Beyond Stylised Facts: Hybrid Detection of Data Anomalies in the Zimbabwe Stock Market Returns

Godfrey Mtunzi[*♣], **Caston Sigauke**[♠], **Edson Mbedzi**[♣], and **Noble Malunguza**[♣]

[♣]National University of Science and Technology (NUST), Zimbabwe
[♠]University of Venda, South Africa

Submitted: November 4, 2025 • Accepted: March 8, 2026

ABSTRACT: Stylised facts are universal, but anomalies may be market-specific. Therefore, this study examines anomalies on the Zimbabwe Stock Exchange (ZSE) to reveal structural or behavioural features unique to frontier markets. Stylised facts are widely observed across global financial markets, regardless of differences in liquidity, regulation, or market structure, and might fail to uncover deeper, market-specific irregularities that may reflect the unique structural and behavioural dynamics of a frontier market. This unexpected similarity suggests that stylised facts alone cannot fully explain frontier markets' return behaviour, exemplified by the ZSE, motivating a deeper search for additional empirical patterns and anomalies. We develop a hybrid anomaly-detection pipeline integrating: Principal Component Analysis (PCA) for denoising and dimensionality reduction, Isolation Forest for point-level deviations, and Recurrent Neural Network (RNN)-based models for contextual anomalies. Hidden Markov Models for collective anomalies and regime transitions. The framework operates sequentially rather than as isolated models, ensuring coherent anomaly classification. The detected anomalies align with major macroeconomic and policy events in Zimbabwe. The hybrid pipeline provides broader anomaly coverage than individual models. The framework shows clearer detection of structural

---

[*] Corresponding Author. E-mail: godfeymtunzi@gmail.com

DOI: 10.33119/ERFIN.2026.11.1.2

and temporal anomalies which are not evident in standard stylised facts analysis, although recall remains low across models. Findings reveal structural and behavioural patterns embedded in ZSE returns. Results support applications in market surveillance and risk assessment.

JEL classification: C22, C40, C52.

Keywords: Behavioural Dynamics, Data Anomalies, Frontier Markets, Stylised Facts.

1 Introduction

When modelling financial returns, one of the main issues is to do with the quality and integrity of the data (Arjunan, 2022). Therefore, ensuring data accuracy, completeness, consistency, and validity through robust data governance, data cleansing tools, data validation, and data quality monitoring is essential for reliable return modelling. However, even when data quality is assured, financial return series may still exhibit irregularities that cannot be explained by data errors alone, necessitating an examination of deeper structural and behavioural anomalies within the market. According to Crépey et al. (2022), irregularities in financial time series data normally unfold in the form of data anomalies that induce non-conformity to the stylised facts of financial returns, miscalibration of the models used to develop financial products, and the quantification and management of risk, among other challenges. Xia et al. (2022) emphasised the importance of anomaly detection, as it is part of data preprocessing that finds and removes anomalous data and improves data quality and integrity, thus improving models' performance.

Anomalies often persist even after standard preprocessing (Saurabh and Verma, 2023). These deviations require deeper investigation beyond routine model checks. Unlike market anomalies, which reflect deviations from the efficient market hypothesis (EMH), data anomalies relate to irregularities in the dataset itself (Murekachiro, 2025). According to Bouattour and Martinez (2019), data anomalies concern dataset quality, whereas market anomalies concern price behaviour. These anomalies may signal inefficiencies or behavioural biases in markets and can be exploited by investors for profit. According to Foorthuis (2021), the most pronounced data anomalies in financial data are structural breaks, outliers, missing values, duplicates, timestamp errors, and incomplete data. Data quality issues such as incompleteness, short horizons, and narrow stock universes can lead to three well-documented sources of biases, i.e., look-ahead bias, survivorship bias, and data-snooping bias (Li et al., 2025). Look-ahead bias occurs when unavailable future information is used in modelling, while survivorship bias arises when only surviving firms are included. Bailey et al. (2016) showed that repeated experimentation on the same dataset leads to data-snooping bias and overfitting.

Frontier markets such as Zimbabwe exhibit structural breaks and regime shifts that require flexible anomaly-detection tools (Pinto and Sobreiro, 2024). A hybrid approach is needed to capture point, contextual, and regime-driven deviations. The study focuses on uncovering patterns not visible through stylised-fact diagnostics alone. Dimensionality reduction is often required in financial anomaly detection. The curse of dimensionality introduces substantial limitations to the manipulation, interpretation, and analysis of high-dimensional datasets, thereby complicating the detection of behavioural patterns within these expanded feature spaces (Pinto and Sobreiro, 2024).

This study addresses this gap by examining anomalies through a regime-aware and temporally sensitive lens. The analysis highlights structural and behavioural features unique to the ZSE (Bakumenko and Elragal, 2022). In financial modelling, the challenge is not only in identifying if the anomalies have already occurred but also lies in predicting them (Shabir and Pearl, 2024). A key challenge is distinguishing structural anomalies from noise in volatile frontier-market environments. Temporal models help anticipate emerging deviations in return dynamics. Financial data from different markets come in different forms, structures, and timestamps, making anomaly detection more complex. Not every anomaly reflects a data error; distinguishing genuine errors from meaningful deviations is challenging, and even small errors can undermine data reliability (Bakumenko and Elragal, 2022). The research question focuses on how detected anomalies relate to Zimbabwe's market dynamics, and builds a sequential hybrid anomaly-detection framework. The framework integrates PCA, Isolation Forest, GRU, and HMMs. Each model targets a different anomaly type: point, contextual, or collective. The approach is tailored to the characteristics of frontier-market return behaviour.

The remaining part of this study is outlined as follows: Section 2 reviews the literature and conceptual foundations. Section 3 presents the data and methodological framework, Section 4 reports the empirical findings, Section 5 discusses implications and limitations, while Section 6 concludes the study.

2 Literature review

Anomaly detection in time series has become a central topic in computational finance, especially for modelling financial returns. Financial returns are inherently discrete, yet often treated as continuous for analytical convenience (Campbell et al., 1998). Under this continuous-time treatment, identifying deviations from normal return behaviour becomes essential and detecting such departures forms a core methodological concern in anomaly identification. Anomalies usually arise from statistical variations and strong outliers, and their typology is foundational to AD (Foorthuis, 2021). More often, true anomalies are data variables that deviate from the underlying data-generating process, and they are commonly classified into point, contextual, and collective anomalies. Other classifications of anomalies in finance include additive anomalies (short-lived), transitional anomalies (gradual decay), level-shift anomalies (structural jumps), and innovational anomalies (trend or seasonal disturbances) (Chandola et al., 2009). According to Foorthuis (2021), these typologies extend further to between-sequence anomalies such as shift, amplitude, and shape outliers (Foorthuis, 2018). Apart from these, regression-based anomalies arise when modelling financial data using regression frameworks. According to Blázquez-García et al. (2021), anomaly detection involves identifying data points or patterns that deviate significantly from expected behaviour. The deviations are typical in manufacturing, healthcare, and finance, where time series data is prevalent.

Classical statistical models such as ARIMA and threshold rules struggle with nonlinear and regime-dependent structures in financial returns (Tsay, 2005; Zhou, 2025; Gupta et al., 2014). According to Tsay (2005), Fischer and Krauss (2018) and Salehi and Rashidi (2018), these limitations motivate the shift toward more flexible anomaly-detection approaches. Unsupervised methods such as PCA and clustering offer dimensionality reduction but fail to capture nonlinear or temporal dependencies (Shyu et al., 2003; Gupta et al., 2014; Crépey et al., 2022; Salehi and Rashidi, 2018; Spoor et al., 2025). These weaknesses highlight the need for models that incorporate both temporal structure and nonlinearity.

Hidden Markov Models (HMMs) offer a solution that addresses this phenomenon (Oelschläger et al., 2024). A study by Li et al. (2017) extended HMMs to multivariate time series by transforming them into univariate representations and integrating fuzzy clustering techniques. The results outperformed conventional statistical models in identifying collective anomalies, reinforcing the model's relevance for regime-sensitive financial applications. Most hybrid AD studies focus on developed markets or industrial systems, leaving frontier-market financial returns underexplored (Ahmed et al., 2016; Chalapathy and Chawla, 2019). Few studies integrate regime-awareness, temporal modelling, and dimensionality reduction within a single framework.

