Evaluation of hierarchical forecasting for substitutable products

Viswanathan, S.; Piplani, R.; Widiarta, Handik

2008

Viswanathan, S., Widiarta, H., & Piplani, R. (2008). Evaluation of hierarchical forecasting for substitutable products. *International Journal of Services and Operations Management*, 4(3), 277-295.

<https://hdl.handle.net/10356/100751>

<https://doi.org/10.1504/IJSOM.2008.017295>

© 2008 Inderscience. This is the author created version of a work that has been peer reviewed and accepted for publication by *International Journal of Services and Operations Management*, Inderscience. It incorporates referee's comments but changes resulting from the publishing process, such as copyediting, structural formatting, may not be reflected in this document. The published version is available at:
[<http://dx.doi.org/10.1504/IJSOM.2008.017295>].

Downloaded on 26 Apr 2023 21:34:39 SGT

Evaluation of hierarchical forecasting for substitutable products

Handik Widiarta and S. Viswanathan*

Nanyang Business School
Nanyang Technological University
639798, Republic of Singapore
E-mail: handik@gmail.com
E-mail: vish@pmail.ntu.edu.sg
*Corresponding author

R. Piplani

School of Mechanical & Aerospace Engineering
Nanyang Technological University
639798, Republic of Singapore
E-mail: mrpiplani@ntu.edu.sg

Abstract: This paper addresses hierarchical forecasting in a production planning environment. Specifically, we examine the relative effectiveness of Top-Down (TD) and Bottom-Up (BU) strategies for forecasting the demand for a substitutable product (which belongs to a family) as well as the demand for the product family under different types of family demand processes. Through a simulation study, it is revealed that the TD strategy consistently outperforms the BU strategy for forecasting product family demand. The relative superiority of the TD strategy further improves by as much as 52% as the product demand variability increases and the degree of substitutability between the products decreases. This phenomenon, however, is not always true for forecasting the demand for the products within the family. In this case, it is found that there are a few situations wherein the BU strategy marginally outperforms the TD strategy, especially when the product demand variability is high and the degree of product substitutability is low.

Keywords: forecasting; hierarchical forecasting; exponential smoothing; substitutable products.

Reference to this paper should be made as follows: Widiarta, H., Viswanathan, S. and Piplani, R. (xxxx) 'Evaluation of hierarchical forecasting for substitutable products', *Int. J. Services and Operations Management*, Vol. X, No. Y, pp.000–000.

Biographical notes: Handik Widiarta obtained his PhD degree from Nanyang Technological University, Singapore. He completed a Bachelor's Degree in Industrial Engineering from Trisakti University Indonesia and a Master's degree in Industrial Management from Catholic University Leuven. His main interests include supply chain management, inventory theory and forecasting. His papers have appeared in *Naval Research Logistics* and the *International Journal of Production Research*.

S. Viswanathan is an Associate Professor of Operations Management at Nanyang Business School, Nanyang Technological University, Singapore. His research interests include inventory management, joint replenishment inventory systems, supply chain coordination, reverse logistics models and hierarchical forecasting. His research has been published in journals such as *Management Science*, *Operations Research*, *Operations Research Letters*, *Naval Research Logistics*, *International Journal of Production Research*, *Journal of the Operational Research Society*, *European Journal of Operational Research* and *International Journal of Physical Distribution & Logistics Management*.

Rajesh Piplani is an Associate Professor in the School of Mechanical and Aerospace Engineering, Nanyang Technological University, Singapore. His research interests are in supply chain coordination, supply chain planning, logistics and production planning. His research has been published in journals such as the *International Journal of Production Research*, *European Journal of Operational Research*, *Naval Research Logistics*, *International Journal of Logistics and Integrated Manufacturing Systems*.

1 Introduction

This paper studies hierarchical forecasting in a production planning environment when the individual products in the family are substitutable. The demand for a particular substitutable product is driven not only by its unique characteristics, but also by the inventory level of other products in the family with similar characteristics. Product substitutability increases the complexity of forecasting and production planning due to its negative effect on the accuracy of the Point-of-Sale (POS) data. But it is a reality that inventory and production planning managers have to contend with. A survey conducted by the Food Marketing Institute (1993), for instance, reported that the majority of typical shoppers (82%–88%) would be willing to buy another size of the same brand, or switch brands, if their favourite brand size was not available on a shopping trip; while the rest indicated that they would not buy any items.

