

Forecasting Market Regimes in an Emerging European Market: A Machine Learning Framework for the Romanian BET Index

Bogdan Ionut ANGHEL*

Bucharest University of Economic Studies, Bucharest, Romania

**Corresponding author, bogdan.anghel@rei.ase.ro*

Radu LUPU

Bucharest University of Economic Studies, Bucharest, Romania

Romanian Academy, Bucharest, Romania¹

radu.lupu@rei.ase.ro

Abstract. *Financial markets do not evolve in a stable or uniform way. They often move through different phases, shaped by changes in risk, returns, volatility, and investor behaviour. These shifts are especially important in emerging markets, where lower liquidity, exposure to global shocks, and domestic vulnerabilities can make transitions from calm periods to stress episodes faster and more severe. This paper examines whether such regimes can be identified and forecast in the Romanian equity market, using the BET index as the main reference. The analysis covers the period from November 2017 to November 2025. In the first stage, a three-state Gaussian Hidden Markov Model is used to identify latent market regimes. The estimated states are then interpreted as calm, stress, and panic regimes based on their conditional realised volatility. In the second stage, the regimes are forecast one day ahead using an XGBoost classifier trained on macro-financial variables, including global risk indicators, cross-asset returns, sovereign yield curve measures, commodity prices, and domestic macroeconomic indicators. The results show strong overall predictive performance. Calm periods are identified with high reliability, while stress regimes are detected with meaningful sensitivity. Panic regimes, however, remain harder to classify because they occur less frequently, which highlights the difficulty of forecasting extreme market conditions in emerging economies.*

Keywords: Market regimes, Hidden Markov model, XGBoost, Emerging markets, BET index, Volatility forecasting, Macro-financial indicators.

Introduction

Financial markets do not usually evolve in a smooth or uniform way. They tend to move through different phases, shaped by changes in risk, return, volatility, and investor behaviour. These phases are commonly described as market regimes. Identifying them is important because sudden changes in volatility, correlations, and downside risk can affect portfolio decisions, risk management, and the monitoring of financial stability.

Regime changes are especially relevant in emerging markets. Compared with developed financial systems, emerging markets often have lower liquidity, stronger sensitivity to global capital flows, and greater exposure to external macro-financial shocks. These characteristics can make shifts from calm periods to stress or crisis episodes faster and more persistent (Égert & Kočenda, 2011). For this reason, traditional linear models may not be sufficient to capture the nonlinear and state-dependent behaviour of emerging equity markets.

¹ Institute for Economic Forecasting

The Romanian equity market, represented in this study by the BET index, is a relevant case of an emerging European market. It has become increasingly connected to global financial markets, while still being influenced by domestic vulnerabilities. Previous studies have demonstrated nonlinear dynamics, volatility clustering, and structural breaks in the BET index, especially during periods of global financial stress (Stoica & Diaconășu, 2014; Voineagu et al., 2019). Romanian equities are also affected by global risk indicators, exchange rate movements, and sovereign risk premia, which suggests that regime changes may be forecast using macro-financial information.

This study uses a hybrid framework that combines regime identification with one-day-ahead regime forecasting. The analysis is guided by three research questions. First, can the Romanian equity market be divided into persistent latent regimes with a clear economic perspective? Second, can regimes be forecast one day ahead using macro-financial variables available in real time? Third, does a ML approach improve predictive performance compared with traditional regime-switching models?

Instead of treating market movements as a simple binary outcome, this paper uses a three-regime structure that distinguishes between calm, stress and panic states. This allows for a gradual interpretation of market deterioration and separates early stress signals from extreme panic episodes, approached also used by Ang & Bekaert, 2002. By combining Hidden Markov Models (HMMs) with XGBoost, the study examines whether regime changes in an emerging equity market can be identified and predicted more effectively.

