

AI Adoption and Firm Innovation: Evidence from Firm-Level Predictive Models

Meral Kagitci¹ and Şener Ali²

¹⁾²⁾ *Bucharest University of Economic Studies, Bucharest, Romania.*

E-mail: meral.kagitci@ase.ro; E-mail: alisener20@stud.ase.ro

Please cite this paper as:

Kagitci, M., & Ali, Ş. (2026). AI Adoption and Firm Innovation: Evidence from Firm-Level Predictive Models. In C. Vasiliu, D.-C. Dabija, A. Tziner, D. Pleşea, & V. Dinu (Eds.), *Proceedings of BASIQ International Conference on New Trends in Sustainable Business and Consumption, 12th Edition* (pp. 363–370). Editura ASE.

DOI 10.24818/BASIQ/2026/12/060

Abstract

Building on financial and firm-performance literature that uses firm-level indicators to model business outcomes, this article examines whether artificial intelligence (AI) adoption indicators can predict subsequent firm-level innovation outcomes. It addresses the growing need to understand AI adoption as a dynamic firm-level economic process connected to productivity, investment behaviour, R&D activity, and patent-based innovation. The study constructs a harmonized firm-year panel by integrating AI employment and AI-related software-use data with firm financial statements, industry classifications, productivity measures, patent information, and trademark registrations. The modelling sample focuses on the top 50 ever-adopting firms. Innovation is measured primarily through total patent counts. Three predictive approaches are compared using a chronological training-validation-test split: a dynamic firm fixed-effects model, pooled ridge regression, and XGBoost. The models incorporate lagged AI variables, lagged controls, prior innovation outcomes, moving-average terms, and centered time controls. The results show that pooled ridge regression achieves the strongest predictive performance on the held-out test sample, with the lowest RMSE and MAE and the highest R². XGBoost ranks second, outperforming the firm fixed-effects benchmark but not the regularized linear model. The findings suggest that coefficient shrinkage is particularly effective when modelling correlated lagged predictors in firm-level innovation data. The article contributes by offering a leakage-aware predictive comparison of AI adoption and patent-based innovation using firm-level panel data, rather than claiming a causal effect. The findings may help firms, analysts, and policymakers in evaluating how AI employment, AI software use, R&D intensity, investment, and prior innovation histories can inform innovation forecasting and strategic decision-making.

Keywords

AI adoption; firm-level innovation; panel data methods; predictive modelling

DOI: 10.24818/BASIQ/2026/12/060

Introduction

Artificial intelligence (AI) has become an increasingly important element of firm strategy, investment, and operational transformation. Recent business evidence suggests that organizations are not only experimenting with AI tools but are also incorporating them into core functions such as marketing, IT operations, supply-chain management, decision support, and product-related activities. These developments make AI adoption relevant not only as a technological phenomenon but also as a firm-level economic process whose effects may appear through changes in productivity, investment behaviour, and innovation activity.

A central issue in this context is whether more intensive AI adoption is associated with more predictable subsequent innovation outcomes. This question is important because AI adoption may affect innovation through several channels. AI-skilled employment can increase a firm's capacity to develop, apply, and recombine technical knowledge. AI-related software use may support experimentation, automation, and the processing of large quantities of operational or market information. At the same time, innovation outcomes are rarely immediate. They often depend on prior capabilities, accumulated R&D activity, firm size,

investment intensity, and the persistence of earlier inventive activity. For this reason, a meaningful empirical analysis must treat AI adoption as a dynamic firm-level process rather than as a contemporaneous or purely descriptive indicator.

The aim of the study is to evaluate whether lagged AI adoption indicators, combined with firm-level financial, industry, productivity, patent, and trademark information, can predict subsequent firm-level innovation outcomes. The analysis constructs a firm-year panel by harmonizing raw sources using a common firm identifier and year index. AI adoption is measured through variables such as AI employment, AI employment share, AI knowledge employment, and AI-related software-use indicators. Innovation is measured primarily through patent-based outcomes, including total patent counts, with additional innovation measures aggregated from patent and trademark information. The modelling sample focuses on the top 50 ever-adopting firms selected using information available no later than the validation period, thereby reducing the risk of future information entering the model-building process.

