International Journal of Forecasting xxx (xxxx) xxx

Contents lists available at ScienceDirect

International Journal of Forecasting

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

ELSEVIER

Properties of the reconciled distributions for Gaussian and count forecasts

Lorenzo Zambon [a,b,*], Arianna Agosto [b], Paolo Giudici [b], Giorgio Corani [a]

[a] *Istituto Dalle Molle di Studi sull'Intelligenza Artificiale (IDSIA), Lugano, Switzerland*
[b] *University of Pavia, Italy*

A R T I C L E I N F O

Keywords:
Reconciliation
Hierarchical forecasting
Importance sampling
Intermittent time series
Probabilistic forecasts

A B S T R A C T

Reconciliation enforces coherence between hierarchical forecasts, in order to satisfy a set of linear constraints. While most works focus on the reconciliation of point forecasts, we consider probabilistic reconciliation and we analyze the properties of distributions reconciled via conditioning. We provide a formal analysis of the variance of the reconciled distribution, treating the case of Gaussian and count forecasts separately. We also study the reconciled upper mean in the case of one-level hierarchies, again treating Gaussian and count forecasts separately. We then show experiments on the reconciliation of intermittent time series related to the count of extreme market events. The experiments confirm our theoretical results and show that reconciliation largely improves the performance of probabilistic forecasting.

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

1. Introduction

Hierarchical forecasting requires forecasts to be coherent, i.e., to satisfy a set of linear constraints determined by the structure of the hierarchy. The base forecasts, computed independently on each time series of the hierarchy, are incoherent; *reconciliation* adjusts them to enforce coherence.

Most reconciliation approaches reconcile point forecasts (Di Fonzo & Girolimetto, 2022; Hyndman et al., 2011; Panagiotelis et al., 2021; Wickramasuriya et al., 2019). However, reconciled predictive distributions are required to support decision making (Kolassa, 2022). Panagiotelis et al. (2023) perform probabilistic reconciliation through projection, learning the parameters of the projection via stochastic gradient descent. Two limits of this approach are that it cannot reconcile count variables and it prevents an analytical study of the reconciled distribution. These considerations are also valid for other methods of probabilistic reconciliation (Hollyman et al., 2022; Jeon et al., 2019; Rangapuram et al., 2021; Taieb et al., 2021).

A different approach to probabilistic forecast reconciliation is constituted by reconciliation via conditioning. In the case of Gaussian base forecasts, it yields the reconciled distribution in closed form (Corani et al., 2020), with the same mean and variance as MinT (Wickramasuriya et al., 2019). Reconciliation via conditioning can also be used in the case of count time series, adopting a sampling approach (Corani et al., 2023; Zambon et al., 2024).

In this paper, we study the properties of reconciled distributions obtained via conditioning. In the Gaussian case we prove that, regardless of the amount of incoherence of the base forecasts, reconciliation decreases the variance of each variable of the hierarchy. In contrast, we show that in the discrete case, reconciliation can increase or decrease the variance of the bottom variables, depending on the probability p_c of the base forecasts being coherent.

* Corresponding author at: Istituto Dalle Molle di Studi sull'Intelligenza Artificiale (IDSIA), Lugano, Switzerland.

E-mail addresses: lorenzo.zambon@idsia.ch (L. Zambon), arianna.agosto@unipv.it (A. Agosto), paolo.giudici@unipv.it (P. Giudici), giorgio.corani@idsia.ch (G. Corani).

https://doi.org/10.1016/j.ijforecast.2023.12.004
0169-2070/© 2023 The Author(s). Published by Elsevier B.V. on behalf of International Institute of Forecasters. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

We then analyze the reconciled mean, restricting our analysis to the case of one-level hierarchies. In the Gaussian case, the reconciled upper mean is a combination of the bottom-up mean and the mean of the upper base forecast (Corani et al., 2020; Hollyman et al., 2021); we refer to this as the *combination* effect. However, in the case of count variables, we show that the reconciled mean of the upper time series can be lower than both the bottom-up mean and the base mean; we refer to this as the *concordant-shift* effect, as the reconciled means of all the time series are shifted towards zero. This happens when the base forecast distributions are right-skewed and reconciliation decreases the variance of the forecasts, shortening the right tails of the distributions and pulling the reconciled means of all time series towards zero.

We present experiments on the reconciliation of intermittent time series referring to counts of extreme market events, interpreting them on the basis of our theoretical insights. We show that, with the intermittent time series of our case study, the concordant-shift effect on the mean is more common than the combination effect. We moreover report a major increase in accuracy for probabilistic forecasts after reconciliation, confirming the beneficial effect of reconciliation for forecasting intermittent time series (Athanasopoulos et al., 2017; Corani et al., 2023; Kourentzes & Athanasopoulos, 2021; Zambon et al., 2024).

The paper is organized as follows. In Section 2, we recall probabilistic reconciliation through conditioning, and analyze the reconciled mean and variance of the reconciled distribution in Gaussian and non-Gaussian cases. In Section 3, we present our case study. Finally, the conclusions are in Section 4.

2. Probabilistic forecast reconciliation

Given a hierarchy, we denote by $\mathbf{b} = [b_1, \dots, b_{n_b}]^T$ the vector of bottom variables, and by $\mathbf{u} = [u_1, \dots, u_{n_u}]^T$ the vector of upper variables. To keep the notation simple, we do not show the time index; it is thus understood that all forecasts refer to the time $t+h$. We denote the vector of all the variables by

$$\mathbf{y} = \begin{bmatrix} \mathbf{u} \\ \mathbf{b} \end{bmatrix} \in \mathbb{R}^n.$$

The hierarchy may be expressed as a set of linear constraints:

$$\mathbf{y} = \mathbf{S}\mathbf{b}, \quad \text{with} \quad \mathbf{S} = \left[\begin{array}{c} \mathbf{A} \\ \hline \mathbf{I} \end{array}\right], \tag{1}$$

where $\mathbf{I} \in \mathbb{R}^{n_b \times n_b}$ is the identity matrix. $\mathbf{S} \in \mathbb{R}^{n \times n_b}$ is the *summing matrix*, while $\mathbf{A} \in \mathbb{R}^{n_u \times n_b}$ is the *aggregating matrix*. The summing constraints can thus be written as $\mathbf{u} = \mathbf{A}\mathbf{b}$.

We assume the base forecasts to be in the form of predictive distributions. We denote by $\hat{\pi}$ the base forecast distribution for the entire hierarchy, and by $\hat{\pi}_U$ and $\hat{\pi}_B$ the base forecasts for the upper and the bottom variables, respectively. The aim of probabilistic reconciliation is to find a reconciled distribution $\tilde{\pi}$ that gives positive probability only to coherent points. To this end, we first obtain a reconciled bottom distribution $\tilde{\pi}_B$ from the base forecast distribution $\hat{\pi}$. Then, we obtain the reconciled distribution $\tilde{\pi}$ on the entire hierarchy as follows:

$$\tilde{\pi}(\mathbf{u}, \mathbf{b}) = \begin{cases} \tilde{\pi}_B(\mathbf{b}) & \text{if } \mathbf{u} = \mathbf{A}\mathbf{b} \\ 0 & \text{if } \mathbf{u} \neq \mathbf{A}\mathbf{b}, \end{cases}$$

so that the probability of any set of incoherent points is zero.

Probabilistic bottom-up. If we set $\tilde{\pi}_B = \hat{\pi}_B$, we have the probabilistic bottom-up, which ignores the base forecast distribution $\hat{\pi}_U$ of the upper variables of the hierarchy. The reconciled bottom-up distribution $\tilde{\pi}_{bu}$ is thus given by

$$\tilde{\pi}_{bu}(\mathbf{u}, \mathbf{b}) = \begin{cases} \hat{\pi}_B(\mathbf{b}) & \text{if } \mathbf{u} = \mathbf{A}\mathbf{b} \\ 0 & \text{if } \mathbf{u} \neq \mathbf{A}\mathbf{b}. \end{cases} \tag{2}$$

