

Forecasting the Unemployment Rate in Ticino: A Time-Series Approach

Simone Foresti

Università della Svizzera italiana (USI) - Istituto di Ricerche Economiche (IRE)

Methodological report accompanying IRE quarterly publications

April 2026

Abstract

This report documents the time-series methodology used at the Istituto di Ricerche Economiche (IRE) to analyze and forecast the unemployment rate registered by the State Secretariat for Economic Affairs (SECO) in the Canton of Ticino (Switzerland). Using quarterly data from 2004Q1 onward, the study combines SECO unemployment statistics with explanatory variables from the Federal Statistical Office (BFS) and the CREA Institute of Applied Economics (economic output), the KOF Swiss Economic Institute (survey-based sentiment), CHF/EUR exchange rates, and Google Trends, harmonized into a consistent dataset. Predictors are lagged both to reflect real-time data availability and to account for publication delays affecting some series. The objective is to produce quarterly forecasts of the unemployment rate for the two quarters following the latest observed value.

The empirical strategy proceeds in two stages. First, Unobserved Components Models (UCM) decompose unemployment into a stochastic level, trend, and quarterly seasonality, and serve as a screening device to identify informative lagged predictor sets. Second, a grid of Seasonal ARIMA (SARIMA) models is estimated to capture univariate dynamics; the best-performing dynamics are then paired with the selected predictor sets to generate a large set of SARIMAX (Seasonal AutoRegressive Integrated Moving Average with Exogenous regressors) specifications. Candidate models are ranked by log-likelihood, Akaike Information Criterion (AIC), and Bayesian Information Criterion (BIC), and the final shortlist is chosen to balance criterion performance with diversity across information sets.

Seasonality is essential for fit and forecasting; the incremental contribution of auxiliary indicators is modest but systematic. Survey-based labour-market expectations (KOF) and canton-level activity (CREA GDP) deliver the strongest gains beyond the SARIMA structure, with exchange rates and a Google Trends specification retained in the shortlist as complementary information sets. The final forecast is the equally weighted average of a fixed Top-15 shortlist of SARIMA and SARIMAX specifications; predictive uncertainty combines within-model and between-model variance. Forecast errors are centered close to zero with no meaningful serial correlation, supporting the adequacy of the framework.

Contents

1	Introduction	3
1.1	The research institute	3
1.2	The Forecasting Product	3
1.2.1	Definition of Unemployment Rate in this Study	4
2	Data Collection and Sources	4
3	Theoretical Foundations of Time Series Analysis	5
3.1	Unobserved Components Model (Harvey, 1989)	5
3.2	ARIMA Models	6
3.3	SARIMA Models	6
3.4	SARIMAX Models	7
4	Methodology	7
4.1	Preliminary Analysis	7
4.2	Autocorrelation and Partial Autocorrelation Analysis	8
4.3	Stationarity Tests: Augmented Dickey-Fuller (ADF) Test	12
4.4	Residual Analysis through UCM	13
4.5	Explanatory Power of Auxiliary Indicators	16
4.5.1	Real-time Analysis	17
4.5.2	Seasonal UCM with lagged predictors	20
4.6	Selecting SARIMA dynamics (univariate benchmarking)	23
5	SARIMAX Modeling	26
5.1	Evaluation of SARIMAX candidates (lagged predictors)	26
5.1.1	Model comparison: information criteria and log-likelihood	26
5.2	Forecasting	30
5.2.1	Forecast setup and dynamic predictions	31
5.2.2	Forecast-error diagnostics: distribution and serial correlation	34
5.2.3	Benchmarking against UCM residuals (model-selection vs forecasting)	36
6	Conclusion	39
A	Data Sources: Full Documentation	41
A.1	Unemployment Data (Amstat / SECO)	41
A.2	Google Trends	41
A.3	Cantonal GDP Growth (CREA)	42
A.4	Retail Turnover (BFS)	42
A.5	KOF Economic Sentiment and Uncertainty Proxies	42
A.5.1	KOF Business Situation Indicator	42
A.5.2	KOF Economic Barometer	43
A.5.3	KOF Employment Indicator	43
A.6	Exchange Rates (ECB)	43
A.7	GDP and GVA Data (BFS)	43
A.8	SECO Data (Secondary and Tourism Sectors)	43
A.9	OECD Data	44

1 Introduction

1.1 The research institute

The Istituto di Ricerche Economiche (IRE) is a research institute with a long-standing tradition of monitoring and analyzing economic indicators in the Canton of Ticino (Switzerland). Founded in 1961 as Ufficio delle Ricerche Economiche (URE) by the Ticino cantonal administration, it was originally established in Bellinzona under the Department of Finance and Economic Affairs (DFE) to monitor regional economic performance. Now integrated into the Università della Svizzera Italiana (USI) in Lugano, IRE's mission encompasses both policy-oriented academic research and public policy consulting. IRE's academic activities mainly focus on research projects fulfilling its public mandate under agreements between the Canton and the University, managing grant-funded projects, and building research networks. Complementing these efforts, IRE provides public policy consulting and reporting, analyzes regional competitiveness, and offers insights into social, environmental, and economic dynamics. To support its applied research and service activities, IRE is organized around three thematic observatories: the Osservatorio del Turismo (O-Tur), the Osservatorio delle Dinamiche Economiche (O-De), which monitors labor market trends, and the Osservatorio Finanze Pubbliche ed Energia (O-FPE).

1.2 The Forecasting Product

This methodological report supports IRE's mandate to monitor labour-market conditions in the Canton of Ticino and to provide regular forecasts of the unemployment rate registered by the State Secretariat for Economic Affairs (SECO). Its purpose is to document a reproducible and transparent forecasting framework for readers of IRE's quarterly publications. The focus is deliberately operational: the goal is not theoretical interpretation, but robust short-horizon prediction in a setting where data availability and publication lags matter.

The empirical analysis uses quarterly data from 2004Q1 onward. The core outcome is the SECO unemployment rate, complemented by auxiliary indicators selected for timeliness and forecasting relevance (see Section 2 for sources and documentation). All predictors enter the models with pre-specified lags (in quarters) to reflect real-time information sets and ensure that explanatory variables are always available before the unemployment outcome they aim to forecast.

The results presented in this paper are based on 2026Q1 as the last observed quarter; the model-comparison tables reflect estimates through that quarter, and the operational forecast covers 2026Q2 and 2026Q3. The framework is designed to be updated as new quarterly observations become available.

The modelling strategy proceeds in two steps. First, a family of Unobserved Components Models (UCM) is estimated to decompose unemployment into a stochastic level/trend and quarterly seasonality, and to screen candidate exogenous indicator sets using one-step-ahead residual performance. Second, a grid of seasonal ARIMA specifications is estimated to capture strong univariate dynamics; the best-performing dynamics are then combined with the selected lagged indicator sets to generate a large set of SARIMAX (Seasonal Autoregressive Integrated Moving Average with Exogenous regressors) models. Throughout the study, candidate models are compared using likelihood- and information-based criteria as well as measures of the forecasting errors. Final predictions are produced through the simple average of a shortlist of the best SARIMA and SARIMAX models.

Prediction intervals are finally estimated to account for both within-model uncertainty and dispersion across models.

1.2.1 Definition of Unemployment Rate in this Study

The dependent variable used throughout this study is the SECO-registered unemployment rate for the Canton of Ticino. The series is obtained from the labour market statistics platform maintained by SECO (Amstat), which disseminates unemployment statistics at the national, cantonal, and communal levels. The SECO unemployment rate is defined as the share of persons registered as unemployed with the regional employment offices in the total labour force, according to SECO’s standard definition (see Appendix A).

In this paper we start the analysis based on the official *raw* (non-seasonally adjusted) quarterly SECO unemployment rate for Ticino, including all registered unemployed persons irrespective of age, nationality or gender. The analysis uses quarterly data from 2004Q1 onward. The models are designed to produce rolling forecasts on a quarterly basis. This forecasting service has been operational since 2025. For quarterly studies of IRE, the shortlisted model set is kept fixed while the operational forecast is rolled forward as new observations become available. The results in this paper use 2026Q1 (observed at 3.20%) as the last observation, with the forecast horizon covering 2026Q2 and 2026Q3.

2 Data Collection and Sources

The analysis relies on a harmonized quarterly dataset constructed from multiple sources. Data preparation consists of aligning frequencies and release calendars, imposing a consistent lag structure, and handling missing observations so that all predictors are available in real time for forecasting.

Table 1 lists each source, the series used, and its original frequency before aggregation to the quarterly frequency. Full documentation (variable definitions, construction methodology, access links, and notes on availability) is provided in Appendix A.

Table 1: Data sources

Series	Provider
Unemployment rate (Canton of Ticino)	Amstat / SECO
Job-search and unemployment-related search intensity	Google Trends
Cantonal Gross Domestic Product (GDP) growth rate	CREA (Univ. Lausanne)
Retail turnover index	BFS
Business Situation Indicator	KOF (ETH Zurich)
Employment Indicator	KOF (ETH Zurich)
Economic Barometer	KOF (ETH Zurich)
EUR/CHF exchange rate	European Central Bank (ECB)
GDP and Gross Value Added (GVA) by canton	BFS
Secondary and tourism sector indices	SECO via BFS PxWeb
Industrial production index	OECD

Notes: All series are aggregated to the quarterly frequency before entering the analysis. BFS: Bundesamt für Statistik (Federal Statistical Office); CREA: CREA Institute of Applied Economics (Univ. Lausanne); KOF: KOF Swiss Economic Institute (ETH Zurich). See Appendix A for source links, variable definitions, and construction notes.

3 Theoretical Foundations of Time Series Analysis

3.1 Unobserved Components Model (Harvey, 1989)

This preliminary UCM decomposition provides a model-based assessment of whether trend, seasonal, and cyclical dynamics are supported by the data, and it offers a coherent framework to evaluate the marginal contribution of candidate predictors. The resulting evidence is used to motivate a parsimonious SARIMAX specification (e.g., seasonal/non-seasonal differencing and a restricted set of exogenous variables), which is then selected and validated based on likelihood diagnostics and forecasting performance.

A general Unobserved Components Model (UCM) can be written as

$$y_t = \tau_t + \gamma_t + \psi_t + \mathbf{x}_t^\top \boldsymbol{\beta} + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, \sigma_\varepsilon^2), \quad (1)$$

where τ_t is the trend, γ_t the seasonal component, ψ_t the cyclical component, and \mathbf{x}_t a vector of observed covariates.

The trend is specified as a local linear trend,

$$\tau_t = \tau_{t-1} + \nu_{t-1} + \eta_t, \quad \eta_t \sim \mathcal{N}(0, \sigma_\eta^2), \quad (2)$$

$$\nu_t = \nu_{t-1} + \xi_t, \quad \xi_t \sim \mathcal{N}(0, \sigma_\xi^2). \quad (3)$$

For quarterly data ($s = 4$), a stochastic seasonal component can be specified as

$$\gamma_t = - \sum_{j=1}^{s-1} \gamma_{t-j} + \omega_t, \quad \omega_t \sim \mathcal{N}(0, \sigma_\omega^2), \quad (4)$$

which imposes the standard sum-to-zero restriction over one seasonal cycle (up to a disturbance).

The cyclical component can be modeled using Harvey's stochastic cycle,

$$\begin{pmatrix} \psi_t \\ \psi_t^* \end{pmatrix} = \rho \begin{pmatrix} \cos \lambda & \sin \lambda \\ -\sin \lambda & \cos \lambda \end{pmatrix} \begin{pmatrix} \psi_{t-1} \\ \psi_{t-1}^* \end{pmatrix} + \begin{pmatrix} \kappa_t \\ \kappa_t^* \end{pmatrix}, \quad \kappa_t, \kappa_t^* \sim \mathcal{N}(0, \sigma_\kappa^2). \quad (5)$$

Here $\rho \in (0, 1]$ is a damping factor controlling how quickly the cycle reverts toward zero, and $\lambda \in (0, \pi]$ is the cycle frequency (in radians per period), which determines the average cycle length.

Because the trend, seasonal, and cyclical components are latent (never directly observed), the model is cast in state-space form and the unknown variance parameters are estimated by maximum likelihood via the Kalman filter. The filter recursively propagates state estimates forward in time and revises them upon each new observation, generating one-step-ahead prediction errors whose joint likelihood is then maximized. The irregular component ε_t absorbs random shocks not attributable to any structural component and is treated as white noise.

I use UCM to decompose the quarterly unemployment series in Ticino and to identify statistically relevant components. I then incorporate potential explanatory variables into the UCM to assess their significance and impact on the dependent variable. Based on this analysis, I select the most informative variables and components to include in the SARIMAX framework, which is then refined to achieve the best fit and forecasting accuracy.

3.2 ARIMA Models

ARIMA (AutoRegressive Integrated Moving Average) models are widely used for time series analysis and forecasting. They combine three components: autoregression (AR), differencing (I), and moving averages (MA). The general form of an ARIMA model is denoted as $ARIMA(p, d, q)$, where:

- p is the number of lagged observations included in the autoregressive (AR) part;
- d is the number of times the series is differenced to achieve stationarity (integrated part);
- q is the order of the moving-average (MA) part, i.e., the number of lagged forecast errors (innovations) included in the model.

