

Economic Growth and the In-Work Poverty Gap Between Temporary and Permanent Employees in the European Union: An Explainable Machine Learning Approach

Hadda Azhar Boukhris

PhD in Economics, University of Continuing Education (UFC), Algeria

boukhrisazhar123@gmail.com

Volume 61 · Issue 2 · 2026 · Pages 114–137

Article SES-2026-61-2-007

Abstract

Economic expansion does not necessarily narrow poverty inequalities within employment, yet the predictive value of growth for such disparities remains insufficiently examined. This study develops a contract-specific forecasting framework to assess whether real GDP growth improves predictions of the in-work poverty gap between temporary and permanent employees and whether its predictive value depends on labour-market segmentation and social protection. The framework is demonstrated using a simulated panel structured around Eurostat indicator definitions for 27 European Union countries over 2005–2024, comprising 513 lag-aligned country-year observations. Persistence and Elastic Net benchmarks are compared with Random Forest, XGBoost, and a multilayer perceptron using expanding-window out-of-time evaluations for 2015–2024. Matched predictor-ablation and interaction tests identify incremental forecasting gains, while SHAP explains model predictions. Including GDP growth reduces held-out root mean squared error from 0.817 to 0.721 percentage points for Elastic Net and from 0.807 to 0.733 for XGBoost. The multilayer perceptron achieves the lowest pooled error (0.693), although historical poverty disparities dominate feature attribution, and institutional interactions deliver negligible and inconsistent improvements. The study introduces a transparent framework for distinguishing the predictive contribution of economic growth from its presumed distributional consequences.

Keywords: In-work poverty; Economic growth; Labour-market dualisation; Explainable machine learning; European Union

JEL codes: I32; J41; O47; C45

Received

24 April 2026

Accepted

23 June 2026

Published

28 August 2026

1. Introduction

Economic growth may improve aggregate economic performance without narrowing poverty inequalities within the employed population. This discrepancy is particularly relevant to the divide between temporary and permanent employees, whose exposure to income insecurity can differ even within the same national labour market. Whether growth provides meaningful information about the evolution of this divide is therefore a substantive question for inclusive economic development. In-work poverty, defined in relation to equivalised disposable household income, captures a dimension of economic vulnerability that employment rates, wage growth, and aggregate poverty indicators cannot fully represent. Consequently, assessing the inclusiveness of growth requires attention not only to employment creation but also to differences in poverty exposure across employment arrangements.

The economic and social dimensions of poverty have also informed research published in *Studi Economici e Sociali*. Mazzotta et al. (2024) examined regional disparities in poverty in Italy and the role of minimum-income support, illustrating the importance of socioeconomic and territorial heterogeneity. Ortiz González-Conde (2024) addressed minimum-wage regulation in the European context, highlighting the institutional foundations of income protection. Together, these contributions situated poverty and labour-market protection within a broader economic and social-policy framework. However, neither directly investigated whether macroeconomic growth helps predict changes in poverty inequalities between temporary and permanent employees across European countries.

Recent international research has further clarified the complexity of in-work poverty. Barbieri et al. (2024) examined poverty persistence and household employment patterns across Western Europe, while Mussida and Sciulli (2025) demonstrated the relevance of low-paid employment, household earning arrangements, and labour-market conditions to poverty risks. Jara and Simon (2024) assessed income protection for atypical workers through unemployment-insurance arrangements. From a macroeconomic perspective, Lehner et al. (2024) showed how labour-market dualisation was associated with wage dynamics across European economies. These contributions established the importance of employment segmentation, household circumstances, and institutional protection. Yet their findings did not directly determine whether national economic growth added predictive information about the changing poverty differential between temporary and permanent employees.

Advances in machine learning offered a means of examining this unresolved relationship without imposing a strictly linear functional form. Lyons et al. (2025) demonstrated the applicability of machine learning to poverty assessment and social-assistance targeting. Conversely, Corral et al. (2025) showed that machine-learning approaches to poverty mapping could underperform conventional benchmarks, emphasising the importance of credible validation rather than methodological complexity alone. The relevant research gap was therefore not the absence of machine-learning applications to poverty, nor the previously documented disadvantage of temporary employment. It concerned the incremental predictive value of economic

growth for contractual inequalities in in-work poverty, particularly after accounting for historical poverty disparities, labour-market conditions, and social protection.

Against this background, the study addressed one research question: To what extent did real economic growth improve out-of-time prediction of the temporary-permanent in-work poverty gap across European Union countries, and how did this predictive contribution vary with labour-market segmentation and social protection?

Two hypotheses structured the investigation:

H1: Including real GDP growth improved out-of-time predictive accuracy for the temporary-permanent in-work poverty gap relative to otherwise comparable models incorporating historical poverty disparities and socioeconomic predictors.

H2: Models allowing interactions between economic growth, labour-market segmentation, and social protection achieved higher out-of-time predictive accuracy than comparable additive models.

The empirical framework examined Eurostat indicators for the 27 current European Union member states over 2005–2024, subject to the availability and comparability of national observations. The outcome was defined as the percentage-point difference between the in-work at-risk-of-poverty rates of temporary and permanent employees. The analysis compared conventional econometric and regularised regression benchmarks with Random Forest and gradient-boosted tree models. Time-ordered validation, predictor-ablation comparisons, and SHAP-based model interpretation were used to assess predictive performance, the additional information associated with economic growth, and potential nonlinear interactions. The analysis distinguished predictive relationships from causal effects and accounted for the temporal structure of the country-level data.

The study made three contributions. First, it shifted the analytical focus from aggregate in-work poverty to inequalities between employment-contract groups. Second, it evaluated the additional predictive value of economic growth against established socioeconomic predictors and simpler statistical benchmarks. Third, it investigated whether the predictive relevance of growth varied across labour-market and social-protection contexts. These contributions connected macroeconomic performance, employment segmentation, and social vulnerability within a unified empirical framework relevant to inclusive-growth and social-policy debates.

The remainder of the article is organised as follows. Section 2 reviews the relevant theoretical and empirical literature. Section 3 presents the data and methodological framework. Section 4 reports the empirical results and model explanations. Section 5 discusses their economic and social-policy implications, and Section 6 concludes.

