



Contents lists available at ScienceDirect

International Journal of Forecasting

journal homepage: www.elsevier.com/locate/ijforecast

Cross-temporal probabilistic forecast reconciliation: Methodological and practical issues

Daniele Girolimetto^{a,*}, George Athanasopoulos^b, Tommaso Di Fonzo^a, Rob J. Hyndman^b

^a Department of Statistical Sciences, University of Padua, Padova 35121, Italy

^b Department of Econometrics & Business Statistics, Monash University, Clayton VIC 3800, Australia

ARTICLE INFO

Keywords:

Forecast reconciliation
Linearly constrained multiple time series
Cross-temporal
Probabilistic forecasting
GDP
Tourism flows

ABSTRACT

Forecast reconciliation is a post-forecasting process that involves transforming a set of incoherent forecasts into coherent forecasts which satisfy a given set of linear constraints for a multivariate time series. In this paper, we extend the current state-of-the-art cross-sectional probabilistic forecast reconciliation approach to encompass a cross-temporal framework, where temporal constraints are also applied. Our proposed methodology employs both parametric Gaussian and non-parametric bootstrap approaches to draw samples from an incoherent cross-temporal distribution. To improve the estimation of the forecast error covariance matrix, we propose using multi-step residuals, especially in the time dimension where the usual one-step residuals fail. To address high-dimensionality issues, we present four alternatives for the covariance matrix, where we exploit the two-fold nature (cross-sectional and temporal) of the cross-temporal structure, and introduce the idea of overlapping residuals. We assess the effectiveness of the proposed cross-temporal reconciliation approaches through a simulation study that investigates their theoretical and empirical properties and two forecasting experiments, using the Australian GDP and the Australian Tourism Demand datasets. For both applications, the optimal cross-temporal reconciliation approaches significantly outperform the incoherent base forecasts in terms of the continuous ranked probability score and the energy score. Overall, the results highlight the potential of the proposed methods to improve the accuracy of probabilistic forecasts and to address the challenge of integrating disparate scenarios while coherently taking into account short-term operational, medium-term tactical, and long-term strategic planning.

© 2023 The Author(s). Published by Elsevier B.V. on behalf of International Institute of Forecasters. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Forecast reconciliation is a post-forecasting process intended to improve the quality of forecasts for a system of linearly constrained multiple time series (Hyndman et al., 2011, Panagiotelis et al., 2021). There are many fields where forecast reconciliation is useful, such as when

forecasting demand in supply chains with product categories (Punia et al., 2020, Kourentzes & Athanasopoulos, 2021), electricity demand and power generation (Spiliotis et al., 2020, Ben Taieb et al., 2021), GDP and its components (Athanasopoulos et al., 2020), tourist flows across geographic regions and travel purposes (Kourentzes & Athanasopoulos, 2019), and more. Moreover, effective decision-making depends on the support of accurate and coherent forecasts, making the use of forecast reconciliation methods increasingly popular in recent years (Athanasopoulos et al., 2023).

* Corresponding author.

E-mail address: daniele.girolimetto@unipd.it (D. Girolimetto).

Temporal reconciliation is another important aspect of forecast reconciliation that can help organizations to better align their forecasting efforts. This approach consists in reconciling forecasts that are generated at different time horizons, such monthly, quarterly, or annual. For example, a retail company may need to reconcile monthly forecasts of sales with quarterly forecasts of revenue to ensure that they are aligned and consistent.

Classical reconciliation approaches (bottom-up, top-down, middle-out; see [Dunn et al., 1976](#), [Gross & Sohl, 1990](#), [Athanasopoulos et al., 2009](#), respectively) addressed the issue of incoherent forecasts in a cross-sectional hierarchy by forecasting only one level, and using these to generate forecasts for the remaining series. All of these approaches ignore useful information available at other levels ([Pennings & van Dalen, 2017](#)). Recently, hierarchical forecasting ([Fliedner, 2001](#)) has significantly evolved to include modern least-squares-based reconciliation techniques in the cross-sectional framework ([Hyndman et al., 2011](#), [Wickramasuriya et al., 2019](#), [Panagiotelis et al., 2021](#)), later extended to temporal hierarchies ([Athanasopoulos et al., 2017](#), [Nystrup et al., 2020](#)). Obtaining coherent forecasts across both the cross-sectional and temporal dimensions (known as cross-temporal coherence) has been limited to sequential approaches that address each dimension separately ([Kourentzes & Athanasopoulos, 2019](#), [Yagli et al., 2019](#), [Punia et al., 2020](#), [Spiliotis et al., 2020](#)). Recently, [Di Fonzo and Girolimetto \(2023a\)](#) suggested a unified reconciliation step that takes into account both the cross-sectional and temporal dimensions, instead of dealing with them separately, utilizing the entire cross-temporal hierarchy.

However, these cross-temporal works focus on point forecasting, and do not consider distributional or probabilistic forecasts ([Gneiting & Katzfuss, 2014](#)). In the cross-sectional and temporal frameworks, there have been some developments towards probabilistic forecasting, including [Ben Taieb et al. \(2017, 2021\)](#), [Panamtash and Zhou \(2018\)](#), [Jeon et al. \(2019\)](#), [Yang \(2020\)](#), [Yagli et al. \(2020\)](#), [Corani et al. \(2021\)](#), [Corani et al. \(2023\)](#), [Zambon et al. \(2022\)](#), and [Wickramasuriya \(2023\)](#). [Panagiotelis et al. \(2023\)](#) made a significant contribution by formalizing cross-sectional probabilistic reconciliation using the geometric framework for point forecast reconciliation of [Panagiotelis et al. \(2021\)](#). They show how a reconciled forecast can be constructed from an arbitrary base forecast when its density is available and when only a sample can be drawn. They also show that in the case of elliptical distributions, the correct predictive distribution can be recovered via linear reconciliation, regardless of the base forecast location and scale parameters, and derive conditions for this to hold in the special case of reconciliation via projection.

In this paper, we extend cross-sectional probabilistic reconciliation to the cross-temporal case, working on issues related to the two-fold nature of this framework. First, we revise and develop the notation proposed by [Di Fonzo and Girolimetto \(2023a\)](#) to generalize the work of [Panagiotelis et al. \(2023\)](#). This allows us to move from cross-temporal point reconciliation to a probabilistic setting through the generalization of definitions and

theorems well established in the cross-sectional framework. Second, we propose solutions to draw a sample from the base forecast distribution according to either a parametric approach that assumes Gaussianity or a non-parametric approach that bootstraps the base model residuals. Third, we propose some solutions to specific problems that arise when combining the cross-sectional and temporal dimensions. We propose using multi-step residuals to estimate the relationships between different forecast horizons when we deal with temporal levels, since one-step residuals are not suitable for this purpose. To solve high-dimensionality issues we introduce the idea of overlapping residuals and consider alternative forms for constructing the covariance matrix. Fourth, we propose new shrinkage procedures for reconciliation that aim to identify a feasible cross-temporal structure. The algorithms described in this paper are implemented in the FoReco package ([Girolimetto and Di Fonzo, 2023a](#)) for R ([R Core Team, 2022](#)). Furthermore, the online appendix contains complementary materials on methodological and practical issues, and supplementary tables and graphs related to the empirical applications.

The remainder of the paper is structured as follows. In Section 2, we provide a unified notation for the cross-sectional, temporal, and cross-temporal point reconciliation. We generalize the cross-sectional definitions and theorems developed by [Panagiotelis et al. \(2023\)](#) in Section 3, and propose both a parametric Gaussian and a non-parametric bootstrap approach to draw a sample from the base forecast distribution. In Section 4, we analyze the structure of the cross-temporal covariance matrix, proposing four alternative forms, and propose shrinkage approaches for reconciliation. In addition, we explore cross-temporal residuals (overlapping and multi-step) looking at their advantages and limitations. Two empirical applications using the Australian GDP and the Australian Tourism Demand datasets are considered in Sections 5 and 6, respectively.¹ Finally, Section 7 presents conclusions and a future research agenda on this and other related topics.

2. Notation and definitions

Let $\mathbf{y}_t = [y_{1,t}, \dots, y_{i,t}, \dots, y_{n,t}]'$ be an n -variate linearly constrained time series observed at the most temporally disaggregated level, with a seasonality of period m (e.g., $m = 12$ for monthly data, $m = 4$ for quarterly data, $m = 24$ for hourly data). Suppose that the constraints are expressed by linear equations such that ([Di Fonzo & Girolimetto, 2023a](#))

$$\mathbf{C}_{cs}\mathbf{y}_t = \mathbf{0}_{(n_a \times 1)}, \quad t = 1, \dots, T, \quad (1)$$

where \mathbf{C}_{cs} is the $(n_a \times n)$ zero-constraints cross-sectional matrix that can be seen as the coefficient matrix of a linear system with n_a equations and n variables.²

An example is a hierarchical time series where series at upper levels can be expressed by appropriately

¹ A complete set of results is available at the GitHub repository <https://github.com/danigiro/ctprob>.

² [Hyndman \(2022\)](#) and [Girolimetto and Di Fonzo \(2023b\)](#) show that this zero-constrained representation is more general and computationally efficient.

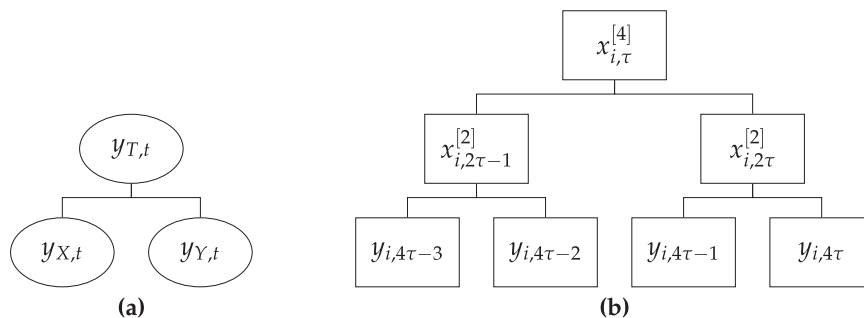


Fig. 1. (a) A simple two-level cross-sectional hierarchy for 3 time series with $n_a = 1$ and $n_b = 2$. (b) A temporal hierarchy for a quarterly series ($m = 4$ and $\kappa = \{4, 2, 1\}$).

summing part or all of the series at the bottom level. Fig. 1(a) shows the two-level hierarchical structure for three linearly constrained time series such that $y_{T,t} = y_{X,t} + y_{Y,t}, \forall t = 1, \dots, T$. Now let $\mathbf{y}_t = [\mathbf{u}'_t \ \mathbf{b}'_t]'$, where $\mathbf{u}_t = [y_{1,t}, \dots, y_{n_a,t}]'$ is the n_a -vector of upper-level time series, and $\mathbf{b}_t = [y_{(n_a+1),t} \ \dots \ y_{n,t}]'$ is the n_b -vector of bottom-level time series with $n = n_a + n_b$. The upper- and lower-level time series are connected by the cross-sectional aggregation matrix \mathbf{A}_{cs} such that $\mathbf{u}_t = \mathbf{A}_{cs}\mathbf{b}_t$. Following Girolimetto and Di Fonzo (2023b), we can always construct a zero-constraints cross-sectional matrix from the aggregation matrix, $\mathbf{C}_{cs} = [\mathbf{I}_{n_a} \ -\mathbf{A}_{cs}]$, where \mathbf{I}_{n_a} is an identity matrix of dimension n_a . Finally, the cross-sectional structural matrix is given by $\mathbf{S}_{cs} = \begin{bmatrix} \mathbf{A}_{cs} \\ \mathbf{I}_{n_b} \end{bmatrix}$, providing the structural representation (Hyndman et al., 2011) $\mathbf{y}_t = \mathbf{S}_{cs}\mathbf{b}_t$. Considering the hierarchical example in Fig. 1(a), we have

$$\mathbf{A}_{cs} = \begin{bmatrix} 1 & 1 \\ 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad \mathbf{C}_{cs} = \begin{bmatrix} 1 & -1 & -1 \end{bmatrix} \quad \text{and}$$

$$\mathbf{S}_{cs} = \begin{bmatrix} 1 & 1 \\ 1 & 0 \\ 0 & 1 \end{bmatrix}.$$

In general, there is no reason for \mathbf{u}_t to be restricted to simple sums of \mathbf{b}_t . Therefore, $\mathbf{A}_{cs} \in \mathbb{R}^{n_a \times n_b}$ may contain any real values, and not only 0s and 1s.

