

DEPARTMENT OF ECONOMETRICS
AND BUSINESS STATISTICS

ISSN 1440-771X

WORKING PAPER SERIES

Cross-temporal Probabilistic Forecast Reconciliation

Daniele Girolimetto, University of Padova, George Athanasopoulos, Monash University, Tommaso Di Fonzo, University of Padova, Rob J Hyndman, Monash University

Cross-temporal Probabilistic Forecast Reconciliation

Daniele Girolimetto

Department of Statistical Sciences, University of Padova

Email: daniele.girolimetto@phd.unipd.it

George Athanasopoulos

Department of Econometrics and Business Statistics, Monash University

Email: george.athanasopoulos@monash.edu

Tommaso Di Fonzo

Department of Statistical Sciences, University of Padova

Email: tommaso.difonzo@unipd.it

Rob J Hyndman

Department of Econometrics and Business Statistics, Monash University

Email: rob.hyndman@monash.edu

30 March 2023

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 evaluate the proposed methods through a detailed simulation study that investigates their theoretical and empirical properties. We further assess the effectiveness of the proposed cross-temporal reconciliation approach by applying it to two empirical forecasting experiments, using the Australian GDP and the Australian Tourism Demand datasets. For both applications, we show that 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, our study expands and unifies the notation for cross-sectional, temporal and cross-temporal reconciliation, thus extending and deepening the probabilistic cross-temporal framework. The results highlight the potential of the proposed cross-temporal forecast reconciliation methods in improving the accuracy of probabilistic forecasting models.

Keywords: Coherent, GDP, Linear constraints, Multivariate time series, Temporal aggregation, Tourism flows

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, Di Fonzo & Girolimetto 2022*d*). There are many fields where forecast reconciliation is useful, such as when forecasting GDP and its components, electricity demand and power generation, demand in supply chains with product categories, tourist flows across geographic regions and travel purpose, and more. Moreover, effective decision-making depends on the support of accurate and coherent forecasts.

Classical reconciliation methods 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. The bottom-up approach (Dunn et al. 1976) starts by generating forecasts at the most disaggregate level and summing these to arrive at the desired forecasts for aggregate levels. On the other hand, the top-down approach (Gross & Sohl 1990) forecasts the most aggregated level and then disaggregates it to lower levels (Fliedner 2001, Athanasopoulos et al. 2009). The middle-out method (Athanasopoulos et al. 2009) combines both approaches by selecting an intermediate level and applies top-down for lower levels and bottom-up for upper levels.

All of these approaches ignore useful information available at other levels (Pennings & van Dalen 2017). Consequently, in the last decade, hierarchical forecasting 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 & Girolimetto (2023*a*) 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 framework, there have been some developments towards probabilistic forecasting including Ben Taieb et al. (2017), Panamtaash & Zhou (2018), Jeon et al. (2019), Ben Taieb et al. (2021), Corani et al. (2021), Corani et al. (2022), 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 proposed by Panagiotelis et al. (2021). They show how a reconciled forecast can be constructed from an arbitrary base forecast when the density of the base forecast 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 & 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 effective and practical 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 methodological contributions described in this paper are implemented in the FoReco package for R (Girolimetto & Di Fonzo 2023).

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. A simulation study is performed in Section 5, to better understand the properties of the methodology. Two empirical applications using the Australian GDP and the Australian Tourism Demand datasets are considered in Sections 6 and 7, respectively¹. Finally, Section 8 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)

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

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

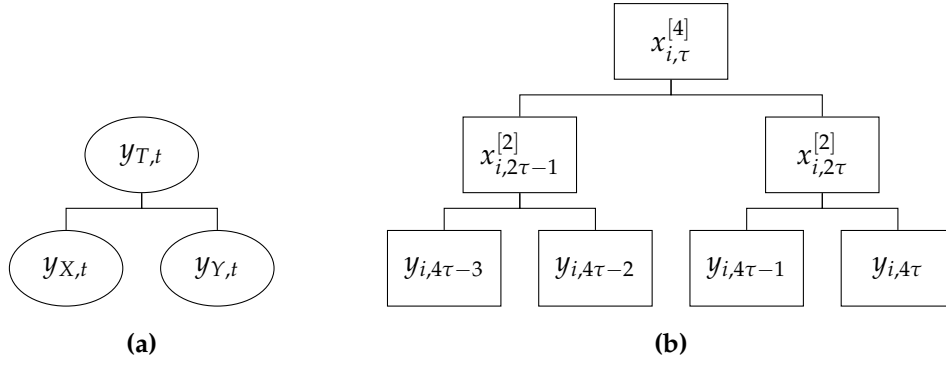


Figure 1: (a) A simple two-level cross-sectional hierarchy for 3 linearly constrained time series with $n = 3$, $n_a = 1$ and $n_b = 2$. (b) A temporal hierarchy for a quarterly series ($m = 4$ and $\mathcal{K} = \{4, 2, 1\}$).

where 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.

An example is a hierarchical time series where series at upper levels can be expressed by appropriately summing part or all of the series at the bottom level. Figure 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 levels 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 A_{cs} such that $\mathbf{u}_t = A_{cs}\mathbf{b}_t$. Following Di Fonzo & Girolimetto (2022d), we can always construct a zero-constraints cross-sectional matrix from the aggregation matrix, $C_{cs} = [I_{n_a} \quad -A_{cs}]$. Finally, the cross-sectional structural matrix is given by $S_{cs} = [A' \quad I_{n_b}]'$, providing the structural representation (Hyndman et al. 2011) $\mathbf{y}_t = S_{cs}\mathbf{b}_t$. Considering the hierarchical example in Figure 1(a), we have $A_{cs} = [1 \quad 1]$, $C_{cs} = [1 \quad -1 \quad -1]$ and $S_{cs} = \begin{bmatrix} 1 & 1 \\ & I_2 \end{bmatrix}$. In general there is no reason for \mathbf{u}_t to be restricted to simple sums of \mathbf{b}_t ; therefore $A_{cs} \in \mathbb{R}^{n_a \times n_b}$ may contain any real values, and not only 0s and 1s.

Considering now the temporal framework, we denote as $\mathcal{K} = \{k_p, k_{p-1}, \dots, k_2, k_1\}$ the set of p factors of m , in descending order, where $k_1 = 1$ and $k_p = m$ (Athanasopoulos et al. 2017). 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 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 $k \in \mathcal{K}$, we stack the observations in the column vector $\mathbf{x}_{i,\tau}^{[k]} = [x_{i,M_k(\tau-1)+1}^{[k]} \quad x_{i,M_k(\tau-1)+2}^{[k]} \quad \dots \quad x_{i,M_k\tau}^{[k]}]'$, and obtain the vector for all the temporal aggregation orders $\mathbf{x}_{i,\tau} = [x_{i,\tau}^{[k_p]} \quad \mathbf{x}_{i,\tau}^{[k_{p-1}]'} \quad \dots \quad \mathbf{x}_{i,\tau}^{[1]'}]'$, $\tau = 1, \dots, N$. The structural representation of the temporal hierarchy (Athanasopoulos et al. 2017) is then $\mathbf{x}_{i,\tau} = S_{te}\mathbf{x}_{i,\tau}^{[1]}$, where $S_{te} = [A'_{te} \quad I_m]'$ is the $[(m + k^*) \times m]$ temporal structural

matrix,

$$\mathbf{A}_{te} = \begin{bmatrix} & \mathbf{1}'_{k_p} & \\ \mathbf{I}_{\frac{m}{k_{p-1}}} & \otimes & \mathbf{1}'_{k_{p-1}} \\ & \vdots & \\ \mathbf{I}_{\frac{m}{k_2}} & \otimes & \mathbf{1}'_{k_2} \end{bmatrix}$$

is the $(k^* \times m)$ temporal aggregation matrix with $k^* = \sum_{k \in \mathcal{K} \setminus \{k_1\}} M_k$, 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 $\mathbf{C}_{te} = [\mathbf{I}_{k^*} \quad -\mathbf{A}_{te}]$ is the $[k^* \times (m+k^*)]$ zero constraints temporal matrix. Figure 1(b) shows the hierarchical representation of a quarterly time series, for which $m = 4$, $\mathcal{K} = \{4, 2, 1\}$, and

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

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 $k \in \mathcal{K}$ and $j = 1, \dots, N_k$, where $\mathbf{x}_j^{[k]} = [\mathbf{u}_j^{[k]'} \quad \mathbf{b}_j^{[k]'}]'$ with $\mathbf{u}_j^{[k]} = [x_{1,j}^{[k]} \quad \dots \quad x_{n_a,j}^{[k]}]'$ is the n_a -vector of upper time series and $\mathbf{b}_j^{[k]} = [x_{(n_a+1),j}^{[k]} \quad \dots \quad x_{n,j}^{[k]}]'$ 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 , whose rows and columns represent, respectively, the cross-sectional and the temporal dimension:

$$\mathbf{X}_\tau = \begin{bmatrix} \mathbf{x}'_{1,\tau} \\ \vdots \\ \mathbf{x}'_{n,\tau} \end{bmatrix} = [\mathbf{X}_\tau^{[k_p]} \quad \dots \quad \mathbf{X}_\tau^{[k_1]}] \quad \text{with} \quad \mathbf{X}_\tau^{[k]} = \begin{bmatrix} \mathbf{U}_\tau^{[k]} \\ \mathbf{B}_\tau^{[k]} \end{bmatrix}, \quad (3)$$

where for any fixed k , $\mathbf{U}_\tau^{[k]}$ is the $(n_a \times N_k)$ matrix grouping the upper time series, $\mathbf{B}_\tau^{[k]}$ is the $(n_b \times N_k)$ matrix grouping the bottom time series. 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 & Girolimetto (2023a) show that the cross-temporal constraints working on the complete set of observations corresponding to time period τ can be expressed in a zero-constrained representation through the full rank $[(n_a m + n k^*) \times n(m+k^*)]$ zero constraints cross-temporal matrix \mathbf{C}_{ct} such that

$$\mathbf{C}_{ct} = \begin{bmatrix} \mathbf{C}_* \\ \mathbf{I}_n \otimes \mathbf{C}_{te} \end{bmatrix} \implies \mathbf{C}_{ct} \mathbf{x}_\tau = \mathbf{0}_{[(n_a m + n k^*) \times 1]} \quad \text{for} \quad \tau = 1, \dots, N, \quad (4)$$

where $\mathbf{x}_\tau = \text{vec}(\mathbf{X}'_\tau) = [\mathbf{x}'_{1,\tau}, \dots, \mathbf{x}'_{n,\tau}]'$, $\mathbf{C}_* = [\mathbf{0}_{(n_b m \times n k^*)} \quad \mathbf{I}_m \otimes \mathbf{C}_{cs}] \mathbf{P}'$, and \mathbf{P} is the commutation matrix (Magnus & Neudecker 2019, p. 54) such that $\mathbf{P} \text{vec}(\mathbf{Y}_\tau) = \text{vec}(\mathbf{Y}'_\tau)$. A structural representation can be considered as well: $\mathbf{x}_\tau = \mathbf{S}_{ct} \mathbf{b}_\tau^{[1]} = s(\mathbf{b}_\tau^{[1]})$, where

$$\mathbf{S}_{ct} = \mathbf{S}_{cs} \otimes \mathbf{S}_{te} \quad (5)$$

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 $\mathbf{b}_\tau^{[1]} = \text{vec}(\mathbf{B}_\tau^{[1]'})$. We observe that, in agreement with Panagiotelis et al. (2021), \mathbf{x}_τ lies in an $(n_b m)$ -dimensional subspace \mathfrak{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} . In Figure 2, we have represented \mathbf{S}_{ct} for 3 linearly constrained quarterly time series, shown in Figure 1.

