Internet Appendix for:

Macroeconomic Fundamentals and the Shape of Sovereign Credit Risk

This Internet Appendix collects supplementary material not included in the published manuscript: implementation details for all forecasting methods (Section A), additional data description (Section B), and additional empirical results (Section C). Cross-references to the main text resolve against the published manuscript.

A Estimation Details

This appendix provides implementation details for the forecasting methods described in Section III.B of the main paper. Sections A.1–A.3 cover the eight linear specifications, Sections A.4–A.6 the five nonlinear specifications, and Section A.7 the time-series cross-validation procedure used to select hyperparameters for every method. Throughout, we use the panel notation established in the main paper: $s_{ij,t}$ for the log CDS spread of country i at maturity j and date t, and $\mathbf{x}_{i,t-1} \in \mathbb{R}^P$ for the real-time predictor vector. Hyperparameters reported in each subsection are tuned over the 12-month training window using the time-series cross-validation procedure in Section III.B of the main paper, with the final estimates frozen before testing.

A.1 Penalized Regressions

We follow the approach outlined by Balduzzi et al. (2023) and estimate a "level+slope" linear regression specification in which the real-time macroeconomic indicators $\mathbf{x}_{i,t-1}$ (plus an intercept) interact with the CDS maturity j of the respective CDS contract for country i,

$$s_{ij,t} = \mathbf{x}'_{i,t-1}\beta + \sum_{j \in \mathcal{J}\setminus\{3\}} \left(\mathbf{x}_{i,t-1} \cdot m_j\right)' \delta_j + \varepsilon_{ij,t}, \tag{1}$$

where $s_{ij,t}$ denotes the logarithm of the CDS spread for country i, maturity j, at time t, m_j is a dummy variable that takes value one if the contract has j-year maturity and zero otherwise, and $\mathcal{J} = \{3, 5, 7, 10\}$. The reference group is the 3-year CDS spread. In the main empirical analysis, we start from Eq. (1) and estimate penalized variants, such as Ridge regression (Hsiang, 1975), Lasso (Tibshirani, 1996), and Elastic Net (Zou and Hastie, 2005).

To prevent the maturity-dummy coefficients from being penalized, we subtract the cross-country mean at each time and maturity before estimation, as discussed in Section III.B:

$$s_{ij,t} - \bar{s}_{\cdot j,t} = (\mathbf{x}_{i,t-1} - \bar{\mathbf{x}}_{\cdot,t-1})'\beta + \sum_{j \in \mathcal{J} \setminus \{3\}} \left((\mathbf{x}_{i,t-1} - \bar{\mathbf{x}}_{\cdot,t-1}) \cdot m_j \right)' \delta_j + \varepsilon'_{ij,t}, \tag{2}$$

where $\bar{s}_{\cdot j,t} = \frac{1}{N_t(j)} \sum_{i=1}^{N_t(j)} s_{ij,t}$ is the cross-country average of the log CDS spread at maturity j and time t, with $N_t(j)$ the number of countries for which maturity-j data are available at t, and $\bar{\mathbf{x}}_{\cdot,t-1} = \frac{1}{N_t(j)} \sum_{i=1}^{N_t(j)} \mathbf{x}_{i,t-1}$ the corresponding predictor average.[1] The demeaned predictions $\widehat{s}^{d}_{ij,t}$ are re-centered at the forecast stage by adding back the maturity-specific training-window mean $\bar{s}^{(\text{train})}_{\cdot j,\cdot}$ computed across i and t within the training window, yielding $\widehat{s}_{ij,t} = \widehat{s}^{d}_{ij,t} + \bar{s}^{(\text{train})}_{\cdot j,\cdot}$.

For Lasso, we search for the optimal L1 penalty α over $[10^{-2}, 10^{3}]$; for Ridge, we search for the optimal L2 penalty λ over $[10^{-2}, 10^{4}]$. In both cases, we select 20 values from the respective ranges using logarithmic spacing. For Elastic Net, we search jointly over (α, λ) using the same ranges. All penalized regressions are estimated on the demeaned inputs without further standardization, since the FRED-MD transformations in Section III.A render the predictors approximately scale-comparable.

For robustness, we also estimate a "level-only" variant of Eq. (2) in which the maturity-interaction terms are omitted. The level-only specification fits a single coefficient vector β across all maturities. Results are reported in Table C.1 in Appendix C.2.

A.2 Latent Factor Models

For the latent-factor specifications, we apply PCA or PLS to the level-only specification, then interact the resulting components with the CDS maturity dummies. The components serve as a condensed representation of the original predictor set, simultaneously achieving dimension reduction and preserving the level+slope structure of Eq. (1). The number of retained components is selected by time-series cross-validation; we fix $K = 5$ throughout the empirical analysis, as this value is selected in the majority of rolling windows.

[1] After cross-country demeaning, the maturity-dummy coefficients are orthogonal to the demeaned intercept component, so penalizing the slope and interaction coefficients does not shrink them.

A.3 Complete Subset Regressions

Complete Subset Regression (CSR) uses equal-weighted combinations of forecasts from all p-variable subsets, where $p \leq P$ and P is the total number of predictors. For a given subset size p, we run the regression of $s_{ij,t}$ on each p-dimensional subset of $\mathbf{x}_{i,t-1}$ and average the resulting coefficient vectors across subsets to obtain the CSR estimator $\widehat{\beta}^{\text{CSR}}$.

As an illustration, the univariate case $p = 1$ has P subset regressions, one per predictor. The CSR forecast is the equal-weighted average $\widehat{s}_{ij,\tau}^{\text{CSR}} = \frac{1}{P}\sum_{k=1}^{P} \mathbf{x}'_{i,\tau-1}\widehat{\beta}_k$, where $\widehat{\beta}_k$ is the P-vector with the least-squares estimate on the k-th predictor in the k-th position and zeros elsewhere. The equal-weighted forecast combination (EWC) is the special case $p = 1$ of CSR.

In our empirical analysis, we treat the subset size p as a hyperparameter. When the number of p-variable subsets exceeds 5,000, we estimate 5,000 subsets drawn without replacement from the full set; otherwise, we estimate every possible subset.

A.4 Tree Ensemble and Regression Trees

At iteration k, XGBoost updates the prediction of the log CDS spread as $\widehat{s}_{ij,t}^{(k)} = \widehat{s}_{ij,t}^{(k-1)} + f_k(\mathbf{x}_{i,t-1})$, with each new tree f_k trained to minimize the regularized objective

$$\mathcal{L}^{(k)} = \sum_{(i,j,t)} \ell\Big(s_{ij,t},\, \widehat{s}_{ij,t}^{(k-1)} + f_k(\mathbf{x}_{i,t-1})\Big) + \Omega(f_k), \qquad \Omega(f_k) = \eta\, T_k + \tfrac{1}{2}\lambda \sum_{l=1}^{T_k} w_l^2, \tag{3}$$

where T_k is the number of leaves in tree k, w_l the weight of leaf l, and (η, λ) control the leaf-count penalty and the L2 shrinkage of leaf weights.[2] Approximating the loss around $\widehat{s}_{ij,t}^{(k-1)}$ by a second-order Taylor expansion and writing I_l for the set of observations assigned to leaf l, the optimal leaf weight and the tree-quality score are

$$w_l^\star = -\frac{\sum_{(i,j,t)\in I_l} g_{ij,t}}{\sum_{(i,j,t)\in I_l} h_{ij,t} + \lambda}, \qquad \widetilde{\mathcal{L}}^{(k)}(q) = -\tfrac{1}{2}\sum_{l=1}^{T_k} \frac{\big(\sum_{(i,j,t)\in I_l} g_{ij,t}\big)^2}{\sum_{(i,j,t)\in I_l} h_{ij,t} + \lambda} + \eta\, T_k, \tag{4}$$

with $g_{ij,t}$ and $h_{ij,t}$ the gradient and Hessian of the loss ℓ at $\widehat{s}_{ij,t}^{(k-1)}$. The sparsity-aware split-finding algorithm of Chen and Guestrin (2016) makes training scalable even when predictors are missing or sparse.

For the XGBoost model, we tune the `number of estimators` (100, 150, 200), the `learning rate` (from 0.05 to 0.30 in increments of 0.05), the `maximum depth of trees` (3, 4, 5), the

[2] The symbol η in Eq. (3) is the XGBoost leaf-count penalty.

leaf-count penalty `gamma` (from 0.0 to 0.4 in increments of 0.1), the `subsample ratio of training instances` (from 0.5 to 0.9 in increments of 0.1), and the `subsample ratio of columns per tree` (0.3, 0.4, 0.5, 0.7).

For the regression tree, we cross-validate `max_depth` (tree depth limited to 3, 5, 7, or 9 levels), `min_samples_split` (minimum of 100, 200, 500, or 1500 observations required to allow a node split), and `min_samples_leaf` (each leaf node must contain at least 500, 1500, or 2000 observations).