Letting $z_t = \Delta^d y_t$, the stationary process z_t follows an $ARMA(p, q)$ model:

$$z_t = c + \sum_{i=1}^p \phi_i z_{t-i} + \varepsilon_t + \sum_{j=1}^q \theta_j \varepsilon_{t-j}. \quad (6)$$

Equivalently, to write ARIMA models compactly, it is convenient to introduce the lag operator L such that $Ly_t = y_{t-1}$, and the difference operator $\Delta = (1 - L)$. Then, an $ARIMA(p, d, q)$ model can be expressed as:

$$\phi(L) \Delta^d y_t = c + \theta(L) \varepsilon_t, \quad \varepsilon_t \sim \text{i.i.d.}(0, \sigma^2), \quad (7)$$

where the autoregressive and moving-average polynomials are defined as

$$\phi(L) = 1 - \phi_1 L - \dots - \phi_p L^p, \quad \theta(L) = 1 + \theta_1 L + \dots + \theta_q L^q.$$

In words, ARIMA models apply differencing Δ^d to remove non-stationary behavior, and then model the resulting stationary series using an $ARMA(p, q)$ structure. When $d > 0$, the constant term c corresponds to a drift component in the original (undifferenced) series.

3.3 SARIMA Models

SARIMA (Seasonal ARIMA) models extend ARIMA models to explicitly account for seasonality in the data. They incorporate both non-seasonal and seasonal autoregressive, differencing, and moving-average components. A SARIMA model is denoted as $(p, d, q) \times (P, D, Q)_s$, where:

- p is the number of non-seasonal autoregressive terms;
- d is the number of non-seasonal differences;
- q is the number of non-seasonal moving-average terms;
- P is the number of seasonal autoregressive terms;
- D is the number of seasonal differences;
- Q is the number of seasonal moving-average terms;

- s is the length of the seasonal cycle (e.g. $s = 4$ for quarterly data).

Using the seasonal difference operator $\Delta_s = (1 - L^s)$, a SARIMA(p, d, q) \times (P, D, Q) $_s$ model can be written as:

$$\Phi(L^s) \phi(L) \Delta^d \Delta_s^D y_t = c + \Theta(L^s) \theta(L) \varepsilon_t, \quad \varepsilon_t \sim \text{i.i.d.}(0, \sigma^2), \quad (8)$$

where the seasonal AR and MA polynomials are:

$$\Phi(L^s) = 1 - \Phi_1 L^s - \dots - \Phi_P L^{Ps}, \quad \Theta(L^s) = 1 + \Theta_1 L^s + \dots + \Theta_Q L^{Qs}.$$

This formulation captures two distinct types of dependence: non-seasonal dynamics through p and q (lags such as $t - 1, t - 2, \dots$) and seasonal dynamics through P and Q (lags such as $t - s, t - 2s, \dots$).

3.4 SARIMAX Models

SARIMAX (Seasonal ARIMA with Exogenous Variables) models further extend SARIMA models by including exogenous predictors. These models are useful when external factors (for example, the exchange rate between the Euro and the Swiss Franc when studying regional unemployment) are believed to influence the time series. The general form is denoted as $(p, d, q) \times (P, D, Q)_s$ with exogenous variables \mathbf{x}_t .

A convenient representation treats the dependent variable as the sum of a regression component plus a SARIMA error process:

$$y_t = \mathbf{x}_t^\top \boldsymbol{\beta} + u_t, \quad (9)$$

$$\Phi(L^s) \phi(L) \Delta^d \Delta_s^D u_t = c + \Theta(L^s) \theta(L) \varepsilon_t, \quad \varepsilon_t \sim \text{i.i.d.}(0, \sigma^2). \quad (10)$$

Here, \mathbf{x}_t is a vector of exogenous variables; in applications, the elements of \mathbf{x}_t are replaced by appropriately lagged values \mathbf{x}_{t-k} to respect real-time data availability and ensure that predictors precede the unemployment outcome they aim to forecast. The coefficient vector $\boldsymbol{\beta}$ measures the marginal effect of each predictor on y_t . As in SARIMA models, the non-seasonal parameters (p, d, q) capture short-run dynamics, while the seasonal parameters $(P, D, Q)_s$ capture periodic dependence at the seasonal frequency.

4 Methodology

4.1 Preliminary Analysis

This study utilizes quarterly unemployment data for the Canton of Ticino spanning the period from 2004 to 2026. The dataset is structured as a time series to investigate its persistence, seasonal dynamics, and predictability across different econometric models. Figure 1 presents the historical evolution of the series.




Figure 1: Evolution of the quarterly unemployment rate in Ticino (CH) over time

Visual inspection suggests a recurring seasonal pattern superimposed on a slowly evolving medium-run component. The series reaches a local minimum around the second and third quarters of 2023. One possible interpretation is that seasonal labor-market dynamics (for example, temporary employment fluctuations in tourism-related activities) may contribute to these within-year movements; however, this interpretation is descriptive and is further assessed using formal time-series diagnostics.

4.2 Autocorrelation and Partial Autocorrelation Analysis

To further investigate persistence and potential seasonal patterns in quarterly unemployment, I inspect the sample Autocorrelation Function (ACF) and Partial Autocorrelation Function (PACF). The ACF measures the linear correlation between the series and its own lags: a slow decay at low lags signals strong persistence, while elevated spikes at multiples of four indicate annual seasonal dependence. The PACF measures the same lag- k correlation but *after* removing the contribution of intermediate lags, making it useful for identifying the order of the autoregressive component. Together, these correlograms provide an initial diagnostic of serial dependence and help motivate a parsimonious dynamic specification before estimating ARIMA/SARIMA-type models.

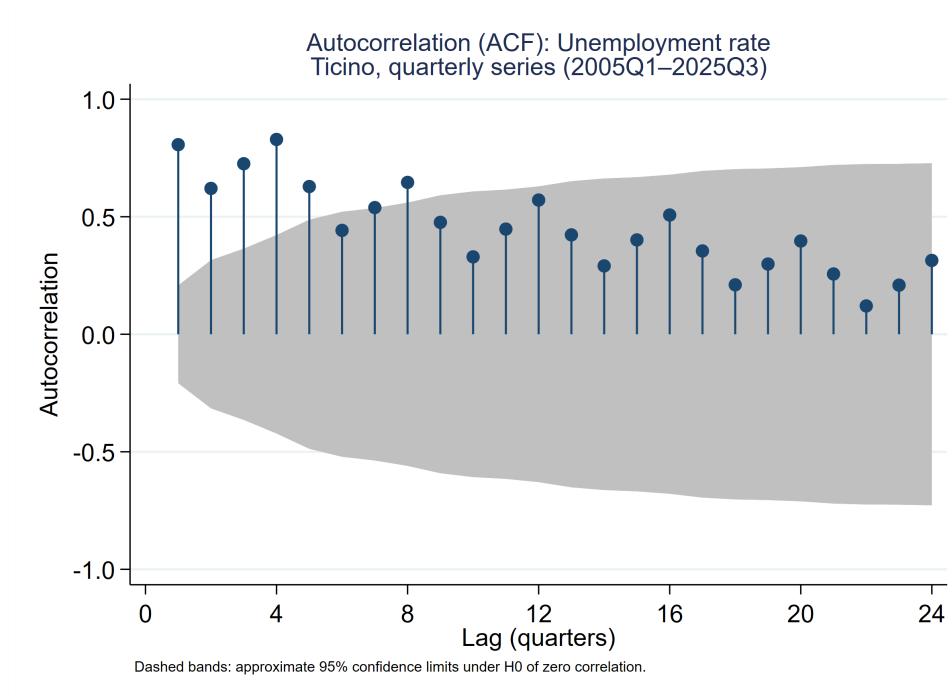


Figure 2: Autocorrelation of the quarterly unemployment rate

Figure 2 reports the ACF of quarterly unemployment. Two features stand out: autocorrelations decay slowly at the first few lags, consistent with strong persistence in the level of the series, and pronounced positive spikes appear at the seasonal lags 4, 8, and 12, indicating annual cyclical dependence. Both patterns motivate a SARIMA specification that captures non-seasonal persistence (via AR or differencing) alongside an explicit seasonal structure.

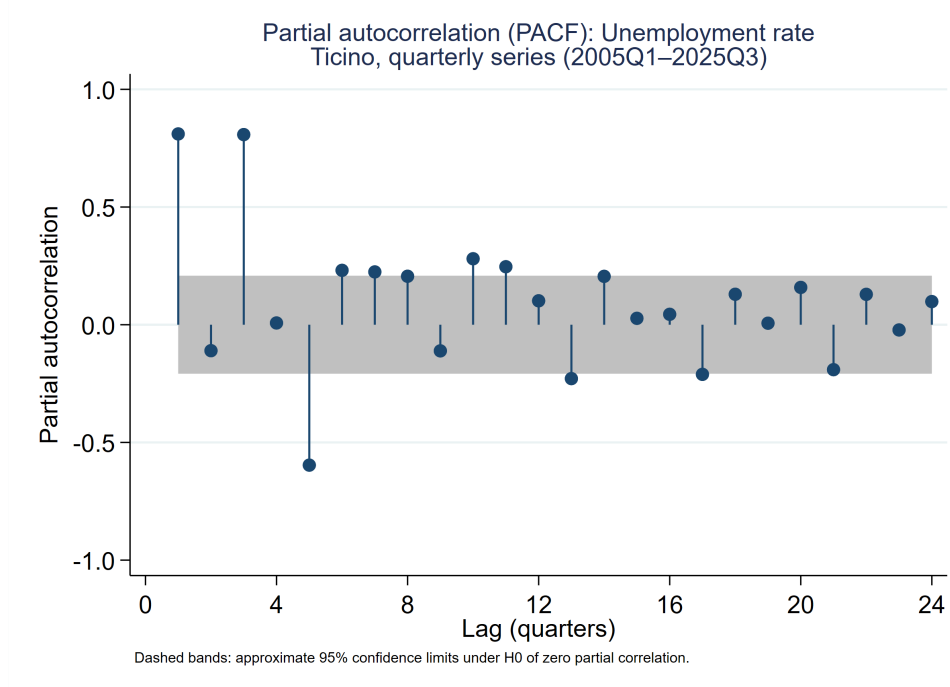


Figure 3: Partial autocorrelation of the quarterly unemployment rate

Figure 3 reports the PACF of quarterly unemployment. A dominant spike at lag 1, with subsequent lags decaying toward the confidence band, is consistent with a strong first-order autoregressive component in the short-run dynamics. A further spike at lag 4 indicates direct seasonal dependence at the annual frequency, reinforcing the case for an explicit seasonal AR term.

Although final model choice should be validated using information criteria and residual diagnostics, ACF and PACF patterns provide a standard heuristic:

- **AR**(p): PACF cuts off after lag p , while ACF typically decays gradually.
- **MA**(q): ACF cuts off after lag q , while PACF typically decays gradually.

Correlograms are therefore used as a screening device for parsimonious dynamic specifications, complemented by likelihood-based model selection and out-of-sample forecasting performance.

Autocorrelation Function (ACF). The autocorrelation function measures the correlation between observations of a time series separated by k periods. For a weakly stationary process $\{y_t\}$ with mean μ , the theoretical autocorrelation at lag k is

$$\rho(k) = \frac{\gamma(k)}{\gamma(0)}, \quad \gamma(k) = \text{Cov}(y_t, y_{t-k}). \quad (11)$$

In empirical applications, $\rho(k)$ is estimated by the sample autocorrelation

$$\hat{\rho}(k) = \frac{\hat{\gamma}(k)}{\hat{\gamma}(0)}, \quad \hat{\gamma}(k) = \frac{1}{T} \sum_{t=k+1}^T (y_t - \bar{y})(y_{t-k} - \bar{y}), \quad (12)$$

where T is the sample size and \bar{y} is the sample mean.

The ACF is typically displayed as a correlogram (a plot of $\hat{\rho}(k)$ against k). Spikes outside the confidence bands indicate autocorrelations that are statistically different from zero under the null hypothesis of no serial correlation, using pointwise confidence intervals.

Econometrically, the ACF provides qualitative guidance on time-series dynamics:

- *Persistence / potential non-stationarity*: a slow decay of the ACF at low lags is often indicative of strong persistence and may suggest a trend component or the need for differencing.
- *Seasonality*: for quarterly data ($s = 4$), pronounced spikes at lags 4, 8, 12, ... are consistent with seasonal dependence at the annual frequency.

Partial Autocorrelation Function (PACF). The partial autocorrelation at lag k measures the correlation between y_t and y_{t-k} *after controlling for* the intermediate lags $y_{t-1}, \dots, y_{t-k+1}$. A convenient definition is obtained from the autoregression

$$y_t = \sum_{j=1}^k \pi_{k,j} y_{t-j} + u_t, \quad (13)$$

where the PACF at lag k is the coefficient on the k -th lag, i.e. $\text{PACF}(k) = \pi_{k,k}$.