Companies have to make sure that they do not overforecast the demand for a product, when its observed demand might be inflated because of the additional demand from the other similar product, and vice versa. Another complexity associated with the product substitutability is that even if the companies, by any means, could know which products are substitutable, it is still difficult to measure the consumer's willingness to buy another product in case his preferred product is out of stock. Many researchers have studied this issue from an inventory management perspective. The most recent ones include Anupindi *et al.* (1998), Smith and Agrawal (2000) and Rajaram and Tang (2001). Typically, papers that study the inventory management of substitutable products adopt a newsboy model framework (Chen and Parlar, 2005).

In this paper, we address the problem of forecasting the demand for substitutable products. Specifically, the relative effectiveness of Top-Down (TD) and Bottom-Up (BU) forecasting strategies in estimating the demand for a substitutable product (which belongs to a family), as well as the demand for a product family under different types of family demand processes, is examined through a simulation study. The TD strategy uses the history of the product family demand to forecast the aggregate demand. This is then

disaggregated based on the historical demand proportion of each individual product in the family. The BU strategy forecasts the product demand individually and combines the results to obtain the forecast for the product family demand. As is common in a production planning environment, it is assumed that exponential smoothing is used as the forecasting technique under both strategies.

The main difference between our study and all the previous research in forecasting is that the focus of investigation is on the impact of (1) the degree of product substitutability and (2) the variability of demand quantity of each product in the family, on the relative performance of the two forecasting strategies. To the best of our knowledge, there has been no study done on this topic. Delurgio and Bhame (1991, p.122) suggested that when substitutions occur, it is important to record the demand history in such a way as to not distort the demand forecasts that will use that history. However, the authors did not provide any clear guideline as to which forecasting strategy should be adopted. Lapide (1998) suggested that the BU strategy is preferred over the TD strategy when the product family is composed of competing items that potentially cannibalise each other's demand. However, the author did not elaborate further on how the change in the degree of product substitutability would affect this preference. Many prior studies such as Widiarta *et al.* (2007), Gross and Sohl (1990), Dangerfield and Morris (1992), Zellner and Tobias (2000), Dekker *et al.* (2004) and Zotteri *et al.* (2005), to name a few, have evaluated the relative effectiveness of TD and BU forecasting strategies. However, to our knowledge, no studies have considered substitutable products in their modelling framework.

Another contribution of this study is the assumption that the product family demand (instead of the product demands as in the other papers) follows a certain time series process, and the individual product demands are derived from the product family demand. This assumption is motivated by the fact that in the retail merchandising industry, the individual product demands are usually more erratic than the product family demand. The fluctuations of several product demands tend to cancel each other out when aggregated together, thus resulting in a smoother family demand. Hence, it is common to find situations wherein the product family demand may be identified to follow a particular time series process, while the individual product demands are simply derived from the product family demand based upon a certain proportion. Besides, for many products that are not regularly purchased and for fashion products, the customer shopping decisions may be based on impulse or personal liking within a family of products, which supports the above assumption.

Our simulation study in this paper reveals that the TD strategy consistently outperforms the BU strategy for forecasting the product family demand. The superiority of the TD strategy to the BU strategy further improves by as much as 52% as the product demand variability increases and the degree of substitutability between the two products decreases. This phenomenon, however, is not always true for forecasting the demand for products within the family. In this case, it is found that there are a few cases where the BU strategy marginally outperforms the TD strategy, especially when the product demand variability is high and the degree of product substitutability is low.

The remainder of this paper is organised as follows: In Section 2, we discuss the design of the simulation experiments. In Section 3, we analyse the results of the simulation study. Finally, in Section 4, we make some concluding remarks.

2 Design of simulation experiments

As mentioned previously, the focus of this research is to examine the impact of the product demand variability and the degree of product substitutability on the performance of TD and BU forecasting strategies. Therefore, in order to obtain meaningful insights, the number of products in the family is restricted to two. Most of the earlier studies have also considered two components in the family (see, for instance, Schwarzkopf *et al.*, 1988; Dangerfield and Morris, 1992). In addition, the demand process for the product family is assumed to follow a particular time series, whereas the demand for each individual product is derived from the product family demand. The demand proportion for each product in the family, p_i ($0 \leq p_i \leq 1$), is assumed to be uniformly distributed, $p_i \sim U(LB, UB)$, where $UB > LB$ and $\tau = UB - LB$. In the rest of the study, τ , the range of variability of the product's proportion in the family, is also referred to as the product demand variability because of its positive correlation to the level of demand uncertainty of each individual product.