For the random forest, we cross-validate the `number of trees` (10, 50, 100, 500), the `maximum depth of trees` (5, 6, 7), the `minimum number of observations` required to split an internal node (100, 200, 500), and the `maximum number of features` evaluated at each split (3, 4, 5).

A.5 Support Vector Regression

We estimate nonlinear forecasts with an ϵ-support vector regression (SVR) model (e.g., Cortes and Vapnik, 1995; Drucker et al., 1996) using a radial basis function (RBF) kernel. To maintain consistency with the linear specifications, we train SVR on cross-sectionally demeaned inputs and targets and recover final forecasts by adding back the maturity-specific training-window mean, following the same procedure as for the penalized regressions.

For country i, maturity $j \in \{3, 5, 7, 10\}$, and time t, let $s_{ij,t}$ be the log CDS spread and $\mathbf{x}_{i,t-1} \in \mathbb{R}^P$ the vector of real-time macroeconomic predictors. We form demeaned pairs by subtracting the cross-country average at the same time and maturity:

$$s^d_{ij,t} = s_{ij,t} - \bar{s}_{\cdot j,t}, \qquad \mathbf{x}^d_{i,t-1} = \mathbf{x}_{i,t-1} - \bar{\mathbf{x}}_{\cdot,t-1}, \tag{5}$$

where $\bar{s}_{\cdot j,t}$ and $\bar{\mathbf{x}}_{\cdot,t-1}$ are defined as in Eq. (2). We recover the final forecasts by adding back the maturity mean computed within the training window:

$$\widehat{s}_{ij,t} = \widehat{s}^d_{ij,t} + \bar{s}^{(\text{train})}_{\cdot j,\cdot}, \qquad \bar{s}^{(\text{train})}_{\cdot j,\cdot} := \frac{1}{|\mathcal{T}_{\text{train}}|} \sum_{t \in \mathcal{T}_{\text{train}}} \bar{s}_{\cdot j,t}. \tag{6}$$

Before fitting SVR, we standardize each predictor to zero mean and unit variance using the training window:

$$\tilde{\mathbf{x}}^d_{i,t-1} = \operatorname{diag}(\widehat{\boldsymbol{\sigma}})^{-1}\big(\mathbf{x}^d_{i,t-1} - \widehat{\boldsymbol{\mu}}\big), \tag{7}$$

with $\widehat{\boldsymbol{\mu}} \in \mathbb{R}^P$ and $\widehat{\boldsymbol{\sigma}} \in \mathbb{R}^P_+$ the per-predictor mean and standard deviation. The standardized design $\tilde{\mathbf{x}}^d_{i,t-1}$ enters SVR as the feature vector.

SVR with RBF Kernel. Let $\{(\tilde{\mathbf{x}}_n, y_n)\}_{n=1}^N$ denote the standardized, demeaned training pairs, with $y_n = s^d_{ij,t}$ and $\tilde{\mathbf{x}}_n = \tilde{\mathbf{x}}^d_{i,t-1}$. We use ϵ-SVR with RBF kernel

$$K_\kappa(u,v) = \exp\big(-\kappa\,\|u-v\|_2^2\big), \tag{8}$$

where κ is the kernel bandwidth. The SVR hyperparameter triple is (C,ϵ,κ). The primal problem is

$$\begin{aligned}
\min_{w,b,\{\xi_n,\xi_n^\star\}} \quad & \tfrac{1}{2}\|w\|_2^2 + C\sum_{n=1}^N(\xi_n+\xi_n^\star) \\
\text{s.t.} \quad & y_n - (w^\top\phi(\tilde{\mathbf{x}}_n)+b) \le \epsilon + \xi_n^\star, \\
& (w^\top\phi(\tilde{\mathbf{x}}_n)+b) - y_n \le \epsilon + \xi_n, \\
& \xi_n, \xi_n^\star \ge 0 \qquad (n=1,\dots,N),
\end{aligned} \tag{9}$$

where ϕ is the implicit feature map of K_κ. Its dual maximizes

$$\max_{\{\alpha_n,\alpha_n^\star\}} -\epsilon\sum_{n=1}^N(\alpha_n+\alpha_n^\star) + \sum_{n=1}^N y_n(\alpha_n^\star-\alpha_n) - \tfrac{1}{2}\sum_{n,m=1}^N(\alpha_n^\star-\alpha_n)(\alpha_m^\star-\alpha_m)\,K_\kappa(\tilde{\mathbf{x}}_n,\tilde{\mathbf{x}}_m), \tag{10}$$

subject to $0 \le \alpha_n, \alpha_n^\star \le C$ and $\sum_{n=1}^N(\alpha_n^\star-\alpha_n)=0$. The resulting predictor is the standard support-vector expansion

$$\widehat{s}^d(\tilde{\mathbf{x}}) = \sum_{n=1}^N(\alpha_n^\star-\alpha_n)\,K_\kappa(\tilde{\mathbf{x}}_n,\tilde{\mathbf{x}}) + b, \tag{11}$$

where only the support vectors (observations with $\alpha_n^\star-\alpha_n \ne 0$) contribute.

For the SVR model, we tune the `regularization parameter` C (0.01, 0.1, 1, 10), the `epsilon-insensitive tube width` ϵ (0.01, 0.1, 0.2, 0.5), and the `RBF kernel bandwidth` κ (`scale`, 10^{-3}, 10^{-2}). The `scale` option corresponds to $\kappa = 1/(P\cdot\mathrm{Var}(\tilde{\mathbf{x}}))$, which simplifies to $\kappa = 1/P$ after standardization. The selected hyperparameter combination $(\widehat{C},\widehat{\epsilon},\widehat{\kappa})$ is applied within each expanding training window to generate demeaned forecasts $\widehat{s}^d_{ij,t}$ via Eq. (11), which are re-centered using Eq. (6).

A.6 Feedforward Neural Network

We implement a feedforward neural network (NN) following the architecture of Gu et al. (2020), adapted to our demeaned panel setting. As with SVR, the NN is trained on demeaned inputs and targets, and final forecasts are obtained by re-adding the maturity-specific training-window means to maintain consistency with the linear specifications.

The NN is a fully connected feedforward network with four hidden layers of sizes 32, 16, 8, and 4, each with `ReLU` activations and `he_normal` weight initialization. To improve numerical stability when training on demeaned inputs, each hidden layer is followed by `batch normalization`. To mitigate overfitting, we apply `dropout` regularization with rate 0.50 in the first three hidden layers and 0.25 in the final hidden layer. The output layer produces a single demeaned prediction through a linear activation.

All inputs are standardized to zero mean and unit variance using a `StandardScaler`. The network is trained with the `Adam` optimizer at `learning rate` 10^{-3} and `batch size` 32, for up to 100 `epochs` with `early stopping` on the validation loss and a patience of 20 epochs. A `ReduceLROnPlateau` scheduler halves the learning rate whenever the validation loss stagnates, with a floor of 10^{-5}. Demeaned out-of-sample predictions $\widehat{s}^{d}_{ij,t}$ are re-centered using the same maturity-mean formula as for SVR (Eq. (6)).

A.7 Cross-Validation Strategy

To cross-validate the hyperparameters of each model, we use a time-series-split approach that preserves the temporal ordering of the data and prevents look-ahead bias (e.g., Gu et al., 2020; Bianchi et al., 2021). Figure A.1 illustrates the procedure.

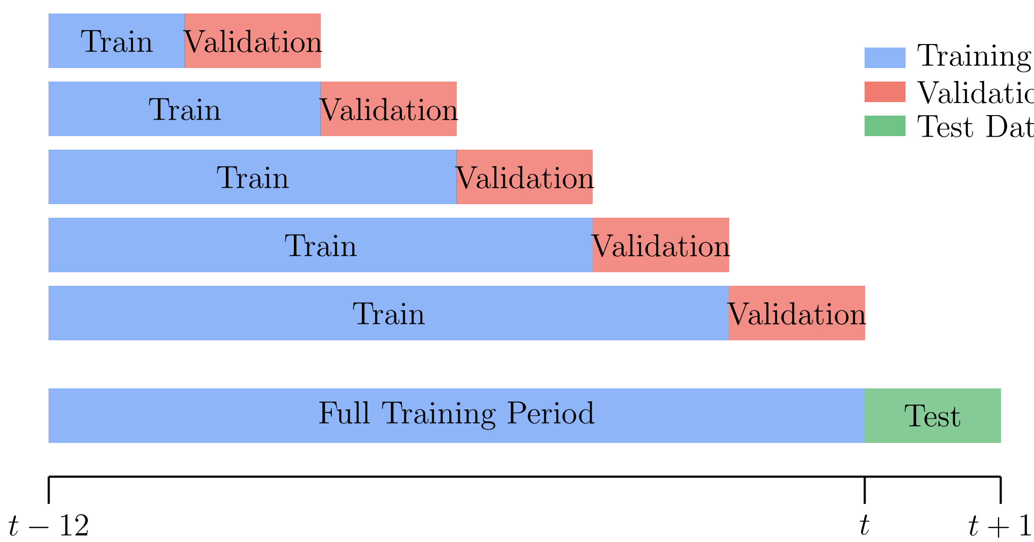

Figure A.1: **Recursive time-series cross-validation.** The figure illustrates the time-series cross-validation strategy with five sequential, non-overlapping validation folds and a final test set. The procedure is used to train each model's hyperparameters, such as penalty terms in penalized regressions or learning rates in XGBoost.