Methodologically, the paper compares three predictive approaches: a dynamic firm fixed-effects model, pooled ridge regression, and XGBoost. The fixed-effects specification accounts for persistent unobserved firm heterogeneity and includes lagged AI variables, lagged controls, lagged innovation outcomes, moving-average terms, and centered time controls. Ridge regression applies coefficient shrinkage to the same high-dimensional lagged feature structure, which is appropriate where predictors are correlated. XGBoost adds a nonlinear benchmark capable of capturing interaction and threshold patterns among AI adoption, firm characteristics, and prior innovation. The models are evaluated through a deterministic temporal split, with training, validation, and test periods separated chronologically.

The main contribution of the paper is therefore not to claim a causal estimate of AI adoption on innovation, but to provide a focused predictive comparison of alternative firm-level modelling strategies. By aligning AI adoption indicators with lagged controls and patent-based outcomes, the study offers evidence on how different model classes perform when predicting innovation among AI-adopting firms. The remainder of the paper reviews relevant literature on firm-level predictive modelling and performance determinants, presents the data-processing and modelling strategy, reports the comparative predictive results, and concludes with implications and limitations.

1. Literature review

The literature relevant to this study covers several connected areas, including firm-level prediction, financial and organizational determinants of performance, and the role of R&D and innovation-related indicators in empirical modelling. Although the studies reviewed in the paper do not all examine AI adoption directly, they provide a methodological basis for constructing a firm-year predictive framework in which AI-related variables are evaluated alongside financial, productivity, and innovation controls. These studies justify the inclusion of firm size, profitability-related controls, investment intensity, R&D activity, lagged outcomes, and model diagnostics in a predictive panel design, while the empirical contribution of the present study is to apply this framework specifically to AI adoption indicators and patent-based innovation outcomes.

One prominent research trajectory demonstrates that structured firm-level data can support prediction beyond descriptive financial analysis. Zhao et al. (2023) compare linear regression, elastic net, random forest, and neural network methods for forecasting stock returns using financial statement indicators. Their study is relevant because it shows that machine learning methods may extract predictive information from firm-level variables more effectively than standard linear models, particularly when indicators related to profit, cash, and growth are included. Tsai et al. (2023) similarly use financial indicators within machine learning-based portfolio construction, comparing random forest, feedforward neural networks, financial graph attention networks, and gated recurrent units. These studies support the methodological choice to compare traditional and machine learning models rather than rely on a single specification.

Beyond predictive modelling, research also examines how firm characteristics and financial ratios are associated with firm value or profitability. Jihadi et al. (2021) use regression analysis to study how profitability, leverage, liquidity, firm size, and corporate social responsibility relate to firm value. Sucuahi and Cambarihan (2016) find that profitability is positively associated with Tobin's Q in a sample of Philippine firms, while Sukmawardini and Ardiansari (2018) examine the relationship between profitability, liquidity, leverage, dividend policy, ownership structure, and price-to-book value. These studies are relevant to the present paper because they support the inclusion of firm-level controls such as size, investment intensity, and past performance.

Another relevant strand of literature examines firm-level factors associated with business performance across different sectors and country contexts. Mijoč (2024) analyzes IT firms in Croatia using ROA and ROE as dependent variables and includes liquidity, debt, financial stability, size, sales growth, EBIT, and lagged profitability indicators. Vintilă and Nenu (2015) study Romanian firms using ROE, ROA, and share price as performance measures, while Bichir et al. (2024) examine technology-sector firms in the United States using ROIC, ROE, and ROA. Lan and Nguyen (2020) analyze firms across six industries and include firm size, capital adequacy, liquidity, leverage, and solvency. These studies show that financial performance is shaped by multiple firm-specific variables and that results often vary by sector, country, sample size, and model specification. This reinforces the need for a modelling design that includes several control variables rather than isolating AI adoption from the broader firm context.

Several reviewed studies also highlight methodological issues that are directly relevant to the present paper. Alarussi and Alhaderi (2018), Daruwala (2023), and Zavalii et al. (2025) use regression-based approaches and diagnostic tests such as multicollinearity, heteroskedasticity, autocorrelation, normality, or Hausman tests. Their findings are often mixed, especially regarding leverage, liquidity, firm size, and other financial indicators. Pervan et al. (2019) extend the analysis of profitability by applying the Generalized Method of Moments and incorporating macroeconomic variables. These works suggest that firm performance models are sensitive to specification, data structure, and temporal dependence. For the present study, this supports the use of lagged predictors, moving-average terms, and a chronological training-validation-test split.

