Universal Library of Business and Economics Research Article

ISSN: 3064-9951 | Volume 3, Issue 1

Open Access | PP: 64-69

DOI: https://doi.org/10.70315/uloap.ulbec.2026.0301008

Forecasting Stock Returns Using State-Space Models and Time-Varying Parameters

Abdelmadjid Laouedj

Quantitative Researcher, Jersey City, USA.

Abstract

The article examines the application of state-space models with time-varying parameters for forecasting stock returns in a non-stationary, regime-switching financial environment. The relevance of the study stems from the fact that observed returns are a hybrid of rare structural breaks and day-to-day noise, which results in traditional regression schemes with fixed coefficients rapidly deteriorating out of sample, especially in sectors characterised by high regulatory and technological uncertainty, such as pharmaceuticals. The aim of the study is to provide a conceptual and methodological justification for a forecasting approach that, from the outset, treats parameter instability as a norm rather than an exception. The scientific novelty lies in a coherent specification of a state-space model in which expected returns, factor loadings and, in extensions, volatility and market regimes are treated as latent states evolving according to probabilistic laws and estimated via recursive filtering, with explicit control of uncertainty, coefficient shrinkage, and time-based validation of the end-to-end pipeline under transaction cost constraints. The main conclusions demonstrate that this approach shifts the focus from guessing a number to constructing an adaptive, properly calibrated representation of risk premia and the range of outcomes, where the key resource is not a universal formula but a procedure for robustly updating inferences in a changing environment. The article is intended for researchers and practitioners in financial markets, as well as pharmaceutical managers and participants in continuing education programmes who rely on return forecasts when assessing risk and the robustness of decisions.

Keywords: *Stock Return Forecasting, State-Space Models, Time-Varying Parameters, Kalman Filter.*

INTRODUCTION

Forecasting stock returns is important not only for professional market participants but also for industry specialists who must make decisions amid financial uncertainty. For pharmaceutical professionals engaged in continuing education, this topic is immediately applicable in several respects: assessing the resilience of companies and sectoral indices; understanding the cost of capital when financing research and development; interpreting market reactions to regulatory events and clinical outcomes; and basic financial literacy for managing personal savings. The key difficulty, however, is that the observed return is a mixture of rare structural shifts and everyday noise; therefore, a good forecast must be able to distinguish changes in fundamental expectations from random fluctuations.

Attempts to predict returns using models with fixed coefficients typically run up against non-stationarity: the relationship between returns and predictors or factors changes over time, and the predictors themselves may work only in certain market phases. Empirical studies emphasise that the economic significance of predictability depends on the horizon, volatility, and parameter uncertainty, and that apparently stable relationships often degrade out of sample (Cederburg et al., 2023). An additional problem is predictor instability and structural breaks, which can lead to the same regression specification yielding opposite conclusions across different time segments (Calonaci et al., 2022). In an educational context, this is fundamental: it is less important for a learner to memorise a universal formula than to master a method that treats variation in the financial environment as the norm rather than the exception.

A natural response to this situation is provided by state-space models with time-varying parameters: instead of assuming constant coefficients, they treat expected returns and factor loadings as latent states that evolve according to a probabilistic law. This formulation allows updating estimates as data arrive and automatically adapting to regime shifts by using recursive filtering (in particular, the Kalman filter) to separate signal from noise (Neslihanoglu

Citation: Abdelmadjid Laouedj, “Forecasting Stock Returns Using State-Space Models and Time-Varying Parameters”, Universal Library of Business and Economics, 2026; 3(1): 64-69. DOI: https://doi.org/10.70315/uloap.ulbec.2026.0301008.

et al., 2021). The practical meaning of this approach is not to guess each subsequent return, but to obtain a robust, continuously updated assessment of how the structure of risk and risk premia changes over time, which is particularly valuable under sectoral focus and limited transferability of results across periods.

MATERIALS AND METHODS

The study draws on a corpus of contemporary empirical and methodological sources that document the central problem of return forecasting: the non-stationarity of relationships and the temporal drift of predictors. As its theoretical foundation, it draws on work showing that the economic significance of predictability depends on horizon and parameter uncertainty, and that patterns that appear stable often deteriorate out of sample (Cederburg et al., 2023). In addition, it incorporates the proposition of predictor instability and structural change, whereby the same regression specification may yield opposite conclusions across different time intervals (Calonaci et al., 2022). At the methodological level, the study follows the probabilistic framework of state-space models, in which observed returns are treated as noisy measurements and expected returns and factor sensitivities as latent states; this allows parameter variability to be formally embedded in the architecture of the forecast itself (Chan & Kroese, 2025).