Reconciliation through conditioning. In this work, we apply reconciliation via conditioning. In order to take into account the base forecasts of all variables, we introduce the random vector

$$\hat{\mathbf{Y}} = \begin{bmatrix} \hat{\mathbf{U}} \\ \hat{\mathbf{B}} \end{bmatrix} \sim \hat{\pi},$$

so that $\hat{\pi}_U$ and $\hat{\pi}_B$ are the distributions of $\hat{\mathbf{U}}$ and $\hat{\mathbf{B}}$, respectively. We then define the reconciled bottom distribution by conditioning on the hierarchy constraints. For discrete distributions, we have

$$\begin{aligned} \tilde{\pi}_B(\mathbf{b}) &:= Prob\left(\hat{\mathbf{B}} = \mathbf{b} \mid \hat{\mathbf{U}} - \mathbf{A}\hat{\mathbf{B}} = 0\right) \\ &= \frac{Prob\left(\hat{\mathbf{B}} = \mathbf{b},\ \hat{\mathbf{U}} - \mathbf{A}\hat{\mathbf{B}} = 0\right)}{Prob\left(\hat{\mathbf{U}} - \mathbf{A}\hat{\mathbf{B}} = 0\right)} \\ &= \frac{Prob\left(\hat{\mathbf{B}} = \mathbf{b},\ \hat{\mathbf{U}} = \mathbf{A}\mathbf{b}\right)}{\sum_{\mathbf{x} \in \mathbb{R}^m} Prob\left(\hat{\mathbf{B}} = \mathbf{x},\ \hat{\mathbf{U}} = \mathbf{A}\mathbf{x}\right)} \\ &= \frac{\hat{\pi}(\mathbf{A}\mathbf{b}, \mathbf{b})}{\sum_{\mathbf{x} \in \mathbb{R}^m} \hat{\pi}(\mathbf{A}\mathbf{x}, \mathbf{x})} \\ &\propto \hat{\pi}(\mathbf{A}\mathbf{b}, \mathbf{b}). \end{aligned} \tag{3}$$

The same formula, $\tilde{\pi}_B(\mathbf{b}) \propto \hat{\pi}(\mathbf{A}\mathbf{b}, \mathbf{b})$, also holds for continuous distributions. We refer to Zambon et al. (2024) for the derivation in the continuous case.

2.1. Gaussian reconciliation

In the Gaussian case, reconciliation via conditioning can be solved in closed form (Corani et al., 2020). Its reconciled mean and variance are numerically equivalent to those of MinT (Wickramasuriya et al., 2019), despite the different derivation.

In Appendix A, we derive in a novel way the reconciliation formulae. They are equivalent to those of Corani et al. (2020) and Wickramasuriya et al. (2019) but more convenient to prove some properties. Let us assume that the base forecasts for the entire hierarchy are multivariate Gaussian:

$$\hat{\mathbf{Y}} = \begin{bmatrix} \hat{\mathbf{U}} \\ \hat{\mathbf{B}} \end{bmatrix} \sim \mathcal{N}\left(\hat{\mathbf{y}},\ \hat{\Sigma}_Y\right), \tag{4}$$

where

$$\hat{\mathbf{y}} = \begin{bmatrix} \hat{\mathbf{u}} \\ \hat{\mathbf{b}} \end{bmatrix}, \quad \hat{\Sigma}_Y = \begin{bmatrix} \hat{\Sigma}_U & \hat{\Sigma}_{UB} \\ \hat{\Sigma}_{UB}^T & \hat{\Sigma}_B \end{bmatrix}.$$

In Appendix A, we show that this is equivalent to the framework of Corani et al. (2020). The covariance matrix of the base forecasts $\widehat{\Sigma}_Y$ can thus be computed in practice as the covariance matrix of the forecasting errors. Assuming $\widehat{\Sigma}_Y$ to be positive definite, the reconciled bottom and upper distributions are multivariate Gaussian:

$$\widetilde{\mathbf{B}} \sim \mathcal{N}\left(\widetilde{\mathbf{b}}, \widetilde{\Sigma}_B\right), \qquad \widetilde{\mathbf{U}} \sim \mathcal{N}\left(\widetilde{\mathbf{u}}, \widetilde{\Sigma}_U\right),$$

where

$$\widetilde{\mathbf{b}} = \widehat{\mathbf{b}} + \left(\widehat{\Sigma}_{UB}^T - \widehat{\Sigma}_B \mathbf{A}^T\right)\mathbf{Q}^{-1}(\mathbf{A}\widehat{\mathbf{b}} - \widehat{\mathbf{u}}), \tag{5}$$

$$\widetilde{\mathbf{u}} = \widehat{\mathbf{u}} + \left(\widehat{\Sigma}_U - \widehat{\Sigma}_{UB} \mathbf{A}^T\right)\mathbf{Q}^{-1}(\mathbf{A}\widehat{\mathbf{b}} - \widehat{\mathbf{u}}), \tag{6}$$

$$\widetilde{\Sigma}_B = \widehat{\Sigma}_B - \left(\widehat{\Sigma}_{UB}^T - \widehat{\Sigma}_B \mathbf{A}^T\right)\mathbf{Q}^{-1}\left(\widehat{\Sigma}_{UB}^T - \widehat{\Sigma}_B \mathbf{A}^T\right)^T, \tag{7}$$

$$\widetilde{\Sigma}_U = \widehat{\Sigma}_U - \left(\widehat{\Sigma}_U - \widehat{\Sigma}_{UB} \mathbf{A}^T\right)\mathbf{Q}^{-1}\left(\widehat{\Sigma}_U - \widehat{\Sigma}_{UB} \mathbf{A}^T\right)^T, \tag{8}$$

and $\mathbf{Q} := \widehat{\Sigma}_U - \widehat{\Sigma}_{UB}\mathbf{A}^T - \mathbf{A}\widehat{\Sigma}_{UB}^T + \mathbf{A}\widehat{\Sigma}_B\mathbf{A}^T$.

Proposition 1 (*Reconciled Gaussian Variance*)**.** *In the Gaussian framework, the variance of each variable decreases after reconciliation. Indeed, for each $i = 1, \ldots, m$ and $j = 1, \ldots, n-m$,*

$$Var(\widetilde{B}_i) \leq Var(\widehat{B}_i),$$
$$Var(\widetilde{U}_j) \leq Var(\widehat{U}_j). \tag{9}$$

The proof is given in Appendix B. We remark that the variance of each variable decreases after reconciliation, regardless of the amount of incoherence. While we consider reconciliation via conditioning, (Wickramasuriya, 2021) provides a similar result for the minT reconciliation. In Wickramasuriya (2021), no Gaussian assumption is made; however, this assumes the unbiasedness of the base forecast, which we do not assume.

Reconciled Gaussian mean. In order to study the reconciled mean, we need some restrictive assumptions (which, however, do apply to the case study of Section 3). In particular, assuming that

- there is only one upper variable, and
- there is no correlation between the bottom and the upper base forecasts,

the reconciled upper mean is a convex combination of $\widehat{u}$ and $\mathbf{A}\widehat{\mathbf{b}}$. Indeed, from (6), we have

$$\widetilde{u} = \frac{\widehat{\sigma}_{bu}^2}{\widehat{\sigma}_U^2 + \widehat{\sigma}_{bu}^2}\,\widehat{u} + \frac{\widehat{\sigma}_U^2}{\widehat{\sigma}_U^2 + \widehat{\sigma}_{bu}^2}\,\mathbf{A}\widehat{\mathbf{b}}, \tag{10}$$

where $\widehat{\sigma}_U^2$ is the variance of the upper base forecast, and $\widehat{\sigma}_{bu}^2 := \mathbf{A}\widehat{\Sigma}_B\mathbf{A}^T \geq 0$ is the variance of the probabilistic bottom-up, defined in (2). The reconciled mean is thus a combination of the base and the bottom-up mean, as already observed (Corani et al., 2020; Hollyman et al., 2021). We call this the *combination effect*.

We can draw an analogy with the Gaussian conjugate model in Bayesian statistics: the posterior variance is always smaller than the prior variance (Gelman, 2011), and the posterior expectation is a convex combination of the prior expectation and the sample mean (Gelman et al., 2013, Ch. 2.5).

2.2. Sampling from the reconciled distribution

In the non-Gaussian case, the reconciled distribution $\widetilde{\pi}$ is in general not available in a parametric form; hence, we need to draw samples from $\widetilde{\pi}$. We follow the approach of Zambon et al. (2024) based on importance sampling (Elvira & Martino, 2021).

Denoting by N the number of samples, the algorithm is as follows:

1. Sample $\left(\widehat{\mathbf{b}}^{(i)}\right)_{i=1,\ldots,N}$ from $\widehat{\pi}_B$.
2. Compute the unnormalized weights $\left(\breve{w}^{(i)}\right)_{i=1,\ldots,N}$ as
$$\breve{w}^{(i)} = \frac{\widehat{\pi}\left(\mathbf{A}\widehat{\mathbf{b}}^{(i)}, \widehat{\mathbf{b}}^{(i)}\right)}{\widehat{\pi}_B\left(\widehat{\mathbf{b}}^{(i)}\right)}.$$
If we assume the bottom and upper base forecasts to be independent, as we do in this paper, we have
$$\breve{w}^{(i)} = \frac{\widehat{\pi}_U\left(\mathbf{A}\widehat{\mathbf{b}}^{(i)}\right)\widehat{\pi}_B\left(\widehat{\mathbf{b}}^{(i)}\right)}{\widehat{\pi}_B\left(\widehat{\mathbf{b}}^{(i)}\right)} = \widehat{\pi}_U\left(\mathbf{A}\widehat{\mathbf{b}}^{(i)}\right).$$
3. Compute the normalized weights as $w^{(i)} = \breve{w}^{(i)} / \sum_h \breve{w}^{(h)}$.
4. Sample $\left(\widetilde{\mathbf{b}}^{(i)}\right)_i$ with replacement from the weighted sample $\left(\widehat{\mathbf{b}}^{(i)}, w^{(i)}\right)_{i=1,\ldots,N}$.