Seasonal differencing diagnostics. Given the evidence of possible seasonality, I also evaluate autocorrelation patterns after seasonal differencing. For quarterly data ($s = 4$), seasonal differencing is defined as

$$\Delta_s y_t = (1 - L^s)y_t = y_t - y_{t-4}. \quad (14)$$

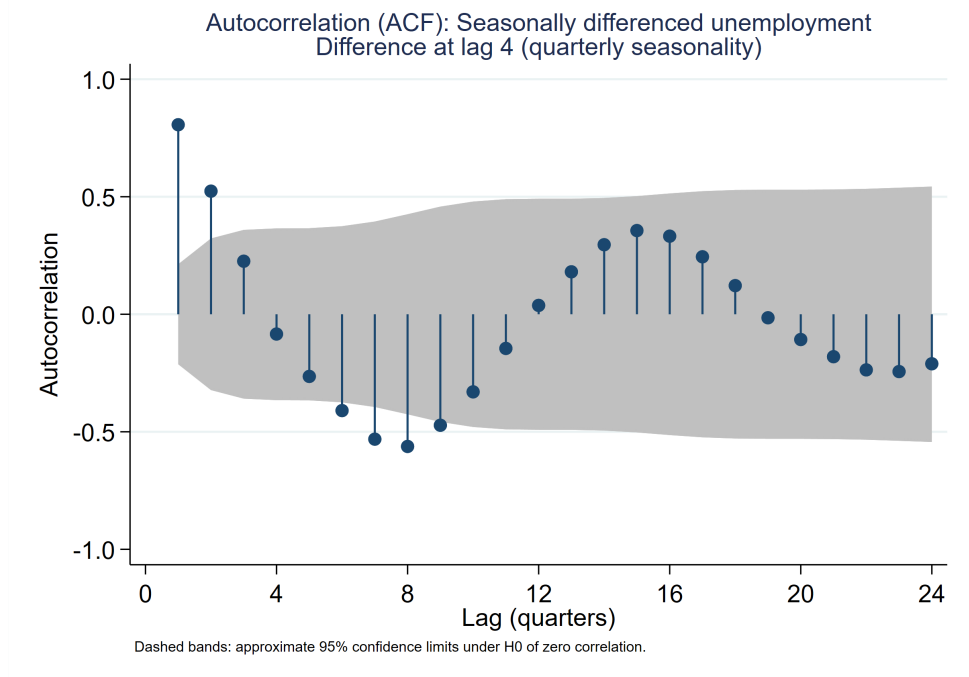


Figure 4: Autocorrelation of seasonally differenced quarterly unemployment

Figure 4 reports the ACF of the seasonally differenced series $\Delta_4 y_t$. The slow decay visible in the raw series is substantially reduced, confirming that seasonal differencing absorbs much of the annual cyclical dependence. A remaining significant spike at lag 4 suggests that a seasonal AR or MA term is still warranted in the SARIMA specification even after differencing.




Figure 5: Partial autocorrelation of seasonally differenced quarterly unemployment

Figure 5 reports the PACF after seasonal differencing. A spike at lag 4 is consistent with seasonal autoregressive behavior, supporting the inclusion of a seasonal AR component at the annual frequency.

4.3 Stationarity Tests: Augmented Dickey-Fuller (ADF) Test

As a preliminary step, I assess stationarity using the Augmented Dickey-Fuller (ADF) unit-root test. The ADF test evaluates the null hypothesis that the series contains a unit root and is therefore non-stationary, against the alternative hypothesis of stationarity around the deterministic components included in the test regression (e.g. a constant and/or a linear trend). In particular, the ADF test is based on estimating the regression

$$\Delta y_t = \alpha + \beta y_{t-1} + \delta t + \sum_{j=1}^k \zeta_j \Delta y_{t-j} + \varepsilon_t, \quad (15)$$

where k denotes the number of lagged differences included to account for serial correlation in the errors. The unit-root null corresponds to

$$H_0 : \beta = 0 \quad (\text{equivalently } \rho = 1),$$

while the stationary alternative corresponds to $H_a : \beta < 0$ (equivalently $\rho < 1$). The reported test statistic is the Dickey-Fuller $Z(t)$ statistic and p-values are computed using MacKinnon's approximate response surface, a regression-based interpolation of simulated critical-value distributions for the ADF test that provides more accurate tail probabilities than standard normal tables given the non-standard distribution of the $Z(t)$ statistic under the unit-root null.

Because quarterly data may exhibit seasonal dynamics, I run the ADF test both on the original series and on a seasonally differenced version of the series. For quarterly data

($s = 4$), seasonal differencing is defined as

$$\Delta_4 y_t = y_t - y_{t-4}.$$

Table 2 reports the ADF $Z(t)$ statistic and MacKinnon approximate p-values across alternative choices of k . On the original series, rejection of the unit-root null is not robust to the number of lagged differences included: the null is rejected only for $k = 0-1$ (p-values around 0.03), while for $k = 2-5$ the p-values are substantially larger (between 0.28 and 0.80). Overall, this pattern suggests that the inference is sensitive to the lag-augmentation choice.

On the seasonally differenced series $\Delta_4 y_t$, the evidence against a unit root becomes substantially stronger: for $k \geq 1$ the null is rejected at conventional significance levels with p-values close to zero, and even at $k = 0$ the rejection remains marginally significant at the 5% level. Taken together, these results indicate that seasonal differencing at lag 4 effectively reduces the persistence in the series and yields a transformed series that is suitable for subsequent ARIMA/SARIMA modeling and forecasting.

Table 2: ADF unit-root tests on quarterly unemployment (Ticino): original and seasonally differenced series

Series	Lags	Obs.	ADF statistic $Z(t)$	MacKinnon p-value
Quarterly unempl. rate	0	88	-3.064	0.0294
Quarterly unempl. rate	1	87	-3.054	0.0302
Quarterly unempl. rate	2	86	-0.859	0.8010
Quarterly unempl. rate	3	85	-0.892	0.7906
Quarterly unempl. rate	4	84	-2.015	0.2800
Quarterly unempl. rate	5	83	-1.540	0.5136
Seasonal diff. Δ_4	0	84	-2.987	0.0361
Seasonal diff. Δ_4	1	83	-4.125	0.0009
Seasonal diff. Δ_4	2	82	-4.470	0.0002
Seasonal diff. Δ_4	3	81	-5.342	0.0000
Seasonal diff. Δ_4	4	80	-4.012	0.0013
Seasonal diff. Δ_4	5	79	-4.923	0.0000

Notes: The ADF test is performed under the null hypothesis that the series follows a random walk without a unit root. To eliminate quarterly seasonal fluctuations, I compute the seasonally differenced series by subtracting the observation from four quarters prior: $\Delta_4 y_t = y_t - y_{t-4}$. The number of observations decreases with the number of lags and (for the differenced series) due to the loss of the first 4 quarters.

4.4 Residual Analysis through UCM

The Unobserved Components Model (UCM). To complement the unit-root evidence, I estimate an Unobserved Components Model (UCM) with a stochastic (random-walk) level, both *without* and *with* a quarterly seasonal component of period $s = 4$.

The baseline specification is the random-walk level (RW level) model:

$$y_t = \mu_t, \tag{16}$$

$$\mu_t = \mu_{t-1} + \eta_t, \quad \eta_t \sim \text{i.i.d.}(0, \sigma_\eta^2), \tag{17}$$

while the seasonal specification augments the measurement equation with a stochastic seasonal component:

$$y_t = \mu_t + \gamma_t, \quad (18)$$

$$\mu_t = \mu_{t-1} + \eta_t, \quad (19)$$

$$\gamma_t = - \sum_{j=1}^{s-1} \gamma_{t-j} + \omega_t, \quad \omega_t \sim \text{i.i.d.}(0, \sigma_\omega^2), \quad (20)$$

where $s = 4$ for quarterly data. Figure 6 compares the resulting *one-step-ahead prediction errors*, that is,

$$e_t = y_t - \hat{y}_{t|t-1}, \quad (21)$$

which are expressed in the same units as the dependent variable (percentage points of the unemployment rate). Positive values indicate *under*-prediction (actual unemployment above the model forecast), while negative values indicate *over*-prediction.

Relative to the baseline model, adding seasonality absorbs systematic intra-year movements and substantially reduces residual dispersion (Table 3). In addition, the improvement in log likelihood and information criteria in Table 4 indicates that seasonal dynamics are an important feature of the Ticino unemployment series. Large residual spikes in specific periods (e.g. the Global Financial Crisis and the COVID-19 shock) highlight episodes where unemployment adjusts abruptly in ways that are not well captured by smooth trend and seasonal components alone.




Figure 6: UCM one-step-ahead residuals (RW level vs. RW level + seasonal(4))

Table 3: Residual diagnostics (one-step-ahead prediction errors)

	Non-seasonal UCM	Seasonal UCM
Non-missing one-step-ahead residuals (N)	84	81
Mean	-0.0242	-0.0283
Std. dev.	0.5448	0.2297
Variance	0.2968	0.05277
RMSE ($\sqrt{\mathbb{E}[e_t^2]}$)	0.5421	0.2300
Min	-1.3667	-0.5037
Max	0.9000	1.1750
Skewness	-0.5312	1.7980
Kurtosis	2.4396	11.0510
Variance reduction vs. non-seasonal	-	82.2%
RMSE reduction vs. non-seasonal	-	57.6%

Notes: Residuals are one-step-ahead prediction errors. The Seasonal UCM has fewer residuals (81 vs. 84) simply because it requires more initial observations to estimate the starting seasonal patterns before it can generate its first prediction. This is a standard initialization effect and does not reflect a change in the underlying dataset reported in Table 4.

Table 4: UCM specifications and in-sample fit (2005Q1–2026Q1)

	RW level	RW level + seasonal(4)
Sample (quarters)	2005Q1–2026Q1	2005Q1–2026Q1
N (estimation)	85	85
Log likelihood (ℓ)	-68.6729	-0.6159
Akaike Information Criterion (AIC) ($-2\ell + 2k$)	139.3457	5.2318
Bayesian Information Criterion (BIC) ($-2\ell + k \ln N$)	141.7884	10.1171
$\widehat{\text{Var}}(\eta_t)$ (level shock)	0.2938	0.04011
$\widehat{\text{Var}}(\omega_t)$ (seasonal shock)	-	0.000500

Notes: k denotes the number of estimated parameters used in the information criteria. In these specifications, the estimated parameters are the state-disturbance variances (one for the baseline; two when adding seasonal(4)). Negative log-likelihood values are not computational errors; in maximum-likelihood estimation the log-likelihood can take any real value depending on model fit and data scale.

4.5 Explanatory Power of Auxiliary Indicators

Core predictors. Following extensive diagnostic testing of correlation structures and UCM residuals across the full dataset, I retained the covariate groups demonstrating the highest predictive power. In what follows, *BSI* denotes the KOF Business Situation Indicator (a quarterly survey measure of current business conditions); *search intensity* denotes the Google Trends index of job-search and unemployment-related queries originating in the Canton; and *local level* denotes the stochastic trend component of the UCM—a random walk process tracking the evolving mean of the unemployment series. Due to publication delays, the effective lag for regional GDP is consistently larger than for high-frequency indicators such as exchange rates. Based on preliminary performance, I retain the following specifications for further testing: (below, t indicates contemporaneous values and $t - k$ indicates a lag of k quarters):

Table 5: UCM information sets retained for further testing.

Code	Covariates added to the UCM baseline (trend + seasonality)
a	Baseline (local level + quarterly seasonality)
b	CHF/EUR exchange rate (t)
c	CHF/EUR exchange rate ($t - 1$)
d	Google Trends search intensity (t)
e	Google Trends search intensity ($t - 1$)
f	CREA cantonal GDP – Ticino ($t - 1$)
g	CREA cantonal GDP – Ticino ($t - 2$)
h	Exchange rate (t) + search intensity (t)
i	Exchange rate (t) + CREA GDP ($t - 1$)
j	Exchange rate ($t - 1$) + CREA GDP ($t - 1$)
k	Exchange rate (t) + CREA GDP ($t - 2$)
l	Exchange rate ($t - 1$) + CREA GDP ($t - 2$)
m2	KOF employment expectations (t)
m3	KOF business situation indicator (BSI) (t)
m4	KOF Economic Barometer (t)
n2	CREA GDP ($t - 1$) + KOF employment expectations (t)
n3	CREA GDP ($t - 1$) + BSI (t)
n4	CREA GDP ($t - 1$) + KOF Barometer (t)
o2	KOF employment expectations ($t - 1$)
o3	BSI ($t - 1$)
o4	KOF Barometer ($t - 1$)
p2	CREA GDP ($t - 1$) + KOF employment expectations ($t - 1$)
p3	CREA GDP ($t - 1$) + BSI ($t - 1$)
p4	CREA GDP ($t - 1$) + KOF Barometer ($t - 1$)
q2	CREA GDP ($t - 1$) + EUR/CHF exchange rate (t) + KOF expectations ($t - 1$)

Continued on next page

Table 5 (*continued*)

Code	Covariates added to the UCM baseline
q3	CREA GDP ($t - 1$) + EUR/CHF exchange rate (t) + BSI ($t - 1$)
r2	CREA GDP ($t - 1$) + EUR/CHF exchange rate ($t, t - 1$) + KOF expectations ($t - 1$)
r3	CREA GDP ($t - 1$) + EUR/CHF exchange rate ($t, t - 1$) + BSI ($t - 1$)
s2	CREA GDP ($t - 1$) + EUR/CHF exchange rate (t) + KOF expectations (t)
s3	CREA GDP ($t - 1$) + EUR/CHF exchange rate (t) + BSI (t)
t2	CREA GDP ($t - 1$) + EUR/CHF exchange rate ($t - 1$) + KOF expectations ($t - 1$)
t3	CREA GDP ($t - 1$) + EUR/CHF exchange rate ($t - 1$) + BSI ($t - 1$)
v	Exchange rate (t) + search intensity (t) + CREA GDP ($t - 2$)
x2	Annual GDP proxy (t) + KOF expectations (t)
x3	Annual GDP proxy (t) + BSI (t)
y	Exchange rate (t) + search intensity (t) + annual GDP proxy (t)
z	Exchange rate ($t - 1$) + search intensity (t) + annual GDP proxy (t)

High-dimensional specifications combining many regressors frequently fail to converge or yield unfeasible standard errors. Likewise, the inclusion of tourism-specific variables in the UCM does not significantly improve predictive power for unemployment, as the lagged coefficients lack statistical significance. Consequently, both the high-dimensional multi-disaggregation specifications and the tourism-related variable specifications are excluded from the final forecast comparisons.

