CIRRELT

CIRRELT-2024-02

Assessing Data Collection Strategies for VMI under Intermittent Demand

Corey Ducharme
Bruno Agard
Martin Trépanier

January 2024

Bureau de Montréal
Université de Montréal
C.P. 6128, succ. Centre-Ville
Montréal (Québec) H3C 3J7
Tél : 1-514-343-7575
Télécopie : 1-514-343-7121

Bureau de Québec
Université Laval,
2325, rue de la Terrasse
Pavillon Palasis-Prince, local 2415
Québec (Québec) G1V 0A6
Tél : 1-418-656-2073
Télécopie : 1-418-656-2624

Assessing Data Collection Strategies for VMI under Intermittent Demand

Corey Ducharme[1,*], Bruno Agard[1,2], Martin Trépanier[1,2]

[1] Laboratoire en Intelligence des Données (LID) and Department of Mathematical and Industrial Engineering, Polytechnique Montréal

[2] Interuniversity Research Centre on Enterprise Networks, Logistics and Transportation (CIRRELT)

Abstract. Collecting a customer's point-of-sale demand data can alleviate the problem of intermittent demand by providing less intermittent data for demand forecasts. This also allows for collaborative supply chain arrangement such as Vendor Managed Inventory. However, data collection technologies can be cost prohibitive and information-sharing arrangements amongst supply chain members, unreliable. When obtaining point-of-sale demand data is not feasible, specialised forecasting methods have been proposed to tackle the problem using improved modeling of intermittent time series. For supply chains faced with intermittent demand, decisions about an information strategy are complex and cost-benefit investigations are rare. In this paper, we present a simulation of the supply chain of a real supplier currently operating a non-periodic stochastic order-up-to-level Vendor Managed Inventory arrangement with multiple customers. The proposed simulation framework allows for comparison between two different information strategies: the point-of-sale telemetry demand data versus the historical deliveries demand data. Simulations are for targeted service levels and attempt to minimize both the number of deliveries and the required safety stock. Results are measured in terms of exact safety level, inventory stock, and the number of deliveries to achieve a targeted service level across the entire supply chain. The influence of product lead time is also explored. Our results show that collecting a customer's point-of-sale demand data offers significant savings in terms of customer inventory stock and lesser savings in terms of the number of deliveries.

Keywords: Intermittent demand, supply chain management, vendor managed inventory, simulation, demand data.

Acknowledgements. The authors would like to acknowledge our industrial partner and the Natural Sciences and Engineering Research Council of Canada (NSERC) for funding under grant RDCPJ 492021-15, RGPIN-2019-04723. We also thank the members of the Data Intelligence Laboratory of Polytechnique Montreal for their constructive comments that helped the authors to make this research possible.

---

* Corresponding author: corey.ducharme@polymtl.ca

Dépôt légal – Bibliothèque et Archives nationales du Québec
Bibliothèque et Archives Canada, 2024

1 Introduction

Intermittent demand in supply chain management remains a pervasive challenge for industries (Nikolopoulos, 2021). Unfortunately, this problem has received less academic attention compared to forecasting of fast moving (non-intermittent) items (Syntetos et al., 2016), even though intermittent items account for substantial proportions of stock value (Johnston et al., 2003). Furthermore, this problem will become increasingly common as industries have shown a willingness to use granular data in their supply chain operations, whether that be by forecasting individual clients or products (Willemain et al., 1994). The more granular the data, the more intermittent the time series can be (Bartezzaghi et al., 1999).

Vendor Managed Inventory (VMI) offers an appealing solution as a customer's direct point-of-sale demand data contains less noise and possesses a higher sampling frequency, while still being of the desired granularity (Murray et al., 2018a). Furthermore, VMI allows for improved inventory management, which in turn translates to cost savings (Cao and Zhang, 2011; Jung et al., 2005; Zhou et al., 2017).

However, acquiring the customer's direct demand data through technological means can be cost prohibitive (Jung et al., 2005) and information sharing often requires proof that the cost improvements are worth the efforts for both suppliers and their customers (Kembro and Näslund, 2014). Information sharing arrangements are inherently risky (Colicchia et al., 2019) and improvements are unknown until the arrangement has been finalized. Such risks imply that suppliers may have to continue operating a VMI without data sharing for some of its customers. Thus, there is a need for more case studies that compare the performance increase and possible risk mitigation to help guide industry practitioners when making these choices.