Methodologically, the study implements a state-space formulation with time-varying parameters as a compromise between the interpretability of factor logic and the need to adapt to regime shifts. Within this scheme, the parameters of the regression component are treated as latent processes (for example, a drift resembling a random walk or a more inert dynamic), and estimates are updated by recursive filtering, primarily via the Kalman filter, to separate the signal of fundamental expectations from everyday noise while simultaneously obtaining a quantitative measure of estimation uncertainty (Neslihanoglu et al., 2021). For interpretation of results and comparison with alternatives, the research design incorporates the context of competing forecasting strands: the high accuracy of nonlinear methods is achievable through complex interactions among predictors, but is sensitive to regime shifts (Gu et al., 2020; Suárez-Cetrulo et al., 2023), whereas models of conditional heteroscedasticity are primarily useful for describing the dynamics of risk and volatility clustering, which is important for realistic assessment of predictive bands (Fang & Han, 2024).

RESULTS AND DISCUSSION

The most common line of return forecasting begins with linear predictive regressions and factor models, in which returns are explained by the market premium and a set of systematic factors, with parameters assumed to be constant. From a teaching perspective, this is convenient: coefficients can be interpreted as the contributions of risk sources, and the capital asset pricing model, together with the Fama–French multifactor constructions, provides a basic language for discussing risk premia. Empirically, however, this simplicity often proves fragile: the quality of factor explanation varies markedly across markets and periods, and in stress phases, parameters and factor significance can shift to such an extent that the conventional interpretation ceases to hold. These effects are illustrated, for example, in studies showing that multifactor models behave differently across developed and emerging markets and deteriorate during crisis episodes (Mosoeu & Kodongo, 2020).

Another branch comprises time-series models that relate returns to their own dynamics: autoregressions and schemes combining autoregressive and moving-average components, where emphasis is placed on dependence of current returns on past values and random disturbances. Their practical strength tends to manifest not so much in stable forecasting of mean returns as in describing variability and uncertainty: models of conditional heteroscedasticity and their extensions are well suited to the fact that risk condenses and dissipates in clusters. In applied tasks, this means that such structures often contribute more to assessing the range of possible outcomes than to accurately predicting the direction and magnitude of future returns. For this reason, contemporary work increasingly combines parameter dynamics with variance dynamics in order to model structural shifts while maintaining sensitivity to changes in risk (Fang & Han, 2024).

Machine learning models broaden the space of possibilities by incorporating nonlinearities and predictor interactions, as well as by working with high-dimensional sets of firm characteristics and macroeconomic variables. A large comparative study shows that such methods can yield substantial gains in forecasting risk premia relative to traditional linear approaches, precisely because they capture complex combinations of signals that a fixed-form regression overlooks (Gu et al., 2020). This flexibility, however, has a flip side that is crucial in educational settings: when distributions shift, and regimes change, even highly sophisticated algorithms begin to learn on a past that no longer exists, and therefore require disciplined drift control, adaptive training windows and explicit consideration of regime structure in the data. This vulnerability is discussed in literature reviews on forecasting under regime change, as well as in recent studies that frame the trade-off between model complexity and environmental non-stationarity as the central problem of return prediction (Suárez-Cetrulo et al., 2023).

The intuition of state-space models allows these approaches to be combined within a single framework: the observations are measured returns, and the states are latent quantities that govern their formation and change over time. As stated, one may consider expected returns as a smoothly evolving component, factor loadings as time-varying parameters, and latent variance as an internal indicator of the current risk level; the task then reduces to sequentially refining these states as new data arrive (Chan & Kroese, 2025). Filtering and smoothing algorithms (in particular, the Kalman filter) provide not only point estimates, but also a probabilistic description of uncertainty, which is essential for decision-making and especially useful for pharmaceutical specialists who need to understand why relationships break down at regime boundaries and how to update inferences correctly when new events occur.

The basic construct linking the preceding review of approaches to the practical teaching task rests on a regression model in which returns are viewed as outcomes of factor influences and random noise, but factor contributions are not fixed once and for all. The observed return at each moment is represented as the sum of an intercept, weighted factors and a measurement disturbance; at the same time, the weights themselves are understood as latent quantities. This formulation is useful in the educational environment of pharmaceutical specialists because it preserves the intuitive appeal of factor-based explanations while removing the unrealistic requirement of invariant relationships, which is often violated in response to news about regulatory decisions, clinical results, and revisions to market expectations.