4.5.1 Real-time Analysis

This subsection presents an in-sample contemporaneous analysis, where the explanatory variables are included in the model at the same quarter (t) at which a specific unemployment rate is observed (e.g., the unemployment rate of 2025-Q1 is observed together with the currency exchange rate data of the same quarter). This regression quantifies how effectively x_t explains the variation in y_t . While this highlights the structural co-movement between indicators and unemployment, it does not yet constitute a genuine forecasting test, as future economic variables are not observed yet at the moment of a forecast. A rigorous real-time evaluation accounting for data availability is therefore performed in the subsequent section.

Model comparison. UCM specifications involve latent components that are never directly observed, so they are cast in state-space form and estimated via the Kalman filter. The filter propagates state estimates forward each period and revises them as each new observation arrives; the resulting one-step-ahead prediction errors are used to evaluate the log-likelihood, which is then maximized over the unknown parameters. Fit is assessed using the log likelihood, Wald tests for the significance of β , and the distribution of one-step-ahead prediction errors. For comparability across specifications, I report RMSE computed from one-step-ahead residuals and I summarize predictive accuracy

using residual mean across predictions in training, as well as:

$$\text{MAE} = \frac{1}{N_{\text{res}}} \sum_{t=1}^{N_{\text{res}}} |\hat{e}_t|, \quad (22)$$

$$\text{RMSE} = \sqrt{\frac{1}{N_{\text{res}}} \sum_{t=1}^{N_{\text{res}}} \hat{e}_t^2}. \quad (23)$$

$$\text{where } e_t = y_t - \hat{y}_{t|t-1}. \quad (24)$$

Here, *Mean Absolute Error* (MAE) measures the average absolute size of the one-step-ahead prediction errors, while *Root Mean Squared Error* (RMSE) gives relatively greater weight to larger forecast misses because errors enter in squared form before averaging. In what follows, lower values of either statistic indicate better predictive performance.

Models are in fact ranked primarily by RMSE and secondarily by MAE. Table 6 reports the full ranking for all specifications. Because some indicators start later in the sample (and lagging removes initial observations), the number of usable one-step-ahead residuals differs across models; comparisons are therefore most informative within the same data-availability window.

Table 6: UCM residual comparison (real-time exercise)

Code	Variables included (lags in quarters)	MAE	RMSE	Mean
s2	GDP (CREA, Ticino) ($t-3$), EUR/CHF ($t-2$), KOF exp. ($t-2$)	0.130	0.158	-0.025
x3	avgGDPyear ($t-2$), BSI ($t-2$)	0.142	0.182	-0.028
x2	avgGDPyear ($t-2$), KOF exp. ($t-2$)	0.145	0.186	-0.013
n3	GDP ($t-3$), BSI ($t-2$)	0.142	0.187	-0.037
m3	BSI ($t-2$)	0.146	0.190	-0.036
n2	GDP (CREA, Ticino) ($t-3$), KOF exp. ($t-2$)	0.148	0.190	-0.019
m2	KOF exp. ($t-2$)	0.149	0.191	-0.019
r2	GDP ($t-3$), EUR/CHF ($t-2, t-3$), KOF exp. ($t-3$)	0.148	0.197	-0.012
q2	GDP ($t-3$), EUR/CHF ($t-2$), KOF exp. ($t-3$)	0.148	0.197	-0.012
p2	GDP ($t-3$), KOF exp. ($t-3$)	0.147	0.198	-0.019
t2	GDP ($t-3$), EUR/CHF ($t-3$), KOF exp. ($t-3$)	0.148	0.199	-0.017
s3	GDP ($t-3$), EUR/CHF ($t-2$), BSI ($t-2$)	0.148	0.203	-0.026
q3	GDP ($t-3$), EUR/CHF ($t-2$), BSI ($t-3$)	0.148	0.203	-0.026
r3	GDP ($t-3$), EUR/CHF ($t-2, t-3$), BSI ($t-3$)	0.148	0.203	-0.025
t3	GDP ($t-3$), EUR/CHF ($t-3$), BSI ($t-3$)	0.150	0.205	-0.033
p3	GDP ($t-3$), BSI ($t-3$)	0.150	0.206	-0.040
p4	GDP ($t-3$), KOF barometer ($t-3$)	0.157	0.212	-0.023
i	CHF/EUR ($t-2$), GDP ($t-3$)	0.154	0.214	-0.015
j	CHF/EUR ($t-3$), GDP ($t-3$)	0.155	0.216	-0.023
f	GDP (CREA, Ticino) ($t-3$)	0.155	0.216	-0.022
n4	GDP ($t-3$), KOF barometer ($t-2$)	0.156	0.217	-0.022
m4	KOF barometer ($t-2$)	0.160	0.219	-0.022
v	CHF/EUR ($t-2$), Google Trends ($t-2$), GDP ($t-4$)	0.154	0.220	-0.013
z	CHF/EUR ($t-3$), Google Trends ($t-2$), avgGDPyear ($t-2$)	0.154	0.220	-0.018
y	CHF/EUR ($t-2$), Google Trends ($t-2$), avgGDPyear ($t-2$)	0.153	0.220	-0.014
h	CHF/EUR ($t-2$), Google Trends ($t-2$)	0.155	0.221	-0.014
d	Google Trends ($t-2$)	0.156	0.221	-0.016
o2	KOF exp. ($t-3$)	0.156	0.222	-0.020
k	CHF/EUR ($t-2$), GDP ($t-4$)	0.158	0.223	-0.020
l	CHF/EUR ($t-3$), GDP ($t-4$)	0.157	0.223	-0.020
g	GDP (CREA, Ticino) ($t-4$)	0.158	0.223	-0.022
b	CHF/EUR ($t-2$)	0.159	0.224	-0.020
o4	KOF barometer ($t-3$)	0.159	0.224	-0.022
c	CHF/EUR ($t-3$)	0.159	0.225	-0.022
a	Baseline (local level + seasonality only)	0.159	0.225	-0.022
e	Google Trends ($t-3$)	0.159	0.227	-0.021
o3	BSI ($t-3$)	0.170	0.241	-0.041

Notes. Rows are ordered by RMSE (primary) and MAE (secondary). Unless explicitly indicated, predictors are lagged two quarters ($t-2$). GDP (CREA, Ticino) is included at lag three by convention ($t-3$) (and lag four when stated).

Residual comparison during the COVID shock (2020–2021). Still within the contemporaneous (real-time) exercise reported in Table 6, Figure 7 zooms into 2020Q1–2021Q4 and compares one-step-ahead residuals from a set of seasonal UCM specifications where indicators enter at the same quarter as unemployment (t). Relative to the baseline model, including economic indicators such as GDP, KOF employment expectations, and the Business Situation Indicator tends to reduce residual dispersion around the shock quarters, indicating an improved in-sample fit in this extreme period. An analogous comparison using strictly lagged predictors (enforcing real-time timing constraints) follows in Section 4.5.2.




Figure 7: Seasonal UCM one-step-ahead residuals: COVID window (2020–2021)

4.5.2 Seasonal UCM with lagged predictors

The purpose of this screening step is to identify a small set of robust indicator sets that improve short-horizon predictability under real-time timing constraints. For each specification, I compute one-step-ahead prediction errors. Because the seasonal component has period 4, the one-step-ahead residual series becomes available only after the initial seasonal states are formed, implying that the number of non-missing residuals is smaller than the estimation sample. Models are again ranked primarily by RMSE and secondarily by MAE.

Across lagged specifications, the strongest and most stable predictor is KOF employment expectations: its coefficient is negative and statistically significant at two- and three-quarter lags, consistent with improving expectations anticipating lower unemployment. Google Trends also shows predictive content at medium lags, with higher search intensity preceding higher unemployment. By contrast, GDP and exchange-rate terms are often insignificant once latent level and seasonality are controlled for, although GDP can contribute modestly when combined with survey indicators.

Across lagged specifications, the strongest predictor is KOF employment expectations: its coefficient is negative and statistically significant at both two- and three-quarter lags,

consistent with improving expectations anticipating lower unemployment. Google Trends at two-quarter lag is positive and significant, suggesting search intensity rises ahead of higher unemployment. GDP and exchange-rate terms are at times insignificant once the latent level and seasonality are controlled for, although including GDP alongside KOF_exp slightly improves residual RMSE.

Table 7: UCM residual comparison (a-z), lagged-predictor exercise

Code	Variables included (lags in quarters)	MAE	RMSE	Mean
s2	GDP (CREA, Ticino) ($t-3$), EUR/CHF ($t-2$), KOF exp. ($t-2$)	0.150	0.217	-0.022
n2	GDP (CREA, Ticino) ($t-3$), KOF exp. ($t-2$)	0.150	0.218	-0.019
r2	GDP ($t-3$), EUR/CHF ($t-2, t-3$), KOF exp. ($t-3$)	0.152	0.218	-0.025
q2	GDP ($t-3$), EUR/CHF ($t-2$), KOF exp. ($t-3$)	0.152	0.218	-0.025
z	CHF/EUR ($t-3$), Google Trends ($t-2$), avgGDPyear ($t-2$)	0.152	0.218	-0.037
v	CHF/EUR ($t-2$), Google Trends ($t-2$), GDP ($t-4$)	0.152	0.219	-0.037
h	CHF/EUR ($t-2$), Google Trends ($t-2$)	0.153	0.219	-0.037
y	CHF/EUR ($t-2$), Google Trends ($t-2$), avgGDPyear ($t-2$)	0.152	0.219	-0.037
d	Google Trends ($t-2$)	0.153	0.219	-0.034
t2	GDP ($t-3$), EUR/CHF ($t-3$), KOF exp. ($t-3$)	0.152	0.219	-0.021
p2	GDP ($t-3$), KOF exp. ($t-3$)	0.152	0.219	-0.019
x2	avgGDPyear ($t-2$), KOF exp. ($t-2$)	0.151	0.220	-0.019
j	CHF/EUR ($t-3$), GDP ($t-3$)	0.157	0.221	-0.025
m2	KOF exp. ($t-2$)	0.152	0.221	-0.019
i	CHF/EUR ($t-2$), GDP ($t-3$)	0.157	0.221	-0.024
p4	GDP ($t-3$), KOF barometer ($t-3$)	0.157	0.221	-0.022
n4	GDP ($t-3$), KOF barometer ($t-2$)	0.156	0.221	-0.022
f	GDP (CREA, Ticino) ($t-3$)	0.156	0.221	-0.022
o2	KOF exp. ($t-3$)	0.153	0.221	-0.018
o4	KOF barometer ($t-3$)	0.158	0.221	-0.022
c	CHF/EUR ($t-3$)	0.158	0.223	-0.027
l	CHF/EUR ($t-3$), GDP ($t-4$)	0.158	0.223	-0.027
r3	GDP ($t-3$), EUR/CHF ($t-2, t-3$), BSI ($t-3$)	0.151	0.223	-0.037
q3	GDP ($t-3$), EUR/CHF ($t-2$), BSI ($t-3$)	0.152	0.224	-0.038
t3	GDP ($t-3$), EUR/CHF ($t-3$), BSI ($t-3$)	0.152	0.224	-0.036
p3	GDP ($t-3$), BSI ($t-3$)	0.152	0.224	-0.037
k	CHF/EUR ($t-2$), GDP ($t-4$)	0.159	0.224	-0.024
b	CHF/EUR ($t-2$)	0.159	0.224	-0.024
m4	KOF barometer ($t-2$)	0.158	0.224	-0.021
a	Baseline (local level + seasonality only)	0.159	0.225	-0.022
g	GDP (CREA, Ticino) ($t-4$)	0.159	0.225	-0.022
o3	BSI ($t-3$)	0.151	0.225	-0.036
e	Google Trends ($t-3$)	0.162	0.228	-0.023
s3	GDP ($t-3$), EUR/CHF ($t-2$), BSI ($t-2$)	0.169	0.239	-0.040
n3	GDP ($t-3$), BSI ($t-2$)	0.169	0.239	-0.040
x3	avgGDPyear ($t-2$), BSI ($t-2$)	0.171	0.243	-0.040
m3	BSI ($t-2$)	0.172	0.243	-0.040

Notes. Rows are ordered by RMSE (primary) and MAE (secondary). Unless explicitly indicated, predictors are lagged two quarters ($t-2$). GDP (CREA, Ticino) is included at lag three by convention ($t-3$) (and lag four when stated).