Temporal anomaly detection methods aim to capture deviations that unfold over time, including contextual and sequence-dependent anomalies (Hundman et al., 2018; Malhotra et al., 2015). Regime-aware anomaly detection methods identify structural shifts and latent state transitions in financial time series (Hamilton, 1989; Ang and Bekaert, 2002). Recent advances in AD studies have proposed new and hybrid methods, including machine learning, deep learning and their combinations, even with traditional statistical techniques. However, many hybrid AD studies overemphasise point anomalies while under-addressing contextual and collective anomalies (Gupta et al., 2014; Campos et al., 2016; Chandola et al., 2009). According to Malhotra et al. (2015), temporal dependencies are critical for both collective and contextual outlier detection, as they may appear statistically normal or abnormal, but when analysed through a regime-specific lens, more information beyond stylised facts can be revealed. In financial markets, bear, bull and transitional regimes are inherent; therefore, there is a need to adapt methodologies that capture these temporal structures (Hamilton, 1989; Ang and Bekaert, 2002). According to Ang and Bekaert (2002), regime-switching models, such as Markov-Switching models in financial time series, are essential as they capture temporal structures like bull, bear, and transitional phases. Nonlinear dimensionality-reduction methods such as Kernel PCA and PCA-NN improve anomaly detection but remain underused in financial return modelling (Liang et al., 2023; Sakurada and Yairi, 2014; Aggarwal, 2015). Prediction-based contextual anomaly detection methods such as CAD offer improved recall and robustness across domains (Choi and Zaiane, 2015).

CAD demonstrated a significant improvement in recall from 7% to 33% compared to traditional statistical techniques such as kNN and Random Walk baselines, without compromising precision (Malhotra et al., 2015; Hundman et al., 2018). According to Choi and Zaiane (2015), if anomalies are correctly identified, most data mining problems in different domains become easier to solve. This task has been difficult as the majority of techniques do not capture temporal data dependencies or latent factors of the data under consideration (Gupta et al., 2014). According to Zhao et al. (2019) and Malhotra et al. (2015), contextually, deep learning models, particularly GRU and LSTM architectures, have shown promise in industrial and control systems. Tang et al. (2023) proposed a GRU-based interpretable multivariate AD model for industrial safety, balancing performance and interpretability. This is promising when applied to complex financial market data to capture facts beyond stylised facts.

Density-based clustering methods such as Density-Based Spatial Clustering of Applications with Noise (DBSCAN) can detect structural irregularities but lack temporal sensitivity (Ester et al., 1996; Shiraj et al., 2024; Baragona and Battaglia, 2007).

For multivariate datasets, as most financial problems involve multi-column datasets, Adams et al. (2019) proposed a strategy that detects multivariate points and collective anomalies by modelling cross-sectional dependence structures in financial time series. Prajesha and Veni (2023) proposed IF-RSPCA, a hybrid of Isolation Forest and PCA with Robust Scaling, for predicting anxiety disorders from time series data. Their results showed superior performance over standard Isolation Forest, highlighting the value of combining dimensionality reduction with robust anomaly scoring (Gao et al., 2019). Despite progress in hybrid anomaly detection, frontier-market applications remain scarce (Chalapathy and Chawla, 2019). Existing studies rarely integrate temporal, cross-sectional, and regime-sensitive components into a unified framework. Most hybrid AD applications focus on developed markets, industrial systems, or non-financial domains, with limited attention to frontier-market equity returns. Existing frameworks seldom combine: (i) dimensionality reduction for high-dimensional financial panels, (ii) temporal models for contextual anomalies, and (iii) regime-sensitive models for collective anomalies and state shifts. Differently, this study adopts a hybrid, regime-aware anomaly detection framework tailored to the Zimbabwe Stock Exchange.

3 Methodology

3.1 Data

The Zimbabwe Stock Exchange (ZSE) All Share Index (ASI) and its listed constituents' daily stock prices data from *Investing.com* is used. Although this vendor is not an official exchange source, it provides consistent historical coverage for frontier markets where official archives are often incomplete and costly. Daily stock prices are cross-validated with ZSE bulletins and press releases, and vendor triangulation is used to reduce source-specific errors. All discrepancies are documented and retain only consistent series. The primary sample spans March 2022 to present, April 2025, corresponding to a period of post-COVID monetary instability and several policy-driven shocks.

There are earlier crises in Zimbabwe; however, they are excluded due to incompatible monetary regimes. Including earlier crises would mix fundamentally different monetary systems, distorting volatility structure and compromising comparability across regimes. April 2025 was the most recent point at which complete and internally consistent multi-asset data were available, and although later data has since been released, the modelling framework was calibrated on the dataset available at the time; extending the sample retrospectively would introduce missingness, mix structurally different volatility regimes, and require full re-estimation of all models, thresholds, and regime classifications, breaking comparability with the validated pipeline.

The main quantity of interest is the vector of daily log returns for the ZSE ASI and its constituent stocks. Let the universe of listed securities be denoted by $\mathcal{U}$. Missing data are handled conservatively by treating missingness as Missing At Random (MAR) due to vendor omissions. Imputation is also avoided to prevent artificial smoothing and asset-day pairs, not full rows are dropped. Robustness is tested under Last Case Forward Observation (LCFO) where the last observed value for a unit is a reasonable proxy for the missing value. Also missingness is treated as if the underlying process is stable after the last observation and listwise deletion where any observation with any missing value in the variables used in the model. The price and return vectors are defined as

$$\mathbf{P}_t = \begin{bmatrix} P_t^{(1)} \\ P_t^{(2)} \\ \vdots \\ P_t^{(n)} \end{bmatrix}, \qquad \mathbf{r}_t = \begin{bmatrix} r_t^{(1)} \\ r_t^{(2)} \\ \vdots \\ r_t^{(n)} \end{bmatrix}, \tag{1}$$

$$r_t^{(i)} = \log\left(\frac{P_t^{(i)}}{P_{t-1}^{(i)}}\right), \qquad \mathbf{r}_t = \log(\mathbf{P}_t) - \log(\mathbf{P}_{t-1}). \tag{2}$$

where $\mathbf{P}_t$ is the price vector and $\mathbf{r}_t$ is the corresponding log-return vector at time t.

3.2 Feature importance analysis

A set of candidate explanatory features $\{X_i\}_{i=1}^{p}$ is constructed to capture cross-sectional and temporal variation in the ASI. X_{it} is defined as feature i at time t, y_t as ASI return at time t. All features are aligned to a common timestamp index. Three complementary importance metrics are used: correlation, PCA loadings, and Random Forest impurity reduction. We combine metrics by equal weighting to avoid structural bias, and each captures a different relevance dimension, i.e., correlation for linear co-movement, PCA loadings for variance structure, and Random Forest for non-linear predictive power. Equal weighting prevents over-reliance on any single structural assumption and aligns with ensemble-style feature scoring commonly used in high-dimensional financial datasets. Empirical studies show that no single feature-importance metric is stable in high-dimensional or correlated financial data (Chandrashekar and Sahin, 2014). Consistent with evidence from Gao et al. (2019) and Liang et al. (2023), hybrid scoring improves robustness and ranking stability relative to single-metric approaches. This aligns with a simple internal ablation check conducted in this study, which showed that the aggregated score produced more stable feature rankings than any individual metric.

In what follows, $i = 1, \dots, p$ indexes features, $t = 1, \dots, n$ indexes time, X_{it} denotes the value of feature i at time t, $\bar{X}_i$ is the sample mean of feature i, and $\bar{y}$ is the sample mean of ASI returns y_t. These definitions apply throughout all the equations.

3.2.1 Correlation analysis

$$r_i = \frac{\sum_{t=1}^{n}(X_{it} - \bar{X}_i)(y_t - \bar{y})}{\sqrt{\sum_{t=1}^{n}(X_{it} - \bar{X}_i)^2}\sqrt{\sum_{t=1}^{n}(y_t - \bar{y})^2}}, \qquad I_i^{\text{Corr}} = |r_i|. \tag{3}$$

The correlation score I_i^{Corr} measures the linear co-movement between feature i and the ASI. In the feature importance analysis, this score is computed for all p features and later normalised for aggregation.

3.2.2 PCA for importance analysis

$$\mathbf{X} \in \mathbb{R}^{n \times p}, \qquad \mathbf{\Sigma} = \frac{1}{n-1}\mathbf{X}^{\top}\mathbf{X}, \qquad \mathbf{\Sigma} = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^{\top}, \tag{4}$$

where $\mathbf{X}$ denotes the matrix of standardised features, $\mathbf{\Sigma}$ denotes the sample covariance matrix, $\mathbf{V}$ contains eigenvectors (PCA loadings), and $\mathbf{\Lambda}$ contains eigenvalues. Feature importance is computed as the absolute loading of each feature on the leading principal components, normalised to $[0, 1]$.