Three time series (*i.e.*, two stationary and one nonstationary) are considered to represent the demand process of the product family. Let $D(t)$ be the family demand in period t . The three processes for the family demand are as follows:

- 1 First-order Moving Average [MA(1)]: $D(t) = \mu + \varepsilon(t) - \theta\varepsilon(t-1)$ (1)

- 2 First-order Autoregressive [AR(1)]: $D(t) = (1 - \phi)\mu + \phi D(t-1) + \varepsilon(t)$ (2)

- 3 Integrated Moving Average of order 1 [IMA(1, 1)]:

$$D(t) = D(t-1) + \varepsilon(t) - \theta\varepsilon(t-1) \quad (3)$$

where, $|\theta| < 1$, $|\phi| < 1$, μ is the expected demand, and $\varepsilon(t)$ is a random variable that is normally distributed with zero mean ($E[\varepsilon(t)] = 0$), variance σ^2 and zero autocovariance ($\text{Cov}[\varepsilon(t), \varepsilon(t-k)] = 0$, for all $k \neq 0$). Note that coefficients θ and ϕ in Equation (1) to Equation (3) introduce the serial correlation of the demand process, *i.e.*, demand for period t is dependent on demand for the earlier periods. Therefore, in the rest of the paper, we identify θ and ϕ as the coefficients of the serial correlation term. The simple exponential smoothing (which is known for its wide usage in the industry due to its simplicity and robustness (Dekker *et al.*, 2004)) is used as the forecasting technique under both the TD and BU strategies.

The degree of product substitutability (or the substitutability ratio) for product i is denoted as β_{ij} , where $0 \leq \beta_{ij} \leq 1$. This represents a deterministic portion of the unsatisfied/excess demand for product i that is passed to product j , provided that product j has sufficient inventory. Clearly, $\beta_{ij} = 0$ implies that none of the excess demand for product i would be passed to and fulfilled by product j . This represents the case when the consumer would rather wait for his/her preferred product than buy the alternative. For simplicity, we assume that the two products have the same level of importance to the consumers, meaning that the degree of substitutability from products i to j is equal to that from products j to i (*i.e.*, $\beta_{ij} = \beta_{ji} = \beta$).

A detailed explanation on the behaviour of the consumers to purchase substitutable products is provided in the Appendix. In summary, product substitution takes place only if part of the substituted demand can be fulfilled with the on-hand inventory of the alternative product. If this condition is not satisfied, no product substitution happens, and

thus, the entire unmet demand will be back-ordered to the respective product. In marketing management, this mechanism is known as *consumer-controlled* substitution. Note that no part of the unmet demand is lost. In other words, all excess demands are assumed to be completely back-ordered. If there are lost sales, then the forecasting would be much more difficult due to the complexity in estimating the unobservable lost sales.

The sequence of activities for the retailer in each time period of the simulation run is as follows: First, the shipments from the outside supplier (shipped L periods ago, where L is the replenishment lead time) are received into the inventory. Next, the demand for the two products arises and is immediately satisfied, if there are adequate stocks, or backlogged, otherwise. Then, the backlogged demand from the previous periods, if any, is satisfied, if there is stock. It is assumed that partial fulfilment of demand is allowed. The demand forecasts (based on either TD or BU strategies) are then updated, followed by computation of the safety stock and the net requirement. Finally, the order schedule using the lot-for-lot method is calculated and orders corresponding to the current period are placed with the outside supplier, who delivers the products after a fixed lead time.

Note that for the case of forecasting the product demands, simple exponential smoothing is used to apportion the aggregate forecast $[F(t)]$ from the TD strategy to each individual product (McLeavey and Narasimhan, 1985, pp.69–71). Mathematically, it is written as follows:

$$F_1(t) = \left[\hat{\alpha} \frac{D_1(t-1)}{D_1(t-1) + D_2(t-1)} + (1 - \hat{\alpha}) \frac{F_1(t-1)}{F(t-1)} \right] F(t) \quad (4)$$

$$F_2(t) = F(t) - F_1(t) \quad (5)$$

where $F_i(t)$ is the demand forecast for product i at time t (for $i = 1, 2$), $F(t-1)$ is the aggregate (family) forecast at time $t-1$, and $\hat{\alpha}$ is a fixed smoothing constant. Notice that $\hat{\alpha}$ used here in the disaggregation or allocation method is different from α_i (for $i = 1, 2$) and α used in the forecasting process. α_i is the smoothing constant used for forecasting the demand for products 1 and 2 under the BU strategy, while α is the smoothing constant used for forecasting the product family demand under the TD strategy.

