2016 IEEE International Conference on Big Data (Big Data)

Forecast UPC-Level FMCG Demand, Part III: Grouped Reconciliation

Dazhi Yang*, Gary S. W. Goh, Siwei Jiang and Allan N. Zhang
Singapore Institute of Manufacturing Technology (SIMTech)
*Agency for Science, Technology and Research (A*STAR)*
Singapore, Singapore
*Email: *yangdz@simtech.a-star.edu.sg, yangdazhi.nus@gmail.com*

***Abstract*—Coordination across a supply chain creates win-win situation for all players in that supply chain; we address the benefits, in terms of forecast accuracy, of reconciling demand forecasts across a supply chain. In Part III of this three-part paper, we continue our discussion on optimal reconciliation of forecasts. Two contributions are made in this paper: 1) the grouped reconciliation technique is used to address the forecast inconsistency in situations when more than one hierarchy can be defined in a supply chain, and 2) minimum trace (MinT) estimator is used to further improve the reconciliation accuracy on top of the weighted least square (WLS) approach, which was used in the earlier parts of this three-part paper. Following the earlier works, the same set of fast moving consumer goods data is used here. The current results are compared to the previous ones. It is shown that the MinT reconciliation technique outperforms the WLS approach, which has been previously identified as the best reconciliation technique for the data from the bottled juice category in the Dominick's Finer Food dataset.**

***Keywords*-FMCG; forecasting; grouped reconciliation; minimum trace**

I. Introduction

From an operational point of view, coordination across a supply chain often creates win-win situation for all players in that supply chain. Such coordination not only refers to the flow of materials, but also refers to the flow of information. Integrating demand related information across a supply chain is an interesting topic in modern sales and operations planning. In this paper, we continue our discussion on demand forecasting for fast moving consumer goods (FMCG); a pair of companion papers (Part I and Part II) published earlier contained background information on this topic [1], [2].

FMCG demand can be considered as a hierarchy; the number of levels in the hierarchy and the aggregation/disaggregation interpretation may vary accordingly. As an example, consider the case where a manufacturer ships several products to various distributors, the individual distributor's demand of a particular product aggregates up in the hierarchy to the total (manufacturer level) demand of that product. As both the manufacturer and distributors ought to produce their own demand forecasts based on their "optimal" approaches, may them be econometrics, machine learning or even qualitative forecasting methods, the forecasts in the hierarchy are most likely to be aggregate inconsistent. In other words, the distributors' demand forecasts do not necessarily add up to the manufacturer's forecasts. The main goal of Ref. [2] was to address this forecast aggregate inconsistency through reconciliation (see below).

Although various forecast reconciliation techniques discussed in Ref. [2] could address aggregate inconsistency by revising all forecasts in a hierarchy, the hierarchical reconciliation does not explain the disagreement in forecasts made using different hierarchies. The hierarchy in the above-mentioned example is known to have a product grouping (see top plot of Figure 1); the middle level is represented by the total demand of each product. Similarly, an alternative hierarchy could be defined based on geographical grouping (see bottom plot of Figure 1). If two sets of reconciled forecasts are made using these two hierarchies, there is little reason, if not no reason, to believe that these forecasts would be identical. This motivates the discussion in this paper, i.e., grouped time series forecast reconciliation. However, before we discuss that, a brief review of the earlier companion papers is provided.

A. Review of Part I and Part II

In Ref. [1], several visualization techniques were introduced specifically for FMCG demand time series. In a big data environment, thousands, or even millions, of demand time series co-exist and interact with each other. Some traditional time series visualization approaches, such as overlaying various time series on a same plot (the spaghetti plot), are not scalable. On this point, kite diagrams were used to visualize a moderate number of time series; a few hundreds of time series can be compared at once. Furthermore, additional information such as missing data points and data points with zero sales can also be visualized in parallel. Ref. [1] also considered a PCA-based plot to visualize a large number of time series based on time series features (e.g., trend, seasonality, linearity). As features of a time series can be designed based on needs, this visualization technique can be used for outlier detection and forecasting model selection. It is well known that exogenous factors such as price and promotional information influence demand; another visualization using step functions was proposed to identify factors that correlate with demand. One of the main goals of Part I is to conduct exploratory analyses on the Dominick's Finer Food (DFF) dataset; and thus we designed our data preprocessing steps specific to the DFF dataset. After preprocessing, demand, price and promotion index time series, each containing 377 weekly data points, from 37 UPCs were prepared for Part II.