2.1 Optimal point forecast reconciliation

Let $\hat{\mathbf{x}}_h = \text{vec}(\hat{\mathbf{X}}'_h)$, $h = 1, \dots, H$, be the h -step ahead base forecasts (however obtained) with error covariance matrix given by $\mathbf{W}_h = \text{Var}(\hat{\mathbf{x}}_h - \mathbf{x})$, where H is the forecast horizon for

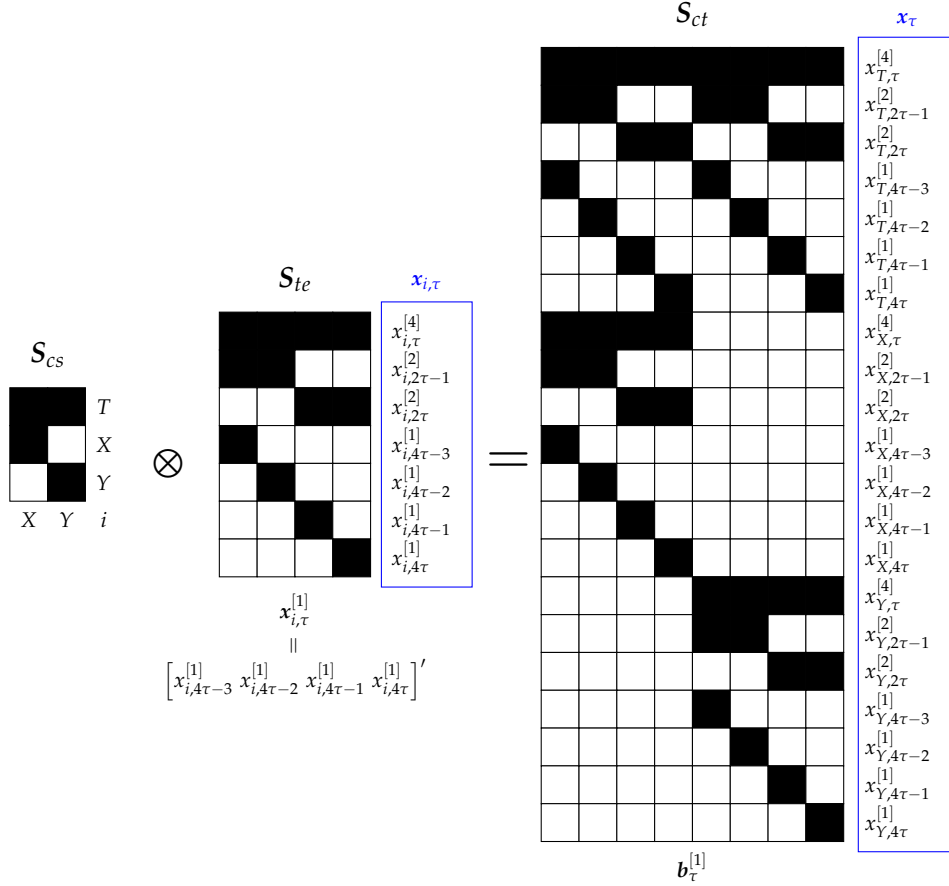


Figure 2: Representation of the cross-temporal summing matrix $\mathbf{S}_{ct} = \mathbf{S}_{cs} \otimes \mathbf{S}_{te}$ defined in equation (5) for a system of 3 linearly constrained quarterly time series (see Figure 1). The white cells indicate a zero, while the black ones a non-zero value.

the most temporally aggregated time series. Denote

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

where $\widehat{\mathbf{U}}_h^{[k]}$ is the $(n_a \times M_k)$ matrix grouping the upper time series and $\widehat{\mathbf{B}}_h^{[k]}$ is the $(n_b \times M_k)$ matrix grouping the bottom time series for a given temporal aggregation order k . The matrix $\widehat{\mathbf{X}}_h$, organized as \mathbf{X}_τ in expression (3), contains incoherent forecasts, such as $\mathbf{C}_{ct}\widehat{\mathbf{x}}_h \neq \mathbf{0}_{[(n_a m + n_b k^*) \times 1]}$ with $h = 1, \dots, H$ and $\widehat{\mathbf{x}}_h = \text{vec}(\widehat{\mathbf{X}}_h')$. 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 aims to adjust the base forecast $\widehat{\mathbf{x}}_h$ by finding a mapping $\psi : \mathbb{R}^{n(m+k^*)} \rightarrow \mathfrak{s}$ such that $\widetilde{\mathbf{x}}_h = \psi(\widehat{\mathbf{x}}_h)$, where $\widetilde{\mathbf{x}}_h \in \mathfrak{s}$ 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)

$$\widetilde{\mathbf{x}}_h = \psi(\widehat{\mathbf{x}}_h) = \mathbf{M}\widehat{\mathbf{x}}_h, \quad (6)$$

where $\mathbf{M} = \mathbf{I}_{n(m+k^*)} - \boldsymbol{\Omega}_{ct}\mathbf{C}'_{ct}(\mathbf{C}_{ct}\boldsymbol{\Omega}_{ct}\mathbf{C}'_{ct})^{-1}\mathbf{C}_{ct}$, for a positive definite matrix $\boldsymbol{\Omega}_{ct}$, and $\widetilde{\mathbf{x}}_h = \text{vec}(\widetilde{\mathbf{X}}_h')$. According to Wickramasuriya et al. (2019) showed that the minimum variance linear unbiased reconciled forecasts, satisfying the unbiased condition $E(\widetilde{\mathbf{x}}_h - \mathbf{x}_h) = 0$, has solution (6) when $\boldsymbol{\Omega}_{ct} = \mathbf{W}_h$.

Alternatively, the cross-temporal reconciled forecasts $\widetilde{\mathbf{X}}_h$ may be found according to the structural approach proposed by Hyndman et al. (2011) for the cross-sectional framework, yielding $\widetilde{\mathbf{x}}_h = \mathbf{S}_{ct}\mathbf{G}\widehat{\mathbf{x}}_h$ for some matrix \mathbf{G} . Wickramasuriya et al. (2019) showed that this leads to a solution equivalent to the cross-temporally reconciled forecasts in (6), given by

$$\widetilde{\mathbf{x}}_h = \psi(\widehat{\mathbf{x}}_h) = (s \circ g)(\widehat{\mathbf{x}}_h) = \mathbf{S}_{ct}\mathbf{G}\widehat{\mathbf{x}}_h, \quad (7)$$

where $\mathbf{G} = (\mathbf{S}'_{ct}\boldsymbol{\Omega}_{ct}^{-1}\mathbf{S}_{ct})^{-1}\mathbf{S}'_{ct}\boldsymbol{\Omega}_{ct}^{-1}$, and $\mathbf{M} = \mathbf{S}_{ct}\mathbf{G}$. 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 Appendix A we report some cross-sectional, temporal and cross-temporal approximations for the covariance matrix to be used in (6) and (7).

2.2 Cross-temporal bottom-up forecast reconciliation

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

Figure 3 provides a visual representation of partly bottom-up in a two-step cross-temporal reconciliation approach. On the left (Figure 3a), 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 (Figure 3b), we first compute temporally reconciled forecasts for the most disaggregated cross-sectional level, and then apply the cross-sectional bottom-up. We denote these two-step reconciliation approaches, respectively, as $\text{ct}(\text{rec}_{te}, \text{bu}_{cs})$, and $\text{ct}(\text{rec}_{cs}, \text{bu}_{te})$, 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 $\text{ct}(\text{bu}_{cs}, \text{bu}_{te}) = \text{ct}(\text{bu}_{te}, \text{bu}_{cs}) = \text{ct}(\text{bu})$.




Figure 3: A visual representation of partly bottom up starting from (3a) cross-sectionally reconciled forecasts for the temporal order 1 ($\tilde{\mathbf{U}}^{[1]}$ and $\tilde{\mathbf{B}}^{[1]}$) followed by temporal bottom-up, and (3b) temporally reconciled forecasts of the cross-sectional bottom time series ($\tilde{\mathbf{B}}^{[k]}$, $k \in \mathcal{K}$) 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.

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. (2022) 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 $(s, \mathcal{F}_s, \check{\nu})$ satisfying the following property: $\check{\nu}(s(\mathcal{B})) = \nu(\mathcal{B}), \forall \mathcal{B} \in \mathcal{F}_{\mathbb{R}^{n_b m}}$.

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{x}_t) 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 \mathfrak{s} with σ -algebra $\mathcal{F}_{\mathfrak{s}}$ satisfying

$$\tilde{\nu}(\mathcal{A}) = \hat{\nu}(\psi^{-1}(\mathcal{A})), \quad \forall \mathcal{A} \in \mathcal{F}_{\mathfrak{s}}, \quad (8)$$

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 (7). With the following result, we can separate the mechanism used to generate the base forecasts samples from the reconciliation phase.

Theorem 3.1 (Cross-temporal reconciled samples). Suppose that $(\hat{x}_1, \dots, \hat{x}_L)$ is a sample drawn from a (cross-temporal) incoherent probability measure $\hat{\nu}$. Then $(\tilde{x}_1, \dots, \tilde{x}_L)$, where $\tilde{x}_\ell = \psi(\hat{x}_\ell)$ and $\ell = 1, \dots, L$, is a sample drawn from the (cross-temporal) reconciled probability measure $\tilde{\nu}$ defined in (8).

Proof. See Theorem 4.5 from 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). It means that a sample from the reconciled distribution can be obtained by reconciling each member of a sample from the incoherent distribution.

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. (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 may extend² this results to the cross-temporal case. To obtain a reconciled forecast using the multivariate normal distribution, we start with a base forecast distributed as $\mathcal{N}(\hat{x}, \Sigma)$, where \hat{x}_h is the mean vector and Σ is the covariance matrix of the base forecasts. The reconciled forecast distribution is then given by $\mathcal{N}(\tilde{x}, \tilde{\Omega})$, where

$$\tilde{x} = M\hat{x} \quad \text{and} \quad \tilde{\Omega} = M\Sigma M', \quad (9)$$

where M is the projection matrix defined in (6). Note that if we assume that $\Sigma = \Omega_{ct}$, then the covariance matrix in (9) simplifies to $\tilde{\Omega} = M\Omega_{ct}$. In the cross-temporal case, sensibly estimating the covariance matrix Σ can be difficult because we need to simultaneously consider both the temporal and cross-sectional structures. This requires many parameters to be estimated, 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 matrix³. These challenges will be explored in more depth in the following sections.

²We assume $H = 1$ and simplify the notation by removing the h suffix without loss of generality

³In particular, some temporal covariances are fixed to zero (see the online appendix for more details).