The output $\left(\widetilde{\mathbf{b}}^{(i)}\right)_i$ is an unweighted sample from the reconciled distribution $\widetilde{\pi}_B$. The algorithm assumes the base forecasts of the upper and the bottom variables to be conditionally independent, i.e., to be independent given the observations available up to time t.

Computationally, this algorithm is suitable for the reconciliation of small hierarchies such those considered in this paper (a single upper variable). Larger hierarchies can be reconciled using an extension of this algorithm, called bottom-up importance sampling (BUIS, Zambon et al., 2024). The BUIS algorithm is implemented in the *R* package *bayesRecon* (Azzimonti et al., 2023).

2.3. Reconciled variance of discrete variables

In the case of discrete variables, the variance of the reconciled distribution of the bottom variables can be larger than the variance of the base distribution. This is a major difference from Gaussian reconciliation. In particular, this happens if the probability $p_c := Prob\left(\widehat{\mathbf{U}} = \mathbf{A}\widehat{\mathbf{B}}\right)$ of the base forecasts being coherent is small.

Proposition 2. *Let us assume $\widehat{\mathbf{B}}$ and $\widehat{\mathbf{U}}$ to be discrete random variables with $p_c > 0$. Then, for any $j = 1, \ldots, m$,*

$$Var\left[\widetilde{B}_j\right] = \frac{Var\left(\widehat{B}_j\right) - (1-p_c)\,Var\left[\widehat{B}_j \mid \widehat{\mathbf{U}} \neq \mathbf{A}\widehat{\mathbf{B}}\right] - p_c(1-p_c)(a-b)^2}{p_c}, \tag{11}$$

where $a := \mathbb{E}\left[\widehat{B}_j \mid \widehat{\mathbf{U}} \neq \mathbf{A}\widehat{\mathbf{B}}\right]$ and $b := \mathbb{E}\left[\widehat{B}_j \mid \widehat{\mathbf{U}} = \mathbf{A}\widehat{\mathbf{B}}\right]$.

The proof is in Appendix C. According to Eq. (11), a small p_c can result in a large variance of the reconciled bottom distributions. We can draw yet another analogy

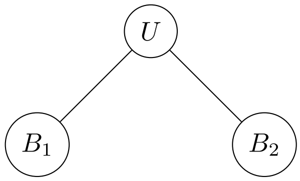

Fig. 1. A minimal hierarchy.

Table 1
Reconciliations that increase the variance of the bottom variables (grey background). For the Bernoulli case, we set $\widehat{p}_1 = 0.3$, $\widehat{p}_2 = 0.2$, and $\widehat{\mathbf{q}} = [0.1, 0.2, 0.7]$, and after reconciliation, we obtain $\widetilde{p}_1 = 0.52$, $\widetilde{p}_2 = 0.40$, and $\widetilde{\mathbf{q}} \approx [0.32, 0.44, 0.24]$.

	Mean		Variance		
	base	*reconc*	*base*	*reconc*	Δ
Bernoulli					
B_1	0.3	0.52	0.21	0.25	0.04
B_2	0.2	0.40	0.16	0.24	0.08
U	1.6	0.92	0.44	0.56	0.12
Poisson					
B_1	0.5	0.97	0.5	0.81	0.31
B_2	0.8	1.56	0.8	1.13	0.33
U	6.0	2.53	6.0	1.41	−4.59

with Bayesian statistics: in non-Gaussian models, the posterior variance can be larger than the prior variance if we condition on observations that are conflicting with the prior beliefs (Gelman et al., 2013, Ch. 2.2; Gelman, 2011).

Reconciling a minimal hierarchy. We now illustrate the increase of variance after reconciliation on the minimal hierarchy of Fig. 1. Consider the following independent base forecasts:

$$\widehat{B}_1 \sim \text{Bernoulli}(\widehat{p}_1), \quad \widehat{B}_2 \sim \text{Bernoulli}(\widehat{p}_2),$$

$$\widehat{U} = \begin{cases} 0 & \text{prob} = \widehat{q}_0 \\ 1 & \text{prob} = \widehat{q}_1 \\ 2 & \text{prob} = \widehat{q}_2, \end{cases}$$

where $\widehat{p}_1, \widehat{p}_2 \in [0, 1]$, $\widehat{q}_0, \widehat{q}_1, \widehat{q}_2 \in [0, 1]$, and $\widehat{q}_0 + \widehat{q}_1 + \widehat{q}_2 = 1$. In Appendix D, we analytically derive the expression of the parameters $\widetilde{p}_1, \widetilde{p}_2$, and $\widetilde{\mathbf{q}}$ of the reconciled distribution.

We set $\widehat{p}_1 = 0.3$, $\widehat{p}_2 = 0.2$, and $\widehat{\mathbf{q}} = [0.1, 0.2, 0.7]$, which induce large incoherence. The probability of coherence is $p_c = 0.17$, computed as follows:

$$p_c = \sum_{\substack{u, b_1, b_2: \\ u = b_1 + b_2}} \widehat{\pi}(u, b_1, b_2). \tag{12}$$

In this case, the variance of all variables increases after reconciliation (Table 1).

Poisson base forecast. We now consider the case of Poisson independent base forecasts:

$$\widehat{B}_1 \sim Poi(\widehat{\lambda}_1), \quad \widehat{B}_2 \sim Poi(\widehat{\lambda}_2), \quad \widehat{U} \sim Poi(\widehat{\lambda}_u),$$

with $\widehat{\lambda}_1, \widehat{\lambda}_2, \widehat{\lambda}_u > 0$. We obtain a small probability of coherence by setting $\widehat{\lambda}_1 = 0.5$, $\widehat{\lambda}_2 = 0.8$, and $\widehat{\lambda}_u = 6.0$, which results in $p_c = 0.03$. We perform reconciliation via importance sampling (Section 2.2). The variance of the bottom variables, computed using samples, increases after reconciliation (Table 1).

2.4. Reconciled mean in the non-Gaussian case

In some cases, the reconciled mean of the upper variable can be lower than *both* the mean of its base forecast and the bottom-up mean. We call this the *concordant-shift effect*. This happens when reconciliation decreases the variance of base forecasts that are right-skewed, which is often the case with count variables: the right tail is shortened, lowering the expected values.

To show the concordant-shift effect on the minimal hierarchy (Fig. 1), we consider independent Poisson base forecasts with $\widehat{\lambda}_1 = 0.5$, $\widehat{\lambda}_2 = 0.8$, and $\widehat{\lambda}_u = 1.5$. Thus, all base forecasts convey information of low counts. Reconciliation fuses the base forecasts, emphasizing the tendency towards zero: the mean of all the variables decreases after reconciliation (Table 2). The reconciled distribution of the upper variable has a lower mean than both the base and bottom-up distributions (Fig. 2(a)). In contrast, we can induce the combination effect (Fig. 2(b)) by setting $\widehat{\lambda}_1 = 5$, $\widehat{\lambda}_2 = 7$, and $\widehat{\lambda}_u = 18$: in this case, the base forecasts are less skewed, and p_c is smaller (Table 2).

3. Case study: Modeling extreme market events

Credit default swaps (CDSs) are financial instruments that guarantee insurance against the possible default of a given company (called a reference company) to the buyer. The CDS price is a function of the probability of default estimated by the market for that company. Thus, a high value of the CDS spread corresponds to an increase in the risk of a company default.

Following (Raunig & Scheicher, 2011), an extreme market event takes place when the value of the CDS spread on a given day exceeds the 90-th percentile of its distribution in the last trading year. In particular, we forecast extreme market events for the companies of the Euro Stoxx 50 index (https://www.stoxx.com/) in the period 2005–2018, which includes 3508 trading days. As in Agosto et al. (2020), we consider the 29 companies included in the index having a regularly quoted CDS, and we divide them into five economic sectors: Financial (FIN), Information and Communication Technology (ICT), Manufacturing (MFG), Energy (ENG), and Trade (TRD). Following (Agosto, 2022), we start from the CDS spread time series retrieved from Bloomberg, and we count the daily number of extreme events for each sector, obtaining five daily time series. The series include 3508 data points each; they have low counts and a high frequency of zeros (Table 3). They are all intermittent, with large average inter-demand intervals (ADIs). The count time series data are available in the *R* package *bayesRecon* (Azzimonti et al., 2023).