2. Literature Review and Conceptual Framework

2.1. Economic Growth, Poverty Measurement, and the Distribution of Economic Gains

The central difficulty in linking economic growth to poverty outcomes lies in distinguishing aggregate economic expansion from the distribution of welfare gains. Growth in real output may coexist with heterogeneous changes in household income, employment security, and relative poverty exposure. This distinction is particularly consequential for in-work poverty, where the poverty threshold is defined relative to national median equivalised disposable income rather than a fixed level of individual earnings. An improvement in economic activity therefore need not produce a corresponding reduction in the relative poverty disadvantage of specific employment groups.

Recent evidence has challenged the assumption that aggregate output growth provides a sufficiently informative representation of household welfare dynamics. Wu et al. (2024) demonstrated that the transmission of economic growth to household welfare differed substantially across developing economies, with structural and institutional conditions influencing the strength of this relationship. Although their analysis concerned developing countries rather than European labour markets, its broader analytical implication was that aggregate growth and household welfare cannot be treated as interchangeable outcomes. The relevant transmission process must instead be examined at the level at which economic disadvantage is experienced.

This distinction also emerged in research published in *Studi Economici e Sociali*. Mazzotta et al. (2024) analysed regional poverty disparities in Italy and the contribution of minimum-income support to observed poverty patterns. Their analysis highlighted the relevance of territorial heterogeneity and redistribution, but did not establish whether aggregate economic growth predicted differences in poverty exposure across employment-contract categories. The unresolved issue was therefore not whether economic conditions and redistributive arrangements mattered, but whether growth contained information about the evolution of a clearly defined inequality within employment.

A further complication arises from the possibility that national economic aggregates predict broad welfare distributions more effectively than particular distributional gaps. Mahler et al. (2026) developed a parsimonious approach to predicting national income and consumption distributions, showing that a limited set of socioeconomic indicators could achieve accuracy comparable to substantially more complex models. Their findings challenged the presumption that additional predictors or methodological complexity necessarily improved poverty-related predictions. However, predicting aggregate income distributions is analytically different from predicting the difference between poverty risks associated with temporary and permanent employment.

These contributions exposed a measurement and prediction problem: information useful for explaining general welfare conditions may not be equally informative about

inequality between employment groups. The temporary-permanent in-work poverty gap thus represents a distinct outcome whose relationship with macroeconomic growth cannot be inferred mechanically from aggregate poverty trends.

2.2. Labour-Market Dualisation, Household Protection, and Conditional Poverty Risks

Labour-market segmentation provided the second foundation of the conceptual framework. Temporary and permanent employment differ in contractual stability, employment continuity, and exposure to labour-market adjustment. Nevertheless, the relationship between contractual status and poverty is neither exclusively individual nor institutionally uniform. It depends on the interaction between earnings, household resources, employment trajectories, and social protection.

Van Lancker (2012) established that temporary employment was associated with greater poverty exposure in a comparative European setting, while also showing that household circumstances and welfare regimes conditioned this disadvantage. Lohmann (2009) similarly demonstrated that social transfers and labour-market institutions influenced the incidence of in-work poverty. Taken together, these studies challenged explanations that attributed contractual poverty differences solely to wage inequalities. They established the importance of institutions and household circumstances, but did not determine how the poverty differential between contractual groups evolved in relation to subsequent macroeconomic growth.

Longitudinal evidence subsequently complicated the relationship between employment status and economic insecurity. Barbieri et al. (2024) showed that in-work poverty reflected household employment arrangements, social stratification, and persistence over time. Mussida and Sciulli (2025) further identified the relevance of low-paid work and partners' labour-market conditions for household poverty risks. These studies suggested that contractual instability and low earnings should not be interpreted as equivalent to household poverty. Multiple-earner households may cushion individual disadvantages, whereas limited household work intensity may sustain poverty despite an individual's employment.

More recent research reinforced this distinction. Aprea et al. (2026) demonstrated that low individual earnings did not necessarily translate into household poverty because income pooling and redistribution differed across European welfare systems. Meanwhile, D'Aguanno et al. (2026) showed that similar aggregate in-work poverty rates could conceal markedly different transitions between employment, poverty, and material deprivation. These findings introduced an important limitation for country-level analysis: a change in the observed poverty gap may reflect both changing economic circumstances and changes in the composition of the employed population. Consequently, country-level predictions cannot be interpreted as evidence of individual transitions out of poverty.

The relationship between contractual arrangements and poverty also extends beyond contemporaneous earnings. Friedrich and Teichler (2025) found that temporary

employment was associated with increased material deprivation in Germany, particularly among lower-income households. At the macroeconomic level, Lehner et al. (2024) demonstrated that labour-market dualisation influenced wage dynamics, including the wage growth of permanent employees. These findings indicate that segmentation may operate through both direct employment insecurity and broader wage-setting relationships. Importantly, the poverty gap could therefore change through developments affecting either employment group, rather than exclusively through changes among temporary workers.

Institutional protection further complicates the expected relationship. Ortiz González-Conde (2024), examined minimum-wage regulation within the European context, illustrating the relevance of national regulatory arrangements for income protection. Jara and Simon (2024) demonstrated how unemployment-insurance design could alter the financial vulnerability of atypical workers. However, the institutional literature did not establish a uniform protective effect across all policy instruments. Perugini and Pompei (2026), analysing reforms across European countries, found a particularly important role for family policies, while other labour-market and social-protection reforms displayed more limited estimated effects.

These findings generated competing expectations about the relationship between economic growth and contractual poverty inequalities. Economic expansion could improve employment continuity and earnings opportunities for vulnerable workers, potentially narrowing the gap. Alternatively, stronger employment or income gains among permanent employees, persistent household vulnerabilities, or changes in workforce composition could sustain or widen it. Social protection could moderate these relationships, but its relevance would depend on institutional coverage, policy design, and the characteristics of the affected households.

The resulting theoretical implication was not a predetermined positive or negative association between growth and the poverty gap. Rather, it was the possibility that this association varied across labour-market and institutional contexts. Whether such heterogeneity provided useful predictive information remained an empirical question.

2.3. The Predictive Frontier: Nonlinearity, Model Comparison, and the Remaining Research Gap

The growing use of flexible statistical methods in poverty research created an opportunity to investigate relationships that conventional additive specifications might not adequately represent. However, recent methodological developments also established that model complexity was not itself a scientific contribution.

