



Point and probabilistic forecast reconciliation for general linearly constrained multiple time series

Daniele Girolimetto¹ · Tommaso Di Fonzo¹

Accepted: 18 November 2023 / Published online: 21 December 2023
© The Author(s) 2023

Abstract

Forecast reconciliation is the post-forecasting process aimed to revise a set of incoherent base forecasts into coherent forecasts in line with given data structures. Most of the point and probabilistic regression-based forecast reconciliation results ground on the so called “structural representation” and on the related unconstrained generalized least squares reconciliation formula. However, the structural representation naturally applies to genuine hierarchical/grouped time series, where the top- and bottom-level variables are uniquely identified. When a general linearly constrained multiple time series is considered, the forecast reconciliation is naturally expressed according to a projection approach. While it is well known that the classic structural reconciliation formula is equivalent to its projection approach counterpart, so far it is not completely understood if and how a structural-like reconciliation formula may be derived for a general linearly constrained multiple time series. Such an expression would permit to extend reconciliation definitions, theorems and results in a straightforward manner. In this paper, we show that for general linearly constrained multiple time series it is possible to express the reconciliation formula according to a “structural-like” approach that keeps distinct free and constrained, instead of bottom and upper (aggregated), variables, establish the probabilistic forecast reconciliation framework, and apply these findings to obtain fully reconciled point and probabilistic forecasts for the aggregates of the Australian *GDP* from income and expenditure sides, and for the European Area *GDP* disaggregated by income, expenditure and output sides and by 19 countries.

Keywords Linearly constrained multiple time series · Hierarchical/grouped time series · Point and probabilistic forecast reconciliation · Quarterly national accounts · *GDP*

✉ Daniele Girolimetto
daniele.girolimetto@unipd.it

Tommaso Di Fonzo
difonzo@stat.unipd.it

¹ Department of Statistical Sciences, University of Padova, Via Cesare Battisti 241, 35121 Padova, Italy

1 Introduction

Starting from Hyndman et al. (2011), regression-based forecast reconciliation has become an hot topic in the forecasting literature (van Erven and Cugliari 2015; Wickramasuriya et al. 2019; Wickramasuriya 2021; Panagiotelis et al. 2021; Jeon et al. 2019; Ben Taieb and Koo 2019; Ben Taieb et al. 2021; Panagiotelis et al. 2023; Wickramasuriya 2023). By forecast reconciliation, we mean a post-forecasting procedure (Di Fonzo and Girolimetto 2023) in which previously (and however) generated incoherent predictions (called base forecasts) for all the components of a multiple time series are adjusted in order to fulfill some externally given linear constraints. These reconciled forecasts are said to be coherent. In many real world applications, forecasts of a large collection of time series have a natural organization according to a hierarchical structure. More precisely, a system is classified as hierarchical when series are created by aggregating others in a tree shape. When the system is formed by a unique tree, then the collection is called “hierarchical time series” (Hyndman et al. 2011). On the other side, when various hierarchies share the same series at both the most aggregated and disaggregated levels (top and bottom level, respectively), we face a “grouped time series” (Hyndman et al. 2016).

In the field of forecast reconciliation, the bottom-up and top-down approaches are among the earliest and best known ones. Bottom-up forecasting (Dunn et al. 1976) uses only forecasts at the bottom level to obtain all the reconciled forecasts. In contrast, top-down forecasting (Gross and Sohl 1990) uses only the forecasts at the highest aggregated level. Having observed that neither forecasting approach uses all the available information, Hyndman et al. (2011) developed a regression-based reconciliation approach consisting in (i) forecasting all the series with no regard for the constraints, and (ii) using then a regression model to optimally combining these (base) forecasts in order to produce coherent forecasts. This approach has witnessed a continuous growth of the related literature (Hyndman et al. 2016; Athanasopoulos et al. 2020; Wickramasuriya 2021; Panagiotelis et al. 2023), that in most cases grounds on the so-called structural representation of a hierarchical/grouped time series (Athanasopoulos et al. 2009), in which the variables are classified either bottom if they belong to the most disaggregated level, or upper if they are obtained by summing the lower levels’ variables. In mathematical terms, upper and bottom variables are linked by an aggregation matrix, which describes how the upper series are obtained from the bottom ones (Hyndman et al. 2011). This representation is directly related to the hierarchical structure, where the series are naturally classifiable and an aggregation matrix may be obtained with little effort. Theoretical aspects for point and probabilistic forecast reconciliation have been developed using a structural representation by Panagiotelis et al. (2021) and Panagiotelis et al. (2023), respectively (see also Wickramasuriya 2021, 2023).

However, it can be shown (van Erven and Cugliari 2015; Wickramasuriya et al. 2019; Bisaglia et al. 2020; Di Fonzo and Girolimetto 2023) that reconciled forecasts may be obtained as the solution to a linearly constrained quadratic optimization

problem,¹ that does not require any “upper and bottom” classification of the involved variables. This approach grounds on a zero-constrained representation of the linearly constrained multiple time series (Di Fonzo and Girolimetto 2023) describing the relationships linking all the individual series in the system. For a genuine hierarchical/grouped time series, where the top and bottom level variables are uniquely identified, it is easy to express the relationship between structural and zero-constrained representations. Wickramasuriya et al. (2019) show that the structural and the corresponding projection reconciliation approaches produce the same final reconciled forecasts. Unfortunately, given a zero-constrained representation of a linearly constrained multiple time series, finding the corresponding structural representation is not trivial (Di Fonzo and Girolimetto 2023), and one of the objective of the present paper is to fill this gap.

Most of the forecast reconciliation approaches proposed in the literature refer to genuine hierarchical/grouped time series, that do not take into account the full spectrum of possible cases encountered in real-life situations. As pointed out by Panagiotelis et al. (2021), “*concepts such as coherence and reconciliation (...) require the data to have only two important characteristics: the first is that they are multivariate, and the second is that they adhere to linear constraints*”. Using a novel geometric interpretation, Panagiotelis et al. (2021) develop definitions and a formulation for linearly constrained multiple time series within a general framework and not just for simple summation hierarchical relationships. However, their results still ground on the structural representation valid only for genuine hierarchical/grouped time series, and this also holds for the probabilistic forecast reconciliation approach developed by Panagiotelis et al. (2023). Nevertheless, the point we wish to stress here is that when working with general linear constraints and many variables, the interchangeability between structural and zero-constrained representations, easy to recover for a genuine hierarchical/grouped time series, is no more always straightforward.

In this paper we address a number of open issues in point and probabilistic cross-sectional forecast reconciliation for general linearly constrained multiple time series. First, we introduce a general linearly constrained multiple time series by exploiting its analogies with an homogeneous linear system. Second, we show that the classical hierarchical representation for a multiple time series is a simple special case of the general representation. Third, we show that if it is possible to express a general linearly constrained multiple time series according to a “structural-like” representation, we can easily achieve the formulation for point and probabilistic regression-based reconciled forecasts using a linear combination matrix, with elements in \mathbb{R} , that is the natural extension of the aggregation matrix used in the structural reconciliation approach, with elements only in $\{0, 1\}$. When the distinction between bottom and upper variables is no longer meaningful, we adopt a classification involving free and constrained variables, respectively, and show how to obtain a structural-like representation, possibly using well known linear algebra

¹ This approach dates back to the seminal paper by Stone et al. (1942) on the least squares adjustment of noisy data with accounting constraints (Byron 1978, 1979). Recent applications to the reconciliation of systems of seasonally adjusted time series are provided by Di Fonzo and Marini (2011, 2015).

techniques, such as the Reduced Row Echelon Form and the QR decomposition (Golub and Van Loan 1996; Meyer 2000).

The remainder of the paper is structured as follows. In Sect. 2, we present the notation and the main results about the point forecast reconciliation for a genuine hierarchical/grouped time series and in the general case. We define the structural-like representation for a general linearly constrained multiple time series by distinguishing between free and constrained variables, instead of bottom and upper, and use this result in the probabilistic forecast reconciliation framework set out by Panagiotelis et al. (2023). In Sect. 3, we show how to obtain the linear combination matrix for the structural-like representation in a computationally efficient way². Two empirical applications are presented in Sect. 4. First, we extend the forecast reconciliation experiment for the Australian *GDP* originally developed by Athanasopoulos et al. (2020) and Bisaglia et al. (2020), to obtain both point and probabilistic *GDP* forecasts simultaneously coherent with their disaggregate counterpart forecasts from income and expenditure sides. Second, point and probabilistic forecasts are obtained for the European Area *GDP* from income, expenditure and output sides, geographically disaggregated by 19 component countries, where the large number of series and constraints makes a full row rank zero-constraints matrix difficult to build. Section 5 contains conclusions.

2 Cross-sectional forecast reconciliation of a linearly constrained multiple time series

Let \mathbf{x}_t be a n -dimensional linearly constrained multiple time series such that all the values for $t = 1, \dots$ (either observed or not) lie in the *coherent* linear subspace $\mathcal{S} \in \mathbb{R}^n$, $\mathbf{x}_t \in \mathcal{S} \forall t > 0$ (Panagiotelis et al. 2021), which means that all linear constraints are satisfied at time t . These constraints can be represented through linear equations and grouped as a (rectangular) linear system. Following Meyer (2000) and Leon (2015), let $x_{1,t}, \dots, x_{n,t}$ be the observations of n individual time series at a given time $t = 1, \dots, T$. An homogeneous linear system of p equations in the n variables present in \mathbf{x}_t can be written as

$$\begin{cases} \gamma_{1,1} x_{1,t} + \gamma_{1,2} x_{2,t} + \dots + \gamma_{1,n} x_{n,t} = 0 \\ \vdots \quad \quad \quad \vdots \quad \quad \quad \vdots \quad \quad \quad \vdots \quad \quad \quad \vdots \quad \quad \quad \vdots \\ \gamma_{p,1} x_{1,t} + \gamma_{p,2} x_{2,t} + \dots + \gamma_{p,n} x_{n,t} = 0 \end{cases},$$

where the γ_{ij} 's are real-valued coefficients. This system can be expressed in matrix form as

$$\mathbf{\Gamma} \mathbf{x}_t = \mathbf{0}_{(p \times 1)}, \quad (1)$$

² The procedures used in this paper are implemented in the R package `FOReco` (Girolimetto and Di Fonzo 2023). A complete set of results is available at the GitHub repository: <https://github.com/danigiro/mtsreco>.

