ECONOMETRICS
EKONOMETRIA
Advances in Applied Data Analysis

Econometrics. Ekonometria.
Advances in Applied Data Analysis
Year 2026, Vol. 30, No. 2
ISSN 2449-9994
journals.uew.pl/eada/

A Quantum-Like Tensor State Model for Bivariate Time Series in Forestry and Residential Construction

Jan Kotlarz

Nicolaus Copernicus University in Toruń, Poland

e-mail: jan.kotlarz@umk.pl

ORCID: 0000-0002-8212-7798

Quote as: Kotlarz, J. (2026). A Quantum-Like Tensor State Model for Bivariate Time Series in Forestry and Residential Construction. *Econometrics. Ekonometria. Advances in Applied Data Analysis*, *30*(2), 26-49

DOI: 10.15611/eada.2026.2.03

JEL: C32, C33, C63, C73, C74

Abstract

Aim: To develop and empirically evaluate a quantum-like tensor state model for bivariate time series, focusing on the joint dynamics of timber harvesting and residential construction across Polish regions.

Methodology: Two related processes were encoded in a four-dimensional complex Hilbert space as a tensor product state whose residual dynamics were governed by a parametrised two-qubit unitary operator acting as a nonlinear second-stage correction to pooled OLS forecasts. The model was estimated on panel data for 16 voivodeships over 2005–2025 and compared with linear, autoregressive, polynomial, and tree-based residual benchmarks using MSE and MAE.

Results: The tensor model substantially reduced global forecast errors for harvesting amplitudes relative to all benchmarks and achieved competitive performance for construction, with statistically significant gains over linear and low-order nonlinear specifications and win rates above 80–90% of regions for harvesting, while its accuracy was broadly comparable to a shallow regression tree for construction.

Implications and recommendations: The findings indicate that quantum-inspired tensor state models can serve as practically useful tools for forecasting and interpreting time-varying cross-sector dependence in regional panels, supporting planning and risk assessment in forestry construction systems. Future research should extend the framework to multisector settings, richer entangling kernels, and partially regionalised operators, and explore applications to other domains with nonlinear, time-varying co-movement.

Originality/value: This study provides one of the first applications of a quantum-like two-qubit tensor state model to classical economic time series, demonstrating that a low-dimensional unitary evolution can yield forecast accuracy at least comparable to strong classical benchmarks while offering a compact, interpretable representation of joint dynamics and entanglement-like interactions between sectors.

Keywords: quantum inspired models, tensor state, bivariate time series, forestry, residential construction

1. Introduction

This study was motivated by the need to model the joint evolution of two economically relevant processes, in this case in two sectors: 1) forestry and 2) residential construction. In many regions, changes in timber resource acquisition and housing construction activity are reported as closely linked, with implications for planning, investment and risk management. Kaputa (2004) observed that the supply of wood-based materials in Poland is primarily driven by demand from downstream industries such as furniture and construction, with building activity directly shaping the use of sawnwood, panels, and other semi-finished wood products. Ratajczak (2014) reported that in Poland roughly 60% of wood materials are ultimately consumed in the building sector (about 40% directly in construction and 20% indirectly via furniture and interior fittings), implying that changes in housing demand are transmitted into timber demand and, in the longer run, into harvesting levels. Dróżdż (2022) documented that, between 2000 and 2020, production of sawnwood and wood-based panels closely followed demand from construction and furniture manufacturers. Beyond the Polish context, Perez-Garcia et al. (2005) showed for US residential construction that alternative structural designs (wood vs. steel or concrete) imply markedly different wood requirements and thus different pressures on forest resources over the full regeneration-to-construction life cycle. On the other hand, the level of timber harvesting in Poland is regulated by ten-year forest management plans, and its magnitude is shaped by factors other than current construction demand, including those arising from sustainable forest management. Rykowski (2012), for example, discussed how environmental and silvicultural conditions – such as genetic variability, fertilisation, and cutting regimes – jointly determine wood production in forests, while Kotlarz and Bejger (2024) noted that climate change may also significantly affect the future availability of wood.

Classical multivariate time series models, such as vector autoregressions (e.g. Mokhtarzadeh (2021), Banaś et al. (2022)) and dynamic regression models (e.g. Hetemäki et al. (2004)), usually described each component through linear combinations of past values and possibly exogenous variables. While these approaches are flexible, they tend to encode dependence between series only via linear cross–-lag terms and do not offer an explicit representation of the joint state of the system at each time point. Song (2006) pointed out the limitations of standard time series models, particularly their reliance on linearity assumptions and stationarity conditions. Those issues are particularly relevant in the present setting, since high construction activity at year t can stimulate timber harvesting in the short run, but may also deplete available resources over longer horizons, leading to lower harvesting in later periods. Conversely, high harvesting does not necessarily induce higher construction activity, whereas low harvesting levels directly constrain construction. This implies that questions regarding synchronisation, lead–lag structure and time-varying qualitative dependence between the processes remain largely implicit in these frameworks. As a result, they may struggle to capture changes in the qualitative relation between the processes, for example shifts between phases where forestry leads or lags construction in different ways.

Quantum-inspired modelling offers an alternative viewpoint in which a system is represented by a state in a Hilbert space and its evolution is described by operators acting on that state. In the present setting, the entire market composed of the forestry and residential construction sectors is described

by a single joint state, which captures their combined configuration at each point in time and evolves jointly, allowing changes in the nature of their interaction to be represented directly through the dynamics of that state. The practical applicability of quantum-inspired modelling frameworks is currently under active investigation, particularly in the context of intelligent data processing and market analysis, where such approaches are increasingly viewed as viable tools for capturing complex dynamic interactions (cf. Mazzaro (2025), Kapoor and Garg (2026)). In this framework, superposition and entanglement provide richer ways to encode dependence than traditional latent variables. For joint time series, this suggests modelling the pair of processes through a shared state whose evolution can capture nonlinear and dynamically changing interactions that go beyond simple linear cross–-effects.

In this article, a quantum-like tensor state model for bivariate time series is proposed. The two processes are encoded into a tensor product state in a four-dimensional complex Hilbert space, representing their joint configuration at each time point. The transition from one period (year) to the next is governed by a parametrised unitary 4×4 operator, constructed using a KAK – inspired decomposition into local single-qubit rotations and an entangling core, and observable next-period amplitudes for each process are obtained by aggregating components of the evolved tensor state. The quantum-like tensor state model proposed in this work bears conceptual similarity to the approaches of Vipulananthan et al. (2024) and Udvarnoki and Fáth (2025) in that all three frameworks encode classical time series into tensor – or Hilbert-space representations – allowing the joint dynamics of multiple processes to be captured through a unified, high-dimensional state.

The aim of this study was to develop and study the proposed tensor state model in an applied stochastic modelling setting. The model was formally introduced. The estimation procedure, based on minimising a joint loss for both channels, was described and practical implementation aspects discussed. The performance of the model was then assessed through a simulation study and an empirical application to Polish forestry and residential construction data, and comparisons were made with linear, lag-based and simpler quantum-inspired rotation models, showing how the estimated tensor states and unitary parameters can be used to interpret heterogeneous joint dynamics across regions.

2. Data Motivating the Methodology

The data used in this study were obtained from the Local Data Bank (*Bank Danych Lokalnych*, BDL) maintained by Statistics Poland (GUS), which aggregates official statistics from various administrative sources (Główny Urząd Statystyczny (2026). Timber harvesting originated from annual forestry reports compiled on the basis of information submitted by the State Forests (*Lasy Państwowe*) and private forest owners, and records the volume of timber extracted (in cubic meters of roundwood). Residential construction data were derived from annual building completion reports collected by local authorities and aggregated by GUS. In both cases, the data were disaggregated by Poland's main administrative units (16 voivodeships), so that for each sector, 16 annual time series from 2005 to 2025 were observed.[1]

In a first step, both raw series were restricted to the years from 2005 to 2025 and filtered to exclude national aggregates, retaining only observations for the 16 voivodeships. For each voivodeship $v \in \{1, \dots, 16\}$, let

[1] In this paper, each voivodeship is identified by a four-letter code: DOLN = Dolnośląskie, KUJA = Kujawsko-Pomorskie, LUBE = Lubelskie, LUBU = Lubuskie, ŁÓDZ = Łódzkie, MAŁO = Małopolskie, MAZO = Mazowieckie, OPOL = Opolskie, PODK = Podkarpackie, PODL = Podlaskie, POMO = Pomorskie, SLĄS = Śląskie, ŚWIĘ = Świętokrzyskie, WARM = Warmińsko-Mazurskie, WIEL = Wielkopolskie, ZACH = Zachodniopomorskie.

$$H_v = \big(H_v(2005), H_v(2006), \dots, H_v(2025)\big)$$

denote the annual series of timber harvesting volumes, and let

$$B_v = \big(B_v(2005), B_v(2006), \dots, B_v(2025)\big)$$

denote the corresponding series of residential construction activity. To put all regions and both sectors on a comparable scale, simple max-scaling was then applied separately within each voivodeship and sector. Specifically, letting

$$H_v^{\max} = \max_{t\in\{2005,\dots,2025\}} H_v(t)$$

$$B_v^{\max} = \max_{t\in\{2005,\dots,2025\}} B_v(t)$$

and the normalised series were defined as

$$\tilde{H}_v(t) = \frac{H_v(t)}{H_v^{\max}} \qquad \tilde{B}_v(t) = \frac{B_v(t)}{B_v^{\max}} \qquad t = 2005, \dots, 2025. \tag{1}$$

This transformation (Equation 1) preserved the temporal shape and relative variability of each regional series, while mapping all observations into the unit interval $[0,1]$ (see Figure 1).