3.2.3 Aggregation and ranking

$$\tilde{I}_i^{(m)} = \frac{I_i^{(m)} - \min_j I_j^{(m)}}{\max_j I_j^{(m)} - \min_j I_j^{(m)}}, \qquad I_i^{\mathrm{Avg}} = \frac{1}{3}\left(\tilde{I}_i^{\mathrm{Corr}} + \tilde{I}_i^{\mathrm{PCA}} + \tilde{I}_i^{\mathrm{RF}}\right). \tag{5}$$

The normalisation step rescales each metric $m \in \{\mathrm{Corr}, \mathrm{PCA}, \mathrm{RF}\}$ to a common $[0, 1]$ range. The aggregated score I_i^{Avg} is then used to rank features. This procedure stabilises rankings across structurally different importance metrics.

3.3 Principal component analysis

PCA is used both for feature-importance scoring and as a baseline anomaly detector. PCA loadings are used only for ranking features, PCA z-scores are used as a baseline anomaly detector, and no PCA-derived components are reused across stages. A z-score baseline is defined as:

$$z_t = \left| \frac{r_t - \mu_r}{\sigma_r} \right|, \tag{6}$$

where r_t is the ASI return at time t, μ_r denotes the sample mean, and σ_r denotes the sample standard deviation. An anomaly is flagged when z_t exceeds a high-quantile threshold (95th–99th percentile).

3.4 Isolation forest (iForest)

Isolation Forest isolates observations by recursively partitioning the feature space. Thresholds are selected using multiple rules: 95–99 percentile cut-offs, elbow point on sorted anomaly scores, empirical tail-index thresholds, and combined thresholds to reduce sensitivity. The anomaly score is defined as

$$s(\mathbf{x}, \psi) = 2^{-\frac{E[h(\mathbf{x})]}{c(\psi)}}, \tag{7}$$

where $\mathbf{x}$ denotes the observation vector, $h(\mathbf{x})$ denotes the path length required to isolate $\mathbf{x}$, $E[h(\mathbf{x})]$ is its expected value, and $c(\psi)$ is a normalising constant depending on subsample size ψ. Higher $s(\mathbf{x})$ indicates stronger anomaly likelihood. An observation is classified as anomalous when $s(\mathbf{x})$ exceeds the combined threshold.

3.5 Gated recurrent unit (GRU) model

We implement a gated recurrent unit (GRU) sequence-to-sequence architecture for contextual anomaly detection. The network comprises two GRU layers with 32 units each and a dropout rate of 0.2. Training uses the Adam optimiser with a learning rate of 0.001, batch size of 32, and a 20-day input window. An 80/20 train-validation split is used, with early stopping (patience = 10). The encoder–decoder architecture is given by

$$\mathbf{h}_t = f_\theta(\mathbf{x}_t, \mathbf{h}_{t-1}), \qquad \mathbf{z} = \mathbf{h}_T, \tag{8}$$

$$\hat{\mathbf{x}}_t = g_\phi(\mathbf{z}, \hat{\mathbf{x}}_{t-1}), \tag{9}$$

$$A(\mathbf{X}) = \frac{1}{T}\sum_{t=1}^{T} \|\mathbf{x}_t - \hat{\mathbf{x}}_t\|^2, \tag{10}$$

where $\mathbf{x}_t$ denotes the observed return vector, $\hat{\mathbf{x}}_t$ denotes the GRU prediction, and $\mathbf{h}_t$ denotes the hidden state. The anomaly score is the reconstruction error $A(\mathbf{X})$. A point is flagged as anomalous when its error exceeds the 99^{th} percentile of training-set errors.

3.6 Hidden Markov models (HMMs)

We employ hidden Markov models (HMMs) to capture latent regime dynamics in asset returns. The number of regimes is selected using the Bayesian Information Criterion (BIC) evaluated over $K \in \{2, 3, 4, 5\}$, with the optimal specification yielding $K = 3$, corresponding to low-, medium-, and high-volatility states. We adopt multivariate Gaussian emission distributions to accommodate cross-asset co-movement. The model is specified as

$$\mathbb{P}(s_{t+1} \mid s_t) = A_{s_t, s_{t+1}}. \tag{11}$$

$$\mathbf{x}_t \mid s_t = k \sim \mathcal{N}(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k). \tag{12}$$

$$\hat{s}_{1:T} = \arg\max_{s_{1:T}} \mathbb{P}(s_{1:T} \mid \mathbf{X}, \theta), \tag{13}$$

where s_t denotes the latent regime, $A_{s_t, s_{t+1}}$ denotes the transition probability, and $(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)$ are regime-specific parameters.

Anomalies are detected when (i) the likelihood $\mathbb{P}(\mathbf{x}_t \mid s_t)$ falls below the empirical 1% tail, or (ii) regime transitions occur with probability below the estimated transition-matrix support.

3.7 Model performance evaluation

Ground-truth labels are constructed from documented macroeconomic and policy events. We employ a ±3-day event window to accommodate short-run information-diffusion lags. Robustness is assessed under alternative ±1-day and ±5-day windows.

Baseline comparisons use three standard anomaly-detection benchmarks: a PCA-based z-score baseline, a rolling 3σ residual threshold, and an AR(1) residual anomaly baseline.

For performance evaluation, AUROC, AUPRC, precision, recall, and F1 are employed:

$$\text{AUROC} = \int_0^1 \text{TPR}(x)\,dx, \tag{14}$$

$$\text{AUPRC} = \sum_{i=1}^{n} (r_i - r_{i-1}) p_i, \tag{15}$$

$$\text{Precision} = \frac{TP}{TP+FP}, \quad \text{Recall} = \frac{TP}{TP+FN}, \quad \text{F1} = \frac{2\cdot\text{Precision}\cdot\text{Recall}}{\text{Precision}+\text{Recall}}. \tag{16}$$

All metrics are computed using the model-specific anomaly labels and the event-based ground-truth labels.

4 Results

In this section, we present the results obtained from a hybrid framework anomaly study that integrated PCA, Isolation Forest, and GRU for robust anomaly detection on the Zimbabwe Stock Exchange financial returns dataset. The primary sample spans March 2022 to the present, i.e., March 2025, matching the period used in the methodological calibration. Table 1 shows the ZSE stock returns data information.

Table 1: ZSE stock returns data information

Information	Non-Null Count	Dtype
Index	752	09/03/2022–31/03/2025
32 Columns	752	float64
Memory Usage	4,128	+kB

31 stocks and the main market index, ZSE ASI, were collected from the Investing.com platform, reflecting sparse stocks for diversification, which is a serious concern for portfolio optimisation. However, most of the analysis is done on the market index, the ZSE ASI, which is the barometer of the Zimbabwean financial markets. The returns of ZSE ALSI are shown in Figure 1.

The plot shows long periods of stable behaviour followed by sudden volatility spikes, which are characteristic region where market regime shift as a result of underlying policy induced shocks. The sharp, irregular fluctuation that appeared in mid-March 2024 suggests a structural break or unusual market activity that anomaly-detection models should capture. The analysis provides a strong basis for detecting abnormal return patterns rather than for outlier detection.

Results from feature importance analysis on the ZSE financial returns dataset revealed the most important 10 features based on the average scores of Correlation, PCA and Random Forest Analysis. This is shown in Table 2.