Base forecasts and hierarchical structure. We organize the time series into a hierarchy with five bottom and one upper time series: the five economic sectors constitute the bottom level, while their sum constitutes the top level (Fig. 3).

We compute the base forecasts for the bottom time series (counts in the different sectors) using the model of Agosto (2022), whose predictive distribution is a multivariate negative binomial with a static vector of dispersion parameters and a time-varying vector of location parameters following score-driven dynamics (Harvey, 2013).

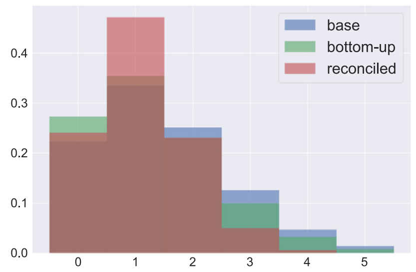

(a) Concordant-shift effect.

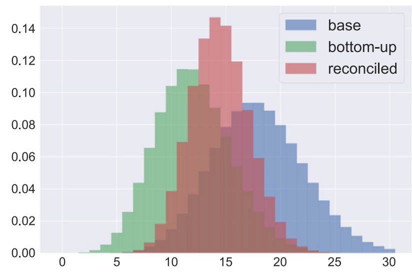

(b) Combination effect.

Fig. 2. Base, bottom-up, and reconciled probability mass functions of U in the case of Poisson base forecasts: (a) $\hat{\lambda}_1 = 0.5$, $\hat{\lambda}_2 = 0.8$, $\hat{\lambda}_u = 1.5$; (b) $\hat{\lambda}_1 = 5$, $\hat{\lambda}_2 = 7$, $\hat{\lambda}_u = 18$.

Table 2
Mean before and after reconciliation; the base forecasts are Poisson.

	Concordant-shift effect $p_c = 0.25$			**Combination effect** $p_c = 0.04$		
	base mean	*rec. mean*	Δ	*base mean*	*rec. mean*	Δ
B_1	0.50	0.43	−0.07	5.00	6.02	1.02
B_2	0.80	0.68	−0.12	7.00	8.43	1.43
U	1.50	1.11	−0.39	18.0	14.44	−3.56

Table 3
Main characteristics of the count time series. A time series is considered intermittent if its ADI is >1.32 (Syntetos & Boylan, 2005).

Sector	Time series	Mean	Proportion of zeros	ADI
FIN	10	0.96	0.73	3.8
ICT	4	0.38	0.79	4.7
MFG	7	0.67	0.73	3.7
ENG	5	0.48	0.80	4.9
TRD	3	0.29	0.81	5.2

The model of Agosto (2022) extends to the multivariate case the modeling approach of Agosto et al. (2016), who proposed a Poisson autoregressive model with exogenous covariates (PARX) to measure systemic risk in corporate default dynamics. According to Agosto (2022), the predictive distribution computed at time t for the count of time $t+1$ in sector i is as follows:

$$\hat{y}_{i,t+1} \sim NB(\mu_{i,t+1}, \alpha_i), \tag{13}$$

where $\boldsymbol{\mu}_t$ is the $k \times 1$ vector of location parameters, k is the number of sectors, and $\alpha_i \geq 0$. The model assumes the following dynamics:

$$\log(\boldsymbol{\mu}_{t+1}) = \mathbf{C} + \mathbf{D}\log\boldsymbol{\mu}_t + \mathbf{E}\frac{\mathbf{y}_t - \boldsymbol{\mu}_t}{\boldsymbol{\alpha}^T\boldsymbol{\mu}_t + 1}, \tag{14}$$

where $\mathbf{C}$ is a $k \times 1$ vector, while $\mathbf{D}$ and $\mathbf{E}$ are $k \times k$ matrices (see Agosto, 2022 for detailed properties and estimation details). Thus, the predicted event count in a given sector is a function of the past expectations ($\boldsymbol{\mu}_t$) and forecast errors ($\mathbf{y}_t - \boldsymbol{\mu}_t$) in the same sector and in the other sectors. The base forecasts for the upper time series are computed by fitting a univariate version of the model (Blasques et al., 2023) on the aggregate count time series.

The base forecasts for the bottom series measure financial risks at the sector level, accounting for the dependencies between sectors, through a shock propagation mechanism, in addition to an autoregressive component. The base upper forecasts express instead a measure of aggregate systemic financial risk.

3.1. Experimental procedure

As in Agosto (2022), we estimate the model parameters through in-sample maximum likelihood. We then compute the one-day-ahead base forecasts for time $t+1$ by conditioning the model on the counts observed up to time t.

We reconcile the in-sample base forecasts using importance sampling (Section 2.2). We perform 3508 reconciliations, each time drawing $N = 100,000$ samples from the reconciled distribution; each reconciliation is almost instantaneous (<0.1 s). We make available the code to reproduce the results as a vignette of the *R* package *bayesRecon* (Azzimonti et al., 2023).

As recommended by Panagiotelis et al. (2023), we compare base and reconciled distributions through the energy score (ES, Székely & Rizzo, 2013):

$$ES(P_t, \mathbf{y}_t) = \mathbb{E}_{P_t}\left[\|\mathbf{y}_t - \mathbf{s}_t\|^{\beta}\right] - \frac{1}{2}\mathbb{E}_{P_t}\left[\|\mathbf{s}_t - \mathbf{s}_t'\|^{\beta}\right],$$

where P_t is the forecast distribution on the whole hierarchy, $\mathbf{s}_t$ and $\mathbf{s}_t'$ are a pair of independent random variables distributed as P_t, and $\mathbf{y}_t$ is the vector of the actual values of all the time series at time t. We compute the ES, with $\beta = 2$, using the sampling approach of Wickramasuriya (2023). We compute the ES of the joint predictive distribution of the upper and bottom time series; we thus have a single ES for the entire hierarchy.

Moreover, we evaluate the prediction intervals using the interval score (IS, Gneiting & Raftery, 2007):

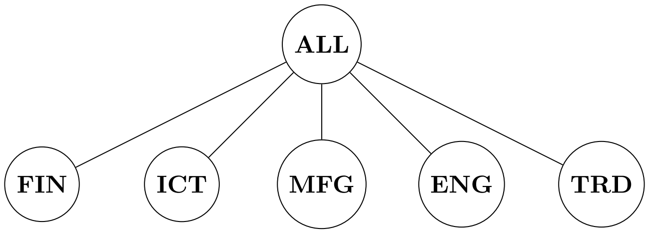

Fig. 3. Hierarchical structure of the extreme events time series.

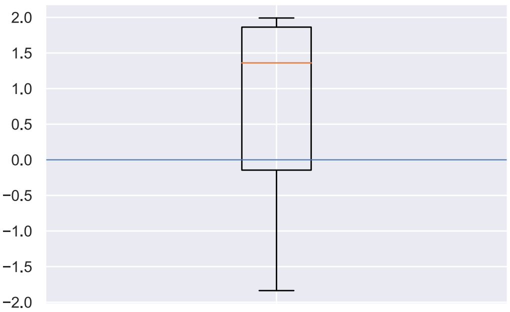

Fig. 4. Boxplot of the skill scores on ES, over 3508 reconciliations.

Table 4
Average skill scores for the different time series. Positive values indicate an improvement of the reconciled forecasts over the base forecasts. The ES is computed with respect to the joint distribution on the entire hierarchy.

	ALL	*FIN*	*ICT*	*MFG*	*ENG*	*TRD*
ES	0.89					
IS	0.87	1.15	0.20	1.07	0.22	0.18
SE	0.82	1.10	1.11	1.07	1.12	1.11
AE	−0.02	−0.02	−0.02	−0.02	−0.02	−0.02

Table 5
Width of the 90% prediction interval, averaged over 3508 reconciliations.

	ALL	*FIN*	*ICT*	*MFG*	*ENG*	*TRD*
base	6.9	3.2	1.1	1.9	1.3	0.9
reconc.	3.3	2.1	0.9	1.2	1.0	0.8

Table 6
Coverage of the 90% prediction intervals, assessed on 3508 reconciliations.

	ALL	*FIN*	*ICT*	*MFG*	*ENG*	*TRD*
base	96%	98%	98%	98%	98%	99%
reconc.	91%	95%	97%	97%	97%	98%

$$\mathrm{IS}_{\alpha}(l_t, u_t; y_t) = (u_t - l_t) + \frac{2}{\alpha}(l_t - y_t)\mathbb{1}(y_t < l_t) + \frac{2}{\alpha}(y_t - u_t)\mathbb{1}(y_t > u_t),$$

where $\alpha \in (0, 1)$, l_t and u_t are the lower and upper bounds of the $(1-\alpha) \times 100\%$ prediction interval, and y_t is the actual value of the time series at time t. We use $\alpha = 0.1$; i.e. we score prediction intervals whose nominal coverage is 90%.