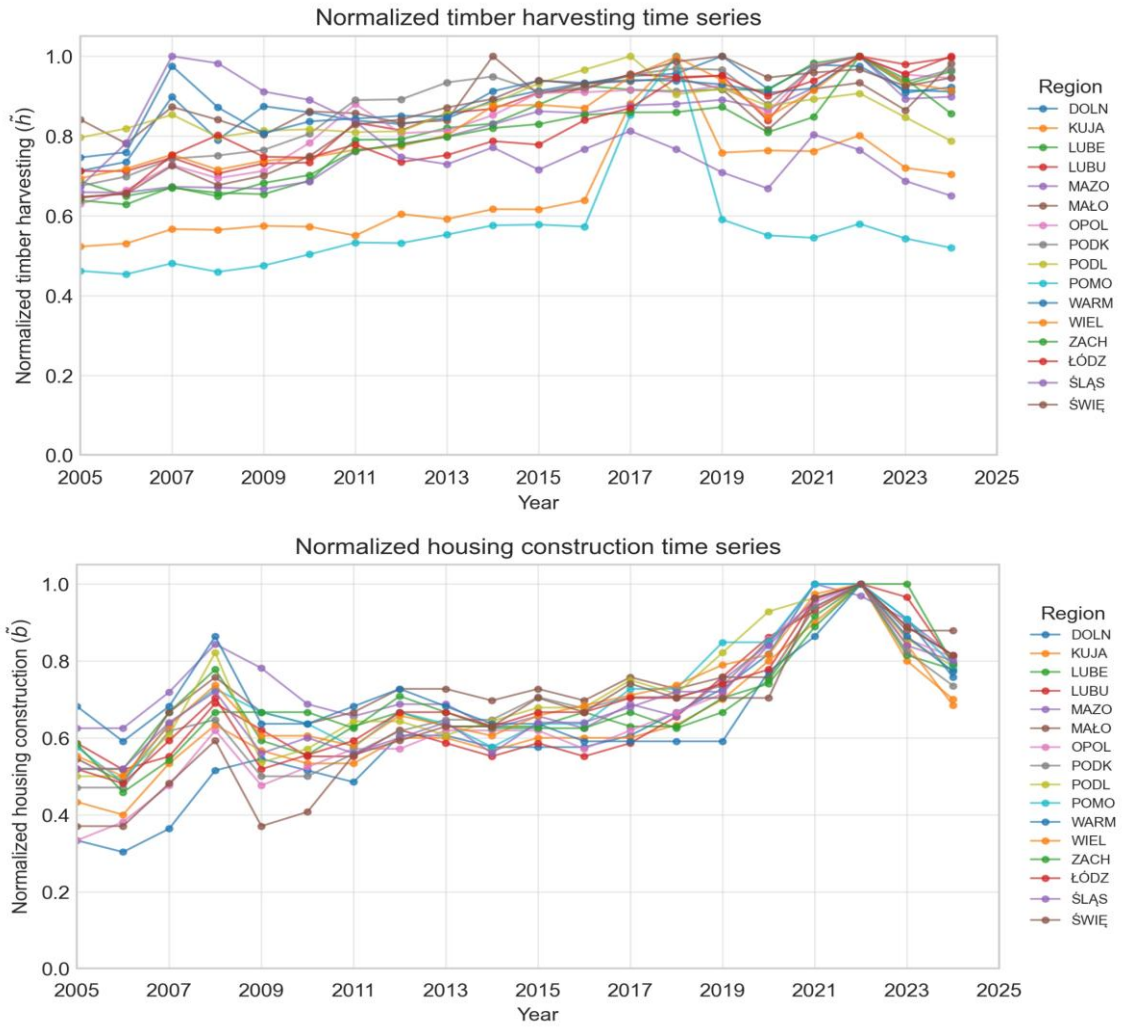

Fig. 1. Normalised source data for timber harvesting and residential construction across 16 Polish voivodeships

Source: author's calculations based on Statistics Poland (GUS) Local Data Bank (Bank Danych Lokalnych, BDL) data.

Stationarity and trend properties: To gain a first impression of the temporal properties of the regional series, the stationarity of the $\widetilde{H}_v(t)$ and $\widetilde{B}_v(t)$ data was examined using two complementary unit root and stationarity tests. Specifically, for each voivodeship v the Augmented Dickey-Fuller (ADF) test was applied to assess the presence of a unit root, and the Kwiatkowski-Phillips-Schmidt-Shin (KPSS) test was used to examine the null of (trend) stationarity. In about three-quarters of the regions and for both variables, the ADF test yielded p values above 0.2, so that the unit–root null could not be rejected, while the KPSS statistics were significant at conventional 5% levels (p–values below 0.05), leading to a rejection of the null of stationarity. Taken together, these results indicated that, even after max-scaling, the level series $\widetilde{H}_v(t)$ and $\widetilde{B}_v(t)$ were predominantly non-stationary and exhibited pronounced trends or persistent changes in level over the sample period, with this pattern being particularly strong for residential construction.

Turning points: Figure 2 presents the mean time series of normalised $\widetilde{H}_v(t)$ and $\widetilde{B}_v(t)$ over all 16 voivodeships for each year. The faint lines show the annual cross-sectional means and the bold curves represent 3-year moving averages. Both sectors exhibit a pronounced upward trajectory over most of the sample: average harvesting rose almost monotonically from around 0.68 in 2005 to a plateau of 0.90 after 2018, whereas construction increased more modestly until roughly 2014, and then accelerated markedly, reaching a peak close to 1.0 in 2022. The smoothed series revealed turning points for timber harvesting in 2018-2019, where the trend flattened and slightly reversed. For construction this was around 2022, where the boom phase ended and the mean normalised level declined. Comparing the slopes of the two curves demonstrated that the dynamics of $\widetilde{H}$ dominated in the first half of the period (roughly 2005–2014), when harvesting grew faster and stayed well above the construction series, while from 2015 onwards the growth of $\widetilde{B}$ became steeper, culminating in the sharp post-2020 upswing during which construction temporarily overtook harvesting before both series converged again towards the end of the sample.

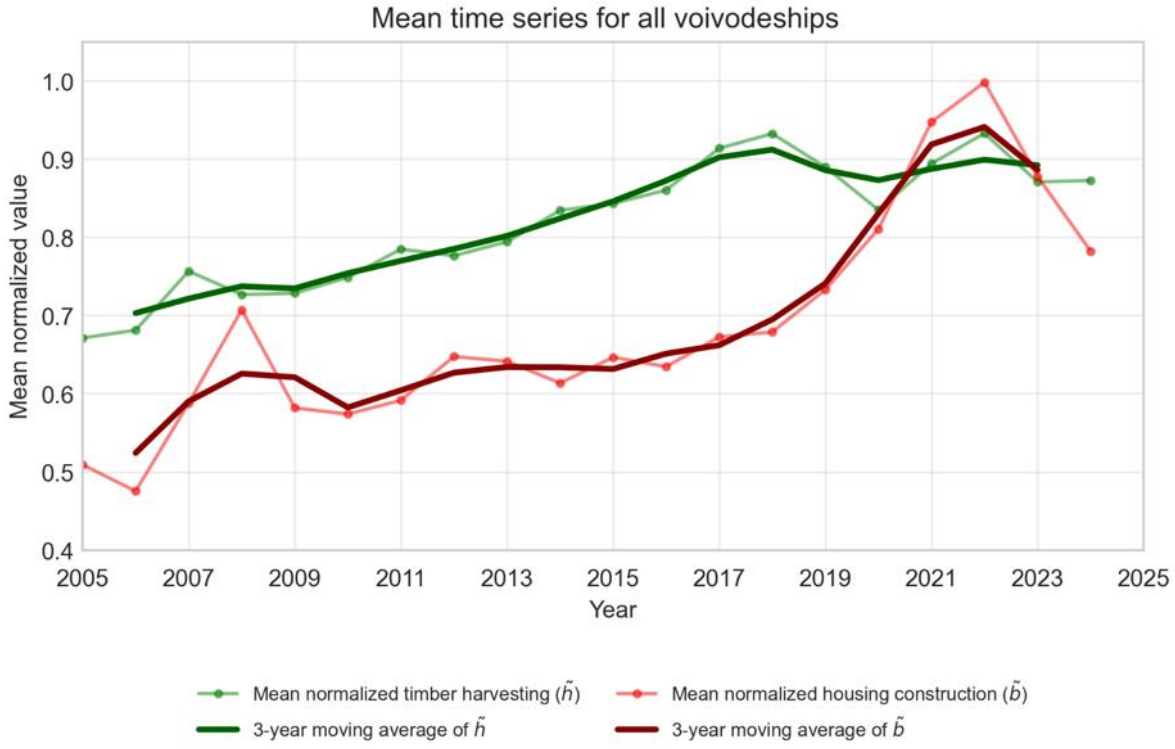

Fig. 2. Mean normalised timber harvesting ($\widetilde{H}$, green) and residential construction ($\widetilde{B}$, red) across Polish voivodeships, together with 3-year moving averages

Source: author's calculations based on Statistics Poland (GUS) Local Data Bank (Bank Danych Lokalnych, BDL) data.

Joint dynamics: The descriptive evidence above indicated that the average normalised series for timber harvesting $\widetilde{H}$ and housing construction $\tilde{B}$ shared a common long-term upward trend, whereas their short-term dynamics and co–movement were far from stable. The scatter plot of annual means $\left(\bar{\tilde{H}}_t, \bar{\tilde{B}}_t\right)$ suggested a broadly positive association, yet for comparable levels of $\bar{\tilde{H}}_t$ the corresponding values of $\bar{\tilde{B}}_t$ could differ substantially across the years, while the rolling correlation between the two series alternated between high positive and clearly negative values. This time-varying strength and even sign of the dependence implies that simple linear models with constant coefficients are unlikely to capture the qualitative changes in the relationship between harvesting and construction, which motivates the search for a more flexible, explicitly dynamic representation of their joint evolution.

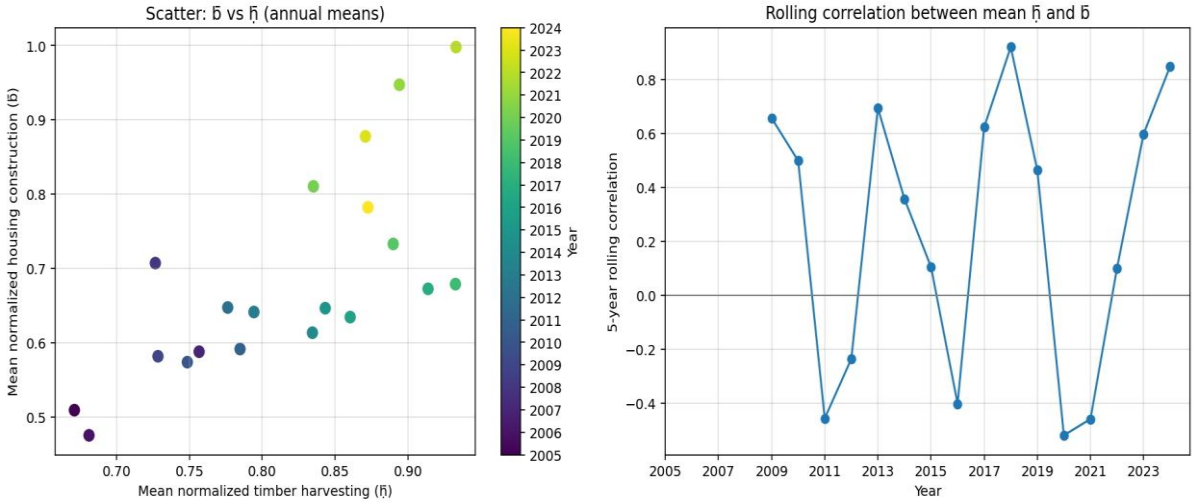

Fig. 3. Relationship between mean normalised timber harvesting ($\widetilde{H}$) and housing construction ($\tilde{B}$) at national level. The left panel shows a scatter plot of annual cross-sectional means $\left(\bar{\tilde{H}}_t, \bar{\tilde{B}}_t\right)$. The right panel shows the 5-year rolling correlation between $\bar{\tilde{H}}_t$ and $\bar{\tilde{B}}_t$, which oscillated between strongly positive and negative values over the sample, indicating that the strength and even the sign of the short-run co-movement between harvesting and construction changed markedly across subperiods

Source: author's calculations based on Statistics Poland (GUS) Local Data Bank (Bank Danych Lokalnych, BDL) data.

Taken together, these empirical patterns point to non-stationary trends, time-varying co-movement and occasional decoupling between harvesting and construction, which are difficult to accommodate within standard linear time series models with fixed coefficients. This motivates the use of a more flexible framework that can represent the joint state of the two processes and allow their dependence structure to evolve over time, as in the quantum-inspired model introduced in the next section.

Specifically, the subsequent analysis was guided by the following hypotheses:

- (i) the qualitative dependence between timber harvesting and residential construction is time-varying, exhibiting phases with different signs and strengths of short-run co-movement and changing lead-lag patterns;
- (ii) such regime-like shifts in dependence are not adequately captured by standard linear multivariate models with constant parameters, but can be represented by a low-dimensional joint state evolving under a unitary two-qubit operator;
- (iii) under the null of equal predictive accuracy, the tensor 4D model and the benchmark specifications (static linear, lag, quantum-inspired 2D) achieve the same expected forecast loss for harvesting and construction amplitudes, so that their mean loss differentials in terms of MSE and MAE are zero;
- (iv) under the working alternative, the tensor 4D model attains strictly lower average forecast errors (for example in terms of MSE or MAE) than at least one of the benchmark models in each

channel, and actually outperforms all the benchmarks in a non-negligible fraction of regions, as measured by negative mean loss differentials.