In Ref. [2], three hierarchical reconciliation techniques, namely, the bottom-up (BU), optimal reconciliation with ordinary least squares (OLS) and weighted least squares (WLS) estimators, were used to reconcile the base forecasts generated by various forecasting algorithms. The term *base forecast* is used to refer to those forecasts generated (independently) by different players in the supply chain. Let the vector of the h-step-ahead base forecasts be $\widehat{\boldsymbol{Y}}_t(h)$, all reconciliation method can be written as [3]:

$$\widetilde{\boldsymbol{Y}}_t(h) = \boldsymbol{S}\boldsymbol{P}\widehat{\boldsymbol{Y}}_t(h), \tag{1}$$

where $\widetilde{\boldsymbol{Y}}_t(h)$ is the vector of the final revised forecasts; $\boldsymbol{S}$ is the summing matrix (more details below); and $\boldsymbol{P}$ is a matrix of choice, indicating various reconciliation techniques. While the bottom-up reconciliation treats forecasts at upper levels as the arithmetic sum of bottom-level forecasts, the optimal reconciliation techniques consider the regression [4]:

$$\widehat{\boldsymbol{Y}}_t(h) = \boldsymbol{S}\boldsymbol{\beta}_t(h) + \boldsymbol{\varepsilon}_h, \tag{2}$$

where $\boldsymbol{\beta}_t(h)$ is the unknown mean of the most disaggregate series at the bottom level and $\boldsymbol{\varepsilon}_h$ is the reconciliation error. Based on this regression, the optimal (minimum variance unbiased) estimator is given by the generalized least squares (GLS):

$$\boldsymbol{\beta}_t^{\mathrm{GLS}}(h) = (\boldsymbol{S}^\top \boldsymbol{\Sigma}_h^\dagger \boldsymbol{S})^{-1} \boldsymbol{S}^\top \boldsymbol{\Sigma}_h^\dagger \widehat{\boldsymbol{Y}}_t(h), \tag{3}$$

where $\boldsymbol{\Sigma}_h^\dagger$ is a generalized inverse of the unknown reconciliation error covariance $\boldsymbol{\Sigma}_h$. As $\boldsymbol{\Sigma}_h$ is not known and had been shown to be not identifiable [5], OLS and WLS were used to estimate $\boldsymbol{\beta}_t(h)$. The two estimators are given as:

$$\boldsymbol{\beta}_t^{\mathrm{OLS}}(h) = (\boldsymbol{S}^\top \boldsymbol{S})^{-1} \boldsymbol{S}^\top \widehat{\boldsymbol{Y}}_t(h) \tag{4}$$

and

$$\boldsymbol{\beta}_t^{\mathrm{WLS}}(h) = (\boldsymbol{S}^\top \widehat{\boldsymbol{W}}_{1,D}^{-1} \boldsymbol{S})^{-1} \boldsymbol{S}^\top \widehat{\boldsymbol{W}}_{1,D}^{-1} \widehat{\boldsymbol{Y}}_t(h), \tag{5}$$

respectively, where $\widehat{\boldsymbol{W}}_{1,D}$ is a diagonal matrix made of the sample variances of base forecast errors. Once the $\boldsymbol{\beta}_t(h)$ estimate is obtained, the reconciled forecasts can be evaluated by substituting the estimate into Eq. (2). Using the data of the 37 UPCs arranged in Ref. [1], Ref. [2] compared the forecast accuracy of BU, OLS and WLS reconciliations, as well as several univariate benchmarking models. It was found, in terms of the mean absolute percentage error (MAPE), the WLS optimal reconciliation performed the best among all methods.

B. Contributions

As mentioned earlier, one of the goals of this paper is to introduce the grouped reconciliation methods, so that the forecasts produced using different hierarchies (product or geographical groupings) can be consistent. While hierarchical forecasting has been applied to a variety of problems [6]–[9], the grouped time series forecasting is less known. To the best of our knowledge, such grouped reconciliations have only been applied to forecast infant mortality rate [10] and Australian labour market [11]. In this paper, we consider such grouped reconciliations for FMCG demand forecasting.

Another contribution of this paper is on the reconciliation technique itself. Instead of using OLS and WLS on grouped time series, another estimator is considered, namely, the minimum trace (MinT) estimator. This estimator minimizes the trace of the error[1] covariance, thus making it optimal in terms of minimum variance. In a later section, both MinT and WLS reconciliations will be used to forecast both grouped and hierarchical FMCG demand. The rest of the paper is organized as follows: Section II introduces formulation of the grouped reconciliation and various estimators for the reconciliations. Section III updates the results of our case study, namely, the 1-week-ahead forecast for the FMCG demand dataset that was used in Refs. [1], [2]. Conclusions follow at the end.

II. Grouped Reconciliation

In this section, grouped reconciliation methods will be discussed based on the illustrative example shown in Figure 1. In hierarchical reconciliation, the summing matrix $\boldsymbol{S}$ contains the structural information of a hierarchy. Similarly, in grouped reconciliation, the summing matrix describes the structures of several possible groupings, simultaneously. Based on the two hierarchies shown in Figure 1, we define a vector containing the bottom-level data:

$$\boldsymbol{b}_t = (y_{\mathrm{U1S1},t}, y_{\mathrm{U1S2},t}, \cdots, y_{\mathrm{U3S3},t})^\top, \tag{6}$$

where $y_{\mathrm{U}i\mathrm{S}j,t}$ denotes the demand at time t for UPC i at store j. We can then write the following:

$$\boldsymbol{Y}_t = \boldsymbol{S}\boldsymbol{b}_t \tag{7}$$

[1]The error here refers to the h-step-ahead reconciled forecast error. Several other errors will be discussed in Section II.