Considering now the temporal framework, we denote as $\kappa = \{k_p, k_{p-1}, \dots, k_2, k_1\}$ the set of p factors of m , in descending order, where $k_1 = 1$ and $k_p = m$ (Athanasopoulos et al., 2017). For example, for quarterly time series $m = 4$, $p = 3$, and $\kappa = \{4, 2, 1\}$. Given a factor k of m , and assuming that $T = Nm$ (where N is the length of the most temporally aggregated version of the series), we can construct a temporally aggregated version of the time series of a single variable $\{y_{i,t}\}_{t=1, \dots, T}$, through the non-overlapping sums of its k successive values, which has a

$$\text{seasonal period equal to } M_k = \frac{m}{k}: x_{i,j}^{[k]} = \sum_{t=(j-1)k+1}^{jk} y_{i,t},$$

where $j = 1, \dots, N_k, i = 1, \dots, n, N_k = \frac{T}{k}$, and $x_{i,j}^{[1]} = y_{i,t}$. Define τ as the observation index of the most aggregate level k_p . For a fixed temporal aggregation order $k \in \kappa$, we stack the observations in the column vector $\mathbf{x}_{i,\tau}^{[k]} = [x_{i,M_k(\tau-1)+1}^{[k]} \ x_{i,M_k(\tau-1)+2}^{[k]} \ \dots \ x_{i,M_k\tau}^{[k]}]'$, and obtain the

vector for all the temporal aggregation orders $\mathbf{x}_{i,\tau} = [x_{i,\tau}^{[k_p]} \ \mathbf{x}_{i,\tau}^{[k_{p-1}]} \ \dots \ \mathbf{x}_{i,\tau}^{[1]}]'$, $\tau = 1, \dots, N$. The structural representation of the temporal hierarchy (Athanasopoulos et al., 2017) is then $\mathbf{x}_{i,\tau} = \mathbf{S}_{te}\mathbf{x}_{i,\tau}^{[1]}$, where $\mathbf{S}_{te} = \begin{bmatrix} \mathbf{A}_{te} \\ \mathbf{I}_m \end{bmatrix}$ is the $[(m + k^*) \times m]$ temporal structural matrix, $\mathbf{A}_{te} = \begin{bmatrix} \mathbf{1}_{k_p} & \mathbf{I}_{\frac{m}{k_{p-1}}} \otimes \mathbf{1}_{k_{p-1}} & \dots & \mathbf{I}_{\frac{m}{k_2}} \otimes \mathbf{1}_{k_2} \end{bmatrix}'$ is the $(k^* \times m)$ temporal aggregation matrix with $k^* = \sum_{k \in \kappa \setminus \{k_1\}} M_k$, the number of upper time series of the temporal hierarchy, $\mathbf{1}_{k_p}$ is a $(k_p \times 1)$ vector of all ones, and \otimes is the Kronecker product. For each series $x_{i,\tau}, i = 1, \dots, n$, we have also the zero-constrained representation

$$\mathbf{C}_{te}\mathbf{x}_{i,\tau} = \mathbf{0}_{[k^* \times (m+k^*)]}, \quad \tau = 1, \dots, N, \quad i = 1, \dots, n, \quad (2)$$

where $\mathbf{C}_{te} = [\mathbf{I}_{k^*} \ -\mathbf{A}_{te}]$ is the $[k^* \times (m + k^*)]$ zero-constraints temporal matrix. Fig. 1(b) shows the hierarchical representation of a quarterly time series, for which $m = 4, \kappa = \{4, 2, 1\}$, and

$$\mathbf{A}_{te} = \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \end{bmatrix}, \quad \mathbf{S}_{te} = \begin{bmatrix} \mathbf{A}_{te} \\ \mathbf{I}_4 \end{bmatrix} \quad \text{and}$$

$$\mathbf{C}_{te} = \begin{bmatrix} 1 & 0 & 0 & -1 & -1 & -1 & -1 \\ 0 & 1 & 0 & -1 & -1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & -1 & -1 \end{bmatrix}.$$

When we temporally aggregate each series, the cross-sectional constraints for the most temporally disaggregated series (1) hold for all the temporal aggregation orders such that $\mathbf{C}_{cs}\mathbf{x}_j^{[k]} = \mathbf{0}_{(n_a \times 1)}$ for $k \in \kappa$ and $j = 1, \dots, N_k$, where $\mathbf{x}_j^{[k]} = [\mathbf{u}_j^{[k]'} \ \mathbf{b}_j^{[k]'}]'$ with $\mathbf{u}_j^{[k]} = [x_{1,j}^{[k]} \ \dots \ x_{n_a,j}^{[k]}]'$ is the n_a -vector of upper time series, and $\mathbf{b}_j^{[k]} = [x_{(n_a+1),j}^{[k]} \ \dots \ x_{n,j}^{[k]}]'$ is the n_b -vector of bottom time series in the temporal hierarchy.

To include both cross-sectional and temporal constraints at the same time in a unified framework, we stack the series into a $[n \times (m + k^*)]$ matrix \mathbf{X}_τ , where we recall that n, m , and k^* respectively represent the total number of time series, the seasonal period, and the number of upper time series of the temporal hierarchy. The rows and columns represent the cross-sectional and temporal

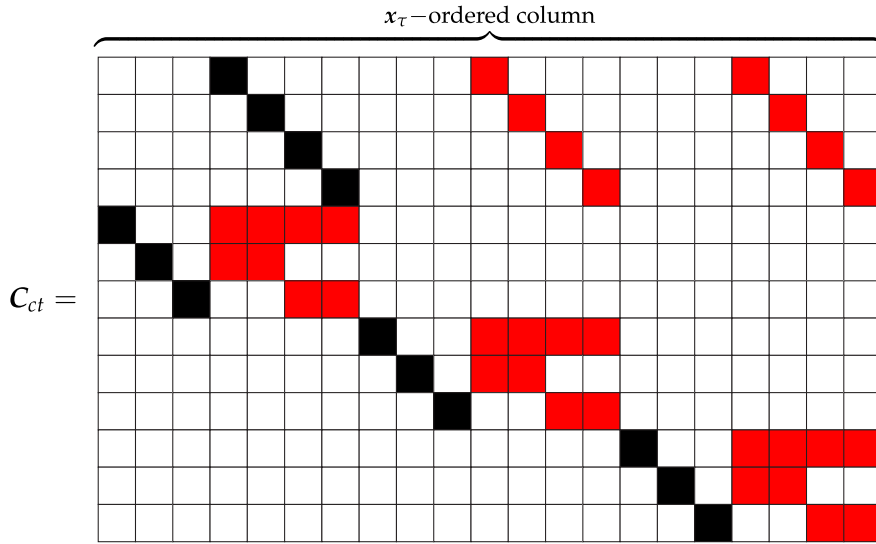


Fig. 2. Visual representation of the zero-constraints cross-temporal matrix C_{ct} defined in (3) for a system of 3 linearly constrained quarterly time series (see Fig. 1). The four upper rows describe the cross-sectional constraints (one for each quarter), the remaining rows the temporal constraints (one for each of the three time series). Colours legend: 0s in white, 1s in black, -1 s in red.

dimensions, respectively:

$$\mathbf{X}_\tau = \begin{bmatrix} \mathbf{x}'_{1,\tau} \\ \vdots \\ \mathbf{x}'_{n,\tau} \end{bmatrix} = \begin{bmatrix} \mathbf{U}_\tau^{[k_p]} & \mathbf{U}_\tau^{[k_p-1]} & \dots & \mathbf{U}_\tau^{[1]} \\ \mathbf{B}_\tau^{[k_p]} & \mathbf{B}_\tau^{[k_p-1]} & \dots & \mathbf{B}_\tau^{[1]} \end{bmatrix},$$

where for any fixed k , $\mathbf{U}_\tau^{[k]}$ is the $(n_a \times N_k)$ matrix grouping the upper time series, and $\mathbf{B}_\tau^{[k]}$ is the $(n_b \times N_k)$ matrix grouping the bottom time series. For example, for the cross-temporal structure of Fig. 1, we have for $\tau = 1$

$$\mathbf{X}_1 = \begin{bmatrix} x_{T,1}^{[4]} & x_{T,1}^{[2]} & x_{T,2}^{[2]} & y_{T,1} & y_{T,2} & y_{T,3} & y_{T,4} \\ x_{X,1}^{[4]} & x_{X,1}^{[2]} & x_{X,2}^{[2]} & y_{X,1} & y_{X,2} & y_{X,3} & y_{X,4} \\ x_{Y,1}^{[4]} & x_{Y,1}^{[2]} & x_{Y,2}^{[2]} & y_{Y,1} & y_{Y,2} & y_{Y,3} & y_{Y,4} \end{bmatrix},$$

for $\tau = 2$

$$\mathbf{X}_2 = \begin{bmatrix} x_{T,2}^{[4]} & x_{T,3}^{[2]} & x_{T,4}^{[2]} & y_{T,5} & y_{T,6} & y_{T,7} & y_{T,8} \\ x_{X,2}^{[4]} & x_{X,3}^{[2]} & x_{X,4}^{[2]} & y_{X,5} & y_{X,6} & y_{X,7} & y_{X,8} \\ x_{Y,2}^{[4]} & x_{Y,3}^{[2]} & x_{Y,4}^{[2]} & y_{Y,5} & y_{Y,6} & y_{Y,7} & y_{Y,8} \end{bmatrix},$$

and so on. Further, $\mathbf{C}_{cs}\mathbf{X}_\tau = \mathbf{0}_{[n_a \times (m+k^*)]}$ and $\mathbf{C}_{te}\mathbf{X}'_\tau = \mathbf{0}_{(k^* \times n)}$. We can consider the cross-temporal framework as a generalization of the cross-sectional and temporal frameworks that simultaneously takes into account both types of constraints. The cross-sectional reconciliation approach proposed by Hyndman et al. (2011) can be obtained by assuming $m = 1$, while the temporal one (Athanasopoulos et al., 2017) is obtained when $n = 1$ (with $n_a = 0$ and $n_b = 1$).

Di Fonzo and Girolimetto (2023a) show that the cross-temporal constraints working on the complete set of observations corresponding to time period $\tau = 1, \dots, N$ can be expressed in a zero-constrained representation through the full rank $[(n_a m + nk^*) \times n(m + k^*)]$

zero-constraints cross-temporal matrix C_{ct} , such that

$$\mathbf{C}_{ct} = \begin{bmatrix} \mathbf{C}_* \\ \mathbf{I}_n \otimes \mathbf{C}_{te} \end{bmatrix} \implies \mathbf{C}_{ct}\mathbf{x}_\tau = \mathbf{0}_{[(n_a m + nk^*) \times 1]}, \quad (3)$$

where $\mathbf{x}_\tau = \text{vec}(\mathbf{X}'_\tau) = [\mathbf{x}'_{1,\tau}, \dots, \mathbf{x}'_{n,\tau}]'$, $\mathbf{C}_* = [\mathbf{0}_{(n_a m \times nk^*)} \mathbf{I}_m \otimes \mathbf{C}_{cs}] \mathbf{P}'$, \mathbf{P} is the commutation matrix (Magnus & Neudecker, 2019, p. 54) such that $\mathbf{P}\text{vec}(\mathbf{X}_\tau) = \text{vec}(\mathbf{X}'_\tau)$, and the operator $\text{vec}(\cdot)$ converts a matrix into a vector. Fig. 2 shows a visual example for the zero-constraints cross-temporal matrix. A structural representation can be considered as well: $\mathbf{x}_\tau = \mathbf{S}_{ct}\mathbf{b}_\tau^{[1]} = s(\mathbf{b}_\tau^{[1]})$, where

$$\mathbf{S}_{ct} = \mathbf{S}_{cs} \otimes \mathbf{S}_{te} \quad (4)$$

is the $[n(k^* + m) \times n_b m]$ cross-temporal summation matrix, $s : \mathbb{R}^{n_b m} \rightarrow \mathbb{R}^{n(m+k^*)}$ is the operator describing the pre-multiplication by \mathbf{S}_{ct} , and $\mathbf{b}_\tau^{[1]} = \text{vec}(\mathbf{B}_\tau^{[1]'})$. In Fig. 3, we represent \mathbf{S}_{ct} for a system of three linearly constrained quarterly time series (see Fig. 1). In agreement with Panagiotelis et al. (2021), \mathbf{x}_τ lies in an $(n_b m)$ -dimensional subspace s_{ct} of $\mathbb{R}^{n(k^*+m)}$, which we refer to as the cross-temporal coherent subspace, spanned by the columns of \mathbf{S}_{ct} .

2.1. Optimal point forecast reconciliation

For $h = 1, \dots, H$, let

$$\widehat{\mathbf{X}}_h = \begin{bmatrix} \widehat{\mathbf{x}}_{1,h} \\ \vdots \\ \widehat{\mathbf{x}}_{n,h} \end{bmatrix} = \begin{bmatrix} \widehat{\mathbf{U}}_h^{[m]} & \dots & \widehat{\mathbf{U}}_h^{[k]} & \dots & \widehat{\mathbf{U}}_h^{[1]} \\ \widehat{\mathbf{B}}_h^{[m]} & \dots & \widehat{\mathbf{B}}_h^{[k]} & \dots & \widehat{\mathbf{B}}_h^{[1]} \end{bmatrix},$$

be the h -step-ahead base forecasts, where $\widehat{\mathbf{U}}_h^{[k]}$ is the $(n_a \times M_k)$ matrix grouping the upper time series, $\widehat{\mathbf{B}}_h^{[k]}$ is the $(n_b \times M_k)$ matrix grouping the bottom time series for a given temporal aggregation order k , and H is the forecast




Fig. 3. Visual representation of the cross-temporal summation matrix $S_{ct} = S_{cs} \otimes S_{te}$ defined in (4) for a system of 3 linearly constrained quarterly time series (see Fig. 1). Colours legend: 0s in white, 1s in black.

horizon for the most temporally aggregated time series. Based on the example in Fig. 1 for $H = 1$, we have that

$$\hat{\mathbf{x}}_1 = \begin{bmatrix} \hat{x}_{T,1}^{[4]} & \hat{x}_{T,1}^{[2]} & \hat{x}_{T,2}^{[2]} & \hat{y}_{T,1} & \hat{y}_{T,2} & \hat{y}_{T,3} & \hat{y}_{T,4} \\ \hat{x}_{X,1}^{[4]} & \hat{x}_{X,1}^{[2]} & \hat{x}_{X,2}^{[2]} & \hat{y}_{X,1} & \hat{y}_{X,2} & \hat{y}_{X,3} & \hat{y}_{X,4} \\ \hat{x}_{Y,1}^{[4]} & \hat{x}_{Y,1}^{[2]} & \hat{x}_{Y,2}^{[2]} & \hat{y}_{Y,1} & \hat{y}_{Y,2} & \hat{y}_{Y,3} & \hat{y}_{Y,4} \end{bmatrix}.$$

The matrix $\hat{\mathbf{x}}_h$ contains incoherent forecasts, such as $C_{ct}\hat{\mathbf{x}}_h \neq \mathbf{0}_{[(n+m+nk^*) \times 1]}$ with $h = 1, \dots, H$ and $\hat{\mathbf{x}}_h = \text{vec}(\hat{\mathbf{X}}_h)$. In this framework, the definition for forecast reconciliation in the cross-sectional framework given by Panagiotelis et al. (2021) can be generalized as follows:

Definition 2.1. Forecast reconciliation adjusts the base forecast $\hat{\mathbf{x}}_h$ by finding a mapping $\psi : \mathbb{R}^{n(m+k^*)} \rightarrow \mathfrak{s}$ such that $\tilde{\mathbf{x}}_h = \psi(\hat{\mathbf{x}}_h)$, where $\tilde{\mathbf{x}}_h \in \mathfrak{s}$ is the vector of the reconciled forecasts.