Figure 1: ZSE returns post COVID-19

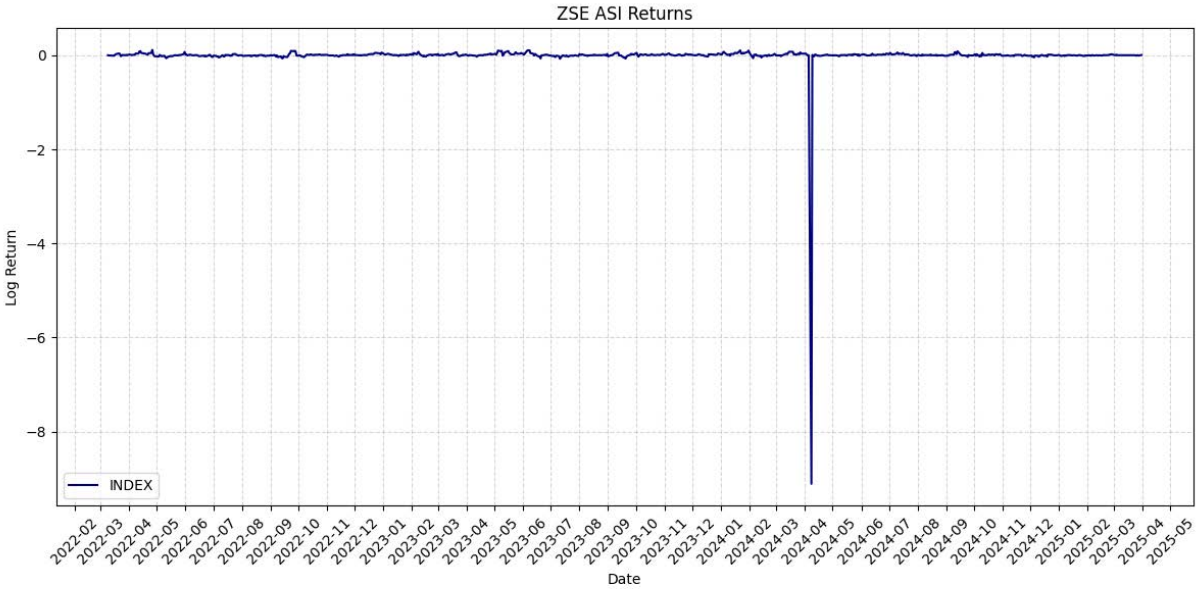

Table 2: Feature importance analysis

Stock/Feature	Correlation	PCA	Random Forest	Average Score
nmbzprice	0.004765	1.000000	0.000013	0.334926
tangaprice	0.000676	0.954302	0.000028	0.318335
unifreprice	0.002044	0.951422	0.000027	0.317831
zbprice	0.011657	0.939696	0.000016	0.317123
hippoprice	0.004508	0.931219	0.000015	0.311914
dairbprice	0.001744	0.931642	0.000014	0.311133
batprice	0.003938	0.919086	0.000016	0.307680
fideltprice	0.002021	0.920825	0.000040	0.307629
fbcprice	0.004738	0.917025	0.000023	0.307262
aristonprice	0.001341	0.905166	0.000025	0.302177

PCA scores dominate across features, indicating that variance-driven structure carries most of the explanatory power for detecting abnormal behaviour. Simple correlations and tree-based splits contribute far less to identifying anomalies in this dataset. The consistently low Random Forest importance values suggest that anomalies are subtle and not easily captured through nonlinear feature interactions.

These patterns reinforce the need for sequence-based or latent-state models capable of capturing temporal and structural dependencies.

For linear dependence, correlation analysis between the top features themselves and the index was undertaken, and the resulting correlation heatmap is presented in Figure 2.

It can be observed that problems of dimensionality, noise and multicollinearity have been dealt with during feature importance analysis. The scatter plot showing distributional structure of the same features in Figure 3 illustrates distributional spread and nonlinear clustering among top features. These patterns complement the heatmap by

Figure 2: Correlation between 10 most important features

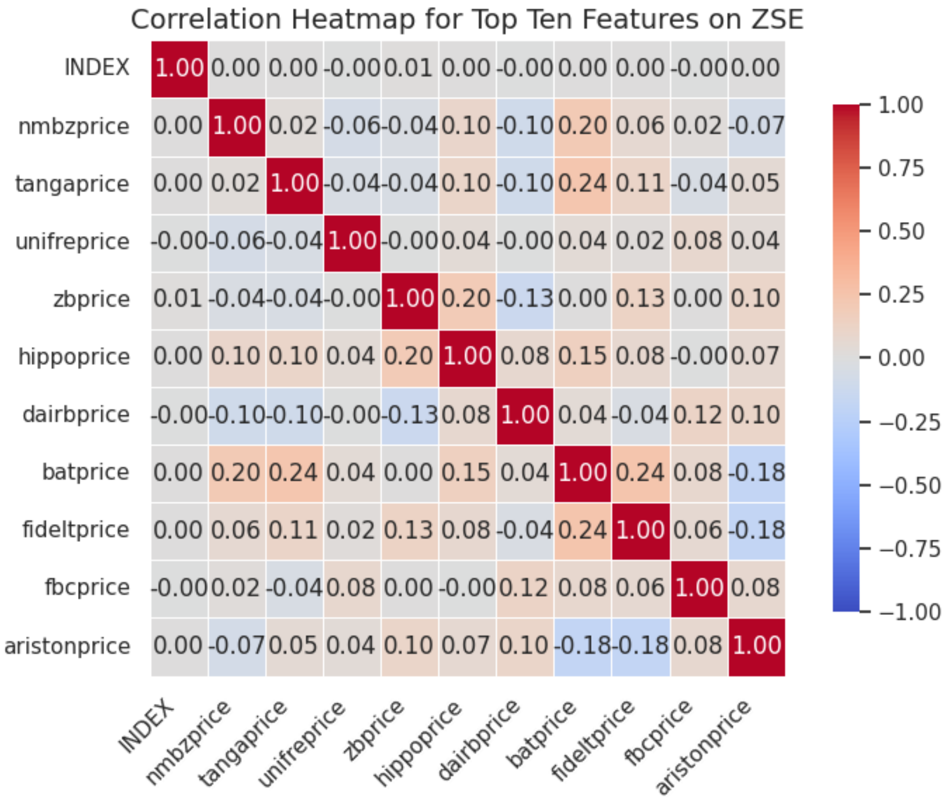

	INDEX	nmbzprice	tangaprice	unifreprice	zbprice	hippoprice	dairbprice	batprice	fideltprice	fbcprice	aristonprice
INDEX	1.00	0.00	0.00	-0.00	0.01	0.00	-0.00	0.00	0.00	-0.00	0.00
nmbzprice	0.00	1.00	0.02	-0.06	-0.04	0.10	-0.10	0.20	0.06	0.02	-0.07
tangaprice	0.00	0.02	1.00	-0.04	-0.04	0.10	-0.10	0.24	0.11	-0.04	0.05
unifreprice	-0.00	-0.06	-0.04	1.00	-0.00	0.04	-0.00	0.04	0.02	0.08	0.04
zbprice	0.01	-0.04	-0.04	-0.00	1.00	0.20	-0.13	0.00	0.13	0.00	0.10
hippoprice	0.00	0.10	0.10	0.04	0.20	1.00	0.08	0.15	0.08	-0.00	0.07
dairbprice	-0.00	-0.10	-0.10	-0.00	-0.13	0.08	1.00	0.04	-0.04	0.12	0.10
batprice	0.00	0.20	0.24	0.04	0.00	0.15	0.04	1.00	0.24	0.08	-0.18
fideltprice	0.00	0.06	0.11	0.02	0.13	0.08	-0.04	0.24	1.00	0.06	-0.18
fbcprice	-0.00	0.02	-0.04	0.08	0.00	-0.00	0.12	0.08	0.06	1.00	0.08
aristonprice	0.00	-0.07	0.05	0.04	0.10	0.07	0.10	-0.18	-0.18	0.08	1.00

revealing structure not visible in correlations.

Scatter plot patterns confirm the presence of nonlinear structure relevant for anomaly detection. PCA-based multivariate structure remains the primary driver of anomaly scores. The PCA accumulated variance plot in Figure 4 reveals that multicollinearity and noisy data problems have been addressed, which was crucial before feeding the data into anomaly detectors.

Only 25 PCA components are needed to retain 95% of the variance. Using 25 features is sufficient for modelling without sacrificing too much information. PCA-transformed features provide a consistent reduced space for all anomaly detectors. This ensures comparability across point, contextual, and collective anomaly types.

IForest's anomaly score fluctuated over time, with spikes above the threshold indicating potential anomalies. Ground truth events in the Zimbabwean markets aligned with high anomaly scores. Several spikes coincide with major 2024 policy and currency events. These alignments indicate sensitivity to macro-driven disruptions.

Some point anomalies flagged by Isolation Forest do not coincide with documented macroeconomic or policy events. These may reflect data irregularities or unrecorded market events. Such cases are interpreted cautiously as potential data-quality signals.

PCA and GRU-based contextual anomaly detection model performance is shown in Figure 6.