Literature review

Hamilton (1989) established the foundation for using Markov regime-switching models to analyze nonlinear macroeconomic dynamics. The main idea is that economic variables can act differently across hidden states, with transitions between these states defined by Markov probabilities. This allows us to see economic cycles not as one stable process but as a series of regimes with distinct traits. Hamilton and Susmel (1994) extended this line of research by introducing the SWARCH model, which combines conditional heteroskedasticity with regime switching. Their results show that stock market volatility tends to move between low-volatility periods and crisis-like states. Dueker (1997) further developed the approach through Markov-switching GARCH models, showing that volatility is mean-reverting, but that this process differs across regimes. Marcucci (2005) also finds that MS-GARCH models perform better than standard single-regime GARCH models when forecasting realised volatility.

More recent research has used HMM to identify periods of financial stress and systemic risk. For example, Kremer (2018), documents clear regime shifts in United States equity market volatility and shows that high-volatility regimes become more persistent during crisis periods. Other studies have combined probabilistic regime-switching models with ML methods in order to improve early warning signals and capture more complex market patterns (Akioyamen et al., 2021).

On the other hand, literature shows that regime-switching models are able to distinguish between calm and turbulent market conditions with reasonable accuracy (Christiansen et al., 2012; Kole & van Dijk, 2017). However, in emerging markets, these dynamics are often more difficult to model and forecast. Walid et al., 2011 shows that lower liquidity, higher macroeconomic instability, and stronger exposure to global shocks can make transitions between regimes more sudden and more persistent. Evidence from Central and Eastern European markets points in the same direction. These economies tend to show higher volatility, more frequent structural problems and stronger sensitivity to global risk factors than Western European markets (Égert & Kočenda, 2011). In the case of Romania, Voineagu et al. (2019) find evidence of regime-dependent volatility

transmission in the BET index, while Stoica and Diaconăşu (2014) identify nonlinear dynamics connected to global financial cycles. Taken together, these findings support the use of flexible nonlinear models for analysing regime-dependent behaviour in emerging equity markets.

Not the last, recent studies have increasingly turned to ML methods for regime identification and forecasting. These methods are useful when the relationships between variables are nonlinear or when the number of predictors is large. Nystrup et al. (2017) show that tree-based models can perform better than traditional regime-switching approaches in detecting market states. Similarly, Kopa and D'Ecclesia (2020) report that gradient boosting methods improve the forecasting of downside risk in emerging stock markets.

Methodology

The study covers the period from November 2017 to November 2025 and is based on daily data. The Romanian equity market is represented by the BET index, which is used as the main market variable. The set of predictors includes a wide range of domestic and international risk indicators, such as implied volatility indices, namely VIX, V2TX and MOVE, Romanian sovereign credit default swap spreads, EUR/RON exchange rate movements, returns from equities, commodities, government bonds and exchange-traded funds, as well as domestic macroeconomic variables. Since some macroeconomic indicators are available only at a monthly frequency, they are converted to daily frequency through forward filling. This means that the latest available value is carried forward until a new observation is released, so that the model only uses information that would have been available to investors at that point in time, in line with Ang and Bekaert (2002).

For model estimation and evaluation, the dataset is divided chronologically into a training sample and a testing sample, using a 70/30 split. The training sample is used to estimate the HMM for regime identification and to train the XGBoost classifier for regime forecasting. The remaining observations are kept only for final test evaluation. This chronological approach is important because it avoids look-ahead bias and gives a more realistic assessment of forecasting performance. The test sample is then used to evaluate overall classification accuracy, the model's ability to distinguish between regimes, and the economic relevance of predicted stress and panic signals.

Market regimes are identified using a Gaussian HMM. The model is estimated on BET returns, 21-day realised volatility, and a selected group of risk indicators (Hamilton, 1989). It assumes that observed market variables are influenced by hidden states that change over time according to a first-order Markov chain. This structure allows the model to capture both persistent market phases and sudden transitions between them. Each hidden state is described by a multivariate Gaussian distribution with its own covariance structure, which allows volatility levels and relationships between variables to differ across regimes. The parameters are estimated by maximum likelihood using the Expectation–Maximisation algorithm, with a larger number of iterations to support convergence and obtain a stable sequence of inferred states.