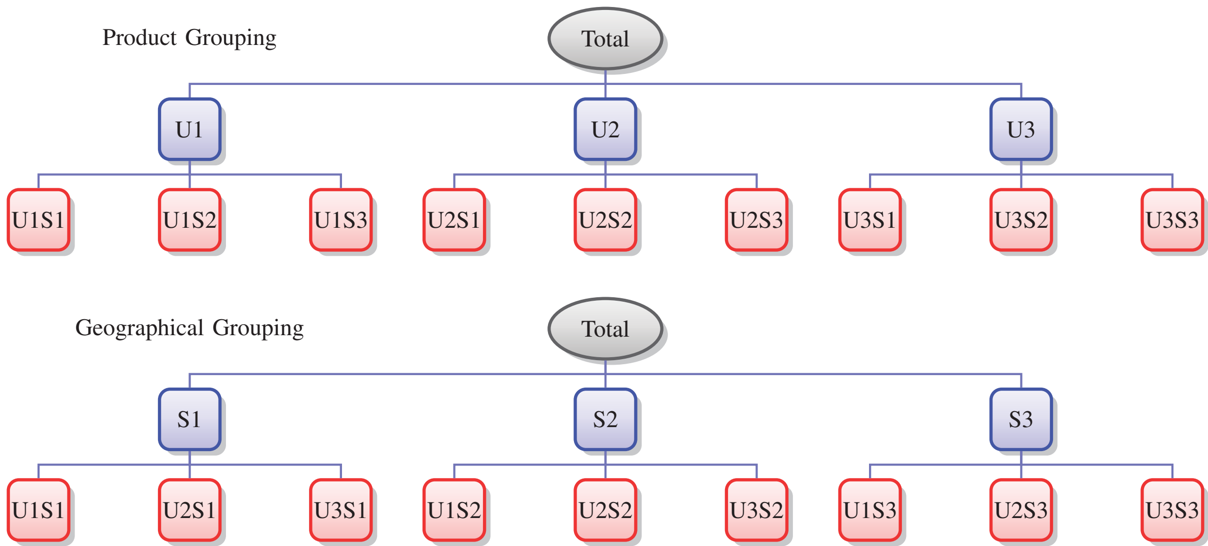

Figure 1. The hierarchy of demand time series can be grouped based on UPC (product grouping) or store (geographical grouping).

where

$$\boldsymbol{S} = \left(\begin{array}{ccccccccc} 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ \hline 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 \\ \hline 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 \\ \hline 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{array}\right) \tag{8}$$

and

$$\boldsymbol{Y}_t = (y_{\mathrm{total},t}, y_{\mathrm{U1},t}, y_{\mathrm{U2},t}, y_{\mathrm{U3},t}, y_{\mathrm{S1},t}, y_{\mathrm{S2},t}, y_{\mathrm{S3},t}, \boldsymbol{b}_t^\top)^\top. \tag{9}$$

It can be seen that the summing matrix of grouped time series contains additional rows than the summing matrix of hierarchical time series. From the above equations, we can see how the bottom-level series aggregate up according to the ones and zeros in $\boldsymbol{S}$. Using the grouped reconciliation, the reconciled forecasts will not only be aggregate consistent for a single hierarchy, but for all hierarchies described by the summing matrix. When the supply chain structure gets more complex, and more hierarchies can be defined (such as manufacturer grouping), the summing matrix can be expanded accordingly.

The regression approach for optimal hierarchical forecasts reconciliation shown in Eq. (2) is also applicable to grouped reconciliation. In other words, the three reconciliation techniques described in Ref. [2] can be used "as is" in this work. Nevertheless, the literature of hierarchical/grouped time series forecasting has advanced since the publication of Refs. [1] and [2]; we update those methodological improvements in this work as well.

A. The MinT estimator

The error term in Eq. (2), ε_h, is known as the reconciliation error; it describes the aggregate inconsistency observed in the base forecasts. However, Ref. [4] confused this error with the h-step-ahead reconciled forecast error [5]. On this point, two other errors need to be defined, namely, the h-step-ahead reconciled forecast error:

$$\widetilde{\boldsymbol{e}}_t(h) = \boldsymbol{Y}_{t+h} - \widetilde{\boldsymbol{Y}}_t(h), \tag{10}$$

and the h-step-ahead base forecast error:

$$\widehat{\boldsymbol{e}}_t(h) = \boldsymbol{Y}_{t+h} - \widehat{\boldsymbol{Y}}_t(h). \tag{11}$$

In Ref. [5], the authors expressed the covariance of $\widetilde{\boldsymbol{e}}_t(h)$ in terms of the covariance of $\widehat{\boldsymbol{e}}_t(h)$ through the following lemma:

Lemma 1 (Wickramasuriya): For any $\boldsymbol{P}$ such that $\boldsymbol{SPS} = \boldsymbol{S}$, the covariance matrix of the h-step-ahead reconciled forecast errors is given by

$$\mathtt{Var}[\boldsymbol{Y}_{t+h} - \widetilde{\boldsymbol{Y}}_t(h)] = \boldsymbol{SPW}_h\boldsymbol{P}^\top\boldsymbol{S}^\top \tag{12}$$

where $\widetilde{\boldsymbol{Y}}_t(h)$ is given by Eq. (1) and $\boldsymbol{W}_h = \mathbb{E}[\widehat{\boldsymbol{e}}_t(h)\widehat{\boldsymbol{e}}_t^\top(h)]$ is the covariance matrix of the h-step-ahead base forecast errors. (*end of lemma*)

In addition, the authors also proposed the following theorem:

Theorem 2 (Minimum trace): Let $\boldsymbol{W}_h$ be the positive definite covariance matrix of the h-step-ahead base forecast error. Then the optimal reconciliation matrix, which minimizes the trace of $\boldsymbol{SPW}_h\boldsymbol{P}^\top\boldsymbol{S}^\top$ such that $\boldsymbol{SPS} = \boldsymbol{S}$, is given by

$$\boldsymbol{P} = \left(\boldsymbol{S}^\top\boldsymbol{W}_h^{-1}\boldsymbol{S}\right)^{-1}\boldsymbol{S}^\top\boldsymbol{W}_h^{-1}. \tag{13}$$

(*end of theorem*)

As the aim of the reconciliation is to make better forecast, reconciled forecasts that minimize the sum of the forecast variances (trace of $\mathtt{Var}[\widetilde{\boldsymbol{e}}_t(h)]$) are desired. The optimal reconciliation approach described in Theorem 2 is referred to as the MinT approach. MinT gives the best (minimum variance) linear unbiased reconciled forecasts. Together with Lemma 1, $\boldsymbol{P}$ can be estimated in terms of the summing matrix $\boldsymbol{S}$ and the covariance matrix of $\widehat{\boldsymbol{e}}_t(h)$, namely, $\boldsymbol{W}_h$. Consequently, the reconciled forecasts using MinT are given by:

$$\widetilde{\boldsymbol{Y}}_t^{\mathrm{MinT}}(h) = \boldsymbol{S}\left(\boldsymbol{S}^\top\boldsymbol{W}_h^{-1}\boldsymbol{S}\right)^{-1}\boldsymbol{S}^\top\boldsymbol{W}_h^{-1}\widehat{\boldsymbol{Y}}_t(h). \tag{14}$$

B. The alternative MinT estimators

Although $\boldsymbol{W}_h$ does not suffer from a lack of identifiability, it is nevertheless difficult to estimate. In Ref. [5], several alternative MinT estimators were discussed. We consider two alternatives in this paper.

If we set $\boldsymbol{W}_h = k_h\mathrm{diag}(\widehat{\boldsymbol{W}}_1) = k_h\widehat{\boldsymbol{W}}_{1,D}$, where k_h is positive and $\widehat{\boldsymbol{W}}_1$ is the sample estimator of the in-sample 1-step-ahead base forecast error covariance, the MinT estimator is equivalent to the WLS estimator described earlier. In other words, the WLS reconciled forecasts are given by:

$$\widetilde{\boldsymbol{Y}}_t^{\mathrm{WLS}}(h) = \boldsymbol{S}\left(\boldsymbol{S}^\top\widehat{\boldsymbol{W}}_{1,D}^{-1}\boldsymbol{S}\right)^{-1}\boldsymbol{S}^\top\widehat{\boldsymbol{W}}_{1,D}^{-1}\widehat{\boldsymbol{Y}}_t(h). \tag{15}$$

As WLS estimator was shown to outperform OLS estimator using the particular dataset in Ref. [2], OLS is not included in this paper.

The second alternative MinT estimator considers shrinkage. Let $\boldsymbol{W}_h = k_h\widehat{\boldsymbol{W}}_{1,D}^*$, where $\widehat{\boldsymbol{W}}_{1,D}^* = \lambda_D\widehat{\boldsymbol{W}}_{1,D} + (1-\lambda_D)\widehat{\boldsymbol{W}}_1$, the shrinkage estimator shrinks the off-diagonal entries of $\widehat{\boldsymbol{W}}_1$, while the diagonal terms remain unchanged. This approach could be used in situation where the sample covariance estimate is not positive definite. In our case, λ_D follows the shrinkage parameter proposed in Ref. [12]. The reconciled forecasts in this case is given by:

$$\widetilde{\boldsymbol{Y}}_t^{\mathrm{MinT*}}(h) = \boldsymbol{S}\left(\boldsymbol{S}^\top\widehat{\boldsymbol{W}}_{1,D}^{*-1}\boldsymbol{S}\right)^{-1}\boldsymbol{S}^\top\widehat{\boldsymbol{W}}_{1,D}^{*-1}\widehat{\boldsymbol{Y}}_t(h), \tag{16}$$

and the methods is referred to as MinT* reconciliation.