For a given forecast horizon $h = 1, \dots, H$, the mapping ψ may be defined as a projection onto \mathfrak{s} given by (Panagiotelis et al., 2021, Di Fonzo & Girolimetto, 2023a):

$$\tilde{\mathbf{x}}_h = \psi(\hat{\mathbf{x}}_h) = \mathbf{M}\hat{\mathbf{x}}_h, \tag{5}$$

where $\mathbf{M} = \mathbf{I}_{n(m+k^*)} - \Omega_{ct}C'_{ct}(C_{ct}\Omega_{ct}C'_{ct})^{-1}C_{ct}$ for a positive definite matrix Ω_{ct} , and $\tilde{\mathbf{x}}_h = \text{vec}(\tilde{\mathbf{X}}_h)$. Wickramasuriya et al. (2019) showed that the minimum variance

linear unbiased reconciled forecasts satisfying the unbiasedness condition $E(\tilde{\mathbf{x}}_h - \mathbf{x}_h) = \mathbf{0}$ has solution (5) when $\Omega_{ct} = \text{Var}(\hat{\mathbf{x}}_h - \mathbf{x}_h)$.

Alternatively, the cross-temporal reconciled forecasts $\tilde{\mathbf{x}}_h$ may be found according to the structural approach proposed by Hyndman et al. (2011) for the cross-sectional framework, yielding $\tilde{\mathbf{x}}_h = S_{ct}\mathbf{G}\hat{\mathbf{x}}_h$ for some matrix \mathbf{G} . Wickramasuriya et al. (2019) showed that this leads to a solution equivalent to the cross-temporally reconciled forecasts in (5), given by

$$\tilde{\mathbf{x}}_h = \psi(\hat{\mathbf{x}}_h) = (s \circ g)(\hat{\mathbf{x}}_h) = S_{ct}\mathbf{G}\hat{\mathbf{x}}_h, \tag{6}$$

where $\mathbf{G} = (S'_{ct}\Omega_{ct}^{-1}S_{ct})^{-1}S'_{ct}\Omega_{ct}^{-1}$, and $\mathbf{M} = S_{ct}\mathbf{G}$. In this case, ψ is the composition of two transformations, say $s \circ g$, where $g : \mathbb{R}^{n(m+k^*)} \rightarrow \mathbb{R}^{n_b m}$ is a continuous function. In Online Appendix A, we report some cross-sectional, temporal, and cross-temporal approximations for the covariance matrix to be used in (5) and (6).

2.2. Cross-temporal bottom-up forecast reconciliation

The classic bottom-up approach (Dunn et al., 1976, Dangerfield & Morris, 1992) consists in simply summing up the base forecasts of the most disaggregated level in the hierarchy to obtain forecasts of the upper-level




Fig. 4. A visual representation of partly bottom-up starting from (a) cross-sectionally reconciled forecasts for the temporal order 1 ($\tilde{U}^{[1]}$ and $\tilde{B}^{[1]}$) followed by temporal bottom-up, and (b) temporally reconciled forecasts of the cross-sectional bottom time series ($\tilde{B}^{[k]}$, $k \in \mathcal{K}$) followed by cross-sectional bottom-up. The blue background indicates generating reconciled forecasts along one dimension, while the pink background indicates the forecasts obtained using bottom-up along the other.

series. To reduce the computational cost involved in optimal cross-temporal reconciliation, we may be interested in applying a reconciliation along only one dimension (cross-sectional or temporal) and reconstructing the cross-temporal structure using a partly bottom-up approach (Di Fonzo & Girolimetto, 2022a, 2023b, Sanguri et al., 2022).

Fig. 4 provides a visual representation of partly bottom-up in a two-step cross-temporal reconciliation approach. On the left (Fig. 4(a)), we first compute the cross-sectionally reconciled forecasts at the highest frequency ($k = 1$) and then apply temporal bottom-up to obtain coherent cross-temporal forecasts. On the right (Fig. 4(b)), we first compute temporally reconciled forecasts for the most disaggregated cross-sectional level, and then apply the cross-sectional bottom-up. We respectively denote these two-step reconciliation approaches as $ct(rec_{te}, bu_{cs})$ and $ct(rec_{cs}, bu_{te})$, where ‘ rec_{te} ’ and ‘ rec_{cs} ’ denote a forecast reconciliation approach in the temporal and cross-sectional dimensions, and ‘ bu_{cs} ’ and ‘ bu_{te} ’ denote using bottom-up in the cross-sectional and temporal dimensions, respectively. It is worth noting that the simple cross-temporal bottom-up approach corresponds to $ct(bu_{cs}, bu_{te}) = ct(bu_{te}, bu_{cs}) = ct(bu)$.

3. Probabilistic forecast reconciliation

To introduce the idea of coherence and probabilistic forecast reconciliation, we adapt the notations and the formal definitions introduced in Wickramasuriya (2023) and Panagiotelis et al. (2023) for the cross-sectional probabilistic case. These definitions can also be generalized to the cross-temporal framework by following the approach developed by Corani et al. (2023) for count data. However, in this paper we only focus on the continuous case.

Our aim is to extend these definitions to *cross-temporal coherent probabilistic forecasts* and *cross-temporal probabilistic forecast reconciliation*. Let $(\mathbb{R}^{n_b m}, \mathcal{F}_{\mathbb{R}^{n_b m}}, \nu)$ be a probability space for the bottom time series $\mathbf{b}_\tau^{[1]}$, where $\mathcal{F}_{\mathbb{R}^{n_b m}}$ is the Borel σ -algebra on $\mathbb{R}^{n_b m}$. Then a σ -algebra \mathcal{F}_s can be constructed from the collection of sets $s(\mathcal{B})$ for all $\mathcal{B} \in \mathcal{F}_{\mathbb{R}^{n_b m}}$.

Definition 3.1 (Cross-Temporal Coherent Probabilistic Forecasts). Given the probability space $(\mathbb{R}^{n_b m}, \mathcal{F}_{\mathbb{R}^{n_b m}}, \nu)$, we define the coherent probability space as the triple $(s, \mathcal{F}_s, \tilde{\nu})$

satisfying the following property: $\tilde{\nu}(s(\mathcal{B})) = \nu(\mathcal{B}), \forall \mathcal{B} \in \mathcal{F}_{\mathbb{R}^{n_b m}}$.

Let $(\mathbb{R}^{n(m+k^*)}, \mathcal{F}_{\mathbb{R}^{n(m+k^*)}}, \hat{\nu})$ be a probability space referring to the incoherent probabilistic forecast $(\hat{\mathbf{x}}_n)$ for all the n series in the system at any temporal aggregation order $k \in \mathcal{K}$.

Definition 3.2 (Cross-Temporal Probabilistic Forecast Reconciliation). The reconciled probability measure of $\hat{\nu}$ with respect to ψ is a probability measure $\tilde{\nu}$ on s with σ -algebra \mathcal{F}_s satisfying

$$\tilde{\nu}(\mathcal{A}) = \hat{\nu}(\psi^{-1}(\mathcal{A})), \quad \forall \mathcal{A} \in \mathcal{F}_s, \quad (7)$$

where $\psi^{-1}(\mathcal{A}) = \{x \in \mathbb{R}^{n(m+k^*)} : \psi(x) \in \mathcal{A}\}$ denotes the pre-image of \mathcal{A} .

The map ψ may be obtained as the composition $s \circ g$, as for the cross-temporal point reconciliation (6).

Theorem 3.1 (Cross-Temporal Reconciled Samples). Suppose that $(\hat{\mathbf{x}}_1, \dots, \hat{\mathbf{x}}_L)$ is a sample drawn from a (cross-temporal) incoherent probability measure $\hat{\nu}$. Then $(\tilde{\mathbf{x}}_1, \dots, \tilde{\mathbf{x}}_L)$, where $\tilde{\mathbf{x}}_\ell = \psi(\hat{\mathbf{x}}_\ell)$, and $\ell = 1, \dots, L$, is a sample drawn from the (cross-temporal) reconciled probability measure $\tilde{\nu}$ defined in (7).

Proof. See Theorem 4.5 in Panagiotelis et al. (2023) using Definition 3.2. \square

Theorem 3.1 is the cross-temporal extension of Theorem 4.5 in Panagiotelis et al. (2023), valid only for the cross-sectional case. It means that a sample from the reconciled distribution can be obtained by reconciling each member of a sample from the incoherent distribution. With this result, we can separate the mechanism used to generate the base forecasts samples from the reconciliation phase.

3.1. Parametric framework: Gaussian reconciliation

It is possible to obtain a reconciled probabilistic forecast analytically for some parametric distributions, such as the multivariate normal (Corani et al., 2021, Eckert et al., 2021, Panagiotelis et al., 2023, Wickramasuriya, 2023). In the cross-sectional framework, Panagiotelis et al.

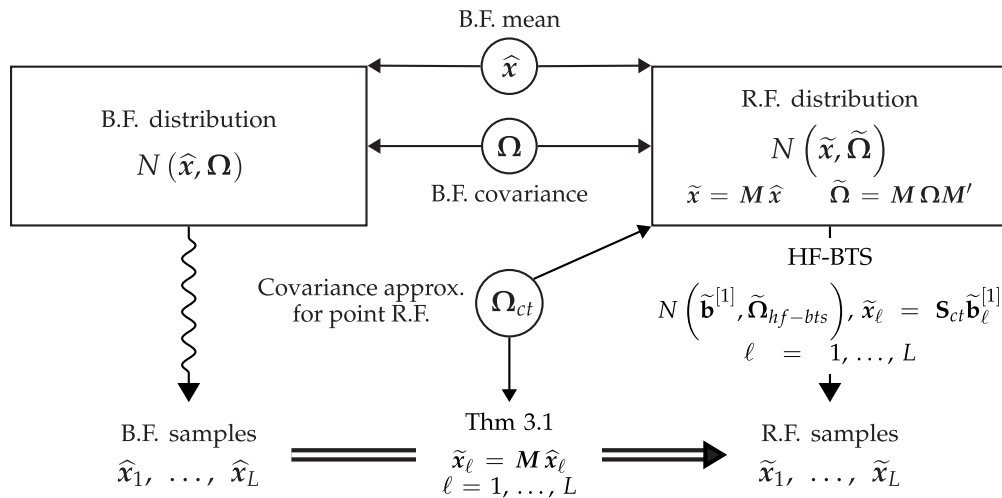


Fig. 5. Overview of cross-temporal forecast reconciliation in the Gaussian framework: two different but equivalent ways of obtaining reconciled forecast samples, as described in Section 3.1. The acronyms R.F. and B.F. stand for Reconciled and Base Forecasts, respectively. HF-BTS stands for High Frequency Bottom Time Series.

(2023) show that, starting from an elliptical distribution for the base forecasts, the reconciled forecast distribution is also elliptical. Using the results shown in Section 2, we extend³ this result to the cross-temporal case. To obtain a reconciled forecast using the multivariate normal distribution, we start with a base forecast distributed as $\mathcal{N}(\hat{\mathbf{x}}, \Omega)$, where $\hat{\mathbf{x}}$ is the mean vector and Ω is the covariance matrix of the base forecasts. Using standard results for the Gaussian case, the reconciled forecast distribution is given by $\mathcal{N}(\tilde{\mathbf{x}}, \tilde{\Omega})$, where

$$\tilde{\mathbf{x}} = \mathbf{M}\hat{\mathbf{x}} \quad \text{and} \quad \tilde{\Omega} = \mathbf{M}\Omega\mathbf{M}', \tag{8}$$

and where \mathbf{M} is the projection matrix defined in (5). Note that if we assume that $\Omega = \Omega_{ct}$ (see the projection matrices in (5) and (6)), then the covariance matrix in (8) simplifies to $\tilde{\Omega} = \mathbf{M}\Omega_{ct}$. In the cross-temporal case, sensibly estimating the covariance matrix Ω can be difficult because we need to consider the temporal and cross-sectional structures simultaneously. This requires estimating many parameters, which can be challenging in practice. Additionally, naively using one-step residuals to estimate the cross-temporal correlation structure can lead to an inappropriate estimate of the covariance matrix.⁴ These challenges are explored in more depth in the following sections.