To easily interpret the results from an economic perspective, the model is restricted to three latent states. These states are labelled after estimation as calm, stress, and panic regimes, based on their average realised volatility. More specifically, the mean 21-day realised volatility is calculated for each hidden state, and the states are then ordered from the lowest to the highest volatility. The state with the lowest volatility is labelled calm, the intermediate state is labelled stress, and the highest-volatility state is labelled panic. This way of ordering regimes is consistent with the empirical finance literature, where realised volatility is commonly used to separate different market conditions (Kritzman et al., 2012). The resulting classification offers a data-driven picture of historical market behaviour and is also used as the target variable in the supervised forecasting stage.

The predictive variables are designed to capture persistence, nonlinear spillovers, and possible early warning signals. The feature set includes realised volatility over several horizons, exponentially weighted volatility measures, higher-order moments, jump indicators, momentum variables, yield curve level, slope and curvature factors, macroeconomic changes and surprise proxies, as well as lagged returns from other asset classes. Together, these variables allow the model to account for both global risk conditions and domestic macro-financial developments.

In the final step, the regime for time $t+1$, as inferred by the HMM, is forecast using only information available at time t . Extreme Gradient Boosting is used for the multinomial classification task because it can capture nonlinear relationships, work with correlated predictors, and limit overfitting through regularisation (Chen & Guestrin, 2016). Training and validation follow a rolling chronological procedure, which preserves the time order of the data and keeps the evaluation strictly out of sample.

Table 1 shows that calm market conditions account for most observations in the sample, while stress and panic regimes appear less often. This class imbalance reflects the fact that tranquil periods tend to last longer in the Romanian equity market. At the same time, it creates a challenge for regime forecasting, because the limited number of stress and panic observations makes these states harder to identify and evaluate with precision.

Table 1. Distribution of market regimes in the full sample

Regime	Number of observations
Calm	2001
Stress	59
Panic	36

Source: Authors' own research results.

As shown in Figure 1, the Hidden Markov Model identifies long calm periods during phases of relatively stable market growth. These periods are interrupted by episodes of stress and by shorter panic regimes, which usually coincide with sharp market declines and sudden increases in volatility. The pattern suggests that calm regimes are highly persistent, while stress and panic states occur more occasionally. Overall, this supports the use of a three-regime classification, as it captures both normal market conditions and periods of market deterioration.




Figure 1. Market regimes identified using HMM

Source: Authors' own research results.

Alongside traditional regime-switching models, recent studies have increasingly used gradient boosting methods to forecast financial market regimes. Since the work of Chen and

Guestrin (2016), Extreme Gradient Boosting, or XGBoost, has become a widely used machine learning method in financial econometrics. Its popularity comes from its ability to work with large and mixed datasets, estimate models efficiently, and reduce overfitting through regularisation.

XGBoost is especially useful in financial time-series applications, where relationships between variables are often nonlinear and may change across market conditions. In such settings, standard parametric models can be too restrictive. Several empirical studies show that gradient boosting methods perform well in forecasting stock prices, asset returns, and related financial variables because they can capture complex interactions between predictors (Xu et al., 2022; Tan et al., 2023; Esmailizade et al., 2024). Another advantage is that the boosting procedure, which improves the model step by step, can deal relatively well with noisy data, missing values, and correlated predictors. These issues are common in macro-financial datasets.

Evidence from European equity markets also supports the use of XGBoost in this context. Anghel and Lupu (2025) show that gradient boosting models improve stock market forecasts when macro-financial transition risk indicators are included. Their findings suggest that XGBoost can capture both regime-dependent patterns and forward-looking risk signals in integrated financial markets. The study *Can Transition Risk Drive European Stock Market Predictions? An XGBoost Approach* provides a more detailed discussion of the forecasting mechanism and feature engineering strategy used in this type of framework (Anghel & Lupu, 2025).

After estimating the initial XGBoost model with the full set of predictors, an additional feature selection step is used. The purpose of this step is to make the model simpler, easier to interpret, and more stable. Although XGBoost can handle irrelevant or correlated variables better than many traditional models, too many redundant predictors can still make the results harder to explain and may reduce out-of-sample performance, especially when the sample is not very large.