III. Update of the previous results

Following Refs. [1] and [2], data from the Dominick's database is used; the dataset is provided by the James M. Kilts Center, University of Chicago Booth School of Business with a collaborative effort by the Dominick's Finer Food (DFF).

A. Previous results

Data of 37 products (UPC/product-level time series) from the bottled juice category (BJC) were arranged in Ref. [1]. With a total of 1971 bottom-level time series, a three-level hierarchy[2] was formed in Ref. [2]. The total number of time series in this hierarchy is thus 2009 (1+37+1971). Using the two base forecast methods, or more specifically, autoregressive distributed lag (ADL) model [13] for level 2 and simple exponential smoothing (SES) [14] for levels 0 and 1, the performance of three hierarchical reconciliation techniques (BU, OLS and WLS) was evaluated in terms of MAPE. As for each time series, the total number of weekly demand observation is 377, data from the first 200 weeks were used for model fitting, the remaining 177 weeks of data were used to compute true out-of-sample MAPE with a rolling forecast horizon (see Figure. 2). It was found that on level 1 (UPC-level), WLS reconciliation performs the best.

In addition to the overall MAPE, some error analyses were also performed in Ref. [2]. In particular, it was found that when the demand time series contains large surges (due to promotional events and price drops), SES fails to identify these demand surges, and thus produces high MAPE (as high as 262%). As the results of SES will be used during reconciliation, and the large errors will propagate to the forecasts of other series, data of 4 UPCs with MAPE $> 50\%$ are excluded from the computation in this paper. In other words, we perform grouped reconciliation based on data of 33 UPCs; the total number of bottom-level series is 1712. These 4 UPCs can be considered as time series anomalies. We note that these time series anomalies in general should not be included in a forecast system; removing unusual time series or data points prior to forecasting is a commonly adopted approach for the industry. On this point, in a big data environment, scalable time series anomaly detection algorithms [15] can be employed. The anomalous series can then be treated separately with other appropriate forecasting methods. We hope to include such discussions in a future work.

B. Results for grouped reconciliation

In this section, 1-step-ahead (weekly) forecasts are performed and evaluated at the UPC level. If the DFF dataset

[2] This is sometimes referred to as a two-level hierarchy, as the top level is considered as level 0.

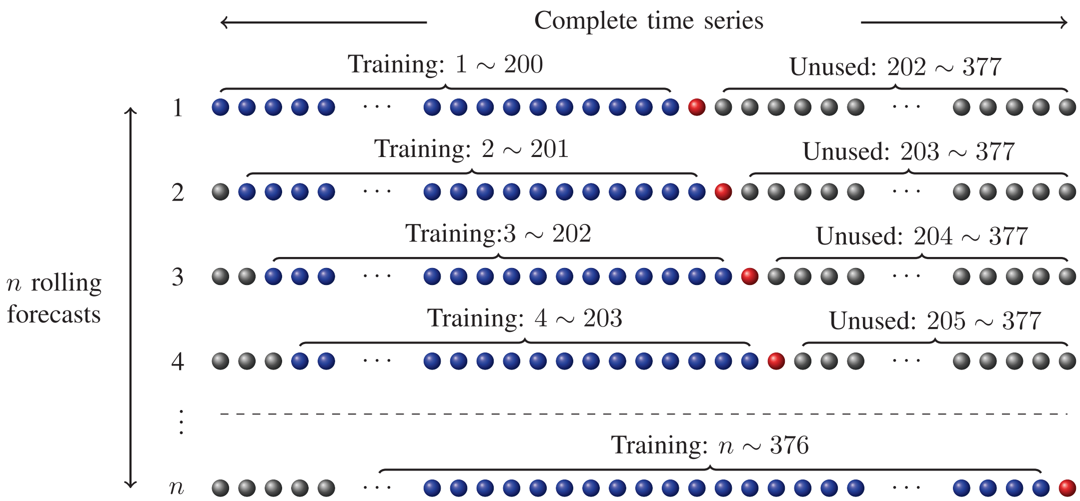

Figure 2. Experimental design: 177 ($n = 177$ in this case) rolling forecasts are performed for each time series. The data point being forecast is shown in red. The size of training window (shown in blue) is fixed at 200. This scheme was used in Ref. [2].