Finally, we evaluate the point forecasts by measuring the squared error (SE) and the absolute error (AE):

$$\mathrm{SE} = \left(y_t - \hat{y}_{t|t-1}\right)^2,$$
$$\mathrm{AE} = \left|y_t - \hat{y}_{t|t-1}\right|,$$

where $\hat{y}_{t|t-1}$ denotes the optimal point forecast. The optimal point forecast depends on the error measure: it is the median of the predictive distribution for AE, and the expected value for SE (Kolassa, 2016, 2020).

Skill score. The skill score measures the improvement of the reconciled forecasts over the base forecasts. For example, the skill score for AE is defined as

$$\mathrm{skill}(reconc, base) = \frac{\mathrm{AE}(base) - \mathrm{AE}(reconc)}{\left(\mathrm{AE}(base) + \mathrm{AE}(reconc)\right)/2}. \tag{15}$$

For all indicators, a positive skill score implies an improvement of the reconciled forecast compared to the base forecasts, and vice versa. The skill score defined in (15) is symmetric, allowing us to fairly compare base and reconciled forecasts. For instance, a skill score of 1 implies that the loss function has been reduced by three times: $(3-1)/((3+1)/2) = 2/2 = 1$. Analogously, a skill score of -1 implies a three-fold worsening of the loss function. Moreover, the skill score is bounded between -2 and 2. Di Fonzo and Girolimetto (2022) compare competing approaches by computing the geometric mean of the indicators. However, this is not suitable for count time series, where the SE and the AE are often zero.

In Table 4, we report the skill scores averaged over the 3508 reconciliations. Reconciliation largely improves the ES (with an average skill score of 0.89) and the IS (average skill score ranging between 0.2 and 1.1, depending on the chosen time series). The boxplot of the skill scores on ES on each day (Fig. 4) confirms the improvement due to reconciliation. As a further insight, reconciliation reduces by 15%–50% the width of the prediction intervals (Table 5) without compromising their coverage (Table 6). We observe large skill scores (0.8–1) on the SE. On the other hand, the skill score on AE is close to 0. Indeed, often the median of the predictive distribution is already 0, and it does not change after reconciliation.

Hence, in our experiments, reconciliation yields a major improvement over the base forecasts, confirming the positive effect of probabilistic reconciliation (Corani et al., 2023; Zambon et al., 2024) and point forecast reconciliation (Kourentzes & Athanasopoulos, 2021) for intermittent time series.

Table 7
Mean absolute error for the upper time series "ALL".

	optimal	*coherent*	*base*
ALL	1.10	1.24	1.33

Coherent vs. optimal point forecast for the upper time series. Even if we have a reconciled distribution, only the SE-optimal point forecasts are coherent. The AE-optimal point forecasts, i.e., the medians of the reconciled distributions, are instead generally incoherent (Kolassa, 2022).

Hence, besides optimal point forecasts, we evaluate coherent point forecasts: we take the sum of the medians of the reconciled bottom distributions and use it as upper point forecast. We compute the mean absolute error for the upper time series obtained using the AE-optimal, the coherent, and the base point forecasts (Table 7). As expected, the AE worsens by imposing coherence rather than optimality, yet it remains better than the AE of the base forecasts.

3.2. Analysis of the reconciled mean and variance

In most reconciliations (96%), we observe the concordant-shift effect, i.e., the reconciled upper mean is lower than both the bottom-up and the base upper mean. In Fig. 5(a), we show an example. Reconciliation largely reduces the variance of "ALL" and "FIN", shortening the right tail of the distribution. Within the bottom time series, reconciliation mostly affects FIN, which is characterized by the largest counts and overdispersion. The predictive distribution for other time series, such as ICT (shown in the figure), is less affected by reconciliation, being characterized by lower counts and variability. The reduction of the variance shifts the expected values towards zero, since the base distributions have a positive skewness.

In Fig. 5(b), we show an example of the combination effect. The reconciled distribution of "ALL" is a combination between its base forecast and the probabilistic bottom-up. Reconciliation decreases the expected values of the bottom time series while increasing the expected value of the upper time series.

The variance of all the variables decreases in most cases (97%). However, in Fig. 6, we show an example in which the variance of the bottom variables increases after reconciliation because of the large incoherence of the base forecasts. The forecast of the upper time series, which is characterized by large uncertainty, is instead sharply shifted towards smaller values.

4. Conclusions

We proved that reconciliation via conditioning decreases the variance of all variables in Gaussian frameworks, while it can increase or decrease the variance of discrete variables, depending on the incoherence of the base forecasts. We also discussed two different effects that reconciliation can have on the upper mean. The empirical analysis of time series of extreme market events confirmed our theoretical insights. We leave as future research the study of the theoretical properties of the reconciliation of other continuous non-Gaussian distributions, including skewed ones.

Declaration of competing interest

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

Acknowledgments

This work was partially funded by the Swiss National Science Foundation (grant 200021_212164/1), by the Hasler foundation, Switzerland (project 23057), and by the European Union's Horizon 2020 Research and Innovation Programme (grant 101016233).

Appendix A. Reconciled distribution in the gaussian case

Bottom distribution. Let us define $\mathbf{T} \in \mathbb{R}^{n \times n}$ as

$$\mathbf{T} = \begin{bmatrix} \mathbf{0} & \mathbf{I}_m \\ \mathbf{I}_{n-m} & -\mathbf{A} \end{bmatrix},$$

and let $\mathbf{Z} := \mathbf{T}\widehat{\mathbf{Y}}$. Hence, $\mathbf{Z}$ is Gaussian:

$$\mathbf{Z} \sim \mathcal{N}\left(\mathbf{T}\hat{\mathbf{y}},\ \mathbf{T}\widehat{\Sigma}_Y\mathbf{T}^T\right). \tag{A.1}$$

We have

$$\mathbf{T}\hat{\mathbf{y}} = \begin{bmatrix} \hat{\mathbf{b}} \\ \hat{\mathbf{u}} - \mathbf{A}\hat{\mathbf{b}} \end{bmatrix},$$

$$\mathbf{T}\widehat{\Sigma}_Y\mathbf{T}^T = \begin{bmatrix} \widehat{\Sigma}_B & \widehat{\Sigma}_{UB}^T - \widehat{\Sigma}_B\mathbf{A}^T \\ \widehat{\Sigma}_{UB} - \mathbf{A}\widehat{\Sigma}_B & \mathbf{Q} \end{bmatrix}, \tag{A.2}$$

where $\mathbf{Q} = \widehat{\Sigma}_U - \widehat{\Sigma}_{UB}\mathbf{A}^T - \mathbf{A}\widehat{\Sigma}_{UB}^T + \mathbf{A}\widehat{\Sigma}_B\mathbf{A}^T$. Since

$$\mathbf{Z} = \begin{bmatrix} \widehat{\mathbf{B}} \\ \widehat{\mathbf{U}} - \mathbf{A}\widehat{\mathbf{B}} \end{bmatrix} =: \begin{bmatrix} \mathbf{Z}_1 \\ \mathbf{Z}_2 \end{bmatrix},$$

the reconciled bottom distribution is given by the conditional distribution of $\mathbf{Z}_1$ given $\mathbf{Z}_2 = 0$. Since the covariance matrix $\widehat{\Sigma}_Y$ is assumed to be positive definite, the covariance matrix $\mathbf{Q}$ of $\mathbf{Z}_2$ is also positive definite; hence, we have

$$\mathbf{Z}_1 \mid \mathbf{Z}_2 = 0 \sim \mathcal{N}\left(\widetilde{\mathbf{b}},\ \widetilde{\Sigma}_B\right),$$

where

$$\widetilde{\mathbf{b}} = \hat{\mathbf{b}} + \left(\widehat{\Sigma}_{UB}^T - \widehat{\Sigma}_B\mathbf{A}^T\right)\mathbf{Q}^{-1}(\mathbf{A}\hat{\mathbf{b}} - \hat{\mathbf{u}}), \tag{A.3}$$

$$\widetilde{\Sigma}_B = \widehat{\Sigma}_B - \left(\widehat{\Sigma}_{UB}^T - \widehat{\Sigma}_B\mathbf{A}^T\right)\mathbf{Q}^{-1}\left(\widehat{\Sigma}_{UB}^T - \widehat{\Sigma}_B\mathbf{A}^T\right)^T. \tag{A.4}$$