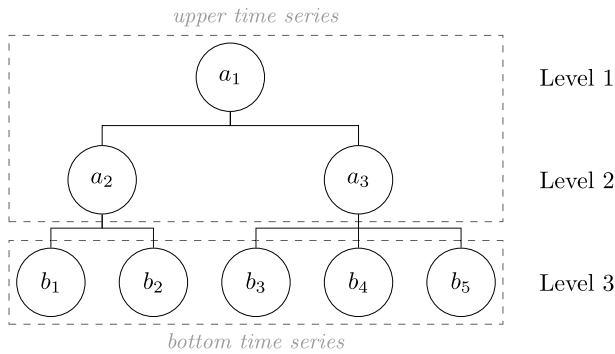


Fig. 1 A simple three-level hierarchical structure for a linearly constrained multiple time series

where $\Gamma \in \mathbb{R}^{(p \times n)}$ is the coefficient matrix

$$\Gamma = \begin{bmatrix} \gamma_{1,1} & \gamma_{1,2} & \cdots & \gamma_{1,n} \\ \vdots & \vdots & \ddots & \vdots \\ \gamma_{p,1} & \gamma_{p,2} & \cdots & \gamma_{p,n} \end{bmatrix}.$$

Expression (1) is called “zero-constrained representation” of a linearly constrained multiple time series (Di Fonzo and Girolimetto 2023).

2.1 Forecast reconciliation of a genuine hierarchical time series

In Fig. 1 it is shown a simple, genuine hierarchical time series (Athanasopoulos et al. 2009; Hyndman et al. 2011; Athanasopoulos et al. 2020; Hyndman and Athanasopoulos 2021), that can be seen as a particular case of a linearly constrained multiple time series. This hierarchical structure is defined only by simple summation constraints,

$$\begin{aligned} a_1 &= b_1 + b_2 + b_3 + b_4 + b_5 \\ a_2 &= b_1 + b_2 \\ a_3 &= b_3 + b_4 + b_5 \end{aligned}, \tag{2}$$

that can be easily transformed into a zero-constrained representation $\Gamma \mathbf{x} = \mathbf{0}_{(8 \times 1)}$, with $\mathbf{x} = [a_1 \ a_2 \ a_3 \ b_1 \ b_2 \ b_3 \ b_4 \ b_5]'$ and $\mathbf{a}' \ \mathbf{b}'$,

$$\Gamma = \begin{bmatrix} 1 & 0 & 0 & -1 & -1 & -1 & -1 & -1 \\ 0 & 1 & 0 & -1 & -1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & -1 & -1 & -1 \end{bmatrix} = [\mathbf{I}_3 \quad -\mathbf{A}],$$

such as the $\gamma_{i,j}$'s coefficients are only -1, 0 and 1, and $\mathbf{a} = \mathbf{A}\mathbf{b}$.

In general, let $\mathbf{b}_t = [b_{1,t} \ \dots \ b_{n_b,t}]' \in \mathbb{R}^{(n_b \times 1)}$ and $\mathbf{a}_t = [a_{1,t} \ \dots \ a_{n_a,t}]' \in \mathbb{R}^{(n_a \times 1)}$, $t = 1, \dots, T$, with $n = n_a + n_b$, be the T vectors containing the bottom and the aggregated series, respectively, of a hierarchy. All series can be collected in the T vectors

$$y_t = Px_t = \begin{bmatrix} a_t \\ b_t \end{bmatrix} \in \mathbb{R}^{(n \times 1)},$$

where $P \in \{0, 1\}^{(n \times n)}$ is a permutation matrix used to appropriately re-order the original vector x_t . If the classification as upper or bottom of the single time series in x_t is known in advance, we assume $x_t = y_t = [a'_t \ b'_t]'$, i.e. $P = I_n$ (no permutation of the original vector is needed). Moreover, also the linear combination (aggregation) matrix $A \in \{0, 1\}^{(n_a \times n_b)}$ describing the summation constraints linking the upper time series to the bottom ones, $a_t = Ab_t$, is assumed known in advance. Thus the “structural representation” is simply given by (Athanasopoulos et al. 2009)

$$y_t = Sb_t \quad \text{with} \quad S = \begin{bmatrix} A \\ I_{n_b} \end{bmatrix}, \tag{3}$$

where $S \in \{0, 1\}^{(n \times n_a)}$ is the structural (summation) matrix.

Suppose now we have the vector $\hat{y}_h = [\hat{a}'_h \ \hat{b}'_h]'$ $\in \mathbb{R}^{(n \times 1)}$ of unbiased and incoherent (i.e., $\hat{y}_h \neq S\hat{b}_h$) base forecasts for the n variables of the linearly constrained series y_t for the forecast horizon h . Hyndman et al. (2011) use the structural representation (3) to obtain the reconciled forecasts \tilde{y}_h as

$$\tilde{y}_h = SG\hat{y}_h, \quad G = (S'W^{-1}S)^{-1}S'W^{-1}, \tag{4}$$

where W is a $(n \times n)$ p.d. matrix assumed known and \hat{y}_h (\tilde{y}_h) is the vector containing the base (reconciled) forecasts at forecast horizon h . Some alternative matrices W have been proposed in the literature for the cross-sectional forecast reconciliation case³:

- identity (ols): $\hat{W}_{ols} = I_n$ (Hyndman et al. 2011),
- series variance (wls): $\hat{W}_{wls} = I_n \odot \hat{W}_{sam}$ (Hyndman et al. 2016),
- MinT-shr (shr): $\hat{W}_{shr} = \hat{\lambda}\hat{W}_{wls} + (1-\hat{\lambda})\hat{W}_{sam}$ (Wickramasuriya et al. 2019),
- MinT-sam (sam): $\hat{W}_{sam} = \frac{1}{T} \sum_{t=1}^T \hat{e}_t\hat{e}'_t$ is the covariance matrix of the one-step ahead in-sample forecast errors \hat{e}_t (Wickramasuriya et al. 2019),

where the symbol \odot denotes the Hadamard product, and $\hat{\lambda}$ is an estimated shrinkage coefficient (Ledoit and Wolf 2004).

The structural representation (3) may be transformed into the equivalent zero-constrained representation $a_t - Ab_t = \mathbf{0}_{(n_a \times 1)}$, that is (Wickramasuriya et al. 2019)

$$Cy_t = \mathbf{0}_{(n_a \times 1)} \quad \text{with} \quad C = [I_{n_a} \ -A]. \tag{5}$$

³ Dealing with the uncertainty in base forecasts, and their error covariance matrices, is of primary interest to establish characteristics and properties of the reconciled forecasts. This topic is left to future research.

$C \in \{-1, 0, 1\}^{(n_a \times n)}$ is a full row-rank zero constraints matrix used to obtain the point reconciled forecasts according to the *projection approach* (van Erven and Cugliari 2015; Wickramasuriya et al. 2019; Di Fonzo and Girolimetto 2023):

$$\tilde{y}_h = M\hat{y}_h, \quad M = I_n - WC'(CWC')^{-1}C. \tag{6}$$

Structural and zero-constrained representations can be interchangeable depending on which of the two is more convenient to use, allowing for greater flexibility in the calculation of the reconciled forecasts. For, the zero-constrained representation appears to be less computational intensive: in equation (4) two matrices must be inverted, one of size $(n \times n)$ and the other $(n_b \times n_b)$, whereas only the inversion of an $(n_a \times n_a)$ matrix is required in formula (6).

2.2 The general case: zero-constrained and structural-like representations

In the hierarchical/grouped cross-sectional forecast reconciliation there is a natural distinction between upper and bottom time series that leads to the construction of the matrix C as in (5), where the first n_a columns refer to the upper and the remaining $n_b (= n - n_a)$ to the bottom variables, respectively. The time series in these two sets are categorized logically: all those related to the last level belong to the second group, the rest to the first. Most of the forecast reconciliation approaches proposed in the literature refer to genuine hierarchical/grouped time series and its structural representation. However, these do not take into account the full spectrum of possible cases encountered in real life situations. For a general linearly constrained multiple time series $(x_t, t = 1, \dots, T)$, the classification between upper and bottom variables might not be meaningful, prompting us rather to look for two new sets: the **constrained** variables, denoted as $c_t \in \mathbb{R}^{(n_c \times 1)}$, and the **free** (unconstrained) variables, denoted as $u_t \in \mathbb{R}^{(n_u \times 1)}$, with $n = n_c + n_u$, such that $y_t = Px_t = [c'_t \ u'_t]'$, and

$$c_t = Au_t. \tag{7}$$

In this general framework, $A \in \mathbb{R}^{(n_c \times n_u)}$ is the *linear combination* matrix associated to the linearly constrained multiple time series $x_t = P'y_t$, that can be thus expressed via the “structural-like representation”

$$y_t = Px_t = \begin{bmatrix} c_t \\ u_t \end{bmatrix} = Su_t \quad \text{with} \quad S = \begin{bmatrix} A \\ I_{n_u} \end{bmatrix}, \tag{8}$$

where $S \in \mathbb{R}^{(n_c \times n_u)}$ is the structural-like matrix.⁴ It is worth noting that a full-rank zero-constraints matrix $C \in \mathbb{R}^{(n_c \times n)}$ like in expression (5) may be easily obtained by expression (8) and used for the full-rank zero constrained representation $Cy_t = CPx_t = 0_{(n_c \times 1)}$. Therefore, even for a general, possibly not genuine

⁴ Unlike the structural (summation) matrix in (3), that describes the simple summation relationships valid for genuine hierarchical/grouped time series, and has only elements in $\{0, 1\}$, in expression (8) S consists of real coefficients, appropriately organized to highlight the links between constrained and free variables.




Fig. 2 A general linearly constrained structure: two hierarchies sharing only the same top-level series

hierarchical/grouped, linearly constrained multiple time series, either the structural (4) or the projection (6) approaches may be used to perform the point forecast reconciliation of incoherent base forecasts.