Moreover, the following hypotheses were formulated regarding cross-regional dominance patterns and the conditions under which the tensor model was expected to perform best:

- (v) the proportion of regions in which the tensor 4D model attained the lowest forecast loss (MSE and MAE) among all the considered models was significantly higher than the value that would be expected if all four specifications were equally likely to be the best;
- (vi) the relative gains of the tensor 4D model over linear and low-dimensional quantum-inspired benchmarks, as captured by regional loss differentials and win rates, are largest in regions or subperiods characterised by strongly time-varying and non-linear dependence between harvesting and construction, as reflected for example in high variability and frequent sign changes of rolling correlations or other measures of non-linearity.

3. Quantum-Inspired Tensor State Model

Recent work in quantum machine learning and numerical analysis has shown that quantum-inspired representations of multivariate data, based on tensor-network or density-matrix formalisms, can capture complex temporal and cross-sectional structure while remaining efficiently emulable on classical hardware (García-Ripoll (2021), Viqueira et al. (2025), Fellner et al. (2024)). Building on these ideas, a quantum-inspired tensor state model with a two-stage architecture is specified: a global linear OLS component, estimated once on the pooled panel and shared across all regions, and a shared two--qubit residual dynamics acting on the remaining nonlinear structure rather than on the raw sectoral levels. In particular, the proposed architecture models the joint evolution of harvesting and construction residuals in a low-rank two-qubit state shared across all regions, aiming to combine the expressive power of quantum recurrent dynamics with the robustness and regularisation of a classical two-stage procedure.

3.1. Encoding Marginal Processes into Amplitudes

At each time t, and for each voivodeship v, normalised sectoral intensities $\widetilde{H}_v(t) \in [0,1]$ for timber harvesting and $\tilde{B}_v(t) \in [0,1]$ for residential construction were observed. In the first stage these rescaled levels entered global ordinary least squares (OLS) models for the next-period amplitudes of harvesting and construction. Thus the following specifications were estimated:

$$a_F^{\mathrm{OLS}}(t+1,v) = \beta_{F,0} + \beta_{F,1}\,\widetilde{H}_v(t) + \beta_{F,2}\,\tilde{B}_v(t), \tag{2}$$

$$a_B^{\mathrm{OLS}}(t+1,v) = \beta_{B,0} + \beta_{B,1}\,\widetilde{H}_v(t) + \beta_{B,2}\,\tilde{B}_v(t), \tag{3}$$

with a single set of coefficients β_F and β_B shared across all voivodeships. The empirical next-period amplitudes were defined as square roots of the observed normalised levels:

$$a_F^{\mathrm{emp}}(t+1,v) = \sqrt{\widetilde{H}_v(t+1)}, \qquad a_B^{\mathrm{emp}}(t+1,v) = \sqrt{\tilde{B}_v(t+1)}.$$

The corresponding OLS residuals,

$$\varepsilon_F(t+1,v) = a_F^{\mathrm{emp}}(t+1,v) - a_F^{\mathrm{OLS}}(t+1,v),$$

$$\varepsilon_B(t+1,v) = a_B^{\mathrm{emp}}(t+1,v) - a_B^{\mathrm{OLS}}(t+1,v),$$

were the primary objects of interest for the quantum-inspired component: the tensor state and its unitary evolution were used to model the nonlinear, cross-sectional dependence structure in $(\varepsilon_F, \varepsilon_B)$, conditional on the global linear effects already captured by Equation 2 and Equation 3.

To embed the residuals in a two-qubit state, they were mapped to pseudo-probabilities in $[0,1]$ using an affine rescaling over the full panel

$$r_F(t+1,v) = \frac{\varepsilon_F(t+1,v) - r_{F,\min}}{r_{F,\max} - r_{F,\min}}, \qquad r_B(t+1,v) = \frac{\varepsilon_B(t+1,v) - r_{B,\min}}{r_{B,\max} - r_{B,\min}},$$

where $r_{F,\min}$, $r_{F,\max}$, $r_{B,\min}$, and $r_{B,\max}$ denote the empirical minima and maxima of the residuals across all regions and periods. The resulting quantities $r_F, r_B \in [0,1]$ were interpreted as marginal probabilities for the 'residual qubits'. Their one-qubit amplitude vectors were then obtained by taking positive square roots and being completed to unit length:

$$\mathbf{a}_F^{(r)}(t+1,v) = \begin{pmatrix} a_{F,r}^{(1)}(t+1,v) \\ a_{F,r}^{(2)}(t+1,v) \end{pmatrix} = \begin{pmatrix} \sqrt{r_F(t+1,v)} \\ \sqrt{1 - r_F(t+1,v)} \end{pmatrix},$$

$$\mathbf{a}_B^{(r)}(t+1,v) = \begin{pmatrix} a_{B,r}^{(1)}(t+1,v) \\ a_{B,r}^{(2)}(t+1,v) \end{pmatrix} = \begin{pmatrix} \sqrt{r_B(t+1,v)} \\ \sqrt{1 - r_B(t+1,v)} \end{pmatrix},$$

so that $\| \mathbf{a}_F^{(r)}(t+1,v) \|^2 = \| \mathbf{a}_B^{(r)}(t+1,v) \|^2 = 1$ for all t and v. In contrast to the original specification, which operated directly on sectoral levels, the tensor model here was explicitly framed as a nonlinear second-stage correction to the global OLS baseline.

3.2. Tensor Product State Representation

The joint two-qubit residual state for region v at time $t+1$ was constructed as the tensor product of the marginal residual amplitude vectors,

$$\left|\psi_v^{(r)}(t+1)\right\rangle = \left|F_v^{(r)}(t+1)\right\rangle \otimes \left|B_v^{(r)}(t+1)\right\rangle = \mathbf{a}_F^{(r)}(t+1,v) \otimes \mathbf{a}_B^{(r)}(t+1,v),$$

which can be written in the computational basis $\{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}$ as

$$\left|\psi_v^{(r)}(t+1)\right\rangle = \begin{pmatrix} \psi_{00,v}^{(r)}(t+1) \\ \psi_{01,v}^{(r)}(t+1) \\ \psi_{10,v}^{(r)}(t+1) \\ \psi_{11,v}^{(r)}(t+1) \end{pmatrix} = \begin{pmatrix} a_{F,r}^{(1)}(t+1,v) a_{B,r}^{(1)}(t+1,v) \\ a_{F,r}^{(1)}(t+1,v) a_{B,r}^{(2)}(t+1,v) \\ a_{F,r}^{(2)}(t+1,v) a_{B,r}^{(1)}(t+1,v) \\ a_{F,r}^{(2)}(t+1,v) a_{B,r}^{(2)}(t+1,v) \end{pmatrix}. \tag{4}$$

By construction the residual state defined in the Equation 4 was normalised, $\| \left|\psi_v^{(r)}(t+1)\right\rangle \|^2 = 1$, and encoded the full four–dimensional dependence structure of the joint residuals.

3.3. Unitary Two-Qubit Evolution on Residuals

The joint dynamics of the residual state over one period was generated by a fixed 4×4 unitary matrix $U(\boldsymbol{\theta})$ of KAK-inspired form, shared across all the regions and time periods. To obtain a parsimonious parametrisation compatible with the short time series and to mitigate overfitting, the unitary was factorised as

$$U(\boldsymbol{\theta}) = \mathrm{e}^{\mathrm{i}\phi} (V \otimes V)\, U_{\mathrm{ent}}(\alpha)\, (U \otimes U), \tag{5}$$

where the local single–qubit gates were[2]:

[2] In much of the literature the pre-entanglement rotation is denoted by b and the post–entanglement rotation by e; in this paper, however, the latter is written as η in order to avoid confusion with e, which denotes Euler's number in complex exponentials.

$$U = R_y(b) \qquad V = R_y(\eta),$$

and

$$R_y(\vartheta) = \begin{pmatrix} \cos\left(\frac{\vartheta}{2}\right) & -\sin\left(\frac{\vartheta}{2}\right) \\ \sin\left(\frac{\vartheta}{2}\right) & \cos\left(\frac{\vartheta}{2}\right) \end{pmatrix},$$

while the entangling kernel was implemented as a single Pauli coupling

$$U_{\text{ent}}(\alpha) = \mathrm{e}^{-\mathrm{i}\alpha\, Z\otimes Z} = \cos(\alpha)\, I_4 - \mathrm{i}\sin(\alpha)\,(Z \otimes Z),$$

using identity $(Z \otimes Z)^2 = I_4$. Collecting all the parameters into

$$\boldsymbol{\theta} = (b, \eta, \alpha, \phi)^{\top},$$

the one-step evolution of the residual tensor state was

$$\left|\psi_v^{(r)}(t+2)\right\rangle = U(\boldsymbol{\theta})\, \left|\psi_v^{(r)}(t+1)\right\rangle, \qquad v = 1, \ldots, 16. \tag{6}$$

In the estimation presented in Equation 6, the parameter vector $\boldsymbol{\theta}$ was regularised using an ℓ_2 penalty of magnitude between 10^{-3} and 10^{-2}, shrinking the effective rotation angles and global phase towards zero and prevented the unitary from interpolating the residuals perfectly in-sample. This ensured that, even though the tensor model was substantially more flexible than linear or low-dimensional nonlinear benchmarks, its fitted dynamics remained stable and predictive rather than degenerate.

3.4. Observation Operators and Induced Residual Dynamics

Given the evolved residual state $\left|\psi_v^{(r)}(t+2)\right\rangle$ with components

$$\left|\psi_v^{(r)}(t+2)\right\rangle = \begin{pmatrix} \psi_{00,v}^{(r)}(t+2) \\ \psi_{01,v}^{(r)}(t+2) \\ \psi_{10,v}^{(r)}(t+2) \\ \psi_{11,v}^{(r)}(t+2) \end{pmatrix},$$

the next-period pseudo-probabilities for the residual qubits were recovered by summing over the appropriate basis states. The probability that the harvesting residual qubit was in state $|1\rangle$ (irrespective of the construction residual) was

$$r_F(t+2, v) = \left|\psi_{10,v}^{(r)}(t+2)\right|^2 + \left|\psi_{11,v}^{(r)}(t+2)\right|^2,$$

while the probability that the construction residual qubit in state $|1\rangle$ was

$$r_B(t+2, v) = \left|\psi_{01,v}^{(r)}(t+2)\right|^2 + \left|\psi_{11,v}^{(r)}(t+2)\right|^2.$$

Equivalently, these can be written as expectation values of diagonal observation operators O_F and O_B,

$$O_F = \mathrm{diag}(0,0,1,1) r_F(t+2, v) = \left\langle\psi_v^{(r)}(t+2)\right| O_F \left|\psi_v^{(r)}(t+2)\right\rangle,$$
$$O_B = \mathrm{diag}(0,1,0,1) r_B(t+2, v) = \left\langle\psi_v^{(r)}(t+2)\right| O_B \left|\psi_v^{(r)}(t+2)\right\rangle.$$

These pseudo-probabilities were then mapped back to residual amplitudes,

$$\varepsilon_F^{\text{tensor}}(t+2,v) = r_F(t+2,v)\left(r_{F,\max} - r_{F,\min}\right) + r_{F,\min},$$
$$\varepsilon_B^{\text{tensor}}(t+2,v) = r_B(t+2,v)\left(r_{B,\max} - r_{B,\min}\right) + r_{B,\min},$$

and combined with the OLS baselines to obtain full next-period amplitude forecasts,

$$\hat{a}_F(t+2,v) = a_F^{\text{OLS}}(t+2,v) + \varepsilon_F^{\text{tensor}}(t+2,v),$$
$$\hat{a}_B(t+2,v) = a_B^{\text{OLS}}(t+2,v) + \varepsilon_B^{\text{tensor}}(t+2,v).$$

This closed a nonlinear induced residual dynamics $\left(\varepsilon_F(t+1,v), \varepsilon_B(t+1,v)\right) \mapsto \left(\varepsilon_F(t+2,v), \varepsilon_B(t+2,v)\right)$ generated by the fixed global unitary $U(\boldsymbol{\theta})$, with the overall prediction for sectoral activity given by the sum of the linear OLS component and the quantum-inspired tensor correction.