Upper distribution. Since $\widetilde{\mathbf{U}} = \mathbf{A}\widetilde{\mathbf{B}}$, we have that $\widetilde{\mathbf{U}} \sim \mathcal{N}\left(\widetilde{\mathbf{u}},\ \widetilde{\Sigma}_U\right)$, with

$$\widetilde{\mathbf{u}} = \mathbf{A}\widetilde{\mathbf{b}}, \qquad \widetilde{\Sigma}_U = \mathbf{A}\widetilde{\Sigma}_B\mathbf{A}^T. \tag{A.5}$$

If we define $\mathbf{D} := \widehat{\Sigma}_U - \widehat{\Sigma}_{UB}\mathbf{A}^T$, from (A.3) and (A.5), we have

$$\begin{aligned}\widetilde{\mathbf{u}} &= \mathbf{A}\hat{\mathbf{b}} + \left(\mathbf{A}\widehat{\Sigma}_{UB}^T - \mathbf{A}\widehat{\Sigma}_B\mathbf{A}^T\right)\mathbf{Q}^{-1}(\mathbf{A}\hat{\mathbf{b}} - \hat{\mathbf{u}}) \\ &= \mathbf{A}\hat{\mathbf{b}} + (\mathbf{D} - \mathbf{Q})\,\mathbf{Q}^{-1}(\mathbf{A}\hat{\mathbf{b}} - \hat{\mathbf{u}})\end{aligned}$$

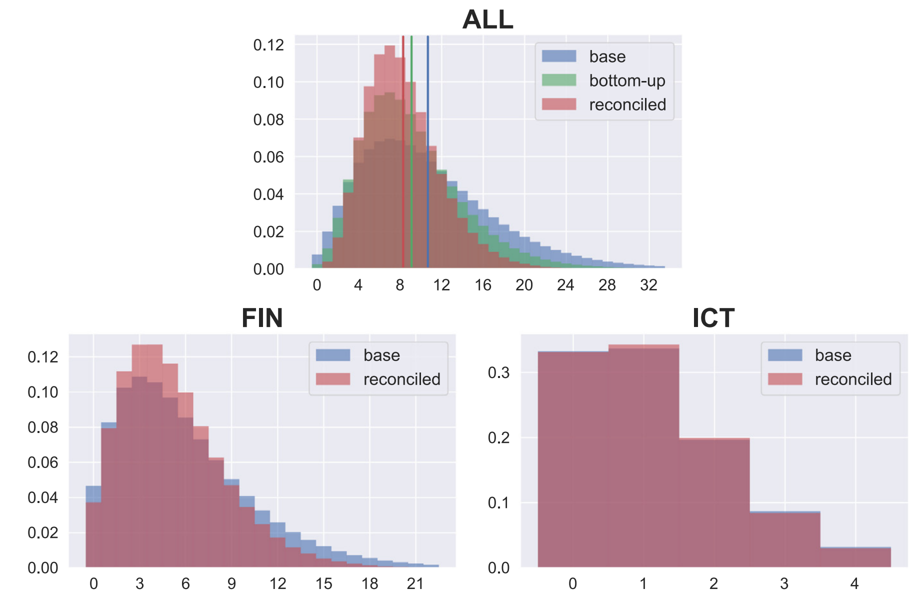

(a) Concordant-shift effect. The reduction of variance of the asymmetric distribution shortens the right tail and decreases the mean.

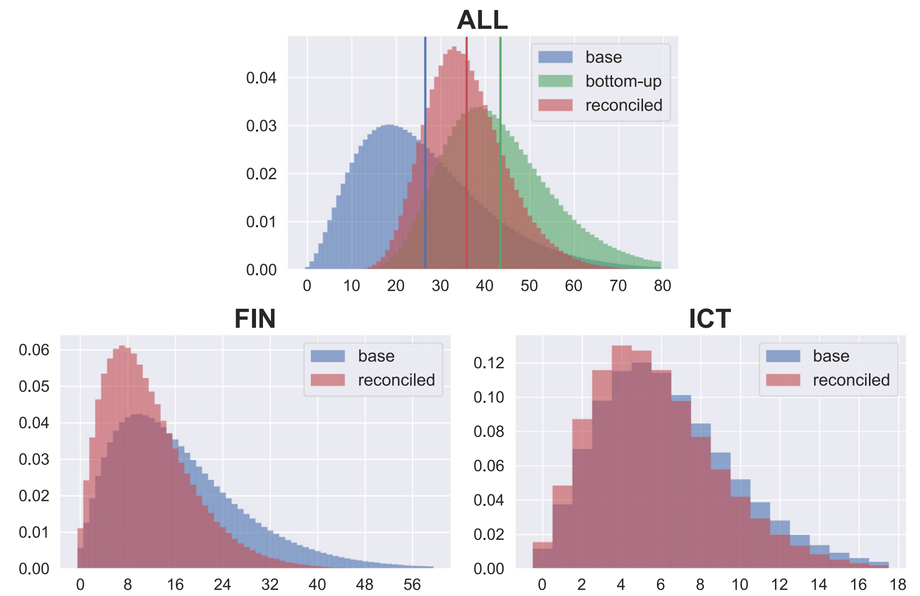

(b) Combination effect

Fig. 5. Concordant-shift effect vs. combination effect. We plot the base and reconciled probability mass functions for "ALL", "FIN", and "ICT". For "ALL", we also show the bottom-up pmf. The vertical lines correspond to the means of the distributions.

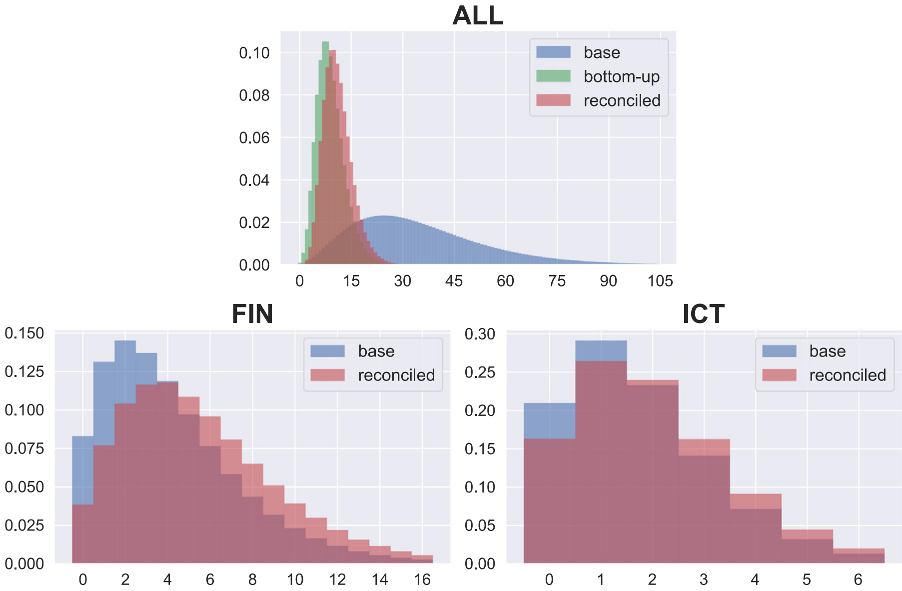

Fig. 6. An example of reconciliation which increases the variance of the bottom forecasts, due to the large incoherence of the base forecasts (note the x-axis of "ALL"). Variance of "FIN": 11.4 (base), 14.3 (reconciled). Variance of "ICT": 2.2 (base), 2.6 (reconciled).

$$= \mathbf{A}\hat{\mathbf{b}} + \mathbf{D}\mathbf{Q}^{-1}(\mathbf{A}\hat{\mathbf{b}} - \hat{\mathbf{u}}) - (\mathbf{A}\hat{\mathbf{b}} - \hat{\mathbf{u}})$$
$$= \hat{\mathbf{u}} + \left(\hat{\Sigma}_U - \hat{\Sigma}_{UB}\mathbf{A}^T\right)\mathbf{Q}^{-1}(\mathbf{A}\hat{\mathbf{b}} - \hat{\mathbf{u}}).$$

Moreover, from (A.4) and (A.5),

$$\begin{aligned}\widetilde{\Sigma}_U &= \mathbf{A}\hat{\Sigma}_B\mathbf{A}^T - \left(\mathbf{A}\hat{\Sigma}_{UB}^T - \mathbf{A}\hat{\Sigma}_B\mathbf{A}^T\right)\mathbf{Q}^{-1}\left(\mathbf{A}\hat{\Sigma}_{UB}^T - \mathbf{A}\hat{\Sigma}_B\mathbf{A}^T\right)^T\\ &= \mathbf{A}\hat{\Sigma}_B\mathbf{A}^T - (\mathbf{D} - \mathbf{Q})\,\mathbf{Q}^{-1}\left(\mathbf{D}^T - \mathbf{Q}\right)\\ &= \mathbf{A}\hat{\Sigma}_B\mathbf{A}^T - \mathbf{D}\mathbf{Q}^{-1}\mathbf{D}^T + \mathbf{D} + \mathbf{D}^T - \mathbf{Q}\\ &= \hat{\Sigma}_U - \left(\hat{\Sigma}_U - \hat{\Sigma}_{UB}\mathbf{A}^T\right)\mathbf{Q}^{-1}\left(\hat{\Sigma}_U - \hat{\Sigma}_{UB}\mathbf{A}^T\right)^T.\end{aligned}$$

Relation to Corani et al. (2020). In Corani et al. (2020), the reconciliation framework was defined in the following way:

- $\mathbf{B}' \sim \mathcal{N}\left(\hat{\mathbf{b}}, \hat{\Sigma}_B\right)$,
- $\hat{\mathbf{U}}' = \mathbf{A}\mathbf{B}' + \boldsymbol{\epsilon}^u$,

where the error term $\boldsymbol{\epsilon}^u$ was assumed to have zero mean and to be jointly Gaussian with $\mathbf{B}'$. The reconciled distribution on the bottom time series was then obtained via a Bayes' update by conditioning on $\hat{\mathbf{U}}' = \hat{\mathbf{u}}$, the point forecast for the upper time series.