For each fold, the model is trained on the earlier block and validated on the subsequent block. This preserves the temporal dependence structure of the data and ensures that validation errors reflect each model's ability to generalize to subsequent, unseen observations. Hyperparameters are selected to minimize the average validation mean squared error across the five folds. Once selected, they remain fixed throughout the testing window, as described in Section III.B.

B Supplementary Data Description

B.1 CDS Contract Specifications

Table B.1 reports the restructuring clause, currency denomination, and earliest available date for each country-maturity pair. All 25 sovereigns trade under the Complete Restructuring (CR) clause, except Canada, which uses the No Restructuring (XR) clause. All contracts are denominated in US dollars, except the United States, whose sovereign CDS trades in euros. The panel is unbalanced: five countries (Belgium, Denmark, Germany, Italy, and the Netherlands) are missing one or two maturities, and several countries enter the sample after the panel start date of 15 January 2008. The most notable late entry is Canada, which becomes available only in November 2017.

B.2 Correlation Structure

Panel A of Figure B.1 reports the per-country correlation between each macroeconomic variable and the 5-year CDS spread. Three patterns stand out. First, business and consumer confidence indicators (BCI, CCI) and consumer price inflation (CPI) are negatively correlated with CDS spreads across most countries in the panel, consistent with sentiment and activity-slowdown signals co-moving with rising credit risk. Second, the interest-rate term-structure slope (Spread) and the debt-to-GDP ratio (Debt/GDP) are positively correlated with CDS spreads, with the strongest positive values concentrated in SWEAP countries and advanced emerging markets — consistent with these variables carrying elevated sovereign-risk content when debt sustainability is a concern. Third, the magnitude and even the sign of the correlations vary substantially across countries for most variables, including the real-activity indicators (GDP, PCE, GCE, GFC) and the global financial factors (VIX, IGREA). This advocates for quite heterogeneous country-level relationships.

Panel B reports the average pairwise correlation among the 23 country-specific macroeconomic variables (excluding VIX and IGREA, which are globally uniform). Real-activity variables (GDP, EGS, GCE, GFC, IGS, PCE) form a dense block of positive correlations, reflecting the common business-cycle component. Confidence indicators (BCI, CCI) and CPI

Country	Clause	CCY	3Y	5Y	7Y	10Y
SWEAP						
Ireland	CR	USD	2008-01	2008-01	2008-01	2008-01
Italy	CR	USD	2004-10	2004-10	—	2004-10
Portugal	CR	USD	2006-03	2004-10	2004-10	2004-10
Spain	CR	USD	2005-07	2004-10	2004-10	2004-10
EU ex-SWEAP						
Austria	CR	USD	2007-05	2004-10	2004-10	2004-10
Belgium	CR	USD	—	2004-10	2004-10	2004-10
Estonia	CR	USD	2006-10	2006-02	2006-02	2006-02
Finland	CR	USD	2007-08	2006-02	2006-02	2006-02
France	CR	USD	2007-05	2005-08	2005-08	2005-08
Germany	CR	USD	2007-05	2005-01	—	2004-10
Netherlands	CR	USD	—	2005-09	—	2005-09
Slovak Republic	CR	USD	2004-10	2004-10	2004-10	2004-10
Slovenia	CR	USD	2004-11	2004-11	2010-09†	2006-01
OECD non-EU						
Canada	XR	USD	2017-11†	2017-11†	2017-11†	2017-11†
Chile	CR	USD	2004-10	2004-10	2004-10	2004-10
Czech Republic	CR	USD	2009-07†	2009-04†	2009-04†	2009-04†
Denmark	CR	USD	—	2009-06†	—	2009-06†
Hungary	CR	USD	2004-10	2004-10	2008-07†	2004-10
Japan	CR	USD	2007-07	2004-10	2004-10	2004-10
Korea	CR	USD	2004-10	2004-10	2004-10	2004-10
Mexico	CR	USD	2004-10	2004-10	2004-10	2004-10
Poland	CR	USD	2004-10	2004-10	2004-10	2004-10
Sweden	CR	USD	2007-11	2006-02	2006-02	2006-02
United Kingdom	CR	USD	2007-11	2006-04	2007-11	2007-11
United States	CR	EUR	2007-12	2007-12	2007-12	2007-12

Table B.1: **CDS contract details.** The table reports the restructuring clause, currency denomination, and earliest available date (year-month) for each country-maturity pair. CR = Complete Restructuring; XR = No Restructuring. "—" indicates that the maturity is not available in Bloomberg for the respective sovereign. Entries marked † indicate country-maturity pairs that enter the sample after the panel start date of 15 February 2008; the panel is therefore unbalanced along both the country and maturity dimensions. Source: Bloomberg.

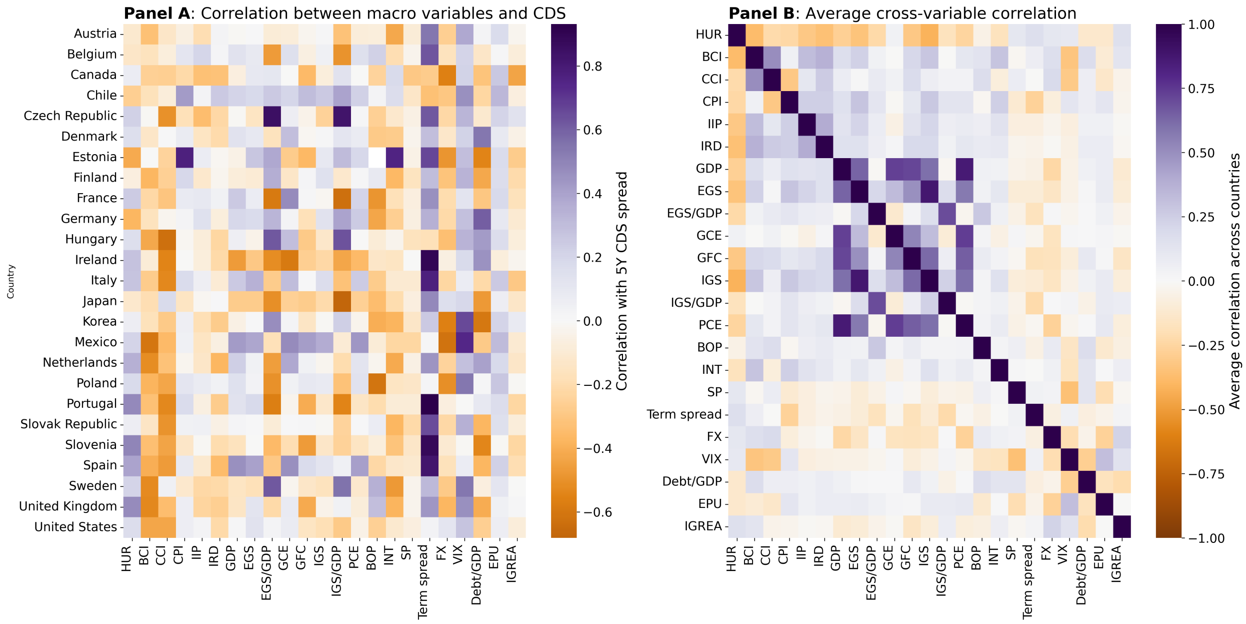

Figure B.1: **Sample correlation.** The figure reports the sample correlation between macroeconomic variables and CDS premiums per country (left panel) and the correlation among macroeconomic variables averaged across countries (right panel). The sample period is daily from 15 January 2008 to 15 December 2024.

form a second, weaker cluster. Global financial and policy-uncertainty variables (Spread, FX, VIX, Debt/GDP, EPU, IGREA) correlate weakly with each other and the real-activity block, indicating that these are distinct sources of variation rather than proxies for the same underlying factor.