In each experiment, the simulation was carried out for a total of 1040 periods. The first 52 periods were used to obtain the initial parameters. The second part, containing 260 periods, was used to determine the optimal smoothing constants α_i and α (by minimising the Root of Mean Squared Error (RMSE) for this part). Note that in the BU strategy, the smoothing constant α_i was optimised for each product demand individually. Finally, the last 728 periods were used to compile the performance statistics of the two forecasting strategies.

The study was carried out for a wide range of statistical parameters. Values of problem parameters that were used are provided in Table 1. For all the problems, the forecast error (RMSE) values of both TD and BU strategies were computed and compared using the ratio:

$$\nabla = RMSE(TD)/RMSE(BU). \quad (6)$$

A value of $\nabla < 1$ implies that the TD strategy is superior to the BU strategy, whereas a value of $\nabla > 1$ indicates the contrary. For a particular set of parameters, the simulation was run for different values of the safety stock until a safety stock that achieved a 95% fill rate ($FR = 95\%$) was obtained. The value of ∇ was then calculated corresponding to this level of safety stock.

Table 1 Problem parameters used in the experiment

<i>Parameter</i>	<i>Value</i>
Expected demand of the product family (μ)	500 units
Lower bound ^a of the product's proportion in the family (LB)	0.3
Delivery lead time from the external supplier to the retailer (L)	1 week
Grid size used to optimise the smoothing constant	0.01
Smoothing constant used in the proration mechanism ($\hat{\alpha}$)	0.2
Actual fill rate for each product (FR)	95%
Variability ^b of the product's proportion in the family (τ)	1 From 0.1 to 0.7 in increments of 0.1 2 From 0.05 to 0.4 in increments of 0.05
Demand variance of the product family (σ^2)	1 900 units ² 2 90 000 units ²
Coefficient for the MA(1) and IMA(1, 1) processes (θ)	From -0.9 to 0.9 in increments of 0.2
Coefficient for the AR(1) process (ϕ)	From -0.9 to 0.9 in increments of 0.2
Degree of product substitutability (β)	From 0.0 to 1.0 in increments of 0.2

Notes: ^aThe upper bound of the product's proportion (UB) = $LB + \tau$.

^bThere are two sets of data for each category in order to ensure that the simulation study encompasses all possible values of demand correlation between the two products. The combination of low demand variance and high variability of the product's proportion results in a negative demand correlation between the two products, and vice versa.

3 Analysis of the results

We now discuss in detail the relative performance of the two forecasting strategies under different experimental scenarios. The discussion is divided into two subsections. Section 3.1, which refers to Figures 1 to 3, evaluates the relative effectiveness of TD and BU strategies for forecasting the product demands; whereas Section 3.2, referring to Figures 4 to 6, discusses the case of forecasting the product family demand. Each figure consists of several graphs, each of which plots the simulation results for a specific demand process. In all the graphs, the average performance ratio ∇ is plotted on the y-axis and is studied with respect to two particular parameters. Note that BO denotes back-orders and represents a special case where the excess demand for a particular product is completely backlogged and is not satisfied by the other product (*i.e.*, $\beta = 0$). Therefore, in this case the observed demand used in the forecasting process is always equal to the actual demand.

3.1 Forecasting product level demands

When the effectiveness of the TD and BU strategies for forecasting the product demands is compared, it is generally found from all the graphs in Figure 1 that the relative superiority of the TD strategy over the BU strategy increases with the degree of product substitutability (β). In Figure 1(c), for instance, the average of the performance ratio ∇ under an IMA(1, 1) demand process appears to decrease by as much as 8.5% as the two products become more substitutable with each other.

Figure 1 The relative performance of the TD over the BU strategy for forecasting product demands (different β)

(a) MA(1) demand process

(b) AR(1) demand process

Figure 1 The relative performance of the TD over the BU strategy for forecasting product demands (different β) (continued)

(c) IMA(1, 1) demand process

The possible explanation for this phenomenon is that when the two products are highly substitutable, the distortion of the *observed* demand from the *real* demand for each individual product becomes more significant. This is primarily due to the fact that any excess demand for a particular product, say product 1, would not be immediately visible to the retailer as the consumer might try to satisfy a portion β of this excess demand by buying the alternative product 2. The higher the substitutability ratio of a product, the lower the exact amount of the excess demand or the excess inventory known to the retailer. Therefore, the consequence of this information distortion to the performance of the BU strategy is very obvious: for the above example, the retailer is likely to underforecast the demand for product 1 and to overforecast that for product 2.

