

PRZEGLĄD STATYSTYCZNY STATISTICAL REVIEW

Vol. 73 No. 3 2026

GŁÓWNY URZĄD STATYSTYCZNY
STATISTICS POLAND

INFORMATION FOR AUTHORS

Przegląd Statystyczny. *Statistical Review* publishes original research papers on theoretical and empirical topics in statistics, econometrics, mathematical economics, operational research, decision science and data analysis. The submitted manuscripts should significantly advance theoretical aspects of these fields or explore their practical applications. Manuscripts presenting important results of research projects are particularly welcome. Review papers, shorter papers reporting on major conferences in the field, and reviews of seminal monographs are eligible for publication based on the Editor-in-Chief's decision.

Since 1st May 2019, the journal has been publishing articles in English.

Any spelling style is acceptable as long as it is consistent within the manuscript.

All papers should be submitted to the journal through the Editorial System (<https://www.editorialsystem.com/pst>).

For details of the submission process and editorial requirements, please visit <https://ps.stat.gov.pl/ForAuthors>.

PRZEGLĄD STATYSTYCZNY STATISTICAL REVIEW

Vol. 73 No. 3 2026

ADVISORY BOARD

Krzysztof Jajuga – Chairman (Wrocław University of Economics and Business, Poland), Marek Gruszczyński (SGH Warsaw School of Economics, Poland), Tadeusz Kufel (Nicolaus Copernicus University in Toruń, Poland), Igor G. Mantsurov (Kyiv National Economic University, Ukraine), Jacek Osiewalski (Krakow University of Economics, Poland), D. Stephen G. Pollock (University of Leicester, United Kingdom), Sven Schreiber (Macroeconomic Policy Institute, Germany), Mirosław Szreder (University of Gdańsk, Poland), Matti Virén (University of Turku, Finland), Aleksander Welfe (University of Łódź, Poland), Janusz Wywił (University of Economics in Katowice, Poland)

EDITORIAL BOARD

Editor-in-Chief: Krzysztof Echaust (Poznań University of Economics and Business, Poland)
Co-Editors: Piotr Fiszeder (Nicolaus Copernicus University in Toruń, Poland), Michał Jakubczyk (SGH Warsaw School of Economics, Poland), Bogumił Kamiński (SGH Warsaw School of Economics, Poland), Gábor Dávid Kiss (University of Szeged, Hungary), Aleksandra Łuczak (Poznań University of Life Sciences, Poland), Silvana Musti (University of Foggia, Italy), Maciej Nowak (University of Economics in Katowice, Poland), Monika Papież (Krakow University of Economics, Poland), Emilia Tomczyk (SGH Warsaw School of Economics, Poland), Łukasz Woźny (SGH Warsaw School of Economics, Poland)

EDITORIAL OFFICE ADDRESS

Statistics Poland (GUS), al. Niepodległości 208, 00-925 Warsaw, Poland

Language editing: Scientific Publications Division, Analyses and Dissemination Department, Statistics Poland
Technical editing and typesetting: Statistical Publishing Establishment, Statistical Computing Center – team supervised by Mariusz Męcina



Centrum Informatyki
Statystycznej

Printed and bound by: Statistical Computing Center
al. Niepodległości 208, 00-925 Warsaw, Poland, cis.stat.gov.pl

Website: ps.stat.gov.pl

© Copyright by Główny Urząd Statystyczny and the authors, some rights reserved. CC BY-SA 4.0 licence



ISSN 0033-2372
e-ISSN 2657-9545
Index 371262

Information on the sales of the journal: Statistical Computing Center
Phone no.: +48 22 608 32 10, +48 22 608 38 10

Order no. 227/2026

CONTENTS

Andrzej Geise

Artificial intelligence and energy intensity nexus in European Union economies:
The moderating role of digital transformation **1**

Krzysztof Siegieda

Convergence to multiple price equilibria in algorithmic collusion **36**

Emilia Tomczyk

The impact of the COVID-19 shock on response distributions in business tendency
surveys: An application of the Kullback–Leibler divergence **56**

Ewa Konarzewska-Gubała

In memory of Professor Władysław Bukietyński (1928–1996) **74**

Artificial intelligence and energy intensity nexus in European Union economies: The moderating role of digital transformation

Andrzej Geise^a

Abstract. Reduction of energy intensity is the central challenge to the European Union's green and digital transition. The role of artificial intelligence innovation in this process has still not been fully analyzed. In particular, there is limited empirical evidence on whether digital maturity conditions the relationship between AI innovation and energy intensity across EU economies. Our study examines the impact of the AI innovation capacity on energy intensity in the EU, with digitalization moderating the effect. Using static panel data models for EU countries over the period of 2017–2022, the analysis employs fixed-effects models and proxies AI innovation capacity by AI-related patents. The main results show that AI innovation capacity is negatively associated with energy intensity, suggesting that knowledge-based AI innovation may contribute to more energy-efficient economic performance. However, this relationship is heterogeneous and depends on the level of digital development. EU economies with a level of digitalization above the median (more digitally advanced) can better translate AI-related innovation into energy-efficiency outcomes, whereas in economies with the level of digitalization index below the median (less digitally advanced), the effect is weaker and less robust. These findings highlight the need to integrate AI policy, digitalization strategies and energy-efficiency measures into the EU's twin transition agenda.

Keywords: artificial intelligence, energy intensity, digitalization, innovation, panel models

JEL: Q43, Q48, C23

1. Introduction

Rapid advancement of artificial intelligence (AI) is actively reshaping the challenges of socio-economic development across economies. The use of AI continues to expand across sectors, including industry, transport, energy, services, trade, healthcare, education (Ghahramani et al., 2020; Zhai et al., 2021), and markets – mostly financial, capital, commodity and derivatives (Abe & Nakayama, 2018; Cao, 2022). In the above-mentioned sectors, AI is primarily used to optimize a broad range of decisions. From the point of view of AI adaptability and its potential benefits, increased energy consumption represents a crucial challenge, where AI solutions can deliver the greatest improvements in energy efficiency and the sustainability of the economy.

^a Nicolaus Copernicus University in Toruń, Faculty of Economic Sciences and Management, Institute of Economics and Finance, Department of Econometrics and Statistics, ul. Gagarina 13A, 87-100 Toruń, Poland, e-mail: a.geise@umk.pl, ORCID: <https://orcid.org/0000-0003-4581-5484>.

Within the framework of production theory, energy, together with capital and labor, constitutes the fundamental input in the production function (Stern, 2011). At the same time, energy consumption based primarily on fossil fuels (coal and crude oil) generates substantial CO₂ emissions, which drive global warming, accelerate climate change, and contribute to the degradation of ecosystems and biodiversity (Bölük & Mert, 2014; Intergovernmental Panel of Climate Change [IPCC], 2023; Jie & Rabnawaz, 2024; Le, 2022; Shayanmehr et al., 2020; Zakari et al., 2022). As energy demand continues to rise, so do greenhouse gas emissions. Therefore, coordinated activity to limit emissions is necessary. In the EU, energy production and use account for more than 75% of the greenhouse gas emissions (European Commission, 2022). Reducing CO₂ emissions is essential to prevent further environmental degradation and climate change, and this reduction can be achieved by improving energy efficiency, lowering the economy's energy intensity, shifting to low-carbon energy sources, and promoting sustainable consumption patterns (Altın, 2024; Lei et al., 2022).

To achieve energy efficiency and advance Sustainable Development Goals (SDGs), politicians need to address key issues related to energy sources and environmental constraints. A low-carbon economic transformation requires a transition from the conventional, energy-intensive growth model into one that prioritizes green energy, technological innovation and digitalization as the core of sustainable development (Rame et al., 2024). In the European Union, shifting the economy's energy-intensive model and reducing energy intensity are directly linked to climate policy objectives that aim to decouple economic growth from energy consumption and reduce greenhouse gas emissions (European Commission, 2022). By lowering energy intensity, member states can progress towards their binding targets for emission reductions and renewable energy shares while maintaining economic competitiveness.

This article focuses on the main question: how does the AI innovation capacity affect the energy intensity of the economy in EU countries, and whether the digital transition is accelerating these effects. To answer this question, we designed an empirical study and a set of several specific goals:

- Goal 1: To review and present the existing definitions of artificial intelligence and to discuss the indicators that can be used to proxy AI development, adoption, and innovation capacity in empirical research;
- Goal 2: To assess the impact of AI innovation capacity on energy intensity in EU countries with regard to digital maturity (low vs high digitalization level across EU members);

- Goal 3: To assess the significance of accelerating AI innovation capacity through digitalization for energy intensity in EU member states during and after the pandemic.

This study examines the effect of AI innovation capacity (measured by the number of AI-related patents) on energy intensity in EU member states taking into account the moderating role of the level of digitalization. We built panel-data models for 27 EU countries from 2017 to 2022.

Our article contributes to the existing literature in several ways. First, in terms of theory, we systematize the concepts of artificial intelligence applied to empirical research, offering clear guidelines as to which variables can capture the development and implementation of AI innovation at the national level. This approach helps to fill a gap in the existing literature, where AI has often been narrowly addressed, only as a part of digitalization. Second, in the methodological section, we analyze AI in a wider context of digital transformation, incorporating a comprehensive set of control variables. These include fundamental factors of production (capital and labor), indicators reflecting the structure of energy consumption, and the factors of an innovative economy, providing a fuller picture of the determinants of energy intensity. Third, the empirical contribution of the study lies in examining a dataset covering 27 EU member states over the period of 2017–2022. While most existing research has focused on the role of digitalization, this article treats AI as the primary factor in the analysis, highlighting its distinct and independent impact on energy intensity. Finally, we emphasize AI's dual role both as a technological tool that improves energy efficiency and a catalyst of broader structural transformation. We further provide evidence for policymakers that AI requires dedicated strategies rather than being addressed solely within general digitalization policies.

2. Literature review – energy intensity, AI and digitalization

2.1. Artificial intelligence and proxy variables

The first references to artificial intelligence appeared in the literature in the 1950s, when AI was defined as the science of creating intelligent machines, with a focus on programmability and the simulation of human-like capabilities (McCarthy et al., 2006). Since the 1950s, the definition of AI has evolved from intuitive descriptions to more precise and system-oriented formulations.

In empirical research on the economic and environmental impact of AI, the challenge is to use valid proxy variables that capture the essence of AI, which is defined in both functional and capability-oriented terms. The term artificial intelligence refers to “a machine-based system that can, for a given set of human-defined

objectives, make predictions, recommendations, or decisions influencing real or virtual environments”, thereby highlighting task-oriented applications (Organisation for Economic Co-operation and Development [OECD], 2019). This definition has been widely adopted in policy frameworks, such as the European Union AI Act. We define AI as “the capability of an engineered system to acquire, process, and apply knowledge and skills”, reflecting a broader technological potential applicable across systems and industries (ISO/IEC 22989:2022). In both cases, these abstract definitions are not directly measurable and thus must be operationalized through observable indicators.

To distinguish between the two concepts, empirical studies typically differentiate between narrow AI, which involves domain-specific applications (e.g. robotics, predictive algorithms), and general AI, which denotes human-level flexibility and reasoning. It should be observed that general AI is currently beyond real-world implementation. Given this distinction and the lack of standardized AI metrics, researchers often resort to proxy variables. Artificial intelligence is commonly proxied in empirical research by indicators such as the number or stock of installed industrial robots (Acemoglu & Restrepo, 2022; Gao & Wang, 2025; Graetz & Michaels, 2018; Huang et al., 2022; Lee et al., 2024; J. Liu et al., 2022; L. Liu et al., 2021; Xie & Wang, 2024; Zhang et al., 2022). In this context, AI refers to automated, intelligent executive systems deployed in industry, often associated with robotization driven by algorithmic decision-making. Such robot-based measures tend to reflect the extent of intelligent manufacturing rather than offering a comprehensive view of AI development.

A more holistic perspective is provided by AI patent data, which captures AI as a knowledge-based technology and represents innovation capacity, a field of innovative activity marked by the generation of novel technological solutions. This conceptualization is widely applied in the economics of innovation and knowledge, particularly in regional studies assessing technological capabilities and the diffusion of general-purpose technologies – GPTs (Biggi et al., 2025; Chun & Hwang, 2025; Kopka & Grashof, 2022; Podrecca et al., 2024; Xie & Wang, 2024). Still, some limitations remain, as many studies rely on narrow keyword sets in identifying AI-related patents.

More systemic approaches employ AI concept indices that capture the technological, institutional, and market dimensions of AI development, reflecting its multifaceted nature (Lyu et al., 2025). Given this heterogeneity, the choice of the AI proxy is critical, as it directly affects the interpretation of AI’s role in energy intensity and the robustness of policy-relevant conclusions. An overview of proxy variables used in previous empirical studies is provided in the Appendix, Table A2.

Proxy variables for AI differ substantially in scope and interpretation, ranging from measures of technological adoption (e.g. industrial robots) to indicators of knowledge creation or policy commitment, such as AI-related patents or national AI strategies. In this study, AI is proxied by the number of AI-related patents, reflecting its role as a knowledge-based innovation activity.

2.2. AI and energy improvements in the economy

The pathway to fostering a sustainable energy policy requires integrating digital technologies and artificial intelligence, reducing energy intensity across sectors, and enhancing collaboration among stakeholders. Recent literature has confirmed that energy intensity is driven by technological change, which has mostly been ICT (information and communication technology) and AI in recent years. So, the increased use of AI contributing to technological change may have an impact on energy intensity (Liu et al., 2021). Also, given the close interrelationships among energy intensity, energy efficiency, energy consumption and energy transition, the literature review takes into consideration all aspects relevant to the empirical study (see Table A1).

The positive impact of AI use in various sectors has been confirmed by several researchers (Gao & Wang, 2025; Kopka & Grashof, 2022; Kwilinski et al., 2023; Lee et al., 2024; X. Liu et al., 2021; Lyu et al., 2025; Melnychenko, 2019; Nishant et al., 2020; Xie & Wang, 2024; Zhang et al., 2022). AI can enhance energy efficiency and reduce energy intensity by optimizing energy consumption patterns (Nishant et al., 2020), improving performance within energy-intensive industries (Kopka & Grashof, 2022), and monitoring equipment performance in real time (Morand et al., 2024). It can also foresee failures before they occur (Ucar et al., 2024; Yadav et al., 2024), predict the volume of energy generated from renewable resources (Şerban & Lytras, 2020), and reduce reliance on fossil fuels for energy backup (International Energy Agency [IEA], 2025), all of which contribute to decreasing overall energy intensity.

The relationship between AI and energy intensity is further explored in the literature. Liu et al. (2021) provided empirical evidence showing that AI adoption significantly reduced energy intensity in China's industrial sector. Their study shows that AI not only enhances output but also cuts down energy consumption, particularly in industries with high initial energy intensity. Also, AI mitigates energy vulnerability by accelerating structural changes in industry to improve energy efficiency and reduce energy intensity, catalyzing green innovation and optimizing governance capabilities (Gao & Wang, 2025; L. Liu et al., 2021). Similarly, Zhang et al. (2022) found that the deployment of industrial intelligence reduces energy intensity by improving processes and optimizing resource use, thereby corroborating previous findings (Zhang et al., 2022). Such a reduction in energy intensity is vital for advancing industrial sustainability and alleviating environmental concerns.

AI systems have also demonstrated a strong ability to improve energy efficiency across the economy and various sectors. For instance, Luan et al. (2022) highlighted the critical role of industrial robots in enhancing energy efficiency and observed their positive influence on air pollution, showing that improved manufacturing processes can lead to lower energy consumption and better environmental outcomes (Luan et al., 2022). Kwilinski et al. (2023) emphasized the transformative impact of digitalization on energy efficiency in the European Union, asserting that technologies such as smart grids and data analytics can significantly optimize energy use and minimize waste (Kwilinski et al., 2023). These advancements underscore the potential of AI-driven solutions to facilitate energy-efficient practices key to sustainable development.

AI plays a crucial role in optimizing energy consumption patterns. Dzwigol et al. (2024) explore how digitalization shifts energy consumption frameworks across EU countries, demonstrating that technological innovations lead to substantial reductions in overall energy use while allowing a more efficient distribution. Ghahramani et al. (2020) emphasize that AI-driven smart manufacturing processes can lead to reductions in energy consumption by optimizing operational workflows and allowing real-time adjustments based on data analytics. These adjustments ensure resources are used more wisely, thus contributing to greater sustainability. Furthermore, AI is instrumental in facilitating broader energy transition towards sustainable and renewable energy systems. Researchers also underline the potential of AI to support carbon-neutrality efforts by optimizing industrial emissions-reduction strategies through advanced predictive analytics and operational enhancements. AI is actively promoting reductions in carbon intensity, and this effect is further strengthened by green total factor productivity and energy efficiency (Li et al., 2022; Nishant et al., 2020). This capacity for data-driven decision-making is pivotal as global energy systems are moving increasingly towards incorporating renewables, and they need AI to manage the intricacies of this transition effectively.

3. Research Design and Methodology

3.1. Study Design

The main variables used in the study include: energy intensity, AI innovation capacity, and the level of digitalization of economies.


Energy intensity is the dependent variable, measured as primary energy use per unit of GDP, expressed in 2015 purchasing power parity. This indicator captures the amount of primary energy required to produce one unit of economic output and therefore reflects the energy efficiency of an economy.

The AI innovation capacity is proxied by the number of AI-related patent applications filed with the European Patent Office (EPO). In the EPO statistics, artificial intelligence is treated as a technological area within computer technologies and includes AI-related subfields such as neural networks, machine learning, pattern recognition, and image and video recognition (OECD, 2025). The variable used in our study measures AI-related patenting activity and thus serves as a proxy for AI innovation capacity, rather than as a direct measure of AI implementation. The list of Cooperative Patent Classification (CPC) subclasses and keywords used to identify AI-related patents is provided in the report referenced above, titled *Identifying emerging AI technologies using patent data: A semi-automated approach* by OECD (2025).

Digital transformation is measured using the DESI index, which provides a synthetic measure of digital maturity across EU economies. The index captures several dimensions of digital development, including human capital, connectivity, digital technology integration, and digital public services (European Commission, n.d.). In this study, the aggregate DESI index is used to capture the overall level of digital transformation and to assess its moderating role in the relationship between AI innovation capacity and energy intensity.

The empirical analysis covers the 27 member states of the European Union and is based on annual data for the period from 2017 to 2022. A visual overview of the study’s design is presented in Figure 1.

Figure 1. Research diagram



Source: author’s work.

The empirical design of this study integrates two dimensions related to the level and dynamics of digitalization across EU member states in order to examine the impact of knowledge-based AI proxy on energy intensity (Figure 1). The horizontal dimension captures cross-country variation by distinguishing two clusters of countries based on their DESI scores – those above and below the median level of digital development. The classification of economies into the EU-14 and EU-13 groups is time-invariant. Countries were assigned to these groups on the basis of their average DESI over the analyzed period, meaning they do not move between groups from year to year. This allows the assessment of whether the AI–energy intensity relationship differs across the levels of digital infrastructure and technological readiness.

The vertical dimension addresses temporal dynamics by segmenting the study period into two samples (sample 1: 2017–2019; sample 2: 2017–2022), with a structural break in 2020, coinciding with the COVID-19 pandemic and the associated acceleration in digital adoption (e.g. remote work, digital services). This division allows an assessment of whether the pace of digital transformation influences the strength and direction of the AI–energy intensity nexus over time.

Three research hypotheses were formulated on the basis of the planned study:

- *H1*: AI innovation capacity contributes to reducing energy intensity across EU economies;
- *H2*: The relationship between AI innovation capacity and energy intensity is moderated by the level of digital development, which means that the marginal effect of AI innovation capacity on energy intensity changes along the varying level of DESI;
- *H3*: The pandemic-induced structural shift amplifies the moderating effect of digital development on the relationship between AI innovation capacity and energy intensity due to the accelerated digital transformation.

3.2. Theoretical framework

The theoretical framework used in the analysis is based on Schumpeterian endogenous growth theory. This theory provides the basics for explaining energy intensity as a function of innovation-driven technological and structural changes. Endogenous growth theories posit that technological advancement stems from innovative activities undertaken by firms and entrepreneurs, driven by profit incentives. The Schumpeterian perspective highlights that innovation enhances productivity and fosters long-term economic growth, while the process of “creative destruction” enables the replacement of obsolete technologies with new ones, thereby facilitating sectoral transformation and sustained development (Aghion et al., 2013, 2015a, 2015b).

The Schumpeterian framework emphasizes that innovation fosters knowledge spillovers, which drive technological progress and productivity growth. They are beneficial only when firms and countries have sufficient absorptive capacity to assimilate and utilize external knowledge (Aghion & Jaravel, 2015).

In the context of digital economy, AI can be interpreted as a frontier technology that expands the economy's knowledge and shifts the technological frontier outward. AI-based innovations improve predictive capabilities, process automation, resource optimization and decision-making efficiency. In energy-intensive sectors, these mechanisms may substantially improve the efficiency of transforming energy inputs into economic output. Therefore, AI innovation might be viewed as a modern "Schumpeterian" engine of technological advancement that has a direct impact on energy efficiency.

Following the endogenous growth framework, aggregate output in country i at time t is represented by an augmented Cobb-Douglas production function:

$$Y_{it} = A_{it} K_{it}^{\alpha} L_{it}^{\beta} \vartheta_{it}^{\delta}, \quad (1)$$

where Y denotes the total output, K represents physical capital, L stands for labor input, A denotes the economy's technological stock, and ϑ represents other production determinants.

According to Aghion and Jaravel (2015), technological progress is assumed to depend on both innovation intensity and absorptive capacity. Consequently, the evolution of technological knowledge (growth rate of technology) may be expressed as:

$$\frac{\dot{A}_{it}}{A_{it}} = g = \psi AI_{it} \phi_{it}, \quad (2)$$

where AI denotes the innovation capacity of artificial intelligence proxied by AI-related patents, ψ stands for the productivity of innovation activities, and ϕ represents the absorptive capacity of an economy.

Consistent with the theory of knowledge spillovers, the absorptive capacity determines the extent to which economies can internalize external technological knowledge. In the digital era, absorptive capacity increasingly depends on the advancement of digital infrastructure, digital skills, data accessibility and institutional readiness (Aghion & Jaravel, 2015). Therefore, this study assumes that absorptive capacity is a function of digital transformation:

$$\phi_{it} = f(D_{it}), \quad (3)$$

where D represents the level of digital transformation.