3.5. Estimation and Practical Implementation

The tensor model was estimated in the amplitude domain, but applied to the residual component left after removing a global linear effect. For each voivodeship and period, the empirical next-period amplitudes $a_F^{\text{emp}}(t+1,v)$ and $a_B^{\text{emp}}(t+1,v)$ were formed as square roots of the normalised harvesting and construction levels, and decomposed into an OLS baseline plus residuals as in (2)–residuals definition. The tensor unitary $U(\boldsymbol{\theta})$ was then estimated by matching the model-implied residual amplitudes to their empirical counterparts, using a pooled loss that averaged mean squared errors across both channels F and B. In the main specification, the objective was augmented with an ℓ_2 penalty on parameter vector $\boldsymbol{\theta}$ of magnitude between 10^{-3} and 10^{-2}, which stabilised the fit and prevented the two-qubit dynamics from perfectly interpolating the in-sample residuals. The resulting forecasts were obtained by adding the tensor-based residual corrections back to the global OLS predictions, and evaluated with standard regression metrics (MSE, RMSE, MAE, R^2) for amplitudes a_F and a_B, so that all benchmark and tensor models were compared on a common one-step-ahead accuracy scale.

4. Benchmark Models

To assess the incremental value of the quantum-inspired tensor specification, its forecasting performance is compared against a set of parsimonious benchmark models defined in the same amplitude domain. All the benchmarks were estimated globally on the pooled panel of voivodeships and time periods, using the same target variables as the tensor model and the same one-step-ahead forecasting horizon. The focus was on specifications that were simple, transparent and directly comparable to the tensor residual architecture.

4.1. Baseline OLS Model on Amplitudes

As the first and most transparent benchmark, global ordinary least squares (OLS) regressions were used to map current-period normalised sectoral levels into next-period amplitudes. For harvesting:

$$a_F^{\text{emp}}(t+1,v) = \beta_{F,0} + \beta_{F,1}\,\widetilde{H}_v(t) + \beta_{F,2}\,\widetilde{B}_v(t) + u_{F,t+1,v}$$

and analogously for construction:

$$a_B^{\text{emp}}(t+1,v) = \beta_{B,0} + \beta_{B,1}\,\widetilde{H}_v(t) + \beta_{B,2}\,\widetilde{B}_v(t) + u_{B,t+1,v},$$

where a_F^{emp} and a_B^{emp} were square roots of the observed normalised levels at $t+1$, and $\widetilde{H}_v(t)$, $\widetilde{B}_v(t)$ were the period-t covariates. Coefficients β_F and β_B were estimated once on the pooled sample and reused for all the regions and periods. These OLS models served two roles: they provided a strong linear benchmark in their own right and defined the residuals that were subsequently modelled by the tensor architecture.

4.2. Residual-Based Benchmarks

To ensure that any gains attributed to the tensor model were not simply capturing trivial patterns in the OLS residuals, several residual-based benchmark specifications were introduced. All of them operated in the same residual amplitude domain as the tensor model, and produced one-step-ahead corrections that were added back to the OLS baseline.

Let $\varepsilon_F(t+1,v)$ and $\varepsilon_B(t+1,v)$ denote the OLS residuals for harvesting and construction amplitudes, respectively, as in the residuals definition. The following residual-based benchmarks were considered:

Constant residual correction (CONST_resid). A simple 'no-dynamics' benchmark retains the OLS forecasts and assumes that residuals are constant over time. Denoting by $\bar{\varepsilon}_F$ and $\bar{\varepsilon}_B$ the global sample means of the OLS residuals, the amplitude forecasts are

$$\hat{a}_F^{\text{const}}(t+1,v) = a_F^{\text{OLS}}(t+1,v) + \bar{\varepsilon}_F,$$
$$\hat{a}_B^{\text{const}}(t+1,v) = a_B^{\text{OLS}}(t+1,v) + \bar{\varepsilon}_B.$$

This specification checks whether a simple global bias correction is sufficient.

AR(1) residual dynamics (AR1_resid). To capture linear persistence in residuals, a global autoregression of order one is used,

$$\varepsilon_F(t+1,v) = \gamma_{F,0} + \gamma_{F,1}\,\varepsilon_F(t,v) + \eta_{F,t+1,v},$$
$$\varepsilon_B(t+1,v) = \gamma_{B,0} + \gamma_{B,1}\,\varepsilon_B(t,v) + \eta_{B,t+1,v},$$

estimated on the pooled residual panel. The corresponding forecasts $\hat{\varepsilon}_F^{\text{AR}}$ and $\hat{\varepsilon}_B^{\text{AR}}$ are then added to the OLS baselines,

$$\hat{a}_F^{\text{AR}}(t+1,v) = a_F^{\text{OLS}}(t+1,v) + \hat{\varepsilon}_F^{\text{AR}}(t+1,v),$$
$$\hat{a}_B^{\text{AR}}(t+1,v) = a_B^{\text{OLS}}(t+1,v) + \hat{\varepsilon}_B^{\text{AR}}(t+1,v).$$

This benchmark captures simple linear persistence without cross-sectional interactions between harvesting and construction.

Polynomial residual model (POLY_resid). To allow for mild nonlinearity and interaction with current-sector conditions in a low-dimensional way, a global polynomial regression of residuals on their own lag and contemporaneous sectoral levels is employed. For harvesting residuals:

$$\varepsilon_F(t+1,v) = f_F\left(\varepsilon_F(t,v), \widetilde{H}_v(t), \tilde{B}_v(t)\right) + \upsilon_{F,t+1,v},$$

where f_F is a second-order polynomial in its arguments (including squared terms and simple interactions). An analogous specification is used for $\varepsilon_B(t+1,v)$. The resulting forecasts $\hat{\varepsilon}_F^{\text{poly}}$ and $\hat{\varepsilon}_B^{\text{poly}}$ are combined with the OLS baseline as

$$\hat{a}_F^{\text{poly}}(t+1,v) = a_F^{\text{OLS}}(t+1,v) + \hat{\varepsilon}_F^{\text{poly}}(t+1,v),$$
$$\hat{a}_B^{\text{poly}}(t+1,v) = a_B^{\text{OLS}}(t+1,v) + \hat{\varepsilon}_B^{\text{poly}}(t+1,v).$$

This benchmark provides a parsimonious nonlinear alternative capable of capturing curvature and simple interaction effects in the residual processes, while remaining straightforward to estimate and interpret.

Tree-based residual model (TREE_resid). As a nonparametric reference, a shallow regression tree is applied to the residuals. For each channel $X \in \{F, B\}$, a decision tree of limited depth is fitted to predict $\varepsilon_X(t+1,v)$ from $\left(\varepsilon_X(t,v), \widetilde{H}_v(t), \tilde{B}_v(t)\right)$. Denoting the tree-based forecasts by $\hat{\varepsilon}_F^{\text{tree}}$ and $\hat{\varepsilon}_B^{\text{tree}}$, the corresponding amplitude forecasts are

$$\hat{a}_F^{\text{tree}}(t+1,v) = a_F^{\text{OLS}}(t+1,v) + \hat{\varepsilon}_F^{\text{tree}}(t+1,v),$$
$$\hat{a}_B^{\text{tree}}(t+1,v) = a_B^{\text{OLS}}(t+1,v) + \hat{\varepsilon}_B^{\text{tree}}(t+1,v).$$

Since the trees are deliberately kept shallow (small maximum depth and minimum leaf size), this benchmark captures simple nonlinear threshold effects in the residuals without approaching the flexibility and overfitting risk of full-scale machine learning models.

5. Results for Benchmark and Tensor Models

5.1. Baseline OLS Amplitude Model

As the first-stage global benchmark, pooled ordinary least squares (OLS) models were estimated that mapped current-period normalised sectoral levels into next-period amplitudes for harvesting and construction. For each voivodeship v and year t, the baseline specifications were

$$a_F^{\text{OLS}}(t+1,v) = \beta_{F,0} + \beta_{F,1}\, f_t + \beta_{F,2}\, b_t,$$
$$a_B^{\text{OLS}}(t+1,v) = \beta_{B,0} + \beta_{B,1}\, f_t + \beta_{B,2}\, b_t,$$

where f_t and b_t denote the max-scaled harvesting and construction levels in period t, and $a_F^{\text{OLS}}(t+1,v)$, $a_B^{\text{OLS}}(t+1,v)$ were the corresponding one-step-ahead amplitudes.

The estimated parameter values were

$$a_F^{\text{OLS}}(t+1,v) = 0.506352 + 0.516891\, f_t - 0.032277\, b_t,$$
$$a_B^{\text{OLS}}(t+1,v) = 0.465996 + 0.066719\, f_t + 0.453803\, b_t.$$

In the harvesting equation, the next-period amplitude was driven mainly by its own current level f_t, with only a small negative contribution from construction b_t, whereas in the construction equation the dominant effect came from b_t, and harvesting f_t played only a minor role. These global OLS models provided the baseline forecasts to which the tensor-based residual corrections were subsequently added.

5.2. Overall Forecast Accuracy on Amplitudes

The one-step-ahead forecast accuracy of the tensor 4D model was first compared with that of the residual-based benchmarks on the common amplitude scale for harvesting and construction. Forecast errors were evaluated using mean squared error (MSE, see Equation 7) and mean absolute error (MAE, see Equation 8) computed over the pooled panel of voivodeships and years, separately for harvesting a_F and construction amplitudes a_B. For each model *m* and channel *X* the following measures were reported:

$$\text{MSE}_X^{(m)} = \frac{1}{N}\sum_{t,v}\left(a_X^{\text{emp}}(t+1,v) - \hat{a}_X^{(m)}(t+1,v)\right)^2, \tag{7}$$

$$\text{MAE}_X^{(m)} = \frac{1}{N}\sum_{t,v}\left|a_X^{\text{emp}}(t+1,v) - \hat{a}_X^{(m)}(t+1,v)\right|, \tag{8}$$

where N denotes the total number of one-step-ahead predictions in the evaluation sample. Table 1 presents the global average errors for all six models. For harvesting amplitudes, the tensor specification clearly improved upon both purely linear and nonlinear benchmarks: its MSE of 0.000669 and MAE of 0.01773 were substantially lower than those of the best competing residual model (TREE_resid, MSE 0.000947, MAE 0.02125) and markedly better than the baseline OLS_base (MSE 0.001122, MAE 0.02267). For construction amplitudes, the tensor model again attained the lowest global errors, with MSE 0.001440 and MAE 0.03083 compared with, for example, 0.001569 and 0.02943 for TREE_resid and 0.002231 and 0.03658 for OLS_base, indicating sizeable gains relative to simple linear baselines and smaller but still non--negligible improvements over the strongest nonlinear benchmark.

