horizon superiority. This analysis examined whether forecast horizon could serve as an additional governance signal.

Fifth, adaptive-learning behavior was analyzed by comparing the available closed outcomes with the pre-specified initial authorization requirement of 100 and recording the Learning Governor's decision. This analysis evaluated whether GDFI regulated model adaptation through evidence sufficiency rather than allowing automatic updates based on short-term outcomes.

6. RESULTS AND DISCUSSION

This section presents the empirical findings from the GDFI evaluation across five dimensions: governance selectivity, decision quality, portfolio-level risk control, forecast-horizon performance, and evidence-driven adaptive learning. The findings show that GDFI functions as a selective decision-intelligence framework rather than an unrestricted forecasting pipeline.

6.1 Governance Funnel Effectiveness

The first evaluation examined how forecast-generated opportunities moved through the governance and execution pipeline. Table 5 summarizes the core governance funnel, including forecast opportunities, mature evaluations, approved opportunities, and rejected signals. Figure 2 provides a broader process-level view of the same pipeline by also showing downstream realized-outcome tracking used for learning-governor evidence accumulation.

Stage	Count	% of Forecast Opportunities
Forecast Opportunities	480	100.0%
Mature Evaluations	200	41.7%
Approved Opportunities	58	12.1%
Rejected Signals	142	29.6%

Table 5. Progression of opportunities through the governance and execution pipeline.

As shown in Table 5, the framework generated 480 forecast opportunities, of which 200 had reached their required T+4 or T+5 outcome horizons by the frozen reporting cutoff and therefore formed the mature evaluation population. The remaining 280 opportunities had not yet reached their applicable outcome horizons and were excluded from realized-outcome comparisons rather than classified as successes or failures. Following governance evaluation, 58 mature opportunities were approved for execution and 142 were suppressed. Thus, only 12.1% of all forecast opportunities and 29.0% of mature opportunities received execution authorization.

ISSN: 2473-4901
v12 n6752

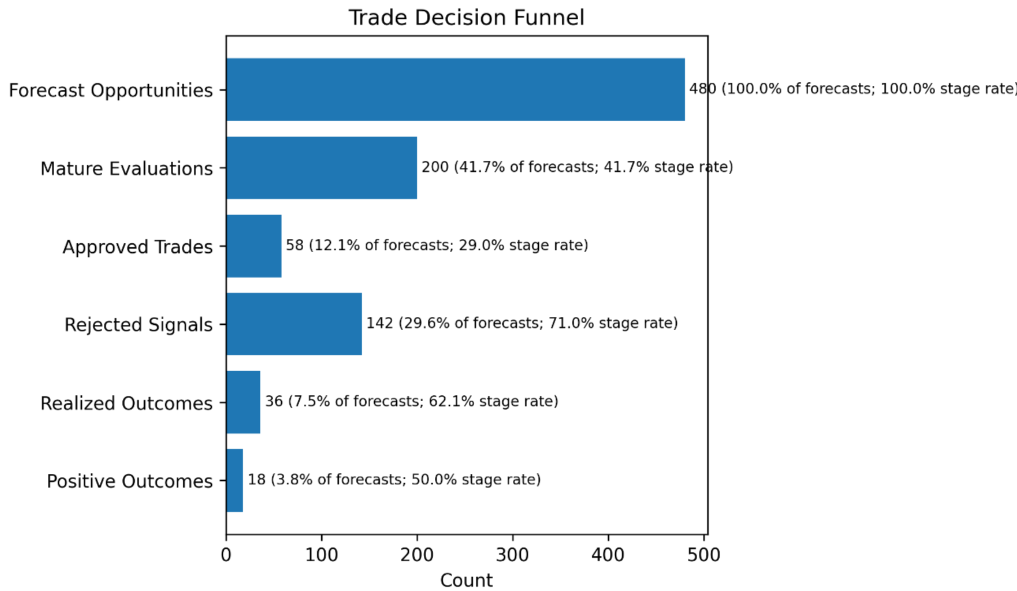

Figure 2. Trade Decision and Outcome Tracking Funnel.

6.2 Decision Quality Improvement Through Governance

The second evaluation examines whether governance approval is associated with stronger realized decision quality under a controlled ex-post evaluation framework. To ensure a fair comparison, governance-approved opportunities and governance-suppressed candidates were evaluated after the same forecast horizon matured, using identical horizon-specific outcome rules and holding-period assumptions. Approved opportunities represent candidates authorized for paper-trading execution, whereas suppressed candidates were not executed but were retained as counterfactual observations and evaluated under the same ex-post outcome rules.

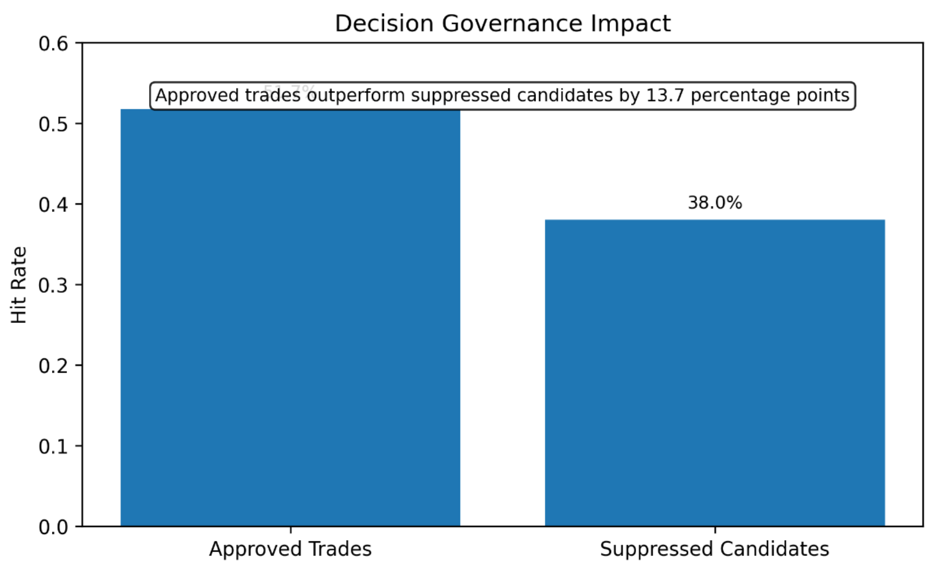

Figure 3. Decision Governance Impact.

As shown in Figure 3, governance-approved opportunities achieved a realized hit rate of 51.7% (30 of 58; Wilson 95% confidence interval [CI]: 39.2%–64.1%), compared with 38.0% (54 of 142; Wilson 95% CI: 30.5%–46.2%) for suppressed candidates. This corresponded to an observed governance lift

 ISSN: 2473-4901
v12 n6752

of 13.7 percentage points, with a Newcombe 95% CI from −1.3 to 28.2 percentage points. Because the confidence interval for the difference in proportions included zero, the observed hit-rate lift was interpreted as descriptive evidence of governance selectivity rather than a statistically significant difference in hit rates.

Metric	Value
Approved Hit Rate	51.7%
Suppressed Hit Rate	38.0%
Governance Lift	+13.7 percentage points
Mean PnL Difference	+0.000298
Permutation p-value	0.001
Cohen's d	0.747
Cohen’s d, 95% CI	0.432–1.060

Table 6. Comparison of realized ex-post outcomes between governance-approved opportunities and suppressed candidates with matured forecast horizons.

For the permutation test, the null hypothesis stated that governance labels were exchangeable with respect to realized PnL under the common ex-post evaluation rules. The two-sided alternative hypothesis stated that the mean realized PnL differed between governance-approved and suppressed opportunities. The observed mean realized-PnL difference was +0.000298, and the two-sided permutation test based on 5,000 permutations produced $p = 0.001$. Cohen’s d was 0.747, with a 95% noncentral-t CI from 0.432 to 1.060, indicating a positive moderate-to-large standardized difference between the two realized-PnL distributions.

Overall, governance approval was associated with stronger realized outcomes. The hit-rate lift was positive but statistically uncertain, whereas the continuous realized-PnL comparison provided stronger inferential evidence of separation. Because suppressed outcomes were counterfactual ex-post evaluations rather than executed paper-trade returns, these findings support governance selectivity but do not establish causal or universally generalizable performance improvement.

6.3 Predictive Throttling and Portfolio-Level Risk Control

The third evaluation examined whether predictive throttling was associated with stronger portfolio-level risk control. Figure 4 compares GDFI with an unrestricted equal-weight benchmark under the same universe, matured evaluation window, opportunity set, holding-period rules, and outcome-maturity rules. The benchmark participated in all eligible mature opportunities, whereas GDFI applied selective authorization and controlled allocation. Returns are reported in normalized portfolio-return units; because gross exposure was not matched, the comparison reflects selective participation and allocation rather than exposure-matched strategy superiority.