The dynamics of the weights are specified by a separate state equation that describes the evolution of the coefficients over time. In the simplest case, each coefficient changes smoothly as a random walk, that is, the new level equals the previous one plus a small random increment; in a more inertial variant, the coefficient tends to revert to its own past value, as in a process with memory in which abrupt jumps are dampened. These variants differ in how they permit adaptation: a random walk allows faster adjustment to structural breaks, whereas a process with memory more strongly prevents the model from mistaking noise for signal. The choice of such a mechanism is, in essence, a choice of how strongly one believes in the smoothness of changes in market regularities over horizons relevant for practice.

Hyperparameters play a key role by governing the relative freedom of the coefficients and the level of noise. The variance parameters in the state equation determine the speed of drift: the higher they are, the more boldly the model allows factor contributions to change and the faster it can respond to regime shifts; the lower they are, the more the weights resemble traditional constant coefficients. The observation noise parameters determine which portion of short-term fluctuations is treated as measurement noise and which as a useful signal. Taken together, these settings define a trade-off between adaptation and robustness that is especially critical in applied educational tasks: excessive flexibility leads to overfitting, whereas excessive rigidity results in belated responses to new conditions.

Estimation and filtering in this model are based on sequential updating of latent states as new data become available. The Kalman filter iteratively estimates current coefficients and expected returns, updating the estimate of uncertainty (the model's confidence in the current state) as it corrects for the observed return. This uncertainty estimate is not merely a by-product, since it provides critical information about the strength of the signal being observed, and the degree to which stronger risk dislike or more cautious behavior is warranted. In contrast, Kalman smoothing leverages the entire path and can better estimate the past state, filtering residual noise and spikes, and allowing one to see when factor sensitivities changed and by how much.

Formulating the forecasting problem naturally follows from the model's probabilistic nature. Forecasts can be made one step ahead or multiple steps ahead, but as the horizon lengthens, the role of state and noise uncertainty increases rapidly, so that a point forecast becomes less informative than a forecast expressed as a distribution of possible outcomes. In an educational framework, this helps reorient expectations: the goal is not to guess a number, but to assess the range, asymmetry and probability of adverse scenarios. Furthermore, the target variable can be redefined to better match the practical decision: instead of the return itself, one may forecast the sign of the change, the excess of return over a selected benchmark, or the risk of rare but severe losses, which is closer to real-world tasks of risk management and robustness of financial plans. The baseline model's algorithm is shown in Figure 1.