A simple example of a linearly constrained multiple time series that cannot be expressed as a genuine hierarchical/grouped structure is shown in Fig. 2. In this case the variable X is at the top of two distinct hierarchies, that do not share the same bottom-level variables. The aggregation relationships between the upper variables X and A, and the bottom ones A1, A2, B, C, and D are given by:

$$\begin{aligned}
 X &= A1 + A2 + B \\
 X &= C + D \\
 A &= A1 + A2
 \end{aligned} \tag{9}$$

In this case, both the zero-constrained and the structural-like representations are found in a rather straightforward manner. We consider $\{A2, B, C, D\}$ as free variables, such that $y_t = x_t$ (i.e., $P = I_4$), with $c_t = [x_{X,t} \ x_{A,t} \ x_{A1,t}]'$ and $u_t = [x_{A2,t} \ x_{B,t} \ x_{C,t} \ x_{D,t}]'$. Thus, the coefficient matrix of the zero-constrained representation (1) is

$$\Gamma = \begin{bmatrix} 1 & 0 & -1 & -1 & -1 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & -1 & -1 \\ 0 & 1 & -1 & -1 & 0 & 0 & 0 \end{bmatrix}.$$

It is immediate to check that the system of linear constraints (9) may be re-written as

$$\begin{aligned}
 X &= C + D \\
 A &= -B + C + D \\
 A1 &= -A2 - B + C + D
 \end{aligned}$$

that is $Cx_t = \mathbf{0}_{(3 \times 1)}$, where

$$C = [I_3 \ -A] = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & -1 & -1 \\ 0 & 1 & 0 & 0 & 1 & -1 & -1 \\ 0 & 0 & 1 & 1 & 1 & -1 & -1 \end{bmatrix} \text{ and } A = \begin{bmatrix} 0 & 0 & 1 & 1 \\ 0 & -1 & 1 & 1 \\ -1 & -1 & 1 & 1 \end{bmatrix}.$$

The structural-like representation of the general linearly constrained multiple time series in Fig. 2 is then $\mathbf{y}_t = \mathbf{S}\mathbf{u}_t$, with $\mathbf{S} = [\mathbf{A}' \quad \mathbf{I}_3]'$.

It should be mentioned, however, that for medium/large systems (with many constraints and/or variables), manually operating on the constraints could be time-consuming and challenging. In Sect. 3 we will show a general technique to derive the linear combination matrix \mathbf{A} from a general zero-constraints matrix $\mathbf{\Gamma}$.

2.3 Probabilistic forecast reconciliation for general linearly constrained multiple time series

So far we have dealt with only point forecasting, but if one wishes to account for forecast uncertainty, probabilistic forecasts should be considered, as they - in the form of probability distributions over future quantities or events - measure the uncertainty in forecasts and are an important component of optimal decision making (Gneiting and Katzfuss 2014).

Representation (8) states that \mathbf{y}_t lies in an n -dimensional subspace of \mathbb{R}^n spanned by the columns of \mathbf{S} , called “coherent subspace” and denoted by \mathcal{S} (Panagiotelis et al. 2023). Now, let $\mathcal{F}_{\mathbb{R}^{n_u}}$ be the Borel σ -algebra on \mathbb{R}^{n_u} , $(\mathbb{R}^{n_u}, \mathcal{F}_{\mathbb{R}^{n_u}}, \nu)$ a probability space for the free variables, and $s : \mathbb{R}^{n_u} \rightarrow \mathbb{R}^n$ a continuous mapping matrix. Then a σ -algebra $\mathcal{F}_{\mathcal{S}}$ can be constructed from the collection of sets $s(\mathcal{B})$ for all $\mathcal{B} \in \mathcal{F}_{\mathbb{R}^{n_u}}$.

Definition 1 (*Coherent probabilistic forecast for a linearly constrained multiple time series*) Given the triple $(\mathbb{R}^{n_u}, \mathcal{F}_{\mathbb{R}^{n_u}}, \nu)$, we define a coherent probability triple $(\mathcal{S}, \mathcal{F}_{\mathcal{S}}, \check{\nu})$ such that $\check{\nu}(s(\mathcal{B})) = \nu(\mathcal{B})$, $\forall \mathcal{B} \in \mathcal{F}_{\mathbb{R}^{n_u}}$.

In order to extend forecast reconciliation to the probabilistic setting, let $(\mathbb{R}^n, \mathcal{F}_{\mathbb{R}^n}, \hat{\nu})$ be a probability triple characterizing base (incoherent) probabilistic forecasts for all n series, and let $\psi : \mathbb{R}^{n_u} \rightarrow \mathbb{R}^n$ be a continuous mapping function defined by Panagiotelis et al. (2023) as the composition of two transformations, $\text{so} \circ g$, where $g : \mathbb{R}^n \rightarrow \mathbb{R}^{n_u}$ is a continuous function corresponding to matrix \mathbf{G} in equation (4).

Definition 2 (*Probabilistic forecast reconciliation for a linearly constrained multiple time series*) The reconciled probability measure of $\hat{\nu}$ with respect to ψ is a probability measure $\check{\nu}$ on \mathcal{S} with σ -algebra $\mathcal{F}_{\mathcal{S}}$ such that

$$\check{\nu}(\mathcal{A}) = \hat{\nu}(\psi^{-1}(\mathcal{A})), \quad \forall \mathcal{A} \in \mathcal{F}_{\mathcal{S}}, \quad (10)$$

where $\psi^{-1}(\mathcal{A}) = \{\mathbf{x} \in \mathbb{R}^{n_u} : \psi(\mathbf{x}) \in \mathcal{A}\}$ is the pre-image of \mathcal{A} .

We consider two alternative approaches to deal with probabilistic forecast reconciliation according to the above definitions: a parametric framework, where probabilistic forecasts are produced under the assumption that the density function of the forecast errors is known, and a non-parametric framework, where no distributional assumption is made.

2.3.1 Parametric framework: Gaussian reconciliation

A reconciled probabilistic forecast may be obtained analytically for some parametric distributions, such as the multivariate normal (Yagli et al. 2020; Corani et al. 2021; Eckert et al. 2021; Panagiotelis et al. 2023; Wickramasuriya 2023). In particular, if the base forecasts distribution is $\mathcal{N}(\hat{\mathbf{y}}_h, \mathbf{W}_h)$, then the reconciled forecasts distribution is $\mathcal{N}(\tilde{\mathbf{y}}_h, \tilde{\mathbf{W}}_h)$, with

$$\tilde{\mathbf{y}}_h = \mathbf{S} \mathbf{G} \hat{\mathbf{y}}_h \quad \text{and} \quad \tilde{\mathbf{W}}_h = \mathbf{S} \mathbf{G} \mathbf{W}_h \mathbf{G}' \mathbf{S}'.$$

The covariance matrix $\tilde{\mathbf{W}}_h$ deserves special attention. In the simple case assumed by Wickramasuriya et al. (2019), it is $\mathbf{W}_h = k_h \mathbf{W}$, where k_h is a proportionality constant, and the reconciled covariance matrix reduces to (see Appendix A):

$$\tilde{\mathbf{W}}_h = k_h \mathbf{S} \mathbf{G} \mathbf{W}. \quad (11)$$

However, the proportionality assumption along the forecast horizons h may be too restrictive, and computing k_h cannot be an easy task. Thus, three alternative formulations of \mathbf{W}_h , already shown in Sect. 2.1, have been proposed in the forecast reconciliation literature:

- diagonal covariance matrix: $\mathbf{W}_h = \hat{\mathbf{W}}_{\text{wls}}$ (Corani et al. 2021; Panagiotelis et al. 2023);
- sample covariance matrix: $\mathbf{W}_h = \hat{\mathbf{W}}_{\text{sam}}$ (Panagiotelis et al. 2023);
- shrinkage covariance matrix: $\mathbf{W}_h = \hat{\mathbf{W}}_{\text{shr}}$ (Athanasopoulos et al. 2020).

2.3.2 Non-parametric framework: joint bootstrap-based reconciliation

When an analytical expression of the forecast distribution is either unavailable, or relies on unrealistic parametric assumptions, the empirical evaluation of the results may be grounded on reconciled samples (Jeon et al. 2019; Yang 2020; Panagiotelis et al. 2023). At this end, we extend theorem 4.5 in Panagiotelis et al. (2023), originally formulated for genuine hierarchical/grouped time series, to the case of a general linearly constrained multiple time series. This theorem states that, if $(\hat{\mathbf{y}}^{[1]}, \dots, \hat{\mathbf{y}}^{[L]})$ is a sample drawn from an incoherent probability measure $\hat{\nu}$, then $(\tilde{\mathbf{y}}^{[1]}, \dots, \tilde{\mathbf{y}}^{[L]})$, where $\tilde{\mathbf{y}}^{[\ell]} := \psi(\hat{\mathbf{y}}^{[\ell]})$ for $\ell = 1, \dots, L$, is a sample drawn from the reconciled probability measure $\tilde{\nu}$ as defined in (10). According to this result, reconciling each member of a sample obtained from an incoherent distribution yields a sample from the reconciled distribution. As a consequence, coherent probabilistic forecasts may be developed through a post-forecasting mechanism analogous to the point forecast reconciliation setting. For this purpose, the bootstrap procedure by Athanasopoulos et al. (2020) is applied:

1. appropriate univariate models M_i for each series in the system are fitted based on the training data $\{y_{i,t}\}_{t=1}^T, i = 1, \dots, n$, and the one-step-ahead in-sample forecast errors are stacked in an $(n \times T)$ matrix, $\hat{E} = \{\hat{e}_{i,t}\}$;
2. $\hat{y}_{i,h}^{[l]} = f_i(M_i, \hat{e}_{i,h}^{[l]})$ is computed for $h = 1, \dots, H$ and $l = 1, \dots, L$, where $f(\cdot)$ is a function of the fitted model and its associated error, $\hat{y}_{i,h}^{[l]}$ is a sample path simulated for the i -th series, and $\hat{e}_{i,h}^{[l]}$ is the (i, h) -th element of an $(n \times H)$ block bootstrap matrix containing H consecutive columns randomly drawn from \hat{E} ;
3. the optimal reconciliation formula, either according to the structural approach (4) or the projection approach (6), is applied to each $\hat{y}_h^{[l]}$.