ISSN: 2473-4901
v12 n6752

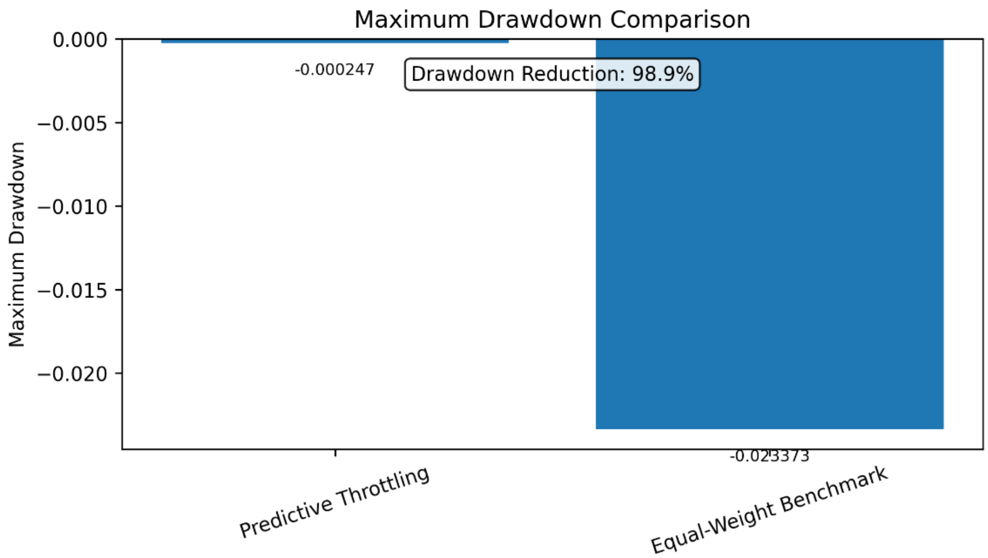

Figure 4. Maximum drawdown comparison between GDFI predictive throttling and an unrestricted equal-weight benchmark over the same asset universe and matured evaluation window.

Predictive throttling was associated with an approximately 98.9% reduction in maximum drawdown, from −0.023373 to −0.000247. Sharpe ratio was calculated as unannualized mean normalized return divided by volatility with a zero risk-free rate. Returns used normalized portfolio-return units under a zero-cost paper-trading convention; transaction-cost and slippage fields were zero in the archived snapshot, and market impact was not modeled.

Metric	Predictive Throttling Framework	Equal-Weight Benchmark
Total Return	0.000201	0.001639
Mean Return	0.000015	0.000146
Volatility	0.000097	0.006553
Sharpe Ratio	0.160	0.022
Maximum Drawdown (MDD)	-0.000247	-0.023373

Table 7. Portfolio-level comparison between the predictive-throttling framework and an equal-weight benchmark under a common paper-trading evaluation framework.

As shown in Table 7, the equal-weight benchmark produced higher return but substantially greater volatility and drawdown, whereas GDFI exhibited lower volatility and a higher Sharpe ratio. The approximately 98.9% drawdown reduction was accompanied by lower absolute return, indicating a risk–return and participation tradeoff. Because exposure was not matched, the findings support downside-risk control rather than exposure-matched investment superiority.

6.4 Forecast Horizon Performance

The fourth evaluation examined whether governance-aware execution quality varied across forecast horizons. T+4 and T+5 denote four- and five-trading-day horizons after the prediction date. Figure 5 compares approved-opportunity hit rate and maximum drawdown. A hit required positive realized directional PnL after horizon maturity under the common zero-cost convention; identical evaluation rules were applied except for maturity length.

ISSN: 2473-4901
v12 n6752

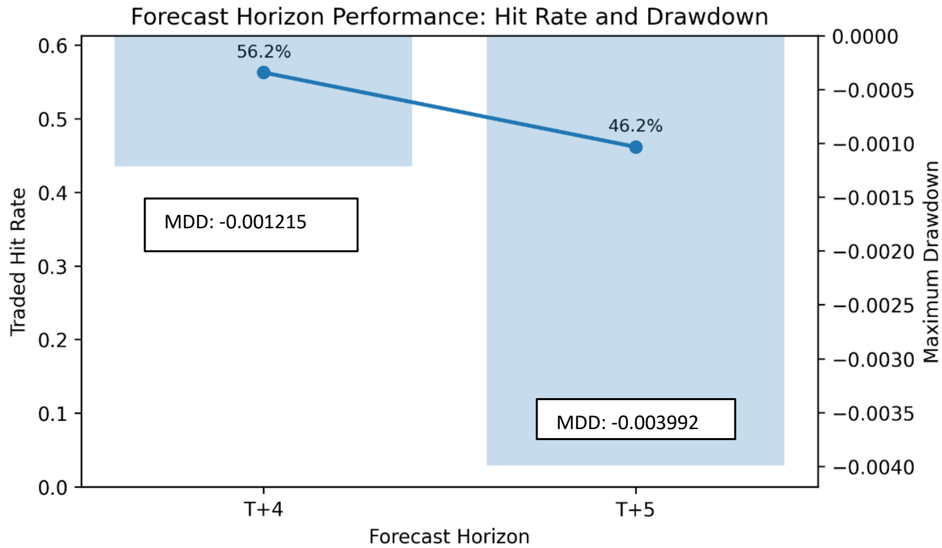

Figure 5. Forecast-horizon performance based on trade-level hit rate and maximum drawdown after four-trading-day and five-trading-day horizon maturity.

Figure 5 shows that the T+4 horizon exhibited stronger observed performance than the T+5 horizon. The T+4 horizon achieved a trade-level hit rate of 56.2% and maximum drawdown of −0.001215, whereas the T+5 horizon achieved a trade-level hit rate of 46.2% and maximum drawdown of −0.003992. This pattern suggests that realized execution quality varied across forecast horizons within the evaluated snapshot. Figure 6 presents the Sharpe-ratio comparison across forecast horizons.

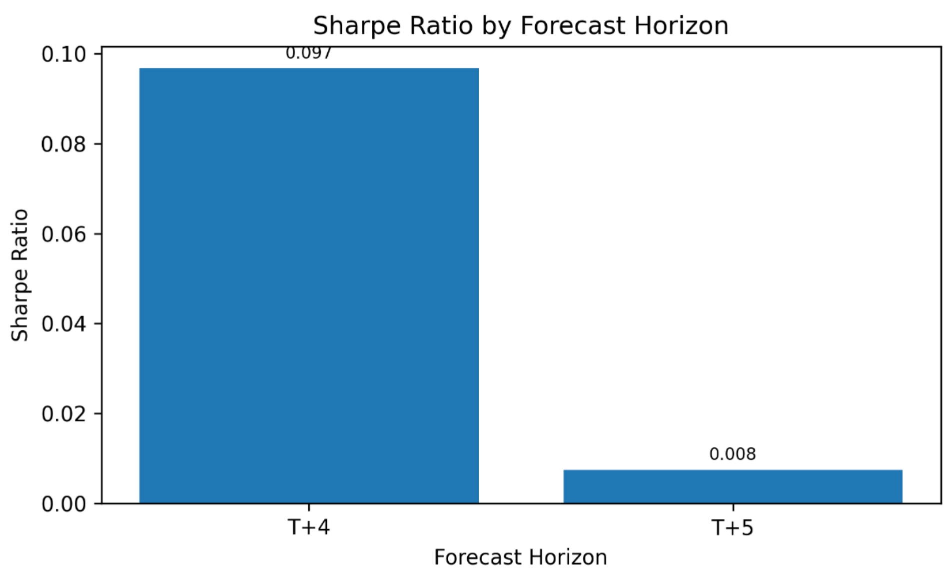

Figure 6. Sharpe Ratio by Forecast Horizon.

As shown in Figure 6, the T+4 horizon also achieved stronger risk-adjusted performance, with a Sharpe ratio of 0.097 compared with 0.008 for T+5.

ISSN: 2473-4901
v12 n6752

Metric	T+4	T+5
Evaluable Opportunities	90	110
Executed Trades	32	26
Mean PnL	+0.000049	+0.000002
Sharpe Ratio	0.097	0.008
Hit Rate	56.2%	46.2%
Maximum Drawdown	-0.001215	-0.003992

Table 8. Performance comparison across forecast horizons using identical ex-post evaluation rules after horizon maturity.

Table 8 shows stronger observed T+4 values across hit rate, mean PnL, Sharpe ratio, and maximum drawdown. However, the approved samples contained only 32 $T+4$ and 26 $T+5$ opportunities; therefore, the comparison is descriptive, and no statistically significant horizon superiority is claimed. Forecast horizon may warrant consideration as an additional governance signal in larger future evaluations.