Lyons et al. (2025) illustrated how machine learning could support multidimensional poverty assessment and assistance targeting by incorporating heterogeneous household and geographical information. Mussida et al. (2026), using a relative-importance framework and cluster analysis for European countries, further demonstrated that the relevance of socioeconomic factors for deprivation differed across national contexts. Together, these studies supported the importance of

examining predictor heterogeneity rather than assuming that socioeconomic characteristics carried the same information in every setting.

Yet flexibility introduced its own methodological risks. Corral et al. (2025) showed that machine-learning approaches to poverty mapping could underperform conventional alternatives when the available predictors contained limited information about the underlying outcome. Mahler et al. (2026) similarly demonstrated the practical value of parsimonious prediction. These contributions challenged a common assumption in applied machine learning: that more elaborate algorithms necessarily produce more economically informative results.

Predictive validity also depends on whether the evaluation design corresponds to the intended forecasting problem. Kapoor and Narayanan (2023) documented how data leakage across machine-learning applications generated overly optimistic performance estimates. In a country-year setting, random observations cannot automatically be treated as independent, since repeated national observations, persistent poverty inequalities, and common shocks create temporal and cross-sectional dependence. Model comparisons must therefore evaluate genuinely held-out periods and prevent future information from entering the training process.

The frontier of research on in-work poverty consequently displayed three complementary strengths but one unresolved intersection. Longitudinal and comparative studies had identified the relevance of employment trajectories, household arrangements, and institutions. Macroeconomic research had established that growth was an imperfect signal of distributional welfare. Machine-learning studies had demonstrated both the potential and the limitations of flexible poverty prediction. What remained insufficiently established across these strands was whether real GDP growth supplied incremental, temporally transferable predictive information about the temporary-permanent in-work poverty gap after accounting for its persistence and established socioeconomic predictors, and whether this information depended on labour-market segmentation and social protection.

This distinction was consequential. A model in which GDP growth appeared important according to an explanation metric would not, by itself, establish that growth improved prediction. The stronger test required comparing otherwise equivalent models with and without growth under the same out-of-time validation procedure. Likewise, observing a nonlinear association would not demonstrate the relevance of institutional interactions unless models allowing these interactions outperformed suitable additive benchmarks. SHAP-based explanations could clarify the structure of fitted predictions, but could not identify causal mechanisms or policy effects.

Accordingly, the study developed a prediction-oriented framework in which the historical poverty gap represented the baseline information set; macroeconomic and socioeconomic indicators formed the extended information set; and flexible interactions provided a separate, testable source of predictive information. This framework supplied the conceptual basis for the two hypotheses introduced in Section

1. It distinguished the incremental contribution of economic growth from the additional contribution of conditional, potentially nonlinear relationships.

The intended contribution was therefore not another demonstration that temporary employees faced poverty disadvantages, nor an assertion that machine learning outperformed econometric methods. It was a disciplined assessment of whether macroeconomic growth improved the prediction of contractual inequalities in economic security, under what observed conditions this information was relevant, and whether the improvement survived comparison with simpler and temporally credible benchmarks. Such evidence would clarify the informational limits of aggregate growth as an indicator of inclusive economic development, while preserving the distinction between predictive associations and causal effects.

3. Data and Empirical Methodology

3.1. Research Design, Data Sources, and Variable Construction

The study adopted a longitudinal, cross-country predictive framework to investigate whether economic growth provided additional information for predicting inequalities in poverty exposure between temporary and permanent employees across European Union countries. The empirical design integrated macroeconomic performance, labour-market segmentation, and social protection within a unified analytical framework combining conventional econometric benchmarks with explainable machine-learning techniques. Rather than assuming that economic expansion necessarily reduced poverty, the study examined whether real GDP growth improved the prediction of the temporary-permanent in-work poverty gap beyond historical poverty disparities and established socioeconomic indicators.

The geographical scope comprised the 27 current European Union member states over 2005–2024. This twenty-year horizon encompassed distinct economic environments, including the global financial crisis, the European sovereign-debt crisis, the COVID-19 pandemic, and the subsequent inflationary episode. These developments provided a relevant temporal setting for evaluating predictive relationships under heterogeneous macroeconomic conditions. The complete analytical design comprised 540 country-year observations, covering 27 countries over 20 years. Following the introduction of one-year lagged variables, the model-ready panel comprised 513 country-year observations for 2006–2024. The initial training period, 2006–2014, contained 243 observations, while the out-of-time evaluation period, 2015–2024, contained 270 observations.

The dependent variable was defined using the European Union Statistics on Income and Living Conditions (EU-SILC) concept of in-work at-risk-of-poverty, which identifies employed individuals living in households whose equivalised disposable income falls below 60% of the national median. This household-based definition captures relative income poverty rather than individual low earnings, thereby accounting conceptually for differences in household resources and income-sharing arrangements. The distinction was essential because temporary employment may increase income

insecurity without necessarily translating into household poverty, depending on household composition and institutional protection.

The temporary-permanent in-work poverty gap was constructed as:

$$GAP_{it} = IWP R_{it}^{TEMP} - IWP R_{it}^{PERM} \quad (1)$$

where GAP_{it} represents the poverty differential in country i and survey year t , $IWP R_{it}^{TEMP}$ denotes the in-work at-risk-of-poverty rate among temporary employees, and $IWP R_{it}^{PERM}$ represents the corresponding rate among permanent employees. The resulting indicator is expressed in percentage points. Positive values indicate greater poverty exposure among temporary employees, while negative values indicate the reverse. Unlike aggregate in-work poverty measures, this indicator directly captures inequality in poverty exposure between employment-contract groups.

Real GDP growth constituted the principal economic predictor and was measured as the annual percentage change in real economic output. Labour-market segmentation was represented by the proportion of temporary employees in total dependent employment, while social protection was proxied by general-government expenditure on social protection as a percentage of GDP. The unemployment rate captured broader labour-market conditions, whereas harmonised consumer-price inflation was incorporated into extended robustness specifications.