3 Building the linear combination matrix A

In the previous section, we limited ourselves to introduce the linear combination matrix A in expression (7), in line with the novel classification distinguishing between free (unconstrained) and constrained variables. In this section we propose two ways of building such a matrix in practice.

First, consider the simple case where there are no redundant constraints ($n_c = p$) and the first n_c columns of Γ are linearly independent, so that $y_t = x_t = \begin{bmatrix} c'_t & u'_t \end{bmatrix}'$ and $\Gamma y_t = \mathbf{0}_{(n_c \times 1)}$. This homogeneous linear system can be written as

$$\Gamma y_t = \mathbf{0}_{(n_c \times 1)} \iff [\Gamma_c \ \Gamma_u] \begin{bmatrix} c_t \\ u_t \end{bmatrix} = \mathbf{0}_{(n_c \times 1)},$$

where $\Gamma_c \in \mathbb{R}^{(n_c \times n_c)}$ contains the coefficients for the constrained variables, and $\Gamma_u \in \mathbb{R}^{(n_c \times n_u)}$ those for the free ones. Thanks to its non-singularity, Γ_c can be used to derive the equivalent zero-constrained representation:

$$\begin{bmatrix} I_{n_c} & (\Gamma_c)^{-1} \Gamma_u \end{bmatrix} \begin{bmatrix} c_t \\ u_t \end{bmatrix} = C y_t = \mathbf{0}_{(n_c \times 1)}, \tag{12}$$

where

$$C = \begin{bmatrix} I_{n_c} & -A \end{bmatrix} \quad \text{and} \quad A = -(\Gamma_c)^{-1} \Gamma_u. \tag{13}$$

In practical situations, mostly if many variables and/or constraints are involved, categorizing variables as either constrained or free may be a challenging task⁵: the goal is to identify a valid set of free variables with invertible coefficient matrix Γ_c .

⁵ The issue of defining a valid set of free variables is studied by Zhang et al. (2023) in the framework of hierarchical/grouped reconciliation with immutable forecasts.

3.1 General (redundant) linear constraints framework

When n is large, it is not always immediate to find a set of non-redundant constraints, so the method shown in expression (12) may be hardly applied in several real-life situations. We propose to overcome these issues by employing standard linear algebra tools, like the Reduced Row Echelon Form or the QR decomposition (Golub and Van Loan 1996; Meyer 2000), that are able to deal with redundant constraints and do not request any a priori classification of the single variables entering the multiple time series.

3.1.1 Reduced Row Echelon Form (rref)

A matrix is said to be in rref (Meyer 2000) if and only if the following three conditions hold:

- it is in row echelon form;
- the pivot in each row is 1;
- all entries above each pivot are 0.

The idea is then very simple: classify as constrained the variables corresponding to the pivot positions of the rref representation coefficient matrix, while the remaining columns form the linear combination matrix A . Usually a rref form is obtained through a Gauss-Jordan elimination (more details in Meyer 2000). So, let $Z \in \mathbb{R}^{(n_c \times n)}$ be the rref of Γ deprived of any possible null rows, then a permutation matrix P can be obtained starting from the pivot columns of Z , such that

$$y_t = Px_t = [x_{\pi_c(1),t} \cdots x_{\pi_c(n_c),t} \ x_{\pi_u(1),t} \cdots x_{\pi_u(n_u),t}],$$

where $\pi_c(i)$, $i = 1, \dots, n_c$, is the position of the i -th pivot column (i.e., one of the columns that identify the constrained variables) and $\pi_u(j)$, $j = 1, \dots, n_u$, is the position of the j -th no-pivot column (i.e., one of the columns associated to the free variables). Then, the linear combination matrix A can be extracted from the expression

$$C = ZP' = [I_{n_c} \ -A].$$

Additional examples can be found in the online appendix.

3.1.2 QR decomposition

Given the $(p \times n)$ coefficient matrix Γ of the zero-constrained representation (1), $\Gamma = QRP$ is a QR decomposition with column pivoting (Lyche 2020), where $Q \in \mathbb{R}^{(p \times p)}$ is a square and orthonormal matrix ($Q'Q = QQ' = I_p$), $P \in \{0, 1\}^{(n \times n)}$ is a permutation matrix, and $R \in \mathbb{R}^{(p \times n)}$ is an upper trapezoidal matrix (Anderson et al. 1992, 1999) such that

$$R = \begin{cases} \begin{bmatrix} R_c & R_u \\ & \end{bmatrix} & \text{if } \Gamma \text{ is full-rank} \\ \begin{bmatrix} R_c & R_u \\ \mathbf{0}_{[(p-n_c) \times n_c]} & \mathbf{0}_{[(p-n_c) \times n_u]} \end{bmatrix} & \text{if } \Gamma \text{ is not full-rank} \end{cases},$$

where $R_c \in \mathbb{R}^{(n_c \times n_c)}$ is upper triangular, and $R_u \in \mathbb{R}^{(n_c \times n_u)}$ is nonsingular (Golub and Van Loan 1996). Applying this decomposition to the homogenous system (1), we obtain

$$\Gamma x_t = QRPx_t = Q\mathbf{R}y_t = \mathbf{0}_{(p \times 1)},$$

that is equivalent to (Meyer 2000)

$$\begin{cases} Qz = \mathbf{0}_{(p \times 1)} \\ \mathbf{R}y_t = z = \mathbf{0}_{(p \times 1)} \end{cases}.$$

Due to the non-singularity of Q , $z = \mathbf{0}_{(p \times 1)}$ is the unique solution to the homogenous system $Qx = \mathbf{0}_{(p \times 1)}$ (Lyche 2020). Then, the homogeneous system (1) representing a general linearly constrained time series may be re-written as⁶ $[R_c \ R_u]y_t = \mathbf{0}_{(n_c \times 1)}$. Finally, from (12) we obtain

$$C = [I_{n_c} \quad -A] \quad \text{and} \quad A = -R_c^{-1}R_u,$$

where R_c is invertible by construction (Golub and Van Loan 1996). It is worth noting that the Pivoted QR decomposition generates a permutation matrix P that “moves” the free variables in x_t to the bottom part of the re-ordered vector y_t , that is $Px_t = y_t = [c'_t \ u'_t]'$.

It should be noted that in both cases (QR and rref), $P = I_n$ if the first n_c columns of Γ are linearly independent. This means that both algorithms start by assuming as constrained and free the variables as they appear in x , whose ordering is then changed only if it is not feasible for the constraints operating on the multivariate time series in equation (1).

4 Empirical applications

In this section we present two macroeconomics applications involving general linearly constrained multiple time series which do not have a genuine hierarchical/grouped structure. In the first case, in the wake of Athanasopoulos et al. (2020), we forecast the Australian *GDP* from income and expenditure sides, for which Bisaglia et al. (2020) already provided a full row-rank C matrix. The second application concerns the European *GDP* disaggregated by three sides (income, expenditure and output) and 19 member countries. In this case, building a full row-rank zero constraints matrix is not an easy task, so we use a QR decomposition (Sect. 3.1). Detailed informations on the variables in either dataset are reported in the online appendix.

⁶ Possible null rows, present if Γ is not full-rank, are removed.




Fig. 3 Linear combination matrix A for the Australian GDP from income and expenditure sides

4.1 Reconciled probabilistic forecasts of the Australian GDP from income and expenditure sides

Athanasopoulos et al. (2020) first considered the reconciliation of point and probabilistic forecasts of the 95 Australian Quarterly National Accounts (QNA) variables that describe the Gross Domestic Product (GDP) at current prices from the income and expenditure sides, interpreted as two distinct hierarchies. In the former case (income), GDP is the top level aggregate of a hierarchy of 15 lower level aggregates with $n_a^I = 6$ and $n_b^I = 10$, whereas in the latter (expenditure), GDP is the top level aggregate of a hierarchy of 79 time series, with $n_a^E = 27$ and $n_b^E = 53$ (for details, Athanasopoulos et al. 2020; Bisaglia et al. 2020; Di Fonzo and Girolimetto 2022, 2023).

Considering these two hierarchies as distinct yields different GDP forecasts depending on the considered disaggregation (either by income or expenditure). The fact that the two hierarchical structures describing the National Accounts share only the same top-level series (GDP), prevents the adoption for the whole set of $n = 95$ distinct variables of the original structural reconciliation approach proposed by Hyndman et al. (2011). However, it is possible to use the results shown so far for a general linearly constrained multiple time series. The homogeneous constraints valid for the variables are described by the following (33×95) matrix Γ (Bisaglia et al. 2020):




Fig. 4 Australian *GDP* empirical one-step-ahead forecast distributions for 2018:Q1, shr joint bootstrap-based reconciliation approach. Empirical Cumulative Distribution Function (left), and Smoothed density (right)

$$\Gamma = \begin{bmatrix} 1 & \mathbf{0}'_5 & -\mathbf{1}'_{10} & \mathbf{0}'_{26} & \mathbf{0}'_{53} \\ 1 & \mathbf{0}'_5 & \mathbf{0}'_{10} & \mathbf{0}'_{26} & -\mathbf{1}'_{53} \\ \mathbf{0}_5 & \mathbf{I}_5 & -\mathbf{A}^I & \mathbf{0}_{5 \times 26} & \mathbf{0}_{5 \times 53} \\ \mathbf{0}_{26} & \mathbf{0}_{26 \times 5} & \mathbf{0}_{26 \times 10} & \mathbf{I}_{26} & -\mathbf{A}^E \end{bmatrix}, \tag{14}$$

where $\mathbf{A}^I \in \{0, 1\}^{(5 \times 10)}$ and $\mathbf{A}^E \in \{0, 1\}^{(26 \times 53)}$ are the aggregation matrices for the income and the expenditure sides, respectively, and Γ has already full row-rank. A structural-like representation of the multiple time series that incorporates both sides' accounting constraints may be obtained by transforming Γ through, for example, the QR technique described in Sect. 3.1. This operation results in a (33×95) matrix $\mathbf{C} = [\mathbf{I}_{33} \quad -\mathbf{A}]$, where \mathbf{A} is the (33×62) linear combination matrix shown in Fig. 3, and $\mathbf{S} = [\mathbf{A}' \quad \mathbf{I}_{62}]'$ is the structural-like matrix (see Sect. 2.2).