Energy intensity is defined as:

$$EI_{it} = \frac{E_{it}}{Y_{it}} = \frac{E_{it}}{A_{it} K_{it}^{\alpha} L_{it}^{\beta} \vartheta_{it}^{\delta}}, \quad (4)$$

$$\ln(EI_{it}) = \ln(E_{it}) - \ln(Y_{it}) = \ln(E_{it}) - \ln(A_{it}) - \alpha \ln(K_{it}) - \beta \ln(L_{it}) - \delta \ln \vartheta_{it}. \quad (5)$$

Equation (4) demonstrates that technological progress reduces energy intensity by increasing the productivity of conventional production factors. Since AI accelerates technological upgrading through automation, predictive maintenance, real-time optimization and improved allocation of productive resources, greater AI innovation capacity should reduce the energy required to generate one unit of economic output.

Moreover, the Schumpeterian literature suggests that the economic benefits of innovation are heterogeneous across countries, depending on their proximity to the technological frontier and their absorptive capacity. Accordingly, the energy-efficiency effects of AI are expected to be stronger in digitally advanced economies, where complementary digital assets facilitate the diffusion and implementation of AI-based technologies (Yuan et al., 2009).

3.3. Conceptual model

Improvements in energy intensity may indirectly result from the adoption of energy-saving innovations, such as artificial intelligence, digital process optimization, and cleaner industrial technologies (Yuan et al., 2009). In this context, AI can be conceptualized as a general-purpose technology capable of reshaping production processes, resource allocation and energy use across sectors, thereby contributing to more energy-efficient economic growth (Acemoglu et al., 2018). Figure 2 reflects this logic by showing that technological progress may increase energy demand through growth, but can also reduce energy intensity when embodied in AI-based efficiency-enhancing solutions.

3.4. Analytical framework

Given the limited availability of data, the initial strategy for modeling relationships focused on static models. Consequently, the pooled OLS estimator for panel data assumes that the units are homogeneous and that the errors are identically distributed across individuals (i) and time (t). However, in practice, these assumptions often prove to be unrealistic. To account for heterogeneity in the panel data analysis, the fixed-effects (FE) estimator is commonly used. In such a model, the individual-specific effects are treated as fixed parameters. Cross-sectional differences are reflected in varying intercept values in this approach. The estimation is performed using a within transformation, which eliminates individual effects by subtracting time-varying variables. This allows consistent parameter estimation on the transformed panel data.

Subtracting time averages $\bar{y}_i = 1/T \sum_t y_{it}$ from the initial model

$$y_{it} = \alpha + x'_{it}\beta + z'_i\gamma + c_i + u_{it} \quad (6)$$

yields the *within model*

$$\dot{y}_{it} = \dot{x}'_{it}\beta + \dot{u}_{it}, \quad (7)$$

while $\dot{y}_{it} = y_{it} - \bar{y}_i$, $\dot{x}_{itk} = x_{itk} - \bar{x}_{ik}$ and $\dot{u}_{it} = u_{it} - \bar{u}_i$. Please note the individual-specific effect c_i , the intercept α , and the time-invariant regressors z_i .

The fixed effects estimator or the within estimator of slope coefficient β estimates the within model by OLS:

$$\hat{\beta}_{fe} = (\ddot{X}'\ddot{X})^{-1}\ddot{X}'\ddot{y}. \quad (8)$$

Please note that parameters α and γ are not estimated by the within estimator.

4. Empirical Research

4.1. Data description and model specification

To explain energy intensity across European Union countries, we incorporate the fundamental production factors, i.e. capital and labor, into control variables that capture the energy sector's characteristics and the economy's structural composition. These controls include fossil fuel, renewable energy consumption, and research and development expenditure (Table 1).


Table 1. Variables and data sources used in empirical research

Variable	Name	Description	Source
<i>Dependent variable</i>			
EI_{it}	Energy intensity	Log of energy intensity level of primary energy (MJ/2015 PPP GDP)	World Bank
<i>Explanatory variables: AI, digitalization and technology development</i>			
AIP_{it}	AI-related patents	Log of number of applicants	EPO (European Patent Office)
$DESI_{it}$	Digital economy and society index	Log of index	Eurostat
ACC_{it}	Acceleration of AI through digitalization	$\log AIP_{it} \times \log DESI_{it}$	Author's calculations
<i>Control variables</i>			
<i>Production factors, energy variables</i>			
K_{it}	Capital	Log of gross fixed capital formation (constant 2015 million US\$)	World Bank
UR_{it}	Labor	Log of unemployment rate, total (% of working-age population, 15-64yo)	World Bank
RD_{it}	Research and development expenditure	Log of % of GDP	World Bank
$FUEL_{it}$	Fossil fuel	Log of fossil fuel energy consumption (% of total)	World Bank
$RENEW_{it}$	Renewable energy consumption	Log of % of total final energy consumption	World Bank

Note. All interaction and nonlinear terms were constructed after logarithmic transformation.

Source: author's work.

Figure 3 (A–C) shows clear differences between the EU-13 and EU-14 groups regarding the changes in the average energy intensity, AI innovation capacity and digital maturity. From 2017 to 2022, energy intensity decreased in both groups, but in the EU-13 consistently more so than in the EU-14. At the same time, the mean AI innovation capacity was significantly higher in the EU-14, indicating a much greater concentration of AI-related patent activity in more digitally advanced economies. By the same token, the DESI index continued to increase in both groups, but in the EU-14 group significantly more so. Overall, Figure 3 suggests that digitally mature economies combine lower energy intensity with higher innovation capacity and digital transformation, whereas the EU-13 economies are characterized by higher energy intensity and weaker AI innovation capacity.

Figure 3. Evolution of key variables in EU-13 and EU-14 economies in 2017–2022

Source: author's work.

This study uses the following model, in which all variables are expressed in logarithms:

1) model without control variables:

$$EI_{it} = \beta_1 AIP_{it} + \varphi_1 DESI_{it} + \varphi_2 DESI_{it}^2 + \gamma_1 ACC_{it} + FE_i + \varepsilon_{it}, \quad (9)$$

2) model with control variables:

$$EI_{it} = \beta_1 AIP_{it} + \varphi_1 DESI_{it} + \varphi_2 DESI_{it}^2 + \gamma_1 ACC_{it} + FE_i + \alpha X_{it} + \varepsilon_{it}, \quad (10)$$

while EI_{it} is the annual average energy intensity of primary energy in country i over a specific period t from 2017 to 2022. $ACC_{it} = AIP_{it} \times DESI_{it}$ is the interaction between AI-related patents (AIP_{it}), and the $DESI_{it}$ index, which captures the acceleration of the AI effect on energy intensity (EI) through the digitalization level, and X_{it} is the set of control variables (see Table 1).

Given the presence of the acceleration factor (ACC_{it}), it is necessary to interpret the relationship between AI and EI in terms of marginal effects. As emphasized by Brambor et al. (2006), multiplicative interaction models imply that the intensity of the effect of one variable depends on the level of another.

Following this approach, the marginal effect of AI innovation capacity is calculated as the partial derivative of energy intensity with respect to the digitalization level. In the baseline specification, the effect is given by Brambor et al. (2006):

$$\frac{\partial EI_{it}}{\partial AIP_{it}} = \beta_1 + \gamma_1 DESI_{it}, \quad (11)$$

where β_1 denotes the coefficient of AI innovation capacity, and γ_1 represents the coefficient of the interaction term.

The use of marginal effects also improves the substantive interpretation of the estimated models. Williams (2012) argues that marginal effects and adjusted predictions help translate regression coefficients into quantities that are easier to interpret in substantive and practical terms. Moreover, graphical presentation of marginal effects with confidence intervals provides a clearer way to assess how the estimated relationship changes across different values of the moderating variable (Williams, 2012).

4.2. Empirical models for energy intensity in EU countries: AI and digitalization in energy intensity models for the pre-2020 period and the whole sample

The models are estimated using an FE framework, which accounts for the unobserved, time-invariant heterogeneity across countries. This approach enables us to identify within-country changes in the AI-energy intensity relationship and to examine the moderating role of digitalization across the EU-27 sample and the EU-14 and EU-13 subsamples. Based on the Wald, Breusch-Pagan and Hausman tests, the fixed-effects estimator was selected over the pooled OLS and random-effects specifications. The corresponding p -values were below the 0.05 significance level, justifying the use of the FE framework. Given the short panel and the strong common time variation in key variables, year fixed effects (two-way FE) could absorb a substantial share of the variation relevant to the research question. Therefore, country fixed effects are used to control for time-invariant heterogeneity, while common shocks such as the COVID-19 pandemic are acknowledged as a limitation.

The results of panel regression models are presented in Table 2 (for the pre-2020 sample) and Table 3 (for the whole sample). In the EU-27, EU-14 and EU-13 models, the DESI index shows generally negative and statistically significant coefficients (φ_1 or φ_2) in both the pre-2020 sample and the whole one. In the pre-2020 sample, the effect of digitalization on energy intensity is linear across all (1)–(6) models. In the whole sample, however, this linear relationship is true for the EU-14 and EU-13 clusters only; the EU-27 model results indicate a non-linear pattern. The EU-27 model for the whole sample exhibits an inverted-U relationship, as indicated by the negative and significant coefficient of φ_2 . Based on the whole sample, we observed that digitalization has a stronger impact on energy intensity in the EU-14 group (−0.808% without controls and −0.614% with controls) than in the EU-13 group (−0.329% without controls and −0.235% with controls). The energy-efficiency benefits of digital transformation may be more visible in digitally mature economies. That might imply that digital technologies, such as smart grids, IoT devices or predictive maintenance

can substantially improve energy optimization, thereby reducing energy consumption and emissions (Khan et al., 2025).

In the pre-2020 sample, the elasticity coefficients of AI innovation capacity (β_1) are negative across all the estimated specifications, i.e. in all the (1)–(6) models. The negative coefficients may serve as initial evidence that increased AI innovation capacity is associated with reduced energy intensity. However, this relationship is statistically significant only in one specification – the EU-27 without controls (model 1). This limitation is particularly relevant because the pre-2020 sample covers only three annual observations (2017–2019), which may reduce the precision of the estimates and inflate standard errors. For this reason, the pre-2020 results should be interpreted with considerable caution and treated mainly as a pre-pandemic benchmark and a starting point for subsequent analyses. The main interpretation of the study is based on the whole-sample estimates and the marginal effects (2017–2022).

Table 2. Regression results of energy intensity models for EU-27, EU-14 and EU-13 for the pre-2020 sample (2017–2019)

Variable	Model (1)	Model (2)	Model (3)	Model (4)	Model (5)	Model (6)
	EU-27	EU-27	EU-14	EU-14	EU-13	EU-13
AIP_{it}	–0.1394 (0.048) ***	–0.1564 (0.1135)	–0.2752 (0.3043)	–0.2586 (0.1937)	–0.3248 (0.1994)	–0.4309 (0.1681)
$DESI_{it}$	–0.5701 (0.0603) ***	–0.2343 (0.1507)	–0.6886 (0.2882) **	–0.4214 (0.2016) *	–0.6301 (0.138) ***	–0.4505 (0.1541) *
ACC_{it}	0.0328 (0.0129) **	0.0382 (0.0306)	0.0677 (0.0785)	0.0689 (0.0502)	0.0889 (0.0572)	0.1210 (0.0474)
RD_{it}	–	–0.0690 (0.1073)	–	–0.1866 (0.146)	–	0.2426 (0.1569)
K_{it}	–	–0.0733 (0.045)	–	–0.0337 (0.0477)	–	–0.0899 (0.1240)
UR_{it}	–	0.0831 (0.0812)	–	0.0828 (0.0929)	–	0.2604 (0.1375)
$FOSIL_{it}$	–	0.6628 (0.1660) ***	–	0.691 (0.1991) ***	–	–0.3109 (0.2376)
$RENEW_{it}$	–	–0.0055 (0.1007)	–	–0.1148 (0.1350)	–	–0.0240 (0.0987)
<i>Constant</i>	3.216 (0.1959) ***	–0.1023 (1.0724)	3.7082 (1.0619) ***	0.4562 (1.2246)	3.3509 (0.4703) ***	4.5750 (1.567) ***
<i>N</i>	27	27	14	14	13	13
<i>T</i>	3	3	3	3	3	3
Estimator	Fixed effect	Fixed effect	Fixed effect	Fixed effect	Fixed effect	Fixed effect

Note. *, **, *** – significance at 10%, 5% and 1%, respectively.

Source: author's calculations.

The elasticity coefficients in the whole sample suggest that AIP_{it} in EU member states is negatively and significantly associated with energy intensity. Considering the

EU-27, a 1% rise in AI-related patents is associated with an average decline of 0.12–0.16% in energy intensity (models 7 and 8). This effect is substantially stronger in EU-14, where a 1% increase in AI-related patents is associated with an average decrease of –0.31% (without controls) and –0.28% (with controls) in EI (models 9–10). This result suggests that EU-14 economies may have greater absorptive capacity and more advanced complementary digital infrastructure. In the EU-13, the elasticity coefficients are smaller than those in the EU-14 models (about –0.13% without controls and –0.15% with controls), and the coefficients (β_1) are only marginally significant after adding controls ($\alpha=10\%$). This might indicate that in economies with weaker digital infrastructure the relationship between AI innovation capacity and energy-efficiency improvements is weaker (Table 3).

At the same time, it should be emphasized that the direct interpretation of these coefficients is limited, as they do not show how the effect of AI innovation capacity changes as the level of digital transformation does. For this reason, further interpretation should focus on marginal effects calculated for different values of DESI, which may indicate that the actual impact of AI innovation capacity is weaker than suggested by the baseline coefficients.

Table 3. Regression results of energy intensity models for EU-27, EU-14 and EU-13 for the whole sample (2017–2022)

Variable	Model (7)	Model (8)	Model (9)	Model (10)	Model (11)	Model (12)
	EU-27	EU-27	EU-14	EU-14	EU-13	EU-13
AIP_{it}	–0.1607 (0.07) **	–0.1194 (0.07) *	–0.3056 (0.126) **	–0.2808 (0.119) **	–0.1262 (0.081)	–0.1516 (0.08) *
$DESI_{it}$	1.7367 (0.51) ***	0.8625 (0.59)	–0.808 (0.128) ***	–0.614 (0.125) ***	–0.329 (0.039) ***	–0.235 (0.06) ***
$DESI^2_{it}$	–0.2995 (0.07) ***	–0.1625 (0.08) **	–	–	–	–
ACC_{it}	0.0417 (0.02) **	0.0323 (0.019) *	0.0825 (0.033) **	0.0783 (0.03) **	0.0325 (0.022)	0.0419 (0.022) *
RD_{it}	–	0.1699 (0.069) **	–	0.1948 (0.106) *	–	0.0604 (0.0796)
K_{it}	–	0.0205 (0.05)	–	0.0359 (0.0653)	–	–0.0271 (0.0887)
UR_{it}	–	0.0983 (0.038) **	–	0.1117 (0.058) *	–	0.1273 (0.047) **
$FOSIL_{it}$	–	0.6301 (0.14) ***	–	0.6045 (0.1619) ***	–	0.4051 (0.2426) *
$RENEW_{it}$	–	–0.0871 (0.061)	–	–0.1485 (0.0864) *	–	–0.034 (0.0797)
Constant	–1.231 (0.896)	–2.772 (1.155) **	4.076 (0.478) ***	0.5378 (1.205)	2.3217 (0.138) ***	0.3757 (1.529)
N	27	27	14	14	13	13
T	6	6	6	6	6	6
Estimator	Fixed effect	Fixed effect	Fixed effect	Fixed effect	Fixed effect	Fixed effect

Note. *, **, *** – significance at 10%, 5% and 1%, respectively.

Source: author's calculations.

The interaction term ACC_{it} in the pre-2020 sample is insignificant. However, in the whole sample, the interaction coefficient (γ_1) becomes statistically significant in specifications for the EU-27 and EU-14, while remaining weaker and less robust for the EU-13 (Table 3). The positive coefficient of the interaction term suggests that, as AI innovation capacity increases and economies become more digitally mature, the initially negative effect of AI innovation capacity on energy intensity weakens and may eventually become positive. This might mean that, at higher levels of digital transformation, the potential efficiency gains from AI-related innovation can be partly offset by the energy requirements of digital and computational infrastructure. Therefore, the full interpretation requires examining marginal effects at different levels of DESI.

Among the control variables, fossil fuel energy consumption is the most consistently significant and positively associated with energy intensity, indicating that economies more reliant on fossil fuels tend to be characterized by higher energy intensity. The remaining controls are generally not statistically significant, especially in the pre-2020 sample, and therefore should not be overinterpreted.

The findings indicate that AI innovation capacity has the potential to reduce energy intensity. However, this potential is not automatic and seems to be conditional on a sufficient level of development and effective sectoral integration of digital transformation.

4.3. Marginal effects of AI innovation capacity on energy intensity across levels of digital transformation

Recognizing that the coefficients in Tables 2 and 3 can only shed limited light on the analyzed relationship, we also present a figure (Figure 4A–C) that graphically illustrates how the marginal effect of AI innovation capacity in EU economies changes across the range of digital transformation. The marginal effects were calculated for the EU-27, EU-14 and EU-13 samples. To facilitate their interpretation, the detailed marginal effects calculated at the 25th (P25), median (P50), and 75th (P75) percentiles of DESI are presented in Table A4 in the Appendix. A 95% confidence interval around the marginal effect enables us to determine the conditions under which AI innovation capacity has a statistically significant effect on energy intensity in the EU economies. If the upper and lower bounds of the confidence interval are both above (or below) the zero line, the marginal effect is statistically significant (Brambor et al., 2006).

Generally, at low levels of DESI, the marginal effects of AI innovation capacity are negative, suggesting that AI-related innovation may contribute to reductions in


energy intensity. However, this effect weakens as digital transformation progresses and, once DESI exceeds a certain threshold, becomes positive, suggesting that at higher levels of digitalization, AI innovation capacity may be associated with higher energy intensity (Figure 4, Table A4).

In the EU-27 model with control variables, the marginal effect of AI innovation capacity on EI equals zero at DESI of approximately 40.2, whereas in the model without control variables, the same result is reached at around 47.4 (Figure 4A). Above these values, the marginal effect turns positive, indicating that at higher levels of digital transformation, AI innovation capacity is associated with higher energy intensity. However, for the EU-27, marginal effects remain statistically insignificant across the 25th and 75th percentiles of DESI. This may partly reflect the heterogeneity of the group, with the average effect masking differences between the more digitally advanced economies and those less digitally advanced. Consequently, by dividing EU economies into those with higher (EU-14) and lower (EU-13) levels of digital development, we can obtain more detailed and consistent results and relationships.

For the EU-13, the marginal effects between the 25th and 75th percentiles are generally negative or near zero, and remain insignificant. This may suggest that economies with lower levels of digital development do not experience a conditional effect of AI innovation capacity on energy intensity, while at the same time moving toward a point where the marginal effect becomes positive.

For the EU-14, the marginal effect becomes positive and statistically significant at high levels of DESI (Figure 4, Table A4). In the specification with control variables (model 10), the marginal effect rises from 0.0151 at the 25th percentile of DESI to 0.0325 at the 75th percentile. At that level, a 1% increase in AI innovation capacity is associated with an approximately 0.0325% increase in energy intensity, *ceteris paribus*. This indicates that in economies with more advanced digital infrastructure, greater capacity for artificial intelligence innovation is associated with increased energy consumption when the digital transformation is already substantial.

Figure 4. Marginal effects of AI innovation capacity on energy intensity, conditional on digital transformation



Note. Dots and dotted vertical lines mark the DESI value where the marginal effect equals zero. The grey area indicates the observed DESI range in the estimation sample.

Source: author's work.

The development and diffusion of AI require digital and computational infrastructure, including data centers, cloud services and increasing computing power, all of which generate additional energy demand. As a result, the energy-saving potential of AI innovation capacity may be partly offset by the energy requirements of this infrastructure, especially if AI technologies are not yet sufficiently mature or widely integrated into energy-efficient production and service processes (Masanet et al., 2020). Such an interpretation is also consistent with the literature, which suggests that digitalization does not necessarily reduce energy demand in a direct or automatic way. Digital technologies may foster energy savings, but they also require ICT infrastructure, data processing capacity and computing power, which may increase energy use, especially in more digitally advanced economies (Lange et al., 2020; Masanet et al., 2020). Similar concerns are raised in the AI literature, where the environmental cost of training and deploying AI systems is increasingly emphasized (Dhar, 2020; Schwartz et al., 2020). Therefore, the positive marginal effects observed at higher DESI levels may be interpreted as the potential efficiency gains from AI being partly offset by the energy requirements of digital infrastructure and the possible rebound effects (Brockway et al., 2021; Gillingham et al., 2016).

4.4. Robustness check

The robustness of the results is tested by constructing additional model specifications to assess the sensitivity of the findings to the assumed linear functional form.

The first alternative form of the model includes a squared term for AI innovation capacity to test for nonlinear effects, while the model takes the following form:

$$EI_{it} = \beta_1 AIP_{it} + \beta_2 AIP_{it}^2 + \varphi_1 DESI_{it} + \gamma_1 ACC_{it} + FE_i + \alpha X_{it} + \varepsilon_{it}. \quad (12)$$

The second alternative model specification examines the lagged influence of AIP and DESI on EI, thereby enabling us to assess the robustness of the results to alternative timing assumptions and to partially address the potential endogeneity concerns. The model has the following form:

$$EI_{it} = \beta_1 AIP_{i,t-1} + \varphi_1 DESI_{i,t-1} + \gamma_1 ACC_{i,t-1} + FE_i + \alpha X_{it} + \varepsilon_{it}. \quad (13)$$

The results of the nonlinear models are presented in Table A5, while the models for the lagged specification are shown in Table A6 in the Appendix. The nonlinear and lagged specifications yield results that are generally consistent with the baseline models (7)–(12). Across both alternative specifications, the coefficient on AI innovation capacity remains negative in all the models, while significant effects are

confirmed for the EU-14 and EU-13 (specifications with control variables) only in the nonlinear form. At the same time, the squared term of AI innovation capacity is statistically insignificant across all nonlinear specifications, providing insufficient empirical support for a nonlinear relationship over the analyzed period. Therefore, the baseline linear specification appears to provide a parsimonious approximation of the relationship. However, this conclusion should be interpreted with caution, as the sample's relatively short time dimension may limit the statistical power of nonlinear tests. The lagged specifications further support the robustness of the main findings, although the estimated delayed effects are weaker, suggesting that the results are not driven solely by the contemporaneous model structure.