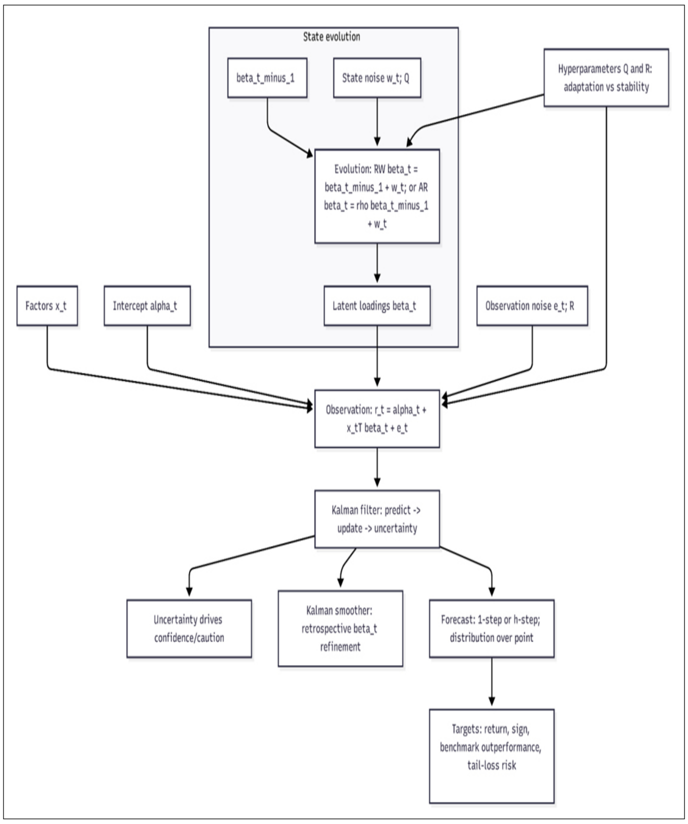

Fig. 1. Algorithm of the baseline model

Once the baseline model has already separated latent expected returns from noise and permits coefficient drift, the next step is to refine how risk itself changes. In reality, the volatility is also not stationary: periods of turbulence alternate with quiet periods, and assigning a single parameter describing the dispersion of the distribution of returns for the whole of this period would give too much confidence to the forecast. Therefore, in the extended version, the observation error variance becomes a latent state with its own dynamics. It reflects the breathing of uncertainty; the model learns to predict not only the expected return but also the return corridor at a given moment. In pharmaceuticals, news about clinical trials, regulatory proceedings, or the risk of a patent falling is often associated with changes to the mean forecast, the probability distribution, and the extent of possible deviations from the mean forecast.

Another realistic complication is that the market behaves as a system in which different regimes operate, and transitions between them do not follow a fixed schedule. Within this logic, parameters and risks depend on a latent regime state that switches probabilistically: for example, one regime may correspond to a phase of elevated uncertainty and dominance of defensive factors, and another to a normalisation phase in which familiar dependencies reappear. The state-space model is then augmented with a switching mechanism, and filtering becomes the simultaneous estimation of coefficients, volatility, and the probability that a particular regime is currently in effect. This dispels the illusion of a single universal rule and makes it possible to formalise what is often described in practice as the market has changed, without reducing everything to ex post explanations.

Finally, contemporary applied problems almost always involve many predictors and not a single return series, but a set of interrelated time series. To prevent overfitting, coefficient shrinkage is introduced, automatically suppressing weak and unstable features while leaving only genuinely useful signals in the dynamics. In parallel, a multivariate formulation allows for the joint dynamics of several stocks, sectors, or portfolios, preserving common latent components and mutual dependencies. This is important when the pharmaceutical sector reacts to common shocks, while individual firms differ in their sensitivity to them. As a result, the extended model becomes not merely more complex, but more disciplined: it simultaneously incorporates time variation in risk, market regime structure, high factor dimensionality and structural linkages within the portfolio and sectoral landscape. The adapted model is shown in Figure 2.