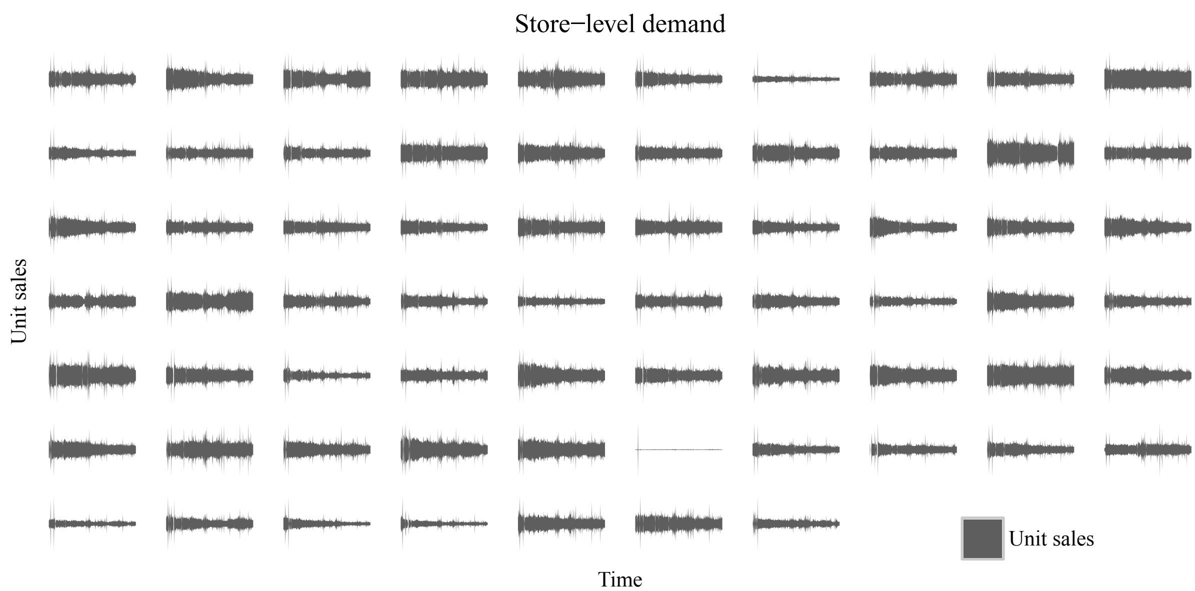

Figure 3. Kite plots (see Ref. [1] for interpretation) for 67 store-level demand series. Each store-level time series is the aggregate of the demand of all BJC products (under consideration) in that store. Frequent and large spikes are observed in these time series.

is modeled as a hierarchy, the optimal (MinT) hierarchical reconciliation can be readily calculated via Eqs. (15) and (16), with $\widehat{\boldsymbol{Y}}_t(1)$ being the vector of 1-step-ahead base forecasts obtained in Ref. [2]. However, when the dataset is described as grouped time series, the reconciliation still requires those forecasts on level which contains demand under a geographical grouping. In other words, to construct $\widehat{\boldsymbol{Y}}_t(1)$ for grouped time series, the forecasts of the total demand in each store (for the 33 UPCs) are needed. The raw DFF dataset contains demand information at 93 stores, however, the filtered dataset only covers 67 stores; data from other stores were filtered in Ref. [1]. Therefore, in the case of grouped time series reconciliation, for each time t, the length of $\widehat{\boldsymbol{Y}}_t(1)$ is 1813 (1+33+67+1712).

The store-level demand can be forecast using any forecasting methods. For instance, we can employ SES to generate base forecasts for these time series. However, as mentioned earlier, SES may generate large errors if the

Table I

FORECAST ERRORS (MAPE) OF 33 SELECTED UPCS FROM THE BOTTLED JUICE CATEGORY. THE ERRORS ARE IN PERCENTAGE.