The performance of the TD strategy, on the other hand, is not really affected by the information distortion caused by the product substitution, as it uses the historical demand of the product family, which is independent of the degree of product substitutability (β). One may argue, though, that it would still be difficult for the TD strategy to properly allocate the aggregate demand forecast to each individual product, especially when the degree of product substitutability is high. This is because the retailer may not be able to identify the actual demand quantity of each individual product. In this experiment, however, exponential smoothing (with $\hat{\alpha} = 0.2$) is used as the disaggregation method for the TD strategy (see Equation (4) and Equation (5)). The advantage of this method is that it considers the most recent distribution of demands, in addition to taking into account the historical distribution of demands. Therefore, the prediction of the product's proportion in the TD strategy is significantly improved.

The impact of product demand variability (τ) on the relative performance of the two forecasting strategies is illustrated in Figure 2. Each value of ∇ plotted in the graph is obtained by averaging the value of ∇ across different values of ϕ or θ , depending on the time series process. As can be seen, the impact of τ on the value of ∇ depends on the degree of product substitutability (β). Specifically, when the value of β is at

low-to-medium level, the increase in τ is found to negatively affect the relative effectiveness of the TD strategy over the BU strategy. For a special case when $\beta = 0$, the RMSE of the TD strategy is even higher than the RMSE of the BU strategy by as much as 4%. This is to be expected, as the value of τ has a direct influence on the volatility level of the product demand. As τ increases, the amplification of the product demand also increases, thus making the allocation process of the aggregate forecast in the TD strategy more difficult.

Figure 2 The relative performance of the TD over the BU strategy for forecasting product demands (different τ)

(a) MA(1) demand process

(b) AR(1) demand process

Figure 2 The relative performance of the TD over the BU strategy for forecasting product demands (different τ) (continued) **(c) IMA(1, 1) demand process**

The negative impact of the increase of τ on the TD strategy's performance, however, reduces as β increases. For a special case when $\beta = 1$, increasing τ is actually found to improve the TD strategy's performance. Intuition suggests that this might be due to the 'mitigation effect' of β , which stabilises the individual product demand (when τ is high). This means that the higher the degree of product substitutability (β), the larger the portion of the excess demand for a particular product that would be passed to the other product, whose on-hand inventory is in surplus. As a result, the overall product demand variability decreases, thus improving the ability of the TD strategy to accurately distribute the aggregate forecast.

Note that when τ is low ($\tau < 0.2$), the change in β has little impact on the relative effectiveness of TD over BU strategies. This is because when τ is low, the product demand becomes more stable and easier to estimate. This results in the reduction in the number of stockout occurrences as well as the (potential) number of product substitution occurrences.

3.2 Forecasting family level demand

In this section, the performance of TD and BU strategies for forecasting product family demand is compared. Unlike the earlier case, from all the figures it is discovered that the TD strategy consistently outperforms the BU strategy by as much as 52%, regardless of the degree of product substitutability (β), the product demand variability (τ) and the coefficient of the serial correlation term (θ and ϕ). Although under some scenarios it is found that $\nabla > 1$ (Figure 3(a)), the relative superiority of the BU strategy over the TD strategy is considered insignificant as it never exceeds more than 1%. One reason for the superior performance of the TD strategy over the BU strategy may be risk-pooling.

Risk-pooling takes advantage of the statistical fact that the variance of the aggregated demand is equal to the sum of the variances and covariances of the individual product demand. In this way, the fluctuation of demand from one source may be offset by that from other sources, thus resulting in a lower forecast error.

Figure 3 The relative performance of the TD over the BU strategy for forecasting family demand (different β)

(a) MA(1) demand process

(b) AR(1) demand process

Figure 3 The relative performance of the TD over the BU strategy for forecasting family demand (different β) (continued) **(c) IMA(1, 1) demand process**

Figure 4 shows that the value of ∇ tends to decrease as τ increases. This confirms the general intuition that a highly volatile product demand reduces the ability of the BU strategy to forecast accurately. That is, the higher the uncertainty of the product demand, the higher the forecast error incurred by the BU strategy. The performance of the TD strategy, however, is unaffected by the changes in τ and β due to the distinct approach to how the product demands are generated in this study (see Section 2).