Radenović et al. (2023) examine R&D expenditures, R&D intensity, and return on R&D capital in relation to profitability, while Erdogan and Yamaltdinova (2019) analyze R&D intensity, firm size, and leverage in relation to ROA and ROE. Although these studies measure profitability rather than patent-based innovation, they provide support for treating R&D intensity and investment-related variables as important controls when modelling firm outcomes. In the present paper, this logic is extended to innovation prediction by including lagged R&D intensity, capital expenditure intensity, prior patent outcomes, and moving averages of patent activity.

Taken together, the reviewed literature supports the use of firm-level quantitative indicators, lagged controls, and comparative modelling strategies, but it also clarifies a specific gap concerning AI adoption and patent-based innovation. The studies centered on profitability, firm value, or portfolio performance are not treated here as evidence that AI adoption causes innovation; rather, they motivate the control structure and predictive modelling choices needed to evaluate AI adoption within a broader firm-year setting. The present study responds to this gap by constructing a leakage-aware firm-year panel in which AI employment and AI-related software-use indicators are linked to subsequent patent-based outcomes, and by comparing firm fixed-effects, pooled ridge regression, and XGBoost on the same lagged predictive feature set. This design reflects the dynamic nature of innovation and allows the predictive role of AI adoption to be assessed relative to established firm-level controls and prior patenting history.

2. Methodology

2.1. Data processing

The study utilizes a comprehensive firm-year panel dataset constructed by harmonizing firm-level AI adoption indicators, specifically AI employment and AI-related software use, with financial statements, productivity measures, and patent and trademark records. This integrated dataset focuses on the top 50 ever-adopting firms and spans a chronological timeframe split into distinct training, validation, and test periods ending in 2017. The data processing pipeline is organized as a two-stage sequence that transforms raw sources into a model-ready firm-year panel. Let the raw input be denoted by:

$$D^{raw} = \{D^{AI}, D^F, D^I, D^P, D^{Pa}, D^T\} \quad (1)$$

where D^{AI} contains firm-year information on AI employment and AI-related software use, D^F contains firm financial statements, D^I contains firm-level industry classifications, D^P contains productivity measures, D^{Pa} contains patent-based innovation measures, and D^T contains trademark registrations. The common identifier is the firm code i , and the time index is the year t . The first processing stage harmonizes these sources and constructs an analytical panel:

$$D^{(1)} = T_1(D^{raw}) \quad (2)$$

In this stage, identifier and year fields are standardized, observations with missing firm identifiers are removed, and duplicate firm-year records are collapsed to a single observation using the first available non-missing value within each key. Furthermore, the core economic variables used later in estimation are constructed. If A_{it} denotes AI employment and L_{it} total employment, the AI employment share is defined as A_{it}/L_{it} whenever $L_{it} > 0$. Similar ratios are computed for AI knowledge employment, AI software use per employee, and AI software variety per employee. Innovation outcomes are aggregated across product and process patent categories to obtain total patent counts, citation-weighted patent totals, and value-weighted patent totals; trademark registrations are similarly summarized at the firm-year level.

The second processing stage converts the broad panel into a leakage-aware modelling sample:

$$D^{proc} = T_2(T_1(D^{raw})) \quad (3)$$

This stage restricts the sample to years up to 2019, ranks firms using information available no later than the end of the validation period, and selects the top 50 ever-adopting firms. The ranking criterion used in the implementation is AI employment share. For the selected firms, the main outcome is defined as:

$$\psi_{it} = \log(1 + \widehat{Patent}_{it}) \quad (4)$$

where $\widehat{Patent}_{it}$ is the total patent count. Then, lagged AI predictors, lagged controls, and lagged or moving average terms are constructed. Concretely, the main AI regressor is the one-period lag of log AI employment, while an alternative AI regressor uses the one-period lag of AI employment share. Additional features include second and third lags. Model-specific sample flags retain only observations with complete information for the corresponding specification. The temporal split is deterministic: years up to 2012 form the training set, 2013–2015 the validation set, and 2016–2017 the test set, to ensure a leakage-aware predictive design that prevents future information from entering the model-building process. This chronological separation respects the path-dependent nature of innovation by requiring the models to predict subsequent outcomes based solely on prior firm history and lagged indicators.