Residual comparison during the COVID shock (2020–2021). Figure 8 zooms into 2020Q1–2021Q4 and compares one-step-ahead residuals from the top-ranked specifications in Table 7. The lagged-predictor designs reduce residual dispersion only slightly relative to the baseline seasonal UCM: improvements are smaller than in the contemporaneous fit exercise because the general-lag timing restricts predictors to information that would have been available at the forecast origin. Nonetheless, models that include KOF employment expectations (alone or with CREA GDP) still tend to track the direction of the COVID-era swings relatively more closely than the univariate model.

Table 8: RMSE comparison: contemporaneous association vs real-time timing

	Timing convention	Best RMSE
Same-quarter indicators	Contemporaneous (x_t explains y_t)	0.187
Lagged indicators	Forecasting (lagged x_{t-k} explain y_t)	0.217

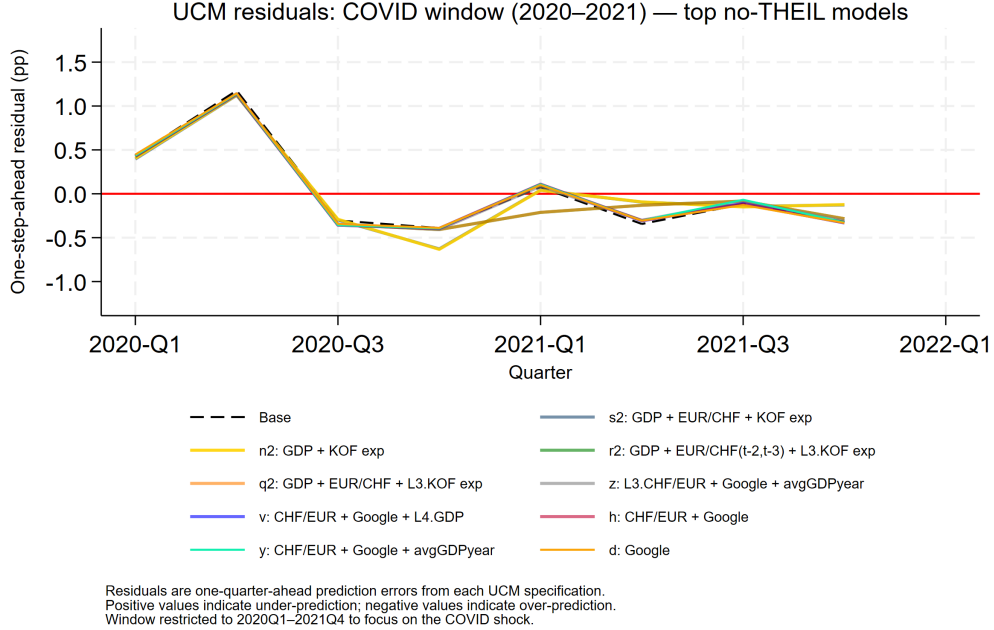


Figure 8: Seasonal UCM one-step-ahead residuals: COVID window (2020–2021), lagged predictors.

Table 8 summarizes this gap: the best specification achieves an RMSE of 0.187 under contemporaneous timing versus 0.217 under the lagged timing constraint, confirming that a share of the explanatory power of auxiliary indicators is contemporaneous rather than purely predictive. Because the real-time (general-lag) exercise shows that auxiliary indicators add only modest incremental predictive power for future unemployment, I next complement the UCM evidence with an explicit time-series forecasting approach. Specifically, I estimate both (seasonal) ARIMA models that rely purely on the internal dynamics of unemployment, and (seasonal) ARIMA-X models that additionally incorporate the best-performing lagged indicators as exogenous regressors. This comparison clarifies whether forecasting gains come primarily from the autoregressive-seasonal structure of unemployment itself, or whether there remains meaningful extra information in lagged macro/sentiment variables once timing constraints are enforced.

Selection of top-10 predictor sets for SARIMAX. The following ten specifications represent the highest-performing covariate combinations identified during the seasonal UCM analysis. These subsets are consequently retained as the exogenous regressors for the subsequent SARIMAX framework:

1. **s2:** GDP (lag) + Currency + KOF exp.

2. **n2**: GDP (lag) + KOF exp.
3. **r2**: GDP (lag) + Currency (current & lag) + KOF exp. (lag)
4. **q2**: GDP (lag) + Currency + KOF exp. (lag)
5. **z**: Currency (lag) + Google Trends + Avg. Yearly GDP
6. **v**: GDP (2-quarter lag) + Currency + Google Trends
7. **h**: Currency + Google Trends
8. **y**: Currency + Google Trends + Avg. Yearly GDP
9. **d**: Google Trends
10. **t2**: GDP (lag) + Currency (lag) + KOF exp. (lag)

4.6 Selecting SARIMA dynamics (univariate benchmarking)

I estimate a grid of univariate SARIMA models for quarterly unemployment (data until 2026Q1) and select a small set of well-performing *time-series backbones* to carry forward into the SARIMAX stage. This step isolates the best internal unemployment dynamics before introducing exogenous predictors.

Model selection is based on the log likelihood and standard information criteria. The log-likelihood (LL) measures in-sample fit, while the Akaike Information Criterion (AIC) and Bayesian Information Criterion (BIC) trade off fit against complexity:

$$\text{AIC} = -2 \ln L + 2k, \quad \text{BIC} = -2 \ln L + k \ln N,$$

where k is the number of estimated parameters and N is the effective sample size. Lower AIC/BIC indicates a better fit–parsimony trade-off, with BIC imposing a stronger penalty for additional parameters. Importantly, when the data transformation changes (specifically, whether or not the series is first-differenced), the likelihood is computed on a different effective sample and scale, so AIC and BIC values cannot be directly compared across the two groups. AIC/BIC comparisons are therefore meaningful only *within* a given transformation and sample, which is why models with $d = 0$ and models with $d = 1$ are presented in separate tables below.

Table 9: Univariate SARIMA screening in levels ($d = 0$), quarterly unemployment (sample 2004Q1–2026Q1)

Model $(p, d, q)(P, D, Q)_4$	AIC	BIC	LL
SARIMA(2, 0, 1)(1, 0, 1)₄	-9.4456	7.9748	11.7228
SARIMA(1, 0, 1)(1, 0, 1)₄	-7.1544	7.7774	9.5772
<i>SARIMA</i> (2, 0, 1)(1, 0, 0) ₄	11.9512	26.8831	0.0244
<i>SARIMA</i> (1, 0, 1)(1, 0, 0) ₄	15.8371	28.2803	-2.9185
<i>SARIMA</i> (2, 0, 1)(0, 0, 1) ₄	72.4850	87.4168	-30.2425
<i>SARIMA</i> (1, 0, 1)(0, 0, 1) ₄	71.2444	83.6876	-30.6222
<i>SARIMA</i> (2, 0, 1)(0, 0, 0) ₄	129.3365	141.7796	-59.6682
<i>SARIMA</i> (1, 0, 1)(0, 0, 0) ₄	127.4522	137.4068	-59.7261

Notes: All models assume quarterly seasonality ($s = 4$). Lower AIC/BIC indicates a better fit–parsimony trade-off; AIC/BIC comparisons are meaningful within this table because all specifications are estimated on the same sample and transformation (levels, $d = 0$). Negative log-likelihood values can arise for poorly fitting specifications and are not computational errors.

Table 10: Univariate SARIMA screening in first differences ($d = 1$), quarterly unemployment (sample 2004Q2–2026Q1)

Model $(p, d, q)(P, D, Q)_4$	AIC	BIC	LL
SARIMA(2, 1, 1)(1, 0, 1)₄	-6.4840	10.8574	10.2420
SARIMA(1, 1, 1)(1, 0, 1)₄	-8.2299	6.6342	10.1149
<i>SARIMA</i> (2, 1, 1)(1, 0, 0) ₄	21.7351	36.5992	-4.8676
<i>SARIMA</i> (1, 1, 1)(1, 0, 0) ₄	19.7565	32.1432	-4.8783
<i>SARIMA</i> (2, 1, 1)(0, 0, 1) ₄	43.7534	58.6174	-15.8767
<i>SARIMA</i> (2, 1, 1)(0, 0, 0) ₄	50.8428	63.2295	-20.4214
<i>SARIMA</i> (1, 1, 1)(0, 0, 1) ₄	82.1810	94.5677	-36.0905
<i>SARIMA</i> (1, 1, 1)(0, 0, 0) ₄	137.5174	147.4267	-64.7587

Notes: All models assume quarterly seasonality ($s = 4$). AIC/BIC comparisons should be interpreted within this table (first differences, $d = 1$), since the transformation changes the likelihood scale and effective sample relative to the levels case. Negative log-likelihood values can arise for poorly fitting specifications and are not computational errors.

Levels models ($d = 0$). Table 9 reports the main candidates estimated on 2004Q1–2026Q1. Seasonal dynamics are crucial: specifications with seasonal ARMA terms dominate non-seasonal models in both information criteria and log likelihood. Among the tested models, AIC favors *SARIMA*(2, 0, 1)(1, 0, 1)₄, while BIC slightly prefers the more parsimonious *SARIMA*(1, 0, 1)(1, 0, 1)₄.

Differenced models ($d = 1$). Table 10 reports the analogous grid estimated on 2004Q2–2026Q1. Within this class, *SARIMA*(1, 1, 1)(1, 0, 1)₄ is preferred by AIC/BIC, while *SARIMA*(2, 1, 1)(1, 0, 1)₄ achieves the best log-likelihood performance and provides a slightly richer autoregressive structure.

Backbones carried into the SARIMAX stage. To keep the subsequent SARIMAX comparison transparent while allowing for alternative treatments of persistence, I there-

fore carry forward such four seasonal *backbones* (all with quarterly seasonality $s = 4$ and seasonal structure $(1, 0, 1)_4$):

$$SARIMA(2, 0, 1)(1, 0, 1)_4, \quad SARIMA(1, 0, 1)(1, 0, 1)_4,$$

$$SARIMA(1, 1, 1)(1, 0, 1)_4, \quad SARIMA(2, 1, 1)(1, 0, 1)_4.$$

These candidates capture the dominant seasonal dependence while spanning both levels ($d = 0$) and differenced ($d = 1$) specifications. Final model choice is then based on real-time forecasting performance once exogenous predictors are introduced under feasible timing conventions.

Historical benchmarking. Tables 11 and 12 report SARIMA screenings conducted on data observed through Q4 2023 and Q1 2024, respectively. The same quarterly seasonal structure $(1, 0, 1)_4$ already emerged as dominant, and all four backbones retained in the current SARIMAX exercise were already among the best-performing candidates. This stability across years supports using these dynamics as the benchmark backbones in the updated (2026) SARIMAX comparison.

I Trimester 2024 Results								
p	d	q	P	D	Q	AIC	BIC	LL
1	1	1	1	0	1	1.585549	15.877709	5.2072255
2	1	1	1	0	1	3.3906005	20.064787	5.3046997
2	0	1	1	0	1	1.8258763	18.500063	6.0870619
1	0	1	1	0	1	3.8237786	18.115938	4.0881107
1	1	1	1	0	0	27.064285	38.974419	-8.5321427
2	1	1	1	0	0	29.041917	43.334076	-8.5209583
2	0	1	1	0	0	20.111725	34.403885	-4.0558624
1	0	1	1	0	0	23.584379	35.494512	-6.7921893
2	1	1	0	0	1	46.720556	61.012715	-17.360278
1	1	1	0	0	1	79.766105	91.676238	-34.883052
2	0	1	0	0	1	70.703099	84.995259	-29.351549
1	0	1	0	0	1	69.351397	81.26153	-29.675698
2	1	1	0	0	0	51.762734	63.672868	-20.881367
1	1	1	0	0	0	128.98609	138.51419	-60.493044

Table 11: Comparison of SARIMA models (Q1 2024).

Notes: Negative log-likelihood values can arise for poorly fitting specifications and are not computational errors.