Figure 3: Scatter plot of 10 most important features

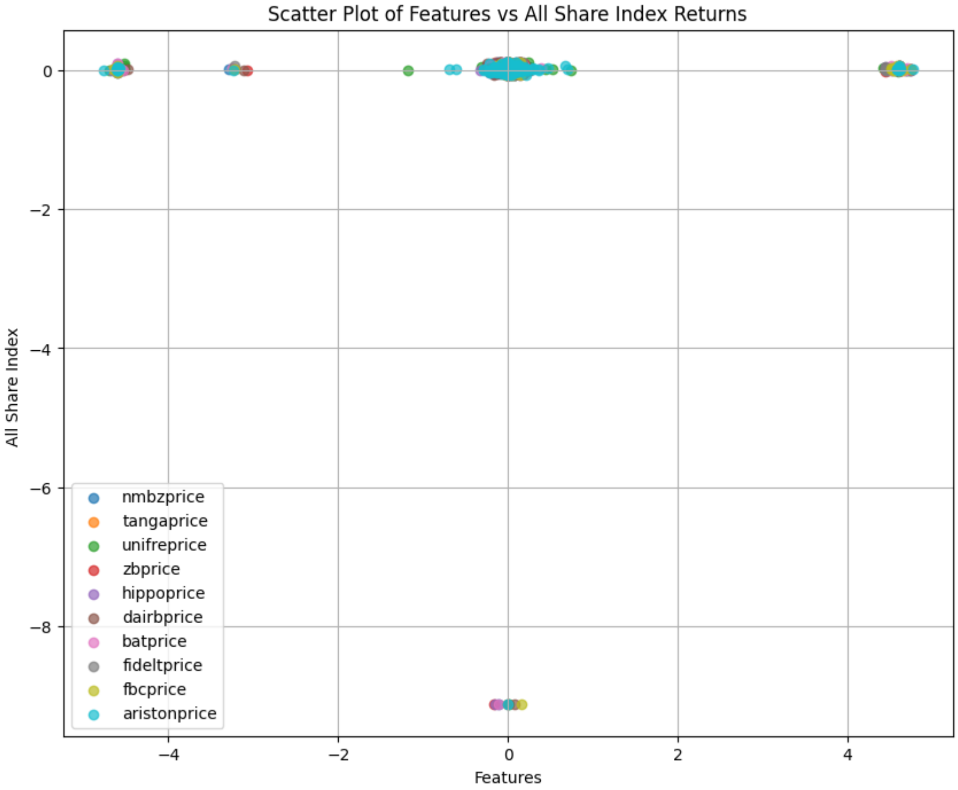

Figure 4: Dimensionality reduction using PCA

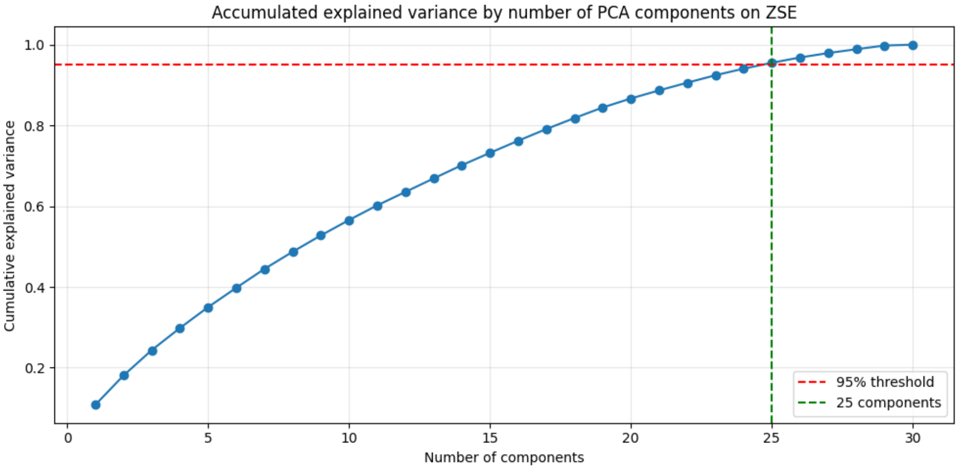

Table 3: Top 10 iForest point anomalies scores and ground truths

Date	Anomaly Score	Possible Ground Truth Event
2024-04-17	0.107960	Market Sentiment to ZIG announcement
2024-05-10	0.104967	Tight monetary policy sentiments
2024-05-29	0.103293	Economic Growth Revised downwards

Figure 5: iForest scores for point anomalies

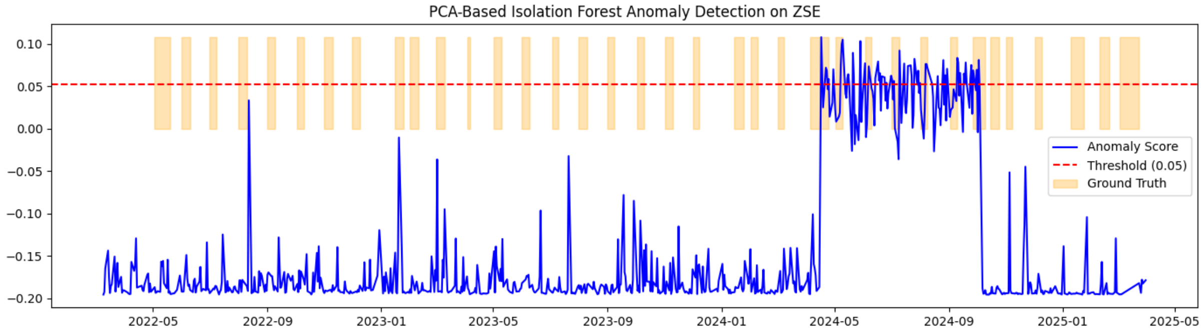

Figure 6: GRU scores for contextual anomalies

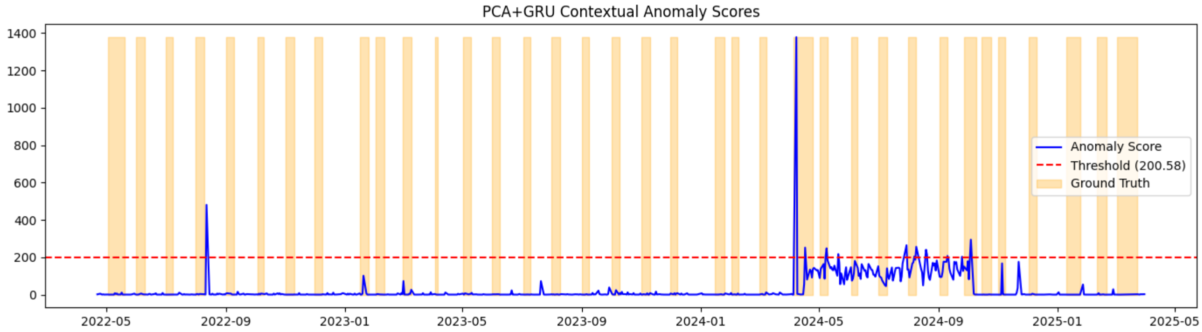

The GRU contextual anomaly scores spiked sharply at several points. Low recall indicates that GRU prioritises precision over broad anomaly coverage. Threshold sensitivity contributes to the under-detection of some true anomalies. Some spikes outside known events may indicate unobserved shocks or data noise.

The PCA-based HMM collective anomaly detection model performance is shown in Figure 7.

Figure 7 presents both the anomaly score series and inferred sequences from the hybrid framework. The upper panel shows the PCA-based anomaly score series, where higher values indicate stronger deviations from normal return behaviour, together with the threshold used to flag collective anomalies. The lower panel corresponds to the HMM model and displays the inferred regime sequence (calm vs. stressed) alongside the ZSE index and documented ground-truth events. HMM identifies transitions between calm and stressed volatility regimes. Several regime transitions align with documented macro events. Table 4 summarises the collective anomalies detected by the Hidden Markov Model (HMM). All identified dates fall within the stressed regime, indicating persistent high-volatility conditions during these periods.