Table 1. Global one-step-ahead forecast accuracy for harvesting (a_F) and construction (a_B) amplitudes, measured by mean squared error (MSE) and mean absolute error (MAE) over all regions and years

	Harvesting a_F		Construction a_B	
Model	MSE	MAE	MSE	MAE
OLS_base	0.001122	0.022667	0.002231	0.036580
CONST_resid	0.001122	0.022667	0.002231	0.036580
AR1_resid	0.001118	0.022673	0.002186	0.035828
POLY_resid	0.000969	0.021386	0.001989	0.033884
TREE_resid	0.000947	0.021254	0.001569	0.029428
TENSOR_resid	0.000669	0.017730	0.001440	0.030831

Source: author's own forecast calculations based on Statistics Poland (GUS) Local Data Bank (Bank Danych Lokalnych, BDL) data.

To formally assess hypotheses (iii) and (iv) on equal versus superior predictive accuracy, mean loss differentials between the tensor model and each benchmark were analysed using the panel of individual forecast errors. For a given channel X and benchmark m, were defined

$$d_{X,t,v}^{(m,\mathrm{MSE})} = e_{X,t,v}^{(m)\,2} - e_{X,t,v}^{(\mathrm{TENSOR})\,2}, \qquad d_{X,t,v}^{(m,\mathrm{MAE})} = \left|e_{X,t,v}^{(m)}\right| - \left|e_{X,t,v}^{(\mathrm{TENSOR})}\right|$$

and tested $H_0: \mathbb{E}\left[d_{X,t,v}^{(m,\cdot)}\right] = 0$ against the one-sided alternative $\mathbb{E}\left[d_{X,t,v}^{(m,\cdot)}\right] > 0$. Table 2 summarises the sample mean $\bar{d}$, t-statistic from a one-sample Welch-type t-test and the associated p-value.

Table 2. Mean loss differentials $\bar{d}$ (benchmark minus tensor), t-statistics and p-values for one-step-ahead forecasts of amplitudes, based on MSE and MAE losses. Positive $\bar{d}$ indicates that the tensor model attains a lower average loss than the benchmark

Channel	Benchmark	Loss	$\bar{d}$	t-stat	p-value	n
F	OLS	MSE	0.000453	3.70	2.5e-04	304
F	OLS	MAE	0.004937	11.75	1.6e-26	304
F	CONST	MSE	0.000453	3.70	2.5e-04	304
F	CONST	MAE	0.004937	11.75	1.6e-26	304
F	AR1	MSE	0.000448	3.89	1.2e-04	304
F	AR1	MAE	0.004943	11.71	2.2e-26	304
F	POLY	MSE	0.000300	3.61	3.6e-04	304
F	POLY	MAE	0.003656	6.18	2.0e-09	304
F	TREE	MSE	0.000278	5.38	1.5e-07	304
F	TREE	MAE	0.003524	6.40	6.0e-10	304
B	OLS	MSE	0.000790	7.12	7.9e-12	304
B	OLS	MAE	0.005749	6.06	4.0e-09	304
B	CONST	MSE	0.000790	7.12	7.9e-12	304
B	CONST	MAE	0.005749	6.06	4.0e-09	304
B	AR1	MSE	0.000746	6.04	4.4e-09	304
B	AR1	MAE	0.004997	5.18	4.0e-07	304
B	POLY	MSE	0.000549	4.44	1.3e-05	304
B	POLY	MAE	0.003053	3.05	2.5e-03	304
B	TREE	MSE	0.000129	1.09	2.8e-01	304
B	TREE	MAE	-0.001403	-1.13	2.6e-01	304

Source: author's own forecast comparison calculations based on Statistics Poland (GUS) Local Data Bank (Bank Danych Lokalnych, BDL) data.

For harvesting, all the mean loss differentials were positive and highly significant across both MSE and MAE, indicating that the tensor model attained systematically lower forecast errors than every benchmark and providing strong evidence in favour of (iv) for channel F. For construction, the tensor model also delivered significantly lower average loss than the linear and low-order nonlinear benchmarks (OLS, CONST, AR1, POLY), but its advantage over the tree-based model was small and not statistically significant for either MSE or MAE, so that equal predictive accuracy with TREE could not be rejected for channel B. Taken together, these results support hypotheses (iii) and (iv) in a nuanced way: the tensor specification offered clear and statistically robust gains for harvesting amplitudes, while for construction it performed at least as well as the strongest nonlinear benchmark and substantially better than the linear baselines, but without a decisive dominance over all competitors.

5.3. Forecast Comparison Tests (Diebold-Mariano and Welch Tests)

Forecast accuracy of the tensor state model relative to its benchmarks was assessed using loss differentials and formal forecast comparison tests. One-step-ahead forecast errors were denoted by $e_{i,t} = \hat{y}_{i,t} - y_t$, with squared–error losses $L_{i,t} = e_{i,t}^2$. For any pair of models (i, j), the loss differential was defined as $d_t = L_{i,t} - L_{j,t}$, and the null hypothesis of equal predictive accuracy H_0: $\mathbb{E}[d_t] = 0$ was tested against a two-sided alternative. In the main panel-wide comparisons (tensor versus OLS and tensor versus the tree benchmark), a Diebold-Mariano-type test based on the sample mean and variance of $\{d_t\}_{t=1}^{T}$ was applied.

As a complementary analysis, variation in relative performance across three subperiods was examined by comparing average improvements in mean squared error (dMSE) over periods A, B and C. For these cross--period comparisons, Welch-type two-sample $t$tests on mean dMSE between subperiods were employed, treating period-level dMSE as independent with potentially unequal variances; these tests were used diagnostically and did not constitute forecast comparison tests in the Diebold–Mariano sense.

5.4. Cross-Regional Dominance and Win-Rates

To complement the global accuracy results, the number of regions in which the tensor model attained the lowest forecast loss among all the considered specifications was examined. For each voivodeship, channel and loss function, the region-specific average loss for each model was computed and the model(s) with the smallest mean loss were identified.

Figure 4 shows the share of voivodeships in which each model achieved the lowest average MSE or MAE, separately for harvesting and construction. For harvesting amplitudes, the tensor model dominated almost uniformly: it obtained the lowest mean MSE in 15 out of the 16 regions (93.75%), and the lowest mean MAE in 14 out of the 16 regions (87.5%), while the tree-based residual benchmark was best in only one region under MSE and in two regions under MAE. In this channel, none of the purely linear or low-order nonlinear benchmarks ever achieved the best regional performance once the tensor model was included, so that the observed tensor win-rate was far above the value of $1/6$ that would have arisen if all six models were equally likely to be highest.

For construction amplitudes, the picture was more balanced. Under MSE, the tensor and tree-based models each achieved the lowest average loss in 8 out of the 16 regions (50% each), while under MAE the tree-based benchmark was highest in 10 regions (62.5%) and the tensor model in 6 regions (37.5%). Thus, although the tree-based residual model retains a slight edge in terms of cross-sectional MAE for construction, the tensor model remained competitive and still won in a substantial fraction of regions, clearly outperforming the linear benchmarks in terms of regional dominance. Overall, these win-rate patterns support hypothesis (v) for harvesting and provide at least partial support for construction: across channels, the tensor model was best in significantly more regions than would be expected under equal-likelihood random dominance, especially when focusing on the harvesting amplitudes where its cross-regional superiority was most pronounced.

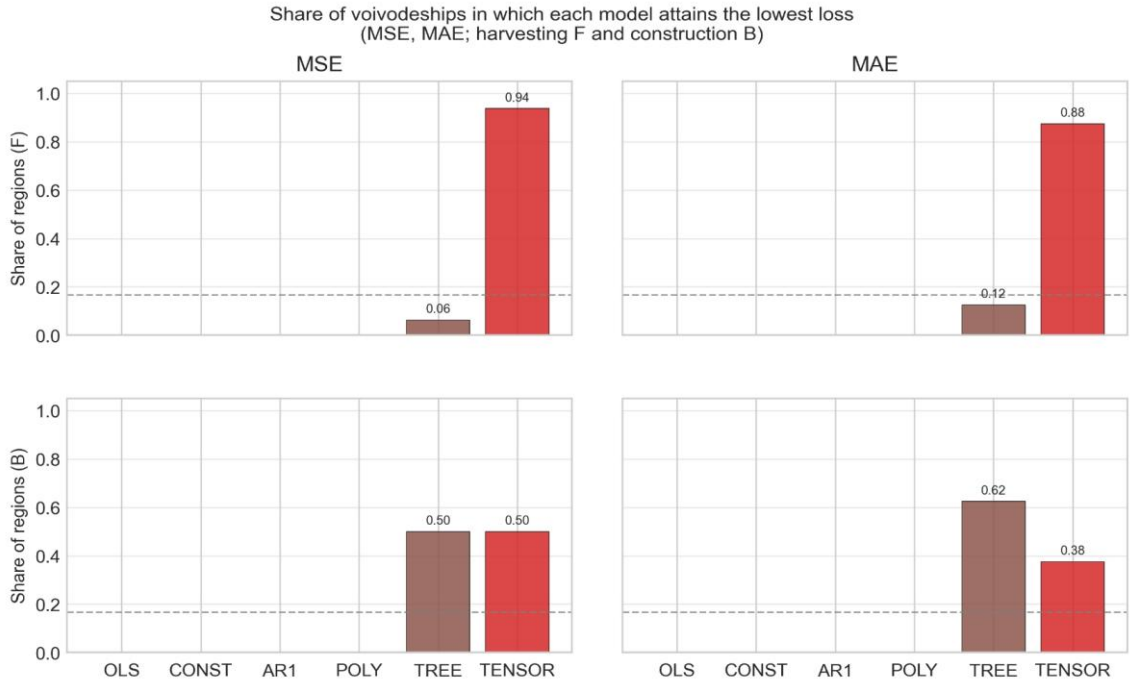

Fig. 4. Share of voivodeships in which each model obtained the lowest average one-step-ahead loss (MSE or MAE) for harvesting (F) and construction (B) amplitudes

Source: author's calculations based on Statistics Poland (GUS) Local Data Bank (Bank Danych Lokalnych, BDL) data.

5.5. Time-Varying Performance Across Periods

5.5.1. Performance Over Early, Middle, and Late Subperiods

To examine how the relative gains of the tensor model evolved over time, the sample was partitioned into three subperiods: early phase $A = 2005\text{-}2010$, middle phase $B = 2011\text{-}2016$, and late phase $C = 2017\text{-}2025$. Phases A and C were characterised by more abrupt changes in the level of construction compared with the more moderate dynamics in B, with the early phase featuring rapid construction growth against a backdrop of almost monotonic increases in harvesting, and the late phase displaying a sharp construction boom on top of a persistently high harvesting plateau (see Figure 2). For each period, channel and benchmark model the mean loss differential in terms of MSE was computed,

$$d_{X,t,v}^{(m)} = e_{X,t,v}^{(m)2} - e_{X,t,v}^{(\mathrm{TENSOR})2}$$

and the average advantage of the tensor model summarised as $\overline{d}_{X,\mathcal{P}}^{(m)}$ over all (t, v) in subperiods A, B and C. The positive values indicated that the tensor specification recorded lower mean squared error than the benchmark.

Figure 5 shows these mean differentials for harvesting and construction, together with t-statistics from one-sample Welch-type tests for $H_0: \mathbb{E}\left[d_{X,\mathcal{P}}^{(m)}\right] = 0$. For harvesting amplitudes, the tensor model consistently improved upon all the benchmarks in each subperiod: its average dMSE against the strongest nonlinear competitor (TREE) was 2.76×10^{-4} in early phase A (t = 3.73, $p \approx 3.6 \times 10^{-4}$), 1.03×10^{-4} in middle phase B (t = 2.53, $p \approx 0.013$), and 4.11×10^{-4} in late phase C (t = 3.79, $p \approx 2.3 \times 10^{-4}$). Similar patterns were obtained for the other benchmarks (OLS, CONST, AR1, POLY), with all mean differentials positive and statistically significant across A, B and C, indicating that the tensor's advantage was both persistent and strongest in the late subperiod when the aggregate dynamics of harvesting and construction became more volatile.