IV Trimester 2023 Results								
p	d	q	P	D	Q	AIC	BIC	LL
1	1	1	1	0	1	1.7605462	15.977233	5.1197269
2	1	1	1	0	1	3.635921	20.222056	5.1820395
2	0	1	1	0	1	1.8258763	18.500063	6.0870619
1	0	1	1	0	1	3.8237786	18.115938	4.0881107
2	1	1	0	0	0	50.513377	62.360617	-20.256689
1	1	1	0	0	0	127.73613	137.21392	-59.868064
2	1	1	1	0	0	28.389297	42.605984	-8.1946486
1	1	1	1	0	0	26.395354	38.242594	-8.1976771
2	0	1	1	0	0	20.111725	34.403885	-4.0558624
1	0	1	1	0	0	23.584379	35.494512	-6.7921893
2	1	1	0	0	1	45.868508	60.085195	-16.934254
1	1	1	0	0	1	78.5495	90.396739	-34.27475
2	0	1	0	0	1	70.703099	84.995259	-29.351549
1	0	1	0	0	1	69.351397	81.26153	-29.675698
2	0	1	0	0	1	120.17997	132.09011	-55.089987
1	0	1	0	0	1	118.19737	127.72547	-55.098683

Table 12: Comparison of SARIMA models (Q4 2023).

Notes: Negative log-likelihood values can arise for poorly fitting specifications and are not computational errors.

5 SARIMAX Modeling

5.1 Evaluation of SARIMAX candidates (lagged predictors)

The SARIMA results above identify strong univariate seasonal dynamics in quarterly unemployment. I now test whether the best-performing lagged explanatory variables in UCM add forecasting content once those dynamics are controlled for. As described beforehand, I estimate seasonal ARIMA models with exogenous regressors (SARIMAX), using as candidate regressors the best-performing lagged predictor sets identified in the UCM screening stage (Section 4.5.2). All predictors are lagged to respect real-time availability. A lag of two quarters is the default for fast-moving series (exchange rates, Google Trends, survey indicators). Canton GDP is included at longer lags (typically 3–4 quarters) to reflect publication delays. The ten retained covariate blocks—S2, N2, R2, Q2, T2, Z, V, H, Y, and D—are defined and listed in Section 4.5.2.

5.1.1 Model comparison: information criteria and log-likelihood

Models are estimated by maximum likelihood. Again, fit is evaluated using log-likelihood (LL), AIC, and BIC. For transparency, Table 13 reports the full candidate set, and I then summarize the best-performing specifications under each criterion.

Backbone	Info. set	N	AIC	BIC	LL
101_101	Univariate	89	-7.154433	7.777385	9.577216
101_101	D	87	-2.706892	14.554465	8.353446
101_101	H	87	-0.945496	18.781769	8.472748
101_101	N2	89	-12.200857	7.708234	14.100429
101_101	Q2	89	-6.784898	15.612829	12.392449
101_101	R2	89	-4.886855	19.999510	12.443427
101_101	S2	89	-10.241136	12.156592	14.120567
101_101	T2	89	-6.767018	15.630710	12.383509
101_101	V	87	0.776557	22.969730	8.611721
101_101	Y	87	1.020947	23.214119	8.489527
101_101	Z	87	1.183703	23.376877	8.408149
111_101	Univariate	88	-8.229856	6.634165	10.114928
111_101	D	86	-3.989383	13.191048	8.994692
111_101	H	86	-2.363563	17.271215	9.181782
111_101	N2	88	-12.816584	7.002110	14.408292
111_101	Q2	88	-6.629465	15.666567	12.314733
111_101	R2	88	-5.315053	19.458315	12.657526
111_101	S2	88	-11.028269	11.267762	14.514134
111_101	T2	88	-5.748455	16.547577	11.874228
111_101	V	86	-0.381569	21.707558	9.190784
111_101	Y	86	-0.504863	21.584263	9.252432
111_101	Z	86	-1.290781	20.798346	9.645391
201_101	Univariate	89	-9.445645	7.974810	11.722822
201_101	D	87	-5.292984	14.434281	10.646492
201_101	H	87	-3.293147	18.900026	10.646574
201_101	N2	89	-15.064470	7.333257	16.532236
201_101	Q2	89	-4.786613	20.099751	12.393306
201_101	R2	89	-2.898146	24.476854	12.449073
201_101	S2	89	-13.076460	11.809904	16.538231
201_101	T2	89	-4.811019	20.075344	12.405509
201_101	V	87	-1.307430	23.351650	10.653715
201_101	Y	87	-1.441598	23.217484	10.720798
201_101	Z	87	-1.646236	23.012846	10.823118
211_101	Univariate	88	-6.483970	10.857388	10.241985
211_101	D	86	41.481979	61.116760	-12.740991
211_101	H	86	-0.529998	21.559128	9.264999
211_101	N2	88	-11.109937	11.186094	14.554968
211_101	Q2	88	-3.094770	21.678598	11.547385
211_101	R2	88	-1.386293	25.864412	11.693147
211_101	S2	88	-9.482477	15.290892	14.741239
211_101	T2	88	-4.272770	20.500597	12.136385
211_101	V	86	42.102001	66.645477	-11.051002
211_101	Y	86	1.336762	25.880236	9.331619
211_101	Z	86	45.255356	69.798828	-12.627679

Table 13: Information criteria for SARIMA (univariate) and SARIMAX candidates (data through 2026Q1). Rows with missing entries correspond to specifications that did not converge or were not retained.

Notes: Negative log-likelihood values can arise for poorly fitting specifications and are not computational errors; in maximum-likelihood estimation the log-likelihood can take any real value.

From criterion-specific rankings to an operational shortlist. Tables 14, 15 and 16 report the fifteen best models by each criterion as a compact summary of the ranking structure. For the final operational selection, however, I widen the admissible pool to all models appearing within the Top-15 of at least one criterion. The size of 15 is chosen to balance fit quality across criteria (AIC, BIC, LL reward different trade-offs) with diversity across information sets, without inflating the pool to the point where poorly fitting specifications enter.

Backbone	Info. set	N	AIC	BIC
201_101	N2	89	-15.064470	7.333257
201_101	S2	89	-13.076460	11.809904
111_101	N2	88	-12.816584	7.002110
101_101	N2	89	-12.200857	7.708234
211_101	N2	88	-11.109937	11.186094
111_101	S2	88	-11.028269	11.267762
101_101	S2	89	-10.241136	12.156592
211_101	S2	88	-9.482477	15.290892
201_101	Univariate	89	-9.445645	7.974810
111_101	Univariate	88	-8.229856	6.634165
101_101	Univariate	89	-7.154433	7.777385
101_101	Q2	89	-6.784898	15.612829
101_101	T2	89	-6.767018	15.630710
111_101	Q2	88	-6.629465	15.666567
211_101	Univariate	88	-6.483970	10.857388

Table 14: Top-15 models by AIC (SARIMA and SARIMAX; data through 2026Q1).

Backbone	Info. set	N	AIC	BIC
111_101	Univariate	88	-8.229856	6.634165
111_101	N2	88	-12.816584	7.002110
201_101	N2	89	-15.064470	7.333257
101_101	N2	89	-12.200857	7.708234
101_101	Univariate	89	-7.154433	7.777385
201_101	Univariate	89	-9.445645	7.974810
211_101	Univariate	88	-6.483970	10.857388
211_101	N2	88	-11.109937	11.186094
111_101	S2	88	-11.028269	11.267762
201_101	S2	89	-13.076460	11.809904
101_101	S2	89	-10.241136	12.156592
111_101	D	86	-3.989383	13.191048
201_101	D	87	-5.292984	14.434281
211_101	S2	88	-9.482477	15.290892
101_101	Q2	89	-6.784898	15.612829

Table 15: Top-15 models by BIC (SARIMA and SARIMAX; data through 2026Q1).

Backbone	Info. set	N	LL
201_101	S2	89	16.538231
201_101	N2	89	16.532236
211_101	S2	88	14.741239
211_101	N2	88	14.554968
111_101	S2	88	14.514134
111_101	N2	88	14.408292
101_101	S2	89	14.120567
101_101	N2	89	14.100429
111_101	R2	88	12.657526
201_101	R2	89	12.449073
101_101	R2	89	12.443427
201_101	T2	89	12.405509
201_101	Q2	89	12.393306
101_101	Q2	89	12.392449
101_101	T2	89	12.383509

Table 16: Top-15 models by Log-Likelihood (SARIMA and SARIMAX; data through 2026Q1).

Some patterns emerge across criteria. First, the *N2* and *S2* information sets, which combine CREA GDP and KOF employment expectations (with or without the EUR/CHF exchange rate), dominate the best-fitting specifications under AIC and LL, confirming that survey-based expectations contain forecasting-relevant information beyond unemployment’s own seasonal dynamics. Second, BIC places more weight on parsimony and, in this updated run, assigns the strongest positions to univariate SARIMA benchmarks, with *N2* remaining the main multivariate information set still present near the top of the ranking. Third, Google-Trends specifications appear in the BIC top-fifteen (ranks 12–13) but not in the AIC or LL rankings; the operational shortlist retains one *D* model to preserve diversity across information sets. Finally, level-estimated backbones do not systematically outperform first-differenced backbones; rather, accuracy depends on the interaction between the backbone and the chosen exogenous information set.

Because AIC, BIC, and LL embody different trade-offs between fit and parsimony, I do not select a single model based on one metric alone. This ensures that the final shortlist does not depend mechanically on one particular penalty choice. Instead, I first widen the admissible pool to models that appear within the Top-15 of at least one criterion, and then retain a fixed Top-15 shortlist that combines strong criterion rankings with diversity across information sets. Using the information-set codes in Table 17, the retained operational shortlist spans the following information sets: UNI, N2, S2, R2, Q2, T2, and D.

The candidate specifications considered in the SARIMAX exercise can be summarised as follows:

- **Univariate SARIMA (UNI):** baseline backbone with no exogenous regressors (pure SARIMA dynamics).
- **D (search intensity):** Google Trends index.
- **N2 (regional output + expectations):** CREA GDP for Ticino and KOF employment expectations.
- **S2 (regional output + exchange rate + expectations):** CREA GDP for Ticino, EUR/CHF exchange rate, and KOF employment expectations.
- **Q2 (longer-lag expectations specification):** CREA GDP for Ticino, EUR/CHF exchange rate, and KOF employment expectations at a longer lag.

Model	LL	AIC	BIC	Avg. Rank
201_101-N2	16.532236	-15.064470	7.333257	2.000
111_101-N2	14.408292	-12.816584	7.002110	3.667
201_101-S2	16.538231	-13.076460	11.809904	4.000
101_101-N2	14.100429	-12.200857	7.708234	5.333
211_101-N2	14.554968	-11.109937	11.186094	5.333
111_101-S2	14.514134	-11.028269	11.267762	6.333
211_101-S2	14.741239	-9.482477	15.290892	7.667
101_101-S2	14.120567	-10.241136	12.156592	8.000
201_101-Univariate	11.722822	-9.445645	7.974810	8.667
111_101-Univariate	10.114928	-8.229856	6.634165	8.667
211_101-Univariate	10.241985	-6.483970	10.857388	11.000
101_101-Q2	12.392449	-6.784898	15.612829	11.333
101_101-T2	12.383509	-6.767018	15.630710	12.333
111_101-R2	12.657526	-5.315053	19.458315	12.667
201_101-D	10.646492	-5.292984	14.434281	13.000

Table 17: Operational Top-15 shortlist used in the ensemble forecast. The set is selected from the union of models appearing in the Top-15 of at least one criterion, balancing criterion performance and diversity across information sets. Avg. Rank is the mean of the three individual criterion ranks and is reported only as a compact visual summary. The retained information sets are: univariate SARIMA dynamics (UNI), CREA GDP with KOF employment expectations (N2), CREA GDP with EUR/CHF and KOF employment expectations (S2), CREA GDP with EUR/CHF at multiple lags and KOF employment expectations (R2), CREA GDP with EUR/CHF and longer-lag KOF employment expectations (Q2), the fully lagged CREA GDP–EUR/CHF–KOF specification (T2), and the Google Trends-only specification (D). All scores refer to the latest screening with data through 2026Q1.

- **R2 (richer exchange-rate specification):** CREA GDP for Ticino, EUR/CHF exchange rate at multiple lags, and KOF employment expectations.
- **T2 (fully lagged specification):** CREA GDP for Ticino, EUR/CHF exchange rate, and KOF employment expectations, all at the longer lag used in the SARIMAX screening.

Timing is expressed relative to the target quarter t : (t) denotes contemporaneous values, while ($t - k$) denotes a lag of k quarters.

5.2 Forecasting

This section reports the forecasting and validation results. Given that multiple models perform similarly well in-sample, I ultimately rely on an *ensemble* based on the average prediction across the fixed Top-15 shortlist. Averaging across well-performing models reduces sensitivity to any single specification while preserving a robust forecast under model uncertainty. For quarterly studies of IRE, the shortlist itself is kept fixed while the operational forecast is rolled forward as new observations become available.

Figure 9 reports the resulting average forecast path from the last observed quarter onward. The operational horizon spans the two quarters immediately following the last observed value.