6.5 Evidence-Driven Adaptive Learning

The final evaluation examined whether the Learning Governor enforced evidence-based adaptation. Figure 7 compares available realized learning outcomes with the minimum evidence threshold for adaptive updates.

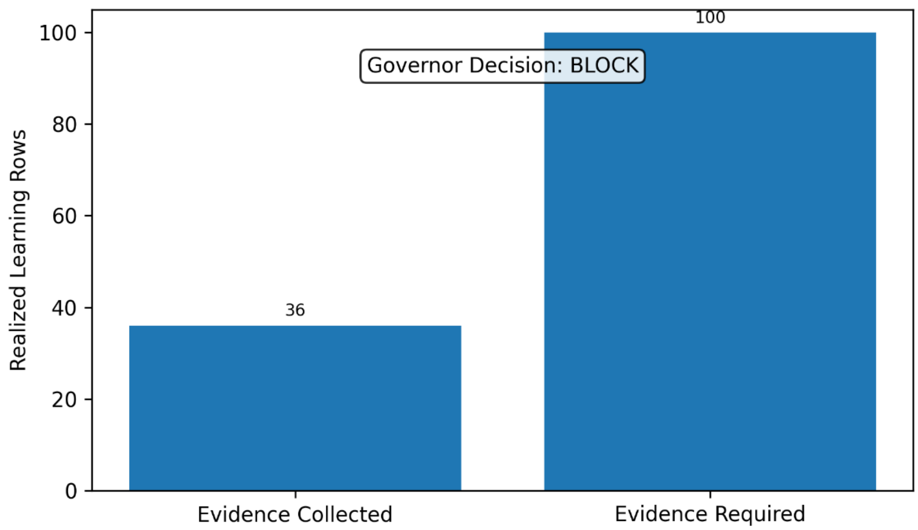

Figure 7. Learning Governor Evidence Requirement.

Figure 7 shows that the system accumulated 36 realized outcomes, which is below the fixed threshold of 100 required by the Learning Governor. This threshold functions as a deterministic data sufficiency rule that gates adaptive updates. As a result, the governor decision was BLOCK. This outcome demonstrates that the system does not automatically retrain or modify thresholds simply because new observations are available. Table 9 summarizes the evidence-sufficiency evaluation for adaptive learning authorization.

Metric	Value
Governor Decision	BLOCK
Approved	False
Realized Outcomes Available	36

ISSN: 2473-4901
v12 n6752

Positive Outcome Rate	50.0%
Mean Realized PnL	0.000141
Production Minimum Outcomes Required	100
New Outcomes Since Last Update	36
Minimum New Outcomes Required	20
Reason	Insufficient realized rows: 36 < 100

Table 9. Evaluation of evidence sufficiency for adaptive learning authorization.

As shown in Table 9, the realized learning set contained 36 outcomes, a 50.0% positive outcome rate, and a positive mean realized PnL of 0.000141. However, the evidence threshold was not satisfied because realized outcomes remained below the pre-specified minimum of 100. The Learning Governor therefore deferred adaptation despite favorable recent performance, prioritizing evidence sufficiency over short-term responsiveness.

The 100-outcome requirement was established as a conservative authorization threshold before the reported evaluation and should not be interpreted as an empirically demonstrated optimal sample size. The additional requirement of 20 new outcomes applies only after a prior update has been approved and serves to prevent repeated modification based on substantially overlapping evidence.

The 36 realized learning outcomes formed a separate operational-learning population and were not an additional subset or sequential stage of the 200-row approved-versus-suppressed evaluation population.

6.6 Discussion

The results indicate that GDFI supported financial decision-making by governing execution authorization and adaptive modification. Approved opportunities were associated with stronger realized outcomes than suppressed candidates retained for counterfactual evaluation, with stronger inferential support from realized PnL than from the descriptive hit-rate lift. Predictive throttling was associated with substantially lower drawdown than unrestricted equal-weight participation, but also with lower absolute return, participation, and exposure. T+4 showed stronger descriptive performance than T+5, while the Learning Governor deferred adaptation because available evidence remained below the pre-specified requirement.

The central insight is that GDFI creates decision-support value by separating forecasting from execution and adaptation authorization. Forecasting identifies candidate opportunities, while governance determines which satisfy quality, utility, exposure, and risk requirements; retained suppressed candidates support counterfactual evaluation. This shifts financial intelligence from a prediction-only process toward a governed decision lifecycle focused on selective participation, downside-risk control, accountability, and evidence-controlled adaptation.

The findings also reveal important tradeoffs. Restrictive governance can reduce downside exposure while suppressing profitable opportunities, reducing participation, and lowering absolute return. Evidence thresholds can protect against premature adaptation while delaying updates when outcomes accumulate slowly. GDFI should therefore be understood as a framework for governing prediction-to-action and adaptation decisions under uncertainty rather than as evidence of universally superior trading performance.

Theoretically, GDFI extends analytics-based decision support by separating prediction generation, execution authorization, and adaptation authorization within one governed lifecycle. Practically, implementation requires reliable point-in-time data, auditable approval and suppression records, retained counterfactual outcomes, and periodic threshold validation. Computational demands are concentrated in offline feature construction and model training, while routine inference can reuse a frozen checkpoint and cached features. Although evaluated in finance, evidence-gated authorization may transfer to other consequential AI-enabled decision settings, subject to domain-specific validation.

ISSN: 2473-4901
v12 n6752

7. CONCLUSION AND FUTURE RESEARCH

This study presented Governance-Driven Financial Intelligence (GDFI), a multi-modal decision-support framework integrating directional-edge forecasting, governance-aware execution, predictive throttling, and evidence-driven adaptive learning. Its central contribution is separating forecast generation from execution authorization: forecasts remain candidate opportunities until governance criteria authorize paper-trading execution.

Within the frozen evaluation snapshot, approved opportunities achieved a 51.7% hit rate versus 38.0% for suppressed candidates, an observed governance lift of 13.7 percentage points. Because its confidence interval included zero, the lift is interpreted as descriptively positive rather than statistically significant. Stronger evidence came from realized PnL ($p = 0.001$; Cohen's $d = 0.747$), indicating a positive moderate-to-large standardized difference between groups. Predictive throttling was associated with an approximately 98.9% reduction in maximum drawdown relative to the unrestricted equal-weight benchmark, but with lower absolute return and participation under the non-exposure-matched design.

The Learning Governor returned BLOCK because available evidence did not meet the pre-specified authorization requirement. Limitations include the frozen evaluation snapshot, ten-equity universe, 200 horizon-matured evaluations, non-exposure-matched benchmark, zero-cost convention, and sequential paper-trading setting. Future research should examine larger outcome populations, longer horizons, broader asset universes and market regimes, transaction costs, slippage, market impact, governance-threshold sensitivity, cumulative-return paths, formal horizon comparisons and stronger exposure- or risk-matched benchmarks.

REFERENCES

Akter, S., Bandara, R. J., Hani, U., Fosso Wamba, S., Foropon, C., & Papadopoulos, T. (2019). Analytics-based decision-making for service systems: A qualitative synthesis and agenda for future research. *International Journal of Information Management, 48*, 85–95. https://doi.org/10.1016/j.ijinfomgt.2019.01.020

Cao, L. (2022). AI in finance: Challenges, techniques, and opportunities. *ACM Computing Surveys, 55*(3), 1–38. https://doi.org/10.1145/3502289

Choowan, P., & Daovisan, H. (2026). Artificial intelligence in data governance for financial decision-making: A systematic review. *Big Data and Cognitive Computing, 10*(1), Article 8. https://doi.org/10.3390/bdcc10010008

Díaz-Rodríguez, N., Del Ser, J., Coeckelbergh, M., López de Prado, M., Herrera-Viedma, E., & Herrera, F. (2023). Connecting the dots in trustworthy artificial intelligence: From AI principles, ethics, and key requirements to responsible AI systems and regulation. *Information Fusion, 99*, Article 101896. https://doi.org/10.1016/j.inffus.2023.101896

Financial Stability Board. (2024, November 14). The financial stability implications of artificial intelligence. *Basel: Financial Stability Board*. https://www.fsb.org/2024/11/the-financial-stability-implications-of-artificial-intelligence/

Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning* (1st ed.). MIT Press.