The marginal effects presented in Table A7 further support the robustness of the baseline findings. For the EU-27 and EU-13 economies, the effects of AI innovation capacity remain close to zero and statistically insignificant across the selected levels of digital transformation. In contrast, the strongest conditional pattern is observed for the EU-14 economies, where the marginal effects are positive and increase along with the DESI levels. So in the controlled nonlinear specification, the effect becomes statistically significant at the 75th percentile of DESI. The lagged specifications yield weaker and insignificant marginal effects. Nevertheless, their direction remains broadly consistent with the baseline and nonlinear models, indicating that the main conclusions are not driven solely by the contemporaneous specification. Overall, the marginal effects confirm the stability of the results across alternative specifications and show that the AI-energy intensity relationship is conditional on the level of digital transformation and differs across groups of economies.

5. Discussion and Conclusions

This study examined the relationship between artificial intelligence innovation capacity and energy intensity in the European Union, with particular attention devoted to the moderating role of digital transformation. The main objective was to determine whether AI-related innovation contributes to reducing energy intensity in EU economies and whether this relationship depends on the level and pace of digitalization. The empirical design made achieving this objective possible by estimating panel models for the EU-27 and for two groups of countries differing in digital maturity: the EU-14 and EU-13.

The results suggest that the relationship between AI innovation capacity and energy intensity is conditional on the level of digital transformation and differs across EU economies. Although the coefficients on AI innovation capacity generally indicate a negative association with energy intensity, the marginal effects reveal a more nuanced picture. In the heterogeneous EU-27 sample, the marginal effects are statistically

insignificant, suggesting that average estimates may mask differences between economies at different stages of digital transformation. In the EU-14, the marginal effect becomes stronger as DESI increases and is statistically significant at higher levels of digitalization; however, its positive sign indicates that AI-related innovation is associated with higher, rather than lower, energy intensity. This might suggest that in digitally advanced economies, the energy demand of digital and computing infrastructure may partly offset potential efficiency gains (Olson, 2024; Solnørdal & Thyholdt, 2019). In the EU-13, the marginal effects remain statistically insignificant, possibly reflecting lower digital maturity and a weaker capacity to translate AI-related innovation into measurable energy-saving outcomes (Mo & Liu, 2026). The results thus provided partial support for the first hypothesis, namely: “AI innovation capacity contributes to reducing energy intensity across EU economies”, and upheld the second hypothesis, which says: “The relationship between AI innovation capacity and energy intensity is moderated by the level of digital development, such that the marginal effect of digitalization (DESI) on energy intensity changes with its level”.

The moderating role of digitalization became more evident after the acceleration of digital transformation associated with the COVID-19 pandemic, but it remains visible only in economies with a higher level of digital maturity (EU-14). The expansion of digital tools, remote work, digital services and enterprise-level technology adoption may have strengthened the relationship between AI-related innovation and energy intensity. However, this effect should not be interpreted as an automatic improvement in energy efficiency, since the diffusion of AI and digital technologies also requires energy-intensive digital and computational infrastructure. That means the third hypothesis, “The pandemic-induced structural shift amplifies the moderating effect of digital development on the relationship between AI innovation capacity and energy intensity, due to accelerated digital transformation”, is partially supported.

This study highlights the importance of energy intensity as a key indicator of energy transformation. Energy intensity reflects not only the amount of energy required to generate the economic output, but also the technological, structural, and organizational capacity of an economy to decouple growth from energy consumption (Croner & Frankovic, 2018; Deichmann et al., 2018; IEA, 2021). Therefore, reducing energy intensity should be considered a central dimension of the green and digital transition.

Recognizing the strategic importance of reducing energy intensity, the EU and its member states have increasingly been embedding energy-efficiency objectives in climate, industrial and innovation policies. The revised Energy Efficiency Directive sets a binding target to reduce final energy consumption by 11.7% by 2030, while national initiatives in countries such as Sweden, Germany, France and Spain link energy-efficiency improvements with industrial modernization, digital energy

management and decarbonization (Directive (EU) 2023/1791 of the European Parliament and of the Council of 13 September 2023 on energy efficiency and amending Regulation (EU) 2023/955 – recast). These developments show that reducing energy intensity is increasingly often treated not only as an energy policy goal, but also as a strategic instrument for strengthening competitiveness, industrial resilience, and environmental sustainability in the low-carbon transition.

Support for AI research, patenting, data infrastructure and industrial AI applications may improve energy performance, especially in energy-intensive sectors. Digital infrastructure is a necessary condition for transforming AI innovation into energy-efficiency gains. Investment in broadband networks, cloud infrastructure, data interoperability, cybersecurity, digital skills and smart energy systems is therefore essential (Colangelo et al., 2025; Olson, 2024). Less digitally advanced EU economies require targeted support to reduce the digital-energy efficiency gap. In particular, policies should support SMEs in adopting AI and digital energy management tools and strengthen the diffusion of technology between the more and the less digitally advanced member states. Policies supporting AI and digital transformation should be aligned with renewable energy deployment, energy-efficient digital infrastructure and rebound-effect monitoring. This highlights the twin transition as an integrated agenda linking AI innovation, digital maturity and energy transformation.

This study has several limitations. First, the empirical analysis is based on a relatively short panel, mainly due to the relatively poor availability of comparable DESI and AI patent data for EU countries. Therefore, the results should be interpreted primarily as short-term conditional associations rather than long-term causal effects. Second, AI innovation capacity is measured using AI-related patents which capture innovation activity but not the actual implementation of AI technologies across sectors. Similarly, the aggregate DESI index reflects the overall digital maturity but does not identify which specific dimensions of digital transformation (human capital, connectivity, integration of digital technology and digital public services) are most relevant for the AI-energy intensity relationship. Third, the potential endogeneity cannot be fully ruled out. Although country fixed effects control for unobserved time-invariant differences between economies, they do not eliminate reverse causality or omitted time-varying factors. For instance, economies with different levels of energy intensity may also differ in their incentives and capacity to invest in AI-related innovation. Therefore, the results should be interpreted as conditional associations rather than definitive causal effects. Finally, there were some common shocks, such as the COVID-19 pandemic, the energy crisis and accelerated digital transformation in the analyzed period, which might have affected the interpretation of the results. These issues suggest that future studies should be expanded to include longer time series,

alternative indicators of digitalization and different approaches, such as dynamic, nonlinear or threshold-based panel models.

References

- Abe, M., & Nakayama, H. (2018). Deep learning for forecasting stock returns in the cross-section. In D. Phung, V. S. Tseng, G. I. Webb, B. Ho, M. Ganji & L. Rashidi (Eds.), *Advances in Knowledge Discovery and Data Mining* (pp. 273–284). Springer. https://doi.org/10.1007/978-3-319-93034-3_22.
- Acemoglu, D., Akcigit, U., Alp, H., Bloom, N., & Kerr, W. (2018). Innovation, Reallocation, and Growth. *The American Economic Review*, 108(11), 3450–3491. <https://doi.org/10.1257/AER.20130470>.
- Acemoglu, D., & Restrepo, P. (2022). Tasks, Automation, and the Rise in U.S. Wage Inequality. *Econometrica*, 90(5), 1973–2016. <https://doi.org/10.3982/ECTA19815>.
- Aghion, P., Akcigit, U., & Howitt, P. (2013). *What Do We Learn From Schumpeterian Growth Theory?* (NBER Working Paper No. 18824). <http://doi.org/10.3386/w18824>.
- Aghion, P., Akcigit, U., & Howitt, P. (2015a). Lessons from Schumpeterian Growth Theory. *The American Economic Review*, 105(5), 94–99. <https://doi.org/10.1257/AER.P20151067>.
- Aghion, P., Akcigit, U., & Howitt, P. (2015b). The Schumpeterian Growth Paradigm. *Annual Review of Economics*, 7(1), 557–575. <https://doi.org/10.1146/annurev-economics-080614-115412>.
- Aghion, P., & Jaravel, X. (2015). Knowledge spillovers, innovation and growth. *The Economic Journal*, 125(583), 533–573. <https://doi.org/10.1111/ECOJ.12199>.
- Altın, H. (2024). The impact of energy efficiency and renewable energy consumption on carbon emissions in G7 countries. *International Journal of Sustainable Engineering*, 17(1), 1–9. <https://doi.org/10.1080/19397038.2024.2319648>.
- Amadei, C., Dosi, C., & Pintus, F. J. (2024). *Energy Intensity and Structural Changes: Does Offshoring Matter?* (Working Paper No. 26). <https://www.econstor.eu/bitstream/10419/307436/1/1910688029.pdf>.
- Biggi, G., Iori, M., Mazzei, J., & Mina, A. (2025). Green intelligence: the AI content of green technologies. *Eurasian Business Review*, 15, 803–840. <https://doi.org/10.1007/s40821-024-00288-1>.
- Bölük, G., & Mert, M. (2014). Fossil & renewable energy consumption, GHGs (greenhouse gases) and economic growth: Evidence from a panel of EU (European Union) countries. *Energy*, 74, 439–446. <https://doi.org/10.1016/J.ENERGY.2014.07.008>.
- Brambor, T., Clark, W. R., & Golder, M. (2006). Understanding Interaction Models: Improving Empirical Analyses. *Political Analysis*, 14(1), 63–82. <https://doi.org/10.1093/PAN/MPI014>.
- Brockway, P. E., Sorrell, S., Semieniuk, G., Heun, M. K., & Court, V. (2021). Energy efficiency and economy-wide rebound effects: A review of the evidence and its implications. *Renewable and Sustainable Energy Reviews*, 141, 1–20. <https://doi.org/10.1016/J.RSER.2021.110781>.
- Cao, L. (2022). AI in Finance: Challenges, Techniques, and Opportunities. *ACM Computing Surveys*, 55(3), 1–38. <https://doi.org/10.1145/3502289>.

- Chun, Y., & Hwang, J. (2025). The Nexus of Artificial Intelligence and Green Innovation: a Cross-Density Analysis at the Country Level. *Journal of the Knowledge Economy*, 16(1), 1688–1716. <https://doi.org/10.1007/s13132-024-02076-8>.
- Colangelo, P., Coskun, A. K., Megrue, J., Roberts, C., Sengupta, S., Sivaram, V., Tiao, E., Vijaykar, A., Williams, C., Wilson, D. C., Records, B., MacFarland, Z., Dreiling, D., Morey, N., Ratnayake, A., & Vairamohan, B. (2025). AI data centres as grid-interactive assets. *Nature Energy*, 11(2), 254–261. <https://doi.org/10.1038/s41560-025-01927-1>.
- Croner, D., & Frankovic, I. (2018). A Structural Decomposition Analysis of Global and National Energy Intensity Trends. *The Energy Journal*, 39(2), 103–122. <https://doi.org/10.5547/01956574.39.2.dcro>.
- Deichmann, U., Reuter, A., Vollmer, S., & Zhang, F. (2018). *Relationship between Energy Intensity and Economic Growth. New Evidence from a Multi-Country Multi-Sector Data Set* (Policy Research Working Paper No. 8322). <https://papers.ssrn.com/sol3/Delivery.cfm/8322.pdf?abstractid=3113873&mirid=1>.
- Dhar, P. (2020). The carbon impact of artificial intelligence. *Nature Machine Intelligence*, 2(8), 423–425. <https://doi.org/10.1038/S42256-020-0219-9>.
- Directive (EU) 2023/1791 of the European Parliament and of the Council of 13 September 2023 on energy efficiency and amending Regulation (EU) 2023/955 (recast) (OJ EU L 231, 20.9.2023).
- Dzwigol, H., Kwilinski, A., Lyulyov, O., & Pimonenko, T. (2024). Digitalization and Energy in Attaining Sustainable Development: Impact on Energy Consumption, Energy Structure, and Energy Intensity. *Energies*, 17(5), 1–17. <https://doi.org/10.3390/en17051213>.
- European Commission. (n.d.). *The Digital Economy and Society Index (DESI) 2022, Methodological Note*. Retrieved 19 July 2026, from <https://digital-strategy.ec.europa.eu/en/policies/desi>.
- European Commission. (2022, April 8). *Energy and the Green Deal*. https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/european-green-deal/energy-and-green-deal_en?utm_source=chatgpt.com.
- Gao, L., & Wang, J. (2025). Can artificial intelligence reduce energy vulnerability? Evidence from an international perspective. *Energy Economics*, 145, 108491. <https://doi.org/10.1016/J.ENERCO.2025.108491>.
- Ghahramani, M., Qiao, Y., Zhou, M. C., Hagan, A. O., & Sweeney, J. (2020). AI-based modeling and data-driven evaluation for smart manufacturing processes. *IEEE/CAA Journal of Automatica Sinica*, 7(4), 1026–1037. <https://doi.org/10.1109/JAS.2020.1003114>.
- Gillingham, K., Rapson, D., & Wagner, G. (2016). The Rebound Effect and Energy Efficiency Policy. *Review of Environmental Economics and Policy*, 10(1), 68–88. <https://doi.org/10.1093/REEP/REV017>.
- Graetz, G., & Michaels, G. (2018). Robots at Work. *The Review of Economics and Statistics*, 100(5), 753–768. https://doi.org/10.1162/rest_a_00754.
- Gyamfi, B. A., Kwakwa, P. A., & Adebayo, T. S. (2023). Energy intensity among European Union countries: the role of renewable energy, income and trade. *International Journal of Energy Sector Management*, 17(4), 801–819. <https://doi.org/10.1108/IJESM-05-2022-0018>.
- Huang, G., He, L. Y., & Lin, X. (2022). Robot adoption and energy performance: Evidence from Chinese industrial firms. *Energy Economics*, 107, 105837. <https://doi.org/10.1016/J.ENERCO.2022.105837>.

- Intergovernmental Panel of Climate Change. (2023). Summary for Policymakers. In H. Lee & J. Romero (Eds.), *Climate Change 2023. Synthesis Report* (pp. 1–34). <https://doi.org/10.59327/IPCC/AR6-9789291691647.001>.
- International Energy Agency. (2021). *Energy intensity. Recent progress*. <https://www.iea.org/reports/sdg7-data-and-projections/energy-intensity>.
- International Energy Agency. (2025). *Energy and AI*. <https://www.iea.org/reports/energy-and-ai>.
- Jie, W., & Rabnawaz, K. (2024). Renewable energy and CO₂ emissions in developing and developed nations: a panel estimate approach. *Frontiers in Environmental Science*, 12, 1–19. <https://doi.org/10.3389/FENV.S.2024.1405001/XML>.
- Khan, N., Khan, A., Khan, D., Riaz, F., Farooq, M., Hadi, N. U., & Khan, M. I. (2025). The impact of digitalization and trade openness on energy efficiency: A comparative analysis across low, middle, and high-income countries. *Energy Conversion and Management: X*, 27, 1–17. <https://doi.org/10.1016/j.ecmx.2025.101089>.
- Kopka, A., & Grashof, N. (2022). Artificial intelligence: Catalyst or barrier on the path to sustainability?. *Technological Forecasting & Social Change*, 175, 121318. <https://doi.org/10.1016/j.techfore.2021.121318>.
- Kwilinski, A., Lyulyov, O., & Pimonenko, T. (2023). The Impact of Digital Business on Energy Efficiency in EU Countries. *Information*, 14(9), 1–20. <https://doi.org/10.3390/info14090480>.
- Lange, S., Pohl, J., & Santarius, T. (2020). Digitalization and energy consumption. Does ICT reduce energy demand?. *Ecological Economics*, 176, 106760. <https://doi.org/10.1016/J.ECOLECON.2020.106760>.
- Le, T. H. (2022). Connectedness between nonrenewable and renewable energy consumption, economic growth and CO₂ emission in Vietnam: New evidence from a wavelet analysis. *Renewable Energy*, 195, 442–454. <https://doi.org/10.1016/J.RENENE.2022.05.083>.
- Lee, C.-C., Fang, Y., Quan, S., & Li, X. (2024). Leveraging the power of artificial intelligence toward the energy transition: The key role of the digital economy. *Energy Economics*, 135, 107654. <https://doi.org/10.1016/J.ENERCO.2024.107654>.
- Lei, W., Xie, Y., Hafeez, M., & Ullah, S. (2022). Assessing the dynamic linkage between energy efficiency, renewable energy consumption, and CO₂ emissions in China. *Environmental Science and Pollution Research*, 29(13), 19540–19552. <https://doi.org/10.1007/S11356-021-17145-7>.
- Levinson, A. (2021). Energy intensity: Deindustrialization, composition, prices, and policies in U.S. states. *Resource and Energy Economics*, 65, 101243. <https://doi.org/10.1016/J.RESENEECO.2021.101243>.
- Li, Y., Zhang, Y., Pan, A., Han, M., & Veglianti, E. (2022). Carbon emission reduction effects of industrial robot applications: Heterogeneity characteristics and influencing mechanisms. *Technology in Society*, 70, 102034. <https://doi.org/10.1016/J.TECHSOC.2022.102034>.
- Liu, J., Liu, L., Qian, Y., & Song, S. (2022). The effect of artificial intelligence on carbon intensity: Evidence from China's industrial sector. *Socio-Economic Planning Sciences*, 83, 101002. <https://doi.org/10.1016/J.SEPS.2020.101002>.
- Liu, L., Yang, K., Fujii, H., & Liu, J. (2021). Artificial intelligence and energy intensity in China's industrial sector: Effect and transmission channel. *Economic Analysis and Policy*, 70, 276–293. <https://doi.org/10.1016/j.eap.2021.03.002>.

- Liu, X., Wang, P., Song, H., & Zeng, X. (2021). Determinants of net primary productivity: Low-carbon development from the perspective of carbon sequestration. *Technological Forecasting and Social Change*, 172, 121006. <https://doi.org/10.1016/J.TECHFORE.2021.121006>.
- Luan, F., Yang, X., Chen, Y., & Regis, P. J. (2022). Industrial robots and air environment: A moderated mediation model of population density and energy consumption. *Sustainable Production and Consumption*, 30, 870–888. <https://doi.org/10.1016/J.SPC.2022.01.015>.
- Lyu, J., Gao, Z., Li, Y., & Zhang, Q. (2025). The two-way street: How AI and clean energy affect each other. *Energy Economics*, 147, 108566. <https://doi.org/10.1016/J.ENERCO.2025.108566>.
- McCarthy, J., Minsky, M. L., Rochester, N. & Shannon, C. E. (2006). A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence, August 31, 1955. *AI Magazine*, 27(4), 12–14. <https://doi.org/10.1609/aimag.v27i4.1904>.
- Masanet, E., Shehabi, A., Lei, N., Smith, S., & Koomey, J. (2020). Recalibrating global data center energy-use estimates. *Science*, 367(6481), 984–986. <https://doi.org/10.1126/SCIENCE.ABA3758>.
- Melnychenko, O. (2019). Application of artificial intelligence in control systems of economic activity. *Virtual Economics*, 2(3), 30–40. [https://doi.org/10.34021/VE.2019.02.03\(3\)](https://doi.org/10.34021/VE.2019.02.03(3)).
- Mo, S., & Liu, X. (2026). How does artificial intelligence affect energy intensity? The role of digital economy. *International Journal of Green Energy*, 23(1), 265–281. <https://doi.org/10.1080/15435075.2025.2552491>.
- Morand, C., Ligozat, A.-L., & Névéol, A. (2024). *How Green Can AI Be? A Study of Trends in Machine Learning Environmental Impacts*. <https://doi.org/10.48550/arXiv.2412.17376>.
- Mulder, P., & de Groot, H. L. F. (2012). Structural change and convergence of energy intensity across OECD countries, 1970–2005. *Energy Economics*, 34(6), 1910–1921. <https://doi.org/10.1016/J.ENERCO.2012.07.023>.
- Nishant, R., Kennedy, M., & Corbett, J. (2020). Artificial intelligence for sustainability: Challenges, opportunities, and a research agenda. *International Journal of Information Management*, 53, 102104. <https://doi.org/10.1016/j.ijinfomgt.2020.102104>.
- Organisation for Economic Co-operation and Development. (2019, June 11). *Artificial Intelligence in Society*. <https://doi.org/10.1787/eedfee77-en>.
- Organisation for Economic Co-operation and Development. (2025). *Identifying emerging AI technologies using patent data: A semi-automated approach*. OECD Publishing. <https://doi.org/10.1787/D17E9A1A-EN>.
- Olson, E. (2024). Digital Transformation and AI in Energy Systems: Applications, Challenges, and the Path Forward. In T. Lynn, P. Rosati, D. Kreps & K. Conboy (Eds.), *Digital Sustainability* (pp. 63–79). Palgrave Macmillan. https://doi.org/10.1007/978-3-031-61749-2_4.
- Podrecca, M., Culot, G., Tavassoli, S., & Orzes, G. (2024). Artificial Intelligence for Climate Change: A Patent Analysis in the Manufacturing Sector. *IEEE Transactions on Engineering Management*, 71, 15005–15024. <https://doi.org/10.1109/TEM.2024.3469370>.
- Rame, R., Purwanto, P., & Sudarno, S. (2024). Industry 5.0 and sustainability: An overview of emerging trends and challenges for a green future. *Innovation and Green Development*, 3(4), 1–20. <https://doi.org/10.1016/J.IGD.2024.100173>.
- Schwartz, R., Dodge, J., Smith, N. A., & Etzioni, O. (2020). Green AI. *Communications of the ACM*, 63(12), 54–63. <https://doi.org/10.1145/3381831>.