B.3 Stationarity Tests

Table B.2 reports panel stationarity tests on daily log CDS spreads for each of the four contract maturities over the 2008–2024 sample. We implement two complementary tests that together bracket the order of integration. The Pesaran (Pesaran, 2007) CIPS test maintains the null of a panel unit root; it is robust to cross-sectional dependence, a first-order feature of sovereign CDS spreads, and invalidates panel tests that assume independence across countries. The Hadri (Hadri, 2000) LM test maintains the null of panel stationarity; it is the panel analog of the KPSS statistic (Kwiatkowski et al., 1992) and provides complementary evidence against the alternative of non-stationarity. We report both panel statistics and the number of countries for which the per-country test statistics reject their respective nulls at the 5% level as a disaggregated diagnostic.

The two tests reject in opposite directions at every maturity. CIPS rejects the panel unit-

Maturity	N	CIPS	# CADF rej.	Hadri Z	# KPSS rej.
3Y	22	-3.559**	19	129.117***	21
5Y	25	-3.374**	20	145.888***	24
7Y	21	-3.437**	17	125.530***	19
10Y	25	-3.711**	20	133.869***	25

Table B.2: **Panel stationarity tests for log CDS spreads.** The table reports panel unit-root and stationarity tests on daily log CDS spreads over 15 February 2008 to 15 December 2024. Column (2) gives the number of countries retained (at least 500 daily observations per country). Column (3) reports the CIPS statistic (Pesaran, 2007) for the null of a panel unit root, computed as the cross-country average of per-country CADF statistics with four augmentation lags. Column (4) reports the number of countries for which the per-country CADF statistic rejects the unit-root null at the 5% level. Column (5) reports the Hadri (Hadri, 2000) LM statistic for the null of panel stationarity, standardized to $N(0,1)$ under the null. Column (6) reports the number of countries for which the per-country KPSS statistic (Kwiatkowski et al., 1992) rejects the stationarity null at the 5% level. Significance: * 10%, ** 5%, *** 1%.

root null at the 5% level for all four maturities, with statistics ranging from -3.37 to -3.71 against a 5% critical value of approximately -2.20; per-country rejections cover 17 to 20 of the 21–25 countries in each maturity panel. Hadri rejects the panel stationarity null at the 1% level for all four maturities, with LM statistics exceeding 125 against an $N(0,1)$ reference distribution; per-country KPSS rejections cover 19 to 25 of the included countries. The combination — decisive rejection of both the strict-unit-root null and the strict-stationary null — is the signature of high persistence with slowly drifting levels rather than either pure random-walk or pure mean-reverting dynamics.

The implication for the forecasting exercise follows the high-persistence characterization in Section II: at the daily horizon, predictable variation in $s_{ij,t}$ is small relative to the country-maturity trailing mean $\bar{s}_{ij}$. The within-country-maturity underperformance documented in Section IV.B is a direct consequence of this property and does not reflect an absence of predictive content in the predictor block more broadly.

C Supplementary Empirical Results

C.1 Yearly R^2_{oos} for sub-groups and long horizons

We present three sets of figures: the 1-month and 3-month heatmaps for each of the three country sub-groups (Figures C.1–C.3); the corresponding 6-month and 12-month heatmaps for the Global panel (Figure C.4); and the 6-month and 12-month heatmaps for each country

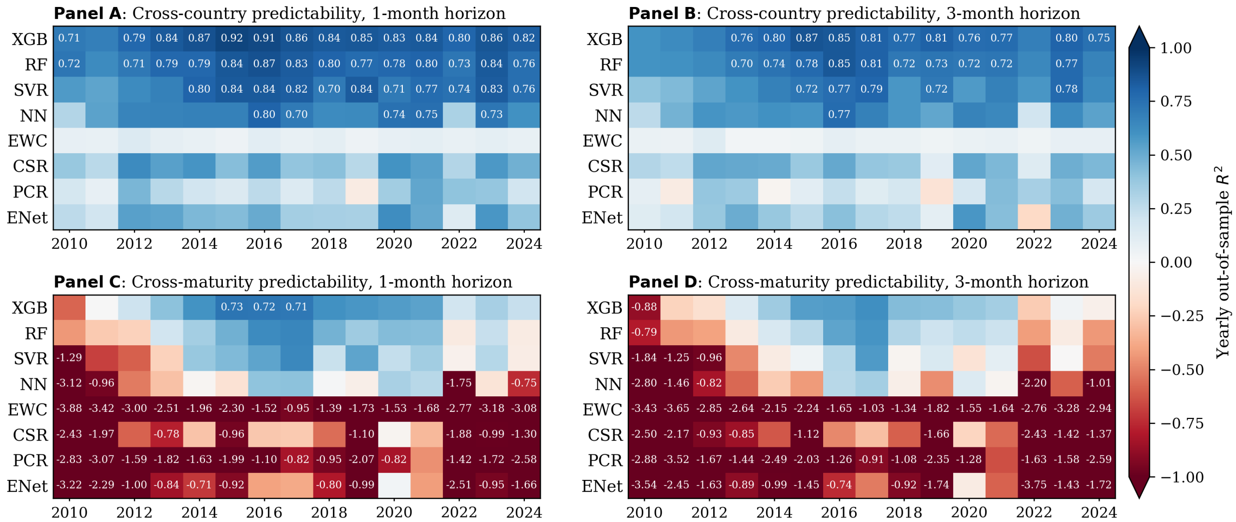

Figure C.1: **Yearly R^2_{oos} across benchmarks and horizons, OECD non-EU panel.** The figure reports the annual average of $R^2_{t,oos}$ for the OECD non-EU country sub-group. Rows: cross-country predictability against $\bar{s}_{\cdot j}$ (top) and cross-maturity predictability against $\bar{s}_{i\cdot}$ (bottom). Columns: 1-month (left) and 3-month (right) testing horizons. Models and formatting follow Figure 2 in the main paper.

sub-group (Figures C.5–C.7). Formatting follows the main-text figure: cells are colored by R^2 using a diverging scale centered at zero, and values with $|R^2| \geq 0.7$ are annotated.

The three sub-group figures confirm the main-text pattern qualitatively: linear methods are broadly positive against $\bar{s}_{\cdot j}$ and systematically negative against $\bar{s}_{i\cdot}$, while XGBoost and Random Forest maintain positive cross-maturity R^2 across most years at the 1-month horizon.

The EWC in the OECD non-EU panel (Figure C.1) is essentially flat at zero across all years, reinforcing the observation that cross-country predictability cannot be obtained from trivial aggregation. Second, the EU ex-SWEAP panel (Figure C.2) shows the cleanest separation between tree-based ensembles and all other methods at the cross-maturity benchmark. Third, the SWEAP panel (Figure C.3) exhibits catastrophic cross-maturity cells for several linear methods in the early years (e.g., EWC below -9 in 2011–2012), illustrating that small-panel linear methods are sensitive to the extreme volatility of the Eurozone crisis years. The within-row ordering of nonlinear methods (XGBoost $\approx$ RF $>$ SVR $>$ NN against $\bar{s}_{i\cdot}$) is preserved across all three sub-groups.

Figure C.4 reports the yearly R^2 pattern at the 6-month and 12-month horizons for the Global panel. The main message is the extension of framing B: linear cross-country predictability degrades as the horizon lengthens, and the 2022–2024 tightening-cycle cells turn visibly red at the 12-month horizon (Elastic Net panel (b) shows -1.08 for 2022 and -0.86 for 2023).

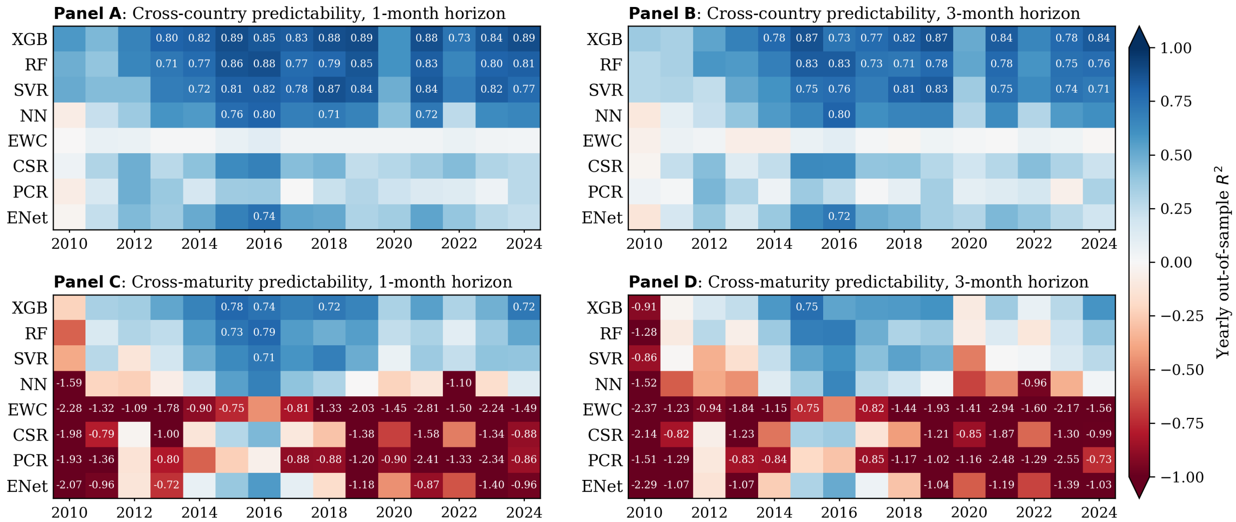

Figure C.2: **Yearly R^2_{oos} across benchmarks and horizons, EU ex-SWEAP panel.** The figure reports the annual average of $R^2_{t,oos}$ for the EU ex-SWEAP country sub-group. Rows, columns, and formatting follow Figure C.1.