Focusing on the computational aspect,⁵ we can take several steps to reduce the time required to obtain simulations from the reconciled forecast distribution. For example, when dealing with a genuine hierarchical structure, it is not necessary to simulate from a normal distribution with a defined covariance matrix for the entire

structure. Instead, we can utilize the properties of elliptical distributions to simulate from the high-frequency bottom time series and then obtain the complete simulation through the \mathbf{S}_{ct} matrix. Furthermore, we do not need to calculate the reconciled mean and variance and generate a new sample if we already have a sample from the normal distribution of the base forecasts; we can simply apply the point forecast reconciliation (5), as outlined in Theorem 3.1. Fig. 5 shows two different but equivalent ways of obtaining reconciled forecast samples: the former from the base distribution through Theorem 3.1, and the latter from the reconciled distribution through the high-frequency bottom time series forecasts $\tilde{\mathbf{b}}^{[1]}$ only. The two rectangles represent the base and reconciled forecast distributions, respectively. Enclosed within circles are the distribution parameters involved in the point forecast reconciliation process, transforming $\hat{\mathbf{x}}$ into $\tilde{\mathbf{x}}$, and Ω into $\tilde{\Omega}$. The wave-like arrows represent the simulation processes, generating both base and reconciled forecast samples. Finally, the bold double arrow ‘ \Rightarrow ’ illustrates the generation of the reconciled forecast distributions, as described in Theorem 3.1.

3.2. Non-parametric framework: Bootstrap reconciliation

Analytical expressions for the base and reconciled forecast distributions are sometimes challenging to obtain. Furthermore, parametric assumptions can be restrictive and unrealistic. We propose a procedure called *cross-temporal joint (block) bootstrap (ctjtb)* to generate samples from the base forecast distributions that preserve cross-temporal relationships. This approach involves drawing samples of all series simultaneously from the most temporally aggregated level, and using the most temporally aggregated level to determine the corresponding time indices for the other levels.

Let $\hat{\mathbf{E}}^{[k]}$ be the $(n \times N_k)$ matrix of the residuals for $k \in \mathcal{K}$. Fig. 6 (on the left) provides a visualization of these

³ We assume $H = 1$ and simplify the notation by removing the h suffix without loss of generality.

⁴ In particular, some temporal covariances are fixed to zero (see Online Appendix C for more details).

⁵ We use two R packages to sample from the base forecast Gaussian distribution: MASS (Venables & Ripley, 2002) and Rfast (Papadakis et al., 2022) in Sections 5 and 6, respectively.

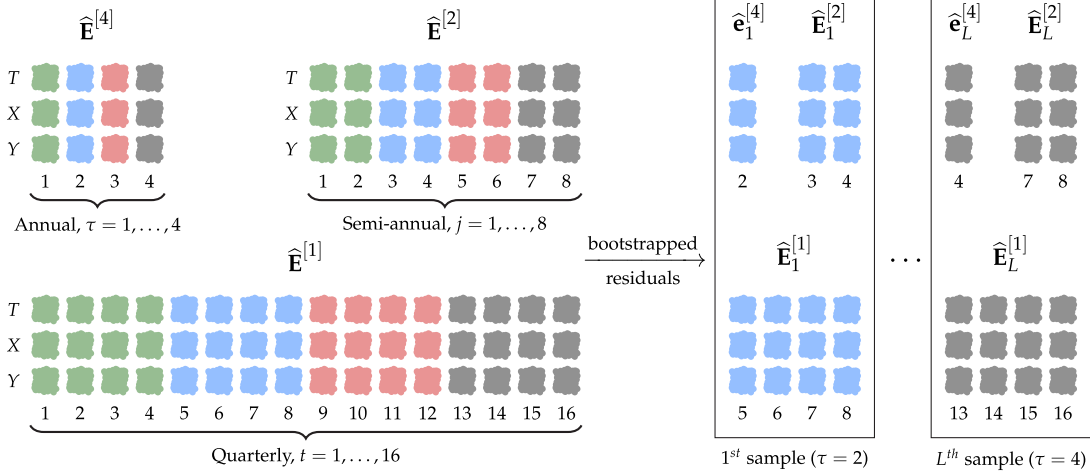


Fig. 6. Example of bootstrapped residuals for 3 linearly constrained quarterly time series (see Fig. 1). On the left there are the residual matrices with 4 years of data ($N = 4$): the green, blue, red and black colors correspond, respectively, to years 1, 2, 3 and 4. On the right the bootstrapped residuals are represented.

matrices and how they are related to each other for the example in Fig. 1. It is assumed that the residuals cover four years ($N = 4$): the green color corresponds to the first year, blue to the second year, and so on. Further, let \mathcal{M}_i be the model used to calculate the base forecasts and residuals for the i th series. Assuming $H = 1$, τ is a random draw with replacement from $1, \dots, N$, and the ℓ th bootstrap incoherent sample is $\hat{\mathbf{x}}_{i,\ell}^{[k]} = f_i(\mathcal{M}_i, \hat{\mathbf{e}}_i^{[k]})$, where $f_i(\cdot)$ depends on the fitted model \mathcal{M}_i . That is, $\hat{\mathbf{x}}_{i,\ell}^{[k]}$ is a sample path simulated for the i th series with error approximated by the corresponding block bootstrapped sample residual $\hat{\mathbf{e}}_i^{[k]}$, the i th row of

$$\hat{\mathbf{E}}_\tau^{[k]} = \begin{bmatrix} \hat{\mathbf{e}}_{1, M_k(\tau-1)+1}^{[k]} & \cdots & \hat{\mathbf{e}}_{1, M_k\tau}^{[k]} \\ \vdots & \ddots & \vdots \\ \hat{\mathbf{e}}_{n, M_k(\tau-1)+1}^{[k]} & \cdots & \hat{\mathbf{e}}_{n, M_k\tau}^{[k]} \end{bmatrix} \quad k \in \mathcal{K}.$$

Fig. 6 (on the right) shows $\hat{\mathbf{E}}_\tau^{[k]}$ for the quarterly cross-temporal hierarchy in Fig. 1.

One of the main advantages of the cross-temporal joint bootstrap is that it allows us to accurately account for the dependence between the different levels of temporal aggregation and not only the cross-sectional dependencies. By sampling residuals from the most temporally aggregated level and using it to determine the indices for the other levels, we can ensure that the bootstrap sample reflects the underlying data distribution. Additionally, the cross-temporal joint bootstrap is easy to implement for many forecasting models, making it a practical and efficient tool. Furthermore, this approach is easily scalable in order to utilize multiple computing power simultaneously for each individual series. This can be especially useful when dealing with large datasets or when trying to speed up the analysis process.

4. Cross-temporal covariance matrix estimation

As the covariance matrix Ω is unknown in practice, a natural estimate is the empirical sample covariance

matrix of the base forecasts $\hat{\Omega}$. In this section, our focus is exclusively on the cross-temporal framework. This means that we have to estimate $r = n(k^* + m)[n(k^* + m) - 1]/2$ different parameters. A possible solution to estimating many parameters when we have fewer observations than r is to construct a shrinkage estimator (Efron, 1975, Efron & Morris, 1975, 1977) using a convex combination of $\hat{\Omega}$ and a diagonal target matrix $\hat{\Omega}_D = \hat{\Omega} \odot \mathbf{I}_{n(k^*+m)}$, such that $\hat{\Omega}_G = \lambda\hat{\Omega}_D + (1 - \lambda)\hat{\Omega}$, where \odot is the Hadamard product, $\lambda \in [0, 1]$ is the shrinkage intensity parameter that can be estimated using the unbiased estimator proposed by Ledoit and Wolf (2004) (see Schäfer & Strimmer, 2005). The linear combination involving these two matrices is referred to as global shrinkage (G), where all off-diagonal elements are shrunk towards zero. $\hat{\Omega}_G$ corresponds to the matrix used by the reconciliation approach *oct(shr)* (Di Fonzo & Girolimetto, 2023a). However, shrinking all off-diagonal elements to zero when we know that the covariance matrix has a cross-sectional and/or temporal structure results in information loss. Therefore, we propose to estimate a smaller matrix, and to use the cross-sectional and/or temporal structure to obtain a better estimator for the covariance matrix of the entire system. Given that $\mathbf{S}_{ct} = \mathbf{S}_{cs} \otimes \mathbf{S}_{te}$, it is possible to express the actual covariance matrix in terms of three smaller matrices, such that

$$\begin{aligned} \tilde{\Omega} &= \mathbf{S}_{ct} \Omega_{hf-bts} \mathbf{S}'_{ct} \\ &= (\mathbf{I}_n \otimes \mathbf{S}_{te}) \Omega_{hf} (\mathbf{I}_n \otimes \mathbf{S}_{te})' \\ &= (\mathbf{S}_{cs} \otimes \mathbf{I}_{m+k^*}) \Omega_{bts} (\mathbf{S}_{cs} \otimes \mathbf{I}_{m+k^*})', \end{aligned} \tag{9}$$

where Ω_{hf-bts} is the $(n_b m \times n_b m)$ covariance matrix for the bottom time series at the temporal aggregation level $k = 1$ (the highest-frequency bottom time series), Ω_{hf} is the $(nm \times nm)$ covariance matrix related to all the high-frequency time series, and Ω_{bts} is the $[n_b(k^* + m) \times n_b(k^* + m)]$ covariance matrix related to bottom time series at any temporal aggregation. Eq. (9) offers three decompositions of the covariance matrix $\tilde{\Omega}$, each characterized by well-defined structures: \mathbf{S}_{ct} capturing cross-temporal, $\mathbf{I}_n \otimes \mathbf{S}_{te}$

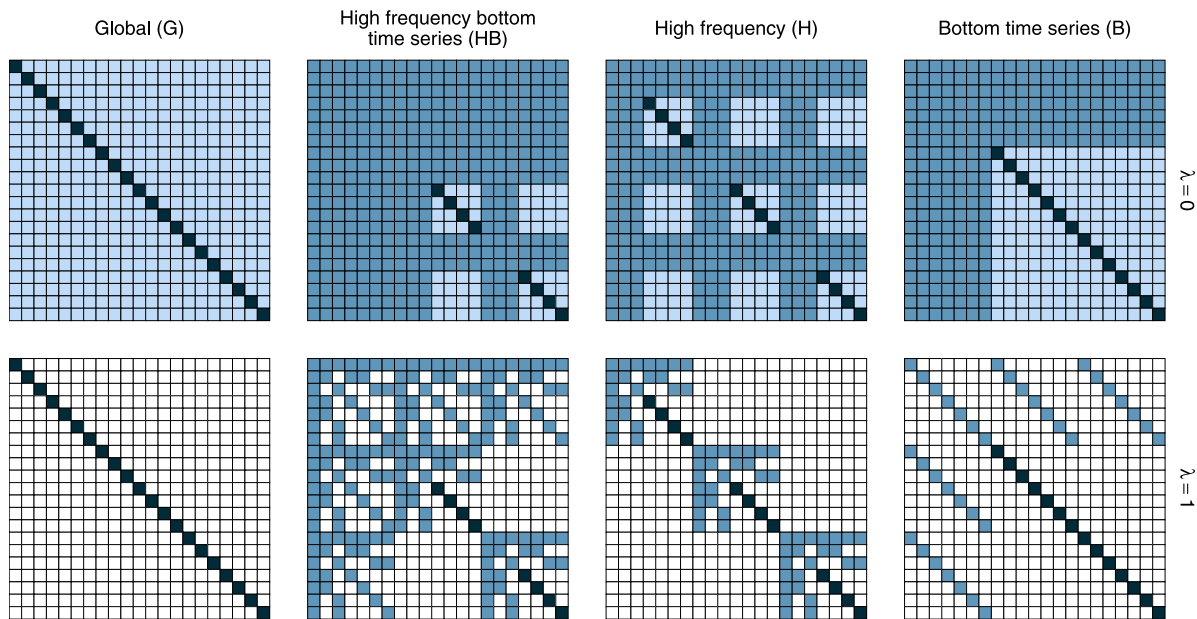


Fig. 7. Representation of four types of covariance matrices that can be obtained from the cross-temporal hierarchical structure (example based on the quarterly series of Fig. 1) for two different values of $\lambda \in \{0, 1\}$, the shrinkage parameter. The entries in black are not modified by shrinkage, the entries in light blue are those actively involved in the shrinkage phase, while the entries in darker blue are derived directly from the cross-sectional and/or temporal structure and hence are not estimated. Additionally, for $\lambda = 1$, the white entries correspond to a zero value.

temporal, and $\mathbf{S}_{cs} \otimes \mathbf{I}_{m+k^*}$ cross-sectional relationships. At the same time, each involves smaller covariance matrices as Ω_{hf-bts} , Ω_{hf} , and Ω_{bts} . Starting from these representations, we propose three different approaches (*HB*, *H*, and *B*, respectively) to approximate $\widehat{\Omega}$.