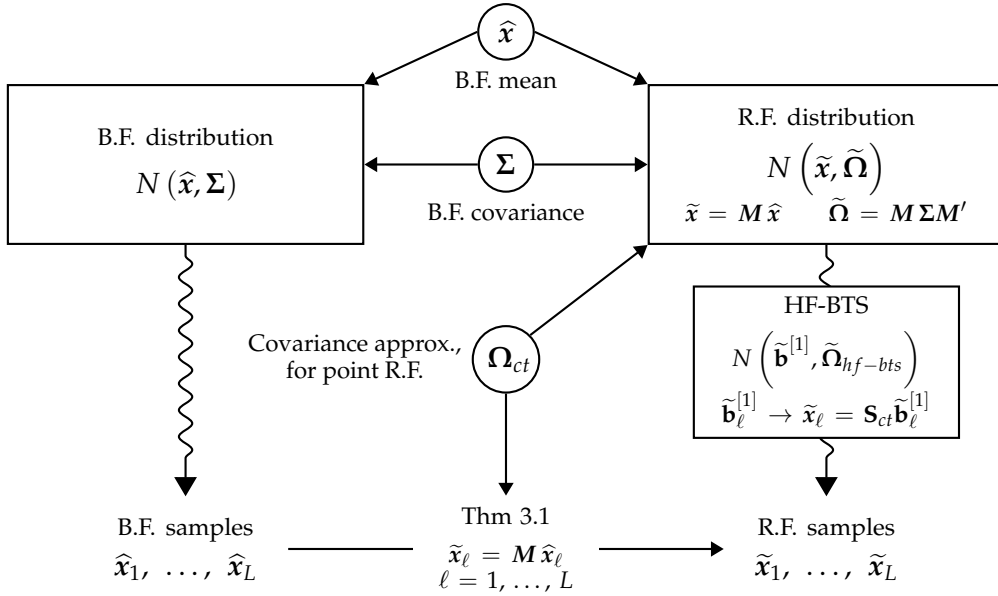


Figure 4: Visual description of cross-temporal forecast reconciliation in the Gaussian framework, as described in Section 3.1. The acronyms R.F and B.F. stand for Reconciled Forecasts and Base Forecasts, respectively. HF-BTS stands for High Frequency Bottom Time Series.

Focusing on the computational aspect⁴, 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 (that share the same top- and bottom-level variables, Di Fonzo & Girolimetto 2022d), 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 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 (6) as outlined in Theorem 3.1. The relationships between base and reconciled forecast distributions and their respective simulations through Theorem 3.1 are depicted in Figure 4.

3.2 Non-parametric framework: bootstrap reconciliation

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

⁴We use two R packages to sample from a the base forecast gaussian distribution: `Rfast` (Papadakis et al. 2022) when the dimension are very big (Section 7) and `MASS` (Venables & Ripley 2002) otherwise (Sections 5 and 6).




Figure 5: Example of bootstrapped residuals for 3 linearly constrained quarterly time series (see Figure 1). On the left there are the residual matrices with 4 years of data ($N = 4$): the green color corresponds to the first year, the blue to the second year, the red to the third year and the black to the fourth year. On the right the bootstrapped residuals are represented.

Let $\hat{\mathbf{E}}^{[k]}$ be the $(n \times N_k)$ matrix of the residuals for $k \in \mathcal{K}$. Figure 5 (on the left) provides a visualization of these matrices and how they are related to each other for the example in Figure 1. It is assumed that the residuals cover four years ($N = 4$): the green color corresponds to the first year, the 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. In this work, we assume \mathcal{M}_i to be a univariate model, however nothing prevents the use of multivariate models, perhaps for different temporal levels or for groups of time 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 univariate 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

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

Figure 5 (on the right) shows $\hat{\mathbf{E}}_{\tau}^{[k]}$, $k \in \{4, 2, 1\}$, for the quarterly cross-temporal hierarchy in Figure 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 in R (R Core Team 2022) using the package `forecast` (Hyndman et al. 2023) for many forecasting models, making it a practical and efficient tool. Furthermore, this approach is easily scalable in order to utilize multiple

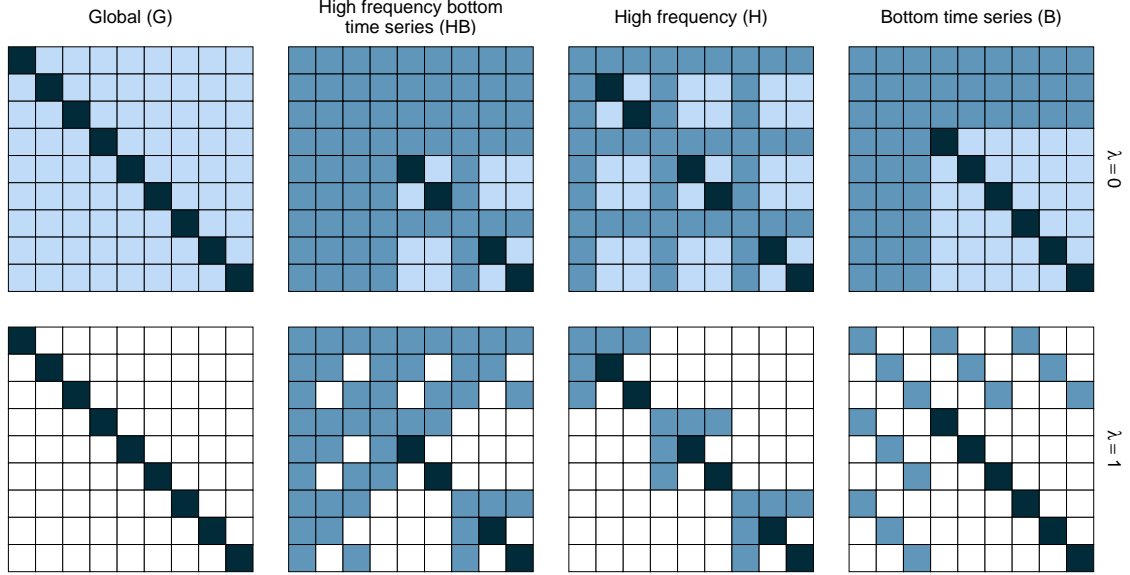


Figure 6: Representation of four types of covariance matrices that can be obtained from the cross-temporal hierarchical structure (3 time series and $m = 2$) for two different values of $\lambda \in \{0, 1\}$, the shrinkage parameter. The cells that are not modified by shrinkage are colored black, those actively involved in the shrinkage phase are colored light blue, and those derived from and not estimated by the base forecasts errors are colored blue. Additionally, for $\lambda = 1$, the cells corresponding to a zero value are colored white.

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 will be 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 $\lambda \in [0, 1]$ is the shrinkage intensity parameter that can be estimate using the unbiased estimator proposed by Ledoit & 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)* shown in Appendix A. 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

$$\Omega = \mathbf{S}_{ct} \Omega_{hf-bts} \mathbf{S}'_{ct} = (\mathbf{I}_n \otimes \mathbf{S}_{te}) \Omega_{hf} (\mathbf{I}_n \otimes \mathbf{S}_{te})' = (\mathbf{S}_{cs} \otimes \mathbf{I}_{m+k^*}) \Omega_{bts} (\mathbf{S}_{cs} \otimes \mathbf{I}_{m+k^*})', \quad (10)$$

Method	# of different parameters	AR(2)	GDP	Tourism
G	$r = \frac{n(k^* + m)[n(k^* + m) - 1]}{2}$	36	221 445	108 052 350
HB	$r_{HB} = \frac{n_b m [n_b m - 1]}{2} < r$	6 (83%)	30 876 (86%)	6 655 776 (94%)
H	$r_{HB} < \frac{nm[nm - 1]}{2} < r$	15 (58%)	72 390 (67%)	19 848 150 (82%)
B	$r_{HB} < \frac{n_b(k^* + m)[n_b(k^* + m) - 1]}{2} < r$	15 (58%)	94 395 (57%)	36 231 328 (66%)

Table 1: Number of different parameters that need to be estimated for the Monte Carlo simulation (AR(2), see Section 5), the Australian GDP (see Section 6) and the Australian Tourism Demand (see Section 7) forecasting experiments: the first one has 3 time series (one upper and two bottom) with temporal aggregation $\mathcal{K} = \{2, 1\}$; the second one has 95 quarterly ($m = 4$ and $k^* = 3$) time series (62 free and 33 constraints, see Di Fonzo & Girolimetto 2022d); the last one has a total of 525 monthly ($m = 12$ and $k^* = 16$) time series (304 bottom and 221 upper). The percentage reductions in the number of parameters compared to the global approach are reported in parentheses.

where $\mathbf{\Omega}_{hf-bts}$ is the $(n_b m \times n_b m)$ covariance matrix for the bottom time series at temporal aggregation level $k = 1$ (highest frequency bottom time series), $\mathbf{\Omega}_{hf}$ is the $(nm \times nm)$ covariance matrix related to all the high frequency time series and $\mathbf{\Omega}_{bts}$ is the $[n_b(k^* + m) \times n_b(k^* + m)]$ covariance matrix related to bottom time series at any temporal aggregation.

Therefore, we can apply the idea of ‘‘Stein-type shrinkage’’ (Efron & Morris 1977) to $\mathbf{\Omega}_{hf-bts}$, $\mathbf{\Omega}_{hf}$ and $\mathbf{\Omega}_{bts}$ by using the corresponding empirical base forecasts residuals estimation. We obtain the following expressions (see the online appendix for details):

- High frequency Bottom time series shrinkage matrix (HB):

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

- High frequency shrinkage matrix (H):

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

- Bottom time series shrinkage matrix (B):

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

where $\widehat{\mathbf{\Omega}}_{l,D} = \mathbf{I}_{n_b m} \odot \widehat{\mathbf{\Omega}}_j$, $l = \{hf-bts, hf, bts\}$, and λ is the shrinkage parameter.

Figure 6 gives some visual insights on the covariance matrices obtainable with $\lambda = 0$ and $\lambda = 1$, respectively, for a simple cross-temporal hierarchical structure with 3 time series and $\mathcal{K} = \{2, 1\}$ (e.g, cross-temporal semi-annual, see the Monte Carlo simulation in Section 5). 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 Monte Carlo simulation (AR2) and for the two forecasting experiment: Australian GDP (see Section 6) and Australian Tourism Demand (see Section 7). In addition, we calculate also the percentage reductions in the number of parameters compared to the global approach, that is: % reduction = $100(1 - r_i/r_G)$ with $i \in \{HB, H, B\}$. As we can see, G involves a considerable number of parameters compared to other procedures. We observe that HB leads

to a decrease of around 85%, whereas the H and B approaches look as a compromise between the previous two. In general, as m and n increase (see the online appendix), using H requires the estimation of a smaller number of parameters.

In our simulations and forecasting experiments, we will be closely analyzing these different constructions with a dual purpose. First, we will use the full covariance matrix ($\lambda = 0$) of the base forecasts to obtain a base forecasts sample of the linearly constrained time series under Gaussianity assumption. Then, we will use the shrinkage versions as approximations of the covariance matrix to be used for reconciliation. This will allow 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, but because of the different temporal dimensions, we need residuals corresponding to different forecast horizons. 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 will be autocorrelated except when $h = 1$.

Following Di Fonzo & 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 $\mathbf{E}_i = [\mathbf{E}_i^{[m]} \ \mathbf{E}_i^{[k_p-1]} \ \dots \ \mathbf{E}_i^{[1]}]$. Then the $[N \times n(m + k^*)]$ cross-temporal residual matrix is given by $\mathbf{E} = [\mathbf{E}_1 \ \mathbf{E}_2 \ \dots \ \mathbf{E}_n]$.