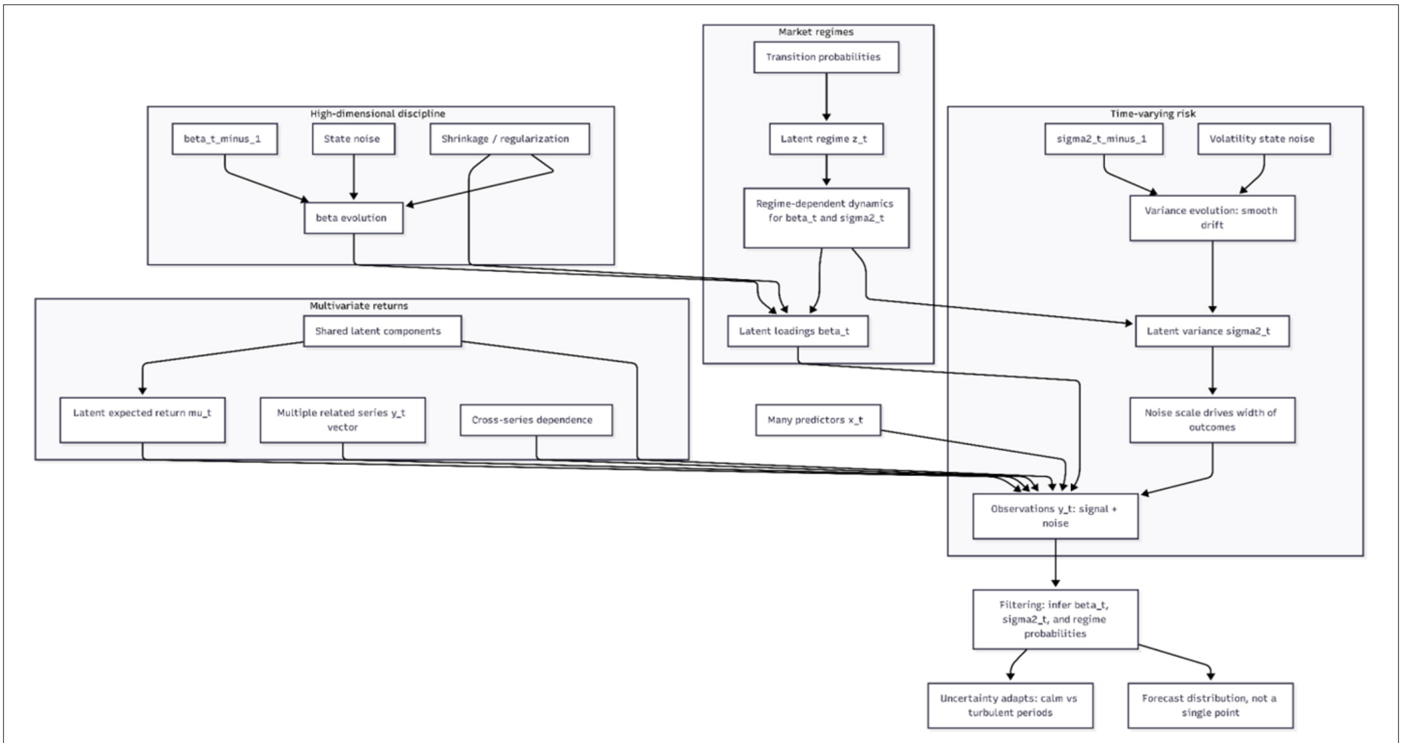

Fig. 2. The algorithm of the adapted model

The practical cycle begins with a considered choice of data, with returns as the target variable and predictors reflecting different layers of causality: market and sectoral factors, macroeconomic conditions, and features extracted from the price series themselves. For pharmaceutical applications, it is also important to carefully account for the calendar of events that alter expectations and risk, as these events often trigger regime shifts and reprogram factor sensitivities. Target preparation involves aligning predictors to the same time grid. The technique includes filling in gaps for missing predictive values, handling outliers, scaling, and applying time-shifting to use only information available at the time of the prediction. The key to the quality of this preprocessing is the alignment of disparate data frequencies. For example, combining daily prices with low-frequency macroeconomic indicators produces spurious signals, leading to overconfidence and poor model performance.

Model evaluation proceeds in chronological order by repeatedly estimating and forecasting using rolling or expanding windows of available data, and by measuring performance on future observations to avoid look-ahead bias. Such a walk-forward exercise is needed not only to compare accuracy, but also to monitor robustness: in models with time-varying parameters, it is important to observe how quickly they adapt and whether adaptation degenerates into a delayed reaction to noise. The final part of the pipeline translates the forecast into a trading or management rule that accounts for transaction costs and rebalancing frequency. Even the best probabilistic forecast will be useless if rebalancing is frequent. To reduce noise, the portfolio turnover is constrained, signals are smoothed, and the across-asset and within-asset persistence and profitability are evaluated under realistic costs and liquidity constraints. The end-to-end pipeline architecture is illustrated in Figure 3.

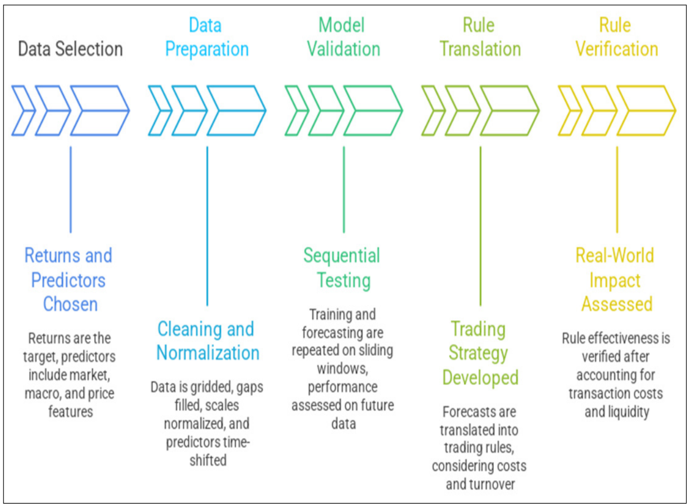

Fig. 3. Architecture of end-to-end pipeline

Quality assessment in such models must capture not only the proximity of forecasts to realised returns, but also the practical usefulness of forecasts and the honesty of stated uncertainty. Accordingly, one simultaneously considers forecast errors in magnitude, the share of correctly predicted directions, the quality of probabilistic distributions, and the ultimate usefulness in risk–return terms when applied in a strategy. In addition, it checks whether the dynamic coefficients have stabilized and not degenerated in reflections of pure noise, whether jumps occurred, which indicate a structural break, and whether the sensitivities are plausible with respect to the market regimes and sectoral events. Sensitivities with respect to the drift and noise parameters are analyzed separately, as this is where the adaptation/overfitting boundary lies. An abrupt change in results with small changes indicates a need for stronger regularization, a simpler specification, or switching to a stronger formulation with regimes and coefficient shrinkage.