The selection of explanatory variables followed the conceptual relationships established in Section 2. Historical poverty disparities accounted for persistence in contractual inequality; real GDP growth represented aggregate economic performance; temporary-employment prevalence captured labour-market dualisation; and social-protection expenditure approximated the institutional environment associated with income security. The unemployment and inflation indicators accounted for complementary macroeconomic conditions. This parsimonious selection restricted unnecessary dimensionality while retaining the principal economic and social factors relevant to the research hypotheses. Table 1 presents the operational definitions, measurement units, analytical roles, and corresponding statistical sources.

The construction of the analytical variables incorporated the distinction between survey years and household-income reference periods. Because EU-SILC household-income information generally referred to the preceding calendar year, the macroeconomic and institutional predictors were aligned with the corresponding income-reference period. Accordingly, GDP growth and other socioeconomic predictors observed in calendar year $t-1$ were associated with the poverty gap reported in survey year t . Country-specific differences in income-reference periods were recognised when establishing temporal comparability.

Table 1. Variable Definitions, Measurement, and Data Sources

Variable	Operational definition	Unit	Analytical role	Eurostat reference
GAP	Temporary minus permanent employees' in-work poverty rate	Percentage points	Dependent variable	Author calculation
IWPR_TEMP	In-work at-risk-of-poverty rate among temporary employees	%	Outcome component	ilc_iw05
IWPR_PERM	In-work at-risk-of-poverty rate among permanent employees	%	Outcome component	ilc_iw05
GDPG	Annual real GDP growth	%	Principal predictor	nama_10_gdp
TEMP	Temporary employees as a proportion of total employees	%	Labour-market segmentation	tesem110
SOC	General-government expenditure on social protection, COFOG GF10	% of GDP	Institutional predictor	gov_10a_exp
UNEMP	Annual unemployment rate, ages 15–74	%	Labour-market control	une_rt_a
INFL	Annual average harmonised consumer-price inflation	%	Robustness control	prc_hicp_ainr

Source: Prepared by the author.

3.2. Econometric Specification and Explainable Machine Learning Framework

The empirical strategy combined interpretable econometric benchmarks with explainable machine-learning algorithms to assess the incremental predictive contribution of economic growth and the relevance of conditional relationships involving labour-market segmentation and social protection. The methodological objective was to distinguish substantive improvements in predictive information from gains attributable solely to increased model flexibility. The econometric benchmark was specified as follows:

$$\begin{aligned} GAP_{it} = & \alpha_i + \rho GAP_{i,t-1} + \beta GDPG_{i,t-1} \\ & + \gamma_1 TEMP_{i,t-1} + \gamma_2 SOC_{i,t-1} + \gamma_3 UNEMP_{i,t-1} \quad (2) \\ & + \theta_1 (GDPG \times TEMP)_{i,t-1} \\ & + \theta_2 (GDPG \times SOC)_{i,t-1} + \varepsilon_{it} \end{aligned}$$

where α_i represents country-specific time-invariant heterogeneity, and ρ captures persistence in contractual poverty inequalities. The coefficient β describes the conditional linear association between real GDP growth and the subsequent observed poverty gap. The coefficients γ_1 , γ_2 , and γ_3 represent the relationships involving temporary-employment prevalence, social-protection expenditure, and unemployment, respectively. The interaction coefficients θ_1 and θ_2 allow the relationship involving GDP

growth to vary across labour-market and institutional conditions, while ε_{it} denotes the disturbance term. The specification served as an interpretable predictive benchmark rather than a structural causal model. Country-specific effects accounted for persistent national heterogeneity, while the lagged dependent variable incorporated historical dependence in poverty inequalities. Potential dynamic-panel bias associated with the inclusion of a lagged outcome was recognised when interpreting individual parameters. Accordingly, the analysis prioritised predictive performance rather than causal identification.

The machine-learning framework incorporated Elastic Net regression, Random Forest, and Extreme Gradient Boosting (XGBoost). Elastic Net provided a regularised linear benchmark capable of limiting estimation instability associated with correlated socioeconomic predictors. Random Forest represented nonlinear relationships through the aggregation of decision trees, while XGBoost employed a regularised gradient-boosting architecture capable of capturing complex interactions under controlled model complexity (Chen & Guestrin, 2016). The modelling architecture distinguished four complementary specifications. The first relied on historical persistence in the poverty gap. The second incorporated labour-market and institutional predictors without GDP growth. The third added GDP growth to an otherwise identical predictor set. The fourth allowed theoretically motivated interactions between economic growth, temporary-employment prevalence, and social protection. These specifications provided a structured comparison of predictive information and functional flexibility.

The first hypothesis, H1, was operationalised through matched predictor-ablation comparisons. Models including real GDP growth were compared with corresponding specifications excluding it, while maintaining identical socioeconomic predictors, training observations, evaluation periods, and hyperparameter-selection procedures. This design isolated the incremental predictive information associated with economic growth. Support for H1 required improvements in out-of-time predictive accuracy rather than a statistically significant coefficient or a high feature-importance score alone. The second hypothesis, H2, concerned the conditional predictive relevance of economic growth. Its evaluation compared additive specifications with models incorporating interactions between GDP growth, temporary employment, and social protection. Within regularised regression, these relationships were represented through explicitly constructed interaction terms. Within tree-based algorithms, restricted and interaction-permitting configurations provided complementary comparisons. This approach distinguished improvements attributable to theoretically relevant interactions from those arising merely from additional algorithmic flexibility.

Model interpretability was incorporated using SHapley Additive exPlanations (SHAP), following Lundberg and Lee (2017). SHAP values decomposed individual predictions into contributions associated with the explanatory variables. Dependence and interaction analyses examined how the model's predictive use of economic growth varied across observed levels of labour-market segmentation and social-protection expenditure. This framework provided a basis for interpreting heterogeneous predictive relationships without imposing a predetermined linear structure.

Importantly, SHAP-based explanations were not interpreted as estimates of causal effects. A substantial predictive contribution associated with GDP growth indicated that the fitted model used information related to growth when generating predictions; it did not establish that economic growth caused changes in poverty inequalities. Similarly, nonlinear relationships involving social protection did not demonstrate that public expenditure causally altered the transmission of growth to poverty. The scientific relevance of these relationships depended on their contribution to credible out-of-time predictions and their stability across specifications. The combined framework therefore distinguished historical persistence, incremental macroeconomic information, and conditional predictive heterogeneity. This separation directly addressed the research gap identified in Section 2 and prevented model complexity from being mistaken for substantive economic contribution.