- Şerban, A. C., & Lytras, M. D. (2020). Artificial Intelligence for Smart Renewable Energy Sector in Europe–Smart Energy Infrastructures for Next Generation Smart Cities. *IEEE Access*, 8, 77364–77377. <https://doi.org/10.1109/ACCESS.2020.2990123>.
- Shayanmehr, S., Henneberry, S. R., Sabouni, M. S., & Foroushani, N. S. (2020). Climate Change and Sustainability of Crop Yield in Dry Regions Food Insecurity. *Sustainability*, 12(23), 1–24. <https://doi.org/10.3390/SU12239890>.
- Solnørdal, M. T., & Thyholdt, S. B. (2019). Absorptive capacity and energy efficiency in manufacturing firms – An empirical analysis in Norway. *Energy Policy*, 132, 987–990. <https://doi.org/10.1016/j.enpol.2019.06.069>.
- Stern, D. I. (2011). The role of energy in economic growth. *Annals of the New York Academy of Sciences*, 1219(1), 26–51. <https://doi.org/10.1111/J.1749-6632.2010.05921.X>.
- Ucar, A., Karakose, M., & Kırımça, N. (2024). Artificial Intelligence for Predictive Maintenance Applications: Key Components, Trustworthiness, and Future Trends. *Applied Sciences*, 14(2), 1–40. <https://doi.org/10.3390/APP14020898>.
- Williams, R. (2012). Using the Margins Command to Estimate and Interpret Adjusted Predictions and Marginal Effects. *The Stata Journal*, 12(2), 308–331. <https://doi.org/10.1177/1536867X1201200209>.
- Xie, T., & Wang, X. (2024). Investigating the nonlinear carbon reduction effect of AI: empirical insights from China’s provincial level. *Frontiers in Environmental Science*, 12, 1–16. <https://doi.org/10.3389/FENVS.2024.1353294>.
- Yadav, D. K., Kaushik, A., & Yadav, N. (2024). Predicting machine failures using machine learning and deep learning algorithms. *Sustainable Manufacturing and Service Economics*, 3, 1–11. <https://doi.org/10.1016/J.SMSE.2024.100029>.
- Yuan, C., Liu, S., & Wu, J. (2009). Research on energy-saving effect of technological progress based on Cobb–Douglas production function. *Energy Policy*, 37(8), 2842–2846. <https://doi.org/10.1016/J.ENPOL.2009.04.025>.
- Zakari, A., Khan, I., Tawiah, V., & Alvarado, R. (2022). Reviewing the ecological footprints of Africa top carbon consumer: a quantile on quantile analysis. *International Journal of Environmental Science and Technology*, 19(11), 11475–11486. <https://doi.org/10.1007/s13762-021-03904-z>.
- Zhai, X., Chu, X., Chai, C. S., Jong, M. S. Y., Istenic, A., Spector, M., Liu, J. B., Yuan, J., & Li, Y. (2021). A Review of Artificial Intelligence (AI) in Education from 2010 to 2020. *Complexity*, (1), 1–18. <https://doi.org/10.1155/2021/8812542>.
- Zhang, X., Liu, P., Zhu, H., & Hackl, M. (2022). The Impact of Industrial Intelligence on Energy Intensity: Evidence from China. *Sustainability*, 14(12), 1–16. <https://doi.org/10.3390/SU14127219>.

Appendix A

Table A1. Literature covering the potential impact of AI on energy intensity, energy efficiency, CO₂ emissions and energy consumption

Author(s)	Country/sector/ research sample	Key findings	Energy changes through			
			EI	EE	EC	CI
(L. Liu et al., 2021)	China/ industrial sectors	Reduction of energy intensity (AI ↑ by 1% then EI ↓ by 0.0244%). Boosting applications of AI can significantly reduce energy intensity by both increasing production value and reducing energy consumption.	X		X	
(Zhang et al., 2022)	China, time series 2008- 2019, OLS and WLS	Increasing industrial AI reduces energy intensity by lowering energy consumption. There is a significant scale effect at the enterprise level: for large enterprises, the AI effect on energy intensity is greater, with larger reductions.	X		X	
(Huang et al., 2022)	Chinese firms' level analysis – panel from 2001 to 2012	Robot adoption in the production process promotes intelligent production and improves energy efficiency, mainly by advancing production technology and increasing output.		X		
(Kopka & Grashof, 2022)	Germany, analysis on the regional level (NUTS-3), panel data for the years 2005-2015	AI's impact on energy use varies by a region's technology and industry. "Dirty industries" consume more energy with AI. Regions with green technology see no significant AI effect. Central regions experience a moderate, decreasing AI impact.			X	
(J. Liu et al., 2022)	Chinese industrial sector (13 categories) - time series 2005-2016	AI adoption significantly reduces carbon intensity, with greater effects in labor- intensive and technology-intensive industries than in capital-intensive sectors. To amplify AI's impact on reducing carbon intensity, diversified policies should be implemented in line with industry characteristics.				X
(Li et al., 2022)	Panel of 35 countries – time span from 1993-2017	The use of industrial robots is actively reducing carbon intensity, and this effect is further enhanced by green total factor productivity and energy efficiency.	X			X
(Gyamfi et al., 2023)	Panel of 26 EU economies for the years 1990-2019	AI contributes to energy efficiency by optimizing the energy use, while promoting the transition towards renewable sources, thereby reducing energy intensity and consumption.	X	X	X	
(Xie & Wang, 2024)	China, panel model for 30 Chinese provinces between 1997 and 2019	The findings indicate that AI technologies have the capacity to substantially reduce carbon emissions, with this effect being particularly evident in regions characterized by advanced digital infrastructure and substantial level of economic development.				X

Table A1. Literature covering the potential AI’s impact on energy intensity, energy efficiency, CO₂ emissions and energy consumption (cont.)

Author(s)	Country/sector/ research sample	Key findings	Energy changes through			
			EI	EE	EC	CI
(Lee et al., 2024)	Panel of 64 countries – time span 2005-2019	Along with the progressing AI development, the scale of energy transformation is gradually growing at the national level. The direct impact of digitalization on the energy transition is negative, but, as a moderating factor, it accelerates the positive influence of AI on it.				X
(Podrecca et al., 2024)	Global dataset of 5,919 granted patents filed between 2011 and 2023	The key findings indicate that while AI patents in manufacturing predominantly support operational efficiency, only a limited share of them directly targets energy management and carbon reduction. This suggests that AI's environmental potential is emerging but has not been fully utilized yet.			X	X
(Gao & Wang, 2025)	Panel of 52 countries – time span 2000-2019	AI mitigates energy vulnerability through accelerating changes in the structure of the industry to improve energy efficiency and reduce energy intensity, catalyzing green innovation technology and optimizing governance capability	X	X		
(Lyu et al., 2025)	China – time series analysis (01.2013-06.2023)	Bi-directional Granger causality relation between AI and CE (clean energy index). AI offers intelligent solutions that enhance efficiency, optimize energy management and foster innovation within the CE sector.		X		

Source: author’s work based on a literature review.

Table A2. Proxy variables describing AI

Type	Autor(s)	Operationalization
Industrial robots / automation-based proxies	(Graetz & Michaels, 2018)	AI-related technological adoption, using country-level robot density data from the International Federation of Robotics (IFR), with robot diffusion serving as a proxy for the influence of intelligent automation on productivity growth.
	(L. Liu et al., 2021)	The authors used AI variables, measured as: (1) the number of the installed industrial robots; (2) the stock of industrial robots.
	(Acemoglu & Restrepo, 2022)	Proxy AI-driven technological change using the stock of industrial robots per thousand workers, treating robots as physical embodiments of task-based automation and intelligent execution systems at a workplace.
	(Zhang et al., 2022)	AI is measured by the stock of industrial robots.
	(J. Liu et al., 2022)	AI is measured by the number of industrial robot installations in China's industrial sectors.
	(Li et al., 2022)	Industrial robot application intensity is a proxy variable of intelligent production.
	(Lee et al., 2024)	AI is measured as the stock of industrial robots.
	(Gao & Wang, 2025)	AI is measured on the basis of: (1) the stock of the robot usage, and (2) the installation of new robots in companies.

Table A2. Proxy variables describing AI (cont.)

Type	Autor(s)	Operationalization
Firm-level AI/robot adoption	(Huang et al., 2022)	To measure the degree of a firm's robot implementation, the researchers used three indicators: (1) a dummy variable to describe whether the firm uses robots in production, (2) the number of robots used by the firm in the production process, and (3) the value of robots used by the firm.
AI indices	(Lyu et al., 2025)	China's AI industry is characterized by the Artificial Intelligence Concept Index, which captures key dimensions of AI development – technology, policy and market implementation.
AI patent-based proxies	(Kopka & Grashof, 2022)	AI knowledge of regions is operationalized using a patent dataset, in which authors calculated the AI share as the cumulative number of AI patents at the regional level, divided by the total number of patent families originating in that region.
	(Xie & Wang, 2024)	AI is measured as (1) the quantity of AI patents, (2) the operational stock of industrial robots, and (3) the increase in the number of industrial robots.
AI-green patent-based proxies	(Podrecca et al., 2024)	The authors use AI-enhanced patent activity in the manufacturing sector, focusing on inventions that integrate AI techniques into climate-oriented process optimization and energy efficiency.
	(Biggi et al., 2025)	AI is proxied by identifying patents that combine artificial intelligence components – such as machine learning and neural networks – with green technologies, using CPC codes and keyword mapping.
	(Chun & Hwang, 2025)	AI is proxied by the density and interaction of national-level AI and green technology patent portfolios, capturing the extent to which AI-related patents intersect with green innovation capacity in each country

Note. The table summarizes selected proxy variables used in the literature to measure AI development, AI adoption and AI innovation capacity. The proxies can be grouped into robot-based measures, firm-level adoption indicators, composite AI indices and patent-based measures.

Source: author's work based on a literature review.

Table A3. Two clusters of EU economies based on the median of the DESI index

Cluster 0 – EU13	Cluster 1 – EU14
Belgium, Bulgaria, Czechia, Cyprus, Greece, Croatia, Hungary, Italy, Latvia, Poland, Portugal, Romania, Slovakia	Austria, Denmark, Estonia, Finland, France, Germany, Ireland, Lithuania, Luxembourg, Malta, Netherlands, Slovenia, Spain, Sweden

Source: author's work.

Table A4. Marginal effects of AI innovation capacity on energy intensity, conditional on digital transformation, based on models (7)–(12)

Model	G	C	P	I_desi	me	se	95% ci
7	EU-27	No	P25	3.6053	-0.0105	0.009	[-0.029; 0.007]
			P50	3.7821	-0.0032	0.009	[-0.020; 0.014]
			P75	3.9259	0.0028	0.009	[-0.015; 0.021]
8	EU-27	Yes	P25	3.6053	-0.0028	0.009	[-0.021; 0.015]
			P50	3.7821	0.0029	0.008	[-0.014; 0.019]
			P75	3.9259	0.0075	0.009	[-0.010; 0.025]
9	EU-14	No	P25	3.7792	0.0063	0.014	[-0.020; 0.033]
			P50	3.886	0.0151	0.014	[-0.013; 0.043]
			P75	4.0021	0.0247	0.016	[-0.006; 0.055]
10	EU-14	Yes	P25	3.7792	0.0151	0.013	[-0.010; 0.039]
			P50	3.886	0.0234	0.013	[-0.002; 0.049]
			P75	4.0021	0.0325	0.014	[0.005; 0.061]
11	EU-13	No	P25	3.4086	-0.0155	0.010	[-0.036; 0.005]
			P50	3.6008	-0.0093	0.009	[-0.027; 0.009]
			P75	3.7081	-0.0058	0.010	[-0.024; 0.013]
12	EU-13	Yes	P25	3.4086	-0.0086	0.010	[-0.028; 0.011]
			P50	3.6008	-0.0006	0.009	[-0.018; 0.017]
			P75	3.7081	0.0039	0.009	[-0.014; 0.022]

Note. Bolded results indicate statistically significant effects, G stands for a group of economies, C for control variables, P for percentiles, and me/se represents marginal effect/standard error.

Source: author's calculations.

Table A5. Regression results of energy intensity models (NL) for the EU-27, EU-14 and the EU-13 for the whole sample, including squared AI innovation capacity (robustness check)

Variable	NL1	NL2	NL3	NL4	NL5	NL6
	EU-27	EU-27	EU-14	EU-14	EU-13	EU-13
$AI P_{i,t}$	-0.1406 (0.1105)	-0.0757 (0.3474)	-0.291** (0.0410)	-0.2515* (0.0560)	-0.1201 (0.1756)	-0.1574* (0.0600)
$AI P^2_{i,t}$	0.0016 (0.6082)	0.0038 (0.1890)	0.0012 (0.7949)	0.0025 (0.5480)	0.0012 (0.8347)	-0.0011 (0.8391)
$DESI_{i,t}$	1.621** (0.014)	0.575 (0.363)	-0.789*** (0.000)	-0.574*** (0.0002)	-0.330*** (0.0000)	-0.233*** (0.0018)
$DESI^2_{i,t}$	-0.282*** (0.0026)	-0.1198 (0.1786)	—	—	—	—
$ACC_{i,t}$	0.0348 (0.1594)	0.0172 (0.4436)	0.0770* (0.0570)	0.0672* (0.0697)	0.0303 (0.2247)	0.0441* (0.0656)
RD_{it}	—	0.181***	—	0.1861*	—	0.057
K_{it}	—	0.024	—	0.0431	—	-0.031
UR_{it}	—	0.099**	—	0.116*	—	0.128**
$FOSIL_{it}$	—	0.641***	—	0.612***	—	0.4142*
$RENEW_{it}$	—	-0.0978	—	-0.149*	—	-0.03
(R ²)	0.9814	0.9856	0.9826	0.9871	0.9793	0.9826

Note. *, ** *** – significance at 10%, 5% and 1%, respectively; p-values in parentheses.

Source: author's calculations.

Table A6. Regression results of energy intensity models with lagged variables (LM) for the EU-27, EU-14 and the EU-13 for the whole sample (robustness check)

Variable	LM-1	LM-2	LM-3	LM-4	LM-5	LM-6
	EU-27	EU-27	EU-14	EU-14	EU-13	EU-13
$AIP_{i,t-1}$	-0.0931 (0.1074)	-0.0517 (0.1002)	-0.1781 (0.1742)	-0.1936 (0.1579)	-0.1111 (0.1187)	-0.1443 (0.1162)
$DESI_{i,t-1}$	1.734** (0.8537)	0.244 (0.8164)	-0.646*** (0.1833)	-0.573*** (0.1723)	-0.363*** (0.0610)	-0.273*** (0.0899)
$DESI^2_{i,t-1}$	-0.301** (0.1217)	-0.0812 (0.1168)	—	—	—	—
$ACC_{i,t}$	0.0222 (0.0287)	0.0131 (0.0266)	0.0416 (0.0459)	0.0517 (0.0414)	0.0324 (0.0328)	0.0427 (0.0321)
RD_{it}	—	0.2784***	—	0.3636**	—	0.1280
K_{it}	—	0.0736	—	0.0748	—	0.0020
UR_{it}	—	0.1242**	—	0.0988	—	0.1546**
$FOSIL_{it}$	—	0.6943***	—	0.6659***	—	0.6215*
$RENEW_{it}$	—	-0.0980	—	-0.1399	—	-0.0271
R^2	0.9825	0.9866	0.9832	0.9876	0.9781	0.9817

Note. *, **, *** – significance at 10%, 5% and 1%, respectively; standard errors in parentheses.

Source: author's calculations.

Table A7. Marginal effects of AI innovation capacity on energy intensity, conditional on digital transformation, based on non-linear models and models with lagged variables

Type			Non-linear models (NL)			Lagged models (LM)		
G	C	P	I_desi	me/se	95% ci	I_desi	me/se	95% ci
EU-27	No	25	3.605	-0.0078 (0.011)	[-0.029; 0.013]	3.579	-0.0135 (0.011)	[-0.036; 0.009]
		50	3.782	-0.0016 (0.009)	[-0.020; 0.017]	3.741	-0.01 (0.011)	[-0.031; 0.011]
		75	3.926	0.0034 (0.009)	[-0.015; 0.022]	3.871	-0.0071 (0.012)	[-0.030; 0.015]
EU-27	Yes	25	3.605	0.0033 (0.01)	[-0.017; 0.023]	3.579	-0.005 (0.011)	[-0.027; 0.017]
		50	3.782	0.0063 (0.009)	[-0.011; 0.023]	3.741	-0.0029 (0.01)	[-0.023; 0.017]
		75	3.926	0.0088 (0.009)	[-0.008; 0.026]	3.871	-0.0012 (0.011)	[-0.022; 0.019]
EU-14	No	25	3.779	0.0076 (0.015)	[-0.021; 0.036]	3.746	-0.0225 (0.018)	[-0.057; 0.012]
		50	3.886	0.0158 (0.014)	[-0.012; 0.044]	3.851	-0.0181 (0.018)	[-0.053; 0.017]
		75	4.002	0.0247 (0.016)	[-0.006; 0.055]	3.946	-0.0141 (0.019)	[-0.052; 0.024]
EU-14	Yes	25	3.779	0.0178 (0.014)	[-0.009; 0.045]	3.746	0.0003 (0.018)	[-0.036; 0.036]
		50	3.886	0.025 (0.014)	[-0.002; 0.052]	3.851	0.0057 (0.018)	[-0.030; 0.042]
		75	4.002	0.0328 (0.015)	[0.004; 0.062]	3.946	0.0106 (0.019)	[-0.027; 0.049]

Table A7. Marginal effects of AI innovation capacity on energy intensity, conditional on digital transformation, based on non-linear models and models with lagged variables (cont.)

Type			Non-linear models (NL)			Lagged models (LM)		
G	C	P	I_desi	me/se	95% ci	I_desi	me/se	95% ci
EU-13	No	25	3.409	-0.0142 (0.015)	[-0.038; 0.01]	3.394	-0.001 (0.013)	[-0.026; 0.024]
		50	3.601	-0.0084 (0.01)	[-0.027; 0.011]	3.569	0.0047 (0.011)	[-0.017; 0.026]
		75	3.708	-0.0051 (0.009)	[-0.023; 0.013]	3.667	0.0079 (0.011)	[-0.014; 0.029]
EU-13	Yes	25	3.409	-0.0096 (0.012)	[-0.032; 0.013]	3.394	0.0007 (0.013)	[-0.025; 0.027]
		50	3.601	-0.0011 (0.01)	[-0.020; 0.018]	3.569	0.0082 (0.012)	[-0.014; 0.031]
		75	3.708	0.0036 (0.01)	[-0.015; 0.022]	3.667	0.0123 (0.012)	[-0.010; 0.035]

Note. Bolded results indicate statistically significant effects, G stands for a group of economies, C for control variables, P for percentiles, and me/se for marginal effect/standard error.

Source: author's calculations.

Convergence to multiple price equilibria in algorithmic collusion

Krzysztof Siegieda^a

Abstract. Simulation studies of algorithmic collusion often summarise results through average prices and profits. Such summaries conceal the variety and distribution of distinct prices, cycles, and collusive regimes. This limitation is salient given that game theory predicts multiple equilibria, implying a distribution rather than a single outcome. This paper examines the heterogeneity hidden by aggregate metrics in a sequential pricing duopoly with Q learning agents. Using an established simulation framework for studying actions of pricing algorithms, the simulations cover three memory configurations: symmetric one period, one period versus none, and one period versus two period memory. K means cluster analysis of steady state prices reveals a spectrum: monopoly focal price equilibria, collusive focal price equilibria at non monopoly grid points and incomplete Edgeworth cycles with supra competitive averages. These regimes are consistent with the multiplicity of the Markov perfect equilibria in the game theory literature. In a symmetric baseline, the monopoly outcome is a minority. Memory asymmetry eliminates or reduces cycling and increases the share of fixed price outcomes, stabilising rather than breaking collusion. Aggregate numbers collapse qualitatively distinct results into a single figure; cluster analysis uncovers the hidden heterogeneity behind those averages and shows that algorithmic coordination spans a spectrum of regimes rather than a single outcome. For the competition policy, instruments that distinguish between pricing regimes are more informative than price level benchmarks alone.

Keywords: algorithmic collusion, cluster analysis, reinforcement learning, Q-learning, pricing algorithms

JEL: K21, L13, L41, C63, C73

1. Introduction

The widespread adoption of algorithmic pricing has fundamentally changed the dynamics of competition in many markets. From petrol retailing to e-commerce, firms increasingly delegate pricing decisions to self-learning algorithms that adjust prices in real time on the basis of the observed market conditions. A growing body of research warns that such algorithms, even without explicit communication, may autonomously learn to coordinate on supra-competitive prices (Calvano et al., 2020; Ezrachi & Stucke, 2016). This possibility – the ‘algorithmic collusion’ – challenges the foundations of the competition policy, which traditionally relies on evidence of explicit agreements between human decision-makers (Harrington, 2018).

^a Student at Collegium of Economic Analysis, SGH Warsaw School of Economics, ul. Madalińskiego 6/8, 02–513 Warszawa, Poland, e-mail: ks142569@student.sgh.waw.pl, ORCID: <https://orcid.org/0009-0009-9839-9948>.

A key insight from the recent literature is that simple reinforcement learning algorithms, such as tabular Q-learning, are capable of converging to collusive outcomes in both simultaneous and sequential pricing games (Calvano et al., 2020; Klein, 2021). Klein (2021), in particular, demonstrates that in a sequential duopoly with a limited set of discrete prices, Q-learning agents consistently coordinate on outcomes that yield average per-period profits well above the competitive benchmark. Siegieda and Zawisza (2026) extend this framework to the asymmetric memory configurations and show that memory asymmetry stabilises the collusion rather than breaks it. However, like much of the existing literature, these studies report outcomes through aggregate metrics – mean prices and mean profits – that compress information and may hide the full distribution of pricing behaviours across simulation runs.

This paper argues that to fully understand the nature and stability of algorithmic collusion, one must look beyond averages and examine the distribution of the outcomes directly. An average price of, for example, 0.42 could reflect a mixture of fundamentally different behaviours: some runs converge to a stable monopoly price, some others to alternative focal prices, and the rest to volatile price cycles whose mean happens to lie in a similar range. From a competition-policy perspective, these distinctions matter. A market trapped in a stable collusive equilibrium at the monopoly price poses a different threat to consumer welfare than a market characterised by Edgeworth-like price cycles with a supra-competitive average. Yet standard aggregate metrics fail to capture this diversity.