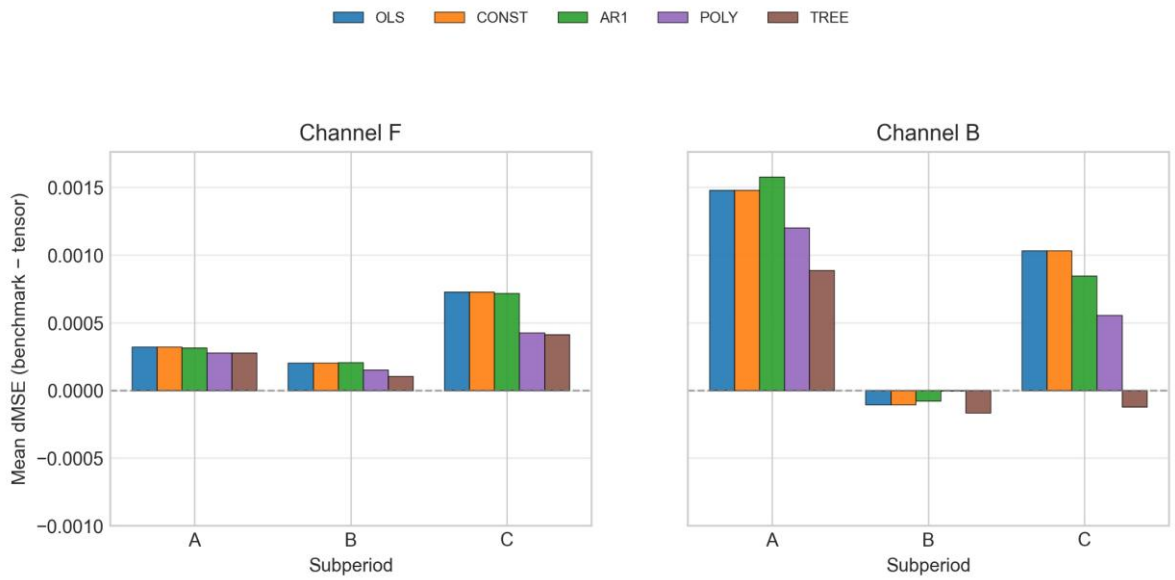

Fig. 5. Mean MSE differentials (benchmark minus tensor) by model and subperiod $A = 2005\text{-}2010$, $B = 2011\text{-}2016$, $C = 2017\text{-}2025$, for harvesting (F) and construction (B) amplitudes. Positive values indicate lower average MSE for the tensor model.

Source: author's calculations based on Statistics Poland (GUS) Local Data Bank (Bank Danych Lokalnych, BDL) data.

For construction amplitudes, the time profile of gains was more heterogeneous. In early phase A, the tensor model clearly outperformed all the benchmarks, with substantial positive dMSE against OLS (0.00148, t = 4.72, $p \approx 1.0 \times 10^{-5}$), AR1 (0.00158, t = 4.08, $p \approx 1.1 \times 10^{-4}$), POLY (0.00120, t = 3.09, $p \approx 2.8 \times 10^{-3}$) and TREE (0.00088, t = 2.80, $p \approx 6.5 \times 10^{-3}$). In middle phase B, however, the tensor's advantage largely vanished or even reversed: the mean dMSE became slightly negative for most benchmarks (e.g. -1.06×10^{-4} vs OLS, -1.65×10^{-4} vs TREE), with t-statistics close to zero or moderately negative and p-values well above conventional significance levels, indicating no robust dominance. In late phase C, the tensor model again regained a clear edge over the linear and low-order nonlinear benchmarks (e.g. dMSE 0.00103 vs OLS, t = 6.89, $p \approx 2.3 \times 10^{-10}$; 0.00085 vs AR1, t = 6.25, $p \approx 5.9 \times 10^{-9}$; 0.00056 vs POLY, t = 3.69, $p \approx 3.3 \times 10^{-4}$), yet its performance relative to the tree-based model became statistically indistinguishable (dMSE -1.23×10^{-4}, t = -0.67, $p \approx 0.50$). In general, these patterns supported hypothesis (vi): the tensor model tended to yield the largest and most significant gains over linear baselines in periods of pronounced structural change (early and late phases), while its advantage over a flexible tree benchmark was most evident in harvesting and less marked for construction.

5.5.2. Performance in Easy vs Hard Dependence Regimes

To complement the subperiod analysis, the tensor model's performance was also contrasted in *easy* versus *hard* dependence regimes, defined on the basis of the rolling correlation between mean normalised harvesting and construction at national level. Periods with low absolute rolling correlation $|\rho_t|$ and stable co-movement were classified as easy, whereas periods with high variability and frequent sign changes of ρ_t were classified as hard, reflecting phases in which the short-run joint dynamics of harvesting and construction were more irregular and nonlinear. For each regime and channel, the focus was on the key comparison between the tensor model and the tree-based residual benchmark, which emerged as the strongest nonlinear competitor in the previous subsections.

Building on the earlier analysis of hypothesis (v), the mean difference in MSE between TREE and TENSOR, $\overline{\Delta\text{MSE}} = \overline{\text{MSE}_{\text{TREE}} - \text{MSE}_{\text{TENSOR}}}$, was summarised separately for easy and hard periods, and equality of means was tested using again Welch's t-test. For harvesting, the average dMSE was positive

in both regimes, indicating that the tensor model delivered lower MSE than the tree benchmark regardless of dependence conditions, and the corresponding p-values were small, confirming that these gains were statistically significant. However, the magnitude of the advantage did not increase markedly in the hard regime, suggesting that while the tensor model was consistently superior, and its relative edge over trees was not strongly amplified by the complexity of the dependence pattern. For construction, the picture was more nuanced: in easy periods the tensor model tended to match or slightly outperform the tree benchmark, whereas in hard periods the mean dMSE became negative and statistically significant, indicating that trees can exploit local threshold structure in construction residuals more effectively when the dependence with harvesting becomes highly unstable.

Taken together, the easy/hard comparison reinforces the view that the tensor architecture was particularly effective at capturing complex residual dynamics in the harvesting channel and remained competitive but not uniformly dominant for construction. In terms of H(vi), the results suggest that the tensor model's relative gains over simple linear and low-order nonlinear benchmarks were indeed largest in the subperiods and regimes characterised by strong non-stationarity and changing co--movement I – especially for harvesting, while its performance gap vis-à-vis a flexible tree model depended more sensitively on the specific structure of nonlinearity in the construction series.

5.6. Regional Heterogeneity and Dependence Patterns

5.6.1. Evidence of Time-Varying Dependence Across Regions

The first indication of time-varying dependence came from the rolling correlation between normalised harvesting and construction at national level, which alternated between strongly positive and clearly negative values and exhibited several pronounced turning points over the sample. This aggregate pattern already suggests that the short-run co-movement between the two sectors was not governed by a single, stable linear relation, but instead went through phases of in-phase and anti-phase behaviour, with occasional decoupling.

To assess how these features manifested themselves across the regions, region-specific rolling correlations between $\widetilde{H}_v(t)$ and $\tilde{B}_v(t)$ were computed and their variability was summarised using the variance of the rolling-correlation path and the number of sign changes per region. Table 3 reports these summary measures for all 16 voivodeships. The distribution was markedly heterogeneous: some regions, such as MAŁO, DOLN and OPOL, displayed comparatively low variance of the rolling correlation (0.17–0.22), consistent with relatively smooth, predominantly positive co-movement. By contrast, regions such as MAZO, PODL, POMO, ŚLĄS and especially WARM, exhibited a much higher variance (often above 0.30 and up to 0.46) and between three and six sign changes, indicating frequent switches between positive and negative short-run correlations.

Table 3. Variance of regional rolling correlation between normalised harvesting and construction, and number of sign changes of the rolling correlation over the sample, by region (four-letter region code, abbreviated voivodeship name)

Region	Var(rolling corr)	# sign changes
DOLN	0.218	9
KUJA	0.265	4
LUBE	0.237	3
LUBU	0.207	5
LÓDZ	0.222	2
MAŁO	0.173	2
MAZO	0.410	2
OPOL	0.175	4

Region	Var(rolling corr)	# sign changes
PODK	0.224	6
PODL	0.316	3
POMO	0.355	3
ŚLĄS	0.383	4
ŚWIĘ	0.257	3
WARM	0.459	4
WIEL	0.236	6
ZACH	0.283	6

Source: author's calculations based on Statistics Poland (GUS) Local Data Bank (Bank Danych Lokalnych, BDL) data.

Complementing this cross-sectional view, the number of times rolling correlations change sign in a given year across all regions was also counted. Table 4 shows the years with the highest total number of sign changes between consecutive observations. The peak occurred in 2020, when nine regions experienced a sign reversal in their rolling correlation between harvesting and construction. Elevated frequencies of sign changes also appeared around 2013–2014 (25–32% of voivodeships per year) and 2016–2022 (32–56%), indicating clusters of years in which many voivodeships simultaneously transitioned between phases of positive and negative short-run co-movement. Taken together, the regional dispersion in variance and sign-change counts, as well as the temporal clustering of sign reversals, provide strong empirical support for hypothesis (i): both the strength and sign of the dependence between harvesting and construction are clearly time-varying and differ substantially across voivodeships and subperiods, making it unlikely that standard linear models with constant coefficients can capture the joint dynamics adequately.

5.6.2. Tensor-Based Representation of Joint Dynamics

The tensor model represents the joint residual dynamics of harvesting and construction through a two--qubit state $|\psi_v(t)\rangle$ whose four complex amplitudes $\left(\psi_{0,v}(t), \dots, \psi_{3,v}(t)\right)$ evolved under a common unitary operator and were observed through sector-specific projectors. For each region this yielded a time series of tensor states that could be related to the observed residual amplitudes $\varepsilon_F(t,v)$ and $\varepsilon_B(t,v)$ via the induced pseudo-probabilities $r_F(t,v)$ and $r_B(t,v)$, providing a compact, four--dimensional representation of the local dependence structure.

Table 4. Years with the highest number of rolling-correlation sign changes across regions

Year	Total sign changes across regions
2020	9
2011	7
2017	7
2013	5
2016	5
2018	5
2021	5
2022	5
2014	4

Source: author's calculations based on Statistics Poland (GUS) Local Data Bank (Bank Danych Lokalnych, BDL) data.

To illustrate how this representation captured heterogeneous joint dynamics, attention focused on Dolnośląskie and Małopolskie. In Dolnośląskie, a particularly clear in-phase configuration with a dominant $|11\rangle$ component occurred around 2021, where the estimated state was

$$|\psi_{\text{DOLN}}(2021)\rangle \approx \begin{pmatrix} -0.0620 - 0.0205\,\mathrm{i} \\ -0.2060 - 0.0604\,\mathrm{i} \\ -0.1784 + 0.0838\,\mathrm{i} \\ -0.7186 + 0.6280\,\mathrm{i} \end{pmatrix},$$

so that

$$|\psi_0| \approx 0.07, \quad |\psi_1| \approx 0.21, \quad |\psi_2| \approx 0.20, \quad |\psi_3| \approx 0.95,$$

and the pseudo-probabilities

$$r_F = |\psi_2|^2 + |\psi_3|^2 \approx 0.95, \qquad r_B = |\psi_1|^2 + |\psi_3|^2 \approx 0.94.$$

Here the state placed almost all of its mass on $|11\rangle$, meaning that, conditional on the residual dynamics, both harvesting and construction qubits were very likely to be in the 'high' state simultaneously, while configurations with mixed signs ($|10\rangle$, $|01\rangle$) or jointly low residuals ($|00\rangle$) had negligible weight. This was a canonical in-phase regime at the level of residuals: positive deviations in harvesting and construction tended to co-occur, and the unitary evolution kept the state concentrated near the $|11\rangle$ corner of the two-qubit Hilbert space.