3.3. Temporal Validation, Hypothesis Testing, and Robustness Assessment

The validation framework accounted for the temporal dependence and repeated national observations inherent in the panel structure. Conventional random splitting was excluded from the principal evaluation strategy because it could place future observations in the training sample while evaluating predictions for earlier periods. Such information leakage can produce overly optimistic estimates of predictive accuracy, particularly in datasets characterised by serial dependence and common macroeconomic shocks (Kapoor & Narayanan, 2023).

An expanding-window temporal validation procedure was adopted. The initial training period covered 2006–2014 and comprised 243 country-year observations in the complete analytical design. The out-of-time evaluation period extended from 2015 to 2024 and comprised 270 observations, distributed across ten annual evaluation periods. At each forecast origin, models were trained using observations from preceding survey years and evaluated against the subsequent held-out year. The training window subsequently expanded chronologically, maintaining the temporal ordering of the data. This validation structure assessed temporal generalisation within European countries already represented in the historical training sample. It did not assume that predictive relationships estimated for these countries could automatically be transferred to previously unseen economies or institutional settings.

All preprocessing operations were restricted to the relevant training periods. Standardisation parameters, predictor-imputation procedures where required, and hyperparameter choices were determined without access to the corresponding held-out outcomes. Model selection employed nested chronological validation within each training sample. Country-specific effects and historical poverty-gap information were likewise restricted to observations available at the relevant forecast origin. Predictive performance was assessed primarily through root mean squared error (RMSE), supplemented by mean absolute error (MAE). Both measures were expressed in percentage points of the temporary-permanent in-work poverty gap. RMSE placed greater weight on relatively large prediction errors, whereas MAE measured the average absolute difference between observed and predicted poverty gaps. Model

comparisons were based on identical held-out observations to prevent differences in sample composition from influencing rankings.

H1 was evaluated by comparing the predictive errors of models with and without GDP growth. Reductions in held-out RMSE and MAE were interpreted as evidence of incremental predictive information when they remained sufficiently consistent across the evaluation horizon. Improvements concentrated in individual countries or isolated macroeconomic episodes were distinguished from those persisting across broader economic conditions.

H2 was examined through comparisons between additive and interaction-augmented models. Evidence supporting the hypothesis required interaction-permitting specifications to improve out-of-time predictive accuracy relative to suitable additive benchmarks. Estimated interaction coefficients and SHAP-based patterns provided additional interpretive information, but were not regarded as sufficient support in the absence of measurable forecasting improvements.

Uncertainty assessment focused on paired forecast-loss differences and their variation across countries and evaluation years. Because repeated national observations were serially dependent and European economies were exposed to common shocks, individual country-year prediction errors were not treated as fully independent. The interpretation consequently emphasised the magnitude, direction, and stability of predictive improvements alongside uncertainty estimates appropriate to the panel structure. Robustness assessment addressed four principal dimensions. First, alternative initial training windows examined whether predictive performance depended on the length and composition of the historical estimation period. Second, the COVID-19 episode was evaluated separately to distinguish general predictive patterns from relationships driven by exceptional economic conditions. Third, leave-one-country-out sensitivity exercises assessed whether the comparative results depended disproportionately on particular national economies. Fourth, expanded specifications incorporated harmonised inflation to examine whether the incremental predictive contribution of economic growth remained informative after accounting for additional macroeconomic conditions. The robustness framework further distinguished sensitivity to model specification from changes in statistical coverage. Common-sample comparisons were required whenever alternative predictor sets or modelling procedures could alter the set of eligible observations. This ensured that differences in predictive performance were evaluated independently of changes in the composition of the analytical sample.

Three substantive limitations governed interpretation. First, the dependent variable measured differences between aggregate poverty rates rather than individual transitions into or out of poverty. Second, social-protection expenditure represented aggregate government spending rather than individual policy eligibility, programme coverage, or benefit adequacy. Third, predictive associations could reflect changes in employment composition, household circumstances, or unobserved socioeconomic factors. Consequently, neither econometric coefficients nor machine-learning explanations were interpreted as causal policy effects. The methodological framework

thus provided a structured assessment of whether aggregate economic growth contributed meaningful additional information about inequalities in poverty exposure between temporary and permanent employees. By comparing historical persistence, socioeconomic predictors, economic growth, and theoretically specified interactions under a common temporal validation strategy, the study established an analytically transparent basis for examining the predictive relevance of macroeconomic performance to employment-related social inequality.

4. Empirical Results

4.1. Descriptive Characteristics of the Analytical Sample

The analytical dataset comprised 513 country-year observations for 27 European Union countries over 2006–2024, following the construction of one-year lagged variables. The initial estimation period contained 243 observations, whereas the out-of-time evaluation period comprised 270 observations distributed equally across ten annual assessment years. Table 2 reports the descriptive statistics for the temporary–permanent in-work poverty gap and the principal socioeconomic predictors. The mean in-work poverty rate among temporary employees was 17.730%, compared with 6.985% among permanent employees. The resulting poverty gap averaged 10.745 percentage points, with a standard deviation of 2.398 percentage points. Its observed values ranged from 5.31 to 18.29 percentage points. Real GDP growth averaged 2.115%, with values ranging from –8.33% to 11.46%. Temporary employment represented 11.125% of dependent employment on average, while social-protection expenditure averaged 17.051% of GDP. The corresponding average unemployment and inflation rates were 7.902% and 2.647%, respectively.

Table 2. Descriptive Statistics of the Analytical Variables

Variable	N	Mean	SD	Min	Median	Max
In-work poverty: temporary (%)	513	17.730	4.401	8.94	17.19	29.40
In-work poverty: permanent (%)	513	6.985	2.271	2.29	6.67	13.50
Poverty gap (pp)	513	10.745	2.398	5.31	10.49	18.29
Real GDP growth (%)	513	2.115	3.204	–8.33	2.17	11.46
Temporary employment (%)	513	11.125	5.690	0.70	10.24	25.16
Social protection (% GDP)	513	17.051	4.211	8.87	16.75	26.34
Unemployment (%)	513	7.902	2.810	2.02	7.32	16.99
Inflation (%)	513	2.647	2.146	–0.65	2.19	10.63

Source: Author's calculations using R.