Therefore, we can apply the idea of Stein-type shrinkage (Efron & Morris, 1977) to Ω_{hf-bts} , Ω_{hf} , and Ω_{bts} by using the corresponding empirical base forecasts residuals estimation. We obtain the following expressions (see Online Appendix B for details):

- *High-frequency bottom time series shrinkage matrix (HB):*

$$\widehat{\Omega}_{HB} = \lambda \mathbf{S}_{ct} \widehat{\Omega}_{hf-bts, D} \mathbf{S}'_{ct} + (1 - \lambda) \mathbf{S}_{ct} \widehat{\Omega}_{hf-bts} \mathbf{S}'_{ct};$$

- *High-frequency shrinkage matrix (H):*

$$\widehat{\Omega}_H = \lambda (\mathbf{I}_n \otimes \mathbf{S}_{te}) \widehat{\Omega}_{hf, D} (\mathbf{I}_n \otimes \mathbf{S}_{te})' + (1 - \lambda) (\mathbf{I}_n \otimes \mathbf{S}_{te}) \widehat{\Omega}_{hf} (\mathbf{I}_n \otimes \mathbf{S}_{te})';$$

- *Bottom time series shrinkage matrix (B):*

$$\widehat{\Omega}_B = \lambda (\mathbf{S}_{cs} \otimes \mathbf{I}_{m+k^*}) \widehat{\Omega}_{bts, D} (\mathbf{S}_{cs} \otimes \mathbf{I}_{m+k^*})' + (1 - \lambda) (\mathbf{S}_{cs} \otimes \mathbf{I}_{m+k^*}) \widehat{\Omega}_{bts} (\mathbf{S}_{cs} \otimes \mathbf{I}_{m+k^*})';$$

where $\widehat{\Omega}_{l, D} = \mathbf{I}_{n_b m} \odot \widehat{\Omega}_j$, $l = \{hf - bts, hf, bts\}$, and λ is the shrinkage parameter. These matrices are not full rank, meaning their inverses, needed to compute the projection to the coherent subspace, do not exist. To address this, a ridge regularization of the form $\widehat{\Omega} + \omega \mathbf{I}$ is used (Marquardt, 1970), where ω is chosen to make the matrix invertible without introducing excessive bias. Fig. 7 gives some visual insights on the covariance matrices obtainable with $\lambda = 0$ and $\lambda = 1$, respectively, for a simple cross-temporal hierarchical structure with three time series and $\mathcal{K} = \{4, 2, 1\}$ (see Fig. 1).

Another important aspect is the number of parameters to be estimated through the residuals of the base forecasts. In Table 1 we report the number of different parameters for the two forecasting experiment: Australian GDP (see Section 5), and Australian Tourism Demand (see Section 6). In addition, we calculate the percentage reductions in the number of parameters compared to the global approach. As we can see, *G* involves a considerably large number of parameters compared to other estimators. *HB* leads to the largest decrease of around 85%, whereas approaches *H* and *B* lie somewhere between *G* and *HB*. In general, as m and n increase, using *H* requires the estimation of fewer parameters than *B*.

It is worth noting that when using the *HB* covariance matrix, we make the assumption that the base error covariance matrix is coherent. This assumption is valid provided that the base forecasts also approximately fulfill constraints (3), which is expected for any reasonable set of forecasts. In addition, with this covariance matrix, the computational complexity of the reconciliation phase is reduced. Specifically, Theorem 4.1 extends Theorem 1 in Hyndman et al. (2011), showing that reconciling using a coherent covariance matrix simplifies to the *ols* approach.

Theorem 4.1. *Let $\widehat{\Omega}_{hf-bts}$ be a $[(n_b m) \times (n_b m)]$ p.d. matrix. Then, using $\Omega_{ct} = \mathbf{S}_{ct} \widehat{\Omega}_{hf-bts} \mathbf{S}'_{ct}$ in the reconciliation formulae (5) and (6) is equivalent to using $\Omega_{ct} = \mathbf{I}_{n(m+k^*)}$ (*ols* approach).*

Proof. See Online Appendix B. □

In the forecasting experiments that follow (and in the simulation in Online Appendix C), we closely analyze

Table 1

Number of different parameters that need to be estimated for the Australian GDP (see Section 5) and the Australian Tourism Demand (see Section 6) forecasting experiments. The percentage reductions in the number of parameters compared to the global approach G are reported in parentheses.

Method	# of different parameters	GDP	Tourism
G	$r = \frac{n(k^* + m)[n(k^* + m) - 1]}{2}$	221 445	108 052 350
B	$r_{HB} < \frac{n_b(k^* + m)[n_b(k^* + m) - 1]}{2} < r$	94 395 (57%)	36 231 328 (66%)
H	$r_{HB} < \frac{nm[nm - 1]}{2} < r$	72 390 (67%)	19 848 150 (82%)
HB	$r_{HB} = \frac{n_b m [n_b m - 1]}{2} < r$	30 876 (86%)	6 655 776 (94%)

these different constructions with a dual purpose. In particular, we use the full covariance matrix ($\lambda = 0$) of the base forecasts to obtain base forecast samples of the linearly constrained time series under Gaussianity. We also use the shrinkage versions as approximations of the covariance matrix to be used for reconciliation (excluding HB; see Theorem 4.1). This allows us to better understand the properties and abilities of each parameterization.

4.1. Multi-step residuals

Model residuals may be used to estimate the covariance matrix in cross-temporal forecast reconciliation. In time series analysis, it is common to use residuals corresponding to one-step-ahead forecasts. However, due to the temporal dimension in our setting, residuals corresponding to different forecast horizons are required. Thus, we define *multi-step residuals* as $e_{i,h,j}^{[k]} = x_{i,j+h}^{[k]} - \hat{x}_{i,j+h|j}^{[k]}$, where $i = 1, \dots, n$, $j = 1, \dots, N_k$, and $\hat{x}_{i,j+h|j}^{[k]}$ is the h -step fitted value, calculated as the h -step-ahead forecast using data up to time j . In general, these residuals are autocorrelated, except when $h = 1$.

Following Di Fonzo and Girolimetto (2023a), we use a matrix organization of the residuals similar to the one for the base forecasts in Section 2.1. Specifically, let N be the total number of observations for the most temporally aggregate time series. Then, the N_k -vectors of multi-step residuals for the temporal aggregation k and the series i , $\mathbf{e}_{i,h}^{[k]} = [e_{i,h,1}^{[k]} \ e_{i,h,2}^{[k]} \ \dots \ e_{i,h,N_k}^{[k]}]'$ with $h = 1, \dots, M_k$, can be organized in matrix form as

$$\mathbf{E}_i^{[k]} = \begin{bmatrix} e_{i,1,1}^{[k]} & e_{i,2,2}^{[k]} & \dots & e_{i,M_k,M_k}^{[k]} \\ \vdots & \vdots & & \vdots \\ e_{i,1,N_k-M_k+1}^{[k]} & e_{i,2,N_k-M_k+2}^{[k]} & \dots & e_{i,M_k,N_k}^{[k]} \end{bmatrix}.$$

Let $\mathbf{E}_i = [\mathbf{E}_i^{[m]} \ \dots \ \mathbf{E}_i^{[1]}]$. Then the $[N \times n(m+k^*)]$ cross-temporal residual matrix is given by $\mathbf{E} = [\mathbf{E}_1 \ \dots \ \mathbf{E}_n]$.

To better understand the properties of the proposed alternatives, a simulation study was performed (the results are shown in Online Appendix C). We studied the effect of combining cross-sectional and temporal aggregations using a simple hierarchy, where the small size and nature of the data generating process make it possible to calculate the true cross-temporal covariance structure exactly, thus providing insights into the nature of the time series data involved in the forecast reconciliation process.

We find that simulating base forecasts from multi-step residuals allows for a more accurate estimation of the covariance matrix and that reconciliation further improves the forecast accuracy.

4.2. Overlapping residuals

Another issue that arises in the case of cross-temporal reconciliation is the low number of available residuals, especially for the higher orders of temporal aggregation. A possible solution is to use residuals calculated using overlapping series by allowing the year to have a varying starting time. To better explain how to calculate overlapping residuals, assume we have a single series $\mathbf{y} = [y_1 \ y_2 \ y_3 \ \dots \ y_{T-1} \ y_T]'$. We can construct k non-overlapping series such that

$$\mathbf{x}^{[k],s} = \left\{ x_j^{[k],s} \right\}_{j=1}^{N_k-s}, \quad \text{where } x_j^{[k],s} = \sum_{t=(j-1)k+s+1}^{jk-s} y_t,$$

with $s = 0, \dots, (k - 1)$. For example, suppose we have a semi-annual series with $k = 2$ and $T = 6$. Then we can construct two annual time series depending on which time is deemed the start of the year:

$$\begin{aligned} \mathbf{x}^{[2],0} &= \left[x_1^{[2],0} \ x_2^{[2],0} \ x_3^{[2],0} \right]' \\ &= \left[y_1 + y_2 \ y_3 + y_4 \ y_5 + y_6 \right]', \end{aligned}$$

and

$$\mathbf{x}^{[2],1} = \left[x_1^{[2],1} \ x_2^{[2],1} \right]' = \left[y_2 + y_3 \ y_4 + y_5 \right]'$$

To calculate overlapping residuals, we propose the following steps:

1. Fit a model to $\mathbf{x}^{[k],0}$ (i.e., select an appropriate model and estimate the model parameters using the available data) and calculate the residuals.
2. Apply the same model in Step 1 to $\mathbf{x}^{[k],s}$ for $s = 1, \dots, k - 1$ without re-estimating the parameters, and calculate the residuals.

The resulting residuals can be used to estimate the covariance matrix in cross-temporal forecast reconciliation. This increases the number of available residuals, particularly when working with higher-frequency observations such as monthly or daily data. It is important to note that this approach assumes that the model used in Step 1 is appropriate for all the different series $\mathbf{x}^{[k],s}$. Some seasonal models will not be appropriate, as the seasonal pattern

Table 2

Cross-temporal reconciliation approaches for the Australian GDP (see Section 5) and the Australian Tourism Demand (see Section 6) forecasting experiments. All the reconciliation procedures are available in FoReco (Girolimetto and Di Fonzo, 2023a).

Label	Description
ct(shr_{cs}, bu_{te})	Partly bottom-up (Section 2.2) starting from cross-sectional reconciled forecasts using the <i>shr</i> approach.
oct(·)	Optimal cross-temporal reconciliation for the <i>struc</i> , <i>wlsv</i> and <i>bdshr</i> approaches. One-step residuals were used with <i>wlsv</i> and <i>bdshr</i> .
oct _n (·)	Optimal cross-temporal reconciliation with multi-step residuals (see Section 4.1) for the approaches presented in Section 4: <i>hshr</i> for high-frequency shrinkage, and <i>bshr</i> for bottom time series shrinkage.
oct _o (·)	Optimal cross-temporal reconciliation with overlapping residuals (see Section 4.2) for the <i>wlsv</i> and <i>bdshr</i> approaches.
oct _{oh} (<i>hshr</i>)	Optimal cross-temporal reconciliation with overlapping and multi-step residuals (see Section 4.1 and 4.2) for the <i>hshr</i> (high-frequency shrinkage) approach presented in Section 4.

will be shifted for different values of *s*. However, this will not affect seasonal ARIMA models, as the seasonality is defined in terms of lags which are unaffected by the value of *s*.

5. Forecasting Australian GDP

The Australian Quarterly National Accounts (QNA) dataset has been widely studied in the literature on forecast reconciliation (Athanasopoulos et al., 2020, Di Fonzo & Girolimetto, 2023a). Building on these results, we now consider cross-temporally reconciled probabilistic forecasts.

We use univariate ARIMA models⁶ to obtain quarterly base forecasts for the $n = 95$ QNA time series, spanning the period 1984:Q4–2018:Q1, defining GDP from both the Income and Expenditure sides. We perform a rolling forecast experiment with an expanding window: the first training sample spans the period 1984:Q4 to 1994:Q3, and the last ends in 2017:Q1, for a total of 91 forecast origins. For the temporal aggregation dimension, we aggregate the quarterly data to both semi-annual and annual. We obtain four-, two-, and one-step-ahead base forecasts from the quarterly, semi-annual, and annual frequencies, respectively; i.e., $\mathcal{K} = \{4, 2, 1\}$.

The base forecast samples in the Gaussian case are obtained using the sample covariance matrices with the global (*G*) and high-frequency (*H*) parameterizations (Section 4), since it is not possible to identify a unique representation for the other cases.⁷ We compare the results obtained using multi-step residuals with and without overlapping, in order to measure the benefit of obtaining overlapping residuals. In the non-parametric case, we use the cross-temporal joint bootstrap (ctjb) presented in Section 3.2. Finally, to reconcile the resulting (1000) base forecasts samples, we applied the following techniques⁸ (see Table 2): ct(shr_{cs}, bu_{te}), ct(wls_{cs}, bu_{te}), oct_o(*wlsv*), oct_o(*bdshr*), and oct_{oh}(*hshr*).

⁶ We use the `auto.arima` function from the R package `forecast` (Hyndman et al., 2023).

⁷ When simultaneously considering hierarchies from the Income and Expenditure sides, the result is a general linearly constrained time series, where bottom and upper time series are not uniquely defined, such that the cross-sectional bottom-up reconciliation approach is unfeasible (Girolimetto & Di Fonzo, 2023b).

⁸ The results with shrunk covariance matrices are available in Online Appendix D.2, where we also report the results obtained using other reconciliation approaches.