By contrast, Małopolskie exhibited a pronounced anti-phase configuration around 2016, with

$$|\psi_{\text{MAŁO}}(2016)\rangle \approx \begin{pmatrix} -0.212 - 0.044\,\mathrm{i} \\ -0.139 - 0.133\,\mathrm{i} \\ -0.198 - 0.441\,\mathrm{i} \\ -0.538 + 0.627\,\mathrm{i} \end{pmatrix},$$

so that

$$|\psi_0| \approx 0.22, \quad |\psi_1| \approx 0.19, \quad |\psi_2| \approx 0.48, \quad |\psi_3| \approx 0.83,$$

and the induced pseudo-probabilities

$$r_F = |\psi_2|^2 + |\psi_3|^2 \approx 0.92, \qquad r_B = |\psi_1|^2 + |\psi_3|^2 \approx 0.72.$$

In this configuration, a high harvesting residual (qubit $F = 1$) was compatible with both high and low construction residuals, since the probability mass conditional on $F = 1$ was split between $|10\rangle$ and $|11\rangle$, whereas a high construction residual (qubit $B = 1$) was much more frequently associated with a high harvesting residual, reflecting the dominance of $|11\rangle$ and the relatively small weight of $|01\rangle$. At the same time, the five-year rolling correlation between r_F and r_B was strongly negative (around -0.85), indicating that within the local window increases in r_F tended to coincide with decreases in r_B. This pattern was characteristic of an anti-phase regime in the residual dynamics: the trajectory of $|\psi_{\text{MAŁO}}(t)\rangle$ increasingly visited configurations in which harvesting and construction residuals moved in opposite directions across successive periods, so that positive deviations in one sector were systematically offset by negative deviations in the other, even though the level of both pseudo-probabilities remained high.

5.7. Estimated Unitary Parameters and Entanglement Strength

The starting point was the reduced four–dimensional specification in which the global unitary was written in Equation 5 with local single-qubit rotations $U_{\text{loc}} = R_y(b)$ and $V_{\text{loc}} = R_y(e)$, and an entangling core $U_{\text{ent}}(\alpha) = \cos(\alpha) I_4 - \mathrm{i}\sin(\alpha)(Z \otimes Z)$. [3] In the empirical application based on panel residuals, the estimated raw parameter vector $\theta_{\text{raw}} = (b_{\text{raw}}, \eta_{\text{raw}}, \alpha_{\text{raw}}, \phi_{\text{raw}})$ was mapped to the effective angles

[3] See Section 2 for the full construction and parametrisation of $U(\theta)$.

$$(b, \eta, \alpha, \phi) = \pi \tanh(\theta_{\text{raw}})$$

so that each component remained in $(-\pi, \pi)$ while the optimiser was free to explore $\mathbb{R}^4$.

For the preferred residual specification, the fitted angles were approximately $b \approx 0.78$, $\eta \approx 1.63$, $\alpha \approx 0.62$ and $\phi \approx 0$, as reported in *results_tensor_operator_params_2d_tensor.csv*. [4] The first two values corresponded to local R_y–rotations of about $45°$ and $94°$, respectively, implying a non–degenerate mixing of the two residual levels in each channel, both before and after the entangling step. The entangling angle $\alpha \approx 0.62$ rad yielded $\sin(\alpha) \approx 0.58$ and $\cos(\alpha) \approx 0.81$, hence $U_{\text{ent}}(\alpha)$ deviated markedly from the identity while remaining far from the maximally entangling regime. In other words, the tensor operator introduced a substantial, yet not extreme, two–way interaction between harvesting and construction residuals, consistent with the presence of nonlinear and time–varying dependence documented in this section. Finally, the global phase ϕ was estimated very close to zero as expected, since global phases did not affect the induced probabilities and therefore did not contribute to the amplitude-based loss.

For the empirically fitted operator, the estimated angles $b \approx 0.7767$, $\eta \approx 1.6313$, $\alpha \approx 0.6152$ and $\phi \approx -7.1 \times 10^{-5}$ were substituted. Numerically,

$$\cos(b/2) \approx 0.923, \quad \sin(b/2) \approx 0.385, \qquad \cos(\eta/2) \approx 0.659, \quad \sin(\eta/2) \approx 0.752$$

and

$$\cos(\alpha) \approx 0.82, \qquad \sin(\alpha) \approx 0.58, \qquad e^{i\phi} \approx 1 - i\,7.1 \times 10^{-5}.$$

The single-qubit rotations became

$$U_{\text{loc}} = R_y(0.7767) \approx \begin{pmatrix} 0.923 & -0.385 \\ 0.385 & 0.923 \end{pmatrix}$$

$$V_{\text{loc}} = R_y(1.6313) \approx \begin{pmatrix} 0.659 & -0.752 \\ 0.752 & 0.659 \end{pmatrix}.$$

The corresponding two-qubit local factors were

$$U_{\text{loc}} \otimes U_{\text{loc}} \approx \begin{pmatrix} 0.852 & -0.355 & -0.355 & 0.148 \\ 0.355 & 0.852 & -0.148 & -0.355 \\ 0.355 & -0.148 & 0.852 & -0.355 \\ 0.148 & 0.355 & 0.355 & 0.852 \end{pmatrix}$$

$$V_{\text{loc}} \otimes V_{\text{loc}} \approx \begin{pmatrix} 0.435 & -0.495 & -0.495 & 0.565 \\ 0.495 & 0.435 & -0.565 & -0.495 \\ 0.495 & -0.565 & 0.435 & -0.495 \\ 0.565 & 0.495 & 0.495 & 0.435 \end{pmatrix}.$$

With $\alpha \approx 0.6152$ the entangling kernel $U_{\text{ent}}(\alpha) = \cos(\alpha) I_4 - i\sin(\alpha)(Z \otimes Z)$ took the diagonal form

$$U_{\text{ent}}(0.6152) \approx \begin{pmatrix} 0.82 - 0.58\,i & 0 & 0 & 0 \\ 0 & 0.82 + 0.58\,i & 0 & 0 \\ 0 & 0 & 0.82 + 0.58\,i & 0 \\ 0 & 0 & 0 & 0.82 - 0.58\,i \end{pmatrix}.$$

Finally, up to the negligible global phase factor $e^{i\phi} \approx 1$, the full fitted two-qubit operator was

$$U_{\text{fitted}} = e^{i\phi}\,(V_{\text{loc}} \otimes V_{\text{loc}})\,U_{\text{ent}}(0.6152)\,(U_{\text{loc}} \otimes U_{\text{loc}}),$$

a concrete 4×4 unitary matrix obtained by multiplying the three explicit blocks above.

In the empirical application this evaluated to

[4] All the code and result files are available in the accompanying Git repository.

$$U_{\text{fitted}} \approx \begin{pmatrix} 0.1050 - 0.4780\,\mathrm{i} & -0.2734 - 0.2177\,\mathrm{i} & -0.2734 - 0.2177\,\mathrm{i} & 0.7116 - 0.0992\,\mathrm{i} \\ 0.2734 - 0.2177\,\mathrm{i} & 0.1051 + 0.4780\,\mathrm{i} & -0.7116 - 0.0991\,\mathrm{i} & -0.2734 + 0.2177\,\mathrm{i} \\ 0.2734 - 0.2177\,\mathrm{i} & -0.7116 - 0.0991\,\mathrm{i} & 0.1051 + 0.4780\,\mathrm{i} & -0.2734 + 0.2177\,\mathrm{i} \\ 0.7116 - 0.0992\,\mathrm{i} & 0.2734 + 0.2177\,\mathrm{i} & 0.2734 + 0.2177\,\mathrm{i} & 0.1050 - 0.4780\,\mathrm{i} \end{pmatrix}$$

which was unitary up to numerical precision and exhibited the expected symmetric structure across the harvesting and construction channels.

Taken together, these values indicate that the fitted unitary exploits both local nonlinear transformations and a genuinely entangling component, rather than collapsing to a purely separable (product) dynamics. This supports hypothesis (ii): the joint behaviour of harvesting and construction residuals can be captured by a low-dimensional tensor state evolving under a single, globally estimated operator whose parameters admit a transparent interpretation in terms of local rotations and interaction strength, and do not exhibit pathological or boundary-driven behaviour.

6. Discussion and Conclusions

A quantum-inspired tensor state model for bivariate time series was evaluated, with a particular focus on the joint dynamics of timber harvesting and residential construction across Polish voivodeships. The proposed specification represents the pair of sectoral processes through a shared two-qubit state whose one-step evolution was governed by a low–dimensional unitary operator. In contrast to conventional linear models with constant coefficients, this framework encoded time–varying dependence directly in the dynamics of the joint state and allowed nonlinear interactions between the two channels to be captured in a compact and interpretable manner. From a broader methodological perspective, the tensor state specification can be seen as a concrete instance of the more general idea of using Hilbert-space or tensor-network representations to model complex sequential structure in classical data, where dependence is encoded in amplitudes and their evolution rather than in linear lag coefficients.[5]

From a descriptive perspective, the empirical evidence strongly supports the working hypothesis that the short–run relation between harvesting and construction was both nonlinear and time-varying. The mean normalised series for both sectors followed broadly increasing trends, but their rolling correlation exhibited pronounced variability, including several sign reversals over the sample. This pattern was difficult to reconcile with a single constant-parameter linear relation and motivated the search for representations that could accommodate alternating phases of positive and negative co–movement, as well as periods of partial decoupling. In this context, the tensor state model provided a natural way to encode such regime-like features by allowing the joint state to explore different regions of the two-qubit configuration space.[6]

The analysis of the fitted unitary parameters indicated that the global tensor operator exploited both local nonlinear transformations and a genuinely entangling component. The estimated angles for the residual specification were moderate in magnitude rather than extreme: the pre-entanglement local rotation $R_y(b)$ performed a full 2π-cycle in roughly eight annual steps, while the post-entanglement rotation $R_y(e)$ completed a full cycle in about four years. This suggests the presence of two characteristic scales in the induced dynamics: a slower component associated with gradual mixing of residual levels in each channel, and a faster component that adjusted amplitudes after the sectors interacted through the entangling core. The entangling angle itself was clearly different from zero but far from the maximally entangling regime, implying that the operator introduced a substantial yet

[5] Related quantum-inspired sequence and time-series models reported similar benefits in terms of their expressive power and interpretability, see e.g. Harvey et al. (2025) for tensor-network sequence processing, and Vipulananthan et al. (2024) for a quantum tensor-network view on time series data.

[6] For related discussion of tensor and density-matrix representations in classical time series see e.g. García--Ripoll (2021), Viqueira et al. (2025), Fellner et al. (2024).

controlled degree of coupling between harvesting and construction residuals, rather than collapsing into either a purely separable or an almost identity-like transformation.

A complementary viewpoint emerged from examining how the fitted operator mapped the four basis configurations $|00\rangle$, $|01\rangle$, $|10\rangle$ and $|11\rangle$, which encoded low/high residual levels in the two sectors. The estimated 4×4 unitary had a highly symmetric structure: images of $|01\rangle$ and $|10\rangle$ shared the same probability profile up to a swap of coordinates, reflecting the symmetric treatment of harvesting and construction in the global operator. More importantly, the two 'extreme' configurations $|00\rangle$ and $|11\rangle$ were approximately mapped into each other with the probability mass slightly above one half, while state-mixed configurations were mapped predominantly into other mixed states. In economic terms, episodes where both residuals were simultaneously very low or very high tended to be driven towards the opposite corner of the configuration space in the next step, whereas asymmetric episodes (one sector high, the other low) were relatively persistent. This pattern was consistent with the idea that simultaneous booms or slumps in both sectors were self–correcting, while temporary imbalances between harvesting and construction could last for several years.