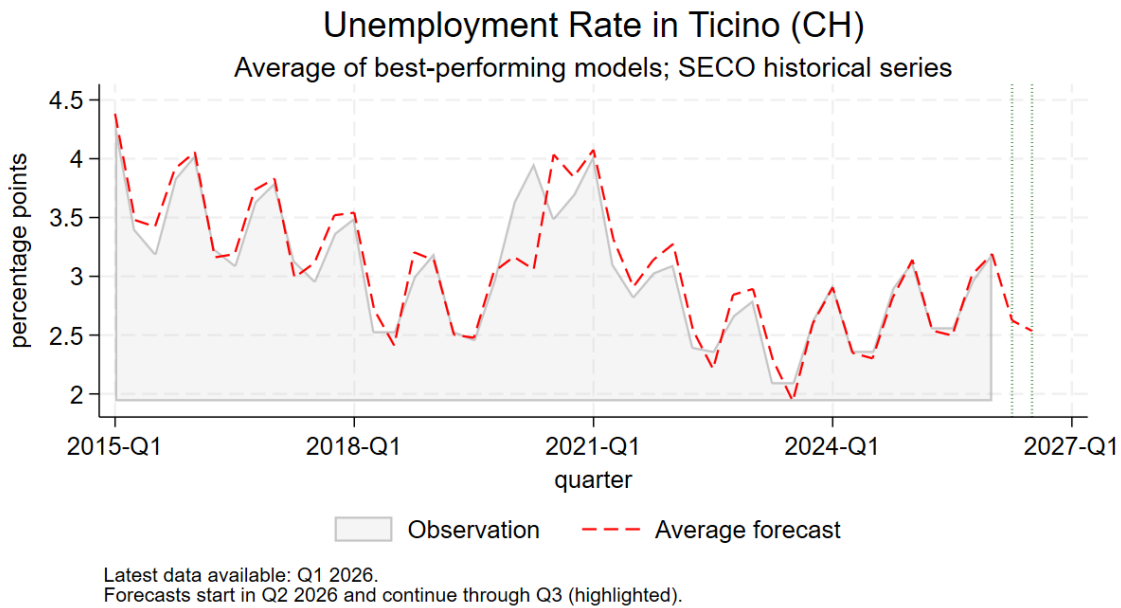


Figure 9: Average of Top-15 forecasts (main specification).

Let $\hat{y}_{m,t}$ denote the forecast from model $m = 1, \dots, M$ at quarter t , where $M = 15$ for our Top-15 set. I define the ensemble (average) forecast as

$$\hat{y}_t^{avg} = \frac{1}{M} \sum_{m=1}^M \hat{y}_{m,t}.$$

The results in this paper are based on 2026Q1 as the last observed quarter (unemployment rate: **3.20%**). Conditional on information available through that quarter, the Top-15 ensemble forecasts **2.62%** for 2026Q2 and **2.54%** for 2026Q3.

5.2.1 Forecast setup and dynamic predictions

All models are estimated on quarterly data (seasonality $s = 4$) after restricting the sample to 2004 onward. Forecasts are generated using Stata's dynamic prediction option, so that multi-step forecasts are produced recursively from the chosen starting quarter. For the results in this paper, the last observed quarter is 2026Q1 and the forecast covers 2026Q2–2026Q3. For SARIMAX models, dynamic forecasts also require the relevant lagged exogenous regressors to be available around the forecast origin.

Figure 10 compares the fitted paths produced by the selected models.




Figure 10: Best-fitting model comparison (forecast origin at 2026Q1).

Backtesting windows. To assess stability across episodes, I replicate the dynamic forecasting exercise using earlier forecast origins: (i) 2013Q1 (long historical backtest), (ii) 2019Q1 (pre-COVID backtest), and (iii) 2020Q1 (COVID zoom). For each selected model, a dynamic forecast path is computed starting from the chosen origin quarter: the model uses observed unemployment up to that point, then projects forward recursively so that each successive quarter is forecast from the model’s own prior predictions rather than from subsequently observed values. For SARIMAX models, the same lagged predictors that would have been available at each origin are used, just as in the operational exercise. The goal is to see how the models propagate forecasts through different historical environments, including the pre-COVID expansion and the COVID shock itself.




Figure 11: Backtest with forecast origin at 2013Q1.




Figure 12: Backtest with forecast origin at 2019Q1.




Figure 13: COVID zoom: dynamic prediction from 2020Q1.

5.2.2 Forecast-error diagnostics: distribution and serial correlation

As the final forecasting object is an *ensemble* obtained by averaging the forecasts from the 15 shortlisted SARIMA/SARIMAX specifications in Table 17, let \hat{y}_t^{avg} denote this average forecast for unemployment in quarter t . The corresponding forecast error is

$$e_t = y_t - \hat{y}_t^{avg},$$

so that $e_t > 0$ indicates that the ensemble *underpredicts* unemployment.

Error distribution. Figure 14 displays the empirical distribution of e_t using a histogram with a Normal overlay and a kernel density estimate. Over the evaluation sample ($N = 89$), the mean error is close to zero ($\bar{e} \approx -0.009$ percentage points), and the dispersion is about 0.24 percentage points. This is not presented as a formal standalone bias test; rather, taken together with the roughly centered histogram and the absence of persistent serial correlation discussed below, it suggests that the average forecast does not exhibit economically meaningful systematic bias over the evaluation sample.

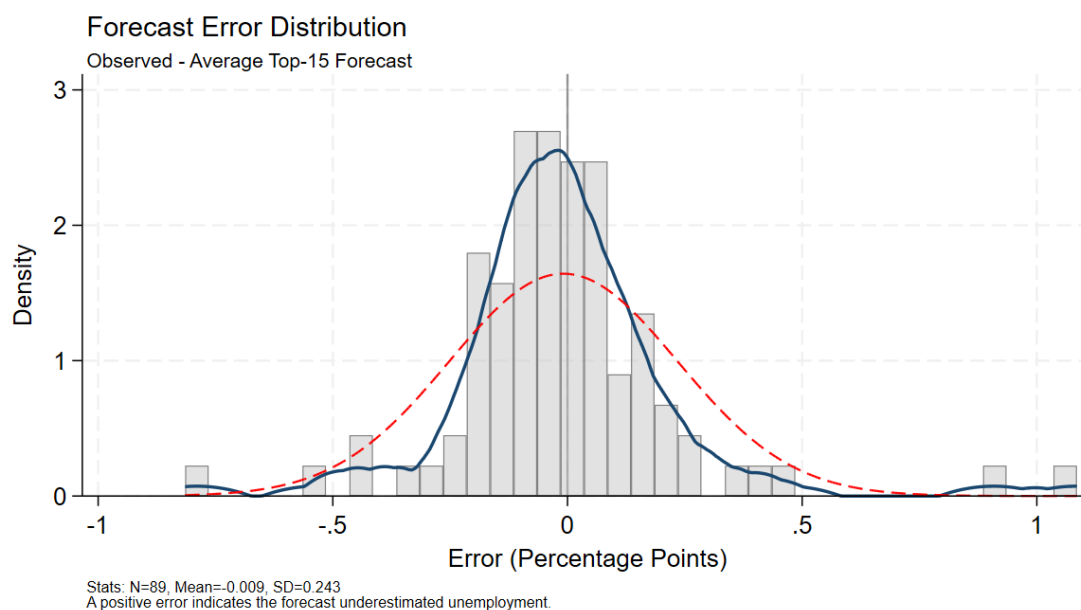


Figure 14: Distribution of average-forecast errors ($e_t = y_t - \hat{y}_t^{avg}$).

Serial correlation: correlogram and Ljung-Box test. A well-specified forecasting model should leave forecast errors that are approximately unpredictable using information available at earlier dates. I therefore assess whether the ensemble errors exhibit remaining serial correlation using two complementary tools:

- **Autocorrelation function (ACF):** that shows *which lags* display correlation, and is particularly useful for detecting leftover seasonality at multiples of four quarters.
- **Ljung–Box portmanteau test:** a joint test of the null hypothesis that autocorrelations are zero up to a chosen lag L , providing a formal “white-noise” check.

With quarterly data, I report $L = 20$ (about five years) as a natural benchmark and $L = 40$ (about ten years) as a stricter robustness check. In both cases the evidence is consistent with white-noise errors, in the sense that the Ljung–Box null of no residual autocorrelation up to the stated lag is not rejected:

$$Q(20) = 9.224, \quad p = 0.980 \quad \text{and} \quad Q(40) = 22.163, \quad p = 0.990.$$

The very high p -values indicate that the observed autocorrelations are jointly small relative to what would be expected under serially uncorrelated forecast errors. Figure 15 complements these tests: ACF spikes remain indeed within the 95% confidence bands, indicating that any residual dependence is weak and not systematic.

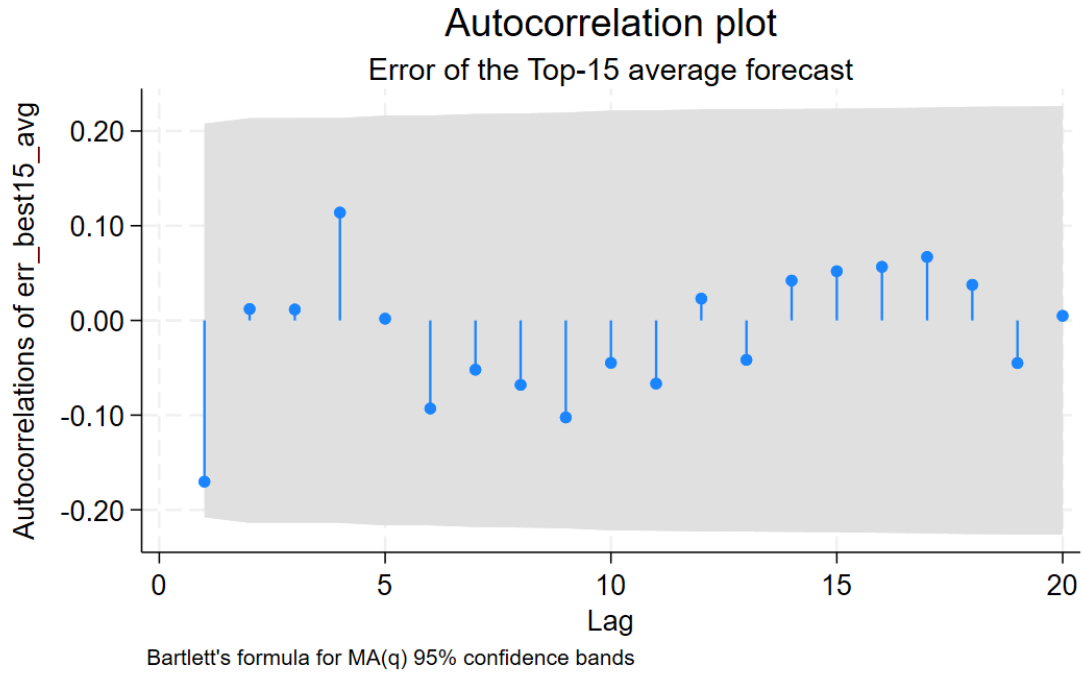


Figure 15: ACF of average-forecast errors. Bands indicate 95% confidence intervals.

5.2.3 Benchmarking against UCM residuals (model-selection vs forecasting)

To connect the SARIMAX forecasting framework with the earlier UCM diagnostics, I compare the ensemble forecast errors to one-step-ahead prediction errors obtained from two UCM specifications: (i) a basic UCM with a random-walk level but no explicit seasonal component, and (ii) a seasonal UCM with a random-walk level and a stochastic quarterly seasonal component of period 4.

Figure 16 overlays the three error series. The ensemble errors are centered near zero and behave more similarly to the seasonal-UCM residuals than to those from the non-seasonal UCM. The largest deviations occur around the COVID period and the 2009 GFC, reflecting abrupt movements that are difficult to capture using smooth trend and seasonal components alone. As Figure 16 shows, the ensemble errors are in several episodes slightly more tightly clustered around zero than those of the seasonal UCM, which is why the ensemble can be interpreted as modestly improving on the seasonal-UCM benchmark in visual one-step-ahead error comparison.

Overall, this comparison supports the modeling strategy adopted in this paper: UCM diagnostics are used primarily for component decomposition and preliminary screening of auxiliary information, whereas final forecasting relies on SARIMA/SARIMAX specifications, which deliver forecast errors closer to white noise.




Figure 16: One-step-ahead errors from basic and seasonal UCM specifications, compared with *ensemble* forecast errors from the Top-15 SARIMA/SARIMAX average.

Prediction interval: within-model and between-model uncertainty. Because the operational forecast is an equally weighted model average, predictive uncertainty arises from two sources: (i) *within-model* forecast uncertainty (each model’s own predictive variance), and (ii) *between-model* uncertainty (the dispersion of point forecasts across models). Let $\hat{y}_{m,t}$ denote the point forecast from model $m = 1, \dots, M$ and let $se_{m,t}^2$ denote its predictive variance. Under the mixture interpretation of model averaging, the Law of Total Variance implies the decomposition

$$V_t = \mathbb{E}_m[\text{Var}(Y_t \mid m, \mathcal{D})] + \text{Var}_m(\mathbb{E}[Y_t \mid m, \mathcal{D}])$$

so that a $(1 - \alpha)$ prediction interval is $\hat{y}_t^{avg} \pm z_{1-\alpha/2}\sqrt{V_t}$.