4.2. Out-of-Time Predictive Performance Across Model Families


Predictive performance was evaluated over the 2015–2024 assessment period, comprising 270 held-out observations. Ten annual expanding-window evaluations were conducted for each model. Table 3 presents the pooled root mean squared error (RMSE) and mean absolute error (MAE), both expressed in percentage points. The multilayer perceptron achieved the lowest pooled forecast errors, with an RMSE of 0.693 and MAE of 0.552 percentage points. Relative to the persistence benchmark, this corresponded to a 15.1% reduction in RMSE. Among the regularised regression specifications, the interaction-augmented Elastic Net achieved an RMSE of 0.716, compared with 0.721 for the model incorporating growth without interactions and 0.817 for the specification excluding GDP growth. The XGBoost specification incorporating targeted interactions achieved an RMSE of 0.727, compared with 0.733 for its additive counterpart and 0.807 for the additive specification excluding economic growth. Random Forest recorded an RMSE of 0.753. Figure 1 presents the annual variation in held-out RMSE and the corresponding pooled rankings. Forecast errors varied across evaluation years, with several specifications exhibiting higher errors during 2020–2021.

Table 3. Comparative Out-of-Time Predictive Performance

Model	Test N	RMSE ↓	MAE ↓
Neural network (MLP 8-4)	270	0.693	0.552
Elastic Net with interactions	270	0.716	0.576
Elastic Net with GDP growth	270	0.721	0.576
XGBoost with targeted interactions	270	0.727	0.579
XGBoost additive with GDP growth	270	0.733	0.581
XGBoost unrestricted	270	0.734	0.581
Random Forest	270	0.753	0.597
XGBoost additive without growth	270	0.807	0.640
Persistence benchmark	270	0.817	0.648
Elastic Net without growth	270	0.817	0.661

Source: Author's calculations using R.


Figure 1. Out-of-Time Predictive Performance, 2015-2024



Source: Author's calculations using R.

A supplementary neural-network specification was estimated using five input variables, two hidden layers containing eight and four neurons, respectively, and one output neuron. Its structure and out-of-time predictions are presented in Figure 2.

Figure 2. Neural Network Architecture and Out-of-Time Predictive Performance: MLP versus XGBoost



Source: Author's calculations using R.

The neural network recorded lower pooled prediction errors than the unrestricted XGBoost specification, whose RMSE and MAE were 0.734 and 0.581 percentage

points, respectively. Both models produced predictions concentrated around the diagonal reference line, with deviations across the observed poverty-gap distribution.

4.3. Incremental Predictive Value, Model Explanations, and Robustness

The incremental contribution of GDP growth and the relevance of conditional predictive relationships were assessed through paired model comparisons. Table 4 reports the change in MAE after introducing the additional information or interaction structure, the number of evaluation years in which prediction accuracy improved, and exploratory two-sided sign-flip test results based on the ten annual paired MAE differences.


Table 4. Incremental Predictive Value and Hypothesis Comparisons

Hypothesis / model comparison	MAE improvement (pp)	Improved years	Sign-flip p
H1: Elastic Net + GDP growth	+0.0854	10/10	0.002
H1: XGBoost + GDP growth	+0.0583	7/10	0.104
H2: Elastic Net + interactions	−0.0003	6/10	0.980
H2: XGBoost + interactions	+0.0023	6/10	0.568

Source: Author's calculations using R.

In the Elastic Net comparison, incorporating GDP growth reduced MAE by 0.0854 percentage points, with improvements recorded in all ten evaluation years. The corresponding exploratory sign-flip test produced a p-value of 0.002. For additive XGBoost, introducing GDP growth reduced MAE by 0.0583 percentage points. Improvements occurred in seven of the ten years, with a sign-flip p-value of 0.104. Thus, the magnitude and temporal consistency of the predictive gain differed between the two model families.

The interaction comparisons produced substantially smaller changes in forecast accuracy. The Elastic Net interaction specification increased MAE by approximately 0.0003 percentage points, whereas the targeted-interaction XGBoost specification reduced MAE by 0.0023 percentage points. Both improvements and deteriorations appeared across the annual evaluation periods. Neither exploratory test provided strong evidence of a consistent improvement from introducing the interactions. The feature-attribution analysis was conducted on held-out predictions generated by the unrestricted XGBoost model. Figure 3 reports the mean absolute SHAP contributions and the corresponding distribution of contributions across the 270 evaluation observations.

Figure 3. Out-of-Time XGBoost Feature Attribution

Source: Author's calculations using R.

Source: Author's calculations using simulated data. SHAP contributions were calculated separately for each held-out year using models fitted to earlier observations. The lagged poverty gap recorded the highest mean absolute SHAP contribution at 1.753 percentage points. Real GDP growth ranked second at 0.243 percentage points, followed by temporary-employment prevalence (0.101), social-protection expenditure (0.046), and unemployment (0.036). The SHAP distribution displayed both positive and negative contributions associated with GDP growth across the evaluated observations.

Figure 4 reports conditional model responses obtained from the unrestricted XGBoost specification estimated using observations through 2023. The predicted poverty gap was evaluated across GDP growth rates from -8% to 8% , with temporary-employment prevalence and social-protection expenditure fixed separately at their 15th and 85th percentiles.


Both panels display declining fitted poverty gaps across increasing GDP growth values. The predicted levels differ between the lower and upper percentiles of temporary employment and social protection. The separation between the curves varies over the evaluated growth range, with a stepwise pattern reflecting the tree-based model structure.

Additional sensitivity assessments used the same out-of-time predictions while restricting the evaluation observations to alternative periods. Excluding 2020-2021 reduced the targeted-interaction XGBoost RMSE from 0.727 to 0.677 percentage points. The MAE gain from adding GDP growth to additive XGBoost was 0.0397 percentage points after excluding these years, compared with 0.0583 for the full evaluation period. When the evaluation was restricted to 2015-2019, the GDP-growth MAE gain was 0.0174 percentage points. For 2022-2024, it was 0.0771 percentage

points. The corresponding interaction-related MAE gains were 0.0016 and 0.0125 percentage points, respectively.