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 $\mathbf{x}^{[k],s} = \{x_j^{[k],s}\}_{j=1}^{N_k-s}$ 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 biannual 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: $\mathbf{x}^{[2],0} = [x_1^{[2],0} \ x_2^{[2],0} \ x_3^{[2],0}]' = [y_1 + y_2 \ y_3 + y_4 \ y_5 + y_6]'$ and $\mathbf{x}^{[2],1} = [x_1^{[2],1} \ x_2^{[2],1}]' = [y_2 + y_3 \ y_4 + y_5]'$. 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 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 Monte Carlo simulation

We study the effect of combining cross-sectional and temporal aggregations, using a simple hierarchy that allows us to effectively visualize the quantities involved, such as the covariance matrix. Additionally, the small size and nature of the data generating process make it possible to exactly calculate the true cross-temporal covariance structure, thus providing insights into the nature of the time series data involved in the forecast reconciliation process.

Consider a 2-level hierarchical structure with three time series (one upper series, A , and two bottom series, B and C) such that the cross-sectional aggregation matrix is $A_{cs} = [1 \ 1]$ ($A = B + C$). The temporal structure we are considering is equivalent to using semi-annual data with $\mathcal{K} = \{2, 1\}$ and $m = 2$. The assumed Data-Generating Processes (DPG) for the semi-annual bottom level series are two AR(2) given by

$$\begin{aligned} y_{B,t} &= \phi_{B,1}y_{B,t-1} + \phi_{B,2}y_{B,t-2} + \varepsilon_{B,t} \\ y_{C,t} &= \phi_{C,1}y_{C,t-1} + \phi_{C,2}y_{C,t-2} + \varepsilon_{C,t} \end{aligned}$$

with parameters⁵ $\boldsymbol{\phi}_B = [\phi_{B,1} \ \phi_{B,2}]' = [1.34 \ -0.74]'$ and $\boldsymbol{\phi}_C = [\phi_{C,1} \ \phi_{C,2}]' = [0.95 \ -0.42]'$. The error $\boldsymbol{\varepsilon}_t = [\varepsilon_{B,t} \ \varepsilon_{C,t}]'$ driving the process is drawn from a multivariate normal distribution with standard deviations simulated from a uniform distribution between 0.5 and 2 and a fixed correlation of -0.8. The cross-sectional error covariance matrix is thus given by

$$\boldsymbol{\Omega}_{cs} = \begin{bmatrix} 0.9 & 0 \\ 0 & 1.8 \end{bmatrix} \begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix} \begin{bmatrix} 0.9 & 0 \\ 0 & 1.8 \end{bmatrix} = \begin{bmatrix} \sigma_B^2 & \sigma_{BC} \\ \sigma_{BC} & \sigma_C^2 \end{bmatrix}.$$

To obtain the remaining series, the bottom series are then cross-temporally aggregated.

For the forecast experiment, the base forecasts are computed using AR models where the order is automatically determined by the algorithm proposed by Hyndman & Khandakar (2008) and implemented in the R package `forecast` (Hyndman et al. 2023), thus allowing for possible mis-specification in the models. The training window length is 500 years, consisting

⁵The $\boldsymbol{\phi}_B$ and $\boldsymbol{\phi}_C$ parameters are estimated from the ‘‘Lynx’’ and ‘‘Hare’’ time series contained in the `pel1t` dataset of the `tsibbledata` package for R (O’Hara-Wild et al. 2022).

Label	Description
$ct(bu)$	Simple cross-temporal bottom-up (Section 2.2).
$ct(\cdot, bu_{te})$	Partly bottom-up (Section 2.2) starting from cross-sectional reconciled forecasts using the <i>shr</i> and <i>wls</i> approaches (Table A.1).
$ct(wlsv_{te}, bu_{cs})$	Partly bottom-up (Section 2.2) starting from temporally reconciled forecasts using the <i>wlsv</i> approach (Table A.1).
$oct(\cdot)$	Optimal cross-temporal reconciliation for the <i>ols</i> , <i>struc</i> , <i>wlsv</i> and <i>bdshr</i> approaches (see Appendix A). One-step residuals were used with <i>wlsv</i> and <i>bdshr</i> .
$oct_h(\cdot)$	Optimal cross-temporal reconciliation with multi-step residuals (see Section 4.1) for the shrinkage approaches presented in Section 4: <i>shr</i> stands for <i>Global shrinkage</i> , <i>hshr</i> for <i>High frequency shrinkage</i> , <i>bshr</i> for <i>bottom time series shrinkage</i> , <i>hbshr</i> for <i>High frequency bottom time series shrinkage</i> .
$oct_o(\cdot)$	Optimal cross-temporal reconciliation with overlapping residuals (see Section 4.2) for the <i>wlsv</i> and <i>bdshr</i> approaches (Appendix A).
$oct_{oh}(\cdot)$	Optimal cross-temporal reconciliation with overlapping and multi-step residuals (see Section 4.1 and 4.2) for the shrinkage approaches presented in Section 4: <i>shr</i> stands for <i>Global shrinkage</i> , <i>hshr</i> for <i>High frequency shrinkage</i> .

Table 2: Cross-temporal reconciliation approaches for the Monte Carlo simulation (see Section 5), the Australian GDP (see Section 6) and the Australian Tourism Demand (see Section 7) forecasting experiments. All the reconciliation procedures are available in the R package *FoReco* (Girolimetto & Di Fonzo 2023).

be able to determine which reconciliation approach and simulation technique produce an accurate estimate of the covariance matrix. Other types of matrix norms were also considered with similar results.

From Table 3, it appears that the reconciled covariance matrices are always closer to the true matrix than the base forecast matrix when using both the Gaussian and the bootstrap approach. Overall, there are no major differences in the findings when using either one-step or multi-step residuals in cross-temporal forecast reconciliation. In fact, using approaches like $oct(bdshr)$, we obtain results that are consistent with approaches such as $oct_h(shr)$, where no temporal and/or cross-sectional correlation assumptions are imposed. It is worth noting that the *HB* covariance matrix when used to calculate the base forecasts samples, is not changed by the reconciliation step (see the online appendix). In conclusion, our results suggest that using multi-step residuals or bootstrap techniques may help find a “good” estimate of the covariance matrix, which can be further improved by the reconciliation.

The accuracy of the probabilistic forecasts is evaluated using the Continuous Ranked Probability Score (CRPS, Gneiting & Katzfuss 2014), given by

$$CRPS(\hat{P}_i, z_i) = \frac{1}{L} \sum_{l=1}^L |x_{i,l} - z_i| - \frac{1}{2L^2} \sum_{l=1}^L \sum_{j=1}^L |x_{i,l} - x_{i,j}|, \quad i = 1, \dots, n, \quad (11)$$

Reconciliation approach	Generation of the base forecasts paths				
	ctjb	Gaussian approach*			
		G	B	H	HB
base	8.260	7.748	6.549	3.409	2.215
ct(<i>bu</i>)	3.195	2.215	2.215	2.215	2.215
ct(<i>shr_{cs}, bu_{te}</i>)	3.202	2.224	2.215	2.224	2.215
ct(<i>wlsv_{te}, bu_{cs}</i>)	3.183	2.188	2.188	2.215	2.215
oct(<i>wlsv</i>)	3.766	3.082	2.191	2.910	2.215
oct(<i>bdshr</i>)	3.203	2.195	2.184	2.224	2.215
oct _{<i>t</i>} (<i>shr</i>)	3.251	2.260	2.202	2.226	2.215
oct _{<i>t</i>} (<i>bshr</i>)	3.602	2.720	2.220	2.756	2.215
oct _{<i>t</i>} (<i>hshr</i>)	4.869	4.138	4.167	2.225	2.215
oct _{<i>t</i>} (<i>hbshr</i>)	5.731	5.085	4.167	2.756	2.215

*The Gaussian method employs a sample covariance with multi-step residuals.

Table 3: Simulation experiment. Frobenius norm between the true and the estimated covariance matrix for different reconciliation approaches and different techniques for simulating the base forecasts. Entries in bold represent the lowest value for each column, while the blue entry represent the global minimum. The reconciliation approaches are described in Table 2.

where $\hat{P}_i(\omega) = \frac{1}{L} \sum_{l=1}^L \mathbf{1}(x_{i,l} \leq \omega)$, $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_L \in \mathbb{R}^n$ is a collection of L random draws from the predictive distribution and $\mathbf{z} \in \mathbb{R}^n$ is the observation vector. CRPS is an index that considers the single series and provides us a marginal evaluation of the approaches. In addition, we employ the Energy Score (ES, Gneiting & Katzfuss 2014), that is the CRPS extension to the multivariate case, to evaluate the forecasting accuracy for the whole system (Panagiotelis et al. 2023, Wickramasuriya 2023),

$$ES(\hat{P}, \mathbf{z}) = \frac{1}{L} \sum_{l=1}^L \|\mathbf{x}_l - \mathbf{z}\|_2 - \frac{1}{2(L-1)} \sum_{i=1}^{L-1} \|\mathbf{x}_i - \mathbf{x}_{i+1}\|_2 \quad (12)$$

where $\|\cdot\|_2$ is the L_2 norm. In particular, we consider the geometric mean of the relative CRPS (Fleming & Wallace 1986), and the relative ES:

$$\text{AvgRelCRPS}_{j,s}^{[k]} = \left(\prod_{i=1}^n \frac{CRPS_{i,j,s}^{[k]}}{CRPS_{i,0,0}^{[k]}} \right)^{\frac{1}{n}} \quad \text{and} \quad \text{RelES}_{j,s}^{[k]} = \frac{ES_{j,s}^{[k]}}{ES_{0,0}^{[k]}}, \quad (13)$$

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 produce by the Bootstrap approach. If we consider all the temporal aggregation orders (i.e. $\forall k \in \{2, 1\}$), the overall accuracy index is given by

$$\text{AvgRelCRPS}_{j,s} = \left(\prod_{\substack{i=1, \dots, n \\ k \in \{1, 2\}}} \frac{CRPS_{i,j,s}^{[k]}}{CRPS_{i,0,0}^{[k]}} \right)^{\frac{1}{n(k^*+m)}}, \quad (14)$$

and

$$\text{AvgRelES}_{j,s} = \left(\prod_{k \in \{1, 2\}} \frac{ES_{j,s}^{[k]}}{ES_{0,0}^{[k]}} \right)^{\frac{1}{(k^*+m)}}. \quad (15)$$

Reconciliation approach	Generation of the base forecasts sample paths									
	ctjb	Gaussian approach*				ctjb	Gaussian approach*			
		G	B	H	HB		G	B	H	HB
	$\forall k \in \{2, 1\}$					$k = 1$				
base	1.000	0.998	0.999	1.002	1.004	1.000	0.998	0.999	0.999	1.000
ct(<i>bu</i>)	0.901	0.900	0.900	0.900	0.900	0.978	0.976	0.976	0.977	0.977
ct(<i>shr_{cs}, bu_{te}</i>)	0.901	0.900	0.899	0.900	0.900	0.977	0.976	0.976	0.976	0.976
ct(<i>wlsv_{te}, bu_{cs}</i>)	0.910	0.916	0.916	0.916	0.917	0.986	0.993	0.993	0.993	0.993
oct(<i>wlsv</i>)	0.922	0.930	0.930	0.930	0.931	0.998	1.006	1.006	1.007	1.007
oct(<i>bdshr</i>)	0.910	0.916	0.915	0.916	0.916	0.986	0.992	0.992	0.993	0.993
oct _{<i>h</i>} (<i>shr</i>)	0.904	0.903	0.902	0.902	0.903	0.980	0.979	0.978	0.979	0.979
oct _{<i>h</i>} (<i>bshr</i>)	0.923	0.922	0.922	0.921	0.922	0.999	0.998	0.998	0.998	0.998
oct _{<i>h</i>} (<i>hshr</i>)	0.974	0.972	0.972	0.974	0.975	1.052	1.050	1.050	1.053	1.053
oct _{<i>h</i>} (<i>hbshr</i>)	0.987	0.985	0.985	0.987	0.989	1.065	1.063	1.064	1.066	1.068
	$k = 2$									
base	1.000	0.998	0.999	1.005	1.008					
ct(<i>bu</i>)	0.831	0.830	0.829	0.829	0.830					
ct(<i>shr_{cs}, bu_{te}</i>)	0.830	0.830	0.829	0.829	0.830					
ct(<i>wlsv_{te}, bu_{cs}</i>)	0.840	0.846	0.844	0.845	0.846					
oct(<i>wlsv</i>)	0.851	0.859	0.859	0.859	0.861					
oct(<i>bdshr</i>)	0.839	0.845	0.844	0.845	0.846					
oct _{<i>h</i>} (<i>shr</i>)	0.834	0.833	0.831	0.832	0.832					
oct _{<i>h</i>} (<i>bshr</i>)	0.852	0.851	0.851	0.851	0.852					
oct _{<i>h</i>} (<i>hshr</i>)	0.902	0.900	0.899	0.901	0.902					
oct _{<i>h</i>} (<i>hbshr</i>)	0.915	0.913	0.913	0.914	0.917					

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

Table 4: Simulation experiment. *AvgRelCRPS* defined in (13) and (14). 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.