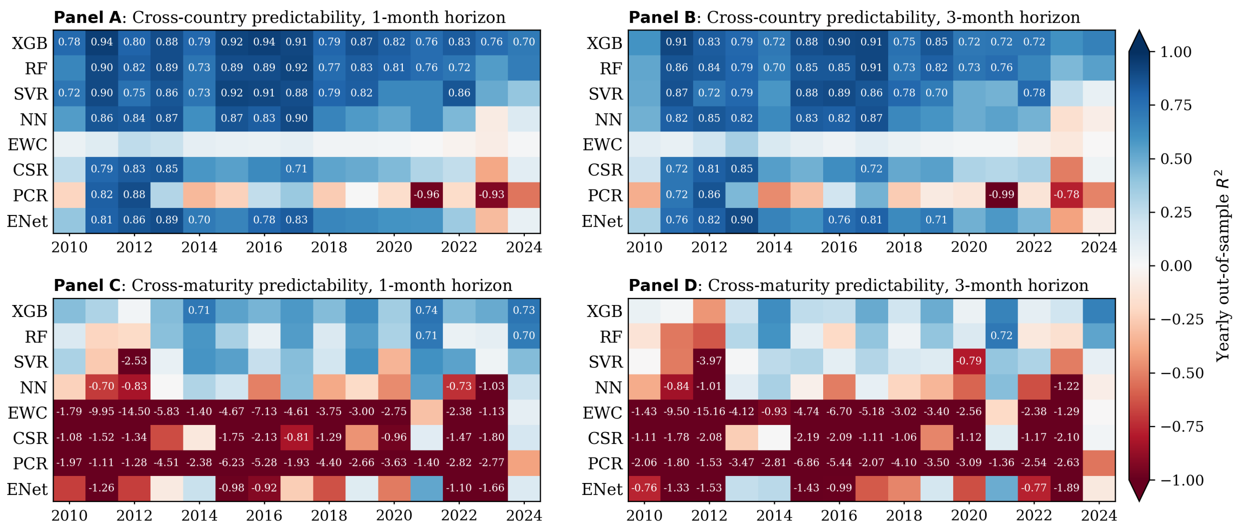

Figure C.3: **Yearly R^2_{oos} across benchmarks and horizons, SWEAP panel.** The figure reports the annual average of $R^2_{t,oos}$ for the SWEAP country sub-group. Rows, columns, and formatting follow Figure C.1.

XGBoost retains positive R^2 against $\bar{s}_{\cdot j}$ across the full 2010–2024 window at both horizons, and the relative ordering of nonlinear methods remains stable. Against $\bar{s}_{i\cdot}$ (bottom row), every method delivers negative R^2 at the 12-month horizon—daily cross-maturity prediction at this horizon lies beyond the reach of any specification in our set, regardless of functional form.

Figures C.5–C.7 confirm that the long-horizon pattern observed on the Global panel holds

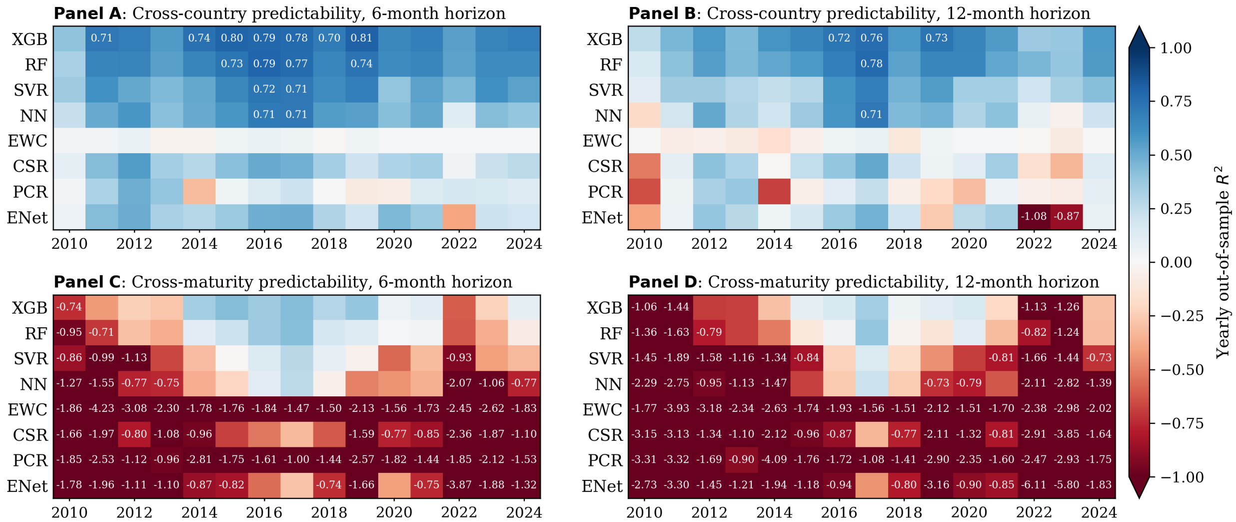

Figure C.4: **Yearly R^2_{oos} across benchmarks at long horizons, Global panel.** The figure reports the annual average of $R^2_{t,oos}$ on the Global panel for the 6-month (left columns) and 12-month (right columns) testing horizons. Rows: cross-country predictability against $\bar{s}_{\cdot j}$ (top) and cross-maturity predictability against $\bar{s}_{i\cdot}$ (bottom). Models and formatting follow Figure C.1.

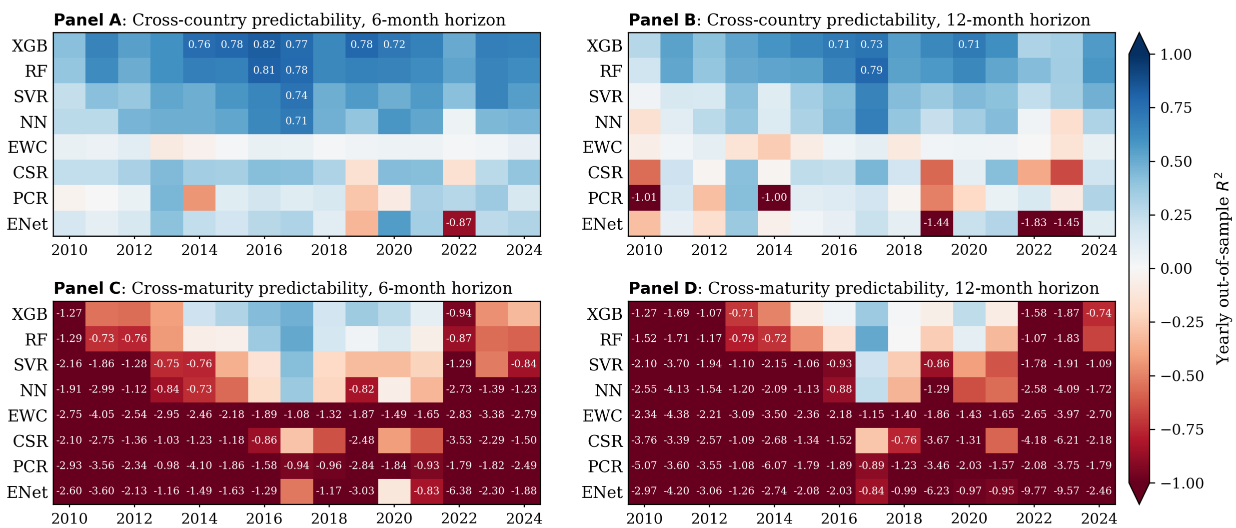

Figure C.5: **Yearly R^2_{oos} across benchmarks at long horizons, OECD non-EU panel.** The figure reports the annual average of $R^2_{t,oos}$ for the OECD non-EU country sub-group at the 6-month (left columns) and 12-month (right columns) testing horizons. Rows, columns, and formatting follow Figure C.4.