An additional evaluation-sample sensitivity calculation excluded one country at a time without refitting the models. Under these exclusions, the additive XGBoost improvement associated with GDP growth ranged from 0.0539 to 0.0634 percentage points. The targeted-interaction improvement ranged from approximately -0.00003 to 0.00449 percentage points. Across the model comparisons, the simulated results showed a measurable incremental predictive contribution from GDP growth, although its strength differed across model families and evaluation periods. By contrast, the introduction of the specified growth interactions produced limited and inconsistent improvements in held-out prediction accuracy.

Figure 4. Conditional Model Response to Economic Growth



Source: Author's calculations using R.

5. Discussion

5.1. Economic Growth and the Persistence of Contractual Poverty Inequalities

The central finding of the illustrative analysis was that economic growth added predictive information about the temporary-permanent in-work poverty gap, but did not displace its historical persistence. Adding GDP growth reduced Elastic Net RMSE from 0.817 to 0.721 percentage points; the corresponding XGBoost comparison declined from 0.807 to 0.733. However, gains varied by algorithm and evaluation period. A variable can improve forecasting without identifying why poverty inequalities arise or demonstrating that growth reduces them.

Barbieri et al. (2024) located in-work poverty within persistent household and employment disadvantages, while D'Aguzzo et al. (2026) showed that similar

aggregate poverty rates can conceal different individual employment-poverty transitions. These findings clarify why the lagged poverty gap dominated model explanations. Yet persistence in a national gap is not equivalent to persistence among the same workers. Changes in household composition and selection into employment may reproduce a stable aggregate disparity even as individuals move between jobs and poverty states.

Moreover, poverty is measured relative to national median equivalised disposable income. Growth can raise real earnings without reducing relative poverty if median income rises faster. Conversely, a narrowing gap need not indicate gains among temporary workers: it may reflect deterioration among permanent employees. Thus, the predictive contribution of GDP growth should be understood as additional information about a distributional indicator, not evidence of inclusive growth itself.

5.2. Labour-Market Segmentation and the Limits of Institutional Proxies

The weak incremental performance of growth interactions posed a more demanding theoretical question. Temporary employment and social protection are plausibly relevant to poverty, but interacting GDP growth with their aggregate measures produced little stable forecasting improvement. This does not negate the importance of labour-market institutions; it exposes the distance between institutional mechanisms and the variables used to represent them.

Van Lancker (2012) documented a poverty disadvantage among temporary employees across Europe. Gebel and Gundert (2023), however, found that entering fixed-term employment from unemployment could reduce poverty risk substantially, comparably to entering permanent work in Germany. These results are not contradictory. The first compares existing employment groups; the second evaluates changes following an employment transition. A temporary job can protect against unemployment-related poverty yet leave its holder relatively disadvantaged among employed people.

Household arrangements further complicate this relationship. Mussida and Sciulli (2025) showed that low pay and household earning positions jointly shape poverty risks. Likewise, Jara and Simon (2024) demonstrated the importance of benefit eligibility for atypical workers, whereas Perugini and Pompei (2026) reported heterogeneous associations across types of social-policy reform. An aggregate social-protection expenditure ratio cannot capture eligibility rules, replacement rates, household targeting, or effective coverage. Consequently, the negligible predictive gains associated with the expenditure interaction provide little basis for claiming that social protection does not matter. They indicate only that this broad measure did not supply reliably useful conditional information in the tested specifications.

5.3. Machine Learning, Evidential Boundaries, and Policy Relevance

The neural network obtained the lowest pooled RMSE (0.693), but its lead over simpler models was limited. Corral et al. (2025) cautioned that sophisticated poverty-prediction methods can underperform traditional approaches when predictors are insufficiently

informative. Accordingly, algorithmic complexity was not itself an empirical contribution. The more defensible contribution was the matched, time-ordered assessment of the value added by economic growth and hypothesised interactions.

SHAP explanations ranked the lagged gap above GDP growth and the remaining socioeconomic indicators. These rankings described how a fitted model allocated predictions, not independent causal effects. Consistent with Kapoor and Narayanan (2023), chronological validation reduced one source of leakage, but shared European shocks, revised data vintages, and only ten annual test periods still restricted inference. The weak H2 result should therefore be reported rather than concealed by attractive dependence plots. For policy, the framework suggests monitoring contractual poverty gaps alongside GDP growth and aggregate employment. It does not identify which reforms would close the gap. More targeted evaluations would require longitudinal household data and detailed measures of contract transitions, benefit access, and working hours.

6. Conclusion

This study examined whether economic growth improved the prediction of the temporary-permanent in-work poverty gap and whether its predictive contribution varied with labour-market segmentation and social protection. By integrating conventional econometric benchmarks with explainable machine-learning methods, the research addressed an important distinction between aggregate economic performance and inequalities in economic security within employment. The analytical framework covered 27 European Union countries over 2005–2024, with 513 country-year observations in the model-ready dataset.

The illustrative results provided three principal findings. First, economic growth contributed additional predictive information beyond historical poverty disparities and established socioeconomic indicators, although the magnitude and stability of this contribution differed across modelling approaches. Second, historical poverty inequalities remained the strongest predictive component, indicating the importance of accounting for persistence when assessing the informational value of macroeconomic conditions. Third, allowing interactions between economic growth, temporary employment, and social-protection expenditure produced limited and inconsistent improvements in predictive accuracy. Accordingly, H1 received model-dependent support, whereas H2 was not consistently supported.

The study's intended contribution extended beyond comparing machine-learning algorithms. It established a framework for evaluating whether aggregate growth contained incremental predictive information about inequalities between employment-contract groups, rather than assuming that economic expansion necessarily translated into more inclusive labour-market outcomes. The combination of predictor-ablation tests, temporal validation, and SHAP-based explanations distinguished predictive relevance from model complexity and causal interpretation.

These distinctions have implications for economic and social-policy analysis. Aggregate GDP growth and employment expansion cannot, by themselves, establish whether improvements in economic performance are distributed equitably across workers. Monitoring contractual poverty disparities alongside macroeconomic indicators could therefore provide a more differentiated assessment of inclusive development. However, the analysis did not identify causal effects of growth or demonstrate that particular social-protection policies reduced poverty inequalities.