In this paper, we adopt a slightly different framework. We assume that the base forecasts are jointly Gaussian:

$$\hat{\mathbf{Y}} = \begin{bmatrix}\hat{\mathbf{U}}\\ \hat{\mathbf{B}}\end{bmatrix} \sim \mathcal{N}\left(\hat{\mathbf{y}}, \hat{\Sigma}_Y\right).$$

Hence, we have $\hat{\mathbf{B}} \sim \mathcal{N}\left(\hat{\mathbf{b}}, \hat{\Sigma}_B\right)$ and $\hat{\mathbf{U}} \sim \mathcal{N}\left(\hat{\mathbf{u}}, \hat{\Sigma}_U\right)$. The reconciled distribution on the bottom time series is obtained by conditioning on $\hat{\mathbf{U}} = \mathbf{A}\hat{\mathbf{B}}$. This is equivalent to the framework of Corani et al. (2020), if we set $\hat{\mathbf{B}} = \mathbf{B}'$ and $\hat{\mathbf{U}} = \hat{\mathbf{u}} - \boldsymbol{\epsilon}^u$. The hypothesis of joint normality of $(\mathbf{B}', \boldsymbol{\epsilon}^u)$ is thus replaced by the hypothesis of joint normality of $(\hat{\mathbf{B}}, \hat{\mathbf{U}})$. Indeed, (A.3) and (A.4) are precisely Eqs. (11) and (12) in Corani et al. (2020), if we set

$$\hat{\Sigma}_{UB} = -k_h\mathbf{M}_1^T, \quad \hat{\Sigma}_B = k_h\hat{\Sigma}_{B,1}, \quad \hat{\Sigma}_U = k_h\hat{\Sigma}_{U,1}. \tag{A.6}$$

Note that in this paper, we only deal with one-step-ahead forecasts (i.e., $h = 1$); hence, $k_h = 1$ and can be omitted. Since the covariance matrices correspond to those of Corani et al. (2020), as shown in (A.6), we conclude that $\hat{\Sigma}_Y$ can be estimated as the covariance matrix of the residuals of the entire hierarchy.

Appendix B. Proof of Proposition 1

First, $\mathbf{Q}$ is positive definite, as it is the covariance matrix of $\hat{\mathbf{U}} - \mathbf{A}\hat{\mathbf{B}}$ (see Appendix A). Hence, $\mathbf{Q}^{-1}$ is also positive definite, and the matrices

$$\mathbf{G} := \left(\hat{\Sigma}_{UB}^T - \hat{\Sigma}_B\mathbf{A}^T\right)\mathbf{Q}^{-1}\left(\hat{\Sigma}_{UB}^T - \hat{\Sigma}_B\mathbf{A}^T\right)^T,$$
$$\mathbf{H} := \left(\hat{\Sigma}_U - \hat{\Sigma}_{UB}\mathbf{A}^T\right)\mathbf{Q}^{-1}\left(\hat{\Sigma}_U - \hat{\Sigma}_{UB}\mathbf{A}^T\right)^T,$$

are positive semi-definite.

From (7), we have that, for each $i = 1, \ldots, m$

$$\mathrm{Var}(\widetilde{B}_i) = \mathrm{Var}(\hat{B}_i) - G_{ii} \le \mathrm{Var}(\hat{B}_i),$$

as $G_{ii} \ge 0$ since the matrix $\mathbf{G}$ is positive semi-definite. Analogously, we have

$$\mathrm{Var}(\widetilde{U}_j) = \mathrm{Var}(\hat{U}_j) - H_{jj} \le \mathrm{Var}(\hat{U}_j),$$

for all $j = 1, \ldots, n - m$.

Appendix C. Proof of Proposition 2

Let us denote $Z := \mathbb{1}_{\{\hat{\mathbf{U}} = \mathbf{A}\hat{\mathbf{B}}\}}$, so that $Z = 1$ when the constraint is satisfied, and 0 otherwise. By the law of total variance (Weiss, 2005), for any $j = 1, \ldots, m$, we have

$$\mathrm{Var}\left(\hat{B}_j\right) = \mathbb{E}_Z\left[\mathrm{Var}\left(\hat{B}_j|Z\right)\right] + \mathrm{Var}_Z\left(\mathbb{E}\left[\hat{B}_j|Z\right]\right). \tag{C.1}$$

Since

$$\mathbb{E}\left[\hat{B}_j|Z\right] = \begin{cases}\mathbb{E}\left[\hat{B}_j|\hat{\mathbf{U}} = \mathbf{A}\hat{\mathbf{B}}\right] & \text{if } Z = 1\\ \mathbb{E}\left[\hat{B}_j|\hat{\mathbf{U}} \neq \mathbf{A}\hat{\mathbf{B}}\right] & \text{if } Z = 0,\end{cases}$$

we have that $\mathbb{E}[\widehat{B}_j|Z] = a+(b-a)Z$, where $a := \mathbb{E}[\widehat{B}_j|\widehat{\mathbf{U}} \neq \mathbf{A}\widehat{\mathbf{B}}]$ and $b := \mathbb{E}[\widehat{B}_j|\widehat{\mathbf{U}} = \mathbf{A}\widehat{\mathbf{B}}]$. Note that $Z \sim \text{Bernoulli}(p_c)$, with $p_c := Prob(\widehat{\mathbf{U}} = \mathbf{A}\widehat{\mathbf{B}})$. Hence,

$$\begin{aligned}\text{Var}_Z\left(\mathbb{E}\left[\widehat{B}_j|Z\right]\right) &= \text{Var}_Z\left(a+(b-a)Z\right)\\ &= (b-a)^2 p_c(1-p_c).\end{aligned} \tag{C.2}$$

Moreover, since

$$\text{Var}\left[\widehat{B}_j|Z\right] = \begin{cases}\text{Var}\left[\widehat{B}_j|\widehat{\mathbf{U}} = \mathbf{A}\widehat{\mathbf{B}}\right] & \text{if } Z=1\\ \text{Var}\left[\widehat{B}_j|\widehat{\mathbf{U}} \neq \mathbf{A}\widehat{\mathbf{B}}\right] & \text{if } Z=0,\end{cases}$$

we have

$$\begin{aligned}\mathbb{E}_Z\left[\text{Var}\left(\widehat{B}_j|Z\right)\right] &= p_c \,\text{Var}\left[\widehat{B}_j|\widehat{\mathbf{U}} = \mathbf{A}\widehat{\mathbf{B}}\right]\\ &\quad + (1-p_c)\,\text{Var}\left[\widehat{B}_j|\widehat{\mathbf{U}} \neq \mathbf{A}\widehat{\mathbf{B}}\right].\end{aligned} \tag{C.3}$$

From (C.1), (C.2), and (C.3), we obtain

$$\begin{aligned}\text{Var}\left(\widehat{B}_j\right) &= p_c \,\text{Var}\left[\widehat{B}_j|\widehat{\mathbf{U}} = \mathbf{A}\widehat{\mathbf{B}}\right] + (1-p_c)\,\text{Var}\left[\widehat{B}_j|\widehat{\mathbf{U}} \neq \mathbf{A}\widehat{\mathbf{B}}\right]\\ &\quad + p_c(1-p_c)(a-b)^2,\end{aligned}$$

from which

$$\begin{aligned}&\text{Var}\left[\widetilde{B}_j\right] = \text{Var}\left[\widehat{B}_j|\widehat{\mathbf{U}} = \mathbf{A}\widehat{\mathbf{B}}\right]\\ &= \frac{\text{Var}\left(\widehat{B}_j\right) - (1-p_c)\,\text{Var}\left[\widehat{B}_j|\widehat{\mathbf{U}} \neq \mathbf{A}\widehat{\mathbf{B}}\right] - p_c(1-p_c)(a-b)^2}{p_c}.\end{aligned}$$

Appendix D. Bernoulli example

Let us denote by $\widehat{\pi}_1$ and $\widehat{\pi}_2$ the probability mass functions of $\widehat{B}_1$ and $\widehat{B}_2$, so that $\widehat{\pi}_1(0) = 1-\widehat{p}_1$, $\widehat{\pi}_1(1) = \widehat{p}_1$, and $\widehat{\pi}_1(k) = 0$ for any $k \neq 0, 1$. The probability mass function $\widehat{\pi}_U$ of $\widehat{U}$ is defined as $\widehat{\pi}_U(0) = \widehat{q}_0$, $\widehat{\pi}_U(1) = \widehat{q}_1$, $\widehat{\pi}_U(2) = \widehat{q}_2$, and $\widehat{\pi}_U(k) = 0$ for any $k \neq 0, 1, 2$.