Table 5 reports the transition probabilities between the three HMM states on the ZSE. State 0 represents a calm regime with a high self-transition probability (0.8909), indicating strong persistence. State 2 has no self-transition and always transitions to State 1, reflecting a short-lived shock regime that rapidly evolves into the more persistent

Figure 7: HMM scores for collective anomalies and corresponding regime shifts

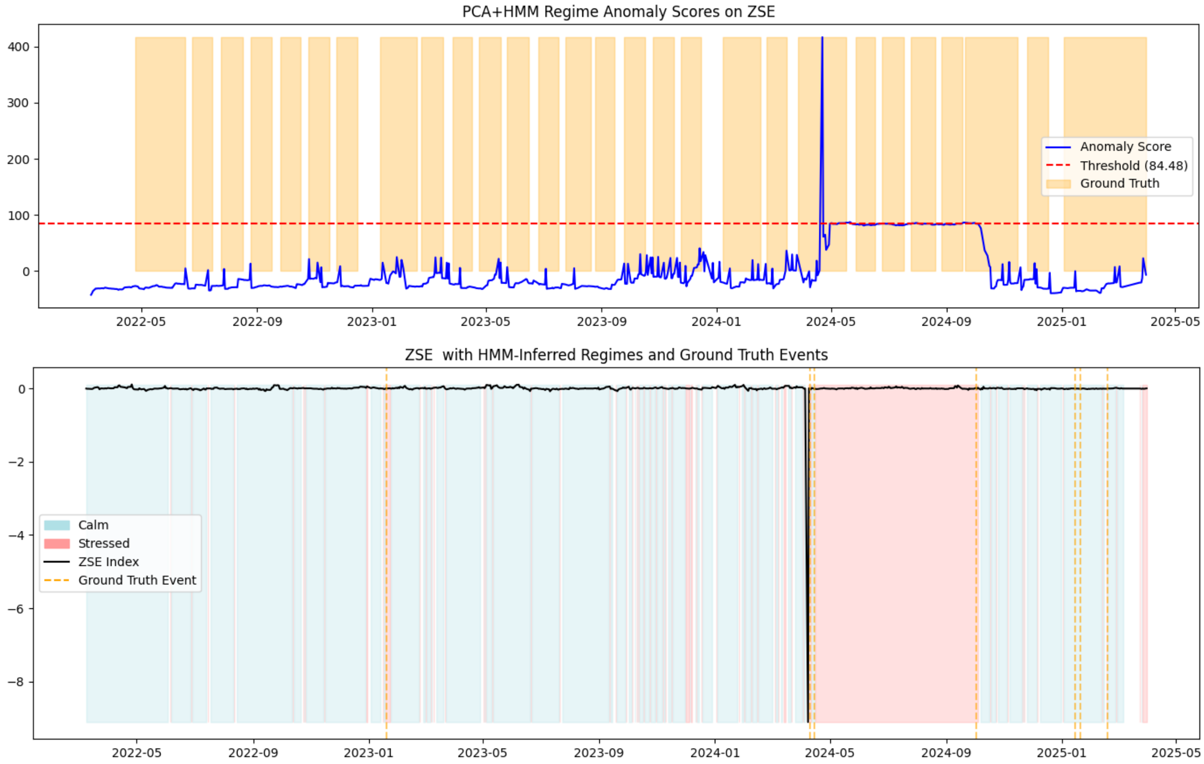

Table 4: Collective anomalies, regime shifts summary on ZSE

Date	Anomaly Score	Ground Truth	HMM State	Regime Type
2024-04-22	416.657008	1	1	Stressed
2024-05-21	87.215225	0	1	Stressed
2024-05-22	86.643075	0	1	Stressed
2024-09-19	86.614547	0	1	Stressed
2024-09-20	86.332617	1	1	Stressed
2024-05-17	86.320767	1	1	Stressed
2024-05-20	86.106676	0	1	Stressed
2024-09-30	85.713521	1	1	Stressed
2024-07-30	85.703041	1	1	Stressed
2024-09-23	85.693947	1	1	Stressed
2024-09-18	85.636873	0	1	Stressed
2024-07-31	85.472327	1	1	Stressed
2024-10-01	85.435082	1	1	Stressed
2024-05-16	85.423908	1	1	Stressed
2024-08-01	85.419189	1	1	Stressed

stressed state.

Average anomaly scores decrease from calm to stressed regimes, indicating rising instability. The absence of self-transition in State 2 suggests that this regime is short-lived and reactive, with abrupt structural breaks that quickly transition into the stressed regime rather than representing a stable long-run state. This is further illustrated in Table 6.

Table 7 summarises the performance of the PCA baseline and the three PCA-based hy-

Table 5: Transition probabilities between states on the ZSE

States	To State 0	To State 1	To State 2
State 0	0.8909	0.1073	0.0018
State 1	0.3141	0.6859	0.0000
State 2	0.0000	1.0000	0.0000

Table 6: Regime summary table on the ZSE

States	Avg Anomaly Score	Regime Type	Self-Transition Prob
State 0	−19.5024	Calm	0.8909
State 1	48.3982	Stressed	0.6859
State 2	−0.3304	Stressed	0.0000

brid anomaly detection models. The metrics include threshold, recall, precision, AUROC, AUPRC, and F1 score.

Table 7: Summary results: model performance vs known Zimbabwean anomalies including PCA baseline

Model	Threshold	Recall	Precision	AUROC	AUPRC	F1
PCA Baseline (Z-score)	0.3084	0.0220	0.7500	0.4591	0.3587	0.0427
PCA+Isolation Forest	0.0522	0.0769	0.3962	0.5226	0.3952	0.1288
PCA+GRU	200.5758	0.0293	0.5333	0.5154	0.4069	0.0556
PCA+HMM	84.4831	0.0526	0.8158	0.5472	0.7995	0.0987

The PCA baseline exhibits very low recall, indicating limited sensitivity to known anomalies. Combination of PCA and Isolation Forest improves recall and overall balance but remains conservative. The combination of PCA and GRU shows moderate precision but low recall due to strict reconstruction-error thresholds while PCA and HMM achieves the highest precision and AUPRC, reflecting strong alignment with documented macroeconomic and policy-driven anomalies.

The combination of PCA and GRU, which targets contextual anomalies, exhibits moderate precision (53.33%) but low recall (2.93%), reflecting the difficulty of detecting temporal context shifts in a frontier-market dataset with structural breaks and limited depth. All hybrid models provide incremental value beyond a simple statistical detector, but each does so along different dimensions.

Figure 8 presents the elbow plots used to select anomaly-detection thresholds for each model in the hybrid framework. Subfigure (a) corresponds to the PCA model and shows the elbow structure of the PCA-based anomaly score distribution; Subfigure (b) corresponds to the Isolation Forest model and displays the elbow pattern of the iForest anomaly scores; Subfigure (c) corresponds to the GRU model and shows the elbow plot of the GRU reconstruction-error scores; and Subfigure (d) corresponds to the HMM model and illustrates the elbow behaviour of the regime-likelihood scores. Elbow plots for the PCA base-

Figure 8: Elbow scores for PCA baseline, Isolation Forest, GRU, and HMM anomaly scores on the ZSE

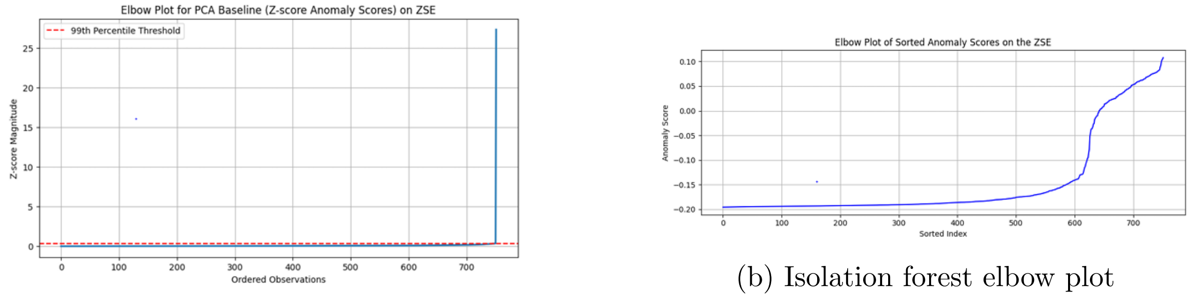

(a) PCA baseline (Z-score) elbow plot

(b) Isolation forest elbow plot

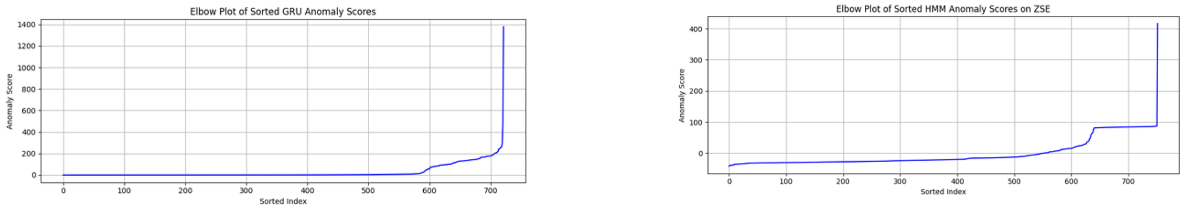

(c) GRU anomaly scores elbow plot

(d) HMM anomaly scores elbow plot

line, Isolation Forest, GRU, and HMM models all exhibit a clear inflexion point around the 700th ordered observation, separating low-level noise from high anomaly scores. This supports the use of the elbow method as a principled thresholding heuristic. Sensitivity checks around the elbow region (±50 observations) confirm that while absolute precision and recall values shift with the threshold, the relative ranking of models remains stable. The combination of PCA and HMM yields the highest precision, PCA and Isolation Forest yields the highest recall, and PCA with GRU yields intermediate performance. The PCA baseline shows a flatter distribution with a late, sharp rise, consistent with a univariate z-score detector. Overall, the elbow-based thresholding approach provides a transparent and replicable method for selecting operating points across models, and the qualitative conclusions about model behaviour and comparative performance remain robust.