Guidotti, R., Monreale, A., Ruggieri, S., Turini, F., Giannotti, F., & Pedreschi, D. (2018). A survey of methods for explaining black box models. *ACM Computing Surveys, 51*(5), 1–42. https://doi.org/10.1145/3236009

Krauss, C., Do, X. A., & Huck, N. (2017). Deep neural networks, gradient-boosted trees, random forests: Statistical arbitrage on the S&P 500. *European Journal of Operational Research, 259*(2), 689–702. http://doi.org/10.1016/j.ejor.2016.10.031

ISSN: 2473-4901
v12 n6752

Kumbure, M. M., Lohrmann, C., Luukka, P., & Porras, J. (2022). Machine learning techniques and data for stock market forecasting: A literature review. Expert Systems with Applications, 197, 116659. https://doi.org/10.1016/j.eswa.2022.116659

Papagiannidis, E., Mikalef, P., & Conboy, K. (2025). Responsible artificial intelligence governance: A review and research framework. *The Journal of Strategic Information Systems, 34*(2), Article 101885. https://doi.org/10.1016/j.jsis.2024.101885

Parisi, G. I., Kemker, R., Part, J. L., Kanan, C., & Wermter, S. (2019). Continual lifelong learning with neural networks: A review. *Neural Networks, 113*, 54–71. https://doi.org/10.1016/j.neunet.2019.01.012

Provost, F., & Fawcett, T. (2013). *Data science for business: What you need to know about data mining and data-analytic thinking*. O'Reilly Media.

Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence, 1*(5), 206–215. https://doi.org/10.1038/s42256-019-0048-x

Shmueli, G., & Koppius, O. R. (2011). Predictive analytics in information systems research. MIS Quarterly, 35(3), 553–572. https://doi.org/10.2307/23042796

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

U.S. Department of the Treasury. (2024, December 19). *Artificial intelligence in financial services: Report on the uses, opportunities, and risks of artificial intelligence in the financial services sector*. U.S. Department of the Treasury.

Vuković, D. B., Dekpo-Adza, S., & Matović, S. (2025). AI integration in financial services: A systematic review of trends and regulatory challenges. *Humanities and Social Sciences Communications*, *12*, Article 562, https://doi.org/10.1057/s41599-025-04850-8

Wamba, S. F., Gunasekaran, A., Akter, S., Ren, S. J., Dubey, R., & Childe, S. J. (2017). Big data analytics and firm performance: Effects of dynamic capabilities. *Journal of Business Research, 70*, 356–365. https://doi.org/10.1016/j.jbusres.2016.08.009

ISSN: 2473-4901
v12 n6752

Appendix A. Implementation Snapshot of Governance Authorization Logic

This implementation snapshot shows the final production governance gate used to convert forecast opportunities into approved or suppressed execution decisions. The code applies hard execution vetoes, utility thresholds, directional-edge thresholds, trade-quality requirements, and exposure-related controls before assigning a final execution decision. Sensitive configuration details, credentials, and system-specific paths are omitted.

Algorithm A1. Hard Execution Governor and Threshold Evaluation

***Input*:** Forecast opportunity, raw-trade indicator, proposed action, calibration status, directional-edge score, utility-Q score, ML-quality score, and governance thresholds.

***Output*:** Hard-veto status, effective utility threshold, and provisional governance decision.

1. Verify that the forecast opportunity contains a valid raw trade and actionable direction.
2. Assign a non-overridable hard veto when the upstream calibration layer records hard suppression, the proposed action is missing or represents no trade, or no valid raw trade exists.
3. When a calibration warning is active, increase the effective utility requirement to 1.10 times the base utility threshold; otherwise, retain the base utility threshold.
4. For opportunities without a hard veto, compare the directional-edge, utility-Q, and ML-quality scores with their applicable governance thresholds.
5. Assign provisional approval only when all required thresholds are satisfied. Otherwise, suppress the opportunity and record the failed governance criterion.

Hard-veto decisions cannot be reversed by downstream approval logic.

Appendix A2: Final Approval/Suppression Decision Logic

Input: Provisionally approved opportunities, governance scores, maximum daily approvals, and fallback-approval requirements.

Output: Final production decision and governance explanation for every opportunity.

1. Rank provisionally approved opportunities in descending order of utility-Q score, directional-edge score, and ML-quality score.
2. Retain opportunities up to the maximum permitted number of daily approvals and suppress any excess opportunities under the maximum-trades constraint.
3. When no opportunity remains approved and fallback authorization is enabled, consider only opportunities that contain a valid raw trade, have no hard veto, show positive directional edge and utility, and satisfy the fallback ML-quality threshold.
4. Rank eligible fallback opportunities using utility-Q, directional edge, and ML quality, and authorize only the highest-ranked eligible opportunity.
5. Record the final production-trade indicator, final approval or suppression decision, and corresponding governance explanation for every evaluated opportunity.
6. Apply the subsequent allocation, correlation, gross-exposure, and cooldown controls before paper-trading execution authorization**.**

ISSN: 2473-4901
v12 n6752

Appendix B. Feature Correlation Analysis

Feature redundancy was examined using within-ticker Pearson correlations among 153 numeric prediction-time covariates. This supplementary analysis used a preserved extended modeling snapshot containing 20,760 ticker-date observations from March 7, 2018, through June 9, 2026. Each covariate was standardized separately within ticker before the pooled correlations were calculated. Identifier and indexing fields, forecasting targets, realized outcomes, prediction errors, governance outputs, and decision labels were excluded. The analysis was diagnostic and did not alter the frozen TFT checkpoint, governance parameters, or performance results reported for the primary evaluation dataset.

Feature 1	Feature 2	Pearson r	Pairwise n
close	low	0.9997	20,760
open	high	0.9997	20,760
close	high	0.9996	20,760
open	low	0.9996	20,760
high	low	0.9995	20,760
close	open	0.9992	20,760
sentiment_shock_20	decayed_sentiment_shock	0.9979	18,684
qqq_close	xlk_close	0.9958	20,760
spy_close	qqq_close	0.9944	20,760
volume_z_20	volume_shock_score	0.9933	20,760
spy_close	xlk_close	0.9920	20,760
news_sentiment_recent	news_polarity_imbalance	0.9881	18,684
spy_return_1	graph_market_spillover	0.9787	20,760
qqq_rv20	xlk_rv20	0.9774	20,760
qqq_return_1	xlk_return_1	0.9735	20,760

Table B1. Strongest within-ticker Pearson correlations among prediction-time numeric covariates

***Note*.** Correlations were calculated using pairwise-complete observations after standardizing each covariate within ticker. Pairwise sample sizes vary because some news-derived variables were not available for every ticker-date observation. Of 11,628 evaluated feature pairs, 73, or approximately 0.63%, exhibited an absolute Pearson correlation of at least 0.90. The strongest relationships primarily reflected common underlying information, including OHLC prices, related benchmark instruments, sentiment transformations, volume-shock measures, and graph-spillover indicators. Correlated variables were retained because the frozen TFT model employed learned variable selection; no post-evaluation feature removal, model retraining, or checkpoint modification was performed.

 ISSN: 2473-4901
v12 n6752

Appendix C. Detailed Frozen TFT Variable Dictionary

Table C1 documents the 168-variable schema associated with the frozen Phase-3 Temporal Fusion Transformer. The schema comprises 159 time-varying model variables recorded in the preserved TimeSeriesDataSet parameters, the date-alignment key, the ticker identifier, the known time index, the forecasting target, the observation-weight field, and four TimeSeriesDataSet-generated alignment or normalization fields. The Role column distinguishes forecast inputs from identifiers, targets, weights, and generated support fields.

Table C1. Definitions and construction methods for the 168-variable frozen TFT modeling schema.

No.	Variable	Domain	Role	Definition	Construction / Source
1	**date**	Identifier / Time	Identifier / alignment key	Normalized trading date used to align all ticker-level, benchmark, macroeconomic, fundamental, event, and news observations.	Converted to a timezone-naive daily timestamp and used as the temporal merge key; retained for auditability rather than supplied as a numeric TFT input.
2	**ticker**	Identifier / Time	Static categorical identifier	Equity symbol identifying the modeled company.	Upper-cased and stripped; supplied to the TFT as the static categorical group identifier.
3	**time_idx**	Identifier / Time	Known time index	Sequential integer time index within each ticker series.	Assigned by cumulative row count after sorting each ticker by date; supplied as the known real time index.
4	**close**	Price / Technical	Time-varying model input	Daily closing price of the equity.	Downloaded from Yahoo Finance through yfinance; sorted by ticker and date.
5	**open**	Price / Technical	Time-varying model input	Daily opening price of the equity.	Downloaded from Yahoo Finance through yfinance.
6	**high**	Price / Technical	Time-varying model input	Highest traded price during the daily session.	Downloaded from Yahoo Finance through yfinance.