The accuracy of the probabilistic forecasts is evaluated using the continuous ranked probability score (CRPS; see Matheson & Winkler, 1976, Gneiting & Katzfuss, 2014), which is an index that considers the single series and provides us with a marginal evaluation of the approaches. In addition, we employ the energy score (ES; see Gneiting & Katzfuss, 2014), which is the CRPS extension to the multivariate case, to evaluate the forecasting accuracy for the whole system (Panagiotelis et al., 2023, Wickramasuriya, 2023). In particular, we consider the geometric mean of the relative CRPS (Fleming & Wallace, 1986) and the relative ES:

$$\overline{\text{RelCRPS}}_{j,s}^{[k]} = \left(\prod_{i=1}^n \frac{\text{CRPS}_{i,j,s}^{[k]}}{\text{CRPS}_{i,0,0}^{[k]}} \right)^{\frac{1}{n}} \quad \text{and} \quad \text{RelES}_{j,s}^{[k]} = \frac{\text{ES}_{j,s}^{[k]}}{\text{ES}_{0,0}^{[k]}}$$

where *j* denotes the reconciliation approach, and *s* indicates the approach used to simulate the base forecasts. As a reference approach ($s = 0$ and $j = 0$), we consider the base forecasts produced by the bootstrap approach. If we consider all the temporal aggregation orders (i.e. $\forall k \in \mathcal{K}$), the overall accuracy indices are given by

$$\overline{\text{RelCRPS}}_{j,s} = \left(\prod_{\substack{i=1, \dots, n \\ k \in \mathcal{K}}} \frac{\text{CRPS}_{i,j,s}^{[k]}}{\text{CRPS}_{i,0,0}^{[k]}} \right)^{\frac{1}{n(k^*+m)}} \quad \text{and} \quad \overline{\text{RelES}}_{j,s} = \left(\prod_{k \in \mathcal{K}} \frac{\text{ES}_{j,s}^{[k]}}{\text{ES}_{0,0}^{[k]}} \right)^{\frac{1}{(k^*+m)}}$$

5.1. Results

Forecasting accuracy indices based on the CRPS and ES are presented in Table 3. As a benchmark approach, we use the base forecasts calculated using the bootstrap method. For base forecasts, we find that using a parametric approach with the normal distribution performs better than the non-parametric bootstrap approach. This is likely due to the limited number of residuals available for bootstrapping, which does not allow for sufficient exploration of the data. Directly specifying diagonal covariance matrices seems to be more effective than shrinking

Table 3

RelCRPS and ES ratio indices defined in (10) and (11) for the Australian QNA dataset. Approaches performing worse than the benchmark (bootstrap base forecasts, ctjb) are highlighted in red, the best for each column is marked in bold, and the overall lowest value is highlighted in blue. The reconciliation approaches are described in Table 2.

Reconciliation approach	Base forecasts' sample approach									
	ctjb	Gaussian approach ^a				ctjb	Gaussian approach ^a			
		G _h	H _h	G _{oh}	H _{oh}		G _h	H _h	G _{oh}	H _{oh}
<i>RelCRPS</i>										
$\forall k \in \{4, 2, 1\}$										
base	1.000	0.979	0.995	0.968	0.976	1.000	0.988	0.971	0.971	
ct(<i>shr_{cs}</i> , <i>bu_{te}</i>)	0.937	0.956	0.956	0.976	0.976	0.992	1.008	1.008	1.029	
ct(<i>wls_{cs}</i> , <i>bu_{te}</i>)	0.930	0.917	0.917	0.898	0.898	0.986	0.974	0.975	0.956	
oct _o (<i>wls_v</i>)	0.926	0.911	0.912	0.896	0.895	0.984	0.971	0.970	0.954	
oct _o (<i>bdshr</i>)	0.978	0.964	0.946	0.952	0.930	1.034	1.016	1.003	0.989	
oct _{oh} (<i>hshr</i>)	1.006	0.983	1.009	0.974	1.001	1.068	1.046	1.059	1.034	
$k = 2$										
base	1.000	0.984	0.993	0.968	0.976	1.000	0.966	1.004	0.981	
ct(<i>shr_{cs}</i> , <i>bu_{te}</i>)	0.949	0.966	0.966	0.987	0.987	0.874	0.896	0.896	0.914	
ct(<i>wls_{cs}</i> , <i>bu_{te}</i>)	0.942	0.928	0.928	0.909	0.909	0.866	0.853	0.853	0.834	
oct _o (<i>wls_v</i>)	0.938	0.921	0.923	0.907	0.906	0.860	0.847	0.848	0.832	
oct _o (<i>bdshr</i>)	0.991	0.974	0.957	0.964	0.942	0.914	0.905	0.883	0.865	
oct _{oh} (<i>hshr</i>)	1.021	0.996	1.021	0.987	1.016	0.934	0.912	0.951	0.904	
$k = 4$										
$\forall k \in \{4, 2, 1\}$										
base	1.000	0.970	0.988	0.960	0.970	1.000	0.977	0.965	0.965	
ct(<i>shr_{cs}</i> , <i>bu_{te}</i>)	0.897	0.944	0.944	0.973	0.973	0.964	1.001	1.001	1.033	
ct(<i>wls_{cs}</i> , <i>bu_{te}</i>)	0.886	0.880	0.880	0.860	0.860	0.954	0.944	0.945	0.928	
oct _o (<i>wls_v</i>)	0.891	0.879	0.881	0.864	0.864	0.958	0.945	0.945	0.931	
oct _o (<i>bdshr</i>)	0.940	0.928	0.910	0.918	0.895	1.004	0.986	0.971	0.980	
oct _{oh} (<i>hshr</i>)	0.986	0.968	0.999	0.959	0.992	1.053	1.034	1.049	1.024	
$k = 2$										
base	1.000	0.972	0.985	0.959	0.969	1.000	0.959	1.000	0.976	
ct(<i>shr_{cs}</i> , <i>bu_{te}</i>)	0.915	0.961	0.960	0.991	0.991	0.818	0.874	0.874	0.899	
ct(<i>wls_{cs}</i> , <i>bu_{te}</i>)	0.904	0.896	0.896	0.877	0.877	0.807	0.805	0.805	0.782	
oct _o (<i>wls_v</i>)	0.908	0.895	0.898	0.881	0.882	0.812	0.802	0.806	0.786	
oct _o (<i>bdshr</i>)	0.960	0.947	0.929	0.938	0.915	0.860	0.856	0.836	0.841	
oct _{oh} (<i>hshr</i>)	1.007	0.988	1.017	0.979	1.014	0.904	0.888	0.934	0.881	

^a The Gaussian method employs a sample covariance matrix: G_h and H_h use multi-step residuals and G_{oh} and H_{oh} use overlapping and multi-step residuals.

to a target covariance matrix. Among all the procedures, ct(*wls_{cs}*, *bu_{te}*) and oct_o(*wls_v*) show the greatest relative gains. In contrast, oct_{oh}(*hshr*) does not show much improvement. Furthermore, the greatest improvements are observed for higher temporal aggregation levels.

We utilize the non-parametric Friedman test and the post hoc multiple-comparison-with-the-best (MCB) Nemenyi test (Koning et al., 2005, Kourentzes & Athanasopoulos, 2019, Makridakis et al., 2022, Kourentzes, 2022) to determine whether the forecasting performances of the different techniques are significantly different from one another. Fig. 8 presents the MCB using the CRPS. The probabilistic forecasts from ct(*wls_{cs}*, *bu_{te}*) and oct_o(*wls_v*) are significantly better than the base forecasts at any level of aggregation. Unlike the application on the Australian Tourism Demand dataset (see Section 6), in this case, one of the partly bottom-up approaches is not significantly worse than the optimal approach.

Overall, we find that using overlapping residuals almost always leads to a greater improvement in terms of both the ES and CRPS. Forecasts at the most aggregated level (year) seem to benefit the most from reconciliation, and using one-step overlapping residuals appears to be sufficient to improve forecasts if the generation of the base forecasts sample paths takes into account the multi-step structure.

6. Forecasting Australian Tourism Demand

The Australian Tourism Demand dataset (Wickramasuriya et al., 2019) measures the number of nights Australians spent away from home. It includes 228 monthly observations of Visitor Nights (VNs) from January 1998 to December 2016, and has a cross-sectional grouped structure based on a geographic hierarchy crossed by purpose of travel. The geographic hierarchy comprises seven states, 27 zones, and 76 regions, for a total of 111 nested geographic divisions. Six of these zones are each formed by a single region, resulting in 105 unique nodes in the hierarchy. The purpose of travel comprises four categories: holiday, visiting friends and relatives, business, and other. To avoid redundancies (Di Fonzo & Girolimetto, 2022b), 24 nodes are not considered, resulting in an unbalanced hierarchy of 525 unique nodes instead of the theoretical 555 with duplicated nodes. The dataset includes the 304 bottom series, which are aggregated into 221 upper time series. Table 4 omits duplicated entries and updates the information in Table 7 from Wickramasuriya et al. (2019). The data can be temporally aggregated into two, three, four, six, or 12 months ($\mathcal{K} = \{12, 4, 3, 2, 1\}$).

We perform a rolling forecast experiment with an expanding window. The process begins by using the first




Fig. 8. MCB Nemenyi test for the Australian QNA dataset using CRPS at different temporal aggregation levels for the Gaussian (using overlapping and multi-step residuals, H) and non-parametric bootstrap approaches. In each panel, the Friedman test p -value is reported in the lower-right corner. The mean rank of each approach is shown to the right of its name. Statistically significant differences in performance are indicated if the intervals of two forecast reconciliation procedures do not overlap. Thus, approaches that do not overlap with the blue interval are considered significantly worse than the best, and vice versa.

Table 4
Grouped time series for the Australian Tourism Demand dataset.

	Number of series		
	GD	PT	Tot.
Australia	1	4	5
States	7	28	35
Zones ^a	21	84	105
Regions	76	304	380
Total	105	420	525

^a Six Zones with only one region are included in 'Regions'. 'GD': geographic division; 'PT': purpose of travel.

10 years, from January 1998 to December 2008, to generate forecasts for the entire following year (2009). Then, the training set is increased by one month. This process is repeated until the last training set is used (January 1998 to December 2015) with a total of 85 different test sets. For the temporal aggregation dimension, we aggregate the monthly data up to annual data. We obtain twelve-,

six-, four-, three-, two-, and one-step-ahead base forecasts from the monthly data and the aggregation over 2, 3, 4, 6, and 12 months. ETS models selected by minimizing the AICc criterion (Hyndman et al., 2023) are fitted to the log-transformed data, with the resulting base forecasts being back-transformed to produce non-negative forecasts (Wickramasuriya et al., 2020).

The (1000) base forecast samples are obtained using the Gaussian approach with sample⁹ covariance matrices (Section 4) using multi-step residuals¹⁰ and the bootstrap approach (Section 3.2). For reconciliation, six different approaches are adopted (see Table 2): $ct(shr_{cs}, bu_{te})$, $oct(struc)$, $oct(wlsv)$, $oct(bdshr)$, $oct_h(bshr)$, and $oct_h(hshr)$.

Negative forecasts may be produced during the reconciliation phase (Wickramasuriya et al., 2020, Di Fonzo

⁹ The results with shrunk covariance matrices are available in Online Appendix E.2, where we also report the results obtained using other reconciliation approaches.

¹⁰ We do not include overlapping, as we are unable to correctly determine the residuals for the overlapping series using ETS models (see Section 4.2).

Table 5

RelCRPS and ES ratio indices defined in (10) and (11) for the Australian Tourism Demand dataset. Approaches performing worse than the benchmark (bootstrap base forecasts, ctjb) are highlighted in red, the best for each column is marked in bold, and the overall lowest value is highlighted in blue. The reconciliation approaches are described in Table 2.

Reconciliation approach	Base forecasts' sample approach									
	ctjb	Gaussian approach ^a				ctjb	Gaussian approach ^a			
		G	B	H	HB		G	B	H	HB
<i>RelCRPS</i>										
		$\forall k \in \{12, 6, 4, 3, 2, 1\}$					$k = 1$			
base	1.000	0.971	0.971	0.973	0.973	1.000	0.972	0.972	0.972	0.972
ct(<i>shr_{cs}</i> , <i>bu_{te}</i>)	1.057	0.974	0.969	0.974	0.969	0.976	0.963	0.962	0.963	0.962
oct(<i>struc</i>)	0.982	0.962	0.961	0.961	0.959	0.970	0.963	0.963	0.963	0.963
oct(<i>wlsv</i>)	0.987	0.959	0.959	0.958	0.957	0.952	0.957	0.957	0.957	0.957
oct(<i>bdshr</i>)	0.975	0.956	0.953	0.952	0.951	0.949	0.955	0.953	0.954	0.954
oct _n (<i>bshr</i>)	0.994	1.018	1.020	1.016	1.019	0.988	1.007	1.013	1.006	1.012
oct _n (<i>hshr</i>)	0.969	0.993	0.993	0.990	0.991	0.953	0.977	0.977	0.979	0.979
		$k = 3$					$k = 12$			
base	1.000	0.971	0.971	0.972	0.973	1.000	0.968	0.967	0.969	0.969
ct(<i>shr_{cs}</i> , <i>bu_{te}</i>)	1.041	0.977	0.974	0.977	0.974	1.163	0.977	0.965	0.977	0.965
oct(<i>struc</i>)	0.986	0.967	0.966	0.966	0.965	0.982	0.951	0.949	0.947	0.943
oct(<i>wlsv</i>)	0.983	0.963	0.962	0.962	0.962	1.025	0.954	0.953	0.949	0.947
oct(<i>bdshr</i>)	0.972	0.960	0.958	0.957	0.957	1.002	0.950	0.944	0.939	0.935
oct _n (<i>bshr</i>)	0.999	1.021	1.022	1.018	1.022	0.987	1.024	1.021	1.021	1.019
oct _n (<i>hshr</i>)	0.971	0.994	0.994	0.992	0.993	0.978	1.003	1.005	0.996	0.997
<i>ES ratio indices</i>										
		$\forall k \in \{12, 6, 4, 3, 2, 1\}$					$k = 1$			
base	1.000	0.956	0.955	0.958	0.951	1.000	0.952	0.950	0.952	0.950
ct(<i>shr_{cs}</i> , <i>bu_{te}</i>)	1.243	0.886	0.879	0.886	0.879	1.098	0.929	0.928	0.930	0.927
oct(<i>struc</i>)	1.085	0.917	0.915	0.916	0.912	1.027	0.943	0.942	0.943	0.942
oct(<i>wlsv</i>)	1.132	0.933	0.929	0.931	0.927	1.050	0.951	0.949	0.950	0.949
oct(<i>bdshr</i>)	1.047	0.904	0.897	0.897	0.891	1.009	0.936	0.933	0.934	0.931
oct _n (<i>bshr</i>)	0.931	0.867	0.866	0.863	0.860	0.965	0.927	0.927	0.925	0.923
oct _n (<i>hshr</i>)	1.081	0.935	0.931	0.935	0.927	1.028	0.952	0.951	0.952	0.950
		$k = 3$					$k = 12$			
base	1.000	0.961	0.958	0.960	0.955	1.000	0.942	0.947	0.951	0.937
ct(<i>shr_{cs}</i> , <i>bu_{te}</i>)	1.245	0.911	0.904	0.911	0.904	1.326	0.779	0.767	0.777	0.766
oct(<i>struc</i>)	1.096	0.939	0.936	0.938	0.933	1.077	0.826	0.822	0.823	0.818
oct(<i>wlsv</i>)	1.142	0.953	0.949	0.951	0.946	1.149	0.851	0.845	0.847	0.840
oct(<i>bdshr</i>)	1.060	0.926	0.920	0.921	0.915	1.021	0.808	0.796	0.796	0.787
oct _n (<i>bshr</i>)	0.954	0.895	0.895	0.892	0.887	0.833	0.741	0.741	0.737	0.735
oct _n (<i>hshr</i>)	1.093	0.955	0.951	0.956	0.949	1.066	0.851	0.846	0.848	0.838

^a The Gaussian method employs a sample covariance matrix and includes four techniques (G, B, H, HB) with multi-step residuals.