For performance evaluation, the combination of PCA and HMM achieves the highest precision, reflecting strong regime-sensitive detection. Low recall shows that the model favours precision over coverage. The combination of PCA and IForest captures the widest range of anomalies. However, the model is sensitive to local deviations, including subtle data shifts, and PCA and HMM-based model results' regime transitions align with documented macro events. All hybrid models outperform the PCA baseline along different dimensions.

Recall remains low because the framework is deliberately conservative, prioritising precision in a frontier-market environment with sparse and noisy ground-truth labels. This behaviour is consistent with anomaly-detection literature, where high-precision, low-recall trade-offs are common in surveillance and early-warning applications.

5 Discussion

31 stocks represent a data sparsity challenge of the listed firms on the Zimbabwe Stock Exchange, which is a diversification concern for freely available stock market data in Zimbabwe for market development and depth. Feature analysis identified the 10 most influential features by combining correlation, PCA loadings, and Random Forest importance into a unified ranking. Each of these measures captures a different aspect of relevance, including linear dependence, variance contribution, and non-linear predictive power. This combined ranking is used as an exploratory tool to prioritise variables for subsequent anomaly detection. The aggregation reflects the complementary strengths of linear, variance-based, and non-linear importance measures in high-dimensional financial data. This approach follows ensemble feature selection practices that aim to stabilise rankings in noisy environments.

PCA identified 25 components or stock market index constituents are needed to retain 95% of the variance or explain the behaviour of ZSE ASI returns. This means that they are sufficient for modelling the Zimbabwe Stock Exchange behaviour without sacrificing too much information. It also balances dimension reduction, information retention and risk of over-fitting in our models. The 25-dimensional PCA-transformed feature space was used as the common input to all anomaly detection models in the framework. Isolation Forest, GRU, and HMM were each trained on these PCA components to ensure consistency in the underlying representation. This design reduces multicollinearity and noise while maintaining a comparable feature basis across point, contextual, and collective anomaly detection. Using a shared PCA representation ensured that differences in performance arose from model structure rather than from feature construction. This pipeline design integrates dimensionality reduction and anomaly detection into a coherent modelling sequence.

It can be observed that problems of dimensionality, noise and multicollinearity have been dealt with during feature importance analysis. The observed effects of PCA are consistent with previous studies, as in Crépey et al. (2022). Improved results from IForest's anomaly score fluctuated over time, with spikes above the 0.05 threshold indicating potential anomalies aligned with ground truths. Several high Isolation Forest anomaly scores coincide with documented macroeconomic and policy events in Zimbabwe. These coincidences indicate that the model is sensitive to some major disruptions in the market. Low recall indicates that IForest prioritises precision over broad anomaly coverage. Threshold sensitivity contributes to under-detection of some known events.

A subset of the largest Isolation Forest anomaly scores overlaps with known episodes of currency reform, drought, and policy tightening. This overlap suggests that the model can highlight periods of heightened stress in the market. Detected anomalies represent high-confidence signals rather than full event coverage. Some missed events may reflect subtle or gradual market adjustments.

Isolation Forest flagged 10% of the dataset as anomalies, with 38% linked to genuine ground truths or macroeconomic shocks such as currency reforms and droughts.

The GRU contextual anomaly detector identifies several periods of heightened temporal stress, many of which coincide with documented macroeconomic disturbances. This confirms that the combination of PCA and GRU configuration is sensitive to short-horizon structural changes in return dynamics. GRU emphasises precision and flags only the strongest temporal deviations. Many weaker contextual anomalies fall below the selected threshold. Its detections should therefore be interpreted as highlighting candidate periods of interest rather than providing comprehensive anomaly coverage.

The GRU contextual anomaly scores spiked sharply at several points, often exceeding the threshold of 200.58, signalling potential anomalies. Some of the largest GRU contextual anomaly scores occur around major macroeconomic and policy events. This pattern indicates that the PCA and GRU configuration can detect certain episodes of heightened temporal stress in returns.

GRU's conservative thresholding results in selective detection of temporal shocks.

The GRU-based contextual anomaly detector highlights several periods that coincide with documented shocks to the currency and policy environment. These detections suggest that the model is responsive to shifts in the temporal structure of returns during turbulent episodes. GRU provides targeted signals rather than full anomaly coverage.

It is best interpreted as a complementary detector within the hybrid framework.

Out-of-event spikes may reflect unrecorded shocks or data irregularities. These cases highlight the need for improved data validation. Overall, the hidden Markov model (HMM) was used for collective anomaly detection in the form of regime shifts. The PCA and HMM specification identifies several periods characterised by shifts in the inferred latent state of the market. Some of these regime shifts occur close to major macroeconomic and policy events, indicating partial alignment with known structural changes.

HMM results depend strongly on the number of latent states. Regime transitions should be interpreted as structural indicators rather than precise event matches. The HMM-based regime classification reveals transitions between calm and stressed states that overlap with episodes of elevated volatility and policy uncertainty. These transitions provide a useful summary of collective anomalies in the index-level dynamics. Threshold and state-selection choices influence regime boundaries. Results should be interpreted as descriptive rather than causal.

The introduction of a PCA-only z-score baseline clarifies the incremental value of the hybrid framework. The baseline achieves modest recall and AUROC, reflecting the limitations of univariate statistical detectors in capturing multiscale market disruptions. Relative to this baseline, the combination of PCA and HMM delivers substantially higher precision and AUPRC, demonstrating improved ability to isolate structurally meaningful anomalies. PCA and Isolation Forest combination improves anomaly coverage, while the combination of PCA and GRU captures temporal context unavailable to the baseline. These comparisons confirm that the hybrid framework adds explanatory depth beyond simple statistical rules, particularly in distinguishing point, contextual, and collective anomalies. Regime shifts align with several documented macro events, supporting the model's structural sensitivity. Not all transitions match events, reflecting both data limitations and model uncertainty. A key limitation is that ground-truth anomalies are derived from documented macroeconomic events, which may introduce confirmation bias and understate anomalies not linked to known policy or market shocks.

The proposed framework can complement existing volatility, residual-based surveillance and standard stylised facts analysis tools by adding a multiscale view of structural and behavioural patterns that goes beyond stylised facts analysis. Its outputs could be integrated into regulatory dashboards as an additional layer of diagnostic information rather than as a standalone decision rule.

Future work should benchmark the hybrid framework against additional baselines, including volatility-based detectors and ARIMA/GARCH residual filters, to further assess whether the added complexity yields material gains in detection performance. A more granular analysis of false positives, distinguishing data artefacts from unrecorded events, would strengthen the interpretability of anomaly classifications. Extending the dataset to include more data and alternative data sources would address concerns about representativeness and enhance robustness.

6 Conclusion

This study proposes a PCA-based hybrid anomaly detection framework that combined point (Isolation Forest), contextual (GRU), and collective (HMM) anomaly detection to the Zimbabwe Stock Exchange, showing that anomalies in returns may reflect both statistical irregularities and underlying structural dynamics. The results show partial alignment between detected anomalies and documented macroeconomic events, but consistently low recall across models limits the strength of these conclusions, a pattern that reflects the complex and heterogeneous feature distributions characteristic of frontier markets. The introduction of a PCA-only z-score baseline clarifies the incremental value added by the hybrid framework. Relative to this baseline, the hybrid models demonstrate improved precision, temporal sensitivity, and regime-awareness, confirming that the added structure provides analytical benefits beyond simple statistical rules. The framework strengthens its ability to detect potential market risks. This extends beyond the stylised facts confirmed on the ZSE in previous studies. Although predictive performance declines over longer horizons, the model still outperforms traditional benchmarks. Additional baselines must be considered in future, including volatility-based detectors and ARIMA/GARCH residual filters, to determine whether the added complexity yields material improvements in anomaly detection performance. Expanding the dataset across more securities and longer horizons would improve representativeness and robustness.

Acknowledgments

I would like to express my sincere gratitude to the promoter for their invaluable guidance and support throughout this research. Special thanks to NUST, Department of Finance, for providing the necessary resources and facilities. Constructive feedback from anonymous reviewers greatly improved the quality of this manuscript.