 ISSN: 2473-4901
v12 n6752

7	**low**	Price / Technical	Time-varying model input	Lowest traded price during the daily session.	Downloaded from Yahoo Finance through yfinance.
8	**volume**	Price / Technical	Time-varying model input	Daily traded share volume.	Downloaded from Yahoo Finance through yfinance.
9	**return_1**	Price / Technical	Time-varying model input	1-session percentage change in the equity closing price.	Computed as close_t / close_(t-1) - 1 within ticker.
10	**return_5**	Price / Technical	Time-varying model input	5-session percentage change in the equity closing price.	Computed as close_t / close_(t-5) - 1 within ticker.
11	**return_10**	Price / Technical	Time-varying model input	10-session percentage change in the equity closing price.	Computed as close_t / close_(t-10) - 1 within ticker.
12	**log_close**	Price / Technical	Time-varying model input	Natural logarithm of the closing price.	Computed as ln(max(close, 1e-8)).
13	**hl_spread**	Price / Technical	Time-varying model input	Intraday high-low range relative to the close.	Computed as (high - low) / close.
14	**oc_spread**	Price / Technical	Time-varying model input	Open-to-close session return.	Computed as (close - open) / open.
15	**rsi_14**	Price / Technical	Time-varying model input	Fourteen-session Relative Strength Index.	Computed with the ta RSIIndicator using a 14-session window.
16	**macd**	Price / Technical	Time-varying model input	Moving Average Convergence Divergence line.	Computed with the ta MACD implementation using its standard fast and slow exponential-moving-average settings.

ISSN: 2473-4901
v12 n6752

17	**realized_vol_10**	Price / Technical	Time-varying model input	10-session realized volatility proxy based on daily returns.	Computed as the rolling 10-session standard deviation of return_1 within ticker.
18	**realized_vol_20**	Price / Technical	Time-varying model input	20-session realized volatility proxy based on daily returns.	Computed as the rolling 20-session standard deviation of return_1 within ticker.
19	**atr_14**	Price / Technical	Time-varying model input	Fourteen-session Average True Range.	Computed from high, low, and close using the ta AverageTrueRange indicator with window=14.
20	**volume_ma_20**	Price / Technical	Time-varying model input	Twenty-session moving average of trading volume.	Computed as the rolling 20-session mean of volume.
21	**volume_z_20**	Price / Technical	Time-varying model input	Twenty-session standardized volume deviation.	Computed as (volume - rolling_mean_20) / rolling_std_20.
22	**spy_close**	Benchmark / Market Context	Time-varying model input	Daily closing level of SPDR S&P 500 ETF Trust.	Downloaded through yfinance and merged to each equity by trading date.
23	**spy_return_1**	Benchmark / Market Context	Time-varying model input	1-session percentage return of SPDR S&P 500 ETF Trust.	Computed from the benchmark closing level as close_t / close_(t-1) - 1.
24	**spy_return_5**	Benchmark / Market Context	Time-varying model input	5-session percentage return of SPDR S&P 500 ETF Trust.	Computed from the benchmark closing level as close_t / close_(t-5) - 1.
25	**spy_return_10**	Benchmark / Market Context	Time-varying model input	10-session percentage return of SPDR S&P 500 ETF Trust.	Computed from the benchmark closing level as close_t / close_(t-10) - 1.

ISSN: 2473-4901
v12 n6752

26	**spy_rv20**	Benchmark / Market Context	Time-varying model input	Twenty-session realized-volatility proxy for SPDR S&P 500 ETF Trust.	Computed as the rolling 20-session standard deviation of the benchmark one-session return.
27	**qqq_close**	Benchmark / Market Context	Time-varying model input	Daily closing level of Invesco QQQ ETF.	Downloaded through yfinance and merged to each equity by trading date.
28	**qqq_return_1**	Benchmark / Market Context	Time-varying model input	1-session percentage return of Invesco QQQ ETF.	Computed from the benchmark closing level as close_t / close_(t-1) - 1.
29	**qqq_return_5**	Benchmark / Market Context	Time-varying model input	5-session percentage return of Invesco QQQ ETF.	Computed from the benchmark closing level as close_t / close_(t-5) - 1.
30	**qqq_return_10**	Benchmark / Market Context	Time-varying model input	10-session percentage return of Invesco QQQ ETF.	Computed from the benchmark closing level as close_t / close_(t-10) - 1.
31	**qqq_rv20**	Benchmark / Market Context	Time-varying model input	Twenty-session realized-volatility proxy for Invesco QQQ ETF.	Computed as the rolling 20-session standard deviation of the benchmark one-session return.
32	**iwm_close**	Benchmark / Market Context	Time-varying model input	Daily closing level of iShares Russell 2000 ETF.	Downloaded through yfinance and merged to each equity by trading date.
33	**iwm_return_1**	Benchmark / Market Context	Time-varying model input	1-session percentage return of iShares Russell 2000 ETF.	Computed from the benchmark closing level as close_t / close_(t-1) - 1.
34	**iwm_return_5**	Benchmark / Market Context	Time-varying model input	5-session percentage return of iShares Russell 2000 ETF.	Computed from the benchmark closing level as close_t / close_(t-5) - 1.

 ISSN: 2473-4901
v12 n6752

35	**iwm_return_10**	Benchmark / Market Context	Time-varying model input	10-session percentage return of iShares Russell 2000 ETF.	Computed from the benchmark closing level as close_t / close_(t-10) - 1.
36	**iwm_rv20**	Benchmark / Market Context	Time-varying model input	Twenty-session realized-volatility proxy for iShares Russell 2000 ETF.	Computed as the rolling 20-session standard deviation of the benchmark one-session return.
37	**xlk_close**	Benchmark / Market Context	Time-varying model input	Daily closing level of Technology Select Sector SPDR Fund.	Downloaded through yfinance and merged to each equity by trading date.
38	**xlk_return_1**	Benchmark / Market Context	Time-varying model input	1-session percentage return of Technology Select Sector SPDR Fund.	Computed from the benchmark closing level as close_t / close_(t-1) - 1.
39	**xlk_return_5**	Benchmark / Market Context	Time-varying model input	5-session percentage return of Technology Select Sector SPDR Fund.	Computed from the benchmark closing level as close_t / close_(t-5) - 1.
40	**xlk_return_10**	Benchmark / Market Context	Time-varying model input	10-session percentage return of Technology Select Sector SPDR Fund.	Computed from the benchmark closing level as close_t / close_(t-10) - 1.
41	**xlk_rv20**	Benchmark / Market Context	Time-varying model input	Twenty-session realized-volatility proxy for Technology Select Sector	Computed as the rolling 20-session standard deviation of the benchmark one-session return.

ISSN: 2473-4901
v12 n6752

				SPDR Fund.	
42	**xlf_close**	Benchmark / Market Context	Time-varying model input	Daily closing level of Financial Select Sector SPDR Fund.	Downloaded through yfinance and merged to each equity by trading date.
43	**xlf_return_1**	Benchmark / Market Context	Time-varying model input	1-session percentage return of Financial Select Sector SPDR Fund.	Computed from the benchmark closing level as close_t / close_(t-1) - 1.
44	**xlf_return_5**	Benchmark / Market Context	Time-varying model input	5-session percentage return of Financial Select Sector SPDR Fund.	Computed from the benchmark closing level as close_t / close_(t-5) - 1.
45	**xlf_return_10**	Benchmark / Market Context	Time-varying model input	10-session percentage return of Financial Select Sector SPDR Fund.	Computed from the benchmark closing level as close_t / close_(t-10) - 1.
46	**xlf_rv20**	Benchmark / Market Context	Time-varying model input	Twenty-session realized-volatility proxy for Financial Select Sector SPDR Fund.	Computed as the rolling 20-session standard deviation of the benchmark one-session return.
47	**xle_close**	Benchmark / Market Context	Time-varying model input	Daily closing level of Energy Select Sector SPDR Fund.	Downloaded through yfinance and merged to each equity by trading date.
48	**xle_return_1**	Benchmark / Market Context	Time-varying	1-session percentage return of	Computed from the benchmark closing level as close_t / close_(t-1) - 1.