In terms of forecasting performance, the tensor state model generally attained mean squared errors for harvesting amplitudes that were at least competitive with, and often modestly lower than those of static linear, lag-based and reduced quantum-inspired benchmarks. The gains were most pronounced in regions and subperiods characterised by strong time-varying and nonlinear dependence, as indicated by large swings and sign changes in rolling correlations. For construction amplitudes the improvements were smaller and less systematic, suggesting that a simple linear mapping already provided a reasonable approximation for that channel. In terms of forecasting performance, the tensor state model generally attained mean squared errors for harvesting amplitudes at least competitive with, and often modestly lower than those of static linear, lag-based and reduced quantum-inspired benchmarks. The gains were most pronounced in regions and subperiods characterised by a strong time-varying and nonlinear dependence, as indicated by large swings and sign changes in rolling correlations. For construction amplitudes the improvements were smaller and less systematic, suggesting that a simple linear mapping already provided a reasonable approximation for that channel. As indicated in Table 2, mean loss differentials (benchmark minus tensor) were positive for almost all the model pairs, indicating lower average MSE and MAE for the tensor specification, especially in the harvesting channel. Formal Diebold–Mariano tests based on these one–step–ahead loss differentials did not reject the null of equal predictive accuracy between the tensor model and its benchmarks at conventional significance levels, implying that the observed reductions in mean squared error, while economically meaningful in some cases, were not statistically large relative to sampling variability on a short panel (see section 5.3. for details). Taken together, these results support the conclusion that the tensor specification was at least competitive with classical alternatives and can deliver incremental accuracy in settings where cross-sector interactions were particularly complex, yet should be viewed primarily as an interpretative and structural tool rather than a universally superior forecasting engine.

An additional strength of the proposed approach lies in the interpretability of its building blocks. The joint state encoded, for each region and year, a probability-like distribution over the four basic harvesting/construction configurations, and the global unitary operator summarised how these configurations tended to evolve over time. Since the operator was shared across voivodeships and estimated under mild regularisation, this provided a concise description of cross-regional regularities in the propagation of shocks and in the adjustment from extreme to more typical configurations. At the same time, heterogeneity across regions was captured through their distinct state trajectories, which could be visualised and analysed using the tensor components or derived marginal quantities such as residual 'activation' probabilities. This offers a novel way of exploring spatially heterogeneous but structurally related dynamics in multi-sector regional panels.

There are, however, several limitations that should be acknowledged. First, the empirical study was based on relatively short annual time series for a small number of regions, which limited the statistical power to detect subtle differences in forecast accuracy across competing models and restricts the

complexity of the unitary parametrisation that could be reliably identified. Second, the use of a single global operator for all voivodeships implied that genuinely region-specific interaction mechanisms were only captured indirectly through differences in state trajectories; richer specifications with partially regionalised gates could, in principle, provide a better fit at the cost of additional parameters. Third, the model operated on max-scaled levels and amplitude-domain residuals, hence translating the behaviour of the tensor dynamics back into statements about physical volumes of wood and housing completions required careful reinterpretation and, ideally, complementary economic analysis.

These limitations point to several promising directions for future research. One avenue is to extend the framework to multilevel or hierarchical tensor models in which some components of the unitary operator are shared across regions while others are allowed to vary, thereby balancing parsimony and local flexibility. Another is to enrich the entangling kernel by including additional two-qubit couplings or by embedding the current two-qubit block into a larger tensor-network architecture that can accommodate more sectors or auxiliary macroeconomic indicators. A third direction concerns the integration of the tensor state dynamics with established econometric structures, such as vector autoregressions or state-space models, either by using the quantum-like state as a nonlinear latent factor or by coupling classical and quantum-inspired components in hybrid forecasting schemes. Finally, applications to higher-frequency data and to other domains where nonlinear, time-varying dependence is suspected, e.g. energy markets, financial-real sector linkages or climate–economy interactions, would help to clarify the scope and robustness of the proposed methodology.

In conclusion, the quantum-inspired tensor state model offered a conceptually appealing and computationally tractable way to represent the joint evolution of related economic time series. While its forecasting gains over simple benchmarks were modest on the data considered here, the model delivered a compact, interpretable description of time-varying dependence, including nonlinear adjustment between extreme and mixed configurations, and provided a flexible foundation for further methodological and applied work at the interface of quantum-inspired modelling and classical econometrics.

References

Banaś, J., Šafařík, D., Utnik-Banaś, K., & Hlaváčková, P. (2022). Identifying long-run and short-run relationships in the European union softwood market. *Forest Policy and Economics*, *143*, 102821. https://doi.org/10.2139/ssrn.4143158

Dróżdż, K. (2022). Dynamika rynku drzewnego w Polsce w latach 2000-2020. *Academic Review of Business and Economics*, *2*, 1-19. https://doi.org/10.22367/arbe.2022.02.01

Fellner, T., Kreplin, D., Tovey, S., Holm, C., & Main, J. (2024). *Quantum machine learning for time series prediction* [Doctoral dissertation, University OF Stuttgart].

García-Ripoll, J. J. (2021). Quantum-inspired algorithms for multivariate analysis: From interpolation to partial differential equations. *Quantum*, *5*, 431. https://doi.org/10.22331/q-2021-04-15-431

Główny Urząd Statystyczny. (2026). Bank danych lokalnych (local data bank) [*GUS, Statistics Poland*. Retrieved January 22, 2026, from https://bdl.stat.gov.pl/bdl/start

Harvey, C., Yeung, R., & Meichanetzidis, K. (2025). Sequence processing with quantum-inspired tensor networks. *Scientific Reports*, *15*(1), 7155. https://dx.doi.org/10.1038/s41598-024-84295-2

Hetemäki, L., Hänninen, R., & Toppinen, A. (2004). Short-term forecasting models for the Finnish forest sector: Lumber exports and sawlog demand. *Forest Science*, *50*(4), 461-472. https://doi.org/10.1093/forestscience/50.4.461

Kapoor, S., & Garg, V. (2026). Market trends in quantum-inspired soft computing for intelligent data processing. *Quantum-Inspired Approaches for Intelligent Data Processing*, 173-200. https://doi.org/10.1002/9781394336449.ch8

Kaputa, V. (2004). Rynek materiałów drzewnych w Polsce. *Intercathedra*, *20*, 74-78.

Kotlarz, J., & Bejger, S. (2024). Estimation of the short-term impact of climate-change-related factors on wood supply in Poland in 2023–2025. *Forests*, *15*(1), 108. https://doi.org/10.3390/f15010108

Mazzaro, G. (2025). *A quantum-enhanced regime-switching model for financial time series forecasting* [Doctoral dissertation, Politecnico di Torino].

Mokhtarzadeh, F. (2021). A global vector autoregression model for softwood lumber trade. In: *International trade in forest products: Lumber trade disputes, models and examples* (pp. 174-193). CABI Wallingford UK.

Perez-Garcia, J., Lippke, B., Briggs, D., Wilson, J. B., Bowyer, J., & Meil, J. (2005). The environmental performance of renewable building materials in the context of residential construction. *Wood and Fiber Science*, 3-17.

Ratajczak, E. (2014). Zrównoważona gospodarka zasobami surowca drzewnego w Polsce. *Konsumpcja i Rozwój*, *2*(7), 15-27.

Rykowski, K. (2012). Czynniki środowiska przyrodniczego determinujące produkcję drewna. *Przyrodnicze i gospodarcze aspekty produkcji oraz wykorzystania drewna–stan obecny i prognoza*, 47.

Song, N. (2006). Structural and forecasting softwood lumber models with a time series approach. *LSU Doctoral Dissertations. 3049*. Louisiana State University; Agricultural & Mechanical College. https://repository.lsu.edu/gradschool_dissertations/3049

Udvarnoki, Z., & Fáth, G. (2025). Quantum-inspired models for classical time series. *Machine Learning and Knowledge Extraction*, *7*(2), 44. https://doi.org/10.3390/make7020044

Vipulananthan, P., Premaratne, K., Sarkar, D., & Murthi, M. N. (2024). A quantum tensor network-based viewpoint for modeling and analysis of time series data. *2024 IEEE International Conference on Knowledge Graph (ICKG)*, 378-387.

Viqueira, J. D., Faílde, D., Juane, M. M., Gómez, A., & Mera, D. (2025). Density matrix emulation of quantum recurrent neural networks for multivariate time series prediction. *Machine Learning: Science and Technology*, *6*(1), 015023. https://doi.org/10.1088/2632-2153/ad9431

Quasi-kwantowy model stanów tensorowych dla dwuwymiarowych szeregów czasowych w leśnictwie oraz budownictwie mieszkaniowym

Streszczenie

Cel: Celem artykułu jest opracowanie i empiryczna ocena kwantowo inspirowanego tensorowego modelu stanu dla dwuwymiarowych szeregów czasowych opisujących łączne zmiany pozyskania drewna i budownictwa mieszkaniowego w polskich regionach.

Metodyka: Zaproponowany model koduje dwa procesy w czterowymiarowej przestrzeni Hilberta jako stan tensorowy, którego dynamika resztowa jest generowana przez parametryzowany dwukubitowy operator unitarnej ewolucji pełniący funkcję nieliniowej korekty drugiego stopnia względem zagregowanych modeli OLS; estymacja opiera się na panelu danych obejmujących 16 województw w latach 2005–2025 oraz na porównaniu z liniowymi, autoregresyjnymi, wielomianowymi i drzewiastymi modelami reszt przy użyciu miar MSE i MAE.

Wyniki: Tensorowy model istotnie obniża globalne błędy prognozowania amplitud pozyskania drewna względem wszystkich benchmarków oraz osiąga konkurencyjne wyniki dla budownictwa mieszkaniowego, zapewniając statystycznie istotne zyski względem modeli liniowych i niskowymiarowych modeli nieliniowych oraz wysokie odsetki „wygranych" regionów w kanale leśnym, przy porównywalnej skuteczności względem płytkiego drzewa regresyjnego w kanale budownictwa.

Implikacje i rekomendacje: Uzyskane wyniki wskazują, że kwantowo inspirowane tensorowe modele stanu mogą stanowić praktyczne narzędzie do prognozowania i interpretacji zmiennej w czasie współzależności między sektorami w panelach regionalnych, wspierając planowanie i ocenę ryzyka w powiązanym systemie leśnictwo–budownictwo; dalsze badania powinny rozszerzać ramę na wielosektorowe układy, bogatsze jądra splątania oraz częściowo zregionalizowane operatory, a także testować jej zastosowania w innych dziedzinach charakteryzujących się nieliniową i zmienną w czasie współzależnością.

Oryginalność/wartość: Praca stanowi jedno z pierwszych zastosowań kwantowo podobnego dwukubitowego modelu tensorowego do klasycznych szeregów ekonomicznych, pokazując, że niskowymiarowa unitarna ewolucja może zapewnić dokładność prognoz co najmniej porównywalną z silnymi klasycznymi benchmarkami, a jednocześnie oferować zwartą i interpretowalną reprezentację wspólnej dynamiki oraz interakcji typu „splątania" między sektorami.

Słowa kluczowe: modele kwantowo inspirowane, stan tensorowy, dwuwymiarowe szeregi czasowe, leśnictwo, budownictwo mieszkaniowe