Declaration of competing interest

The authors have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper. We also declare the use of Microsoft Copilot for language refinement and organisation. All analytical decisions, interpretations, and intellectual contributions were made solely by the authors.

References

Adams, J., Hayunga, D., Mansi, S., Reeb, D., and Verardi, V. (2019). Identifying and treating outliers in finance. *Financial Management*, 48(2):345–384.

Aggarwal, C. C. (2015). *Outlier analysis*. Springer, 2nd edition.

Ahmed, M., Mahmood, A. N., and Hu, J. (2016). A survey of network anomaly detection techniques. *Journal of Network and Computer Applications*, 60:19–31.

Ang, A. and Bekaert, G. (2002). Regime changes in interest rates. *Journal of Business & Economic Statistics*, 20(2):163–182.

Arjunan, G. (2022). Enhancing data quality and integrity in machine learning pipelines: Approaches for detecting and mitigating bias. *International Journal of Scientific Research and Management*, 10(9):940–945.

Bailey, D. H., Borwein, J. M., López de Prado, M., and Zhu, Q. J. (2016). The probability of backtest overfitting. *Journal of Computational Finance*, 20(2):39–69.

Bakumenko, A. and Elragal, A. (2022). Detecting anomalies in financial data using machine learning algorithms. *Systems*, 10(5):130.

Baragona, R. and Battaglia, F. (2007). Outliers detection in multivariate time series by independent component analysis. *Neural Computation*, 19(7):1962–1984.

Blázquez-García, A., Conde, A., Mori, U., and Lozano, J. A. (2021). A review on outlier/anomaly detection in time series data. *ACM Computing Surveys*, 54(3):1–33.

Bouattour, M. and Martinez, I. (2019). Efficient market hypothesis: An experimental study with uncertainty and asymmetric information. *Finance Contrôle Stratégie*, 22(4).

Campbell, J. Y., Lo, A. W., MacKinlay, A. C., and Whitelaw, R. F. (1998). The econometrics of financial markets. *Macroeconomic Dynamics*, 2(4):559–562.

Campos, G. O., Zimek, A., Sander, J., Campello, R. J. G. B., Micenková, B., Schubert, E., Assent, I., and Houle, M. E. (2016). On the evaluation of unsupervised outlier detection: Measures, datasets, and an empirical study. *Data Mining and Knowledge Discovery*, 30(4):891–927.

Chalapathy, R. and Chawla, S. (2019). Deep learning for anomaly detection: A survey. arXiv.

Chandola, V., Banerjee, A., and Kumar, V. (2009). Anomaly detection: A survey. *ACM Computing Surveys*, 41(3):1–58.

Chandrashekar, G. and Sahin, F. (2014). A survey on feature selection methods. *Computers & Electrical Engineering*, 40(1):16–28.

Choi, H. D. and Zaiane, O. R. (2015). Time series contextual anomaly detection for detecting market manipulation in stock market. In *2015 IEEE International Conference on Data Science and Advanced Analytics (DSAA)*, pages 1–10. IEEE.

Crépey, S., Lehdili, N., Madhar, N., and Thomas, M. (2022). Anomaly detection in financial time series by principal component analysis and neural networks. *Algorithms*, 15(10):385.

Ester, M., Kriegel, H.-P., Sander, J., and Xu, X. (1996). A density-based algorithm for discovering clusters in large spatial databases with noise. In *Proceedings of the 2nd International Conference on Knowledge Discovery and Data Mining (KDD)*, pages 226–231. AAAI Press.

Fischer, T. and Krauss, C. (2018). Deep learning with long short-term memory networks for financial market predictions. *European Journal of Operational Research*, 270(2):654–669.

Foorthuis, R. (2018). *A typology of data anomalies*, pages 26–38. Springer.

Foorthuis, R. (2021). On the nature and types of anomalies: A review of deviations in data. *International Journal of Data Science and Analytics*, 12(4):297–331.

Gao, R., Zhang, T., Sun, S., and Liu, Z. (2019). Research and improvement of isolation forest in detection of local anomaly points. *Journal of Physics: Conference Series*, 1237(5):052023.

Gupta, M., Gao, J., Aggarwal, C. C., and Han, J. (2014). Outlier detection for temporal data: A survey. *IEEE Transactions on Knowledge and Data Engineering*, 26(9):2250–2267.

Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, 57(2):357–384.

Hundman, K., Constantinou, V., Laporte, C., Colwell, I., and Soderstrom, T. (2018). Detecting spacecraft anomalies using LSTMs and nonparametric dynamic thresholding. arXiv.

Li, F. W. W., Kim, H., Cucuringu, M., and Ma, T. (2025). Can LLM-based financial investing strategies outperform the market in long run? arXiv.

Li, J., Pedrycz, W., and Jamal, I. (2017). Multivariate time series anomaly detection: A framework of hidden Markov models. *Applied Soft Computing*, 60:229–240.

Liang, A., Hu, Y., and Li, G. (2023). The impact of improved PCA method based on anomaly detection on chiller sensor fault detection. *International Journal of Refrigeration*, 155:184–194.

Malhotra, P., Vig, L., Shroff, G., and Agarwal, P. (2015). Long short term memory networks for anomaly detection in time series. In *Proceedings of the 23rd European Symposium on Artificial Neural Networks, Computational Intelligence and Machine Learning (ESANN)*, pages 89–94. i6doc.com.

Murekachiro, D. (2025). Stock market calendar anomalies in sub-Saharan Africa stock markets. *International Journal of Research and Scientific Innovation*, 12(2):42–49.

Oelschläger, L., Adam, T., and Michels, R. (2024). fHMM: Hidden Markov models for financial time series in R. *Journal of Statistical Software*, 109(9):1–25.

Pinto, S. O. and Sobreiro, V. A. (2024). Note on the curse of dimensionality regarding financial anomaly detection. *International Journal of Business Intelligence and Systems Engineering*, 2(1):77–87.

Prajesha, T. and Veni, S. (2023). An efficient outlier detection using isolation forest based on robust scaling and principal component analysis for the prediction of anxiety disorder. *Indian Journal of Science and Technology*, 16(29):2244–2251.

Sakurada, M. and Yairi, T. (2014). Anomaly detection using autoencoders with nonlinear dimensionality reduction. In *Proceedings of the MLSDA 2014 2nd Workshop on Machine Learning for Sensory Data Analysis*, pages 4–11. ACM.

Salehi, M. and Rashidi, L. (2018). A survey on anomaly detection in evolving data. *SIGKDD Explorations*, 20(1):33–47.

Saurabh, P. and Verma, B. (2023). Negative selection in anomaly detection—A survey. *Computer Science Review*, 48:100557.

Shabir, I. and Pearl, J. (2024). Anomaly detection in financial services: The power of data-driven insights. *International Journal of Communication and Information Technology.*

Shiraj, M. M. B., Rahman, M., Al-Imran, M., Liza, M., Murshed, M. M., and Akhter, N. (2024). Anomaly detection in financial time series data via mapper algorithm and DBSCAN clustering. *World Journal of Advanced Engineering Technology and Sciences*, 13(1):70–84.

Shyu, M.-L., Chen, S.-C., Sarinnapakorn, K., and Chang, L. (2003). A novel anomaly detection scheme based on principal component classifier. In *Proceedings of the IEEE Foundations and New Directions of Data Mining Workshop*, pages 172–179. IEEE.

Spoor, J. M., Weber, J., and Ovtcharova, J. (2025). Anomaly detection and prediction evaluation for discrete nonlinear dynamical systems. *Transactions of the Institute of Measurement and Control*, 47(7):1258–1272.

Tang, C., Xu, L., Yang, B., Tang, Y., and Zhao, D. (2023). GRU-based interpretable multivariate time series anomaly detection in industrial control system. *Computers & Security*, 127:103094.

Tsay, R. S. (2005). *Analysis of financial time series.* Wiley, 2nd edition.

Xia, X., Pan, X., Li, N., He, X., Ma, L., Zhang, X., and Ding, N. (2022). GAN-based anomaly detection: A review. *Neurocomputing*, 493:497–535.

Zhao, R., Yan, R., Chen, Z., Mao, K., Wang, P., and Gao, R. X. (2019). Deep learning and its applications to machine health monitoring. *Mechanical Systems and Signal Processing*, 115:213–237.

Zhou, P. (2025). A survey of streaming data anomaly detection in network security. *PeerJ Computer Science*, 11:e3066.