In Tables 4 and 5 are reported the results using these two scores where low values indicate better quality of the forecasts.

A limitation of this simulation setting is that we are using a high number of residuals, which may result in undervaluing the benefit from using the parameterization form of the covariance matrix such as *HB*, *H*, or *B* (see Section 4). Additionally, shrinkage techniques often yield very similar results when we use the corresponding matrix with $\lambda = 0$ (full covariance matrix).

The good performance of the *ct(bu)* approach can be explained by a good quality of the base forecasts at the bottom level for $k = 1$, and therefore it is difficult for the other approaches to correctly adjust them using the somewhat less good forecasts of the higher temporal and cross-sectional levels. This also explains the good performance of *ct(shr_{cs}, bu_{te})*, which by definition only takes into account the information provided by the most temporally disaggregated base forecasts. Looking at the optimal cross-temporal reconciliation approaches, it does not seem to be any advantage in using multi-step residuals to calculate the covariance matrix in the reconciliation step.

In conclusion, we found that simulating base forecasts from multi-step residuals allows us to estimate a covariance matrix close to the true one. Additionally, we observed that reconciliation could be used to further improve the accuracy of these estimates: accurate base

Reconciliation approach	Generation of the base forecasts sample paths									
	ctjb	Gaussian approach*				ctjb	Gaussian approach*			
		G	B	H	HB		G	B	H	HB
	$\forall k \in \{2, 1\}$					$k = 1$				
base	1.000	0.996	0.999	1.000	1.004	1.000	0.997	1.000	0.997	1.000
ct(<i>bu</i>)	0.897	0.895	0.896	0.897	0.895	0.969	0.967	0.967	0.968	0.968
ct(<i>shr_{cs}, bu_{te}</i>)	0.896	0.895	0.895	0.896	0.896	0.968	0.968	0.967	0.968	0.968
ct(<i>wlsv_{te}, bu_{cs}</i>)	0.906	0.912	0.911	0.910	0.912	0.977	0.984	0.983	0.981	0.984
oct(<i>wlsv</i>)	0.916	0.923	0.923	0.923	0.924	0.989	0.994	0.995	0.995	0.997
oct(<i>bdshr</i>)	0.906	0.910	0.910	0.911	0.912	0.977	0.981	0.982	0.983	0.985
oct _{<i>h</i>} (<i>shr</i>)	0.900	0.898	0.898	0.897	0.898	0.971	0.969	0.969	0.969	0.969
oct _{<i>h</i>} (<i>bshr</i>)	0.916	0.914	0.916	0.915	0.916	0.987	0.986	0.987	0.987	0.988
oct _{<i>h</i>} (<i>hshr</i>)	0.967	0.964	0.964	0.966	0.967	1.040	1.036	1.036	1.040	1.040
oct _{<i>h</i>} (<i>hbshr</i>)	0.978	0.975	0.976	0.977	0.980	1.051	1.047	1.049	1.051	1.052
	$k = 2$									
base	1.000	0.996	0.998	1.003	1.008					
ct(<i>bu</i>)	0.831	0.829	0.829	0.830	0.828					
ct(<i>shr_{cs}, bu_{te}</i>)	0.829	0.828	0.829	0.829	0.829					
ct(<i>wlsv_{te}, bu_{cs}</i>)	0.839	0.844	0.844	0.844	0.845					
oct(<i>wlsv</i>)	0.849	0.858	0.856	0.856	0.857					
oct(<i>bdshr</i>)	0.839	0.845	0.843	0.845	0.844					
oct _{<i>h</i>} (<i>shr</i>)	0.835	0.833	0.833	0.831	0.832					
oct _{<i>h</i>} (<i>bshr</i>)	0.850	0.847	0.849	0.849	0.850					
oct _{<i>h</i>} (<i>hshr</i>)	0.900	0.897	0.896	0.897	0.899					
oct _{<i>h</i>} (<i>hbshr</i>)	0.910	0.907	0.908	0.909	0.912					

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

Table 5: Simulation experiment. ES ratio indices defined in (13) and (15). 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.

forecasts for $k = 1$ assist the good performance for bottom-up and optimal cross-temporal reconciliation approaches, such as oct(*wlsv*) and oct(*bdshr*), which perform well in terms of both CRPS and ES.

6 Forecasting Australian GDP

The Australian Quarterly National Accounts (QNA) dataset has been widely used in the literature on forecast reconciliation (Athanasopoulos et al. 2020, Bisaglia et al. 2020, Di Fonzo & Girolimetto 2022c,d, 2023a). Building on these results (see the online appendix), now we consider the 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 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 4-step, 2-step and 1-step

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

ahead base forecasts respectively from the quarterly, semi-annual and annual frequencies, 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) parameterization (Section 4), since it is not possible to identify a unique representation for the other cases⁷. 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 presented in Section 3.2. Finally, to reconcile the resulting (1000) base forecasts samples, we have applied the following techniques⁸ (see Table 2): $ct(shr_{cs}, bu_{te})$, $ct(wls_{cs}, bu_{te})$, $oct_o(wlsv)$, $oct_o(bdshr)$, $oct_{oh}(shr)$ and $oct_{oh}(hshr)$.

6.1 Results

Forecasting accuracy indices based on CRPS and ES are presented in Tables 6 and 7, respectively. 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. Diagonal covariance matrices were found to be more effective than matrices that try to recover the correlation structure through shrinkage. Among all the procedures, $ct(wls_{cs}, bu_{te})$ and $oct_o(wlsv)$ show the greatest relative gains. In contrast, $oct_{oh}(shr)$ and $oct_{oh}(hshr)$ do not show much improvement. Furthermore, the greatest improvements are observed for higher temporal aggregation levels.

Additionally, 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) to determine if the forecasting performances of the different techniques are significantly different from one another. Figure 8 presents the MCB using the CRPS. We found that $ct(wls_{cs}, bu_{te})$ and $oct_o(wlsv)$ are significantly better than the base forecasts at any level of aggregation.

Overall, we find that using overlapping residuals almost always leads to a greater improvement in terms of ES and CRPS, and this is also generally true for reconciliation. 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 take into account the multi-step structure.

7 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

⁷When simultaneously considering Income and Expenditure sides hierarchies, the result is a general linearly constrained time series, where bottom and upper time series are not uniquely defined, making unfeasible the cross-sectional bottom-up reconciliation approach (Di Fonzo & Girolimetto 2022d).

⁸In the online appendix, we show the results with shrunk covariance matrices. We also report the results obtained using one-step residuals in the reconciliation.

Reconciliation approach	Generation of the base forecasts sample paths									
	ctjb	Gaussian approach*				ctjb	Gaussian approach*			
		G_h	H_h	G_{oh}	H_{oh}		G_h	H_h	G_{oh}	H_{oh}
		$\forall k \in \{4, 2, 1\}$					$k = 1$			
base	1.000	0.979	0.995	0.968	0.976	1.000	0.988	0.988	0.971	0.971
ct(shr_{cs}, bu_{te})	0.937	0.956	0.956	0.976	0.976	0.992	1.008	1.008	1.029	1.029
ct(wls_{cs}, bu_{te})	0.930	0.917	0.917	0.898	0.898	0.986	0.974	0.975	0.956	0.956
oct ₀ ($wlsv$)	0.926	0.911	0.912	0.896	0.895	0.984	0.971	0.970	0.954	0.954
oct ₀ ($bdshr$)	0.978	0.964	0.946	0.952	0.930	1.034	1.016	1.003	1.005	0.989
oct _{oh} (shr)	1.102	1.059	1.001	1.094	0.988	1.172	1.109	1.066	1.160	1.059
oct _{oh} ($hshr$)	1.006	0.983	1.009	0.974	1.001	1.068	1.046	1.059	1.034	1.061
		$k = 2$					$k = 4$			
base	1.000	0.984	0.993	0.968	0.976	1.000	0.966	1.004	0.964	0.981
ct(shr_{cs}, bu_{te})	0.949	0.966	0.966	0.987	0.987	0.874	0.896	0.896	0.914	0.914
ct(wls_{cs}, bu_{te})	0.942	0.928	0.928	0.909	0.909	0.866	0.853	0.853	0.834	0.834
oct ₀ ($wlsv$)	0.938	0.921	0.923	0.907	0.906	0.860	0.847	0.848	0.832	0.830
oct ₀ ($bdshr$)	0.991	0.974	0.957	0.964	0.942	0.914	0.905	0.883	0.892	0.865
oct _{oh} (shr)	1.120	1.069	1.013	1.113	1.002	1.020	1.002	0.928	1.015	0.909
oct _{oh} ($hshr$)	1.021	0.996	1.021	0.987	1.016	0.934	0.912	0.951	0.904	0.931

*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.

Table 6: AvgRelCRPS defined in (13) and (14) 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	Generation of the base forecasts sample paths									
	ctjb	Gaussian approach*				ctjb	Gaussian approach*			
		G_h	H_h	G_{oh}	H_{oh}		G_h	H_h	G_{oh}	H_{oh}
		$\forall k \in \{4, 2, 1\}$					$k = 1$			
base	1.000	0.970	0.988	0.960	0.970	1.000	0.977	0.977	0.965	0.965
ct(shr_{cs}, bu_{te})	0.897	0.944	0.944	0.973	0.973	0.964	1.001	1.001	1.033	1.033
ct(wls_{cs}, bu_{te})	0.886	0.880	0.880	0.860	0.860	0.954	0.944	0.945	0.928	0.928
oct ₀ ($wlsv$)	0.891	0.879	0.881	0.864	0.864	0.958	0.945	0.945	0.931	0.931
oct ₀ ($bdshr$)	0.940	0.928	0.910	0.918	0.895	1.004	0.986	0.971	0.980	0.961
oct _{oh} (shr)	1.059	1.015	0.956	1.053	0.945	1.130	1.063	1.019	1.121	1.016
oct _{oh} ($hshr$)	0.986	0.968	0.999	0.959	0.992	1.053	1.034	1.049	1.024	1.055
		$k = 2$					$k = 4$			
base	1.000	0.972	0.985	0.959	0.969	1.000	0.959	1.000	0.957	0.976
ct(shr_{cs}, bu_{te})	0.915	0.961	0.960	0.991	0.991	0.818	0.874	0.874	0.899	0.900
ct(wls_{cs}, bu_{te})	0.904	0.896	0.896	0.877	0.877	0.807	0.805	0.805	0.782	0.783
oct ₀ ($wlsv$)	0.908	0.895	0.898	0.881	0.882	0.812	0.802	0.806	0.786	0.786
oct ₀ ($bdshr$)	0.960	0.947	0.929	0.938	0.915	0.860	0.856	0.836	0.841	0.816
oct _{oh} (shr)	1.082	1.029	0.973	1.076	0.963	0.971	0.954	0.882	0.967	0.861
oct _{oh} ($hshr$)	1.007	0.988	1.017	0.979	1.014	0.904	0.888	0.934	0.881	0.913

*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.