We perform a forecasting experiment as the one designed by Athanasopoulos et al. (2020). Base forecasts from $h = 1$ quarter ahead up to $h = 4$ quarters ahead for all the 95 separate time series have been obtained through simple univariate ARIMA models selected using the `auto.arima` function of the R-package `forecast` (Hyndman and Khandakar 2008). The first training sample is set from 1984:Q4 to 1994:Q3, and a recursive training sample with expanding window length is used, for a total of 94 forecast origins. Finally the reconciled forecasts are obtained using three reconciliation approaches (ols, wls and shr, see Sect. 2.1) through the R package `Foreco` (Girolimetto and Di Fonzo 2023).

In Athanasopoulos et al. (2020) the probabilistic forecasts of the Australian quarterly *GDP* aggregates are separately reconciled from income (\tilde{X}_{GDP}^I) and expenditure (\tilde{X}_{GDP}^E) sides. This means that the empirical forecast distributions \tilde{X}_{GDP}^I and \tilde{X}_{GDP}^E are each coherent (see Sect. 2.3) within its own pertaining side with the other empirical forecast distributions, but in general $\tilde{X}_{GDP}^I \neq \tilde{X}_{GDP}^E$ at any forecast horizon. This circumstance could confuse the user, mostly when the difference between the empirical

forecast distributions is not negligible, as shown in Fig. 4, where the *GDP* empirical forecast distributions from income and expenditure sides for 2018:Q1 are presented along with their fully reconciled counterparts through the shr joint bootstrap-based reconciliation approach (see Sect. 2.3.2).⁷

4.1.1 Point and probabilistic forecasting accuracy

To evaluate the accuracy of the point forecasts we use the Mean Square Error (MSE)⁸:

$$MSE_{j,h} = \frac{1}{nT_h} \sum_{i=1}^n \sum_{t=1}^{T_h} \tilde{e}_{i,j,t+h}^2,$$

where $\tilde{e}_{i,j,t+h} = (\tilde{y}_{i,j,t+h} - y_{i,j,t+h})$ is the the h -step-ahead forecast error using the approach j to forecast the i -th series, $j = 0$ denotes the base forecast (i.e., $\tilde{y}_{i,0,t+h} = \hat{y}_{i,t+h}$), and t is the forecast origin. To assess any improvement in the reconciled forecasts compared to the base ones, we use the *MSE*-skill score:

$$\left(1 - \frac{MSE_{j,h}}{MSE_{0,h}} \right) \times 100.$$

The accuracy of the probabilistic forecasts is evaluated using the Cumulative Rank Probability Score (CRPS, Gneiting and Katzfuss 2014):

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

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 taken from the predictive distribution and $\mathbf{z} \in \mathbb{R}^n$ is the observation vector. In addition, to evaluate the forecasting accuracy for the whole system, we employ the Energy Score (ES), that is the CRPS extension to the multivariate case⁹:

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

⁷ Note that the naive practice of averaging *GDP* forecasts from different sides yields a single forecast, that is though inconsistent with the component variables from both sides.

⁸ The Mean Absolute Scaled Error (MASE) leads to the same conclusions (see the online appendix).

⁹ An alternative to the Energy Score is the Variogram Score (Scheuerer and Hamill 2015), considered in the online appendix, that leads to similar conclusions.

Table 1 MSE and CRPS-skill scores (relative to base forecasts) for the point and probabilistic Australian GDP forecasts from alternative reconciliation approaches

h	Point forecasts			Probabilistic forecasts - CRPS(%)					
	MSE (%)			Bootstrap			Gaussian		
	ols	wls	shr	ols	wls	shr	ols	wls	shr
<i>Income</i>									
1	1.63	1.07	5.41	0.57	-1.26	0.52	0.13	-1.87	-0.45
2	2.54	5.68	6.10	1.46	0.68	-0.50	1.11	1.26	-0.34
3	2.28	7.81	4.56	0.18	-0.64	-1.93	0.14	-0.24	-1.69
4	1.98	9.33	7.04	<i>-0.08</i>	-1.12	-1.35	<i>-0.23</i>	-1.00	-1.39
<i>Expenditure</i>									
1	4.53	0.07	2.48	1.10	0.78	0.08	0.64	-0.83	-0.70
2	5.09	3.90	1.72	2.34	1.25	-1.07	1.68	0.58	-1.95
3	6.96	9.18	6.24	2.42	1.50	-0.38	1.69	1.00	-1.02
4	8.01	11.76	8.34	3.73	2.73	0.72	3.18	2.49	0.20
<i>Fully reconciled</i>									
1	4.59	1.14	4.77	1.13	-0.75	-0.20	<i>-0.59</i>	-3.46	-2.85
2	5.76	6.24	4.76	2.81	1.07	-1.27	1.30	-0.33	-3.02
3	7.31	10.94	8.21	1.99	0.26	-1.46	0.91	-0.62	-2.80
4	7.90	13.24	10.81	2.83	0.75	-0.63	1.86	-0.43	-2.05

Negative values are highlighted in italic, the best for each row is marked in bold

4.1.2 Results

Table 1 shows the MSE and CRPS-skill scores for the GDP point and probabilistic reconciled forecasts. Table 2 presents the MSE and ES-skill scores for all 95 Australian QNA variables from both income and expenditure sides. The ‘Income’ and ‘Expenditure’ panels, respectively, reproduce the results found by Athanasopoulos et al. (2020). The ‘Fully reconciled’ panels show the skill scores for the simultaneously reconciled forecasts. For the point forecasts, all the reconciliation approaches improve forecast accuracy compared to the base forecasts. In detail, shr is almost always the best approach for the one-step-ahead forecasts, whereas wls is competitive for $h \geq 2$. Looking at the probabilistic reconciliation results in Table 1, it is worth noting that for GDP ols outperforms both wls and shr, whatever side and framework (parametric or not) is considered. However, when all 95 variables are considered (Table 2), shr and wls approaches almost always show the best performance. In the Gaussian framework, these results are confirmed for the income side, whereas shr performs poorly when we look at the expenditure side (either fully reconciled or not).

In Fig. 5 are shown the results obtained by the non-parametric Friedman test and the post hoc “Multiple Comparison with the Best” (MCB) Nemenyi test (Koning et al. 2005; Kourentzes and Athanasopoulos 2019; Makridakis et al. 2022) to determine if the forecasting performances of the different techniques are significantly different from one another. In general, wls always falls in the set of the best performing approaches.

Table 2 MSE and ES-skill scores (relative to base forecasts) for the point and probabilistic forecasts from alternative reconciliation approaches (all Australian QNA variables)

h	Point forecasts			Probabilistic forecasts - ES(%)					
	MSE (%)			Bootstrap			Gaussian		
	ols	wls	shr	ols	wls	shr	ols	wls	shr
<i>Income</i>									
1	3.16	6.32	10.55	2.30	4.24	6.15	1.87	3.47	5.15
2	2.58	6.07	8.18	2.14	4.08	4.08	1.62	3.41	3.08
3	2.18	5.81	4.09	1.78	3.29	3.19	1.28	2.67	2.35
4	2.18	6.78	5.51	1.86	3.73	4.42	1.42	3.10	3.67
<i>Income—Fully reconciled</i>									
1	3.78	7.57	8.85	2.77	4.58	4.87	1.30	1.94	2.25
2	2.91	6.12	6.92	2.59	3.95	3.40	1.20	1.59	0.87
3	2.67	6.23	5.57	2.24	3.70	3.17	0.92	1.60	0.79
4	2.87	7.21	6.07	2.65	4.40	3.95	1.44	2.34	1.59
<i>Expenditure</i>									
1	6.50	6.75	8.78	3.71	4.54	4.94	2.20	1.92	2.58
2	4.90	5.50	5.52	2.88	3.04	2.68	1.28	0.67	0.32
3	4.27	6.08	5.65	2.57	2.94	2.29	0.88	0.71	0.09
4	4.01	6.69	5.20	2.43	2.65	1.63	0.77	0.44	-0.65
<i>Expenditure—Fully reconciled</i>									
1	6.51	6.82	9.08	3.76	4.57	4.99	2.19	1.79	2.44
2	5.09	6.24	6.54	3.03	3.48	3.14	1.18	0.82	0.57
3	4.38	6.75	5.94	2.58	3.11	2.29	0.81	0.69	-0.12
4	3.98	7.33	5.82	2.36	2.82	1.81	0.60	0.32	-0.70

Negative values are highlighted in italic, the best for each row is marked in bold

For the bootstrap-based probabilistic reconciled forecasts of the expenditure side variables, wls and shr significantly improves in terms of MSE and CRPS compared to the base forecasts. This result is confirmed in the remaining cases as well.

In conclusion, when income and expenditure sides are simultaneously considered for both point and probabilistic forecasts, forecast reconciliation succeeds in improving the base forecasts of *GDP* and its component aggregates, while preserving the full coherence with the National Accounts constraints.

4.2 Reconciled probabilistic forecasts of the European Area *GDP* from output, income and expenditure sides

In this section, we consider the system of European QNA for the *GDP* at current prices (in euro), with time series spanning the period 2000:Q1-2019:Q4. This system has many variables linked by several, possibly redundant, accounting constraints, such that it is difficult to manually build a system of non-redundant constraints.