 ISSN: 2473-4901
v12 n6752

			model input	Energy Select Sector SPDR Fund.	
49	**xle_return_5**	Benchmark / Market Context	Time-varying model input	5-session percentage return of Energy Select Sector SPDR Fund.	Computed from the benchmark closing level as close_t / close_(t-5) - 1.
50	**xle_return_10**	Benchmark / Market Context	Time-varying model input	10-session percentage return of Energy Select Sector SPDR Fund.	Computed from the benchmark closing level as close_t / close_(t-10) - 1.
51	**xle_rv20**	Benchmark / Market Context	Time-varying model input	Twenty-session realized-volatility proxy for Energy Select Sector SPDR Fund.	Computed as the rolling 20-session standard deviation of the benchmark one-session return.
52	**vix_close**	Benchmark / Market Context	Time-varying model input	Daily closing level of the CBOE Volatility Index.	Downloaded for ^VIX through yfinance and merged by date.
53	**vix_return_1**	Benchmark / Market Context	Time-varying model input	1-session percentage change in the VIX closing level.	Computed as vix_close_t / vix_close_(t-1) - 1.
54	**vix_return_5**	Benchmark / Market Context	Time-varying model input	5-session percentage change in the VIX closing level.	Computed as vix_close_t / vix_close_(t-5) - 1.
55	**vix_rv20**	Benchmark / Market Context	Time-varying model input	Twenty-session variability of daily VIX changes.	Computed as the rolling 20-session standard deviation of vix_return_1.

ISSN: 2473-4901
v12 n6752

56	**vix_z_20**	Benchmark / Market Context	Time-varying model input	Twenty-session z-score of the VIX closing level.	Computed as (vix_close - rolling_mean_20) / rolling_std_20.
57	**vix_spike_flag**	Benchmark / Market Context	Time-varying model input	Indicator of an unusually elevated VIX level.	Set to 1 when vix_z_20 >= 1.5; otherwise 0.
58	**days_since_vix_spike**	Benchmark / Market Context	Time-varying model input	Number of observations since the latest VIX spike.	Counted from the most recent vix_spike_flag=1 row and capped at 30; set to 30 before the first observed spike.
59	**vix_regime_high**	Benchmark / Market Context	Time-varying model input	High-volatility-market regime indicator.	Set to 1 when vix_z_20 >= 1.0; otherwise 0.
60	**excess_return_vs_spy**	Benchmark / Market Context	Time-varying model input	Equity one-session return relative to SPY.	Computed as return_1 - spy_return_1.
61	**rolling_beta_spy_20**	Benchmark / Market Context	Time-varying model input	Twenty-session rolling equity beta relative to SPY.	Computed as rolling_cov(return_1, spy_return_1) / rolling_var(spy_return_1).
62	**excess_return_vs_qqq**	Benchmark / Market Context	Time-varying model input	Equity one-session return relative to QQQ.	Computed as return_1 - qqq_return_1.
63	**fed_funds**	Macroeconomic	Time-varying model input	Federal funds rate.	Downloaded through fredapi from FRED FEDFUNDS; ordered by observation date, forward-filled, and backward-merged to market dates.
64	**cpi**	Macroeconomic	Time-varying model input	Consumer Price Index for All Urban Consumers.	Downloaded through fredapi from FRED CPIAUCSL; ordered by observation date, forward-filled, and backward-merged to market dates.
65	**unemployment**	Macroeconomic	Time-varying model input	Civilian unemployment rate.	Downloaded through fredapi from FRED UNRATE; ordered by observation date, forward-filled, and backward-merged to market dates.

ISSN: 2473-4901
v12 n6752

66	**teny_yield**	Macroeconomic	Time-varying model input	10-year U.S. Treasury yield.	Downloaded through fredapi from FRED GS10; ordered by observation date, forward-filled, and backward-merged to market dates.
67	**twoy_yield**	Macroeconomic	Time-varying model input	2-year U.S. Treasury yield.	Downloaded through fredapi from FRED GS2; ordered by observation date, forward-filled, and backward-merged to market dates.
68	**oil_wti**	Macroeconomic	Time-varying model input	West Texas Intermediate crude-oil price.	Downloaded through fredapi from FRED DCOILWTICO; ordered by observation date, forward-filled, and backward-merged to market dates.
69	**dollar_index_proxy**	Macroeconomic	Time-varying model input	Trade-weighted U.S. dollar index proxy.	Downloaded through fredapi from FRED DTWEXBGS; ordered by observation date, forward-filled, and backward-merged to market dates.
70	**fed_funds_chg_1**	Macroeconomic	Time-varying model input	1-observation percentage change in federal funds rate.	Computed with pct_change(1) on the FRED FEDFUNDS series before as-of alignment to market dates.
71	**fed_funds_chg_5**	Macroeconomic	Time-varying model input	5-observation percentage change in federal funds rate.	Computed with pct_change(5) on the FRED FEDFUNDS series before as-of alignment to market dates.
72	**cpi_chg_1**	Macroeconomic	Time-varying model input	1-observation percentage change in consumer price index for all urban consumers.	Computed with pct_change(1) on the FRED CPIAUCSL series before as-of alignment to market dates.
73	**cpi_chg_5**	Macroeconomic	Time-varying model input	5-observation percentage change in consumer price index for all urban	Computed with pct_change(5) on the FRED CPIAUCSL series before as-of alignment to market dates.

 ISSN: 2473-4901
v12 n6752

				consumers.	
74	**unemployment_chg_1**	Macroeconomic	Time-varying model input	1-observation percentage change in civilian unemployment rate.	Computed with pct_change(1) on the FRED UNRATE series before as-of alignment to market dates.
75	**unemployment_chg_5**	Macroeconomic	Time-varying model input	5-observation percentage change in civilian unemployment rate.	Computed with pct_change(5) on the FRED UNRATE series before as-of alignment to market dates.
76	**teny_yield_chg_1**	Macroeconomic	Time-varying model input	1-observation percentage change in 10-year u.s. treasury yield.	Computed with pct_change(1) on the FRED GS10 series before as-of alignment to market dates.
77	**teny_yield_chg_5**	Macroeconomic	Time-varying model input	5-observation percentage change in 10-year u.s. treasury yield.	Computed with pct_change(5) on the FRED GS10 series before as-of alignment to market dates.
78	**twoy_yield_chg_1**	Macroeconomic	Time-varying model input	1-observation percentage change in 2-year u.s. treasury yield.	Computed with pct_change(1) on the FRED GS2 series before as-of alignment to market dates.
79	**twoy_yield_chg_5**	Macroeconomic	Time-varying model input	5-observation percentage change in 2-year u.s. treasury yield.	Computed with pct_change(5) on the FRED GS2 series before as-of alignment to market dates.
80	**oil_wti_chg_1**	Macroeconomic	Time-varying	1-observation	Computed with pct_change(1) on the FRED DCOILWTICO

ISSN: 2473-4901
v12 n6752

			model input	percentage change in west texas intermediate crude-oil price.	series before as-of alignment to market dates.
81	**oil_wti_chg_5**	Macroeconomic	Time-varying model input	5-observation percentage change in west texas intermediate crude-oil price.	Computed with pct_change(5) on the FRED DCOILWTICO series before as-of alignment to market dates.
82	**dollar_index_proxy_chg_1**	Macroeconomic	Time-varying model input	1-observation percentage change in trade-weighted u.s. dollar index proxy.	Computed with pct_change(1) on the FRED DTWEXBGS series before as-of alignment to market dates.
83	**dollar_index_proxy_chg_5**	Macroeconomic	Time-varying model input	5-observation percentage change in trade-weighted u.s. dollar index proxy.	Computed with pct_change(5) on the FRED DTWEXBGS series before as-of alignment to market dates.
84	**revenue**	Point-in-Time Fundamental	Time-varying model input	Reported company revenue available at the prediction date.	Extracted from SEC companyfacts revenue concepts for 10-K/10-Q filings and amendments; backward-merged using filing availability date and forward-filled within ticker.
85	**net_income**	Point-in-Time Fundamental	Time-varying model input	Reported company net income available at the prediction date.	Extracted from SEC companyfacts NetIncomeLoss or ProfitLoss for 10-K/10-Q filings and amendments; backward-merged using filing availability date and forward-filled within ticker.
86	**operating_income**	Point-in-Time Fundamental	Time-varying model input	Reported operating income available at the	Extracted from SEC companyfacts OperatingIncomeLoss for 10-K/10-Q filings and amendments; backward-

ISSN: 2473-4901
v12 n6752

				prediction date.	merged using filing availability date and forward-filled within ticker.
87	**assets**	Point-in-Time Fundamental	Time-varying model input	Reported total assets available at the prediction date.	Extracted from SEC companyfacts Assets for 10-K/10-Q filings and amendments; backward-merged using filing availability date and forward-filled within ticker.
88	**liabilities**	Point-in-Time Fundamental	Time-varying model input	Reported total liabilities available at the prediction date.	Extracted from SEC companyfacts Liabilities for 10-K/10-Q filings and amendments; backward-merged using filing availability date and forward-filled within ticker.
89	**equity**	Point-in-Time Fundamental	Time-varying model input	Reported stockholders equity available at the prediction date.	Extracted from SEC companyfacts stockholders-equity concepts for 10-K/10-Q filings and amendments; backward-merged using filing availability date and forward-filled within ticker.
90	**eps_diluted**	Point-in-Time Fundamental	Time-varying model input	Reported diluted earnings per share available at the prediction date.	Extracted from SEC companyfacts EarningsPerShareDiluted for 10-K/10-Q filings and amendments; backward-merged using filing availability date and forward-filled within ticker.
91	**debt_to_equity_pti**	Point-in-Time Fundamental	Time-varying model input	Point-in-time liabilities-to-equity ratio.	Computed as liabilities / equity using the latest SEC filing available on or before the market date.
92	**return_on_equity_pti**	Point-in-Time Fundamental	Time-varying model input	Point-in-time return on equity.	Computed as net_income / equity using filing-availability alignment.
93	**operating_margin_pti**	Point-in-Time Fundamental	Time-varying model input	Point-in-time operating profit margin.	Computed as operating_income / revenue using filing-availability alignment.