& Girolimetto, 2022b, 2023b), thus generating unreasonable values (e.g., a negative forecast for tourism demand makes no sense). To overcome this limitation, we applied the simple heuristic proposed by Di Fonzo and Girolimetto (2022a, 2023b). Following Theorem 3.1, we are thus able to obtain reconciled samples respecting non-negativity constraints starting from an incoherent sample simulated from a Gaussian distribution. Finally, to evaluate the performance, we employ the continuous ranked probability score (CRPS), the energy score (ES), and the multiple-comparison-with-the-best (MCB) Nemenyi test, introduced in Sections 5 and 5.1.

6.1. Results

The CRPS and ES indices are shown, respectively, in Table 5 for monthly, quarterly and annual forecasts. These tables are divided by different temporal levels and each column uses a different approach to calculate the base forecasts, referred to as 'base'. The bootstrap method is used as a benchmark to calculate the accuracy indices.

Base forecasts using a Gaussian approach are better in terms of both the CRPS and ES compared to the bootstrap approach (the benchmark). Assumptions of truncated Gaussianity (Gaussian with negative values set to zero) may seem strict, but given the limited number of residuals, it can lead to improved forecasts in terms of CRPS and ES. Bootstrap forecasts suffer from the limited number of available residuals, leading in general to lower forecast accuracy. The Gaussian approach overcomes this limitation and provides better results. Regarding the different covariance matrix estimates for Gaussian base forecasts, there are no big differences. For this reason, using only the high-frequency bottom time series can be useful to estimate fewer parameters and reduce the initial high dimensionality.

In the Gaussian case, partly bottom-up techniques like ct(*shr_{cs}*, *bu_{te}*) lead to better results than the benchmark (bootstrap base forecasts). However, it is not always guaranteed that the improvement is higher than the starting base forecasts (by comparing the value of each column).

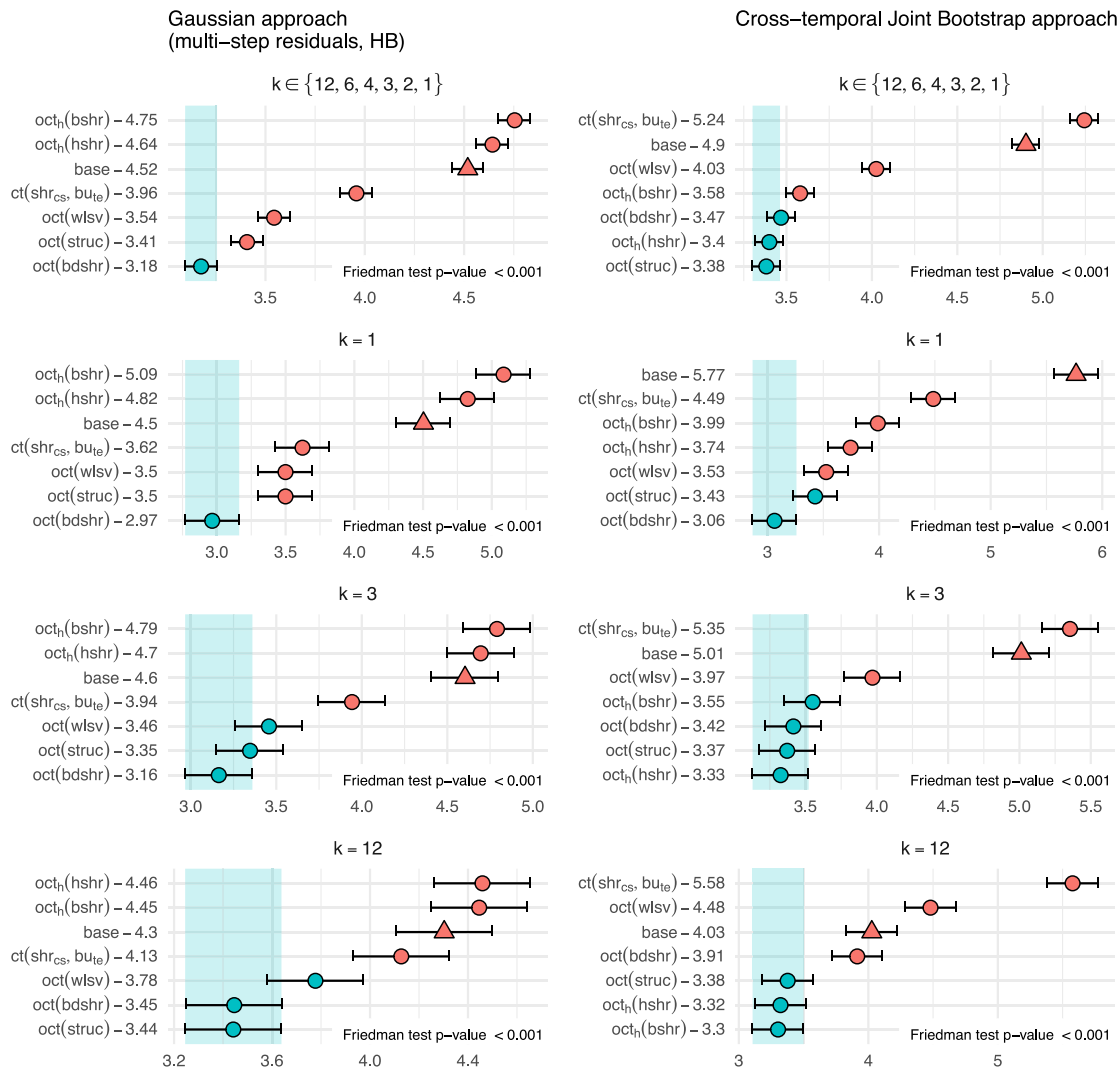


Fig. 9. MCB Nemenyi test for the Australian Tourism Demand dataset using the CRPS at different temporal aggregation levels for the Gaussian (multi-step residuals, HB) and the non-parametric bootstrap approaches. In each panel, the Friedman test p-value is reported in the lower-right corner. The mean rank of each approach is shown to the right of its name. Statistically significant differences in performance are indicated if the intervals of two forecast reconciliation procedures do not overlap. Thus, approaches that do not overlap with the blue interval are considered significantly worse than the best, and vice versa.

This is particularly true for higher levels of temporal aggregation. Overall, oct(bdshr) is almost always the best in terms of the CRPS. The shrinkage approach oct_h(hshr) performs well in the bootstrap case: it is competitive with oct(bdshr) at a lower temporal frequency ($k \in \{2, 1\}$) and it is able to improve for $k \geq 3$. In terms of the ES, oct(bdshr) is still competitive, although it does not always show the best relative performance, like oct_h(bshr). It is also worth noting that oct(struc), which does not make use of residuals, proves to be competitive by consistently improving on the base forecasts in terms of both the CRPS and ES.

Fig. 9 shows the MCB using the CRPS for the Gaussian approach using multi-step residuals (HB) and the non-parametric bootstrap approach. In general, the partly bottom-up procedure improves with respect to the base forecasts at the monthly level, but optimal cross-temporal

procedures are always better. In the bootstrap framework, we can identify a group of three procedures—oct(bdshr), oct(hshr), and oct(struc)—that are almost always in the group of best approaches (denoted by the blue dot). In the Gaussian framework, oct(wlsv), oct(struc), and oct(bdshr) are always significantly better than the base forecasts and equivalent in terms of the results for temporal aggregation orders greater than two. For monthly series, oct(bdshr) is always significantly better than all other approaches.

7. Conclusion

In this paper, we extended the probabilistic reconciliation setting developed by Panagiotelis et al. (2023) for the cross-sectional case to the cross-temporal framework. Through appropriate notation, we showed how

theorems and definitions valid for the cross-sectional case can be reinterpreted and extended. The general notation proposed can help investigate extensions following different probabilistic approaches, such as those in Jeon et al. (2019), Ben Taieb et al. (2021), and Corani et al. (2023). We proposed a Gaussian and a bootstrap approach to simulate the base forecasts that takes into account both cross-sectional and temporal relationships simultaneously, opening the way for further research into cross-temporal probabilistic forecasting.

Moreover, we analyzed the use of residuals, showing that one-step residuals fail to capture the temporal structure, and proposed multi-step residuals that can fully capture the cross-temporal relationships. Due to the high-dimensionality of the cross-temporal setting when dealing with covariance matrices, we proposed four alternative forms to reduce the number of parameters to be estimated, showing that the overlapping residuals may reduce the high-dimensionality burden by increasing the number of available residuals. These ideas warrant further investigation in future works.

Finally, we performed empirical applications on two datasets commonly used in forecast reconciliation research: Australian GDP from Income and Expenditure sides and Australian Tourism Demand. We found that in both cases, optimal cross-temporal reconciliation approaches significantly improved on base forecasts. We also compared these with partly bottom-up techniques that use uni-dimensional reconciliation (either cross-sectional or temporal) and confirmed that simultaneously exploiting both dimensions in reconciliation produces better results, especially at higher levels of temporal aggregation. This was more evident in the Australian Tourism Demand application, where the involved temporal hierarchies are richer, allowing the regression-based forecast reconciliation approaches to capture and exploit more features of the data through the available temporal aggregation levels (Kourentzes et al., 2014, 2017, Kourentzes & Petropoulos, 2016) compared to the partly bottom-up approach. In these two datasets, $oct(wlsv)$ and $oct(bdshr)$ appeared as the two best performing approaches in terms of improving both forecast accuracy and computational efficiency (see the online appendix), thus corroborating the results of Di Fonzo and Girolimetto (2023a) for point forecast reconciliation.

In conclusion, cross-temporal forecast reconciliation is an important tool to improve the accuracy of forecasts while simultaneously ensuring their coherency in both space and time. Furthermore, these techniques can be customized to suit the specific needs of an organization, allowing for the incorporation of relevant domain-specific knowledge (e.g., non-negative constraints) and expertise, ensuring that the resulting forecasts are not only accurate but also coherent and more reliable for decision-making purposes.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

The authors gratefully acknowledge the comments and suggestions provided by two referees and the associate editor. The authors thank Casey Lichtendahl and all the participants to the 2023 IIF Workshop on Forecast Reconciliation in Prato (Italy). Tommaso Di Fonzo and Daniele Girolimetto acknowledge financial support from project PRIN2017 'HiDEA: Advanced Econometrics for High-frequency Data', 2017RSMPPZ. Rob Hyndman acknowledges the support of the Australian government through the Australian Research Council Industrial Transformation Training Centre in Optimisation Technologies, Integrated Methodologies, and Applications (OPTIMA), Project ID IC200100009.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.ijforecast.2023.10.003>.