The National Accounts are a coherent and consistent set of macroeconomic indicators that are used mostly for economic research and forecasting, policy design,

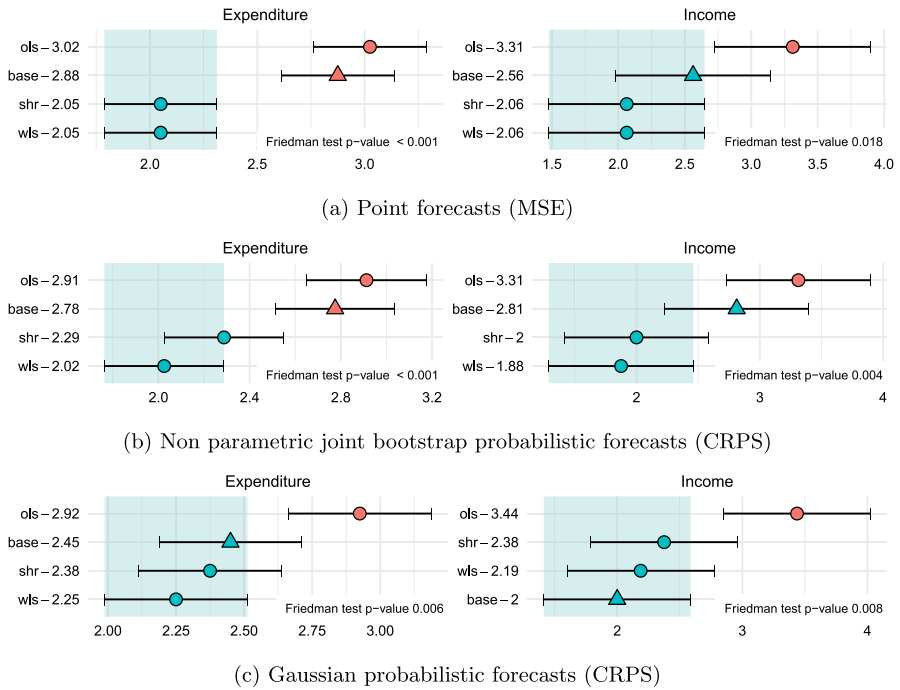


Fig. 5 MCB Nemenyi test for the fully reconciled forecasts of the Australian QNA variables at any forecast horizon. In each panel, the Friedman test p -value is reported in the lower right corner. The mean rank of each approach is shown to the right of its name. Statistical differences in performance are indicated if the intervals of two forecast reconciliation approaches do not overlap. Thus, approaches that do not overlap with the blue interval are considered significantly worse than the best, and vice-versa

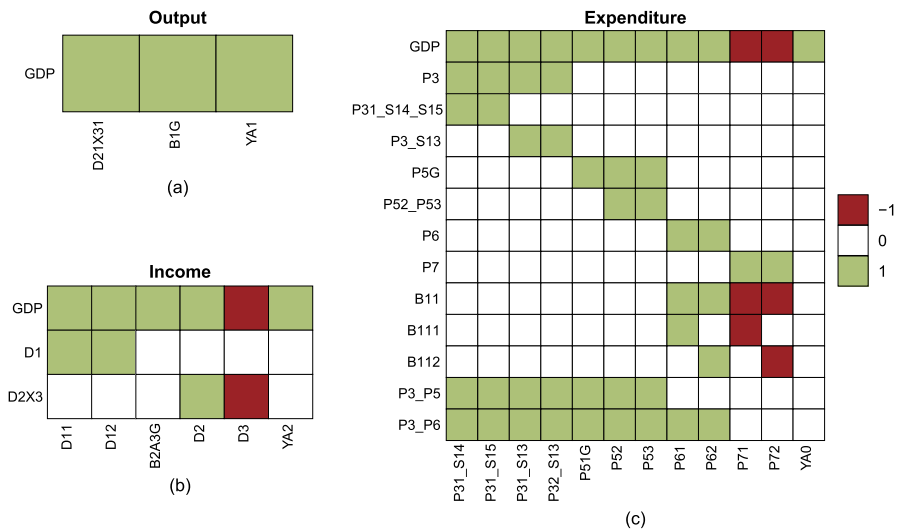


Fig. 6 Linear combination matrices A for the European Area GDP : output side in panel (a), income side in panel (b), expenditure side in panel (c)

and coordination mechanisms. In this dataset, *GDP* is a key macroeconomic quantity that is measured using three main approaches, namely output (or production), income and expenditure. These parallel systems internally present a well-defined hierarchical structure of variables with relevant economic significance, such as Final consumption, on the expenditure side, Gross operating surplus and mixed income on the income side, and Total gross value added on the output side. In the EU countries, the data is processed on the basis of the ESA 2010 classification and are released by Eurostat.¹⁰ We consider the 19 Euro Area member countries (Austria, Belgium, Finland, France, Germany, Ireland, Italy, Luxembourg, Netherlands, Portugal, Spain, Greece, Slovenia, Cyprus, Malta, Slovakia, Estonia, Latvia, and Lithuania) that have been using the euro since 2015. In Fig. 6 we have represented the aggregation matrices describing output, income, and expenditure constraints, respectively: in panel (a), matrix A_O for the output side, in panel (b) matrix A_I for the income side, and in panel (c) matrix A_E for the expenditure side. The zero-constraints coefficient matrix describing the QNA variables for a single country can thus be written as

$$\Gamma_{GDP} = \begin{bmatrix} K_E & -A_E & \mathbf{0}_{(13 \times 6)} & \mathbf{0}_{(13 \times 3)} \\ K_I & \mathbf{0}_{(3 \times 12)} & -A_I & \mathbf{0}_{(3 \times 3)} \\ K_O & \mathbf{0}_{(1 \times 12)} & \mathbf{0}_{(1 \times 6)} & -A_O \end{bmatrix},$$

where K_E , K_I and K_O , respectively, are the following (13×15) , (3×15) and (1×15) matrices:

$$K_E = \begin{bmatrix} 1 & \mathbf{0}'_{(1 \times 12)} & \mathbf{0}'_{(1 \times 2)} \\ \mathbf{0}_{(12 \times 1)} & I_{12} & \mathbf{0}_{(12 \times 2)} \end{bmatrix}, K_I = \begin{bmatrix} 1 & \mathbf{0}'_{(1 \times 12)} & \mathbf{0}'_{(1 \times 2)} \\ \mathbf{0}_{(2 \times 1)} & \mathbf{0}_{(2 \times 12)} & I_2 \end{bmatrix}, K_O = [1 \ \mathbf{0}'_{(1 \times 14)}].$$

This disaggregation is common for almost all European countries, the only differences being related to the presence/absence of an aggregate measuring the statistical discrepancy in each accounting side.¹¹ The (361×720) matrix describing the accounting relationships for the whole EA19 QNA by countries and accounting sides can be written as follows:

$$\Gamma = \begin{bmatrix} \Gamma_{GDP} & \mathbf{0}_{(17 \times 684)} \\ [\mathbf{0}_{(21 \times 15)} \ I_{21}] & -\mathbf{1}'_{19} \otimes [\mathbf{0}_{(21 \times 15)} \ I_{21}] \\ \mathbf{0}_{(323 \times 36)} & I_{19} \otimes \Gamma_{GDP} \end{bmatrix},$$

where \otimes is the Kronecker product, and the top-left portion of Γ refers to the European Area aggregates as a whole. In order to proceed with the calculations, it is necessary to eliminate the columns related to null variables (e.g., the statistical discrepancy aggregate for the countries/sides where it is not contemplated, see

¹⁰ Further information can be found at <https://ec.europa.eu/eurostat/esa2010/> and <https://ec.europa.eu/eurostat/web/national-accounts/data/database>.

¹¹ For 14 countries (Belgium, France, Germany, Italy, Luxembourg, Netherlands, Spain, Greece, Slovenia, Cyprus, Malta, Slovakia, Latvia, and Lithuania), the expenditure, income and output statistical discrepancies are not present. A statistical discrepancy aggregate is present in the output QNA of Portugal and in the expenditure QNA of Finland, Estonia and Austria. Ireland is the only country where a statistical discrepancy aggregate is present in all accounting sides.

Table 3 MSE and ES-skill scores (relative to base forecasts) for the point and probabilistic forecasts from alternative reconciliation approaches (European Area QNA)

	Point forecasts			Probabilistic forecasts - ES(%)					
	MSE (%)			Bootstrap			Gaussian		
	ols	wls	shr	ols	wls	shr	ols	wls	shr
GDP	14.0	41.1	21.6	9.7	25.8	15.7	3.4	20.8	3.7
<i>Sides</i>									
Expenditure	15.8	31.0	28.3	9.1	20.2	16.3	6.8	15.7	8.7
Income	- 2.2	27.3	9.2	- 0.3	15.3	5.6	- 2.7	11.6	- 3.2
Output	0.3	35.0	15.0	1.2	21.6	12.9	- 1.5	17.1	1.3
<i>Countries</i>									
EA19	17.6	37.1	31.6	10.1	23.5	18.1	7.8	19.3	10.0
Austria	<-30	18.5	21.5	- 26.8	12.1	10.0	- 21.6	9.1	5.4
Belgium	- 5.9	13.5	21.0	- 2.7	10.7	10.8	- 3.1	8.3	5.9
Cyprus	<-30	11.7	8.6	<-30	7.1	1.7	<-30	3.6	- 2.8
Estonia	<-30	20.9	26.8	<-30	12.7	11.2	<-30	9.7	6.7
Finland	<-30	24.8	14.7	<-30	15.3	9.5	<-30	13.2	3.4
France	4.4	19.1	- 0.1	3.2	11.7	3.8	0.3	7.9	- 6.8
Germany	15.0	20.6	25.0	13.5	21.6	18.6	7.5	11.1	6.8
Greece	<-30	1.5	<-30	- 12.9	3.9	- 7.1	- 15.1	0.8	- 14.1
Ireland	4.1	6.6	3.7	2.1	5.9	3.0	0.5	2.9	- 0.7
Italy	9.8	12.6	24.0	8.4	10.8	17.1	4.4	4.9	7.8
Latvia	<-30	26.2	18.3	<-30	13.0	8.7	<-30	13.7	5.6
Lithuania	<-30	27.3	28.5	<-30	15.0	12.8	<-30	12.1	8.4
Luxembourg	<-30	6.6	- 10.3	<-30	6.1	- 3.0	<-30	2.0	- 10.0
Malta	<-30	5.4	- 9.9	<-30	4.9	- 7.4	<-30	0.8	- 12.0
Netherlands	7.4	16.9	11.9	6.1	14.1	7.3	3.5	11.7	4.4
Portugal	<-30	11.4	- 8.1	<-30	7.4	- 3.4	<-30	4.1	- 10.1
Slovakia	<-30	21.5	16.4	<-30	10.0	6.9	<-30	9.2	2.0
Slovenia	<-30	24.1	15.5	<-30	13.5	8.3	<-30	11.1	1.8
Spain	6.5	17.7	0.9	5.3	9.2	1.4	1.7	6.4	- 7.3

Negative values are highlighted in italic, the best for each row is marked in bold

footnote 11). Then, to eliminate possible remaining redundant constraints, we apply the QR decomposition (Sect. 3.1). Finally, we obtain the linear combination matrix A , which refers to 311 free and 358 constrained time series, and the full rank zero-constraints matrix $C = [I_{358} \ -A]$.