Since, from (3), the reconciled probability mass function $\widetilde{\pi}_B$ is given by

$$\widetilde{\pi}_B(b_1, b_2) \propto \widehat{\pi}_1(b_1)\widehat{\pi}_2(b_2)\widehat{\pi}_U(b_1+b_2),$$

the reconciled bottom distribution can be expressed as

$$(\widetilde{B}_1, \widetilde{B}_2) = \begin{cases}(0,0) & \text{prob} = (1-\widehat{p}_1)(1-\widehat{p}_2)\widehat{q}_0 \,/\, p_c\\ (1,0) & \text{prob} = \widehat{p}_1(1-\widehat{p}_2)\widehat{q}_1 \,/\, p_c\\ (0,1) & \text{prob} = (1-\widehat{p}_1)\widehat{p}_2\widehat{q}_1 \,/\, p_c\\ (1,1) & \text{prob} = \widehat{p}_1\widehat{p}_2\widehat{q}_2 \,/\, p_c,\end{cases}$$

where the normalizing constant $p_c := (1-\widehat{p}_1)(1-\widehat{p}_2)\widehat{q}_0 + \widehat{p}_1(1-\widehat{p}_2)\widehat{q}_1 + (1-\widehat{p}_1)\widehat{p}_2\widehat{q}_1 + \widehat{p}_1\widehat{p}_2\widehat{q}_2$ is the probability of coherence, defined in (12). Hence,

$$\widetilde{B}_1 \sim Bernoulli(\widetilde{p}_1), \quad \widetilde{B}_2 \sim Bernoulli(\widetilde{p}_2),$$

with

$$\begin{aligned}\widetilde{p}_1 &= \frac{[(1-\widehat{p}_2)\widehat{q}_1 + \widehat{p}_2\widehat{q}_2]\widehat{p}_1}{S},\\ \widetilde{p}_2 &= \frac{[(1-\widehat{p}_1)\widehat{q}_1 + \widehat{p}_1\widehat{q}_2]\widehat{p}_2}{S}.\end{aligned} \tag{D.1}$$

Finally,

$$\widetilde{U} = \begin{cases}0 & \text{prob} = (1-\widehat{p}_1)(1-\widehat{p}_2)\widehat{q}_0/S\\ 1 & \text{prob} = (\widehat{p}_1+\widehat{p}_2-2\widehat{p}_1\widehat{p}_2)\widehat{q}_1/S\\ 2 & \text{prob} = \widehat{p}_1\widehat{p}_2\widehat{q}_2/S.\end{cases} \tag{D.2}$$

References

Agosto, A. (2022). Multivariate score-driven models for count time series to assess financial contagion. http://dx.doi.org/10.2139/ssrn.4119895.
Agosto, A., Ahelegbey, D. F., & Giudici, P. (2020). Tree networks to assess financial contagion. *Economic Modelling*, *85*, 349–366.
Agosto, A., Cavaliere, G., Kristensen, D., & Rahbek, A. (2016). Modeling corporate defaults: Poisson autoregressions with exogenous covariates (PARX). *Journal of Empirical Finance*, *38*, 640–663.
Athanasopoulos, G., Hyndman, R. J., Kourentzes, N., & Petropoulos, F. (2017). Forecasting with temporal hierarchies. *European Journal of Operational Research*, *262*(1), 60–74.
Azzimonti, D., Rubattu, N., Zambon, L., & Corani, G. (2023). bayesRecon: Probabilistic reconciliation via conditioning. R package version 0.1.2.
Blasques, F., Holý, V., & Tomanová, P. (2023). Zero-inflated autoregressive conditional duration model for discrete trade durations with excessive zeros. *Studies in Nonlinear Dynamics & Econometrics*, http://dx.doi.org/10.1515/snde-2022-0008.
Corani, G., Azzimonti, D., Augusto, J. P., & Zaffalon, M. (2020). Probabilistic reconciliation of hierarchical forecast via Bayes' rule. *Vol. 3*, In *Proc. European conf. on machine learning and knowledge discovery in database ECML/PKDD* (pp. 211–226).
Corani, G., Azzimonti, D., & Rubattu, N. (2023). Probabilistic reconciliation of count time series. *International Journal of Forecasting*, http://dx.doi.org/10.1016/j.ijforecast.2023.04.003.
Di Fonzo, T., & Girolimetto, D. (2022). Forecast combination-based forecast reconciliation: Insights and extensions. *International Journal of Forecasting*, http://dx.doi.org/10.1016/j.ijforecast.2022.07.001.
Elvira, V., & Martino, L. (2021). Advances in importance sampling. *Wiley StatsRef-Statistics Reference Online*.
Gelman, A. (2011). Adding more information can make the variance go up (depending on your model). https://statmodeling.stat.columbia.edu/2011/07/20/adding_more_inf/.
Gelman, A., Carlin, J., Stern, H., Dunson, D., Vehtari, A., & Rubin, D. (2013). *Bayesian data analysis* (3rd ed.). Chapman and Hall/CRC.
Gneiting, T., & Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, *102*(477), 359–378.
Harvey, A. C. (2013). *Dynamic models for volatility and heavy tails: with applications to financial and economic time series*. New York: Cambridge University Press.
Hollyman, R., Petropoulos, F., & Tipping, M. E. (2021). Understanding forecast reconciliation. *European Journal of Operational Research*, *294*(1), 149–160.
Hollyman, R., Petropoulos, F., & Tipping, M. E. (2022). Hierarchies everywhere–Managing & measuring uncertainty in hierarchical time series. arXiv:2209.15583.
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.
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.
Kolassa, S. (2016). Evaluating predictive count data distributions in retail sales forecasting. *International Journal of Forecasting*, *32*(3), 788–803.
Kolassa, S. (2020). Why the "best" point forecast depends on the error or accuracy measure. *International Journal of Forecasting*, *36*(1), 208–211.
Kolassa, S. (2022). Do we want coherent hierarchical forecasts, or minimal MAPEs or MAEs? (We won't get both!). *International Journal of Forecasting*, *39*, 1512–1517.
Kourentzes, N., & Athanasopoulos, G. (2021). Elucidate structure in intermittent demand series. *European Journal of Operational Research*, *288*(1), 141–152.
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.

Rangapuram, S. S., Werner, L. D., Benidis, K., Mercado, P., Gasthaus, J., & Januschowski, T. (2021). End-to-end learning of coherent probabilistic forecasts for hierarchical time series. In *Proc. 38th int. conference on machine learning (ICML)* (pp. 8832–8843).

Raunig, B., & Scheicher, M. (2011). A value-at-risk analysis of credit default swaps. *Journal of Risk*, *13*(4), 3–29.

Syntetos, A. A., & Boylan, J. E. (2005). The accuracy of intermittent demand estimates. *International Journal of Forecasting*, *21*(2), 303–314.

Székely, G. J., & Rizzo, M. L. (2013). Energy statistics: A class of statistics based on distances. *Journal of Statistical Planning and Inference*, *143*(8), 1249–1272.

Taieb, S. B., 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.

Weiss, N. A. (2005). *A course in probability*. Pearson.

Wickramasuriya, S. L. (2021). Properties of point forecast reconciliation approaches. arXiv:2103.11129.

Wickramasuriya, S. L. (2023). Probabilistic forecast reconciliation under the Gaussian framework. *Journal of Business & Economic Statistics*, http://dx.doi.org/10.1080/07350015.2023.2181176.

Wickramasuriya, S. L., Athanasopoulos, G., & Hyndman, R. J. (2019). Optimal forecast reconciliation for hierarchical and grouped time series through trace minimization. *Journal of the American Statistical Association*, *114*(526), 804–819.

Zambon, L., Azzimonti, D., & Corani, G. (2024). Efficient probabilistic reconciliation of forecasts for real-valued and count time series. *Statistics and Computing*, *34*(1), article no. 21.