across country sub-groups. In each sub-group, the cross-country R^2 against $\bar{s}_{\cdot j}$ remains positive for most years and all nonlinear methods at the 6-month horizon, but linear methods weaken in the 2022–2024 tightening-cycle cells, with SWEAP showing the sharpest linear fragility at the 12-month horizon (Elastic Net in panel (b) of Figure C.7) and OECD non-EU the most resilient. Against $\bar{s}_{i\cdot}$ at 12-month horizon, all methods deliver negative R^2 across

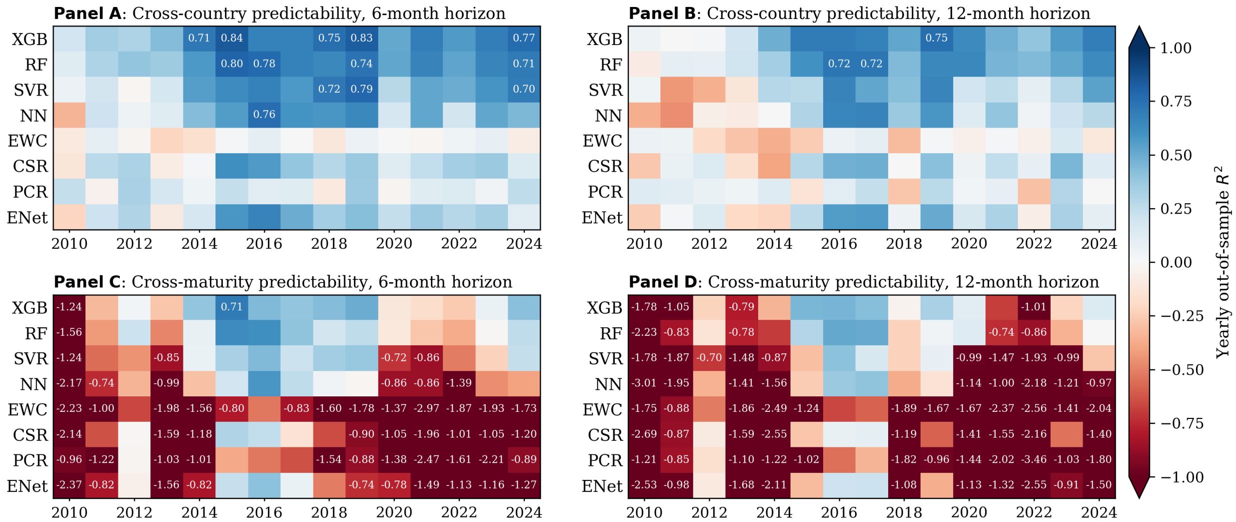

Figure C.6: **Yearly R^2_{oos} at long horizons, EU ex-SWEAP panel.** The figure reports the annual average of $R^2_{t,oos}$ for the EU ex-SWEAP country sub-group at the 6-month and 12-month testing horizons. Rows, columns, and formatting follow Figure C.4.

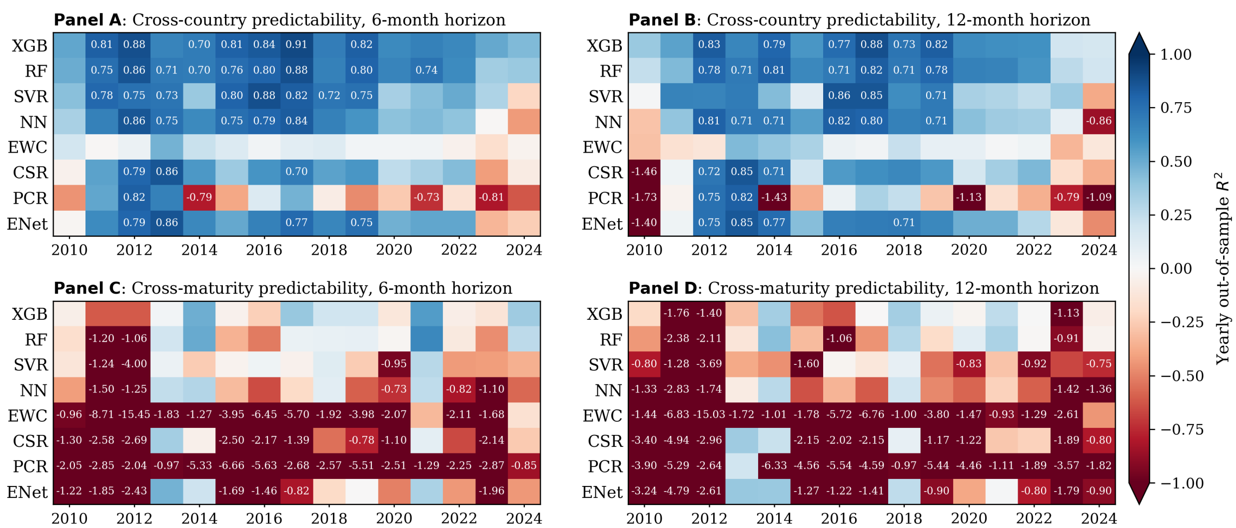

Figure C.7: **Yearly R^2_{oos} at long horizons, SWEAP panel, 2010–2024.** The figure reports the annual average of $R^2_{t,oos}$ for the SWEAP country sub-group at the 6-month and 12-month testing horizons. Rows, columns, and formatting follow Figure C.4.

all country sub-groups, reinforcing the main-text reading that cross-maturity prediction at the 12-month horizon is not within reach of either linear or nonlinear methods on a daily panel of this size.

Horizon	Benchmark	*Level-only linear* OLS	EWC	CSR	Lasso	Ridge	EN	PCR	PLS
1-month	$\bar{s}_{\cdot\cdot}$	0.43	0.19	0.49	0.51	0.47	**0.52**	0.35	0.48
	$\bar{s}_{\cdot j}$	0.35	0.06	0.42	0.43	0.40	**0.45**	0.24	0.40
	$\bar{s}_{i\cdot}$	−1.09	−2.12	−0.85	−0.83	−0.90	**−0.77**	−1.45	−0.97
	$\bar{s}_{ij}$	−5.92	−10.29	−5.49	−5.20	−5.49	**−5.11**	−8.50	−6.04
3-month	$\bar{s}_{\cdot\cdot}$	0.15	0.17	0.43	0.44	0.35	**0.44**	0.30	0.30
	$\bar{s}_{\cdot j}$	0.06	0.04	0.35	0.36	0.26	**0.36**	0.18	0.20
	$\bar{s}_{i\cdot}$	−2.20	−2.13	−1.09	−1.05	−1.37	**−1.02**	−1.58	−1.59
	$\bar{s}_{ij}$	−7.60	−9.86	−5.85	−5.68	−6.39	**−5.65**	−8.48	−7.69
6-month	$\bar{s}_{\cdot\cdot}$	−0.02	0.14	**0.35**	0.35	0.24	0.35	0.22	0.01
	$\bar{s}_{\cdot j}$	−0.13	0.01	**0.26**	0.26	0.13	0.25	0.10	−0.12
	$\bar{s}_{i\cdot}$	−2.48	−2.14	**−1.31**	−1.35	−1.70	−1.34	−1.79	−2.52
	$\bar{s}_{ij}$	−8.55	−9.28	**−6.20**	−6.35	−7.25	−6.37	−8.54	−10.51
12-month	$\bar{s}_{\cdot\cdot}$	−0.27	0.10	**0.22**	0.14	−0.01	0.12	0.08	−0.52
	$\bar{s}_{\cdot j}$	−0.43	−0.03	**0.11**	0.02	−0.15	0.00	−0.06	−0.74
	$\bar{s}_{i\cdot}$	−3.44	−2.23	**−1.79**	−2.14	−2.59	−2.17	−2.23	−4.46
	$\bar{s}_{ij}$	−10.90	−8.66	**−6.99**	−8.05	−9.18	−8.15	−8.84	−16.29

Table C.1: **Four-benchmark aggregate $\overline{R}^2_{oos}$, Global panel, linear "level-only".** The table reports the full-sample average of yearly $R^2_{t,oos}$ (2009–2024) for each level-only linear specification (without maturity interactions), forecasting horizon, and naive benchmark. Benchmarks and horizons are as in Table 1 of the paper; within each row, the best-performing specification is in bold.

C.2 Aggregate Results for Sub-Groups and Linear "Level-Only"

Table C.1 repeats the four-benchmark decomposition for the eight level-only linear specifications on the Global panel. Tables C.2–C.4 report the decomposition separately for the three country sub-groups—OECD non-EU, EU ex-SWEAP, and SWEAP—using the same level+slope linear and nonlinear specifications as the main-text table. The layout of each appendix table mirrors Table 1 in the paper, so the decomposition reading is preserved: scanning a single model column from top to bottom within a horizon block traces the R^2 decomposition as the benchmark becomes more demanding.