Figure 4 The relative performance of the TD over the BU strategy for forecasting family demand (different τ) **(a) MA(1) demand process**

Figure 4 The relative performance of the TD over the BU strategy for forecasting family demand (different τ) (continued) **(b) AR(1) demand process** **(c) IMA(1, 1) demand process**

The impact of the degree of product substitutability (β) on the value of ∇ can be observed in all the figures. It is generally found that ∇ increases with β , especially when τ is high (Figures 4(b) and 4(c)). This might be due to the mitigating effect of β , which reduces the fluctuation of the product demand and thus improves the accuracy of the BU strategy (as in the previous case). When τ is low, however, the positive influence of β on the performance of the BU strategy is not that significant, as most of the demand from the consumers are fairly predictable.

4 Conclusion

In this paper, the performance of TD and BU strategies for forecasting the demand for a product (item) that belongs to a family and the demand for a product line (family) was evaluated. As is common in the retail merchandising industry, the products may be substituted for each other to a certain degree. The demand of the product family was assumed to follow a particular time series process, whereas the demand for each individual product was derived from the product family. Three demand time series were considered: first-order Moving Average [MA(1)], first-order Autoregressive [AR(1)], and Integrated Moving Average of order one [IMA(1, 1)]. Using a simulation approach, the impacts of (1) the degree of product substitutability and (2) the variability of demand quantity of each product in the family, on the relative performance of the two forecasting strategies were investigated.

Simple exponential smoothing, which is known for its reasonable accuracy, is used as the forecasting technique under both strategies and as the disaggregation method for the TD strategy. The forecast performance was measured by the RMSE and the contribution achieved by each strategy was compared by a simple ratio.

The study revealed that the TD strategy, with the benefit of risk-pooling, consistently outperformed the BU strategy for forecasting the product family demand by as much as 52%. The superiority of the TD strategy over the BU strategy further improved as the variability of the product's proportion in the family increased and the degree of substitutability between the two products decreased. This phenomenon, however, was not always true when the performance of the two strategies for forecasting product demands was compared. It was found that there were a few cases wherein the BU strategy marginally outperformed the TD strategy, particularly when the variability of the product's proportion was high and the degree of product substitutability was low.

Acknowledgements

We would like to express our gratitude to the Editor-in-Chief and the anonymous referees for their comments and suggestions on an earlier version of the paper.

References

- Anupindi, R., Dada, M. and Gupta, S. (1998) 'Estimation of consumer demand with stock-out based substitution: an application to vending machine products', *Marketing Science*, Vol. 17, No. 4, pp.406–423.
- Chen, H. and Parlar, M. (2005) 'Dynamic analysis of the newsboy model with early purchase commitments', *International Journal of Services and Operations Management*, Vol. 1, No. 1, pp.56–74.
- Dangerfield, B.J. and Morris, J.S. (1992) 'Top-down or bottom-up: aggregate versus disaggregate extrapolations', *International Journal of Forecasting*, Vol. 8, No. 2, pp.233–241.
- Dekker, M., van Donselaar, K.H. and Ouwehand, D.P. (2004) 'How to use aggregation and combined forecasting to improve seasonal demand forecasts', *International Journal of Production Economics*, Vol. 90, pp.151–167.
- Delurgio, S. and Bhame, C. (1991) *Forecasting Systems for Operations Management*, Illinois: Business One Irwin.

- Food Marketing Institute (1993) *Variety or Duplication: A Process to Know Where You Stand*, The Research Department, Food Marketing Institute, Washington, DC.
- Gross, C.W. and Sohl, J. (1990) 'Disaggregation methods to expedite product line forecasting', *Journal of Forecasting*, Vol. 9, No. 3, pp.233–254.
- Lapide, L. (1998) 'New developments in business forecasting', *The Journal of Business Forecasting Methods & Systems*, Vol. 17, No. 2, pp.28–29.
- McLeavey, D.W. and Narasimhan, S.L. (1985) *Production Planning and Inventory Control*, Boston: Allyn and Bacon, Inc.
- Rajaram, K. and Tang, C.S. (2001) 'The impact of product substitution on retail merchandising', *European Journal of Operational Research*, Vol. 135, pp.582–601.
- Schwarzkopf, A.B., Tersine, R.J. and Morris, J.S. (1988) 'Top-down versus bottom-up forecasting strategies', *International Journal of Production Research*, Vol. 26, No. 11, pp.1833–1843.
- Smith, S.A. and Agrawal, N. (2000) 'Management of multi-item retail inventory systems with demand substitution', *Operations Research*, Vol. 48, No. 1, pp.50–64.
- Widiarta, H., Viswanathan, S. and Piplani, R. (2007) 'On the effectiveness of top-down approach for forecasting autoregressive demands', *Naval Research Logistics*, Vol. 54, pp.176–188.
- Zellner, A. and Tobias, J. (2000) 'A note on aggregation, disaggregation and forecasting performance', *Journal of Forecasting*, Vol. 19, No. 5, pp.457–469.
- Zotteri, G., Kalchschmidt, M. and Caniato, F. (2005) 'The impact of aggregation level on forecasting performance', *International Journal of Production Economics*, Vols. 93–94, pp.479–491.