2.2. Firm fixed-effects model

The first model is a linear firm fixed-effects specification designed for panel data with persistent cross-firm heterogeneity. For firm i in year t , the model can be written as:

$$\psi_{it} = \alpha_i + x'_{it}\beta + \varepsilon_{it} \quad (5)$$

where α_i captures time-invariant firm heterogeneity, x_{it} includes lagged log AI employment, lagged moving averages of AI intensity, lagged firm size, lagged R&D intensity, lagged capital expenditure intensity, lagged values of the patent outcome, lagged moving averages of the patent outcome, and centered linear and quadratic year terms, while ε_{it} is an error term. Thus, the specification is explicitly dynamic and allows innovation to depend both on past AI adoption and on the persistence of prior innovative activity.

Estimation proceeds using the within transformation, which removes the firm-specific intercepts before least-squares estimation. In compact form, the coefficient vector is obtained from

$$\min_{\beta} \sum_{i,t} [(\psi_{it} - \bar{\psi}_i) - (x_{it} - \bar{x}_i)' \beta]^2 \quad (6)$$

Predictions are then formed by combining the estimated coefficients with recovered firm effects. Because the specification includes lagged variables and moving-average terms, estimation is limited to firm-year observations with sufficient prior history to construct all predictors. This framework is appropriate for panel data because it accounts both for persistent unobserved differences across firms and for the dynamic nature of innovation over time.

2.3. Pooled ridge regression

The second model applies pooled ridge regression to the same predictive feature set. The conditional mean is specified as:

$$\hat{\psi}_i = \beta_0 + x'_i\beta \quad (7)$$

where x_i contains the full set of lagged AI variables, lagged controls, lagged patent outcomes, moving average terms, and centered time controls. Unlike the fixed-effects estimator, this model does not include a separate intercept for each firm. Instead, it treats the panel as a pooled sample and controls model complexity through coefficient shrinkage.

The parameters are estimated by minimizing the penalized least-squares objectives

$$\min_{\beta_0, \beta} \sum_{i=1}^N (\psi_i - \beta_0 - x_i' \beta)^2 + \lambda \sum_{j=1}^p \beta_j^2 \quad (8)$$

where λ is the regularization parameter. All predictors are standardized before estimation so that the penalty is comparable across variables measured on different scales. The penalty level is selected on the training sample from a broad grid of candidate values, after which the chosen model is evaluated on the validation sample and refit on the combined training and validation data. This specification is particularly useful in the present setting because the predictive design includes many correlated lagged regressors, for which ridge regularization can reduce estimation variance and improve out-of-sample stability.

2.4. Extreme gradient boosting (XGBoost)

The third model uses XGBoost. Let z_i denote the vector of numeric predictors augmented with firm indicator variables. The model maps inputs into predicted innovation outcomes through an additive ensemble of regression trees:

$$\hat{\psi}_i = \sum_{m=1}^M g_m(z_i) \quad (9)$$

where each $g_m(\cdot)$ is a decision tree and M is the number of boosting iterations. This structure allows the effect of AI adoption to vary nonlinearly with firm size, investment, prior patenting history, and firm identity.

The model is trained by minimizing a regularized squared-error loss, with hyperparameters tuned over the number of trees, tree depth, learning rate, row subsampling, column subsampling, and regularization terms. Candidate specifications are estimated on the training sample and compared on the validation sample using root mean squared error. Early stopping is used during tuning to determine the effective boosting length. The selected parameterization is then refit on the combined training and validation data before final evaluation on the test sample. Relative to the two linear models, XGBoost provides the most flexible approximation, as it can capture threshold effects, nonlinearities, and interaction patterns that may arise in the relationship between AI adoption and subsequent innovation.

2.5. Workflow overview

Figure no. 1 presents the full workflow, beginning with the integration of raw firm-level AI, financial, industry, productivity, patent, and trademark sources, continuing through the two-stage construction of the model-ready panel, and ending with the estimation of three alternative predictive models on a common dataset.




Figure no. 1. Workflow of the methodology

3. Results

The predictive performance of the three evaluated models is assessed using root mean squared error (RMSE), mean absolute error (MAE), and the coefficient of determination (R^2). Lower RMSE and MAE indicate smaller prediction errors, whereas a higher R^2 indicates a better fit to the variation in the observed outcome. Table no. 1 reports the comparative results based on the held-out test sample.