To address this problem, a multi-step feature pruning procedure is applied after the initial model estimation. First, autoregressive variables and multiple lagged versions of the same indicators are removed. This is done to ensure that the regime forecasts are not driven mainly by mechanical persistence in returns or volatility, but by macro-financial information with clearer predictive content.

Second, only one transformation of each base variable is kept where possible. For example, the model avoids including the level, change, and volatility measure of the same variable at the same time. When there is no clear reason to prefer a transformed version, the original variable is retained. This helps prevent the same economic signal from being counted several times.

Third, multicollinearity among the remaining predictors is checked using the Variance Inflation Factor (VIF). The variables with the highest VIF values are then removed step by step until all remaining predictors meet the chosen collinearity threshold. At the same time, the final number of predictors is limited to a maximum of twenty variables.

This procedure produces a smaller and more interpretable set of predictors. It reduces the redundancy created by repeated versions of the same variables and limits the instability that can appear when predictors are highly correlated. By combining the nonlinear forecasting ability of XGBoost with a structured feature selection process, the final model offers a balance between predictive accuracy, transparency, and robustness. This is particularly important in emerging market applications, where macro-financial indicators are often noisy and strongly connected to one another.

Table 2. One-versus-rest AUC by regime

Hyperparameter	Value
Number of estimators (n_estimators)	200
Learning rate (learning_rate)	0.20
Maximum tree depth (max_depth)	6
Minimum child weight (min_child_weight)	3
Subsample ratio (subsample)	0.90
Column subsampling (colsample_bytree)	0.90
Gamma (gamma)	0

Source: Authors' own research results.

Table 2 presents the optimal hyperparameters selected through randomized search with five-fold cross-validation. The chosen configuration balances model flexibility and regularisation, which helps reduce overfitting and supports more reliable out-of-sample regime forecasts.

Results and discussions

This section evaluates the performance of the proposed HMM–XGBoost framework and interprets the results from the perspective of regime forecasting in an emerging equity market (see Figure 2). Since calm, stress, and panic regimes have different implications for risk management and investment decisions, the analysis focuses not only on overall accuracy, but also on how well the model predicts each regime separately.




Figure 2. Confusion matrix for regime classification

Source: Authors' own research results.

The model is evaluated on a hold-out test sample of 496 daily observations. Using the selected hyperparameters, the XGBoost classifier reaches an overall accuracy of 0.9153 and a weighted F1-score of 0.9246. These results suggest strong predictive performance, especially in a sample where calm market conditions are dominant. They also indicate that macro-financial variables contain useful short-term information for regime forecasting when they are modelled with a nonlinear method such as XGBoost (see Table 3).

However, the overall results do not fully show how the model performs across individual regimes. The regime-specific precision and recall scores reveal clear differences between calm,

stress, and panic states. These differences reflect both the economic nature of the regimes and the statistical difficulty created by class imbalance.

Table 3. Regime classification performance in the test sample

Regime	Precision	Recall	F1-score
Calm (0)	0.9739	0.9411	0.9572
Stress (1)	0.1944	0.50	0.28
Panic (2)	0	0	0
Macro average	0.3894	0.4804	0.4124

Source: Authors' own research results.

Calm market conditions are identified with high precision (0.9739) and a recall of 0.9411. This result is expected, since calm periods usually last longer and are characterised by more stable volatility and smoother return movements. The strong performance for calm regimes suggests that the proposed framework can identify periods in which market risk is relatively contained. This is useful for investors, as such periods may allow risk exposure to be maintained, or even increased, without a major rise in downside risk. The result is also in line with previous studies showing that regime-based models tend to perform well during stable market phases (Nystrup et al., 2017; Kritzman et al., 2012).

The stress regime shows a different pattern. Precision is relatively low (0.1944), while recall reaches 0.50. This means that the model correctly identifies around half of the stress observations in the test sample, but it also produces some false stress signals. In practical terms, the model seems to detect stress conditions early, sometimes before the market has fully moved into a stress regime.