Table 7: ES ratio indices defined in (13) and (15) 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.

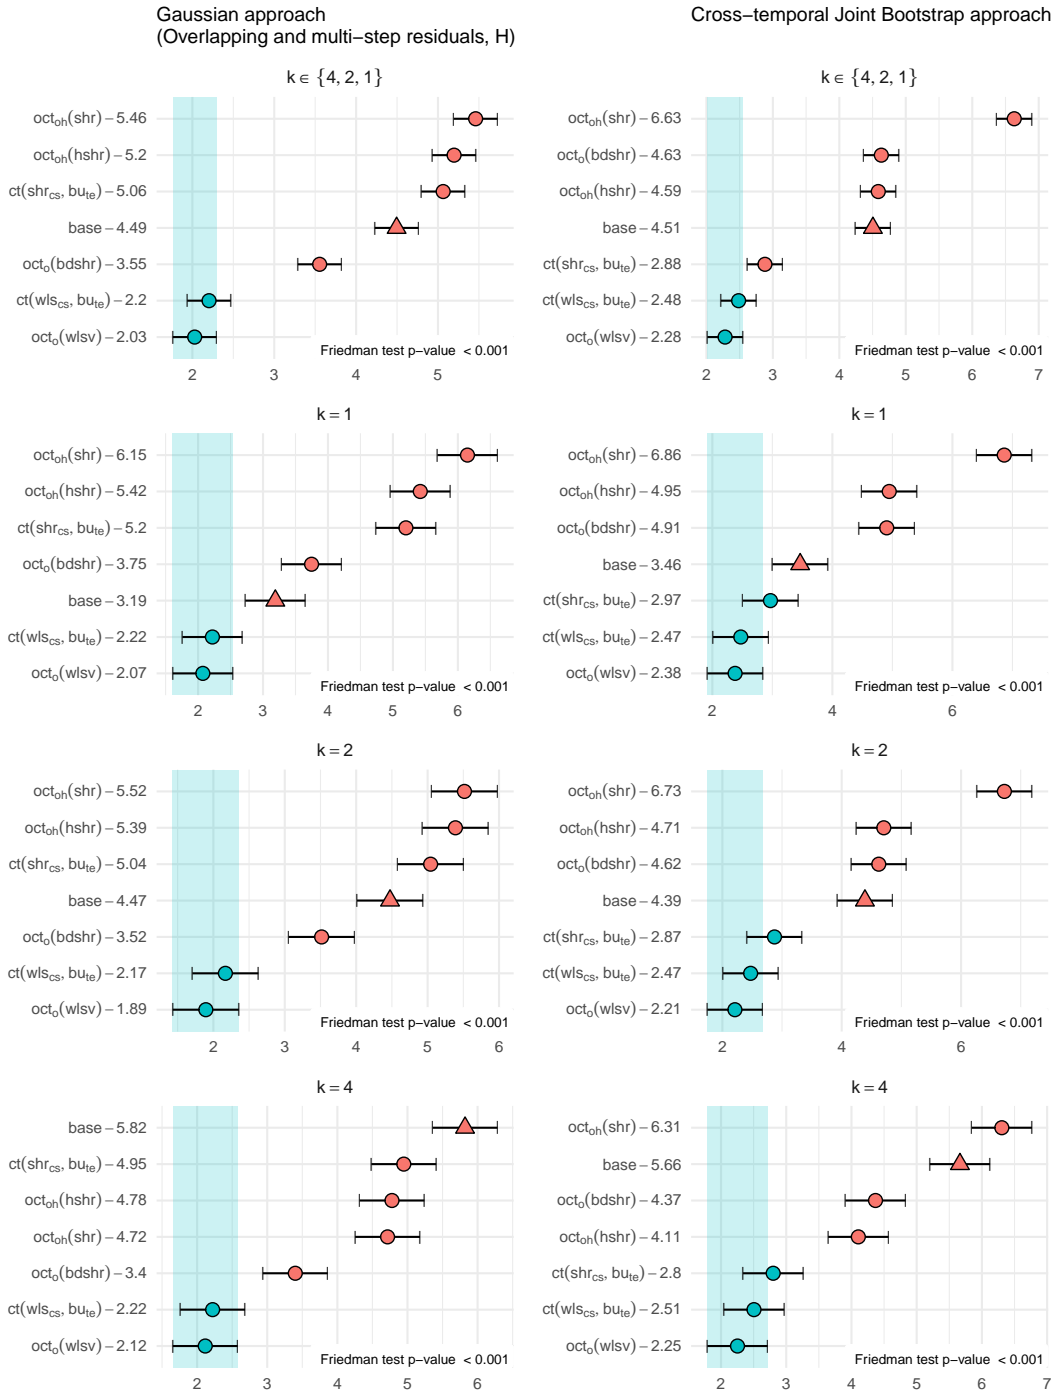


Figure 8: MCB Nemenyi test for the Australian QNA dataset using CRPS at different temporal aggregation levels for the Gaussian (using overlapping and multi-step residuals, H) and the non-parametric bootstrap approaches. In each panel, the Friedman test p-value is reported in the lower right corner. The mean rank of each approach is shown to the right of its name. Statistical 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.

Nights (VN) 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 (see the online appendix) are each formed by a single region, resulting in a total of 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 8 omits duplicated entries and updates the information in Table 7 from Wickramasuriya et al. (2019). This data can be temporally aggregated into 2, 3, 4, 6, or 12 months ($\mathcal{K} = \{12, 4, 3, 2, 1\}$).

The forecasting experiment involves using a recursive training sample with an expanding window. The process begins by using the first 10 years of data, 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 forecast origins. For the temporal aggregation dimension we aggregate the monthly data up to annual data. We obtain 12-step, 6-step, 4-step, 3-step, 2-step and 1-step ahead base forecasts respectively from the monthly data and the aggregation over 2, 3, 4, 6, and 12 months. ETS models selected by minimizing the AICc criterion with the R package `forecast` are fitted to the log-transformed data, with the resulting base forecasts being back-transformed to produce non-negative forecasts, as described in 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, 11 different approaches have been adopted (see Table 2): $ct(bu)$, $ct(shr_{cs}, bu_{te})$, $ct(wlsv_{te}, bu_{cs})$, $ct(ols)$, $oct(struc)$, $oct(wlsv)$, $oct(bdshr)$, $oct_h(hbshr)$, $oct_h(bshr)$, $oct_h(hshr)$, and $oct_h(shr)$.

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

* 6 Zones with only one Region are included in Regions.

Table 8: Grouped time series for the Australian Tourism Demand dataset. GD: Geographic Division; PT: Purpose of Travel.

⁹The results with shrunk covariance matrices are available in the online appendix.

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

Reconciliation approach	Generation of the base forecasts sample paths										
	ctjb	Gaussian approach*				ctjb	Gaussian approach*				
		G	B	H	HB		G	B	H	HB	
		$\forall k \in \{12, 6, 4, 3, 2, 1\}$						$k = 1$			
base	1.000	0.971	0.971	0.973	0.973	1.000	0.972	0.972	0.972	0.972	
ct(<i>bu</i>)	1.321	1.011	1.011	1.011	1.011	1.077	0.983	0.982	0.982	0.982	
ct(<i>shr_{cs}, bu_{te}</i>)	1.057	0.974	0.969	0.974	0.969	0.976	0.963	0.962	0.963	0.962	
ct(<i>wlsv_{te}, bu_{cs}</i>)	1.062	0.974	0.974	0.972	0.972	0.976	0.965	0.965	0.966	0.966	
oct(<i>ols</i>)	0.989	0.989	0.989	0.987	0.987	0.982	0.986	0.988	0.986	0.989	
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 _h (<i>hbshr</i>)	0.989	1.018	1.020	1.016	1.018	0.982	1.004	1.007	1.004	1.009	
oct _h (<i>bshr</i>)	0.994	1.018	1.020	1.016	1.019	0.988	1.007	1.013	1.006	1.012	
oct _h (<i>hshr</i>)	0.969	0.993	0.993	0.990	0.991	0.953	0.977	0.977	0.979	0.979	
oct _h (<i>shr</i>)	1.007	0.980	0.972	0.970	0.970	1.000	0.986	0.977	0.976	0.974	
		$k = 3$						$k = 12$			
base	1.000	0.971	0.971	0.972	0.973	1.000	0.968	0.967	0.969	0.969	
ct(<i>bu</i>)	1.273	1.010	1.010	1.010	1.010	1.675	1.038	1.037	1.037	1.038	
ct(<i>shr_{cs}, bu_{te}</i>)	1.041	0.977	0.974	0.977	0.974	1.163	0.977	0.965	0.977	0.965	
ct(<i>wlsv_{te}, bu_{cs}</i>)	1.046	0.976	0.976	0.974	0.974	1.174	0.978	0.978	0.971	0.971	
oct(<i>ols</i>)	0.994	0.992	0.993	0.991	0.992	0.982	0.982	0.983	0.980	0.975	
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 _h (<i>hbshr</i>)	0.994	1.019	1.021	1.018	1.020	0.982	1.027	1.029	1.024	1.021	
oct _h (<i>bshr</i>)	0.999	1.021	1.022	1.018	1.022	0.987	1.024	1.021	1.021	1.019	
oct _h (<i>hshr</i>)	0.971	0.994	0.994	0.992	0.993	0.978	1.003	1.005	0.996	0.997	
oct _h (<i>shr</i>)	1.009	0.986	0.978	0.976	0.976	1.010	0.963	0.956	0.952	0.952	

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

Table 9: AvgRelCRPS defined in (13) and (14) 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.

Negative forecasts may be produced during the reconciliation phase (Wickramasuriya et al. 2020, Di Fonzo & Girolimetto 2022b, 2023b) thus generating unreasonable values (e.g., a negative forecast for tourism demand makes no sense). To overcome this limitation (see the online appendix), we applied the simple heuristic proposed by Di Fonzo & Girolimetto (2022a, 2023b). Following Theorem 3.1, we are thus able to obtain the reconciled sample respecting non-negativity constraints starting from an incoherent sample simulated from a Gaussian distribution. Finally, to evaluate the performance, we employed the Continuous Ranked Probability Score (CRPS), the Energy Score (ES), and the ‘‘Multiple Comparison with the Best’’ (MCB) Nemenyi test, introduced in Sections 5.1 and 6.1.

7.1 Results

The CRPS and ES indices are shown, respectively, in Tables 9 and 10 for monthly, quarterly and annual forecasts¹¹. 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 was used as a benchmark to calculate the accuracy indices.

¹¹The complete set of results for all temporal aggregation levels is reported in the online appendix.