A rolling forecast experiment with expanding window is performed using ARIMA models to produce the individual series' base forecasts. The first training set is set from 2000:Q1 to 2009:Q4, which gives 40 one-step-ahead, 39 two-step-ahead, 38 three-step-ahead and 37 four-step-ahead ARIMA forecasts, respectively. The used reconciliation approaches are ols, wls and shr, and the forecast accuracy is evaluated through MSE, CRPS and ES indices, as described in Sect. 4.1.

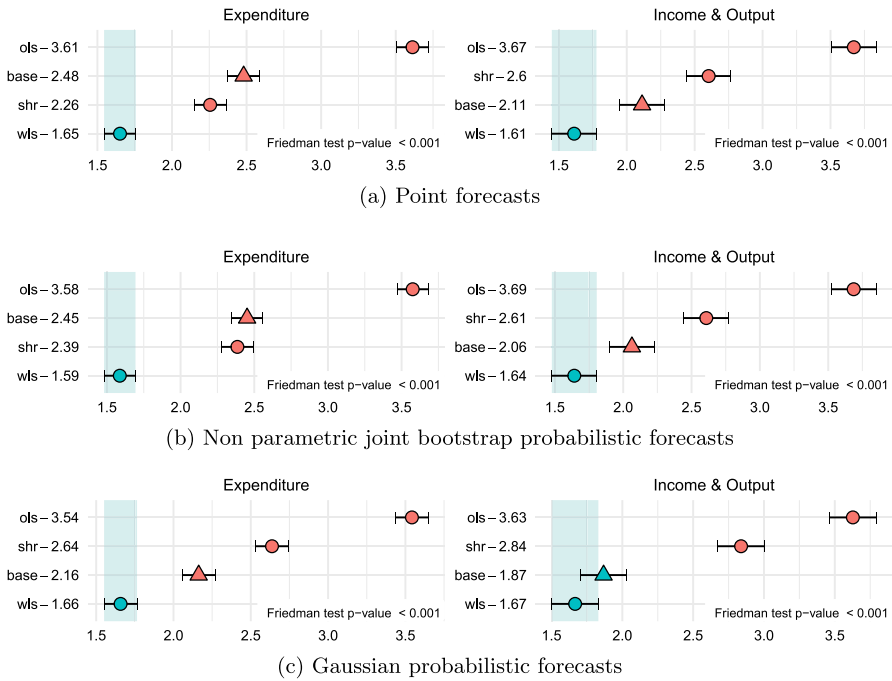


Fig. 7 MCB Nemenyi test for the fully reconciled forecasts of the European Area QNA variables at any forecast horizon. In each panel, the Friedman test p-value is reported in the lower right corner. The mean rank of each approach is shown to the right of its name. Statistical differences in performance are indicated if the intervals of two forecast reconciliation approaches do not overlap. Thus, approaches that do not overlap with the blue interval are considered significantly worse than the best, and vice-versa

4.2.1 Results

Table 3 shows the MSE indices for point forecasts, and the ES indices for probabilistic nonparametric (bootstrap) and parametric (Gaussian) forecasts, respectively. The rows of the table are divided into three parts: the first row shows the results for *GDP*, the second to fourth rows the National Accounts' divisions (income, expenditure, or output sides), while the remaining rows correspond to the 19 countries and EA19. All forecast horizons are considered.¹²

When only *GDP* is considered, any reconciliation approach consistently outperforms the base forecasts, both in the point and probabilistic cases. The wls approach confirms a good performance when we look at the income, expenditure, and output sides, while ols shows the worst performance. When the parametric framework is considered, shr is worse than the base forecast for the income side. At country level, ols is the approach that overall shows the worst performance, with many relative losses in the accuracy indices higher than 30%. It is worth

¹² A disaggregated analysis by forecast horizon is reported in the online appendix.

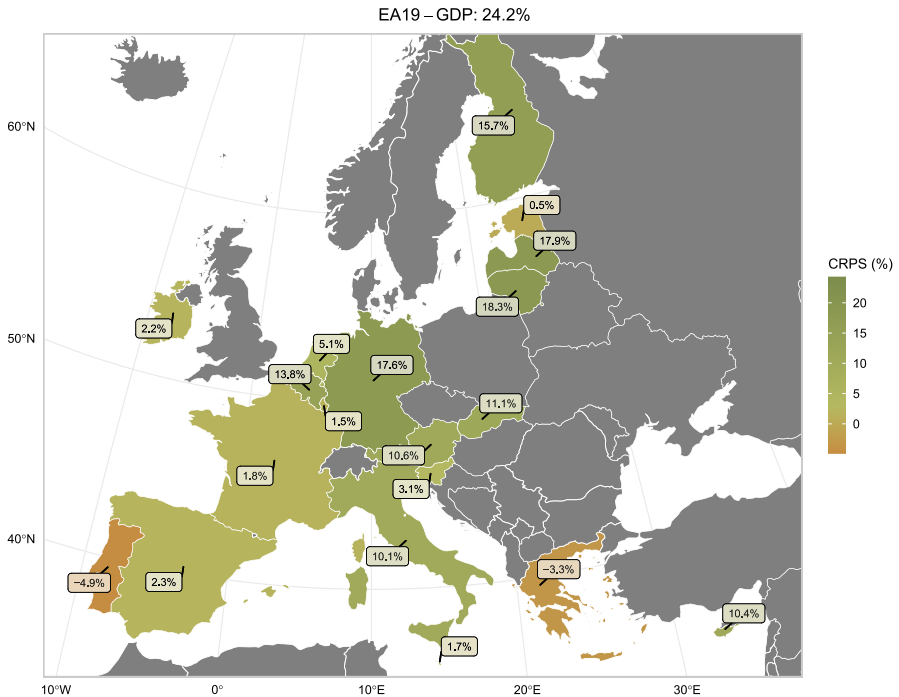


Fig. 8 CRPS-skill scores of the one-step-ahead *GDP* non-parametric joint bootstrap probabilistic reconciled (wls) forecasts for the 19 Euro Area countries

noting that all reconciliation approaches always perform well for the whole Euro Area (EA19). Overall, in this forecasting experiment wls appears to be the most performing reconciliation approach, showing no negative skill score, and improvements higher than 20% and 10% for nonparametric and parametric probabilistic frameworks, respectively.

Figure 7 shows the MCB Nemenyi test at any forecast horizon, distinct by expenditure side and income and output sides, respectively. The results just seen are further confirmed by this alternative forecast assessment tool: it clearly appears that the wls approach almost always significantly improves compared to the base forecast, in terms of both point and probabilistic forecasts.

Finally, a visual evaluation of the accuracy improvement obtained through wls forecast reconciliation, although limited to a single forecast horizon, is offered by Fig. 8, showing the European map with the CRPS skill scores for the one-step-ahead *GDP* non-parametric probabilistic reconciled forecasts. It is worth noting that only for two countries (Greece and Portugal) a decrease is registered (−3.3% and −4.9%, respectively). In all other cases, improvements in the forecasting accuracy are obtained, with Germany and Lithuania leading the way with about 18%. Furthermore, the improvement in the forecasting accuracy for the EA19 – *GDP* is 24.2%, the highest throughout the whole Euro Area.

5 Conclusions

Producing and using coherent information, disaggregated by different characteristics useful for different decision levels, is an important task for any practitioner and quantitative-based decision process. At this end, in this article we aimed to generalize the results valid for the forecast reconciliation of a genuine hierarchical/grouped time series to the case of a general linearly constrained time series, where the distinction between upper and bottom variables, which is typical in the hierarchical setting, is no longer meaningful.

Two motivating examples have been considered, both coming from the National Accounts field, namely the forecast reconciliation of quarterly *GDP* of (i) Australia, disaggregated by income and expenditure variables, and (ii) Euro Area 19, disaggregated by the income, output and expenditure side variables of 19 component countries. In both cases, the structure of the time series involved cannot be represented according to a genuinely hierarchical/grouped scheme, so the standard forecast reconciliation techniques fail in producing a “unique” *GDP* forecast, either point or probabilistic, making it necessary to solve this annoying issue.

We have shown that using well known linear algebra tools, it is always possible to establish a formal connection between the unconstrained GLS structural approach originally developed by Hyndman et al. (2011), and the projection approach to reconciliation dating back to the work by Stone et al. (1942), and then applied to solve different reconciliation problems (Di Fonzo and Marini 2011; van Erven and Cugliari 2015; Wickramasuriya et al. 2019; Di Fonzo and Girolimetto 2023). We propose a new classification of the variables forming the multiple time series as free and constrained, respectively, that can be seen as a generalization of the standard bottom/upper variables classification used in the hierarchical setting. Furthermore, we show techniques for deriving a linear combination matrix describing the relationships between these variables, starting from the coefficient matrix summarizing the (possible redundant) constraints linking the series.

The application of these findings to both point and probabilistic reconciliation techniques proved to be easy to implement and powerful, resulting in significant improvements in the forecasting accuracy of *GDP* and its components in both forecasting experiments.

Appendix A: Derivation of equation (11)

Denote $W_h = k_h W$, where k_h is a proportionality constant and W is the p.d. covariance matrix used in the point forecast reconciliation formula (4). Then:

$$\begin{aligned}
\tilde{W}_h &= SGW_h G' S' \\
&= k_h SGWG' S' \\
&= k_h S (S' W^{-1} S)^{-1} \underbrace{S' W^{-1} W W^{-1} S (S' W^{-1} S)^{-1}}_{I_n} S' \\
&= k_h S (S' W^{-1} S)^{-1} S' \\
&= k_h S \underbrace{(S' W^{-1} S)^{-1} S' W^{-1} W}_{G} = k_h SGW.
\end{aligned}$$

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s10260-023-00738-6>.