Several limitations qualify the findings. The numerical results were obtained from simulated observations and therefore cannot be interpreted as empirical evidence about actual European labour markets. Moreover, country-level poverty indicators cannot identify individual employment transitions, household income adjustments, or the effective coverage of social benefits. The restricted number of annual evaluation periods also limits the precision of comparisons between forecasting models.

Future research should replicate the framework using verified longitudinal Eurostat observations, assess its sensitivity to statistical discontinuities, and incorporate more detailed measures of household circumstances, employment instability, and social-protection eligibility. Individual-level longitudinal data would further permit investigation of the mechanisms underlying aggregate contractual poverty disparities.

The central implication is that the predictive relevance of economic growth must be established independently of its presumed capacity to reduce social inequality. A rigorous assessment of this distinction offers a more precise foundation for understanding when aggregate economic performance provides useful information about the economic security of workers in different contractual positions.

References

- Aprèa, M., Raitano, M., & Subioli, F. (2026). From individual earnings to household disposable income: Linking low pay and poverty risks in EU countries. *Economic Systems*, 50(2), Article 101362. <https://doi.org/10.1016/j.ecosys.2025.101362>
- Barbieri, P., Cutuli, G., & Scherer, S. (2024). In-work poverty in Western Europe. A longitudinal perspective. *European Societies*, 26(4), 1232-1264. <https://doi.org/10.1080/14616696.2024.2307013>
- Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 785-794). Association for Computing Machinery. <https://doi.org/10.1145/2939672.2939785>
- Corral, P., Henderson, H., & Segovia, S. (2025). Poverty mapping in the age of machine learning. *Journal of Development Economics*, 172, Article 103377. <https://doi.org/10.1016/j.jdeveco.2024.103377>
- D'Aguanno, M. C., Guio, A.-C., & Van Kerm, P. (2026). Cross-country evidence on the individual dynamics of in-work poverty in the EU. *Economic Systems*, 50(2), Article 101360. <https://doi.org/10.1016/j.ecosys.2025.101360>

- Friedrich, M., & Teichler, N. (2025). Do temporary employees experience increased material deprivation? Evidence from German panel data. *Journal of European Social Policy*, 35(2), 143–156. <https://doi.org/10.1177/09589287241300011>
- Gebel, M., & Gundert, S. (2023). Changes in income poverty risks at the transition from unemployment to employment: Comparing the short-term and medium-term effects of fixed-term and permanent jobs. *Social Indicators Research*, 167, 507–533. <https://doi.org/10.1007/s11205-023-03118-5>
- Jara, H. X., & Simon, A. (2024). A European unemployment benefit to protect atypical workers? *Social Indicators Research*, 171, 967–986. <https://doi.org/10.1007/s11205-023-03276-6>
- Kapoor, S., & Narayanan, A. (2023). Leakage and the reproducibility crisis in machine-learning-based science. *Patterns*, 4(9), Article 100804. <https://doi.org/10.1016/j.patter.2023.100804>
- Lehner, L., Ramskogler, P., & Riedl, A. (2024). Beggar thy co-worker: Labor market dualization and the wage growth slowdown in Europe. *ILR Review*, 77(5), 659–684. <https://doi.org/10.1177/00197939241248162>
- Lohmann, H. (2009). Welfare states, labour market institutions and the working poor: A comparative analysis of 20 European countries. *European Sociological Review*, 25(4), 489–504. <https://doi.org/10.1093/esr/jcn064>
- Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. *Advances in Neural Information Processing Systems*, 30, 4765–4774. <https://proceedings.neurips.cc/paper/2017/hash/8a20a8621978632d76c43dfd28b67767-Abstract.html>
- Lyons, A. C., Montoya Castano, A., Kass-Hanna, J., Zhang, Y., & Soliman, A. (2025). A machine learning approach to assessing multidimensional poverty and targeting assistance among forcibly displaced populations. *World Development*, 192, Article 107013. <https://doi.org/10.1016/j.worlddev.2025.107013>
- Mahler, D. G., Schoch, M., Lakner, C., & Nguyen, M. C. (2026). A parsimonious approach to predicting income distributions. *Journal of Development Economics*, 180, Article 103695. <https://doi.org/10.1016/j.jdeveco.2025.103695>
- Mazzotta, F., Parisi, L., & Bisogno, M. (2024). Un’analisi della povertà in Italia: Differenze regionali ed effetto del Reddito di Cittadinanza. *Studi Economici e Sociali*, 59(1–2), 51–102.
- Mussida, C., Parisi, M. L., & Pontarollo, N. (2026). Relative-importance analysis of material and social deprivation across European countries. *Social Indicators Research*, 181, Article 16. <https://doi.org/10.1007/s11205-025-03745-0>
- Mussida, C., & Sciulli, D. (2025). Low-pay work and the risk of poverty: A dynamic analysis for European countries. *The Journal of Economic Inequality*, 23, 1173–1196. <https://doi.org/10.1007/s10888-025-09666-9>
- Ortiz González-Conde, F. M. (2024). La experiencia española sobre salario mínimo en el contexto europeo? *Studi Economici e Sociali*, 59(1–2), 129–158.
- Perugini, C., & Pompei, F. (2026). Reforms and in-work poverty in the EU. *Economic Systems*, 50(2), Article 101361. <https://doi.org/10.1016/j.ecosys.2025.101361>
- Van Lancker, W. (2012). The European world of temporary employment: Gendered and poor? *European Societies*, 14(1), 83–111. <https://doi.org/10.1080/14616696.2011.638082>

Wu, H., Atamanov, A., Bundervoet, T., & Paci, P. (2024). Is economic growth less welfare enhancing in Africa? Evidence from the last forty years. *World Development*, 184, Article 106759. <https://doi.org/10.1016/j.worlddev.2024.106759>

How to cite this article

Boukhris, H. A. (2026). Economic Growth and the In-Work Poverty Gap Between Temporary and Permanent Employees in the European Union: An Explainable Machine Learning Approach. *Studi Economici e Sociali: Rivista di vita economica*, 61(2), 114-137. <https://economiciesocialijournal.it/article/SES-2026-61-2-007>