	UPC	**Product name**	**No. of series**	**Univariate methods**		**Hierarchical methods**			**Grouped methods**	
				Persistence	SES	Bottom-up	WLS	MinT	WLS	MinT*
1	7045011484	SUNSWEET PRUNE JUIC 48Oz	55	10.16	10.74	9.02	8.64	8.23	8.64	**8.20**
2	5300015154	REALEMON PLASTIC LEM 45Oz	61	11.17	9.64	9.00	8.92	8.20	8.93	**8.16**
3	5300015108	REALEMON LEMON JUICE 8Oz	64	15.11	13.05	10.89	10.89	11.23	**10.89**	11.26
4	3828103123	HH LEMON JUICE 32Oz	61	16.74	16.92	18.32	17.50	12.75	17.52	**12.58**
5	3828103091	HH APPLE JUICE 128Oz	66	31.81	37.66	20.92	20.86	19.40	20.86	**19.40**
6	3828103025	DOM APPLE JUICE 32Oz	52	39.11	41.96	21.97	**21.68**	27.20	21.70	28.02
7	3120027407	OS PINK GRAPEFRUIT 64Oz	42	29.99	34.58	17.61	17.52	16.15	17.50	**16.13**
8	3120027007	OS GRAPEFRUIT JUICE 64Oz	53	28.16	30.34	12.68	12.59	12.67	**12.59**	12.68
9	3120026134	OS LC CRANRASP 48Oz	52	26.11	24.66	12.26	12.33	12.31	12.32	**12.30**
10	3120021007	OS CRANAPPLE DRINK 64Oz	62	25.18	27.01	12.21	12.12	11.66	12.13	**11.64**
11	3120020035	OS LO CAL CRANBERRY 48Oz	63	21.33	20.29	10.49	10.43	10.22	10.43	**10.20**
12	3120020007	OS CRNBRY JCE COCKTA 64Oz	65	26.15	30.91	11.67	11.51	**11.21**	11.54	11.23
13	3120020005	OS CRANBERRY COCKTA 48Oz	66	34.65	50.54	12.46	12.38	15.03	**12.37**	15.20
14	7045011329	SUNSWEET PRUNE JUICE 32Oz	61	12.28	11.99	12.10	11.70	10.05	11.71	9.98
15	5300015132	REALEMON LEMON JUICE 32Oz	66	16.55	17.49	17.05	16.04	14.48	16.04	14.43
16	4850000193	TROP TWSTR ORGCRAN 46Oz	34	30.35	29.55	16.27	16.20	15.25	16.19	**15.19**
17	4180022700	WELCHS WHITE GRAPE J 64Oz	54	10.49	11.11	16.51	15.15	10.79	15.15	**10.66**
18	4180020750	WELCHS GRAPE JUICE 64Oz	65	11.21	11.11	14.87	13.89	10.23	13.89	**10.16**
19	3828103017	HH APPLE JUICE 64Oz	67	66.32	76.37	25.33	25.26	28.67	25.22	29.07
20	3828103009	HH PRUNE JUICE 40Oz	40	15.72	15.93	13.47	12.91	11.60	12.91	**11.55**
21	3120027005	OS GRAPEFRUIT JUICE 48Oz	39	29.75	30.56	16.69	16.60	**15.69**	16.59	15.72
22	3120026107	OS CRANRASPBERRY DR 64Oz	56	28.70	33.27	12.07	11.97	11.46	11.99	**11.44**
23	3120026105	OS CRANRASPBERRY 48Oz	45	43.58	71.13	20.12	20.34	25.18	**20.08**	25.39
24	3120021005	OS CRANAPPLE DRINK 48Oz	41	41.11	55.02	18.67	18.65	19.72	**18.61**	19.82
25	3828103115	DOM CRANBERRY RSPBRY 48Oz	7	33.58	33.57	27.12	26.60	24.30	26.60	**24.25**
26	3828103033	HH CRANBERRY JUICE C 48Oz	61	23.28	24.27	18.61	18.45	16.38	18.47	**16.34**
27	3828103005	DOM PRUNE JUICE 32Oz	41	16.49	15.89	15.48	14.63	12.37	14.64	**12.30**
28	3120027405	OS PINK GRAPEFRUIT 48Oz	26	31.49	36.20	19.40	19.33	18.88	19.32	**18.92**
29	1480000032	MOTTS APPLE JUICE P 32Oz	62	9.57	**9.07**	11.71	10.60	9.46	10.60	9.35
30	5300015407	REAL LIME JUICE 8Oz	7	24.35	22.16	22.72	21.44	20.89	21.44	**20.86**
31	7045011402	SUNSWEET PRUNE JCE W 40Oz	58	12.53	14.76	11.56	11.00	9.69	11.00	**9.65**
32	3120020000	OS CRANBERRY COCKT 32Oz	59	10.99	9.96	10.50	10.20	9.13	10.20	**9.09**
33	1480051324	MOTTS CLAMATO JUICE 32Oz	61	14.10	**12.69**	15.27	14.75	13.00	14.75	12.95
	overall		1712	24.19	26.98	15.61	15.24	**14.65**	15.24	14.67

demand time series is spiky. In our present case, the store-level demand is in fact spiky, as shown in Figure 3. This is owning to the fact that only 37 products from BJC are considered in the forecasting exercise. For such reason, we use the bottom-up approach for store-level forecasts; the respective base forecasts are simply added together to form these store-level forecasts. It is believed that when the number of bottom-level series increases, the store-level demand will be smoothed out. In those cases, other forecasting methods may be deemed as more appropriate.

At this stage, all base forecasts have been obtained; grouped reconciliations can thus be performed. To benchmark the results of grouped reconciliation, three hierarchical reconciliation techniques, namely, BU, WLS and MinT* reconciliation, are also implemented. The hierarchy used here is based on the product grouping; it follows the top plot in Figure 1. The results of the two univariate forecasting methods are reiterated as well. Table I shows the MAPE of all methods considered in this study.

It is observed from Table I that the grouped reconciliation does not improve forecasts in terms of MAPE as compared to the hierarchical methods. This is due to the fact that no new forecast information was created during the generation of store-level demand forecasts. However, if other methods, besides the simple aggregation of bottom-level series, are used, the results of grouped reconciliation will be different from those of hierarchical reconciliation. Furthermore, if the store-level demand is sufficiently smooth, the grouped reconciliation is expected to out-perform hierarchical reconciliation due to the better forecasts available at the store-level. On the other hand, it is clear that the MinT* approach improves forecast accuracy on top of WLS, in both grouped and hierarchical methods.