Acknowledgements The authors acknowledge financial support from project PRIN2017 “HiDEA: Advanced Econometrics for High-frequency Data”, 2017RSMPZZ.

Funding Open access funding provided by Università degli Studi di Padova.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Anderson E, Bai Z, Bischof C, Blackford S, Demmel J, Dongarra J, Du Croz J, Greenbaum A, Hammerling S, McKenney A et al. (1999) LAPACK Users’ Guide: Third Edition. Software, Environments, and Tools. Society for Industrial and Applied Mathematics
- Anderson E, Bai Z, Dongarra J (1992) Generalized QR factorization and its applications. *Linear Algebra Appl* 162–164:243–271. [https://doi.org/10.1016/0024-3795\(92\)90379-0](https://doi.org/10.1016/0024-3795(92)90379-0)
- Athanasopoulos G, Ahmed RA, Hyndman RJ (2009) Hierarchical forecasts for Australian domestic tourism. *Int J Forecast* 25(1):146–166. <https://doi.org/10.1016/j.ijforecast.2008.07.004>
- Athanasopoulos G, Gamakumara P, Panagiotelis A, Hyndman RJ, Affan M (2020) Hierarchical forecasting. In: Fuleky P (ed) *Macroeconomic forecasting in the era of big data*, vol 52. Springer International Publishing, Cham, pp 689–719. https://doi.org/10.1007/978-3-030-31150-6_21
- Ben Taieb S, Koo B (2019) Regularized regression for hierarchical forecasting without unbiasedness conditions. In: *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, Anchorage AK USA, pp 1337–1347. ACM
- Ben Taieb S, Taylor JW, Hyndman RJ (2021) Hierarchical probabilistic forecasting of electricity demand with smart meter data. *J Am Stat Assoc* 116(533):27–43. <https://doi.org/10.1080/01621459.2020.1736081>

- Bisaglia L, Di Fonzo T, Girolimetto D (2020) Fully reconciled GDP forecasts from income and expenditure sides. In: Pollice A, Salvati N, Schirripa Spagnolo F (eds) Book of short papers SIS 2020, 951–956. Pearson
- Byron RP (1978) The estimation of large social account matrices. *J R Stat Soc A (General)* 142(3):405. <https://doi.org/10.2307/2982515>
- Byron RP (1979) Corrigenda: the estimation of large social account matrices. *J R Stat Soc A (General)* 142(3):405. <https://doi.org/10.2307/2982515>
- Corani G, Azzimonti D, Augusto JPSC, Zaffalon M (2021) Probabilistic reconciliation of hierarchical forecast via Bayes' rule. *Mach Learn Knowled Discover Databases* 12459:211–226. https://doi.org/10.1007/978-3-030-67664-3_13
- Di Fonzo T, Girolimetto D (2022) Fully reconciled probabilistic GDP forecasts from Income and Expenditure sides. In: Balzanella A, Bini M, Cavicchia C, Verde R (eds) Book of short papers SIS 2022, 1376–1381. Pearson
- Di Fonzo T, Girolimetto D (2023) Cross-temporal forecast reconciliation: optimal combination method and heuristic alternatives. *Int J Forecast* 39(1):39–57. <https://doi.org/10.1016/j.ijforecast.2021.08.004>
- Di Fonzo T, Marini M (2011) Simultaneous and two-step reconciliation of systems of time series: methodological and practical issues. *J R Stat Soc C* 60(2):143–164. <https://doi.org/10.1111/j.1467-9876.2010.00733.x>
- Di Fonzo T, Marini M (2015) Reconciliation of systems of time series according to a growth rates preservation principle. *Stat Methods Appl* 24(4):651–669. <https://doi.org/10.1007/s10260-015-0322-y>
- Dunn DM, Williams WH, Dechaine TL (1976) Aggregate versus subaggregate models in local area forecasting. *J Am Stat Assoc* 71(353):68–71. <https://doi.org/10.1080/01621459.1976.10481478>
- Eckert F, Hyndman RJ, Panagiotelis A (2021) Forecasting swiss exports using Bayesian forecast reconciliation. *Eur J Oper Res* 291(2):693–710. <https://doi.org/10.1016/j.ejor.2020.09.046>
- Girolimetto D, Di Fonzo T (2023) FoReco: forecast reconciliation. R package v0.2.6. <https://danigiro.github.io/FoReco/>
- Gneiting T, Katzfuss M (2014) Probabilistic forecasting. *Annual Rev Stat Appl* 1(1):125–151. <https://doi.org/10.1146/annurev-statistics-062713-085831>
- Golub GH, Van Loan CF (1996) Matrix computations. Johns Hopkins University Press, Baltimore
- Gross CW, Sohl JE (1990) Disaggregation methods to expedite product line forecasting. *J Forecast* 9(3):233–254. <https://doi.org/10.1002/for.3980090304>
- Hyndman RJ, Ahmed RA, Athanasopoulos G, Shang HL (2011) Optimal combination forecasts for hierarchical time series. *Computat Stat Data Anal* 55(9):2579–2589. <https://doi.org/10.1016/j.csda.2011.03.006>
- Hyndman RJ, Athanasopoulos G (2021) Forecasting: principles and practice (3rd ed). Melbourne: OTexts. <https://otexts.com/fpp3/>
- Hyndman RJ, Khandakar Y (2008) Automatic time series forecasting: the forecast package for R. *J Stat Softw* 27: 1–22. <https://doi.org/10.18637/jss.v027.i03>
- Hyndman RJ, Lee AJ, Wang E (2016) Fast computation of reconciled forecasts for hierarchical and grouped time series. *Computat Stat Data Anal* 97:16–32. <https://doi.org/10.1016/j.csda.2015.11.007>
- Jeon J, Panagiotelis A, Petropoulos F (2019) Probabilistic forecast reconciliation with applications to wind power and electric load. *Eur J Oper Res* 279(2):364–379. <https://doi.org/10.1016/j.ejor.2019.05.020>
- Koning AJ, Franses PH, Hibon M, Stekler H (2005) The M3 competition: statistical tests of the results. *Int J Forecast* 21(3):397–409. <https://doi.org/10.1016/j.ijforecast.2004.10.003>
- Kourentzes N, Athanasopoulos G (2019) Cross-temporal coherent forecasts for Australian tourism. *Ann Tour Res* 75:393–409. <https://doi.org/10.1016/j.annals.2019.02.001>
- Ledoit O, Wolf M (2004) A well-conditioned estimator for large-dimensional covariance matrices. *J Multivar Anal* 88(2):365–411. [https://doi.org/10.1016/S0047-259X\(03\)00096-4](https://doi.org/10.1016/S0047-259X(03)00096-4)
- Leon SJ (2015) Linear algebra with applications, 9th edn. Pearson, Boston
- Lyche T (2020) Numerical linear algebra and matrix factorizations. Springer, New York
- Makridakis S, Spiliotis E, Assimakopoulos V (2022) M5 accuracy competition: results, findings, and conclusions. *Int J Forecast* 38(4):1346–1364. <https://doi.org/10.1016/j.ijforecast.2021.11.013>
- Meyer CD (2000) Matrix analysis and applied linear algebra. Society for Industrial and Applied Mathematics, Philadelphia

- Panagiotelis A, Athanasopoulos G, Gamakumara P, Hyndman RJ (2021) Forecast reconciliation: a geometric view with new insights on bias correction. *Int J Forecast* 37(1):343–359. <https://doi.org/10.1016/j.ijforecast.2020.06.004>
- Panagiotelis A, Gamakumara P, Athanasopoulos G, Hyndman RJ (2023) Probabilistic forecast reconciliation: properties, evaluation and score optimisation. *Eur J Oper Res* 306(2):693–706. <https://doi.org/10.1016/j.ejor.2022.07.040>
- Scheuerer M, Hamill TM (2015) Variogram-based proper scoring rules for probabilistic forecasts of multivariate quantities. *Mon Weather Rev* 143(4):1321–1334. <https://doi.org/10.1175/MWR-D-14-00269.1>
- Stone R, Champornowne DG, Meade JE (1942) The precision of national income estimates. *Rev Econ Stud* 9(2):111. <https://doi.org/10.2307/2967664>
- van Erven T, Cugliari J (2015) Game-theoretically optimal reconciliation of contemporaneous hierarchical time series forecasts. In: Antoniadis A, Poggi JM, Brossat X (eds) *Modeling and stochastic learning for forecasting in high dimensions*, vol. 217, pp 297–317. Cham, Springer International Publishing. https://doi.org/10.1007/978-3-319-18732-7_15
- Wickramasuriya SL (2021) Properties of point forecast reconciliation approaches. <https://doi.org/10.48550/arXiv.2103.11129>
- Wickramasuriya SL (2023) Probabilistic forecast reconciliation under the gaussian framework. *J Bus Econ Stat*. <https://doi.org/10.1080/07350015.2023.2181176>
- Wickramasuriya SL, Athanasopoulos G, Hyndman RJ (2019) Optimal forecast reconciliation for hierarchical and grouped time series through trace minimization. *J Am Stat Assoc* 114(526):804–819. <https://doi.org/10.1080/01621459.2018.1448825>
- Yagli GM, Yang D, Srinivasan D (2020) Reconciling solar forecasts: probabilistic forecasting with homoscedastic Gaussian errors on a geographical hierarchy. *Sol Energy* 210:59–67. <https://doi.org/10.1016/j.solener.2020.06.005>
- Yang D (2020) Reconciling solar forecasts: probabilistic forecast reconciliation in a nonparametric framework. *Sol Energy* 210:49–58. <https://doi.org/10.1016/j.solener.2020.03.095>
- Zhang B, Kang Y, Panagiotelis A, Li F (2023) Optimal reconciliation with immutable forecasts. *Eur J Oper Res* 308(2):650–660. <https://doi.org/10.1016/j.ejor.2022.11.035>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.