The level-only results in Table C.1 closely track the level+slope pattern in Table 1 of the paper. Elastic Net without maturity interactions achieves an average R^2 of 0.52 against $\bar{s}_{\cdot\cdot}$ at 1M, 0.45 against $\bar{s}_{\cdot j}$, and −0.77 against $\bar{s}_{i\cdot}$; the corresponding level+slope values in the main-text table are 0.52, 0.45, and −0.76. The gain from maturity interactions in linear specifications is quantitatively negligible across all four benchmarks and all four horizons, confirming the observation in prior work (e.g., Balduzzi et al., 2023) that simple maturity-variable interactions do not recover the term-structure nonlinearity captured by nonlinear methods. We focus on the level+slope specifications in the main text because they are the more favorable comparison for the linear class.

The three sub-group tables establish that the four-benchmark pattern of the Global panel is preserved across country groups. In every sub-group and at every horizon, linear specifications deliver positive R^2 against $\bar{s}_{\cdot j}$ and negative R^2 against $\bar{s}_{i\cdot}$; nonlinear specifications—XGBoost in particular—retain positive R^2 on the cross-maturity benchmark at short horizons. In addition to the sign pattern being preserved across sub-groups, two further patterns are worth noting.

First, the OECD non-EU panel (Table C.2) yields the lowest linear cross-country R^2 among the three subgroups (Elastic Net at 1M is 0.40, relative to 0.45 for Global and 0.50 for SWEAP), with modest degradation at longer horizons. The smaller magnitudes relative to Global reflect the lower cross-country variance in this subset—OECD non-EU countries have relatively homogeneous macroeconomic fundamentals—which reduces the denominator of the R^2 statistic with respect to $\bar{s}_{\cdot j}$.

Second, the SWEAP panel (Table C.4) exhibits the largest cross-country R^2 values among the sub-groups at short horizons (XGBoost at 1M against $\bar{s}_{\cdot\cdot}$ is 0.82, comparable to the Global panel). This reflects the pronounced cross-country dispersion introduced by peripheral Eurozone sovereigns, which enlarges the R^2 denominator and makes cross-country prediction visibly rewarding. However, the SWEAP panel also shows increasing linear fragility at long horizons: at the 12-month horizon, Elastic Net's full-sample R^2 against $\bar{s}_{\cdot j}$ remains positive (0.32) but its yearly values collapse sharply in the 2022–2024 tightening window (see Figure C.7). Against $\bar{s}_{i\cdot}$ at 12M, Elastic Net averages -1.22, versus -0.34 for XGBoost, preserving the benchmark-based decomposition pattern observed in the Global panel.

C.3 First-Release vs Vintage Data

A concern in real-time forecasting is whether first-release macroeconomic data contain sufficient signal relative to revised vintages. In this subsection, we compare the out-of-sample R^2 of XGBoost trained on first-release information with that of XGBoost trained on the latest available data vintage for each forecast date. Table C.5 reports the results for one-month, three-month, six-month, and 12-month testing horizons, for the Global panel and for each country sub-group (OECD non-EU, EU ex-SWEAP, SWEAP). The naive benchmark is the cross-country or cross-maturity mean, as in the main text.

		Level + Slope linear								*Nonlinear models*				
Horizon	Benchmark	OLS	EWC	CSR	Lasso	Ridge	EN	PCR	PLS	SVR	RT	RF	XGB	NN
1M	$\bar{s}_{..}$	0.37	0.21	**0.52**	0.46	0.43	0.47	0.37	0.44	0.76	0.68	0.79	**0.83**	0.65
	$\bar{s}_{.j}$	0.28	0.08	**0.45**	0.39	0.35	0.40	0.26	0.35	0.72	0.64	0.76	**0.81**	0.60
	$\bar{s}_{i.}$	−1.71	−2.37	**−1.00**	−1.26	−1.39	−1.20	−1.65	−1.27	−0.00	−0.30	0.16	**0.34**	−0.59
	$\bar{s}_{ij}$	−8.35	−11.99	**−6.33**	−7.20	−7.55	−7.02	−9.92	−8.06	−2.27	−3.36	−1.78	**−0.99**	−4.23
3M	$\bar{s}_{..}$	−0.10	0.19	**0.44**	0.37	0.27	0.38	0.30	0.24	0.66	0.59	0.73	**0.77**	0.58
	$\bar{s}_{.j}$	−0.21	0.06	**0.36**	0.28	0.17	0.28	0.19	0.12	0.62	0.53	0.70	**0.74**	0.53
	$\bar{s}_{i.}$	−4.79	−2.36	**−1.34**	−1.59	−2.14	−1.59	−1.83	−2.02	−0.37	−0.66	−0.06	**0.10**	−0.79
	$\bar{s}_{ij}$	−12.63	−11.30	**−6.93**	−7.99	−9.26	−7.96	−9.90	−10.27	−3.44	−4.56	−2.46	**−1.77**	−4.76
6M	$\bar{s}_{..}$	−0.40	0.16	**0.34**	0.24	0.08	0.23	0.20	−0.10	0.56	0.50	0.67	**0.70**	0.49
	$\bar{s}_{.j}$	−0.54	0.03	**0.24**	0.13	−0.05	0.12	0.08	−0.28	0.51	0.43	0.62	**0.66**	0.43
	$\bar{s}_{i.}$	−4.67	−2.36	**−1.65**	−2.05	−2.74	−2.09	−2.11	−3.21	−0.74	−0.99	−0.30	**−0.17**	−1.06
	$\bar{s}_{ij}$	−14.05	−10.72	**−7.76**	−9.29	−11.02	−9.36	−10.17	−14.70	−4.56	−5.46	−3.15	**−2.55**	−5.51
12M	$\bar{s}_{..}$	−0.77	0.12	**0.14**	−0.07	−0.28	−0.09	0.04	−0.84	0.42	0.39	0.58	**0.60**	0.34
	$\bar{s}_{.j}$	−0.97	−0.02	**0.02**	−0.23	−0.46	−0.25	−0.11	−1.14	0.34	0.30	0.52	**0.55**	0.25
	$\bar{s}_{i.}$	−5.84	−2.46	**−2.37**	−3.27	−4.12	−3.34	−2.66	−6.03	−1.31	−1.45	−0.67	**−0.60**	−1.63
	$\bar{s}_{ij}$	−17.91	−10.11	**−9.45**	−12.39	−14.44	−12.55	−10.81	−25.66	−6.03	−6.32	−3.95	**−3.60**	−6.88

Table C.2: **Four-benchmark aggregate $\overline{R}^2_{oos}$, OECD non-EU panel.** The table reports the full-sample average of yearly $R^2_{t,oos}$ (2009–2024) for the OECD non-EU country sub-group. Benchmarks, horizons, models, and formatting follow Table 1 in the paper. Analogous tables for the Global panel and other country sub-groups are reported in Tables C.3 and C.4.

Horizon	Benchmark	*Level + Slope linear* OLS	EWC	CSR	Lasso	Ridge	EN	PCR	PLS	*Nonlinear models* SVR	RT	RF	XGB	NN
1-month	$\bar{s}_{\cdot\cdot}$	0.39	0.22	0.45	0.47	0.44	**0.48**	0.37	0.43	0.72	0.59	0.73	**0.78**	0.55
	$\bar{s}_{\cdot j}$	0.28	0.04	0.34	0.37	0.32	**0.37**	0.22	0.31	0.67	0.52	0.69	**0.74**	0.47
	$\bar{s}_{i\cdot}$	−0.81	−1.51	−0.73	−0.66	−0.79	**−0.64**	−1.06	−0.77	0.30	−0.10	0.28	**0.44**	−0.23
	$\bar{s}_{ij}$	−7.76	−13.39	−8.20	−7.22	−8.06	**−7.10**	−10.61	−8.22	−2.05	−3.32	−1.87	**−1.16**	−4.11
3-month	$\bar{s}_{\cdot\cdot}$	0.27	0.20	0.40	**0.43**	0.37	0.42	0.34	0.29	0.63	0.52	0.67	**0.70**	0.49
	$\bar{s}_{\cdot j}$	0.14	0.01	0.28	**0.32**	0.24	0.31	0.18	0.13	0.57	0.43	0.61	**0.65**	0.40
	$\bar{s}_{i\cdot}$	−1.12	−1.51	−0.84	−0.75	−0.95	**−0.74**	−1.08	−1.13	0.04	−0.38	0.07	**0.20**	−0.38
	$\bar{s}_{ij}$	−8.08	−12.36	−7.97	−7.08	−8.15	**−7.02**	−9.99	−9.21	−2.92	−4.18	−2.55	**−1.94**	−4.45
6-month	$\bar{s}_{\cdot\cdot}$	0.25	0.18	0.36	**0.39**	0.31	0.38	0.29	0.10	0.53	0.43	0.61	**0.64**	0.43
	$\bar{s}_{\cdot j}$	0.11	−0.01	0.23	**0.27**	0.17	0.26	0.13	−0.11	0.44	0.32	0.54	**0.57**	0.33
	$\bar{s}_{i\cdot}$	−1.18	−1.50	−0.91	**−0.82**	−1.06	−0.83	−1.10	−1.61	−0.25	−0.61	−0.09	**0.01**	−0.58
	$\bar{s}_{ij}$	−8.15	−11.34	−7.65	−6.95	−8.21	**−6.94**	−9.36	−10.98	−3.96	−5.06	−3.04	**−2.63**	−5.04
12-month	$\bar{s}_{\cdot\cdot}$	0.13	0.13	0.29	**0.31**	0.20	0.30	0.24	−0.14	0.35	0.29	0.52	**0.54**	0.31
	$\bar{s}_{\cdot j}$	−0.05	−0.08	0.13	**0.16**	0.03	0.15	0.05	−0.40	0.22	0.15	0.43	**0.46**	0.18
	$\bar{s}_{i\cdot}$	−1.61	−1.58	−1.12	**−1.05**	−1.38	−1.08	−1.27	−2.47	−0.82	−1.00	−0.36	**−0.28**	−0.98
	$\bar{s}_{ij}$	−8.74	−10.31	−7.46	**−7.09**	−8.58	−7.13	−9.03	−14.71	−5.53	−6.19	−3.73	**−3.42**	−6.17