Reconciliation approach	Generation of the base forecasts sample paths									
	ctjb	Gaussian approach*				ctjb	Gaussian approach*			
		G	B	H	HB		G	B	H	HB
		$\forall k \in \{12, 6, 4, 3, 2, 1\}$					$k = 1$			
base	1.000	0.956	0.955	0.958	0.951	1.000	0.952	0.950	0.952	0.950
ct(<i>bu</i>)	2.427	0.983	0.983	0.983	0.983	1.759	0.982	0.982	0.982	0.982
ct(<i>shr_{cs}, bu_{te}</i>)	1.243	0.886	0.879	0.886	0.879	1.098	0.929	0.928	0.930	0.927
ct(<i>wlsv_{te}, bu_{cs}</i>)	1.499	0.977	0.977	0.971	0.972	1.241	0.975	0.975	0.973	0.974
oct(<i>ols</i>)	0.955	0.893	0.891	0.893	0.888	0.975	0.937	0.936	0.936	0.935
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 _h (<i>hbshr</i>)	0.956	0.889	0.886	0.888	0.884	0.975	0.937	0.936	0.937	0.935
oct _h (<i>bshr</i>)	0.931	0.867	0.866	0.863	0.860	0.965	0.927	0.927	0.925	0.923
oct _h (<i>hshr</i>)	1.081	0.935	0.931	0.935	0.927	1.028	0.952	0.951	0.952	0.950
oct _h (<i>shr</i>)	1.068	0.899	0.878	0.875	0.864	1.023	0.935	0.923	0.921	0.916
		$k = 3$					$k = 12$			
base	1.000	0.961	0.958	0.960	0.955	1.000	0.942	0.947	0.951	0.937
ct(<i>bu</i>)	2.428	0.998	0.997	0.997	0.997	2.990	0.922	0.921	0.923	0.923
ct(<i>shr_{cs}, bu_{te}</i>)	1.245	0.911	0.904	0.911	0.904	1.326	0.779	0.767	0.777	0.766
ct(<i>wlsv_{te}, bu_{cs}</i>)	1.500	0.991	0.991	0.986	0.987	1.679	0.917	0.917	0.906	0.908
oct(<i>ols</i>)	0.976	0.918	0.915	0.917	0.912	0.872	0.783	0.784	0.783	0.779
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 _h (<i>hbshr</i>)	0.975	0.915	0.912	0.915	0.909	0.872	0.775	0.772	0.772	0.770
oct _h (<i>bshr</i>)	0.954	0.895	0.895	0.892	0.887	0.833	0.741	0.741	0.737	0.735
oct _h (<i>hshr</i>)	1.093	0.955	0.951	0.956	0.949	1.066	0.851	0.846	0.848	0.838
oct _h (<i>shr</i>)	1.082	0.923	0.903	0.900	0.890	1.043	0.797	0.768	0.764	0.750

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

Table 10: ES ratio indices defined in (13) and (15) 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.

Base forecasts using a Gaussian approach are better in terms of both 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 predictive 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, bottom-up ct(*bu*) and partly bottom-up techniques like ct(*shr_{cs}, bu_{te}*) and ct(*wlsv_{te}, bu_{cs}*) lead to better results than the benchmark (bootstrap base forecasts). However, it's not always guaranteed that the improvement is higher than the starting base forecasts (by comparing the value of each column). This is particularly true for higher levels of temporal aggregation (see online appendix for details). Overall, oct(*bdshr*) in terms of CRPS is almost always the best. The shrinkage approach oct_h(*hshr*) has good performance in the bootstrap case: it is competitive with oct(*bdshr*) at lower temporal frequency ($k \in \{2, 1\}$)

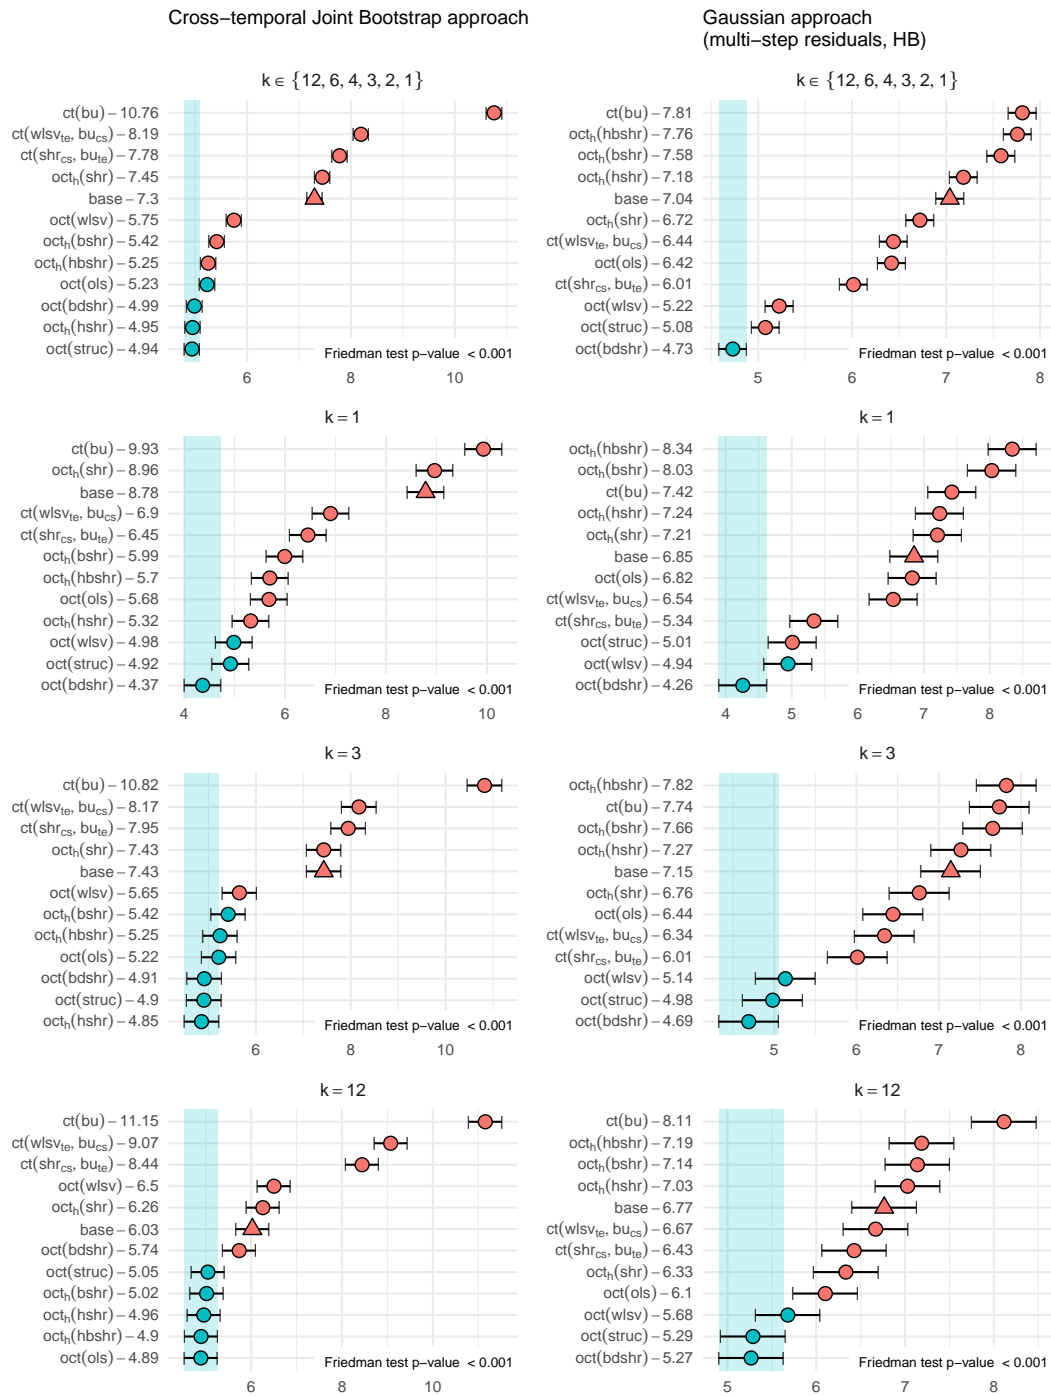


Figure 9: MCB Nemenyi test for the Australian Tourism Demand dataset using CRPS at different temporal aggregation orders for the Gaussian (multi-step residuals, HB) and the non-parametric bootstrap approaches. In each panel, the Friedman test p-value is reported in the lower right corner. The mean rank of each approach is shown to the right of its name. Statistical 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.

and it is able to improve for $k \geq 3$. In terms of ES, $\text{oct}(bdshr)$ is still competitive, although it does not always show the best relative performance. In this case, approaches that attempt to explicitly take into account temporal and cross-sectional relationships, such as $\text{oct}_h(hbshr)$ and $\text{oct}_h(bshr)$, look better. It is also worth noting that techniques that don't make use of residuals like $\text{oct}(ols)$ and $\text{oct}(struc)$ prove to be competitive by consistently improving the base forecasts in terms of both CRPS and ES.

Figure 9 shows the MCB using the CRPS for the Gaussian approach using multi-step residuals (HB) and the non-parametric bootstrap approach. In general, partly bottom-up procedure improves with respect to base forecasts at monthly level, but optimal cross-temporal procedures are always better. In the bootstrap framework, we can identify a group of three procedures, $\text{oct}(bdshr)$, $\text{oct}(hshr)$ and $\text{oct}(struc)$ that are almost always in the group of the best approaches (denoted by the blue dot). In the Gaussian framework, $\text{oct}(wlsv)$, $\text{oct}(struc)$, and $\text{oct}(bdshr)$ are always significantly better than base forecasts and equivalent in terms of results for temporal aggregation orders greater than 2. For monthly series, $\text{oct}(bdshr)$ is always significantly better than all other approaches.

8 Conclusion

In this paper, we extend the probabilistic reconciliation setting developed by Panagiotelis et al. (2023) for the cross-sectional case to the cross-temporal framework. Through appropriate notation, we show 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. (2022). We propose a Gaussian and a bootstrap approach to simulate the base forecasts able to take into account both cross-sectional and temporal relationships simultaneously, opening the way for further research into cross-temporal probabilistic forecasting.

Moreover, we analyze the use of residuals, showing that one-step residuals fail to capture the temporal structure, and we propose multi-step residuals that can fully capture the full cross-temporal relationships. When dealing with covariance matrices (due to the high-dimensionality of the cross-temporal setting), we propose 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. We present a simple simulation to investigate using the different types of residuals in the cross-temporal setting arriving to some useful conclusions. These ideas are worth requiring further investigation in future works.

Finally, we perform empirical applications on two Australian data sets commonly used in forecast reconciliation research: GDP from Income and Expenditure sides and Australian Tourism Demand. We find that in both cases optimal cross-temporal reconciliation approaches significantly improve on base forecasts. We also compare these with partly bottom-up techniques that use uni-dimensional reconciliations (either cross-sectional or temporal) and find confirmation that simultaneously exploiting both dimensions in reconciliation gives

better results, especially at higher levels of temporal aggregation. In conclusion, reconciliation approaches can play an important role to improve the accuracy of forecasts in a probabilistic framework while achieving the important attribute of producing coherent forecasts.

A Covariance approximations

Table A.1 presents some approximations for the cross-sectional and the temporal covariance matrices. Di Fonzo & Girolimetto (2023a) consider the following approximations for the cross-temporal covariance matrix.