Appendix

Definitions of the following notations:

β : Degree of product substitutability, $0 \leq \beta \leq 1$

$D_i(t)$: *Actual* (real) demand for product i in period t (where $i = 1, 2$), in units

$\tilde{D}_i(t)$: *Observed* demand for product i in period t , in units

$I_i(t)$: On-hand inventory for product i in period t (after satisfying the direct demand for product i), in units

$I_i^B(t)$: Beginning inventory for product i in period t (before adding with the order quantity shipped L periods ago), in units (equivalent to $I_i^E(t-1)$)

$I_i^E(t)$: Ending inventory for product i in period t , in units (equivalent to $I_i^B(t+1)$)

$\tilde{I}_i(t)$: Available inventory for product i in period t (before satisfying the direct demand for product i), in units

$q_{i \rightarrow j}$: Excess demand for product i which is satisfied by product j , in units

$SR_i(t)$: Scheduled receipt for product i in period t , in units

Note that a negative inventory level represents back orders. The explanation of the inventory logic begins with the case when $\beta = 0$.

Case 1: The individual products are not substitutable (i.e., $\beta = 0$)

The inventory logic for this case is relatively straightforward, as no excess demand for a particular product is passed to and satisfied by the other product. At time t , the available inventory $[\tilde{I}_i(t)]$ is first computed by summing the incoming supply, which was delivered L periods ago $[SR_i(t)]$, and the ending inventory from the previous period $[I_i^E(t-1)]$. This available inventory is then used to satisfy any (real) demand requested by a consumer $[D_i(t)]$. Whenever the quantity of the demand is larger than the available inventory $[D_i(t) > \tilde{I}_i(t)]$, the remaining demand is completely backlogged and no part of this excess demand will be substituted for by the other product.

Case 2: The individual products are substitutable (i.e., $0 < \beta \leq 1$)

The pseudocode of the inventory logic with two substitutable products is given below:

Step 1: Update the parameters for product 1.

1	$\tilde{I}_1(t) = I_1^B(t) + SR_1(t)$	Update the available inventory
2	$I_1(t) = \tilde{I}_1(t) - D_1(t)$	Update the on-hand inventory
3	$\tilde{D}_1(t) = D_1(t)$	Copy the value of the actual demand to the observed demand

4	If ($I_1(t) > 0$)	<i>If the on-hand inventory is positive</i>
5	{	
6	Set $q_{1 \rightarrow 2} = 0$	<i>No excess demand is passed to product 2</i>
7	$I_1^E(t) = I_1(t)$	<i>Update the ending inventory</i>
8	}	
9	Else	<i>If the on-hand inventory is negative</i>
10	{	
11	If ($\tilde{I}_1(t) < 0$)	<i>If the available inventory is negative</i>
12	$q_{1 \rightarrow 2} = \lceil \beta \times D_1(t) \rceil$	<i>Compute the demand quantity to be passed to product 2</i>
13	Else	<i>If the available inventory is positive</i>
14	$q_{1 \rightarrow 2} = \lceil \beta \times I_1(t) \times (-1) \rceil$	<i>Compute the demand quantity to be passed to product 2</i>
15	$\tilde{D}_1(t) = \tilde{D}_1(t) - q_{1 \rightarrow 2}$	<i>Update the observed demand</i>
16	$I_1^E(t) = I_1(t) + q_{1 \rightarrow 2}$	<i>Update the ending inventory</i>
17	}	

Step 2: Update the parameters for product 2.