ISSN: 2473-4901
v12 n6752

94	**profit_margin_pti**	Point-in-Time Fundamental	Time-varying model input	Point-in-time net profit margin.	Computed as net_income / revenue using filing-availability alignment.
95	**event_shock_score**	Event / Shock	Time-varying model input	Magnitude of the absolute one-session return relative to recent return variability.	Computed as abs(return_1) / rolling_std_20(return_1).
96	**volume_shock_score**	Event / Shock	Time-varying model input	Standardized abnormal trading volume around a possible event.	Computed as (volume - rolling_mean_20(volume)) / rolling_std_20(volume).
97	**gap_shock_score**	Event / Shock	Time-varying model input	Standardized magnitude of the overnight opening gap.	Gap return is (open - prior_close) / prior_close; the feature is abs(gap_return) / rolling_std_20(gap_return).
98	**range_shock_score**	Event / Shock	Time-varying model input	Standardized intraday-range intensity.	Range ratio is (high - low) / close; the feature is range_ratio / rolling_std_20(range_ratio).
99	**is_price_event_proxy**	Event / Shock	Time-varying model input	Binary indicator of an unusually large price/volume, gap/volume, or range/return combination.	Set from three threshold rules combining event_shock_score, volume_shock_score, gap_shock_score, and range_shock_score.

ISSN: 2473-4901
v12 n6752

100	**days_since_event_proxy**	Event / Shock	Time-varying model input	Number of observations since the most recent price-event proxy.	Counted from the latest is_price_event_proxy=1 row and capped at 30.
101	**event_cluster_10**	Event / Shock	Time-varying model input	Recent concentration of price-event proxies.	Computed as the rolling 10-session sum of is_price_event_proxy.
102	**post_event_decay**	Event / Shock	Time-varying model input	Exponentially decaying indicator of recency after a price event.	Computed as exp(-days_since_event_proxy / 5).
103	**is_event_regime**	Event / Shock	Time-varying model input	Composite market-event regime indicator.	Set to 1 for a current price event, a recent strong event shock, or event_cluster_10 >= 3; otherwise 0.
104	**news_count_recent**	News / Sentiment	Time-varying model input	Number of news items for the ticker on the normalized publication date.	Headlines and available bodies were grouped by ticker-date after source normalization.
105	**news_sentiment_recent**	News / Sentiment	Time-varying model input	Mean signed FinBERT sentiment score for ticker-date news.	ProsusAI/FinBERT probabilities were signed positive, negative, or zero for neutral and averaged by ticker-date.
106	**positive_news_count_recent**	News / Sentiment	Time-varying model input	Count of positively classified ticker-date news items.	Headline sentiment > 0.15 was counted as positive.
107	**negative_news_count_recent**	News / Sentiment	Time-varying	Count of negatively	Headline sentiment < -0.15 was counted as negative.

ISSN: 2473-4901
v12 n6752

			model input	classified ticker-date news items.	
108	**news_sentiment_std_recent**	News / Sentiment	Time-varying model input	Within-day dispersion of signed news sentiment.	Computed as the standard deviation of ticker-date FinBERT scores; single-item days were filled with 0.
109	**news_count_recent_3d**	News / Sentiment	Time-varying model input	Rolling 3-session total number of ticker news items.	Computed as the within-ticker rolling 3-session sum of news_count_recent.
110	**news_count_recent_5d**	News / Sentiment	Time-varying model input	Rolling 5-session total number of ticker news items.	Computed as the within-ticker rolling 5-session sum of news_count_recent.
111	**news_sentiment_recent_3d**	News / Sentiment	Time-varying model input	Rolling 3-session mean of daily ticker news sentiment.	Computed as the within-ticker rolling 3-session mean of news_sentiment_recent.
112	**news_sentiment_recent_5d**	News / Sentiment	Time-varying model input	Rolling 5-session mean of daily ticker news sentiment.	Computed as the within-ticker rolling 5-session mean of news_sentiment_recent.
113	**news_burst_z_20**	News / Sentiment	Time-varying model input	Twenty-session z-score of daily news volume.	Computed from news_count_recent using a rolling 20-session mean and standard deviation with at least five observations.
114	**news_sentiment_z_20**	News / Sentiment	Time-varying model input	Twenty-session z-score of daily mean news sentiment.	Computed from news_sentiment_recent using a rolling 20-session mean and standard deviation with at least five observations.
115	**news_sentiment_abs**	News / Sentiment	Time-varying	Absolute magnitude of daily	Computed as abs(news_sentiment_recent).

ISSN: 2473-4901
v12 n6752

			model input	mean news sentiment.	
116	**news_polarity_imbalance**	News / Sentiment	Time-varying model input	Normalized balance between positive and negative news counts.	Computed as (positive_count - negative_count) / news_count, with zero when the denominator is unavailable.
117	**negative_news_flag**	News / Sentiment	Time-varying model input	Indicator that negative items outnumber positive items on a news-active date.	Set to 1 when negative count > positive count and news_count_recent > 0.
118	**positive_news_flag**	News / Sentiment	Time-varying model input	Indicator that positive items outnumber negative items on a news-active date.	Set to 1 when positive count > negative count and news_count_recent > 0.
119	**extreme_negative_news**	News / Sentiment	Time-varying model input	Indicator of strongly negative mean ticker-date sentiment.	Set to 1 when news_sentiment_recent < -0.35.
120	**extreme_positive_news**	News / Sentiment	Time-varying model input	Indicator of strongly positive mean ticker-date sentiment.	Set to 1 when news_sentiment_recent > 0.35.
121	**news_event_regime**	News / Sentiment	Time-varying model input	Composite indicator of an unusually active or	Set to 1 for news_burst_z_20 >= 1.5, positive/negative dominance, or an extreme sentiment flag.

 ISSN: 2473-4901
v12 n6752

				polarized news state.	
122	**news_event_regime_3d**	News / Sentiment	Time-varying model input	Three-session persistence indicator for the news-event regime.	Computed as the rolling 3-session maximum of news_event_regime within ticker.
123	**news_sentiment_x_return_1**	News / Sentiment	Time-varying model input	Interaction between daily news sentiment and the equity one-session return.	Computed as news_sentiment_recent * return_1.
124	**news_sentiment_x_volume_z_20**	News / Sentiment	Time-varying model input	Interaction between daily news sentiment and abnormal volume.	Computed as news_sentiment_recent * volume_z_20.
125	**news_sentiment_x_vix_return_1**	News / Sentiment	Time-varying model input	Interaction between daily news sentiment and the one-session VIX return.	Computed as news_sentiment_recent * vix_return_1.
126	**news_uncertainty_x_volatility**	News / Sentiment	Time-varying model input	Interaction between daily sentiment dispersion and realized equity volatility.	Computed as news_sentiment_std_recent * realized_vol_20.
127	**sentiment_shock_20**	News / Sentiment	Time-varying model input	Unexpected daily news sentiment relative to	Computed as news_sentiment_recent - rolling_mean_20(news_sentiment_recent).

ISSN: 2473-4901
v12 n6752

				its recent baseline.	
128	**negative_sentiment_shock_20**	News / Sentiment	Time-varying model input	Downside-only component of the sentiment shock.	Computed by clipping sentiment_shock_20 at an upper bound of 0.
129	**sentiment_intensity**	News / Sentiment	Time-varying model input	Magnitude of a sentiment surprise weighted by news volume.	Computed as abs(sentiment_shock_20) * news_count_recent.
130	**days_since_event**	News / Sentiment	Time-varying model input	Number of observations since the most recent news-event-regime activation.	Causal within-ticker counter; 0 on event rows and a large sentinel before the first event.
131	**decayed_sentiment_shock**	News / Sentiment	Time-varying model input	Recency-weighted sentiment surprise.	Computed as sentiment_shock_20 * exp(-days_since_event / 3).
132	**stress_x_sentiment_shock**	News / Sentiment	Time-varying model input	Interaction between sentiment surprise and market stress.	Computed as sentiment_shock_20 * vix_z_20.
133	**nominal_regime_flag**	Regime / Stress	Time-varying model input	Indicator of low-stress market conditions.	Set to 1 when VIX, news-burst, connectedness, and peer-shock measures are all below the specified nominal thresholds.
134	**stress_regime_flag**	Regime / Stress	Time-varying model input	Indicator of a stressed or event-driven market condition.	Set to 1 when VIX, news-burst, connectedness, peer-shock, extreme-negative-news, or event-regime conditions exceed the stress rules.