I propose cluster analysis as a systematic method for identifying distinct pricing regimes – groups of simulation runs that share similar dynamic patterns – and apply it to the canonical sequential pricing framework studied by Klein (2021). The approach extracts a small number of features characterising the steady-state pricing behaviour, and then applies K-means clustering to uncover the natural groupings present in the data. Rather than reducing the simulation output to a handful of aggregate statistics, this method preserves the multi-dimensional structure of the results and allows a more comprehensive economic interpretation. Three memory configurations, i.e. symmetric one-period baseline, one-period versus no memory and one-period versus two-period memory are analysed in the study. The analysis starts from the symmetric baseline, where any heterogeneity uncovered arises endogenously from the stochastic learning dynamics rather than from built-in asymmetries. The clustering methodology has a general character and can be extended to finer price grids and alternative learning algorithms in the future research.

2. Literature review

2.1. Foundations of algorithmic collusion

The question whether pricing algorithms can learn to collude without explicit communication has attracted intense scholarly attention since the seminal legal analyses of Ezrachi and Stucke (2016) and Mehra (2016). These early contributions indicated the possibility that autonomous algorithms might achieve outcomes economically equivalent to human cartels yet fall outside the existing antitrust laws because they lack an explicit agreement. Subsequent economic literature provided robust simulation evidence that this concern is justified. Calvano et al. (2020) show that independent Q-learning agents in a simultaneous-move duopoly can learn to sustain supra-competitive prices through reward-punishment strategies, with outcomes satisfying the Nash criterion. Klein (2021) extends this analysis to a sequential pricing framework based on Maskin and Tirole (1988), demonstrating that Q-learning agents with one-period memory consistently achieve supra-competitive profits. Klein, however, does not systematically report the distribution of final prices across runs, nor does he examine whether qualitatively different pricing behaviours underlie the aggregate results. Other studies reinforce these findings in related settings (Abada & Lambin, 2020; Waltman & Kaymak, 2008). A common feature across this strand of literature is the reliance on aggregate performance metrics, which, while informative about central tendencies, provide no information about the distribution of outcomes or the potential existence of distinct behavioural clusters.

2.2. Algorithmic heterogeneity and collusion

A related branch of research examines whether algorithmic collusion persists when the competing algorithms are not identical. Most existing work focuses on heterogeneity across algorithm classes. Den Boer et al. (2022) show that Q-learning agents can be exploited by simpler bandit algorithms, leading to more competitive outcomes. Abada et al. (2024) and Bichler et al. (2024) find that alternative algorithms often outperform Q-learners and push prices downward, suggesting that mixing fundamentally different algorithm types might undermine collusion. Much less is known about heterogeneity within the same algorithm family – for example, when all firms use Q-learning but differ in memory depth. Abada et al. (2025) explicitly identify this as a gap. Siegieda and Zawisza (2026) address it by systematically varying memory depth in a sequential pricing duopoly, showing that memory asymmetry stabilises rather than breaks collusion. This paper builds directly on their work by applying cluster analysis to decompose aggregate results into their constituent pricing regimes and by tracing how regime composition changes across symmetric and asymmetric memory configurations.

2.3. Pricing dynamics and the analysis of regimes

Beyond the collusion-versus-competition dichotomy, several studies examine the finer structure of pricing dynamics. Maskin and Tirole (1988) characterise two types of Markov-perfect equilibria in sequential pricing: focal price equilibria and Edgeworth price cycle equilibria. Empirical research on retail petrol markets shows patterns consistent with Edgeworth cycles, with considerable variation in cycle length and amplitude (Byrne & de Roos, 2019; Eckert, 2013), suggesting that substantial heterogeneity might exist even within a broadly cyclic regime. In the algorithmic collusion literature, Klein (2021) notes that finer price grids produce Edgeworth-like cycles but does not systematically categorise their forms or their frequency in relation to fixed-price outcomes. Likewise Calvano et al. (2020), who treat convergence to different collusive price levels as incidental. This paper addresses this gap by applying cluster analysis to the steady-state pricing data from many simulation runs, identifying the natural groupings that emerge from the learning dynamics and checking whether they correspond to the canonical equilibrium types or to previously undocumented patterns.

3. Model and methodology

3.1. Sequential pricing duopoly

The environment follows the sequential pricing framework of Maskin and Tirole (1988), as adapted by Klein (2021). Two companies produce a homogeneous good at zero marginal cost and compete in an infinite-horizon, discrete-time setting. In each period $t = 1, 2, \dots$, only one of them adjusts its price, while the other firm's price remains fixed. The firms alternate deterministically, so that Firm 1 moves in odd periods and Firm 2 in even periods. The price set by the company inactive in a given period is carried over from the previous period.

Both firms choose prices from a common discrete grid: $P = \{0, \frac{1}{k}, \dots, 1\}$, with $k = 6$ price intervals, yielding the set $\{0, 0.167, 0.333, 0.5, 0.667, 0.833, 1\}$. This is the same grid used by Klein in his baseline specification, and it allows the direct comparability of results.

Demand is allocated according to an undercutting rule. When firm i sets price p_i and firm j sets price p_j , firm i captures the entire market and faces a linear demand schedule if $p_i < p_j$, splits the market equally if $p_i = p_j$, and sells nothing if $p_i > p_j$. Formally:

$$D_i(p_i, p_j) = \begin{cases} 1 - p_i & \text{if } p_i < p_j \\ \frac{1}{2}(1 - p_i) & \text{if } p_i = p_j \\ 0 & \text{if } p_i > p_j \end{cases} \quad (1)$$

Given zero marginal costs, firm i 's stage profit is simply:

$$\pi_i(p_i, p_j) = p_i \cdot D_i(p_i, p_j). \quad (2)$$

The monopoly price on the grid is $p^m = 0.5$, with the associated monopoly profit $\pi^m = 0.125$. Following Klein, I take as the competitive benchmark the average per-period profit of the most competitive Markov-perfect equilibrium (MPE) that corresponds to an Edgeworth price cycle, as described by Maskin and Tirole (1988), which yields approximately 0.061 for $k = 6$.

The alternating-move structure implies that, under the Markov assumption that strategies condition only on payoff-relevant variables, the relevant state for the firm about to move is simply the last price set by its competitor. This one-period memory structure is the natural starting point for the analysis and the focus of this paper.

3.2. Q-learning algorithm

Both companies are modelled as independent Q-learning agents (Watkins & Dayan, 1992). Q-learning is a model-free reinforcement learning algorithm that seeks to estimate the long-run value of taking a particular action in a particular state, without requiring knowledge of the transition dynamics or the competitor's strategy.

Each firm i maintains Q-matrix $Q_i(s, a)$ of the size $|S| \times |P|$, where s in S is the state (the competitor's last price) and a in P is the chosen price. In the baseline symmetric configuration, both firms use one-period memory, so $|S| = |P| = 7$.

After each period in which firm i acts, it updates the Q-value corresponding to the state-action pair (s, a) that was just obtained, according to:

$$Q_i(s, a) \leftarrow (1 - \alpha)Q_i(s, a) + \alpha \left[\pi + \delta \pi' + \delta^2 \max_{a'} Q_i(s', a') \right], \quad (3)$$

where α in $(0,1)$ is the learning rate, δ in $[0,1)$ is the discount factor, π is the immediate profit from action a , π' is the profit earned in the period following the rival's move, and s' is the successor state. This two-step update reflects the sequential nature of the game: after firm i acts, firm j responds, and only then does firm i observe the full consequences of its action.

Action selection follows an ε -greedy policy. With probability ε_t , a firm explores the action space by picking a price uniformly at random from P . With probability $1 - \varepsilon_t$,

it exploits its current knowledge by choosing the action that maximises $Q_i(s, \cdot)$. The exploration probability decays over time according to:

$$\varepsilon_t = (1 - \theta)^t, \quad (4)$$

where decay parameter θ is calibrated so that $\varepsilon_0 = 1$ (full exploration at the start), $\varepsilon_{T/2} = 0.001$ (0.1% exploration at midpoint), and $\varepsilon_T = 10^{-6}$ (negligible exploration at the end).

The simulation parameters are set exactly as in Klein (2021): learning rate $\alpha = 0.3$, discount factor $\delta = 0.95$, and the total horizon of $T = 500,000$ periods. Both firms' Q-tables are initialised to zero. For each experimental condition, I run 1,000 independent simulations with distinct random seeds. The final $M = 50,000$ periods (10% of the horizon) are treated as the steady state and used for all the subsequent analyses.

3.3. Memory configurations

The defining feature of a Q-learning agent in this environment is its state representation – what information it conditions its pricing decision on. Under the Markov assumption that strategies depend only on payoff-relevant variables, the natural state for a firm about to move is the last price set by its competitor. This one-period memory structure is the starting point for most of the algorithmic collusion literature (Calvano et al., 2020; Klein, 2021). However, there is no reason that all firms in a market must use the same state representation. Differences in memory depth, i.e. how many periods of past prices an agent encodes in its state, constitute a natural and policy-relevant form of algorithmic heterogeneity. This paper examines three configurations that span a range from a complete symmetry to two qualitatively different types of memory asymmetry. The three memory configurations examined here are those introduced by Siegieda and Zawisza (2026), whose aggregate results provided the motivation for this cluster-based re-analysis. These are labelled:

- Scenario A: Symmetric one-period memory. Both firms condition their decisions on the competitor's most recent price. For firm i moving at time t , the state is $s_t = p_{j,t-1}$, where j denotes the competitor. The state space for each firm has the size $|S| = |P| = 7$. This is the canonical baseline studied by Klein (2021) and the configuration that yields the largest variety of pricing regimes;
- Scenario B: One-period vs. no memory. Firm 1 uses one-period memory ($s_t = p_{2,t-1}$), while Firm 2 has no memory at all – it observes a constant state regardless of the history of play ($|S_2| = 1$). A memoryless agent cannot distinguish between different competitive situations; it effectively learns a single best price

irrespective of what its competitor does. This configuration tests whether a sophisticated agent can exploit a naive one, or whether the naive agent nonetheless manages to sustain supra-competitive coordination;

- Scenario C: One-period vs. two-period memory. Firm 1 uses one-period memory ($s_t = p_{2,t-1}$), while Firm 2 uses two-period memory, conditioning on both the competitor's last price and its own last price: $s_t = (p_{2,t-1}, p_{1,t})$. The state space for Firm 2 expands to $|S_2| = |P|^2 = 49$, offering more strategic possibilities at the cost of substantially slower learning. This configuration examines whether deeper memory confers a strategic advantage and whether the expanded state space hampers convergence.

The three scenarios together allow a systematic assessment of how the memory heterogeneity – both its presence and its type – shapes the set of the resulting pricing regimes. For each scenario, I simulated 1,000 independent runs, each with a distinct random seed.

3.4. Performance metrics

To characterise the outcomes of each simulation run, I compute several metrics. The first is the average price, a direct measure of market outcomes and a proxy for consumer welfare. For each firm i :

$$\text{avg}_{p_i} = \left(\frac{1}{M}\right) \cdot \sum_{t=T-M+1}^T p_{i,t}. \quad (5)$$

I also report the market-wide average $\text{avg}_p = (\text{avg}_p p_1 + \text{avg}_p p_2) / 2$.

Profitability measures the actual per-period profits obtained in the steady state:

$$\text{avg}_{\pi_i} = \left(\frac{1}{M}\right) \cdot \sum_{t=T-M+1}^T \pi_i(p_{i,t}, p_{j,t}). \quad (6)$$

Optimality captures how close each firm's learnt behaviour is to the best response given the competitor's strategy:

$$\Gamma_i = \frac{Q_i(p_{i,T}, s_T)}{\max_{p \in P} Q_i(p, s_T)}, \quad (7)$$

where Q_i is the firm's final Q-matrix and s_T is the state in the final period. The value of $\Gamma_i = 1$ (within a tolerance of 10^{-5}) indicates that the firm is playing the best response.

A simulation run is classified as a Nash equilibrium if $\Gamma_1 = \Gamma_2 = 1$ simultaneously. Both optimality and the Nash criterion serve as supplementary diagnostics – they help interpret the economic nature of the regimes uncovered by clustering, but they are not the primary focus of the analysis. The optimality measure Γ_i and the Nash-equilibrium criterion are not used as input features for clustering. Clustering is based purely on the price features described in Section 3.5; the optimality measures appear only in the post-clustering analysis (Tables 1–3) to validate whether the identified fixed-price clusters correspond to the Nash equilibria, following Klein (2021).

These aggregate metrics are useful for establishing comparability with Klein (2021) and for providing a first summary of each scenario. However, as argued in Section 1, they collapse the rich heterogeneity of individual runs into single numbers. The cluster analysis described next is designed to overcome this limitation.

3.5. Cluster analysis methodology

The core methodological contribution of this paper is the application of cluster analysis to the steady-state pricing data. Rather than reducing the output of 1,000 simulation runs to a few means, I extract features that capture the structure of each run’s final pricing behaviour and then use a clustering algorithm to partition the runs into groups that share similar characteristics.

The choice of features differs between the symmetric and asymmetric scenarios, reflecting a fundamental difference in the economic interpretation of firm labels.

The symmetric scenario (Scenario A: 1 vs. 1). In the symmetric configuration, both companies are *ex ante* identical – they use the same memory depth, the same learning parameters, and face the same environment. The assignment of the ‘Firm 1’ and ‘Firm 2’ labels is therefore arbitrary, and a given pricing pattern should be recognised as the same regime regardless of which firm happened to set the higher price. For example, a run where Firm 1 consistently charges 0.5 and Firm 2 charges 0.333 is economically distinct from the one where both charge 0.5, but it is the same regime regardless of which of the two companies set the high price. To achieve this label-invariance, I transform the raw price pair (p_1, p_2) into the following order statistics:

$$p_{min,t} = \min(p_{1,t}, p_{2,t}), \quad (8)$$

$$p_{max,t} = \max(p_{1,t}, p_{2,t}), \quad (9)$$

and then average these quantities over the steady-state periods to obtain two features per run:

$$\text{avg_}p_{\min} = \left(\frac{1}{M}\right) \cdot \sum_{t=T-M+1}^T p_{\min,t}, \quad (10)$$

$$\text{avg_}p_{\max} = \left(\frac{1}{M}\right) \cdot \sum_{t=T-M+1}^T p_{\max,t}. \quad (11)$$

This two-dimensional representation summarises the steady-state behaviour of each run using two features. Runs where both firms set the same stable price cluster along the 45-degree line where $\text{avg_}p_{\min} \approx \text{avg_}p_{\max}$. Runs with persistent price dispersion or large-amplitude cycles appear off the diagonal, with $\text{avg_}p_{\min}$ substantially below $\text{avg_}p_{\max}$. The distance between these two values also proxies for the amplitude of price cycles when they are present.

Asymmetric scenarios (Scenarios B and C: 1 vs. 0 and 1 vs. 2). In the asymmetric configurations, the two companies differ in their memory depth and, as a consequence, in their strategic capabilities. The labels ‘Firm 1’ and ‘Firm 2’ are no longer arbitrary – they convey economically meaningful information about which agent has more memory and which has less. A run where the one-period-memory firm charges 0.5 and the zero-memory firm charges 0.333 is a different outcome from the one where the one-period firm charges 0.333 and the zero-memory firm charges 0.5, because the distribution of profits and the strategic roles are reversed. It is therefore appropriate to use the raw firm-specific mean prices directly as features:

$$\text{avg}_{p_1} = \left(\frac{1}{M}\right) \cdot \sum_{t=T-M+1}^T p_{1,t}, \quad (12)$$

$$\text{avg_}p_2 = (1/M) \cdot \sum_{t=T-M+1}^T p_{2,t}. \quad (13)$$

This preserves the identity of each firm and allows the clustering algorithm to identify regimes that differ according to which firm takes the lead in setting higher or lower prices. For Scenario B (1 vs. 0), Firm 1 has one-period memory and Firm 2 has none. For Scenario C (1 vs. 2), Firm 1 has one-period memory and Firm 2 has two-period memory.

For both the symmetric and asymmetric scenarios, the clustering is based solely on the two price features described above, namely the order-statistic averages for

Scenario A and the firm-specific averages for scenarios B and C. This deliberately lean set of features captures the essential economic distinction between the fixed-price coordination (where the two features are equal or nearly equal) and cycling (where they differ), while keeping the dimensionality low enough for robust clustering. The more comprehensive statistics presented in Tables 1–3 are computed after the clusters are formed, and serve solely to characterise each regime and enable the reader to assess the economic interpretation of the groupings.

Standardisation and clustering algorithm. For all the three scenarios, the features are standardised to a zero mean and unit variance before clustering to ensure that each dimension contributes equally to the distance metric:

$$X_{scaled} = (X - \mu) / \sigma. \quad (14)$$

I use the K-means algorithm (MacQueen, 1967) with $n_{init} = 50$ random initialisations to mitigate the sensitivity to the centroid initialisation. The choice of k is guided by a combination of the elbow method (plotting within-cluster sum of squares against k) and the silhouette score (Rousseeuw, 1987), which measures how similar each point is to its own cluster compared to the other clusters. These are standard tools in statistical learning for cluster validation (James et al., 2023, pp. 15–67). The final choice of k also takes into account a qualitative criterion: the resulting clusters must be economically interpretable. Across all the three scenarios, both the elbow and the silhouette diagnostics point toward $k = 4$ or $k = 5$ as the natural number of clusters. I select the value that yields the clearest separation into economically distinct regimes. The specific choice is explained and presented alongside the results for each scenario in the next section. The methodology is implemented in Python; analogous clustering procedures are widely available in other statistical computing environments, including R (Kamiński & Zawisza, 2012). Once each simulation run has been assigned to a cluster, I examine the within-cluster distributions of prices, profits, and dynamic patterns. For each cluster, I report the number of the assigned runs, the mean and standard deviation of prices, the mean profitability, and the share of periods in which both firms set equal prices.

4. Results

This section presents the results of the cluster analysis for each of the three memory configurations. For each scenario, I show the outcome of the K-means clustering and characterise the resulting groups in terms of average prices and profitability. Three qualitatively distinct types of regimes emerge: monopoly focal price (stable coordination on the joint-profit-maximising price of 0.5), collusive focal price (stable

coordination on a supra-competitive price lower than 0.5), and incomplete Edgeworth cycles (persistent price undercutting with truncated, downward phases and supra-competitive averages).

Before examining the cluster compositions, let us note a general feature of the learning dynamics. The exploration probability decays to negligible levels by roughly the midpoint of the simulation horizon ($t \approx 250\,000$), so the steady-state behaviour is exploitation-driven. Which regime a run converges to depends on the stochastic early exploration phase; once the agents lock in, the pattern – whether a fixed price or a cycle – stays unchanged. In the cycling clusters, the cycles are regular and stable throughout the steady state. Thus, the identified clusters represent stable long-run outcomes rather than transient dynamics. Illustrative price and profit trajectories for Scenario A are presented in Figures A1 and A2 in the Appendix, where the cluster-level averages are plotted alongside the scenario-wide mean. These figures confirm that, as stated above, the clusters represent stable long-run outcomes rather than transient dynamics.

4.1. Scenario A: a symmetric one-period memory (1 vs. 1)

The symmetric one-period memory configuration is the canonical baseline studied by Klein (2021). The diagnostics – as described in section 3.5 – point to $k = 4$ clusters. As seen in Table 1, the clusters are clearly separated by the price level and dispersion, and by profitability. Cluster 0 ($n = 295$) is the only group with substantial price dispersion and asymmetric profits, while in Clusters 1–3, near-perfect price equality and symmetric profits close to their theoretical grid-point levels could be observed. These differences reflect distinct pricing regimes that are characterised below.

Table 1. Scenario A: symmetric one-period memory – steady state (fixed or cyclical) prices and profits by cluster

Cluster	n	avg_p1	avg_p2	avg_p_min	avg_p_max	std_p_min	avg_pi	Interpretation
0	295	0.469	0.460	0.302	0.627	0.150	0.094	Incomplete collusive Edgeworth cycles
1	420	0.330	0.330	0.330	0.331	0.027	0.110	Partial collusive focal price (1/3)
2	252	0.500	0.500	0.499	0.501	0.015	0.125	Monopoly focal price (1/2)
3	33	0.667	0.667	0.667	0.667	0.000	0.111	Collusive focal price (2/3)

Source: author's calculations.

Cluster 0 ($n = 295$) is the only group with substantial price dispersion. The large discrepancy between avg_p_min (0.302) and avg_p_max (0.627), combined with the

high standard deviation of p_{min} (0.150) indicates persistent price cycling with truncated downward phases. These cycles are incomplete in the sense that the undercutting phase stops well before prices reach the competitive floor, after which one firm resets the cycle with a large upward jump. The average profit of 0.094 is well above the competitive benchmark of 0.061, confirming that this regime, though dynamic, remains supra-competitive. This regime is labelled incomplete collusive Edgeworth cycles. Cluster 2 ($n = 252$) is the only cluster achieving the monopoly price (0.5) and maximum profits (0.125). This is the monopoly focal price regime. Clusters 1 ($n = 420$) and 3 ($n = 33$) are stable fixed-price regimes at non-monopoly grid points, specifically 0.333 (1/3) and 0.667 (2/3), with near-perfect price equality, very low standard deviations, and supra-competitive profits (0.110 and 0.111, respectively). They are labelled collusive focal price equilibria.

4.2. Scenario B: one-period memory vs. no memory (1 vs. 0)

In this asymmetric configuration, Firm 1 has one-period memory and Firm 2 has no memory. The diagnostics suggest $k = 4$ clusters. Table 2 presents the cluster characteristics. All four clusters show stable fixed-price outcomes with near-zero standard deviations; no Edgeworth cycles appear in this scenario. The most striking feature of Table 2 is the dominance of Cluster 0, which captures 74.2% of all runs and converges to the grid price of 1/6 (0.167). This price yields the profit of 0.069, which is only marginally above the competitive benchmark of 0.061. The monopoly outcome (Cluster 3, with the profit of 0.125) emerges in only 8.4% of runs. Clusters 1 and 2 replicate the collusive focal prices at 1/3 and 2/3 observed in Scenario A, with profits of 0.111 in both cases. The absence of cycles is not a transient phenomenon: the inspection of the full price trajectories confirms that once the agents settle on a fixed price, they never depart from it during the steady state.