From an economic perspective, this behaviour is reasonable. Stress regimes are usually transitional phases rather than clearly defined crisis periods. They often develop gradually and may include mixed signals, such as rising volatility, weaker returns, and changing macro-financial conditions. For risk management, some overprediction of stress may be acceptable, especially for investors who prefer to reduce exposure before market conditions deteriorate further. Therefore, the recall obtained for the stress regime supports the use of the model as an early-warning tool, even though its stress signals are not always precise.

The weakest results are observed for the panic regime. In the test sample, none of the seven panic observations are correctly classified, which leads to zero precision and recall for this state. The confusion matrix shows that panic episodes are mostly classified as calm or stress regimes.

This result mainly reflects a statistical limitation rather than a failure of the overall framework. Panic regimes are rare by nature, and the test sample contains very few such observations. As a result, the model has limited information from which to learn the specific patterns of extreme market episodes. In addition, cross-validation tends to favour performance on the more frequent classes unless the model is explicitly adjusted to penalise mistakes in rare regimes.

This difficulty is common in crisis prediction. Extreme events are infrequent, especially in emerging markets, but they have high economic importance. For this reason, they are difficult to forecast with standard multiclass classification methods. The results therefore suggest that future extensions should consider imbalance-aware techniques, rather than rejecting the HMM-XGBoost framework itself.

Despite the weak classification of panic episodes, the final feature set provides useful economic insights. The most important predictors include domestic realised volatility, BET index changes, sovereign yield curve indicators, commodity price movements, and domestic

macroeconomic variables related to inflation, unemployment, and exports. This shows that regime transitions in the Romanian equity market are influenced by both global financial conditions and domestic macro-financial factors.

The importance of commodity prices and yield curve variables suggests that external shocks and broader financial conditions play a key role in regime changes. At the same time, the relevance of domestic macroeconomic indicators shows that country-specific fundamentals also matter, either by amplifying or reducing the impact of global shocks.

Overall, the results show that the hybrid framework provides economically meaningful regime forecasts. It performs especially well in identifying calm market conditions and offers useful early-warning signals for stress regimes. However, the prediction of panic episodes remains difficult, mainly because these events are rare and underrepresented in the sample.

Conclusion

This paper examined whether market regimes in emerging European market can be identified and forecast using a hybrid econometric and ML framework, with an application to the Romanian BET index. The results of the paper provide clear answers to the research questions set in the introduction part.

First, the findings show that the Romanian equity market can be divided into 3 distinct meaningful regimes. The three-state Gaussian HMM separates market conditions into calm, stress and panic regimes, based on differences in realised volatility. This approach supports the idea that regime-based models describe emerging equity markets more accurately than linear models that assume a single market regime.

Second, the results show that next-day market regimes can be forecast with strong accuracy when macro-financial information is used in a nonlinear ML model. The XGBoost classifier performs well in the test set, especially in identifying calm periods and detecting stress conditions. This suggests that macro, financial and commodities variables provide useful information about regime shifts.

Moreover, the analysis shows that combining HMM with XGBoost adds value compared with relying only on traditional regime-switching models. The HMM provides a clear structure for identifying regimes, while XGBoost improves forecasting accuracy by capturing nonlinear relationships between variables. Taking this into account, the proposed framework combines the interpretability of econometric models with the predictive strength of ML models, too.

The findings have practical implications for risk management and portfolio allocation in emerging markets. Calm regime forecasts can help to support decisions to maintain or increase risk exposure during stable periods. At the same time, stress regime forecasts can act as early warning signals, allowing investors to reduce exposure before market conditions go worse.

The study also has some limitations, weak prediction of panic regimes, caused by the very small number of panic observations in the test sample. Also, this is a common problem in tail-risk forecasting, since extreme events are rare but economically important. Other limitations come from the volatility-based labelling of regimes and from the assumption that the relationship between predictors and regimes remains relatively stable over time.

Future improvements of the actual research could improve the framework in several directions. Imbalance methodologies, cost-sensitive loss functions, or hierarchical classification models may help improve the predictive power of panic regimes. The use of higher-frequency data, liquidity indicators or other volatility measures could also provide better signals during extreme

conditions. Extending the analysis to other emerging market would make it possible to compare regime dynamics across countries.