ISSN: 2473-4901
v12 n6752

135	**mixed_regime_flag**	Regime / Stress	Time-varying model input	Indicator of observations that are neither nominal nor stress.	Set to 1 when nominal_regime_flag=0 and stress_regime_flag=0.
136	**cross_sectional_rank_return_1**	Graph / Relational	Time-varying model input	Ticker percentile rank of the one-session return on the same date.	Computed by date across the modeled equity universe; missing ranks were filled with 0.5.
137	**cross_sectional_rank_vol_20**	Graph / Relational	Time-varying model input	Ticker percentile rank of 20-session realized volatility on the same date.	Computed by date across the modeled equity universe; missing ranks were filled with 0.5.
138	**cross_sectional_rank_rsi_14**	Graph / Relational	Time-varying model input	Ticker percentile rank of RSI on the same date.	Computed by date across the modeled equity universe; missing ranks were filled with 0.5.
139	**graph_avg_corr_20**	Graph / Relational	Time-varying model input	Mean 20-session return correlation between a ticker and all other modeled tickers.	Calculated from the rolling 20-session cross-ticker return correlation matrix, excluding self-correlation.
140	**graph_avg_abs_corr_20**	Graph / Relational	Time-varying model input	Mean absolute 20-session return correlation with peer tickers.	Mean of absolute peer correlations from the rolling 20-session correlation matrix.
141	**graph_corr_dispersion_20**	Graph / Relational	Time-varying model input	Dispersion of 20-session peer-return	Standard deviation of the ticker-to-peer correlations in the rolling 20-session matrix.

ISSN: 2473-4901
v12 n6752

				correlations.	
142	**graph_connectedness_20**	Graph / Relational	Time-varying model input	Fraction of peer tickers with materially strong return correlation.	Computed as the share of peer correlations with absolute value >= 0.30 in the rolling 20-session window.
143	**peer_mean_return_1**	Graph / Relational	Time-varying model input	Mean one-session return of all other modeled tickers.	Cross-sectional leave-one-out mean computed by date.
144	**peer_mean_return_5**	Graph / Relational	Time-varying model input	Mean five-session return of all other modeled tickers.	Cross-sectional leave-one-out mean computed by date.
145	**peer_mean_vol_20**	Graph / Relational	Time-varying model input	Mean 20-session realized volatility of all other modeled tickers.	Cross-sectional leave-one-out mean computed by date.
146	**peer_mean_volume_z_20**	Graph / Relational	Time-varying model input	Mean 20-session volume z-score of all other modeled tickers.	Cross-sectional leave-one-out mean computed by date.
147	**peer_relative_return_1**	Graph / Relational	Time-varying model input	Ticker one-session return relative to the peer mean.	Computed as return_1 - peer_mean_return_1.
148	**peer_relative_return_5**	Graph / Relational	Time-varying model input	Ticker five-session return relative to the peer mean.	Computed as return_5 - peer_mean_return_5.

 ISSN: 2473-4901
v12 n6752

149	**peer_relative_vol_20**	Graph / Relational	Time-varying model input	Ticker realized volatility relative to the peer mean.	Computed as realized_vol_20 - peer_mean_vol_20.
150	**peer_mean_rsi_14**	Graph / Relational	Time-varying model input	Mean RSI of all other modeled tickers.	Cross-sectional leave-one-out mean computed by date.
151	**peer_relative_rsi_14**	Graph / Relational	Time-varying model input	Ticker RSI relative to the peer mean.	Computed as rsi_14 - peer_mean_rsi_14.
152	**graph_peer_negative_breadth**	Graph / Relational	Time-varying model input	Share of peer tickers experiencing a one-session return <= -2%.	Computed as a leave-one-out cross-sectional peer fraction for each date.
153	**graph_peer_positive_breadth**	Graph / Relational	Time-varying model input	Share of peer tickers experiencing a one-session return >= 2%.	Computed as a leave-one-out cross-sectional peer fraction for each date.
154	**graph_peer_event_intensity**	Graph / Relational	Time-varying model input	Share of peers exhibiting an event-like return or volume shock.	Peer event-like flag equals 1 when abs(return_1) >= 2.5% or abs(volume_z_20) >= 2.0; aggregated as a leave-one-out fraction.
155	**graph_peer_shock_breadth**	Graph / Relational	Time-varying model input	Net positive-versus-negative peer shock breadth.	Computed as graph_peer_positive_breadth - graph_peer_negative_breadth.
156	**graph_sector_spillover**	Graph / Relational	Time-varying model input	Sector-context return spillover proxy.	Computed as 0.5 * xlf_return_1 + 0.5 * xlk_return_1.

ISSN: 2473-4901
v12 n6752

157	**graph_market_spillover**	Graph / Relational	Time-varying model input	Broad-market return spillover proxy.	Computed as the mean of spy_return_1, qqq_return_1, and iwm_return_1.
158	**graph_risk_off_pressure**	Graph / Relational	Time-varying model input	Interaction of negative peer breadth with rising market stress.	Computed as negative_peer_breadth * vix_return_1 * (1 + max(vix_z_20, 0)).
159	**graph_connectedness_z_20**	Graph / Relational	Time-varying model input	Within-ticker rolling z-score of graph connectedness.	Computed from graph_connectedness_20 using a 20-session rolling mean and standard deviation with at least five observations.
160	**graph_spillover_z_20**	Graph / Relational	Time-varying model input	Within-ticker rolling z-score of broad-market spillover.	Computed from graph_market_spillover using a 20-session rolling mean and standard deviation with at least five observations.
161	**trade_threshold_5**	Target / Decision Context	Target / decision context	Volatility-adaptive minimum absolute return required for a directional T+5 trade label.	Computed as max(0.005, 0.50 * realized_vol_20 * sqrt(5)).
162	**target_trade_class_5**	Target / Decision Context	Target / decision context	Three-class T+5 directional outcome label: short (-1), no-trade/noise (0), or long (+1).	Assigned from future_return_5 relative to +/-trade_threshold_5 during supervised target construction; retained as historical target context in the frozen dataset schema.
163	**target_edge_return_5**	Target / Decision Context	Forecast target	Signed T+5 return beyond the	For long outcomes: future_return_5 - threshold; for short outcomes:

ISSN: 2473-4901
v12 n6752

				adaptive no-trade threshold.	future_return_5 + threshold; otherwise 0. This is the TFT forecasting target.
164	**sample_weight**	Model Support	Training-weight field	Observation-level training weight emphasizing event, price-shock, news-event, and directional-trade examples.	Initialized at 1.0; raised to 2.5 for event-regime rows and to at least 2.0 for price-event, news-event, or directional-class rows; optional ticker balancing was bounded to 0.75-4.0.
165	**relative_time_idx**	Identifier / Time	Generated model-support field	Relative position of each observation within an encoder-decoder sample.	Generated by PyTorch Forecasting because add_relative_time_idx=True; expresses sequence position relative to the forecast origin.
166	**encoder_length**	Identifier / Time	Generated model-support field	Historical sequence-length context supplied with each sample.	Generated by PyTorch Forecasting because add_encoder_length=True; the study used a fixed 48-step encoder.
167	**target_edge_return_5_center**	Model Support	Generated model-support field	Ticker-specific centering parameter for the forecasting target.	Generated by the GroupNormalizer because add_target_scales=True; estimated from the applicable ticker group.
168	**target_edge_return_5_scale**	Model Support	Generated model-support field	Ticker-specific scaling parameter for the forecasting target.	Generated by the GroupNormalizer because add_target_scales=True; used to normalize target magnitude by ticker.

Note. The date field is retained for temporal alignment and auditability but is not supplied as a numeric TFT predictor. The target, observation weight, and generated normalization fields are documented to make the complete frozen modeling schema reproducible. Macroeconomic series were as-of merged by

ISSN: 2473-4901
v12 n6752

observation date; point-in-time fundamentals were aligned by SEC filing-availability date; news variables were derived from ticker-date FinBERT aggregation; and graph variables were constructed from same-date cross-sectional and trailing-window market information.