Table 2. Scenario B: asymmetric one-period memory vs no memory – steady state (fixed or cyclical) prices and profits by cluster

Cluster	n	avg_p1	avg_p2	avg_p_min	avg_p_max	std_p_min	avg_pi	Interpretation
0	742	0.167	0.167	0.167	0.167	0.000	0.069	Minimal supra-competitive focal price (1/6)
1	164	0.332	0.333	0.332	0.333	0.013	0.111	Collusive focal price (1/3)
2	10	0.667	0.667	0.667	0.667	0.000	0.111	Collusive focal price (2/3)
3	84	0.500	0.500	0.500	0.500	0.001	0.125	Monopoly focal price (1/2)

Source: author's calculations.

The dominance of Cluster 0 is economically significant. Despite the presence of a memoryless agent, the market does not collapse into competitive pricing; instead, it stabilises at the lowest price on the grid that still yields supra-competitive profits. I label this regime the minimal supra-competitive focal price. The mechanism is collusive in the sense that both firms maintain a stable, equal price without undercutting, yet the resulting profit (0.069) is barely distinguishable from the competitive benchmark (0.061). This contrasts with other fixed-price clusters, where profits are substantially higher.

4.3. Scenario C: one-period memory vs. two-period memory (1 vs. 2)

In this configuration, Firm 1 has one-period memory and Firm 2 has two-period memory. The diagnostics support $k = 4$ clusters. Table 3 presents the cluster-wise summary. Cluster 0 ($n = 512$) is the largest and converges to an essentially monopolistic outcome ($\text{avg_pi} = 0.124$), although with a slightly higher variance than in Scenarios A and B, likely due to slower learning caused by Firm 2's larger state space. Cluster 1 ($n = 406$) reproduces the collusive focal price at $1/3$ with symmetric profits (0.111). Cluster 3 ($n = 32$) is a collusive focal price near $2/3$, although with moderately higher variance than its counterparts in Scenario A. Cluster 2 ($n = 50$) stands out. It consists of incomplete Edgeworth cycles with a striking profit asymmetry: the two-period memory agent (Firm 2) earns 0.132, while the profit of the one-period memory agent is only 0.059. The average market profit of 0.096 remains well above the competitive benchmark. This is the only regime where the profit distribution is visibly unequal in favour of one company, and it illustrates how deeper memory can be leveraged to capture a disproportionate share of the supra-competitive surplus within a cycling regime. The cycles in this cluster are persistent throughout the entire steady-state window and do not show any systematic change in amplitude or frequency over time.

Table 3. Scenario C: asymmetric one-period memory vs two-period memory – steady state (fixed or cyclical) prices and profits by cluster

Cluster	n	avg_p1	avg_p2	avg_p_min	avg_p_max	std_p_min	avg_pi	Interpretation
0	512	0.496	0.501	0.494	0.503	0.044	0.124	Monopoly focal price (near 1/2)
1	406	0.333	0.333	0.333	0.334	0.016	0.111	Collusive focal price (1/3)
2	50	0.599	0.344	0.303	0.641	0.141	0.096*	Incomplete Edgeworth cycles
3	32	0.651	0.672	0.648	0.674	0.089	0.007	Collusive focal price (near 2/3)

Note. * In Cluster 2, the two-period memory agent (Firm 2) earns an average profit of 0.132, while the one-period memory agent (Firm 1) earns 0.059. The table reports the average market profit (avg_pi) of 0.096.


Source: author's calculations.

4.4. Cross-scenario comparison

The tables in the preceding subsections characterise each cluster in isolation. To see how the overall landscape of pricing regimes changes with memory asymmetry, I turn to the full distribution of steady-state prices across all the 1,000 runs in each scenario. Figure 1 shows the empirical cumulative distribution function (CDF) of the average market price for each scenario. Figure 2 presents the corresponding side-by-side histograms.

Under the specific parameterisation studied, the CDF of Scenario C lies almost everywhere below that of Scenario B, which in turn lies below that of Scenario A. This near-uniform ordering means that the price distribution in Scenario C first-order stochastically dominates (FOSD) that in Scenario B, and Scenario B stochastically dominates Scenario A. In other words, among the three configurations examined, prices are unambiguously higher in Scenario C than in Scenario B, and in Scenario B than in Scenario A – a ranking that a welfare-minded regulator would interpret as increasing consumer harm (Mas-Colell et al., 1995). The rightward shift of probability mass is equally clear in the histograms: moving from Scenario B (where the dominant outcome is the minimal supra-competitive price of 1/6) through Scenario A (where the supra-competitive price of 1/3 and incomplete cycles together account for over 70 % of runs) to Scenario C (where the monopoly price becomes the majority regime), the distribution visibly concentrates on higher prices.

Figure 1. Empirical CDF of steady-state market price by scenario



Source: author's calculations.

Figure 2. Distribution of steady-state market prices by scenario

Source: author's calculations.

Table 4 quantifies these shifts using the regime typology developed in Sections 4.1–4.3.

Table 4. Regime distribution across scenarios (percentage of runs)

Regime	Scenario A (1v1)	Scenario B (1v0)	Scenario C (1v2)
Monopoly focal price (1/2)	25.2%	8.4%	51.2%
Collusive focal price (1/3)	42.0%	16.4%	40.6%
Collusive focal price (2/3)	3.3%	1.0%	3.2%
Collusive focal price (1/6)	–	74.2%	–
Incomplete Edgeworth cycles	29.5%	–	5.0%

Source: author's calculations.

The three main findings are as follows. Firstly, the symmetric baseline (Scenario A) yields the greatest diversity of regimes: nearly 30% of runs fall into incomplete Edgeworth cycles, only a quarter achieve the monopoly price, and the remaining 45% are split between the supra-competitive prices of 1/3 and 2/3. The commonly reported aggregate statistics thus conceal a heterogeneous mix. Secondly, the memoryless agent in Scenario B eliminates cycles completely and anchors the market to the lowest supra-competitive price of 1/6 in almost three-quarters of the runs. Thirdly, introducing a two-period memory agent (Scenario C) more than doubles the probability of the monopoly outcome compared to the symmetric baseline (from 25.2% to 51.2%) and reduces the frequency of cycles from 29.5% to 5.0%. Memory asymmetry therefore does not merely preserve collusion – it systematically shifts the distribution of prices towards higher, more profitable outcomes, as the stochastic dominance ordering in figures 1 and 2 demonstrates.

5. Discussion

5.1. The spectrum of algorithmic coordination

The cluster analysis shows that the algorithmic collusion is not a single outcome, but a spectrum of at least three pricing regimes. The first of them is the monopoly focal price, where both firms settle on the joint-profit-maximising price of 0.5. In the second, i.e. the collusive focal prices at non-monopoly grid points such as $1/3$ or $2/3$, the coordination is equally stable but profits are lower. The third one, incomplete Edgeworth cycles, consists in firms undercutting each other repeatedly, and yet the average price remaining close to the monopoly level. These regimes are all supra-competitive, but they differ fundamentally in their underlying dynamics and the consequences for the customers' wellbeing. In the symmetric baseline, only about a quarter of runs reach the monopoly outcome; the rest are split between other focal prices and cycles. The binary view of 'collusion or not' that dominates the literature thus misses the variety of ways in which algorithms can coordinate.

5.2. Memory asymmetry as a stabiliser of collusion

Across the three scenarios examined, memory asymmetry does not undermine algorithmic collusion; instead, it stabilises coordination. In the symmetric baseline, only about half of all the runs converge to a fixed-price equilibrium, while the rest follow cycles. When one firm has no memory (Scenario B), every run settles on a fixed price, eliminating cycles completely. When one firm has a two-period memory (Scenario C), fixed-price equilibria account for 95 % of runs, and the monopoly outcome alone captures over half of all runs – more than twice as many as in the symmetric case. Thus, within the parameter space studied, asymmetry pushes the market towards stationary collusion, and deeper memory on one side strongly favours the most profitable equilibrium.

5.3. Why are averages misleading

The standard practice of reporting only the average prices and profits masks the above-described entire structure. In a symmetric scenario, the global average profit of 0.109 is a blend of runs that yield 0.125, 0.110 and 0.094, each corresponding to a qualitatively different regime. A policymaker seeing only the aggregate number would not know if full monopoly collusion occurs in the minority of cases, or if cycling runs produce near-monopoly average prices through a competitive-looking process. Cluster analysis avoids this pitfall by letting the data reveal its own structure. It partitions the simulations into economically meaningful groups without imposing prior assumptions, using only the steady-state price patterns. This method provides a more comprehensive and actionable picture, and I suggest it become a standard robustness check in the future algorithmic-collusion studies.

5.4. Policy implications: detection and monitoring

For the competition authorities, the results imply that diagnosing algorithmic collusion requires looking beyond average price levels. Although authorities do not have access to thousands of repeated simulations, they can obtain historical time series of firm-level prices once an investigation begins – a single realisation of the price-generating process. By discretising the price grid and comparing the observed trajectory of the cluster prototypes identified here, they could assess whether the pattern resembles a stable focal-price regime, an Edgeworth cycle or a competitive benchmark. Different regimes produce different dynamic signatures, and cluster-based monitoring tools would thus allow more targeted interventions than simple price-level comparisons.

Two specific regimes warrant particular attention. One of them is incomplete Edgeworth cycles, which are deceptive in that prices change frequently, giving the impression of competition, yet the average price hovers near the monopoly level and the consumer harm is severe. The other is collusive focal prices at non-monopoly grid points, which are equally stable and harmful in relation to the competitive benchmark; disrupting one focal price may simply cause the algorithms to re-coordinate on another. Hence, the entire grid of possible focal prices – $1/6$, $1/3$, $1/2$, $2/3$ – should be of regulatory concern, not only the monopoly outcome.

5.5. Limitations and future research

This study is confined to a specific parameterisation: a coarse price grid with $k = 6$ intervals, fixed learning parameters, and a single horizon of 500,000 periods. Klein (2021) showed that finer grids ($k = 24$) amplify cycling dynamics. Whether the same regime typology and the stabilising effect of memory asymmetry hold under finer grids, alternative exploration schedules, or other learning algorithms (e.g. deep Q-networks) remains an open question. Additionally, the cluster analysis is descriptive; it does not explain why a given run converges to one regime rather than another. Understanding the determinants of regime selection, perhaps by analysing early-stage learning trajectories, is a promising direction for future research. Finally, a full consumer-wellbeing quantification of the different regimes, incorporating the deadweight loss from supra-competitive prices and the additional costs of price volatility would strengthen policy conclusions.

6. Conclusions

This paper set out to look beyond the aggregate statistics that dominate the algorithmic collusion literature and to show the true heterogeneity of outcomes in a sequential pricing duopoly. Using cluster analysis on the steady-state pricing data of Q-learning agents, I found that algorithmic coordination is not a single monolithic phenomenon, but a spectrum of three distinct regimes: monopoly focal price,

collusive focal price at other grid points, and incomplete Edgeworth cycles. Each regime is supra-competitive, yet each operates through a different mechanism and poses distinct challenges for detection and policy.

A key finding is that introducing memory asymmetry – a realistic form of between-firm heterogeneity – does not break collusion. On the contrary, it stabilises coordination, pushing the market towards fixed-price equilibria and, when one agent has deeper memory, strongly favouring the monopoly outcome. This result challenges the intuition that diverse algorithms automatically restore competition.

Methodologically, the paper demonstrates that cluster analysis is a simple yet powerful tool for revealing the structure hidden in simulation data. It avoids the information loss inherent in averaging and should be adopted as a standard complement to aggregate metrics in future research.

As regards competition policy, the findings suggest that market monitoring should move beyond price-level analysis. Tools that can classify pricing regimes based on their dynamic signatures could help authorities distinguish genuinely competitive markets from those in which algorithms autonomously coordinated. In an era of widespread algorithmic pricing, understanding what kind of collusion occurs in addition to whether it occurs at all will be essential.

References


- Abada, I., Harrington, J. E., Lambin, X., & Meylahn, J. M. (2025). *Algorithmic Collusion: Where are we and where should we be going?* (Working paper). <https://dx.doi.org/10.2139/ssrn.4891033>.
- Abada, I., & Lambin, X. (2020). *Artificial intelligence: Can seemingly collusive outcomes be avoided?* (Working paper). <https://dx.doi.org/10.2139/ssrn.3559308>.
- Abada, I., Lambin, X., & Tchakarov, N. (2024). Collusion by mistake: Does algorithmic sophistication drive supra-competitive profits?. *European Journal of Operational Research*, 318(3), 927–953. <https://doi.org/10.1016/j.ejor.2024.06.006>.
- Bichler, M., Durmann, J., & Oberlechner, M. (2024). *Online optimization algorithms in repeated price competition: Equilibrium learning and algorithmic collusion*. <https://doi.org/10.48550/arXiv.2412.15707>.
- den Boer, A. V., Meylahn, J. M., & Schinkel, M. P. (2022). *Artificial collusion: Examining supracompetitive pricing by Q-learning algorithms* (Amsterdam Law School Research Paper No. 25). <https://dx.doi.org/10.2139/ssrn.4213600>.
- Byrne, D. P., & de Roos, N. (2019). Learning to coordinate: A study in retail gasoline. *The American Economic Review*, 109(2), 591–619. <https://doi.org/10.1257/aer.20170116>.
- Calvano, E., Calzolari, G., Denicolò, V., Harrington, J. E., & Pastorello, S. (2020). Protecting consumers from collusive prices due to AI. *Science*, 370(6520), 1040–1042. <https://doi.org/10.1126/science.abe3796>.
- Calvano, E., Calzolari, G., Denicolò, V., & Pastorello, S. (2020). Artificial intelligence, algorithmic pricing, and collusion. *The American Economic Review*, 110(10), 3267–3297. <https://doi.org/10.1257/aer.20190623>.

- Eckert, A. (2013). Empirical studies of gasoline retailing: A guide to the literature. *Journal of Economic Surveys*, 27(1), 140–166. <https://doi.org/10.1111/j.1467-6419.2011.00698.x>.
- Ezrachi, A., & Stucke, M. E. (2016). *Virtual Competition. The promise and perils of the algorithm-driven economy*. Harvard University Press.
- Harrington, J. E. (2018). Developing competition law for collusion by autonomous artificial agents. *Journal of Competition Law & Economics*, 14(3), 331–363. <https://doi.org/10.1093/joclec/nhy016>.
- James, G., Witten, D., Hastie, T., Tibshirani, R., & Taylor, J. (2023). *An Introduction to Statistical Learning with Applications in Python*. Springer. https://doi.org/10.1007/978-3-031-38747-0_2.
- Kamiński, B., & Zawisza, M. (2012). *Receptury w R. Podręcznik dla ekonomistów*. Oficyna Wydawnicza SGH.
- Klein, T. (2021). Autonomous algorithmic collusion: Q-learning under sequential pricing. *The RAND Journal of Economics*, 52(3), 538–558. <https://doi.org/10.1111/1756-2171.12383>.
- Mas-Colell, A., Whinston, M. D., & Green, J. R. (1995). *Microeconomic Theory*. Oxford University Press.
- MacQueen, J. (1967). Some methods for classification and analysis of multivariate observations. In L. M. Le Cam & J. Neyman (Eds.), *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability* (Vol. 1, pp. 281–297). University of California Press.
- Maskin, E., & Tirole, J. (1988). A theory of dynamic oligopoly, II: Price competition, kinked demand curves, and Edgeworth cycles. *Econometrica*, 56(3), 571–599. <https://doi.org/10.2307/1911701>.
- Mehra, S. K. (2016). Antitrust and the Robo-Seller: Competition in the Time of Algorithms. *Minnesota Law Review*, 100, 1323–1375.
- Rousseeuw, P. J. (1987). Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics*, 20, 53–65. [https://doi.org/10.1016/0377-0427\(87\)90125-7](https://doi.org/10.1016/0377-0427(87)90125-7).
- Siegieda, K., & Zawisza, M. (2026). *Memory Arms Race in Q-learning: How Memory Asymmetry Shapes Algorithmic Collusion in Sequential Pricing* (Working paper). <https://dx.doi.org/10.2139/ssrn.6479701>.
- Waltman, L., & Kaymak, U. (2008). Q-learning agents in a Cournot oligopoly model. *Journal of Economic Dynamics & Control*, 32(10), 3275–3293. <https://doi.org/10.1016/j.jedc.2008.01.003>.
- Watkins, C. J. C. H., & Dayan, P. (1992). Q-learning. *Machine Learning*, 8(3–4), 279–292. <https://doi.org/10.1007/BF00992698>.

Appendix


This appendix presents Figures A1 and A2, showing steady-state price and profit trajectories for the symmetric one-period memory configuration (Scenario A).

Figure A1. Price trajectories for Scenario A (1 vs. 1, symmetric one-period memory)



Source: author's calculations.

Figure A2. Profit trajectories for Scenario A (1 vs. 1, symmetric one-period memory)



Source: author's calculations.

The impact of the COVID-19 shock on response distributions in business tendency surveys: An application of the Kullback–Leibler divergence

Emilia Tomczyk^a

Abstract. This paper applies the Kullback-Leibler divergence (KLD) to quantify discrepancies between expectations and subsequent assessments in business tendency surveys for Polish manufacturing firms. Using monthly survey data from May 2004 to January 2026, it examines how closely firms' expectations align with subsequent assessments across key economic variables and how this relationship evolves during the COVID-19 pandemic. The results indicate that, under normal conditions, expectations and assessments remain closely aligned, suggesting a high degree of consistency. However, during the pandemic, the KLD increased across all variables, with statistically significant rises observed for some of them, indicating a widening gap between expectations and assessments. The strongest effects are found for aggregate economic indicators, particularly the general economic situation and, to a lesser extent, firms' financial standing. On the other hand, employment and inventories exhibit relatively greater stability. Overall, the results demonstrate that the KLD provides a useful distribution-based measure for identifying regime shifts in expectation formation and for assessing the information content of survey data under changing economic conditions.

Keywords: Kullback–Leibler divergence (KLD), business tendency surveys (BTS), survey data, expectations, COVID-19

JEL: C10, D84, E32

1. Introduction

Expectations play a central role in economic decision-making and constitute a key component of business tendency surveys (BTS). These surveys are a valuable tool for monitoring economic dynamics, as they provide timely and forward-looking information about firms' perceptions of current and future economic conditions. Traditionally, the analysis of such data has relied on balance statistics or aggregate indicators, which summarise the direction of change but provide limited insight into the underlying distribution of responses.

This paper adopts an alternative perspective by focusing on the full distribution of aggregated survey responses and applying the Kullback-Leibler divergence (KLD) to quantify shifts in response distributions in BTS. KLD is used as a measure of dissimilarity between expected and reported outcomes, capturing the distance

^a SGH Warsaw School of Economics, Collegium of Economic Analysis, Institute of Econometrics, Applied Econometrics Unit, ul. Madalińskiego 6/8, 02-513 Warszawa, Poland, e-mail: emilia.tomczyk@sgh.waw.pl, ORCID: <https://orcid.org/0000-0002-4565-0352>.

between expectations and subsequent assessments. Rather than directly measuring uncertainty, KLD provides indirect evidence on the difficulty of forming expectations that are supported by actual developments. This is achieved by quantifying the mismatch between expectations and later assessments.

The empirical analysis is based on monthly BTS in manufacturing conducted by the Research Institute for Economic Development (RIED) at the SGH Warsaw School of Economics. The dataset covers the period from May 2004 to January 2026, enabling a comparison between normal economic conditions and the disruptions associated with the COVID-19 pandemic. By examining changes in KLD for several economic variables across time, the study aims to identify patterns in the divergence between expectations and assessments and to examine how this relationship changes during periods of economic instability.

This paper contributes to the literature in three ways. First, it presents a distribution-based framework for analysing expectation-assessment divergence in BTS data, which is a topic that has received limited attention in the literature. Second, it provides empirical evidence illustrating how this divergence responds to systemic shocks. Third, it identifies heterogeneity across economic variables, showing that expectations linked to variables related to the overall economic sentiment and firms' financial standing are particularly sensitive to external shocks.

2. Literature review

The existing literature contains relatively few empirical studies that apply KLD (or any related information-theoretic measures) to analyse probability distributions of expectations and assessments in BTS. Instead, most quantitative analyses of BTS data focus on the statistical properties of balance statistics, the determinants of expectation accuracy and the usefulness of aggregated survey indicators in econometric modelling. In addition, Bayesian approaches are often used to combine survey data with prior information and to account for measurement error and noise.

On the other hand, KLD is typically used in financial econometrics, entropy-based risk modelling and instability detection in time series. To provide two recent examples, Ishizaki and Inoue (2024) employed short-term KLD to identify instability in high-frequency financial time series, while Fan et al. (2025) used the same measure to evaluate the performance of alternative methods for detecting anomalies in financial time series. Applications in Poland include Wędrowska (2011), who employed KLD to study the dynamics of household expenditure structures, and Pawelec (2021), who used a related concept of structural distance to analyse the similarities and differences in expenditure patterns across different household demographic types.

Only a few previous studies employed entropy-based measures, including KLD to evaluate the information content of responses in BTS. In two separate studies, Tomczyk applied entropy and dissimilarity measures to assess the information content of Polish BTS data across four variables: two ownership sectors and four firm-size groups (2011), and to examine how the COVID-19 pandemic influenced the dynamics of expectations and assessments reported in business tendency surveys of Polish manufacturing firms (2023). In addition, Tomczyk and Kowalczyk (2023) used entropy-based methods to measure the degree of consensus among survey respondents.

Despite the limited use of KLD in BTS, this approach relates closely to the established literature on survey expectations, encompassing analyses of disagreement, forecast errors and alternative measures of uncertainty. In the BTS literature, disagreement is typically proxied by cross-sectional dispersion of qualitative responses and interpreted as heterogeneity in firms' expectations (Lahiri & Sheng, 2010; Mokinski et al., 2015). However, such measures capture only selected moments and may not fully reflect the distributional structure of three-category responses, nor distinguish clearly between disagreement and uncertainty (Rich & Tracy, 2021).

A related strand of the literature evaluates expectations expressed in BTS through forecast errors, defined as differences between *ex ante* expectations and *ex post* outcomes. Empirical evidence shows that these errors increase with subjective uncertainty and become more pronounced during periods of volatility (Buchholz et al., 2021; Morikawa, 2023). However, the categorical nature of BTS data limits the applicability of standard point-based error measures.