CONCLUSION

Within the framework of this article, stock return forecasting is conceptualised not as an attempt to capture each subsequent value, but as the task of distinguishing in the observed return that which genuinely alters fundamental expectations from that which remains everyday noise. Against the background of empirically documented non-stationarity of predictive relationships and the out-of-sample degradation of elegant dependencies, models with fixed coefficients prove conceptually vulnerable: factors operate in a fragmented manner, parameters drift, crisis episodes alter the meaning of familiar interpretations, and predictability itself becomes a function of horizon and uncertainty. Therefore, state-space models with time-varying parameters naturally emerge as the methodological core, in which expected returns, factor sensitivities and (in extensions) risk are treated as latent states evolving probabilistically and refined via recursive filtering. This approach disciplines forecasting by forcing simultaneous consideration of adaptation to regime shifts and the inevitable uncertainty in estimates, transforming the forecast of a number into a continuously updated representation of the structure of risk premia and the confidence in this representation.

In applied and educational settings, the value of this formulation is particularly transparent: it preserves the intuitive clarity of the factor language while removing the unrealistic requirement of constant relationships, precisely the requirement that most often breaks down at regime boundaries, during news events, and during revisions of expectations. Parametric dynamics (from random walks to more inertial, memory-based processes) and hyperparameters governing drift speed and noise magnitude define a delicate balance between flexibility and robustness: an overly mobile model readily mistakes noise for signal, while an overly rigid one lags when the market has already changed its logic. Extensions in which volatility becomes a latent state, regime switches are introduced, coefficient shrinkage is applied, and a multivariate structure is used do not merely complicate the description, but enhance its honesty with respect to the data: the market breathes risk, switches between phases and does not forgive high-dimensional naivety without regularisation. The resulting methodological line is closed with an end-to-end pipeline that includes time-based validation, robustness control, translation of forecasts into a managerial rule that accounts for costs, and evaluation of not only accuracy but also practical usefulness and uncertainty calibration.

REFERENCES

1. Calonaci, F., Kapetanios, G., & Price, S. (2022). Stock Returns Predictability with Unstable Predictors. *SSRN*. https://dx.doi.org/10.2139/ssrn.4007703

2. Cederburg, S., Johnson, T. L., & O'Doherty, M. S. (2023). On the Economic Significance of Stock Return Predictability. *Review of Finance*, *27*(2), 619–657. https://doi.org/10.1093/rof/rfac035

3. Chan, J. C. C., & Kroese, D. P. (2025). State Space Models. In: Statistical Modelling and Computation. *Springer Texts in Statistics*. https://doi.org/10.1007/978-1-0716-4132-3_13

4. Fang, Z., & Han, J. (2024). GARCH Model in Volatility Forecasting and Option Pricing. *Computational Economics*, *66*, 3637–3657.

5. Gu, S., Kelly, B., & Xiu, D. (2020). Empirical Asset Pricing via Machine Learning. *The Review of Financial Studies*, *33*(5), 2223–2273. https://doi.org/10.1093/rfs/hhaa009

6. Mosoeu, S., & Kodongo, O. (2020). The Fama-French five-factor model and emerging market equity returns. *The Quarterly Review of Economics and Finance*, *85*, 55–76. https://doi.org/10.1016/j.qref.2020.10.023

7. Neslihanoglu, S., Bekiros, S., & McColl, J. (2021). Multivariate time-varying parameter modelling for stock markets. *Empirical Economics*, *61*, 947–972.

8. Suárez-Cetrulo, A. L., Quintana, D., & Cervantes, A. (2023). Machine Learning for Financial Prediction Under Regime Change Using Technical Analysis: A Systematic Review. *International Journal of Interactive Multimedia and Artificial Intelligence.* https://doi.org/10.9781/ijimai.2023.06.003

Copyright: © 2026 The Author(s). This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.