Table no. 1. Predictive performance of evaluated models on the test sample

Model	RMSE	MAE	R^2
Firm fixed-effects	0.7212	0.5114	0.7102
Pooled ridge regression	0.6266	0.3865	0.7812
XGBoost	0.6711	0.4073	0.7490

The firm fixed-effects model records an RMSE of 0.7212, an MAE of 0.5114, and an R^2 of 0.7102. These values indicate that the model captures a meaningful share of the variation in the transformed patent outcome, but its predictive errors remain relatively large compared with the alternative specifications. In particular, the MAE implies that the average absolute deviation between predicted and observed values is slightly above 0.51 in log-transformed outcome units. The R^2 value suggests that the model retains substantial explanatory power, which is consistent with the inclusion of firm-specific heterogeneity and dynamic predictors. Overall, the fixed-effects model provides a credible benchmark, but its out-of-sample predictive accuracy is the weakest among the three evaluated approaches.

The pooled ridge regression delivers the strongest performance. Its RMSE is 0.6266, its MAE is 0.3865, and its R^2 reaches 0.7812. Relative to the first model, both error measures are materially lower, indicating more accurate out-of-sample predictions, while the higher R^2 shows a better overall fit to the observed pattern of innovation outcomes. The empirical implication is that regularized linear pooling appears to handle the dataset particularly well, likely because shrinkage helps stabilize estimation when multiple lagged and averaged predictors are included simultaneously. From a results perspective, the ridge model combines the smallest prediction errors with the highest explanatory fit, making it the most effective specification on the test sample.

The XGBoost model yields an RMSE of 0.6711, an MAE of 0.4073, and an R^2 of 0.7490. These results place the model between the other two approaches. Compared with the firm fixed-effects model, XGBoost reduces both RMSE and MAE and increases R^2 , indicating superior predictive performance on the held-out data. At the same time, it does not outperform the pooled ridge regression, which attains lower prediction errors and a higher goodness-of-fit measure. The XGBoost results therefore suggest that allowing for nonlinearities and interaction effects improves predictive accuracy relative to the fixed-effects

benchmark, but in this application that added flexibility does not translate into the best overall out-of-sample performance.

A direct comparison across the three models shows a clear ranking on the test sample. Pooled ridge regression performs best overall, as it combines the lowest RMSE (0.6266), the lowest MAE (0.3865), and the highest R^2 (0.7812). The XGBoost model ranks second on all three criteria, with stronger predictive performance than firm fixed-effects but weaker performance than pooled ridge regression. The firm fixed-effects model has the highest prediction errors and the lowest R^2 , making it the least accurate of the evaluated models in out-of-sample prediction. Importantly, there is no trade-off across metrics in this comparison: the same model dominates on RMSE, MAE, and R^2 .

Conclusions

This paper examined whether firm-level indicators of artificial intelligence adoption can predict subsequent innovation outcomes. The study addressed the need to understand AI adoption not only as a technological trend, but also as a dynamic economic process connected to firm capabilities, investment behaviour, R&D activity, and accumulated innovative performance. By focusing on patent-based innovation outcomes, the analysis placed AI adoption within a broader firm-level framework in which prior innovation, financial characteristics, productivity-related variables, and industry information jointly shape predictive performance.

The paper proposed a leakage-aware predictive design based on a harmonized firm-year panel integrating AI employment and AI-related software-use indicators with firm financial statements, industry classifications, productivity measures, patent data, and trademark information. The empirical strategy compared three modelling approaches on a common lagged feature structure: a dynamic firm fixed-effects model, pooled ridge regression, and XGBoost. This design allowed the study to evaluate the predictive contribution of AI adoption indicators while accounting for temporal ordering, prior firm behaviour, and the persistence of innovation outcomes.

The results show that pooled ridge regression achieved the strongest out-of-sample performance on the held-out test sample, recording the lowest RMSE and MAE and the highest R^2 among the evaluated models. XGBoost improved upon the firm fixed-effects benchmark, indicating that nonlinear and interaction-based modelling can add predictive value, but it did not outperform the regularized linear specification. The findings therefore suggest that, in this setting, coefficient shrinkage is particularly effective for stabilizing predictions when multiple correlated lagged AI, financial, and innovation-related variables are included.

The contribution of the study lies in its predictive rather than causal contribution. It demonstrates that AI adoption indicators, when combined with firm-level controls and prior innovation histories, can support the forecasting of patent-based innovation outcomes among AI-adopting firms. For firms, the results suggest that AI-related workforce indicators and AI software-use measures should be interpreted together with R&D intensity, capital expenditure intensity, firm size, productivity, and prior patenting when assessing future innovation trajectories. For analysts, the findings indicate that leakage-aware temporal validation and regularized predictive models can provide a disciplined way to benchmark innovation forecasts in panels with correlated lagged predictors. For policymakers, the results suggest that monitoring AI adoption alongside complementary investment and innovation indicators can improve the interpretation of firm-level innovation capacity.