The present study contributes to a broader body of research on BTS in Poland, which has demonstrated that qualitative survey data provide valuable information for tracking business-cycle developments, forecasting macroeconomic variables and estimating output gaps (Białowolski et al., 2010, 2014; Ratuszny & Walczyk, 2022). Other studies examined the information content of survey responses and the mechanisms underlying the formation of expectations among firms (Białowolski, 2013; Kalinowski & Kokocińska, 2015). This research clearly demonstrates the usefulness of BTS as a source of economic information. However, they focus on forecasting performance, composite indicators or structural relationships between expectations and economic outcomes. Much less is known about the extent to which firms' expectations are subsequently reflected in their assessments, especially during periods of heightened uncertainty and economic turbulence. Understanding this relationship is important, because major shocks may alter not only economic conditions but also the ability of firms to anticipate them.

In Poland, these issues have been explored within a broader research tradition devoted to BTS. An important centre of this kind of research is the RIED, where an

extensive body of work has been devoted to the analysis of survey data. It encompasses a wide range of topics: the formation of firms' expectations, the relationship between assessments and forecasts, the forecasting performance of qualitative survey data and methodological developments, including distribution-based approaches (Bernardelli & Dędyś, 2017). The present study particularly extends the literature on distribution-based methods by demonstrating the practical application of KLD, a well-established information-theoretic measure that has so far been only rarely used in the analysis of survey response distributions.

Polish science has also addressed the formation and rationality of economic expectations. A prominent example is the extensive body of research on inflation expectations conducted within the National Bank of Poland. Among these contributions, Łyziak (2016) examined the behaviour and information properties of inflation expectations in Poland, while also relating them to issues of expectation accuracy and macroeconomic performance. Kowalczyk and Tomczyk (2008) examined the rationality of firms' expectations, focusing on unbiasedness and information efficiency. More recently, Ziemińska (2021) explored alternative approaches to testing expectation-formation mechanisms and the rationality of economic forecasts. Altogether, these studies have contributed to a wider understanding of how economic expectations are formed, revised and evaluated. However, they focus primarily on the properties of expectations, typically rationality, rather than on the degree to which expectations and subsequent assessments reported in BTS correspond with each other.

A related, though still somewhat limited strand of the literature, examines survey-based expectations from the perspective of uncertainty and consensus among respondents (Dudek, 2017; Tomczyk & Kowalczyk, 2023). This line of research includes both traditional uncertainty measures derived from survey and financial market data, as well as more recent approaches based on text-mining and web-scraping techniques (Hołda, 2019). Lately, Kocięcki et al. (2022) analysed subjective expectations and uncertainty, emphasising the role of survey-based measures in capturing the agents' perceptions of future economic developments. While KLD is not a direct measure of uncertainty, its elevated values indicate a growing divergence between expectations and subsequent assessments. The present study contributes to the literature by focusing on the distribution of survey responses, represented by the probabilities of the qualitative response categories.

This paper differs from earlier Polish studies in three aspects. First, rather than examining information content, respondent consensus or forecasting performance, it focuses explicitly on the divergence between expectations and assessments. Second, the analysis is conducted at the level of full response distributions available from a qualitative tendency survey rather than balance statistics or selected summary

measures, capturing the changes in the entire structure of responses. Third, the study draws on a long monthly sample covering different phases of the business cycle, making it possible to examine how the divergence between expectation and assessment evolves across economic regimes, including the COVID-19 shock. Thus, the paper extends the previous applications of entropy-based methods in Polish BTS and demonstrates the usefulness of KLD as a tool for identifying periods in which firms' expectations become less consistent with the reported subsequent assessments.

3. Methods and data

Originally introduced by Kullback and Leibler (1951), KLD has been extensively studied in the literature, with its properties and interpretations discussed in popular textbooks such as those by Cover and Thomas (2006), and Joyce (2011). The standard definition of KLD is as follows:

$$KLD(S_q^n, S_p^n) = \sum_{i=1}^n q_i \log \frac{q_i}{p_i}, \quad (1)$$

where S_p^n stands for *a priori* structure, S_q^n for *a posteriori* structure, and structure S^n itself is defined as a vector: $S^n = [s_1, s_2, \dots, s_n]^T \in \mathbb{R}^n$, whose elements s_i ($i = 1, 2, \dots, n$) fulfil two conditions:

$$0 \leq s_i \leq 1, \quad (2)$$

$$\sum_{i=1}^n s_i = 1. \quad (3)$$

Structure S^n is therefore fully described by a vector of fractions (structure elements) summing to a total of 1. In the context of BTS, *a priori* structure S_p^n reflects the distribution of expectations (typically a discrete distribution over three or five response categories), whereas *a posteriori* structure S_q^n describes the distribution of the assessment (i.e. the observed changes).

One of the key properties of (1) is that it attains its minimum value of zero if and only if the two structures are identical (i.e. $S_p^n = S_q^n$) and it increases to infinity (as it is not bounded from above) with the magnitude of the differences between the structures. Accordingly, (1) can be interpreted as a measure of the divergence between the *a priori* (expected) and *a posteriori* (observed) probability distributions. In this sense, it serves as a measure of structural dissimilarity: the larger its value, the greater

the disparity between the two structures, i.e. the distributions of respondents' answers to questions about their expectations and their subsequent assessments.

It is important to note that KLD does not measure uncertainty *per se*, but rather the divergence between the expected and realised distributions. Higher values of KLD indicate a greater mismatch between expectations and later assessments. While such discrepancies may be associated with heightened uncertainty or increased forecasting difficulty, KLD should not be interpreted as a direct measure of uncertainty. Instead, it is best viewed as a distribution-based indicator of the degree to which the observed outcomes deviate from the previously reported expectations.

KLD was selected over alternative measures such as the Jensen-Shannon divergence, Hellinger distance and total variation distance for both theoretical and substantive reasons. First, KLD has a direct information-theoretic interpretation as the expected information loss when the forecast distribution is used instead of the realised distribution. Second, unlike purely metric measures, KLD explicitly preserves the directional distinction between expectations (*a priori* distributions) and subsequent assessments (*a posteriori* distributions), which is central to the objective of this study. Third, previous comparative reviews of probability-distribution distances identify KLD as one of the most widely used information-theoretic measures for analysing discrepancies between categorical probability distributions and expectation-related processes (Cha, 2007).

The data on assessments and expectations concerning major economic variables have been obtained from the monthly BTS in manufacturing conducted by RIED at the SGH Warsaw School of Economics. The scope of the survey and variants of the answers are presented in Table A in the Appendix. Eight fields of economic activity are evaluated by respondents with respect to the changes observed in the previous month and expectations for the next three to four months. To ensure consistency between the expectations and the subsequent assessments, expectations reported in month t are matched with assessments reported three months later ($t + 3$). This is reflected in the survey design, where the respondents formulate expectations for the next three to four months. Following previous analyses of RIED BTS data (Tomczyk & Kowalczyk, 2023), a three-month period was adopted as the effective expectation horizon.

It should be noted that the analysis compares aggregated distributions of survey responses rather than expectations and the subsequent outcomes reported by the same firms. Consequently, the KLD should be interpreted as a measure of dissimilarity between aggregate expectations and assessment distributions. In addition, assessment responses represent the firms' subsequent subjective evaluations rather than objectively observed outcomes. The measure therefore captures the extent to which

the later assessments differ from the previously reported expectations, rather than forecasting performance in the conventional sense of comparing expectations with outcomes recorded in official statistics.

4. Results of the empirical analysis

This section presents the results of the empirical analysis, examining the distributional behaviour of the KLD across variables and over time. Table 1 presents the basic descriptive statistics of the KLD for the full sample.

Table 1. Basic KLD descriptive statistics for the full sample (May 2004 – January 2026)

Variable	Minimum	Maximum	Mean
Level of production	0.000169	0.428026	0.039396
Level of orders	0.000147	0.407289	0.041480
Level of export orders	0.000048	0.389705	0.038750
Stocks of finished goods	0.000247	0.176382	0.030195
Prices of goods produced	0.000029	0.273608	0.024966
Level of employment	0.000025	0.125837	0.021060
Financial standing	0.000012	0.302973	0.024777
General situation of the economy	0.000185	0.444007	0.045507

Source: author's work based on the RIED database.

Across the full sample, the average KLD values were relatively low, ranging between 0.02 and 0.046, indicating that expectations were, on average, fairly consistent with the subsequent assessments. Minimum values close to zero for all variables confirm that there were many periods when the expectations and the observed outcomes were almost perfectly aligned, which is typical under stable macroeconomic conditions.

However, important differences emerged across the variables. The highest average KLD values were observed for the general economic situation, as well as orders and production. This indicates that firms face considerable difficulties in forming accurate expectations about the overall economy, as it is strongly influenced by external shocks, policy decisions and global developments. Relatively high values for orders and production place these variables in the same category as the general economic situation.

The lowest average KLD values were observed for employment, prices and financial standing. These variables tended to adjust more slowly and were often subject to institutional rigidities (e.g. labour contracts, pricing strategies and financing conditions). As a result, the distribution of the expectations in these domains was relatively well reflected in the distribution of the assessments.

The presence of high maximum KLD values indicates episodic shocks, which are likely associated with major crises. Indeed, most of the maximum values were observed at the onset of the pandemic (for export orders, prices and the general economic situation, the maximum was reached in April 2020, for stocks in May 2020 and for financial standing in July 2020). This suggests that at the beginning of this unprecedented crisis, expectations based on historical patterns failed. Table 2 presents the maximum, minimum and mean values of KLD for the pandemic subsample.

Table 2. Basic KLD descriptive statistics for the pandemic sample (March 2020 – March 2022)

Variable	Minimum	Maximum	Mean
Level of production	0.000507	0.311640	0.054086
Level of orders	0.001220	0.399952	0.058409
Level of export orders	0.004233	0.389705	0.063080
Stocks of finished goods	0.001798	0.176382	0.034339
Prices of goods produced	0.000029	0.273608	0.052187
Level of employment	0.001126	0.119150	0.026147
Financial standing	0.003217	0.302973	0.056791
General situation of the economy	0.007801	0.444007	0.100685

Source: author's work based on the RIED database.


The key result is that the average KLD increases for all variables during the pandemic, reflecting a substantial increase in the divergence between expectations and assessments. The pandemic period was characterised by radical uncertainty (e.g. lockdowns, supply-chain disruptions and policy interventions) and rapid regime shifts. As a result, firms' expectations lost their predictive power and were frequently revised. The KLD captures this as a persistent divergence between expectations and assessments.

The variable most affected is the general economic situation. The average KLD during the pandemic more than doubled relative to its long-run level, indicating that firms were unable to form coherent expectations about the macroeconomic environment. This variable reflects uncertainty related to health policies, fiscal measures, international developments and demand conditions. A similarly strong effect was observed for prices, where the pandemic average KLD also more than doubled compared to the full sample. This suggests a breakdown in price-setting expectations relative to the subsequent assessments: shortages, rising energy prices and transport costs, and policy interventions weakened firms' ability to anticipate price developments. Export orders also exhibited an increase in KLD, indicating that border closures and highly volatile foreign demand made external conditions extremely difficult to predict.

By contrast, the labour market and inventories exhibited a relative resilience during periods of economic turbulence. Although the KLD increased for employment and stocks, the rise was more moderate compared to demand and macroeconomic variables. This may be explained by labour hoarding, government subsidies and job-protection schemes that stabilised employment dynamics. Inventories, in turn, reflected adjustment margins which, at least to some extent, responded automatically to the disruptions. Consequently, expectation formation in these dimensions remained relatively more stable, even under extreme conditions.

The Figure below shows the evolution of the KLD for two variables: the general economic situation and employment. Higher values indicate a greater divergence between firms' expectations and their subsequent assessments.

Figure. Monthly KLD between expectations and subsequent assessments in BTS (May 2004 – January 2026)



Source: author's calculations based on the RIED database.

These two variables were selected because they represent distinct patterns in the data: a strong pandemic effect (general economic situation) and relative stability (employment). Together, they provide a concise illustration of the heterogeneous impact of the COVID-19 shock on the divergence between expectations and assessments.

It should be noted that throughout the sample period, the KLD for the general economic situation was considerably more volatile than for employment. While the most pronounced increase occurred during the COVID-19 period, elevated KLD values were also observed during the global financial crisis in 2008–2010, indicating that firms' expectations regarding broad macroeconomic conditions were particularly sensitive to major economic disturbances. A temporary spike in divergence also occurred around 2023, suggesting that the post-pandemic period did not mean a complete return to stability. To verify whether these patterns are statistically significant, differences in mean KLD values across the sub-periods are assessed using Welch's two-sample *t*-test. To avoid sample overlap and ensure a meaningful comparison, the analysis distinguishes between the pre-pandemic period (May 2004–February 2020, number of observations $N_1 = 190$), the pandemic period (March 2020–March 2022, $N_2 = 25$), and the post-pandemic period (April 2022–January 2026, $N_3 = 46$). A comparison between the full sample and the pandemic subsample (as presented in Tables 1 and 2) would be misleading, as the latter is part of the former and therefore influences its average.

Welch's *t*-test has been selected because the assumption of equal variances was unlikely to hold, as indicated by the differences in the standard deviations across the studied periods. Unlike the classical Student's *t*-test, Welch's test does not impose the assumption of homoscedasticity and is therefore more robust in the presence of unequal variances and sample sizes. The analysis compares non-overlapping time periods, allowing for a meaningful comparison of the mean values across the subperiods. However, since the observations form a time series, they may exhibit serial correlation. As a result, the effective sample size may be smaller than the nominal number of observations due to the lack of independence among observations, potentially affecting the accuracy of the standard errors and *p*-values. The results of Welch's *t*-test should be regarded as approximate and interpreted primarily as descriptive evidence of the differences across subperiods.

Table 3 presents the results of Welch's two-sample *t*-test for equality of means, presenting the *t*-statistics and two-sided *p*-values. The values in bold denote statistical significance at a 0.10 level. Negative *t*-values in the first comparison indicate higher mean KLD during the pandemic relative to the pre-pandemic period.

Table 3. Results of Welch's two-sample *t*-test for equality of means

Variable	Pre-pandemic vs pandemic (<i>t</i>)	<i>p</i> -value	Pandemic vs post- pandemic (<i>t</i>)	<i>p</i> -value
Level of production	-0.793	0.435	1.111	0.276
Level of orders	-0.607	0.549	1.511	0.143
Level of export orders	-1.469	0.153	1.085	0.286
Stocks of finished goods	-0.367	0.716	1.459	0.154
Prices of goods produced	-2.297	0.030	1.010	0.321
Level of employment	-0.785	0.437	0.950	0.348
Financial standing	-2.298	0.029	2.089	0.046
General situation of the economy	-2.557	0.016	1.704	0.099

Source: author's calculations based on the RIED database.

The results of the Welch *t*-tests demonstrate that the increase in the KLD observed during the pandemic period is statistically significant only for a subset of variables. In particular, the null hypothesis of equal means between the pre-pandemic and pandemic periods is rejected for variables related to prices ($t = -2.297$, $p = 0.030$), financial standing ($t = -2.298$, $p = 0.029$) and the general economic situation ($t = -2.557$, $p = 0.016$). These findings suggest that the similarity between the distributions of the expectations and assessments decreased significantly in areas most closely linked to macroeconomic conditions and financial performance.

For the remaining variables – production, orders, export orders, stocks and employment – the differences in mean KLD values between the pre-pandemic and pandemic periods are not statistically significant. Although the average divergence increased for several of these indicators, the changes are not sufficiently large relative to their variability to be distinguished from random fluctuations.

Comparing the pandemic and post-pandemic periods, the results indicate a general weakening of the effect. Only financial standing ($t = 2.089$, $p = 0.046$) remained statistically significant, indicating a persistent mismatch between the expectations and the subsequent assessments regarding firms' financial conditions. For the general economic situation, the difference remained elevated, but it is only marginally significant at a 10% significance level ($t = 1.704$, $p = 0.099$). For all other variables, no statistically significant differences were detected.

Given the time-series nature of the data, the results of the *t*-tests should be interpreted with caution, as serial correlation may affect the estimated standard errors. To address this issue, the mean equality was additionally evaluated using the Newey-West heteroskedasticity and autocorrelation consistent (HAC) framework, which adjusts standard errors for both heteroskedasticity and autocorrelation. The following results are based on comparisons of the mean KLD values across the subsamples.

Statistical significance was assessed using Newey-West HAC standard errors with a lag length of 12, which is commonly adopted for monthly data. The results are shown in Table 4. The values in bold denote statistical significance at a 0.10 level. The negative t -values in the first comparison indicate a higher mean KLD during the pandemic relative to the pre-pandemic period.

Table 4. Results of the Newey-West HAC test for equality of means

Variable	Pre-pandemic vs pandemic (t)	p -value	Pandemic vs post- pandemic (t)	p -value
Level of production	-0.490	0.624	0.700	0.484
Level of orders	-0.426	0.670	1.007	0.314
Level of export orders	-1.030	0.303	0.794	0.427
Stocks of finished goods	-0.197	0.844	0.988	0.323
Prices of goods produced	-1.356	0.175	0.617	0.537
Level of employment	-0.548	0.584	0.602	0.547
Financial standing	-1.745	0.081	1.626	0.104
General situation of the economy	-1.986	0.047	1.335	0.182

Source: author's calculations based on the RIED database.

Financial standing ($t = -1.745$, $p = 0.081$) and general situation of the economy ($t = -1.986$, $p = 0.047$) were the only variables for which the increase in mean KLD during the pandemic remained statistically significant after correcting for heteroskedasticity and autocorrelation at a 0.1 significance level. By contrast, comparisons between the pandemic and post-pandemic periods revealed no statistically significant differences in mean KLD values for any of the analysed variables.

The results of the Newey-West HAC tests largely confirm the main conclusions drawn from Welch's t -tests. The most pronounced and statistically significant increases in KLD during the pandemic were observed for the financial standing and general situation of the economy, indicating that these variables were most strongly affected by the disruption caused by the COVID-19 shock. However, the effect observed for prices in the Welch t -test disappeared once HAC standard errors were employed.

Overall, the two sets of results indicate that the pandemic-related increase in divergence was concentrated in a limited number of variables, primarily those linked to the firms' financial conditions and overall economic sentiment. The more conservative HAC estimates suggest that part of the statistical significance detected by Welch's tests may have been driven by serial correlation in the monthly data; however, this does not alter the substantive conclusions.

On the whole, the results suggest that the deterioration in the alignment between the expectations and assessments during the pandemic was not universal but restricted to certain domains, particularly those associated with macroeconomic conditions and the financial situations of the studied firms. Moreover, the absence of statistically significant differences between the pandemic and post-pandemic periods points to a gradual normalisation of these relationships, although some effects may still have persisted.

5. Discussion

The results obtained using the KLD are partially consistent with the related research. Lahmiri and Bekiros (2020) empirically demonstrated that entropy-based measures are not only effective in capturing increases in uncertainty, nonlinearities and structural breaks in financial markets during the COVID-19 crisis, but they also outperform traditional econometric tools in detecting early-stage instability. Although their study does not rely directly on KLD, it provides important empirical evidence supporting the usefulness of entropy-based measures in analysing uncertainty and expectation dynamics.

Other studies, employing different methodologies, also confirm rising uncertainty and growing difficulties in forming accurate predictions. Castelnovo (2022) showed a systematic survey of the role of uncertainty in business cycles, with particular emphasis on how it surged and propagated during the COVID-19 pandemic, and how this surge compares with previous crises. He found that COVID-19 generated an unprecedented increase in uncertainty and led to greater disagreement among firms and forecasters. Furthermore, the pandemic triggered frequent revisions of expectations and showed a high degree of persistence. The KLD-based results confirmed these findings in terms of the similarity between the distributions of the expectations and subsequent assessments.

Meyer et al. (2022) investigated the nature, magnitude and persistence of uncertainty during the COVID-19 pandemic, particularly focusing on forward-looking uncertainty and firm-level expectations. They showed that COVID-19 caused an extraordinary increase in uncertainty, although short- and long-horizon expectations evolved differently. Moreover, they characterised COVID-19 as a prolonged expectations shock rather than merely a volatility episode. The elevated KLD observed during the pandemic is consistent with their evidence of a breakdown in forecast reliability and structural changes in the information environment faced by firms.

The findings also contribute to the growing Polish literature on BTS. While earlier studies focused primarily on forecasting performance, expectation formation,

information content and respondent consensus, the present analysis highlights a different dimension: the extent to which expectations are subsequently reflected in firms' assessments. The results suggest that this relationship becomes substantially less stable during periods of major economic disruption.

In conclusion, these results suggest that entropy-based measures effectively complement traditional indicators of uncertainty and disagreement, particularly in the context of qualitative survey data.

6. Conclusions

The results indicate that the relationship between expectations and assessments is state-dependent and the effects are stronger and statistically significant for some variables, while weaker for others. Given the relative nature of KLD, expectations and subsequent assessments remain fairly closely aligned in normal conditions, whereas during systemic crises their consistency deteriorates with varying effects across all variables. This pattern, particularly evident during the COVID-19 period, suggests that BTS capture changes in the degree of similarity between the distributions of expectations and subsequent assessments. KLD proves especially useful as a distribution-based measure, revealing changes in belief formation beyond aggregate indicators such as balances or means.

Higher KLD values indicate greater divergence between expectations and subsequent assessments. One plausible interpretation here is that the elevated divergence reflects increased uncertainty, forecasting difficulties or changes in the information environment faced by firms. In descriptive terms, the highest divergence is observed for variables related to demand and macroeconomic conditions, i.e. production, orders, export orders and the general economic situation, indicating that these expectations are the most sensitive to shocks. Statistically significant differences are found for financial standing and the general economic situation. For these variables, systemic crises should be interpreted not only as real shocks but also as expectation shocks that impair firms' ability to form expectations consistent with actual future developments.

By contrast, employment exhibits consistently low divergence throughout the sample. Financial standing and the general situation of the economy show the most pronounced increases in divergence during the pandemic, with both variables displaying statistically significant changes relative to the pre-pandemic period in the Newey-West HAC framework. On the other hand, prices, although characterised by a noticeable increase in average divergence during the pandemic, do not remain statistically significant once serial correlation and heteroskedasticity are taken into account. The smallest effects, both in magnitude and statistical significance, are

observed for employment and stocks, which are firm-level operational indicators and therefore appear to reflect firm-specific rather than aggregate economic conditions.