Here, $m \in \{1, \dots, M\}$ indexes the forecasting model in the ensemble, and \mathcal{D} denotes the information set (the observed training data, up to the forecast origin) used to estimate the models and form predictions.

Thus, the ensemble predictive variance can be rewritten as follows:

$$V_t = \underbrace{\frac{1}{M} \sum_{m=1}^M se_{m,t}^2}_{\text{average within-model variance}} + \underbrace{\text{Var}_m(\hat{y}_{m,t})}_{\text{between-model dispersion}},$$

where $M = 15$ is the number of models in the ensemble. A $(1 - \alpha)$ prediction interval is then given by

$$\hat{y}_t^{avg} \pm z_{1-\alpha/2}\sqrt{V_t}.$$

Figure 17 reports the resulting 95% prediction interval as a fan chart from the last observed quarter onward. Figures 17 and 18 are the two outputs disseminated by IRE in its quarterly public releases.




Figure 17: Top-15 average forecast with 95% prediction interval (within-model plus between-model uncertainty).

Finally, Figure 18 situates each 2026 forecast within the recent same-quarter pattern. The *de-seasonalised* forecasts for the two forecast quarters are placed alongside the corresponding same-quarter observations from 2023–2025, so that the projected seasonal position can be read in the context of how the same quarter has evolved in recent years, rather than against the raw annual path alone. The accompanying prediction intervals make the uncertainty surrounding each projected position explicit.

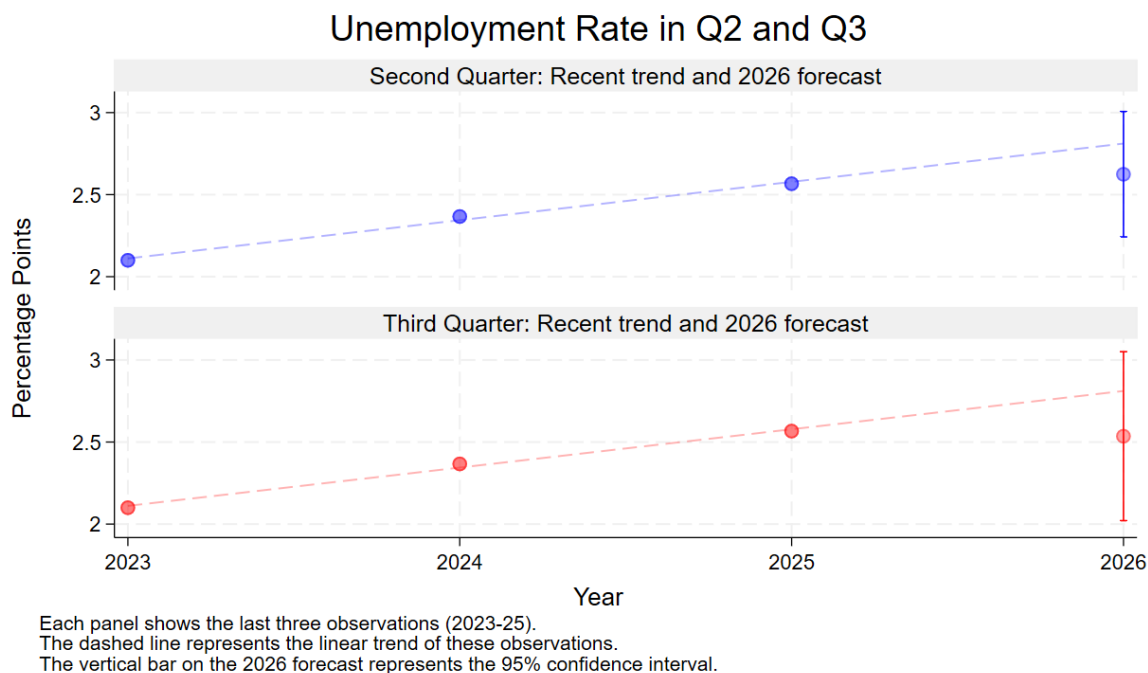


Figure 18: Within-quarter trend (2023–2025) and current-edition forecast for Q2 and Q3, with 95% prediction intervals.

6 Conclusion

This report documents and motivates the time-series framework used at the Istituto di Ricerche Economiche (IRE) to analyze and forecast the SECO-registered unemployment rate in the Canton of Ticino. Using quarterly data from 2004Q1 onward, with the operational forecast conditioned on the observed 2026Q1 outcome, the empirical strategy combines classical time-series diagnostics to build the best-performing forecasting models by integrating a curated set of auxiliary indicators harmonized into a consistent dataset.

The analysis proceeded in two steps. First, unit-root tests, correlograms, and Unobserved Components Models were implemented to confirm that quarterly seasonality was a key feature of unemployment in the region, and to provide a structured way to screen the predictive power of explanatory variables in predicting fluctuations in the unemployment rate. Second, a grid of SARIMA models was estimated to identify strong univariate dynamics, which were then paired with the best-performing lagged predictor sets to generate a set of SARIMAX candidates. Model comparison based on log likelihood, AIC, and BIC yielded a broader pool of high-ranking models, from which the final Top-15 forecasting ensemble was selected so as to retain both strong rankings and diversity across information sets.

Overall, the results indicate that accounting for seasonality is essential for both model fit and forecasting accuracy. While the incremental contribution of explanatory variables is modest on average, it is systematic, with predictive gains concentrated in specifications that combine KOF labour-market expectations and CREA activity measures. Exchange rates and a single Google Trends specification remain in the final ensemble as complementary information sets. Final forecasts are obtained from a fixed Top-15 ensemble, with prediction intervals combining within-model uncertainty and between-model disper-

sion. The results in this paper, conditioning on the observed 2026Q1 outcome of 3.20%, indicate a moderate decline in unemployment to 2.62% in 2026Q2 and 2.54% in 2026Q3. Forecast-error diagnostics show errors centered close to zero and no meaningful remaining serial correlation, supporting the adequacy of the selected approach for IRE's quarterly reporting and applied economic monitoring.

A Data Sources: Full Documentation

This appendix documents each data source used in the analysis, including variable definitions, construction notes, and access links.

A.1 Unemployment Data (Amstat / SECO)

Unemployment data were sourced from Amstat. The original series are available at monthly frequency and with breakdowns by nationality and gender; for the purposes of this study, monthly observations are aggregated to quarterly averages and no demographic disaggregation is used.

- <https://www.amstat.ch/MicroStrategy/>

A.2 Google Trends

Google Trends data includes search interest for a number of job-related queries (specific words or sets of words) and *topics* (covering all terms Google Trends considers part of a coherent research area), by users located in Ticino. The keywords were inputted into Google Trends and supplemented by related search terms suggested by Google Trends itself. They include Italian terms related to job offers, job searches, and unemployment, reflecting the local language and search behavior. The full list, with translations, is as follows:

- “offerte lavoro” – “job offers”
- “lavoro Ticino” – “work [or *jobs*] in Ticino”
- “unia lavoro” – job searches involving the UNIA labour union in Ticino
- “ocst lavoro” – job searches involving the OCST (Organizzazione Cristiano-Sociale Ticinese) labour union
- “lavoro” – “work” or “job”
- “tutti.ch lavoro” – job listings on the Tutti.ch Ticino-based platform
- “cerco lavoro” – “looking for a job”
- “annunci lavoro” – “job advertisements”
- “disoccupazione” – “unemployment”
- “Randstad” – searches related to the Randstad employment agency
- “Adecco” – searches related to the Adecco employment agency
- “Carriera (TOPIC)” – “Career (TOPIC)”
- “Disoccupazione (TOPIC)” – “Unemployment (TOPIC)”
- “Sussidio (TOPIC)” – “Subsidy (TOPIC)”
- “unia” – the UNIA labour union

- “ocst” – the OCST labour union
- “offerta di impiego (TOPIC)” – “job offer (TOPIC)”
- “Linkedin” – searches related to the LinkedIn platform
- “Job room” – searches related to the Job Room platform
- <https://trends.google.com/trends/explore>

A.3 Cantonal GDP Growth (CREA)

The quarterly GDP growth rates for cantons were sourced from the CREA Institute of Applied Economics (Université de Lausanne). This indicator of local economic activity is adjusted for seasonality and major sporting events; the methodology allows quarterly GDP estimates for Swiss cantons and cities from 1997 onward.

- <https://www.unil.ch/crea/home/menuinst/indicateurs/cantonal-gdp-growth.html>

A.4 Retail Turnover (BFS)

Retail turnover data were sourced from the Federal Statistical Office (BFS). The underlying series are calendar- and seasonally-adjusted indices (annual average 2015 = 100), aggregated to the quarterly frequency for consistency with the main dataset.

- https://www.bfs.admin.ch/asset/it/px-x-0603020000_102

A.5 KOF Economic Sentiment and Uncertainty Proxies

To proxy uncertainty and firms’ labour-market expectations, this report uses three quarterly indicators produced by the KOF Swiss Economic Institute that are available over a long sample. Higher values of these indicators signal lower uncertainty and more favourable labour-market conditions.

A.5.1 KOF Business Situation Indicator

The KOF Business Situation Indicator summarises firms’ qualitative assessments of their current business situation. The data are derived from monthly surveys of companies across industry, retail, construction, project engineering, financial, and insurance sectors. Respondents evaluate their situation on a three-point scale (“good”, “satisfactory”, or “poor”); KOF computes a seasonally adjusted balance (percentage “good” minus percentage “poor”) for each sector and aggregates these into a single series for the Swiss economy.

- <https://kof.ethz.ch/en/forecasts-and-indicators/indicators/kof-business-situation-indicator.html>

A.5.2 KOF Economic Barometer

The KOF Economic Barometer is a composite leading indicator tracking the short-term stance of the Swiss business cycle. It is anchored to Swiss GDP growth and draws on a database of several hundred Swiss and international time series, re-evaluated annually. In this analysis, the quarterly average of the monthly barometer is used.

- <https://kof.ethz.ch/en/forecasts-and-indicators/indicators/kof-economic-barometer.html>

A.5.3 KOF Employment Indicator

The KOF Employment Indicator is a survey-based leading indicator of Swiss labour-market developments. It is constructed from business tendency surveys in nine sectors (covering around 85% of private-sector employment). Firms report current staffing levels and three-month hiring intentions; KOF computes seasonally adjusted balances and aggregates them into an overall indicator. Positive values indicate net hiring intentions; negative values indicate net reduction plans. The quarterly version is used in the empirical analysis.

- <https://kof.ethz.ch/en/forecasts-and-indicators/indicators/kof-employment-indicator.html>

A.6 Exchange Rates (ECB)

EUR/CHF exchange rate data were sourced from the European Central Bank (ECB). Daily rates are averaged to the quarterly frequency.

- <https://data.ecb.europa.eu/data/datasets/EXR/EXR.D.CHF.EUR.SP00.A>

A.7 GDP and GVA Data (BFS)

GDP and Gross Value Added (GVA) data by sector for Swiss cantons were sourced from the Federal Statistical Office (BFS). The dataset covers annual observations from 2008 onward and is used to construct an annual GDP proxy for Ticino.

- <https://www.bfs.admin.ch/bfs/en/home/statistics/national-economy/national-accounts/gross-domestic-product-canton.assetdetail.23526492.html>

A.8 SECO Data (Secondary and Tourism Sectors)

Sectoral series disseminated via BFS PxWeb (source: SECO) were collected to proxy developments in the secondary sector and in tourism-related activity in Ticino.

- https://www.pxweb.bfs.admin.ch/pxweb/it/px-x-1003020000_202/px-x-1003020000_202/px-x-1003020000_202.px/table/tableViewLayout2/
- https://www.pxweb.bfs.admin.ch/pxweb/it/px-x-0603010000_102/px-x-0603010000_102/px-x-0603010000_102.px/table/tableViewLayout2/

A.9 OECD Data

Data on industrial production and manufacturing were sourced from the OECD, as indices with base year 2015. The dataset covers quarterly total manufacturing and construction production.

- <https://data.oecd.org/industry/industrial-production.htm>

References

- [1] Box, George E. P., Gwilym M. Jenkins, and Gregory C. Reinsel. *Time Series Analysis: Forecasting and Control*. 4th ed. Hoboken, NJ: Wiley, 2008.
- [2] Hamilton, James D. *Time Series Analysis*. Princeton University Press, 1994.
- [3] Harvey, Andrew C. *Forecasting, Structural Time Series Models and the Kalman Filter*. Cambridge University Press, 1989.
- [4] Hayashi, Fumio. *Econometrics*. Princeton University Press, 2000.
- [5] Kalman, Rudolf E. “A new approach to linear filtering and prediction problems.” *Transactions of the ASME—Journal of Basic Engineering*, Series D, 82: 35–45, 1960.
- [6] Stata Documentation. *ARIMA — ARIMA, ARMAX, and other dynamic regression models*. StataCorp, College Station, TX. <https://www.stata.com/manuals/tsarima.pdf>
- [7] Stata Documentation. *UCM — Unobserved-components model*. StataCorp, College Station, TX. <https://www.stata.com/manuals/tsucm.pdf>