- oct(*ols*) - identity: $\mathbf{\Omega}_{ct} = \mathbf{I}_{n(k^*+m)}$.
- oct(*struc*) - structural: $\mathbf{\Omega}_{ct} = \text{diag}(\mathbf{S}_{ct}\mathbf{1}_{mn_b})$.
- oct(*wlsv*) - series variance scaling: $\mathbf{\Omega}_{ct} = \widehat{\mathbf{\Omega}}_{ct,wlsv}$, a straightforward extension of the series variance scaling matrix presented by Athanasopoulos et al. (2017).
- oct(*bdshr*) - block-diagonal shrunk cross-covariance scaling: $\mathbf{\Omega}_{ct} = \mathbf{P}\widehat{\mathbf{W}}_{ct,shr}^{BD}\mathbf{P}'$, with $\widehat{\mathbf{W}}_{ct,shr}^{BD}$ a block diagonal matrix where each k -block ($k = m, k_{p-1}, \dots, 1$) is $\mathbf{I}_{M_k} \otimes \widehat{\mathbf{W}}_{shr}^{[k]}$, $\widehat{\mathbf{W}}_{shr}^{[k]}$ is the shrunk estimate of the cross-sectional covariance matrix proposed by Wickramasuriya et al. (2019), and \mathbf{P} is the commutation matrix such that $\text{Pvec}(\mathbf{Y}_\tau) = \text{vec}(\mathbf{Y}'_\tau)$.
- oct(*shr*) - MinT-shr: $\mathbf{\Omega}_{ct} = \hat{\lambda}\widehat{\mathbf{\Omega}}_{ct,D} + (1 - \hat{\lambda})\widehat{\mathbf{\Omega}}_{ct}$, where $\hat{\lambda}$ is an estimated shrinkage coefficient (Ledoit & Wolf 2004), $\widehat{\mathbf{\Omega}}_{ct,D} = \mathbf{I}_{n(k^*+m)} \odot \widehat{\mathbf{\Omega}}_{ct}$ with \odot denoting the Hadamard product, and $\widehat{\mathbf{\Omega}}_{ct}$ is the covariance matrix of the cross-temporal one-step ahead in-sample forecast errors.
- oct(*sam*) - MinT-sam: $\mathbf{\Omega}_{ct} = \widehat{\mathbf{\Omega}}_{ct}$.

	Cross-sectional framework	Temporal framework
identity	cs(<i>ols</i>): $\mathbf{W} = \mathbf{I}_n$	te(<i>ols</i>): $\mathbf{\Omega} = \mathbf{I}_{k^*+m}$
structural	cs(<i>struc</i>): $\mathbf{W} = \text{diag}(\mathbf{S}_{cs}\mathbf{1}_{nb})$	te(<i>struc</i>): $\mathbf{\Omega} = \text{diag}(\mathbf{S}_{te}\mathbf{1}_m)$
series variance	cs(<i>wls</i>): $\mathbf{W} = \widehat{\mathbf{W}}_D = \mathbf{I}_n \odot \widehat{\mathbf{W}}$	te(<i>wlsv</i>): $\mathbf{\Omega} = \widehat{\mathbf{\Omega}}_{wlsv}$
MinT-shr	cs(<i>shr</i>): $\mathbf{W} = \hat{\lambda}\widehat{\mathbf{W}}_D + (1 - \hat{\lambda})\widehat{\mathbf{W}}$	te(<i>shr</i>): $\mathbf{\Omega} = \hat{\lambda}\widehat{\mathbf{\Omega}}_D + (1 - \hat{\lambda})\widehat{\mathbf{\Omega}}$
MinT-sam	cs(<i>sam</i>): $\mathbf{W} = \widehat{\mathbf{W}}$	te(<i>sam</i>): $\mathbf{\Omega} = \widehat{\mathbf{\Omega}}$

Note: $\widehat{\mathbf{W}}$ ($\widehat{\mathbf{\Omega}}$) is the covariance matrix of the cross-sectional (temporal) one-step ahead in-sample forecast errors, $\widehat{\mathbf{\Omega}}_{wlsv}$ is a diagonal matrix presented by Athanasopoulos et al. (2017), and $\widehat{\mathbf{\Omega}}_D = \mathbf{I}_{k^*+m} \odot \widehat{\mathbf{\Omega}}$, where \odot denotes the Hadamard product.

Table A.1: Approximations for cross-sectional (\mathbf{W}) and temporal ($\mathbf{\Omega}$) covariance matrices.

Acknowledgments

Tommaso Di Fonzo and Daniele Girolimetto acknowledge financial support from project PRIN2017 “HiDEA: Advanced Econometrics for High-frequency Data”, 2017RSMPPZZ. 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.

References

- Athanasopoulos, G., Ahmed, R. A. & Hyndman, R. J. (2009), 'Hierarchical forecasts for Australian domestic tourism', *International Journal of Forecasting* **25**(1), 146–166.
- 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, Springer International Publishing, Cham, pp. 689–719.
- Athanasopoulos, G., Hyndman, R. J., Kourentzes, N. & Petropoulos, F. (2017), 'Forecasting with temporal hierarchies', *European Journal of Operational Research* **262**(1), 60–74.
- 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', PMLR, pp. 3348–3357.
- 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.
- Bisaglia, L., Fonzo, T. D. & Girolimetto, D. (2020), Fully reconciled GDP forecasts from Income and Expenditure sides, in A. Pollice, N. Salvati & F. Schirripa Spagnolo, eds, 'Book of Short Papers SIS 2020', Pearson, pp. 951–956.
- 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.
- Corani, G., Rubattu, N., Azzimonti, D. & Antonucci, A. (2022), 'Probabilistic Reconciliation of Count Time Series', *arXiv:2207.09322* .
- Dangerfield, B. J. & Morris, J. S. (1992), 'Top-down or bottom-up: Aggregate versus disaggregate extrapolations', *International Journal of Forecasting* **8**(2), 233–241.
- Di Fonzo, T. & Girolimetto, D. (2022a), 'Enhancements in cross-temporal forecast reconciliation, with an application to solar irradiance forecasts', *arXiv:2209.07146* .
- Di Fonzo, T. & Girolimetto, D. (2022b), 'Forecast combination-based forecast reconciliation: Insights and extensions', *International Journal of Forecasting* . in press.
- Di Fonzo, T. & Girolimetto, D. (2022c), Fully reconciled probabilistic GDP forecasts from Income and Expenditure sides, in A. Balzanella, M. Bini, C. Cavicchia & R. Verde, eds, 'Book of Short Papers SIS 2022', Pearson, pp. 1376–1381.
- Di Fonzo, T. & Girolimetto, D. (2022d), 'Point and probabilistic forecast reconciliation for general linearly constrained multiple time series', *Working paper* .
- Di Fonzo, T. & Girolimetto, D. (2023a), 'Cross-temporal forecast reconciliation: Optimal combination method and heuristic alternatives', *International Journal of Forecasting* **39**(1), 39–57.
- Di Fonzo, T. & Girolimetto, D. (2023b), 'Spatio-temporal reconciliation of solar forecasts', *Solar Energy* **251**, 13–29.

- 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.
- Eckert, F., Hyndman, R. J. & Panagiotelis, A. (2021), 'Forecasting Swiss exports using Bayesian forecast reconciliation', *European Journal of Operational Research* **291**(2), 693–710.
- Efron, B. (1975), 'Biased versus unbiased estimation', *Advances in Mathematics* **16**(3), 259–277.
- Efron, B. & Morris, C. (1975), 'Data Analysis Using Stein's Estimator and its Generalizations', *Journal of the American Statistical Association* **70**(350), 311–319.
- 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.
- Fliedner, G. (2001), 'Hierarchical forecasting: Issues and use guidelines', *Industrial Management & Data Systems* **101**(1), 5–12.
- Girolimetto, D. & Di Fonzo, T. (2023), *FoReco: Point Forecast Reconciliation*. R package v0.2.6.
URL: <https://danigiro.github.io/FoReco/>
- Gneiting, T. & Katzfuss, M. (2014), 'Probabilistic Forecasting', *Annual Review of Statistics and Its Application* **1**(1), 125–151.
- Gross, C. W. & Sohl, J. E. (1990), 'Disaggregation methods to expedite product line forecasting', *Journal of Forecasting* **9**(3), 233–254.
- 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.
- Hyndman, R. J., Athanasopoulos, G., Bergmeir, C., Caceres, G., Chhay, L., Kuroptev, K., O'Hara-Wild, M., Petropoulos, F., Razbash, S., Wang, E., Yasmeen, 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.
URL: <https://pkg.robjhyndman.com/forecast/>
- Hyndman, R. J. & Khandakar, Y. (2008), 'Automatic Time Series Forecasting: The forecast Package for R', *Journal of Statistical Software* **27**, 1–22.
- 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.
- 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.
- Kourentzes, N. & Athanasopoulos, G. (2019), 'Cross-temporal coherent forecasts for Australian tourism', *Annals of Tourism Research* **75**, 393–409.
- Ledoit, O. & Wolf, M. (2004), 'A well-conditioned estimator for large-dimensional covariance matrices', *Journal of Multivariate Analysis* **88**(2), 365–411.
- Magnus, J. R. & Neudecker, H. (2019), *Matrix Differential Calculus with Applications in Statistics and Econometrics*, 3rd edn, Wiley, Hoboken, NJ.

- Makridakis, S., Spiliotis, E. & Assimakopoulos, V. (2022), 'M5 accuracy competition: Results, findings, and conclusions', *International Journal of Forecasting* **38**(4), 1346–1364.
- 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.
- O'Hara-Wild, M., Hyndman, R. J., Wang, E. & Godahewa, R. (2022), *tsibbledata: Diverse Datasets for 'tsibble'*. R package v0.4.1.
URL: tsibbledata.tidyverts.org
- 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.
- 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.
- Panamtash, H. & Zhou, Q. (2018), Coherent Probabilistic Solar Power Forecasting, in '2018 IEEE International Conference on Probabilistic Methods Applied to Power Systems (PMAPS)', Boise, ID, USA, 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.
URL: <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.
- 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**, 106796.
- R Core Team (2022), *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria.
URL: www.R-project.org
- 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**, 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).
- Spiliotis, E., Petropoulos, F., Kourentzes, N. & Assimakopoulos, V. (2020), 'Cross-temporal aggregation: Improving the forecast accuracy of hierarchical electricity consumption', *Applied Energy* **261**, 114339.
- Venables, W. N. & Ripley, B. D. (2002), *Modern Applied Statistics with S*, fourth edn, Springer, New York. ISBN 0-387-95457-0.
URL: <https://www.stats.ox.ac.uk/pub/MASS4/>

- Wickramasuriya, S. L. (2023), 'Probabilistic forecast reconciliation under the gaussian framework', *Journal of Business & Economic Statistics*, *in press* .
- 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.
- Wickramasuriya, S. L., Turlach, B. A. & Hyndman, R. J. (2020), 'Optimal non-negative forecast reconciliation', *Statistics and Computing* **30**(5), 1167–1182.
- Yagli, G. M., Yang, D. & Srinivasan, D. (2019), 'Reconciling solar forecasts: Sequential reconciliation', *Solar Energy* **179**, 391–397.
- Zambon, L., Azzimonti, D. & Corani, G. (2022), 'Efficient probabilistic reconciliation of forecasts for real-valued and count time series', *arXiv:2210.02286* .