18	$\tilde{I}_2(t) = I_2^B(t) + SR_2(t)$	<i>Update the available inventory</i>
19	$I_2(t) = \tilde{I}_2(t) - D_2(t)$	<i>Update the on-hand inventory</i>
20	$\tilde{D}_2(t) = D_2(t)$	<i>Equalise the actual and the observed demand</i>
21	If ($I_2(t) > 0$)	<i>If the on-hand inventory is positive</i>
22	{	
23	Set $q_{2 \rightarrow 1} = 0$	<i>No excess demand is passed to product 1</i>
24	If ($q_{1 \rightarrow 2} > 0$)	<i>If there is an excess demand from product 1</i>
25	{	
26	$I_2^E(t) = I_2(t) - q_{1 \rightarrow 2}$	<i>Update the ending inventory</i>
27	$\tilde{D}_2(t) = \tilde{D}_2(t) + q_{1 \rightarrow 2}$	<i>Update the observed demand</i>
28	Return $q_{1 \rightarrow 2} = 0$	
29	}	
30	Else	
31	$I_2^E(t) = I_2(t)$	<i>Update the ending inventory</i>
32	}	
33	Else	<i>If the on-hand inventory is negative</i>

34	{	
35	If $(\tilde{I}_2(t) < 0)$	<i>If the available inventory is negative</i>
36	$q_{2 \rightarrow 1} = \lceil \beta \times D_2(t) \rceil$	<i>Compute the demand quantity to be passed to product 1</i>
37	Else	<i>If the available inventory is positive</i>
38	$q_{2 \rightarrow 1} = \lceil \beta \times I_2(t) \times (-1) \rceil$	<i>Compute the demand quantity to be passed to product 1</i>
39	$\tilde{D}_2(t) = \tilde{D}_2(t) - q_{2 \rightarrow 1}$	<i>Update the observed demand</i>
40	$I_2^E(t) = I_2(t) + q_{2 \rightarrow 1}$	<i>Update the ending inventory</i>
41	If $(q_{1 \rightarrow 2} > 0)$	<i>If there is an excess demand from product 1</i>
42	{	
43	$\tilde{D}_2(t) = \tilde{D}_2(t) + q_{2 \rightarrow 1}$	<i>Update the observed demand</i>
44	$I_2^E(t) = I_2^E(t) - q_{2 \rightarrow 1}$	<i>Update the ending inventory</i>
45	Return $q_{2 \rightarrow 1} = 0$	
46	$\tilde{D}_1(t) = \tilde{D}_1(t) + q_{1 \rightarrow 2}$	<i>Update the observed demand for product 1</i>
47	$I_1^E(t) = I_1^E(t) - q_{1 \rightarrow 2}$	<i>Update the ending inventory for product 1</i>
48	Return $q_{1 \rightarrow 2} = 0$	
49	}	
50	Else	
51	{	
52	$\tilde{D}_1(t) = \tilde{D}_1(t) + q_{2 \rightarrow 1}$	<i>Update the observed demand for product 1</i>
53	$I_1^E(t) = I_1(t) - q_{2 \rightarrow 1}$	<i>Update the on-hand inventory for product 1</i>
54	Return $q_{2 \rightarrow 1} = 0$	
55	}	
56	}	

The above inventory logic can be explained as follows: Suppose that at time t , the sum of the beginning inventory for a particular product, say product 1 $[I_1^B(t)]$, and the scheduled receipt $[SR_1(t)]$ is not sufficient to fully satisfy the (real) demand from a consumer $[D_1(t)]$. A portion of this unsatisfied (real) demand (*i.e.*, $q_{1 \rightarrow 2} = \lceil \beta \times \min\{D_1(t), I_1(t) \times (-1)\} \rceil$, where β is assumed to be known and deterministic, and $\lceil x \rceil$ is the smallest integer $\geq x$) would then be fulfilled from the on-hand inventory of the alternative product 2 $[I_2(t)]$, provided that $I_2(t) > 0$. In this case, if it is found that $q_{1 \rightarrow 2} > I_2(t)$, the remaining unsatisfied demand (*i.e.*, $q_{1 \rightarrow 2} - I_2(t)$) would be back-ordered to product 2. This

mechanism is adopted in order to avoid duplication of demands and also from the fact that it may be impractical for the consumer to have orders for both products 1 and 2 if the two products are qualitatively comparable. The computations of the ending inventory and the new observed demand for alternative product 2 are shown in code lines 26 and 27, respectively.

However, if the on-hand inventory of alternative product 2 is also not enough to even satisfy its own demand ($I_2(t) \leq 0$), the excess demand for each individual product would not be fulfilled by the other product. In other words, substitution by the other product takes place only if at least part of the substituted demand can be satisfied in the current period. The computations of the ending inventory and the new observed demand for the two products under this situation are shown in code lines 43 to 48.