The analysis is subject to limitations consistent with its design. The focus on the top 50 ever-adopting firms and on patent-based innovation outcomes limits the scope of generalization, while the modelling strategy does not identify causal effects. Future research could extend the sample, examine alternative innovation measures such as citation-weighted patents or trademarks, compare additional modelling approaches, and investigate whether predictive relationships differ across industries, firm sizes, or AI adoption intensities.

References

- Alarussi, A. S., & Alhaderi, S. M. (2018). Factors affecting profitability in Malaysia. *Journal of Economic Studies*, 45(3), 442-458. <https://doi.org/10.1108/JES-05-2017-0124>.
- Bichir, A., Mihai, D. G., & Toader, C. I. (2024). Determinants of financial performance evidence from the leading us technology companies. *International Journal of Academic Research in Accounting, Finance and Management Sciences*, 14(3), Pages 1-15. <https://doi.org/10.6007/IJARAFMS/v14-i3/21877>.

- Daruwala, Z. (2023). Influence of financial leverage on corporate profitability: Does it really matter? *International Journal of Economics and Financial Issues*, 13(4), 37-46. <https://doi.org/10.32479/ijefi.14461>.
- Erdogan, M., & Yamaltdinova, A. (2019). A panel study of the impact of r&d on financial performance: Evidence from an emerging market. *Procedia Computer Science*, 158, 541-545. <https://doi.org/10.1016/j.procs.2019.09.087>.
- Jihadi, M. et al., 2021. The effect of liquidity, leverage, and profitability on firm value: Empirical evidence from indonesia. *The Journal of Asian Finance, Economics and Business*, 8(3), 423-431. <https://doi.org/10.13106/JAFEB.2021.VOL8.NO3.0423>.
- Lan, T. N. & Nguyen, C., 2020. The determinants of profitability in listed enterprises: A study from vietnamese stock exchange. *The Journal of Asian Finance, Economics and Business*, 7(1), 47-58. <https://doi.org/10.13106/JAFEB.2020.VOL7.NO1.47>.
- Mijoč, I. (2024). Determinants of profitability of the IT industry in Croatia. *Ekonomski vjesnik*, 37(1), 45-62. <https://doi.org/10.51680/ev.37.1.4>.
- Pervan, M., Pervan, I., & Ćurak, M. (2019). Determinants of firm profitability in the Croatian manufacturing industry: Evidence from dynamic panel analysis. *Economic Research-Ekonomska Istraživanja*, 32(1), 968-981. <https://doi.org/10.1080/1331677X.2019.1583587>.
- Radenović, T., Krstić, B., Janjić, I., & Jovanović-Vujatović, M. (2023). The effects of R&D performance on the profitability of highly innovative companies. *Strategic Management*, 28(3), 34-45. <https://doi.org/10.5937/StraMan2200034R>.
- Sucuahi, W., & Cambarihan, J. M. (2016). Influence of profitability to the firm value of diversified companies in the philippines. *Accounting and Finance Research*, 5(2), p149. <https://doi.org/10.5430/afr.v5n2p149>.
- Sukmawardini, D. & Ardiansari, A., 2018. The Influence of Institutional Ownership, Profitability, Liquidity, Dividend Policy, Debt Policy on Firm Value. *Management Analysis Journal*, 7(2), pp. 211-222.
- Tsai, P.-F., Gao, C.-H., & Yuan, S.-M. (2023). Stock selection using machine learning based on financial ratios. *Mathematics*, 11(23), 4758. <https://doi.org/10.3390/math11234758>.
- Vintilă, G., & Nenu, E. A. (2015). An Analysis of Determinants of Corporate Financial Performance: Evidence from the Bucharest Stock Exchange Listed Companies . *Econjournals*, 732-739.
- Zavali, T., Lehenchuk, S., Chyzhevska, L., & Hrabchuk, I. (2025). Determinants of financial performance in advertising and marketing companies: Evidence from central and eastern european countries. *Journal of Risk and Financial Management*, 18(3), 141. <https://doi.org/10.3390/jrfm18030141>.
- Zhao, C., Yuan, X., Long, J., Jin, L., & Guan, B. (2023). Financial indicators analysis using machine learning: Evidence from Chinese stock market. *Finance Research Letters*, 58, 104590. <https://doi.org/10.1016/j.frl.2023.104590>.