Overall, the paper shows that a hybrid HMM–XGBoost framework can identify and forecast market regimes in an emerging equity market in a meaningful way. The model performs particularly well in stable periods and provides useful signals for stress episodes. By combining regime-switching methods with ML, the study offers both an empirical contribution to nonlinear risk modelling and practical insights for risk management in emerging European financial markets.

Acknowledgments: During the preparation of this work, the authors utilized ChatGPT 5.4 to enhance readability and language clarity. Following the use of this tool, the authors meticulously reviewed and revised the content as necessary and assumes full responsibility for the final content of the publication.

References

- Ang, A., & Bekaert, G. (2002). International asset allocation with regime shifts. *Review of Financial Studies*, 15(4), 1137–1187.
- Anghel, B. I., & Lupu, R. (2025). Can Transition Risk Drive European Stock Market Predictions? An XGBoost Approach.
- Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 785–794.
- Christiansen, C., Kole, E., & van Dijk, D. (2012). Regime switching in stock return volatility. *Journal of Financial Econometrics*, 10(4), 449–470.
- Dueker, M. (1997). Markov switching in GARCH processes and mean-reverting stock - market volatility. *Journal of Business & Economic Statistics*, 15(1), 26–34.
- Égert, B., & Kočenda, E. (2011). Time-varying synchroniz ation of European stock markets. *Empirical Economics*, 40(2), 393–407.
- Esmailzade, H., Khosravi, A., & Nahavandi, S. (2024). Gradient boosting for nonlinear demand and price forecasting. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 54(1), 223–234.
- Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, 57(2), 357–384.
- Hamilton, J. D., & Susmel, R. (1994). Autoregressive conditional heteroskedasticity and changes in regime. *Journal of Econometrics*, 64(1–2), 307–333.
- Kole, E., & van Dijk, D. (2017). How to identify and forecast regimes in financial markets. *Journal of Applied Econometrics*, 32(2), 370–393.
- Kopa, M., & D'Ecclesia, R. L. (2020). Machine learning for downside risk forecasting in emerging stock markets. *European Journal of Operational Research*, 287(3), 1075–1088.
- Kremer, M. (2018). Time-varying systemic risk: Evidence from a Markov switching approach. *Journal of Financial Stability*, 38, 13–29.
- Kritzman, M., Page, S., & Turkington, D. (2012). Regime shifts: Implications for dynamic strategies. *Financial Analysts Journal*, 68(3), 22–39.
- Marcucci, J. (2005). Forecasting stock market volatility with regime-switching GARCH models. *Studies in Nonlinear Dynamics & Econometrics*, 9(4).
- Moore, T., & Wang, P. (2007). Volatility in stock returns for new EU member states. *Journal of International Financial Markets, Institutions & Money*, 17(4), 389–406.

- Nystrup, P., Madsen, H., & Lindström, E. (2017). Regime- based versus static asset allocation: Letting the data speak. *Journal of Portfolio Management*, 43(6), 56–71.
- Stoica, O., & Diaconășu, D. E. (2014). Nonlinear dynamics of the Romanian stock market. *Economic Computation and Economic Cybernetics Studies and Research*, 48(4), 113–129.
- Tan, M., Wang, L., Xu, Z., & Zhang, Y. (2023). Forecasting financial time series using gradient boosting frameworks. *Expert Systems with Applications*, 213, 118879.
- Voineagu, V., Dumitrescu, B., & Roman, A. (2019). Structural changes and volatility transmission in the Romanian stock market. *Romanian Journal of Economic Forecasting*, 22(3), 23–41.
- Walid, C., Chaker, A., Masood, O., & Fry, J. (2011). Stock market volatility and exchange rates in emerging countries. *Economic Modelling*, 28(1–2), 183–193.
- Xu, Y., Zhang, X., & Li, J. (2022). Stock price prediction based on XGBoost and advanced feature engineering. *Applied Soft Computing*, 114, 108078.