References

- Athanasopoulos, G., Ahmed, R. A., & Hyndman, R. J. (2009). Hierarchical forecasts for Australian domestic tourism. *International Journal of Forecasting*, 25(1), 146–166. <http://dx.doi.org/10.1016/j.ijforecast.2008.07.004>.
- Athanasopoulos, G., Gamakumara, P., Panagiotelis, A., Hyndman, R. J., & Affan, M. (2020). Hierarchical forecasting. In P. Fuleky (Ed.), *Macroeconomic forecasting in the era of big data*, vol. 52 (pp. 689–719). Cham: Springer International Publishing, http://dx.doi.org/10.1007/978-3-030-31150-6_21.
- Athanasopoulos, G., Hyndman, R. J., Kourentzes, N., & Panagiotelis, A. (2023). Forecast reconciliation: A review. *International Journal of Forecasting* (in press), <https://www.monash.edu/business/ebs/research/publications/ebs/2023/wp08-2023.pdf>.
- Athanasopoulos, G., Hyndman, R. J., Kourentzes, N., & Petropoulos, F. (2017). Forecasting with temporal hierarchies. *European Journal of Operational Research*, 262(1), 60–74. <http://dx.doi.org/10.1016/j.ejor.2017.02.046>.
- Ben Taieb, S., Taylor, J. W., & Hyndman, R. J. (2017). Coherent probabilistic forecasts for hierarchical time series. In *Proceedings of the 34th international conference on machine learning* (pp. 3348–3357). PMLR.
- Ben Taieb, S., Taylor, J. W., & Hyndman, R. J. (2021). Hierarchical probabilistic forecasting of electricity demand with smart meter data. *Journal of the American Statistical Association*, 116(533), 27–43. <http://dx.doi.org/10.1080/01621459.2020.1736081>.
- Corani, G., Azzimonti, D., Augusto, J. P. S. C., & Zaffalon, M. (2021). Probabilistic reconciliation of hierarchical forecast via Bayes' rule. *Machine Learning and Knowledge Discovery in Databases*, 12459, 211–226. http://dx.doi.org/10.1007/978-3-030-67664-3_13.
- Corani, G., Azzimonti, D., & Rubattu, N. (2023). Probabilistic reconciliation of count time series. *International Journal of Forecasting* (in press). <http://dx.doi.org/10.1016/j.ijforecast.2023.04.003>.
- Dangerfield, B. J., & Morris, J. S. (1992). Top-down or bottom-up: Aggregate versus disaggregate extrapolations. *International Journal of Forecasting*, 8(2), 233–241. [http://dx.doi.org/10.1016/0169-2070\(92\)90121-0](http://dx.doi.org/10.1016/0169-2070(92)90121-0).
- Di Fonzo, T., & Girolimetto, D. (2022a). Enhancements in cross-temporal forecast reconciliation, with an application to solar irradiance forecasts. [arXiv:2209.07146](https://arxiv.org/abs/2209.07146).
- Di Fonzo, T., & Girolimetto, D. (2022b). Forecast combination-based forecast reconciliation: Insights and extensions. *International Journal of Forecasting* (in press). <http://dx.doi.org/10.1016/j.ijforecast.2022.07.001>.

- Di Fonzo, T., & Girolimetto, D. (2023a). Cross-temporal forecast reconciliation: Optimal combination method and heuristic alternatives. *International Journal of Forecasting*, 39(1), 39–57. <http://dx.doi.org/10.1016/j.ijforecast.2021.08.004>.
- Di Fonzo, T., & Girolimetto, D. (2023b). Spatio-temporal reconciliation of solar forecasts. *Solar Energy*, 251, 13–29. <http://dx.doi.org/10.1016/j.solener.2023.01.003>.
- Dunn, D. M., Williams, W. H., & Dechaine, T. L. (1976). Aggregate versus subaggregate models in local area forecasting. *Journal of the American Statistical Association*, 71(353), 68–71. <http://dx.doi.org/10.1080/01621459.1976.10481478>.
- Eckert, F., Hyndman, R. J., & Panagiotelis, A. (2021). Forecasting Swiss exports using Bayesian forecast reconciliation. *European Journal of Operational Research*, 291(2), 693–710. <http://dx.doi.org/10.1016/j.ejor.2020.09.046>.
- Efron, B. (1975). Biased versus unbiased estimation. *Advances in Mathematics*, 16(3), 259–277. [http://dx.doi.org/10.1016/0001-8708\(75\)90114-0](http://dx.doi.org/10.1016/0001-8708(75)90114-0).
- Efron, B., & Morris, C. (1975). Data analysis using Stein's estimator and its generalizations. *Journal of the American Statistical Association*, 70(350), 311–319. <http://dx.doi.org/10.1080/01621459.1975.10479864>.
- Efron, B., & Morris, C. (1977). Stein's paradox in statistics. *Scientific American*, 236(5), 119–127.
- Fleming, P. J., & Wallace, J. J. (1986). How not to lie with statistics: The correct way to summarize benchmark results. *Communications of the ACM*, 29(3), 218–221. <http://dx.doi.org/10.1145/5666.5673>.
- Fliedner, G. (2001). Hierarchical forecasting: Issues and use guidelines. *Industrial Management & Data Systems*, 101(1), 5–12. <http://dx.doi.org/10.1108/02635570110365952>.
- Girolimetto, D., & Di Fonzo, T. (2023a). FoReco: forecast reconciliation. R package version 0.2.6. <https://CRAN.R-project.org/package=FoReco>.
- Girolimetto, D., & Di Fonzo, T. (2023b). Point and probabilistic forecast reconciliation for general linearly constrained multiple time series. [arXiv:2305.05330](https://arxiv.org/abs/2305.05330).
- Gneiting, T., & Katzfuss, M. (2014). Probabilistic forecasting. *Annual Review of Statistics and Its Application*, 1(1), 125–151. <http://dx.doi.org/10.1146/annurev-statistics-062713-085831>.
- Gross, C. W., & Sohl, J. E. (1990). Disaggregation methods to expedite product line forecasting. *Journal of Forecasting*, 9(3), 233–254. <http://dx.doi.org/10.1002/for.3980090304>.
- Hyndman, R. J. (2022). Notation for forecast reconciliation. <https://robjhyndman.com/hyndsight/reconciliation-notation.html>. (Accessed 20 July 2023).
- Hyndman, R. J., Ahmed, R. A., Athanasopoulos, G., & Shang, H. L. (2011). Optimal combination forecasts for hierarchical time series. *Computational Statistics & Data Analysis*, 55(9), 2579–2589. <http://dx.doi.org/10.1016/j.csda.2011.03.006>.
- Hyndman, R. J., Athanasopoulos, G., Bergmeir, C., Caceres, G., Chhay, L., Kuroptev, K., O'Hara-Wild, M., Petropoulos, F., Razbash, S., Wang, E., Yasmeeen, F., Garza, F., Girolimetto, D., Ihaka, R., R Core Team, Reid, D., Shaub, D., Tang, Y., Wang, X., & Zhou, Z. (2023). forecast: Forecasting functions for time series and linear models. R package v8.20. <https://CRAN.R-project.org/package=forecast>.
- Jeon, J., Panagiotelis, A., & Petropoulos, F. (2019). Probabilistic forecast reconciliation with applications to wind power and electric load. *European Journal of Operational Research*, 279(2), 364–379. <http://dx.doi.org/10.1016/j.ejor.2019.05.020>.
- Koning, A. J., Franses, P. H., Hibon, M., & Stekler, H. (2005). The M3 competition: Statistical tests of the results. *International Journal of Forecasting*, 21(3), 397–409. <http://dx.doi.org/10.1016/j.ijforecast.2004.10.003>.
- Kourentzes, N. (2022). tsutils: Time series exploration, modelling and forecasting. R package v0.9.3. <https://CRAN.R-project.org/package=tsutils>.
- Kourentzes, N., & Athanasopoulos, G. (2019). Cross-temporal coherent forecasts for Australian tourism. *Annals of Tourism Research*, 75, 393–409. <http://dx.doi.org/10.1016/j.annals.2019.02.001>.
- Kourentzes, N., & Athanasopoulos, G. (2021). Elucidate structure in intermittent demand series. *European Journal of Operational Research*, 288(1), 141–152. <http://dx.doi.org/10.1016/j.ejor.2020.05.046>.
- Kourentzes, N., & Petropoulos, F. (2016). Forecasting with multivariate temporal aggregation: The case of promotional modelling. *International Journal of Production Economics*, 181, 145–153. <http://dx.doi.org/10.1016/j.ijpe.2015.09.011>.
- Kourentzes, N., Petropoulos, F., & Trapero, F. (2014). Improving forecasting by estimating time series structural components across multiple frequencies. *International Journal of Forecasting*, 30(2), 291–302. <http://dx.doi.org/10.1016/j.ijforecast.2013.09.006>.
- Kourentzes, N., Rostami-Tabar, B., & Barrow, D. K. (2017). Demand forecasting by temporal aggregation: Using optimal or multiple aggregation levels? *Journal of Business Research*, 78, 1–9. <http://dx.doi.org/10.1016/j.jbusres.2017.04.016>.
- Ledoit, O., & Wolf, M. (2004). A well-conditioned estimator for large-dimensional covariance matrices. *Journal of Multivariate Analysis*, 88(2), 365–411. [http://dx.doi.org/10.1016/S0047-259X\(03\)00096-4](http://dx.doi.org/10.1016/S0047-259X(03)00096-4).
- Magnus, J. R., & Neudecker, H. (2019). *Matrix differential calculus with applications in statistics and econometrics* (3rd ed.). Hoboken, NJ: Wiley.
- Makridakis, S., Spiliotis, E., & Assimakopoulos, V. (2022). M5 accuracy competition: Results, findings, and conclusions. *International Journal of Forecasting*, 38(4), 1346–1364. <http://dx.doi.org/10.1016/j.ijforecast.2021.11.013>.
- Marquardt, D. W. (1970). Generalized inverses, ridge regression, biased linear estimation, and nonlinear estimation. *Technometrics*, 12(3), 591–612.
- Matheson, J. E., & Winkler, R. L. (1976). Scoring rules for continuous probability distributions. *Management Science*, 22(10), 1087–1096. <http://dx.doi.org/10.1287/mnsc.22.10.1087>.
- Nystrup, P., Lindström, E., Pinson, P., & Madsen, H. (2020). Temporal hierarchies with autocorrelation for load forecasting. *European Journal of Operational Research*, 280(3), 876–888. <http://dx.doi.org/10.1016/j.ejor.2019.07.061>.
- Panagiotelis, A., Athanasopoulos, G., Gamakumara, P., & Hyndman, R. J. (2021). Forecast reconciliation: A geometric view with new insights on bias correction. *International Journal of Forecasting*, 37(1), 343–359. <http://dx.doi.org/10.1016/j.ijforecast.2020.06.004>.
- Panagiotelis, A., Gamakumara, P., Athanasopoulos, G., & Hyndman, R. J. (2023). Probabilistic forecast reconciliation: Properties, evaluation and score optimisation. *European Journal of Operational Research*, 306(2), 693–706. <http://dx.doi.org/10.1016/j.ejor.2022.07.040>.
- Panamtash, H., & Zhou, Q. (2018). Coherent probabilistic solar power forecasting. In *2018 IEEE international conference on probabilistic methods applied to power systems* (pp. 1–6).
- Papadakis, M., Tsagris, M., Dimitriadis, M., Fafalios, S., Tsamardinos, I., Fasiolo, M., Borboudakis, G., Burkardt, J., Zou, C., Lakiotaki, K., & Chatzipantsiou, C. (2022). Rfast: A collection of efficient and extremely fast R functions. R package version 2.0.6. <https://CRAN.R-project.org/package=Rfast>.
- Pennings, C. L., & van Dalen, J. (2017). Integrated hierarchical forecasting. *European Journal of Operational Research*, 263(2), 412–418. <http://dx.doi.org/10.1016/j.ejor.2017.04.047>.
- Punia, S., Singh, S. P., & Madaan, J. K. (2020). A cross-temporal hierarchical framework and deep learning for supply chain forecasting. *Computers & Industrial Engineering*, 149, Article 106796. <http://dx.doi.org/10.1016/j.cie.2020.106796>.
- R Core Team (2022). *R: A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing.
- Sanguri, K., Shankar, S., Punia, S., & Patra, S. (2022). Hierarchical container throughput forecasting: The value of coherent forecasts in the management of ports operations. *Computers & Industrial Engineering*, 173, Article 108651. <http://dx.doi.org/10.1016/j.cie.2022.108651>.
- Schäfer, J., & Strimmer, K. (2005). A Shrinkage Approach to Large-Scale Covariance Matrix Estimation and Implications for Functional Genomics. *Statistical Applications in Genetics and Molecular Biology*, 4(1), <http://dx.doi.org/10.2202/1544-6115.1175>.
- Spiliotis, E., Petropoulos, F., Kourentzes, N., & Assimakopoulos, V. (2020). Cross-temporal aggregation: Improving the forecast accuracy of hierarchical electricity consumption. *Applied Energy*, 261, Article 114339. <http://dx.doi.org/10.1016/j.apenergy.2019.114339>.

- Venables, W. N., & Ripley, B. D. (2002). *Modern applied statistics with S* (4th ed.). New York: Springer.
- Wickramasuriya, S. L. (2023). Probabilistic forecast reconciliation under the Gaussian framework. *Journal of Business & Economic Statistics* (in press). <http://dx.doi.org/10.1080/07350015.2023.2181176>.
- Wickramasuriya, S. L., Athanasopoulos, G., & Hyndman, R. J. (2019). Optimal forecast reconciliation for hierarchical and grouped time series through trace minimization. *Journal of the American Statistical Association*, 114(526), 804–819. <http://dx.doi.org/10.1080/01621459.2018.1448825>.
- Wickramasuriya, S. L., Turlach, B. A., & Hyndman, R. J. (2020). Optimal non-negative forecast reconciliation. *Statistics and Computing*, 30(5), 1167–1182. <http://dx.doi.org/10.1007/s11222-020-09930-0>.
- Yagli, G. M., Yang, D., & Srinivasan, D. (2019). Reconciling solar forecasts: Sequential reconciliation. *Solar Energy*, 179, 391–397. <http://dx.doi.org/10.1016/j.solener.2018.12.075>.
- Yagli, G. M., Yang, D., & Srinivasan, D. (2020). Reconciling solar forecasts: Probabilistic forecasting with homoscedastic Gaussian errors on a geographical hierarchy. *Solar Energy*, 210, 59–67. <http://dx.doi.org/10.1016/j.solener.2020.06.005>.
- Yang, D. (2020). Reconciling solar forecasts: Probabilistic forecast reconciliation in a nonparametric framework. *Solar Energy*, 210, 49–58. <http://dx.doi.org/10.1016/j.solener.2020.03.095>.
- Zambon, L., Azzimonti, D., & Corani, G. (2022). Efficient probabilistic reconciliation of forecasts for real-valued and count time series. [arXiv:2210.02286](https://arxiv.org/abs/2210.02286).