Table C.3: **Four-benchmark aggregate $\overline{R}^2_{oos}$, EU ex-SWEAP panel.** The table reports the full-sample average of yearly $R^2_{t,oos}$ (2009–2024) for the EU ex-SWEAP country sub-group. Benchmarks, horizons, models, and formatting follow Table 1 in the paper. Analogous tables for the Global panel and other country sub-groups are reported in Tables C.2 and C.4.

Horizon	Benchmark	*Level + Slope linear* OLS	EWC	CSR	Lasso	Ridge	EN	PCR	PLS	*Nonlinear models* SVR	RT	RF	XGB	NN
1-month	$\bar{s}_{..}$	0.44	0.15	0.48	0.52	0.53	**0.57**	0.08	0.30	0.76	0.67	0.79	**0.82**	0.63
	$\bar{s}_{.j}$	0.37	0.06	0.41	0.46	0.46	**0.50**	−0.07	0.20	0.72	0.61	0.76	**0.80**	0.57
	$\bar{s}_{i.}$	−0.71	−4.02	−1.04	−0.65	**−0.57**	−0.60	−2.73	−1.12	0.06	−0.05	0.26	**0.45**	−0.28
	$\bar{s}_{ij}$	**−5.25**	−15.98	−7.44	−5.37	−5.61	−5.48	−16.28	−7.68	−1.75	−2.89	−1.75	**−0.69**	−3.75
3-month	$\bar{s}_{..}$	0.23	0.14	0.41	0.44	0.41	**0.50**	0.05	−0.03	0.67	0.58	0.74	**0.77**	0.60
	$\bar{s}_{.j}$	0.16	0.06	0.34	0.38	0.33	**0.44**	−0.09	−0.14	0.61	0.51	0.71	**0.75**	0.53
	$\bar{s}_{i.}$	−1.27	−3.91	−1.23	−0.85	−0.92	**−0.78**	−2.83	−1.99	−0.27	−0.30	0.07	**0.23**	−0.38
	$\bar{s}_{ij}$	−5.61	−12.80	−6.31	**−4.71**	−5.33	−4.79	−12.59	−8.09	−2.46	−3.37	−1.85	**−1.12**	−3.33
6-month	$\bar{s}_{..}$	0.02	0.14	0.34	0.40	0.28	**0.46**	−0.04	−0.45	0.60	0.49	0.69	**0.71**	0.54
	$\bar{s}_{.j}$	−0.06	0.06	0.27	0.33	0.21	**0.40**	−0.17	−0.57	0.53	0.41	0.65	**0.68**	0.47
	$\bar{s}_{i.}$	−1.75	−3.71	−1.38	−1.00	−1.24	**−0.92**	−3.05	−3.12	−0.53	−0.60	−0.16	**−0.00**	−0.54
	$\bar{s}_{ij}$	−6.05	−10.44	−5.55	**−4.36**	−5.34	−4.41	−10.88	−9.89	−2.74	−4.15	−2.10	**−1.51**	−3.23
12-month	$\bar{s}_{..}$	−0.01	0.12	0.33	**0.41**	0.26	0.39	−0.22	−0.69	0.50	0.40	0.64	**0.65**	0.49
	$\bar{s}_{.j}$	−0.09	0.04	0.26	**0.34**	0.18	0.32	−0.34	−0.84	0.43	0.32	0.60	**0.61**	0.41
	$\bar{s}_{i.}$	−2.06	−3.45	−1.50	**−1.17**	−1.51	−1.22	−3.46	−4.06	−0.82	−1.04	−0.46	**−0.34**	−0.78
	$\bar{s}_{ij}$	−6.47	−8.45	−4.99	**−4.10**	−5.46	−4.26	−10.15	−11.72	−3.04	−4.41	−2.39	**−2.18**	−3.38

Table C.4: **Four-benchmark aggregate $\overline{R}^2_{oos}$, SWEAP panel.** The table reports the full-sample average of yearly $R^2_{t,oos}$ (2009–2024) for the SWEAP country sub-group. Benchmarks, horizons, models, and formatting follow Table 1 in the paper. Analogous tables for the Global panel and other country sub-groups are reported in Tables C.2 and C.3.

	Cross-country ($\bar{s}_{\cdot j}$)		Cross-maturity ($\bar{s}_{i\cdot}$)	
	First release	Vintages	First release	Vintages
1-month testing				
Global	0.890	0.893	0.631	0.651
OECD non-EU	0.905	0.906	0.675	0.681
EU ex-SWEAP	0.846	0.859	0.621	0.660
SWEAP	0.881	0.883	0.362	0.345
3-month testing				
Global	0.808	0.810	0.376	0.389
OECD non-EU	0.825	0.826	0.409	0.419
EU ex-SWEAP	0.747	0.758	0.353	0.400
SWEAP	0.831	0.820	0.096	0.070
6-month testing				
Global	0.723	0.720	0.101	0.116
OECD non-EU	0.746	0.746	0.150	0.161
EU ex-SWEAP	0.644	0.653	0.052	0.145
SWEAP	0.769	0.729	−0.194	−0.341
12-month testing				
Global	0.588	0.590	−0.316	−0.284
OECD non-EU	0.611	0.624	−0.294	−0.225
EU ex-SWEAP	0.487	0.477	−0.367	−0.320
SWEAP	0.663	0.665	−0.596	−0.651

Table C.5: **XGBoost with first-release vs. vintage data.** The table compares the R^2_{oos} of XGBoost trained on first-release macroeconomic data against XGBoost trained on the latest available data vintage. Results are reported for 1-month, 3-month, 6-month, and 12-month testing horizons, for the Global panel and for each country sub-group. The R^2_{oos} values are obtained by averaging the period-by-period $R^2_{t,oos}$ over the full out-of-sample period. The naive benchmark is the average of the log CDS spread across countries (left panel) or across maturities (right panel).

References

Balduzzi, P., Savona, R., and Alessi, L. (2023). Anatomy of a sovereign debt crisis: machine learning, real-time macro fundamentals, and cds spreads. *Journal of Financial Econometrics*, 21(5):1728–1758.

Bianchi, D., Büchner, M., and Tamoni, A. (2021). Bond risk premiums with machine learning. *The Review of Financial Studies*, 34(2):1046–1089.

Chen, T. and Guestrin, C. (2016). Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pages 785–794.

Cortes, C. and Vapnik, V. (1995). Support-vector networks. *Machine learning*, 20(3):273–297.

Drucker, H., Burges, C. J., Kaufman, L., Smola, A., and Vapnik, V. (1996). Support vector regression machines. *Advances in neural information processing systems*, 9.

Gu, S., Kelly, B., and Xiu, D. (2020). Empirical asset pricing via machine learning. *The Review of Financial Studies*, 33(5):2223–2273.

Hadri, K. (2000). Testing for stationarity in heterogeneous panel data. *The econometrics journal*, 3(2):148–161.

Hsiang, T. (1975). A bayesian view on ridge regression. *Journal of the Royal Statistical Society: Series D*, 24(4):267–268.

Kwiatkowski, D., Phillips, P. C., Schmidt, P., and Shin, Y. (1992). Testing the null hypothesis of stationarity against the alternative of a unit root: How sure are we that economic time series have a unit root? *Journal of econometrics*, 54(1-3):159–178.

Pesaran, M. H. (2007). A simple panel unit root test in the presence of cross-section dependence. *Journal of applied econometrics*, 22(2):265–312.

Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 267–288.

Zou, H. and Hastie, T. (2005). Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2):301–320.