IV. CONCLUSION

In Part III of this three-part paper, grouped reconciliation is used to forecast UPC-level FMCG demand. As compared to the hierarchical reconciliation techniques used in Part II, the grouped reconciliation techniques consider additional hierarchy definitions. This is useful when the hierarchical structure in a supply chain is not unique. Although in our case study, the improvements in grouped time series forecasts over the hierarchical forecasts are not observable, the new forecasts are conceptually preferred, as the revised forecasts using grouped reconciliation are aggregate consistent across all hierarchies under consideration. Furthermore, it is hypothesized that when the number of bottom-level time series increases, and when the higher-level time series become more smooth, the advantage of the grouped time series reconciliation would be clear.

Besides using grouped reconciliation, this paper also considers the minimum trace reconciliation technique. Previously in Part II, WLS reconciliation was shown to outperform OLS and other benchmarking methods. It is now shown that MinT* reconciliation could further improve forecasts on top of WLS, making MinT* the best overall reconciliation technique for our data presented in the case study.

Although from the definition itself, MinT is the best (minimum variance) linear unbiased estimator, the accuracy of the method in application can still be limited by other factors, such as the accuracy of base forecasts and the error ratio between forecasts at each level. We hope to discuss such issues in a future paper. More specifically, it is important to know the conditions for the reconciled forecasts to be better than the base forecasts. On top of that, FMCG demand has strong variability due to the intricate market mechanisms. Continuous development of base forecast methods in an FMCG context is also important.

Acknowledgment

This work is partially supported under the A*STAR TSRP fund 1424200021 and Antuit–SIMTech Supply Chain Analytics Lab.

References

[1] D. Yang, G. S. W. Goh, C. Xu, A. N. Zhang, and O. Akcan, "Forecast UPC-level FMCG demand, Part I: Exploratory analysis and visualization," in *Big Data (Big Data), 2015 IEEE International Conference on*, Oct 2015, pp. 2106–2112.

[2] D. Yang, G. S. W. Goh, S. Jiang, A. N. Zhang, and O. Akcan, "Forecast UPC-level FMCG demand, Part II: Hierarchical reconciliation," in *Big Data (Big Data), 2015 IEEE International Conference on*, Oct 2015, pp. 2113–2121.

[3] G. Athanasopoulos, R. A. Ahmed, and R. J. Hyndman, "Hierarchical forecasts for Australian domestic tourism," *International Journal of Forecasting*, vol. 25, no. 1, pp. 146–166, Jan. 2009.

[4] R. J. Hyndman, R. A. Ahmed, G. Athanasopoulos, and H. L. Shang, "Optimal combination forecasts for hierarchical time series," *Computational Statistics & Data Analysis*, vol. 55, no. 9, pp. 2579–2589, Sep. 2011.

[5] S. L. Wickramasuriya, G. Athanasopoulos, and R. J. Hyndman, "Forecasting hierarchical and grouped time series through trace minimization," 2015, working paper 15/15, Department of Econometrics & Business Statistics, Monash University.

[6] V. Almeida, R. Ribeiro, and J. Gama, *Hierarchical Time Series Forecast in Electrical Grids*. Singapore: Springer Singapore, 2016, pp. 995–1005. [Online]. Available: http://dx.doi.org/10.1007/978-981-10-0557-2_95

[7] A. A. Syntetos, Z. Babai, J. E. Boylan, S. Kolassa, and K. Nikolopoulos, "Supply chain forecasting: Theory, practice, their gap and the future," *European Journal of Operational Research*, vol. 252, no. 1, pp. 1–26, 2016.

[8] M. Kremer, E. Siemsen, and D. J. Thomas, "The sum and its parts: Judgmental hierarchical forecasting," *Management Science*, vol. 62, no. 9, pp. 2745–2764, 2016.

[9] S. Moon, C. Hicks, and A. Simpson, "The development of a hierarchical forecasting method for predicting spare parts demand in the South Korean Navy:A case study," *International Journal of Production Economics*, vol. 140, no. 2, pp. 794–802, 2012.

[10] H. L. Shang, "Reconciling forecasts of infant mortality rates at national and sub-national levels: Grouped time-series methods," *Population Research and Policy Review*, pp. 1–30, 2016.

[11] R. J. Hyndman, A. J. Lee, and E. Wang, "Fast computation of reconciled forecasts for hierarchical and grouped time series," *Computational Statistics & Data Analysis*, vol. 97, pp. 16–32, May 2016.

[12] J. Schäfer and K. Strimmer, "A shrinkage approach to large-scale covariance matrix estimation and implications for functional genomics," *Statistical Applications in Genetics and Molecular Biology*, vol. 4, no. 1, pp. 1175–1189, Nov. 2005.

[13] T. Huang, R. Fildes, and D. Soopramanien, "The value of competitive information in forecasting FMCG retail product sales and the variable selection problem," *European Journal of Operational Research*, vol. 237, no. 2, pp. 738–748, 2014.

[14] R. J. Hyndman, A. B. Koehler, J. K. Ord, and R. D. Snyder, *Forecasting with Exponential Smoothing*. Deblik, Berlin, Germany: Springer, 2008.

[15] R. J. Hyndman, E. Wang, and N. Laptev, "Large-scale unusual time series detection," 2015, working paper.