It is worth noting that the KLD employed in this study measures divergence between aggregate distributions of survey expectations and subsequent survey assessments. It does not compare the expectations with the objectively observed outcomes, nor does it track revisions made by the same firms over time. Accordingly, the results should be interpreted as evidence on the consistency of aggregate survey responses rather than on individual forecasting accuracy.

These findings have important implications for both measurement and policy. From a methodological perspective, they highlight the advantages of distribution-based approaches such as KLD, in capturing shifts in expectations and assessments that remain hidden in conventional summary measures. They also highlight the importance of analysing the full distribution of survey responses (at least as captured by the probabilities of qualitative response categories, when more detailed data are unavailable), rather than relying solely on aggregate indicators such as balances.

From a policy perspective, elevated KLD values suggest a weaker correspondence between expectations and subsequent assessments during periods of economic disruption. In this context, effective communication and stabilisation policies may contribute to a closer alignment between expectations and subsequent assessments. This observation is particularly relevant for financial standing and the general economic situation, for which statistically significant increases in KLD are observed.

In light of the above, several avenues for further research emerge. First, the analysis could be extended to other countries and sectors to assess the generality of the observed patterns. Such extensions may be constrained by data availability, as detailed distributional survey responses (shares of positive, neutral and negative answers) are not always published.

Second, KLD could be complemented with other entropy-based or information-theoretic measures to provide a more comprehensive evaluation of the divergence between expectations and subsequent assessments. In particular, future research could examine whether the substantive conclusions are robust to alternative measures of distributional divergence. Such comparisons could help identify the extent to which empirical findings depend on the specific properties of KLD, including its asymmetry, and provide further insights into the advantages and limitations of alternative information-theoretic approaches in the analysis of BTS data.

Third, future studies could use micro-level survey data to examine heterogeneity in expectations across firms, e.g. by size, ownership structure or international exposure.

Fourth, the pandemic period is analysed separately in this study because descriptive evidence indicates that it differs substantially from the rest of the sample, with the highest KLD values for most of the variables observed during this interval.

The post-pandemic period, by contrast, is characterised by several overlapping shocks, including the energy crisis, elevated inflation and the consequences of Russia's invasion of Ukraine, making a clean separation of individual effects difficult. Given the limited number of observations available after 2022, identifying the effects of individual post-pandemic shocks remains a promising avenue for future research.

Finally, the dynamic properties of KLD warrant further investigation in a time-series framework, including its potential role as an early-warning indicator of regime shifts and its relationship with macroeconomic activity. Further research could also explore how policy interventions affect expectation formation and whether they help stabilise the link between expectations and subsequent assessments.

References

- Bernardelli, M., & Dędyś, M. (2017). Mapping the respondents' assessments in business tendency survey using the Viterbi paths. *Prace i Materiały Instytutu Rozwoju Gospodarczego SGH*, (101), 27–43. <https://econjournals.sgh.waw.pl/PiM/article/view/3973/3538>.
- Białowolski, P. (2013). Expectations' Formation in Business Survey Data. *Prace i Materiały Instytutu Rozwoju Gospodarczego SGH*, (93), 5–28. <https://econjournals.sgh.waw.pl/PiM/article/view/3905/3471>.
- Białowolski, P., Kuszewski, T., & Witkowski, B. (2010). Prognozy kombinowane wskaźników makroekonomicznych z użyciem danych z testów koniunktury. *Contemporary Economics*, 4(4), 41–58.
- Białowolski, P., Kuszewski, T., & Witkowski, B. (2014). Bayesian averaging of classical estimates in forecasting macroeconomic indicators with application of business survey data. *Empirica*, 41(1), 53–68. <https://doi.org/10.1007/s10663-013-9227-x>.
- Buchholz, M., Tonzer, L., & Berner, J. (2021). Firm-specific forecast errors and asymmetric investment propensity. *Economic Inquiry*, 60(2), 764–793. <https://doi.org/10.1111/ecin.13045>.
- Castelnuovo, E. (2022). Uncertainty before and during COVID-19: A survey. *Journal of Economic Surveys*, 37(3), 821–864. <https://doi.org/10.1111/joes.12515>.
- Cha, S.-H. (2007). Comprehensive survey on distance/similarity measures between probability density functions. *International Journal of Mathematical Models and Methods in Applied Sciences*, 1(4), 300–307.
- Cover, T. M., & Thomas, J. A. (2006). *Elements of Information Theory* (2nd ed.). John Wiley & Sons. <https://doi.org/10.1002/047174882X>.
- Dudek, S. (2017). Niepewność w opiniach polskich gospodarstw domowych. *Prace i Materiały Instytutu Rozwoju Gospodarczego SGH*, (101), 55–91. <https://econjournals.sgh.waw.pl/PiM/article/view/3976/3540>.
- Fan, Y., Yang, Y., Wang, Z., & Gao, M. (2025). Instability of Financial Time Series Revealed by Irreversibility Analysis. *Entropy*, 27(4), 1–22. <https://doi.org/10.3390/e27040402>.
- Hołda, M. (2019). *Newspaper-based economic uncertainty indices for Poland* (NBP Working Paper No. 310). https://static.nbp.pl/publikacje/materiały-i-studia/310_en.pdf.

- Ishizaki, R., & Inoue, M. (2024). Short-term Kullback–Leibler divergence analysis to extract unstable periods in financial time series. *Evolutionary and Institutional Economics Review*, 21, 227–236. <https://doi.org/10.1007/s40844-024-00284-0>.
- Joyce, J. M. (2011). Kullback–Leibler divergence. In M. Lovrić (Ed.), *International Encyclopedia of Statistical Science* (pp. 720–722). https://doi.org/10.1007/978-3-642-04898-2_327.
- Kalinowski, S., & Kokocińska, M. (2015). Business tendency survey data. Where do the respondents' opinions come from?. *Economics and Business Review*, 1(4), 69–83. <https://doi.org/10.18559/ebr.2015.4.5>.
- Kocięcki, A., Łyziak, T., & Stanisławska, E. (2022). *Subjective Expectations and Uncertainty* (NBP Working Papers No. 345). https://static.nbp.pl/publikacje/materialy-i-studia/345_en.pdf.
- Kowalczyk, B., & Tomczyk, E. (2008). Rationality of Expectations of Industrial Enterprises – Analysis Based on Business Tendency Surveys with Item Non-respons. *Bank i Kredyt*, 39(8), 3–11.
- Kullback, S., & Leibler, R. A. (1951). On information and sufficiency. *The Annals of Mathematical Statistics*, 22(1), 79–86. <https://doi.org/10.1214/aoms/1177729694>.
- Lahiri, K., & Sheng, X. (2010). Measuring forecast uncertainty by disagreement: The missing link. *Journal of Applied Econometrics*, 25(4), 514–538. <https://doi.org/10.1002/jae.1167>.
- Lahmiri, S., & Bekiros, S. (2020). The impact of COVID-19 pandemic upon stability and sequential irregularity of equity and cryptocurrency markets. *Chaos, Solitons & Fractals*, 138, 1–6. <https://doi.org/10.1016/j.chaos.2020.109936>.
- Łyziak, T. (2016). Financial crisis, low inflation environment and short-term inflation expectations in Poland. *Bank i Kredyt*, 47(3), 285–300.
- Meyer, B. H., Mihaylov, E., Barrero, J. M., Davis, S. J., Altig, D., & Bloom, N. (2022). *Pandemic-Era uncertainty* (NBER Working Paper No. 29958). <https://doi.org/10.3386/w29958>.
- Mokinski, F., Sheng, X., & Yang, J. (2015). Measuring Disagreement in Qualitative Expectations. *Journal of Forecasting*, 34(5), 405–426. <https://doi.org/10.1002/for.2340>.
- Morikawa, M. (2023). Firms' subjective uncertainty and forecast errors: Survey evidence from Japan. *Applied Economics Letters*, 30(1), 33–36. <https://doi.org/10.1080/13504851.2021.1971611>.
- Pawełec, A. (2021). Analiza struktury wydatków konsumpcyjnych polskich gospodarstw domowych według typu biologicznego. *Metody Ilościowe w Badaniach Ekonomicznych*, 22(3), 87–96. <https://doi.org/10.22630/MIBE.2021.22.3.8>.
- Ratuszny, E., & Walczyk, K. (2022). Use of data from business tendency surveys in output gap estimation. In S. Białowas (Ed.), *Economic tendency surveys and economic policy – measuring output gaps and growth potentials* (pp. 30–38). Poznań University of Economics and Business. <https://doi.org/10.18559/978-83-8211-129-3/2>.
- Rich, R., & Tracy, J. (2021). A closer look at the behavior of uncertainty and disagreement: Micro evidence from the euro area. *Journal of Money, Credit and Banking*, 53(1), 233–253. <https://doi.org/10.1111/jmcb.12728>.
- Tomczyk, E. (2011). Application of measures of entropy, information content and dissimilarity of structures to business tendency survey data. *Przegląd Statystyczny. Statistical Review*, 58(1–2), 88–101. <https://doi.org/10.59139/ps.2011.01-02.6>.

- Tomczyk, E. (2023). Dynamics of survey responses before and during the pandemic: entropy and dissimilarity measures applied to business tendency survey data. *Statistics in Transition new series*, 24(2), 185–199. <https://doi.org/10.59170/stattrans-2023-027>.
- Tomczyk, E., & Kowalczyk, B. (2023). Consensus in Business Tendency Surveys: Comparison of Alternative Measures. *Gospodarka Narodowa. The Polish Journal of Economics*, (4), 17–29. <https://doi.org/10.33119/GN/174277>.
- Wędrowska, E. (2011). Application of Kullback–Leibler relative entropy for studies on the divergence of household expenditures structures. *Olsztyn Economic Journal*, 6(1), 133–142.
- Ziemińska, P. (2021). Quality of Tests of Expectation Formation for Revised Data. *Central European Journal of Economic Modelling and Econometrics*, 13(4), 405–453. <https://doi.org/10.24425/cejeme.2021.139836>.

Appendix

Table A. Monthly RIED questionnaire in the manufacturing industry

Variable	Observed within the last month	Expected for the next 3–4 months
Level of production	up / unchanged / down	will increase / will remain unchanged / will decrease
Level of orders	up / unchanged / down	will increase / will remain unchanged / will decrease
Level of export orders	up / unchanged / down / not applicable ^b	will increase / will remain unchanged / will decrease / not applicable ^a
Stocks of finished goods	up / unchanged / down	will increase / will remain unchanged / will decrease
Prices of goods produced	up / unchanged / down	will increase / will remain unchanged / will decrease
Level of employment	up / unchanged / down	will increase / will remain unchanged / will decrease
Financial standing	up / unchanged / down	will increase / will remain unchanged / will decrease
General situation of the economy regardless of the situation in your sector and enterprise	up / unchanged / down	will increase / will remain unchanged / will decrease

a 'Not applicable' answers to the Level of export orders question are omitted from the empirical analysis.

Source: RIED database.

In memory of Professor Władysław Bukietyński (1928–1996)



Wrocław University of Economics and Business

Thirty years ago, on 30th October 1996, we paid our final respects to Professor Władysław Bukietyński, the father of the academic community of operational research analysts in Wrocław and patron of the annual National Scientific Conference ‘Methods and Applications of Operational Research’ (Pol. Ogólnopolska Konferencja Naukowa im. Profesora Władysława Bukietyńskiego Metody i Zastosowania Badań Operacyjnych – MZBO). The 44th edition of the event – MZBO 2026 – was organised by the University of Economics in Poznań, and took place two weeks prior to the anniversary of his passing.

September 2026 also marked 50 years since the establishment of the Faculty of Management, Computer Science and Finance at the University of Economics and Business in Wrocław, founded by Professor Bukietyński, who was also its first dean during the subsequent terms from 1976 to 1984. This event provided another opportunity to commemorate this outstanding teacher, scholar, originator, organiser, mentor and demanding reviewer. Professor Bukietyński was also the life and soul of social life, a humanist with a love for mathematics, a great wit and master of sharp repartee, and a friend to lean on in difficult moments.

I was fortunate to witness the multiple roles that Professor Bukietyński, also known as Unek, fulfilled. He did not allow us, namely the staff of the Department of Operational Research (initially named the Department of Mathematical

Programming), to address him as ‘Professor’. His reaction was invariable: ‘Do not call me names’.

In fact, I was one of his students. During my studies in 1963–1968 at the Faculty of Engineering and Industrial Economy of the then Higher School of Economics in Wrocław, I became devoted to probability calculus, mathematical and descriptive statistics, and mathematical programming, all of which were taught by Bukietyński, then holder of a PhD. His didactic *virtuoserie* made me continue my scholarly path in the Department of Statistics (then called the Department of Statistics and Economic Calculus), headed first by Professor Zdzisław Hellwig, and during his extended absence during his tenure in Nigeria by Dr Bukietyński.

The Department of Statistics was established in 1950, at the time when the Higher School of Economics (Pol. Wyższa Szkoła Ekonomiczna – WSE) in Wrocław was founded. The decree of the Council of Ministers of 27th September 1950, transformed the private, vocational Higher School of Commerce (operating from February 1947) into a public educational institution, entitled to grant a Master’s degree in Economics and providing free tuition.

Władysław Bukietyński, thanks to his outstanding knowledge and excellent mathematical skills (he graduated from a mathematics and physics profiled class at the Bolesław Chrobry secondary school in Kłodzko), was appointed Deputy Assistant, and later Junior Assistant in the Department of Statistics, still as a WSE student. The department was headed by Jan Falewicz, a senior lecturer, holding a Master’s degree in engineering, who taught statistics and elements of advanced mathematics. Upon receiving his Master’s degree in 1954, Bukietyński was appointed Assistant Professor at the Department of Statistics, and in 1956, in the newly established Department of Mathematics. There, in 1961, began his long-term scholarly collaboration and friendship with Professor Antoni Smoluk, also a young Master’s degree holder.

Professor Smoluk, in his book entitled *Siedmiu z Ekonometrii. Nauka, historia, anegdoty* [The seven from the Econometrics Department: Science, history, anecdotes], published in 2016 by the WUEB Publishing House, remembers his colleague from those times:

Demanding of himself, he was understanding towards others. His entire life was dedicated to the search for truth, goodness and beauty. He contributed to the creation of the Department of Mathematics at the Higher School of Economics in Wrocław in 1955, and was one of the first three staff members. Gifted with a profound scientific, mathematical mind, he felt most at home in the world of science. At the same time, he was also a great humanist who experienced poetical visions and was able to ‘paint with words’, exploring the nature of language (p. 14).

At that time, he was preparing his PhD dissertation *O niektórych teoretycznych problemach statystycznej analizy popytu* [On some theoretical problems of the statistical analysis of demand], which he defended in 1963.

In 1965, Bukietyński and Smoluk moved to the Department of Statistics, and their shared scientific passions yielded several articles, e.g. *O pewnej klasie funkcji liniowo niezależnych* [On a certain class of linearly independent functions] (1968) and *Uwagi o generatorach liczb pseudolosowych opartych na ciągach Fibonacciego* [Notes on pseudorandom number generators based on Fibonacci sequences] (1969), published in the *Prace Naukowe* [Scientific Works] WSE series. Over a decade later, they authored together an important publication, widely discussed in the economic community, titled *O tym, że równowaga jest ważniejsza niż optymalność* [On equilibrium being more important than optimality] (1982).

The school of mathematical programming officially emerged in Wrocław in the context of the University's institutes reform carried out in 1969. It entailed the establishment of five institutes comprising 29 sections, one of which was the Institute of Methods of Economic Calculus (later the Institute of Economic Cybernetics). The Institute consisted of eight sections, one of which was the Section of Mathematical Programming, managed by the then Dr Bukietyński, who was at the same time nominated Deputy Director of the Institute and served as Assistant Professor. The staff members of the Section of Mathematical Programming included: Ryszard Jasiński, MSc, Stanisław Krawczyk, MSc, Ewa Konarzewska, BEng, MSc, and Urszula Królik, BEng, MSc. Ryszard, Stanisław and myself, exactly in that order, became the first PhD students supervised by Bukietyński, studying different classes of mathematical programming methods and their application in decision-making processes. In my case, this regarded the problems of multicriteria programming. My supervisor took particular pride in the fact that my PhD dissertation from 1975, later published as a book entitled *Programowanie przy wielorakości celów* [Multi-objective programming] (Państwowe Wydawnictwo Naukowe 1980), became the first monograph dedicated to problems of vector optimisation in Poland, and one of the first in the world.

Professor Bukietyński was at that time working on transforming certain problems of optimisation into tasks of discrete programming, important from the perspective of economic practice. This was the subject of his post-doctoral dissertation *Niektóre problemy optymalizacji dyskretnej* [Selected problems of discrete optimisation] from 1975, which defined a certain general class of tasks of the dynamic selection. His paper *O pewnym zagadnieniu optymalizacyjnym* [On a certain optimisation issue] provided a simplified algorithm, allowing for an approximate solution to a generalised selection task, where the generalisation stems from incorporating the timing of subsequent acts of choice. In July 1976, then holding the office of

the Dean of the new Faculty of Management and Computer Science at the Oskar Lange Academy of Economics, Władysław Bukietyński was granted a post-doctoral degree in Economics, specialising in Econometrics, and in 1979, he was nominated Associate Professor.

Despite his many duties as the faculty dean, member of the Rector's Commission for International University Cooperation, and member of the Academic Council for Scientific Research Sections, the Professor managed to pursue his scholarly interests. He achieved this goal by running weekly science seminars and working within scientific research teams to study issues submitted by different economic entities.

The seminars in the Department of Mathematical Programming (which in 1981 became the Department of Operational Research) were conducted in an extraordinary atmosphere created by Professor Bukietyński, who greatly enjoyed the seminar format and lively debates. He was uniquely suited for team work, which yielded a monograph titled *Programowanie matematyczne* [Mathematical programming] (Wydawnictwo Akademii Ekonomicznej we Wrocławiu 1976, 1981), which he edited, and a textbook, *Programowanie matematyczne. Zbiór zadań* [Mathematical programming. A collection of problems] (Polskie Wydawnictwo Ekonomiczne 1980), co-edited with Dr Stanisław Krawczyk. He was known for his surprising and inspiring ideas, as well as for his sharp academic criticism. As noted by Professor Smoluk in his book mentioned above: 'He liked scholarly order and was of a pedantic nature'. Professor Juliusz Siedlecki, the successor of Professor Bukietyński as head of the Department of Operations Research, described him in his memoirs published in the *Przegląd Statystyczny* (1997, vol. 44, no. 2) journal in the following words:

He favoured instant intervention whenever the presentation of ideas made by seminar participants was unclear or erroneous. By asking for explanations or pointing out mistakes made by the presenter, he was often able to suggest a correct resolution to the discussed issues. Whilst being impulsive in scholarly debates, he remained lenient with regards to others' mistakes; however, he showed no mercy to those who obstinately persisted in their errors.

Our seminars focused on methods and applications of the broadly understood Operations Research, a scientific area which was then intensely developed. The meetings attracted growing numbers of participants, among whom were new members of the Department's staff, PhD candidates supervised by the Professor, including those from other universities of economics (mainly the Academy of Economics in Katowice). Professor Bukietyński was profoundly dedicated to educating young scholars, as confirmed by 16 PhD degrees granted to his students.

The seminars often stretched out in time and turned into social gatherings where we could share the knowledge and literature forbidden in the period of the Polish People's Republic regime; for example, we listened to the banned Russian songs of Vysotsky, who opposed the Soviet regime. Unek spoke excellent Russian (also German), spent six months in St Petersburg (then Leningrad) on a scholarship, and had an expert knowledge of Russian literature. We, his colleagues, often wondered about this, knowing that his father was a judge in Brest-on-the-Bug (the Professor's birthplace) and among the Polish officers murdered by the Soviets in Katyń.

Professor Bukietyński was highly open to collaboration with the business sector, which often served as a testing ground for models and algorithms we prepared together. The Professor conducted over 20 such research projects, including those for the meat, timber, copper and electromechanics industries. In the late 1980s and early 1990s, the team working at the Department of Operations Research completed two long-term projects whose vital element was the implementation of IT technologies in supporting decision-making processes. These projects were: CPBP 10.9.VI.6 (1986–1991), *Optymalizacja komputerowa i jej wykorzystanie do zarządzania w warunkach reformy gospodarczej* [Computerised optimisation and its use for management in conditions of economic reform], headed by Professor Władysław Bukietyński, and CPBP 10.7/15 (1987–1991), *Zaprojektowanie, oprogramowanie i uruchomienie podsystemu optymalizacji operatywnych planów produkcji wyrobów mięsnych* [Design, programming and implementation of a subsystem of the optimisation of operational meat production planning], conducted by Ewa Konarzewska-Gubała, BEng, PhD. In the mid-1990s, the scope of research carried out in the Department of Operational Research was expanded by two important areas of theory and application, namely methods of neuron networks and their application, and methods of managing quality, in line with the concept of total quality management.

International collaboration was another important and highly valued aspect of Professor Bukietyński's academic and organisational activity. From 1984 to 1990, he was the Deputy Rector for development and international cooperation, maintaining constant and mutual contacts with the University of Munster, where in 1987, he delivered a series of lectures as Visiting Professor. Other members of his team performed the same role: Professor Stanisław Krawczyk at the University of Marburg (in 1991, 1993 and 1995) and Doctor Ewa Konarzewska-Gubała at the University of Richmond, Virginia (1994). As a team, we were very active within the international academic community, participating in numerous congresses, symposia and seminars in Poland and abroad, always with the strong support of our mentor. I can recall the circumstances related to my application to participate in the International Conference on Multiple Criteria Decision Making in Kyoto. I was invited there by

the organisers who covered the costs of my visit, however I had to apply to our University to finance the journey to Japan. Deputy Rector Bukietyński accepted my request, congratulated me on the invitation, but at the same time I was denied the permission of the Polish United Workers' Party (PZPR) Secretary at the University. Upon hearing the news, Professor grabbed the phone and threatened the Secretary with resignation from his post of Deputy Rector. I went to the conference and it was excellent.

Professor Władysław Bukietyński, for some simply Unek, left us suddenly thirty years ago. Yet, he is still with us, remembered for his mathematician's soul and an unwavering dedication to truth and perfection in all academic endeavours.

Associate Professor Ewa Konarzewska-Gubała