In this paper, we have access to two sources of demand data from an industrial partner operating under a VMI arrangement: a customer's point-of-sale demand data collected by means of telemetry storage containers, and a customer's delivery demand data from the supplier's delivery records. These two sources of demand data represent two different perspectives of the same demand along the supply chain. Both sources of demand data have varying degrees of intermittent behavior.

This paper proposes and executes an experimental design for evaluating these two sources of information in a VMI arrangement. First, we propose a VMI arrangement under which a supplier is entirely responsible for a customer's stock. The supplier, using the available demand data, forecasts and performs customers' replenishments. Secondly, we iteratively simulate the replenishments for each customer over an entire year to determine the lowest number of deliveries and safety stock level required to maintain a desired service level. Lastly, we evaluate the simulations on the inventory performance, the safety stock level, and the number of deliveries. Results are presented for different targeted service levels and different lead times under both demand information scenarios.

The remainder of this paper is organized as follows. Section 2 presents a review of the pertinent literature on VMI and intermittent demand models for both forecasting and inventory management. Section 3 describes the experimental design of the VMI, the simulation, the forecasting model, and the performance measurements. Section 4 presents and discusses the results. Section 5 concludes the paper with managerial recommendations, limits, and opportunities for further research.

2 Literature Review

2.1 Vendor Managed Inventory

Various forms of VMIs have been implemented in multiple industries differing in the type of good, supply chain size, and industry-specific problems (Borade and Sweeney, 2015). Broadly, VMIs are a type of collaborative supply chain management in which the supplier and customer share information that allows the supplier to manage the customer's stock directly. Once implemented, the supplier can independently decide when and how much to deliver based on the shared information (Vigtil, 2007). Thus, information sharing is key to successful VMI (Angulo et al., 2004). Ideally, a supplier has access to downstream information related to product usage, sales plans, and product forecasts (Achabal et al., 2000). This allows suppliers to optimise their own forecasts and their logistic network to ensure the reliability of the supply arrangement between themselves and their customers. Many papers have demonstrated the advantages of information sharing in collaborative supply chain management for both suppliers and their customers (Cao and Zhang, 2011; Jung et al., 2005; Zhou et al., 2017). Despite this, supply chain partners may be unable or unwilling to share data (Holweg et al., 2005; Kembro and Näslund, 2014). In such cases where information is not shared, the supplier must turn to other means to forecast customer demand (Ali et al., 2017). Oftentimes, as a last resort, the supplier will use historical delivery records for demand forecast, but these records may be noisy (Murray et al., 2018a).

Information collection offers an alternative when data is not willingly shared. Telemetry systems can directly measure and relay the downstream information on product usage. Advances in information technology have lowered the cost (Ru et al., 2018) of what was in the past considered a cost-prohibitive solution (Jung et al., 2005). However, information sharing should not be considered an entirely cost-free solution. Determining the cost benefit of information collection is not obvious and remains grounded in empirical evaluations that hopefully match the conditions of an industry pursuing this solution. In terms of cost determining cost benefits of VMI, Ru et al. (2018) show that for the one supplier one customer scenario, the benefit of VMI depends on having low inventory holding costs for customers and high inventory holding costs for suppliers. The opposite result has also been presented (Kim, 2008).

VMI studies that directly researched VMI under intermittent demand are more rare. Wu and Hsu (2008) proposed a configurable bill-of-materials to reduce logistic costs for spare parts, but their proposed approach was both computationally costly and slow. In our case study, there is no spare parts bill-of-materials, as each customer receives a single product. Scala et al. (2013) proposed an alternative measure for the lead time demand based on adjusting the lead time following a simulation to minimize inventory costs, but the use case for their method was for extremely infrequent spare parts in the order of less than one transaction per year, which is significantly lower than our available case study. Fu and Chien (2019) showed improvements of intermittent demand forecasting model for vendor managed inventories over traditional models but did not explore the influence of the input demand data.

2.2 Intermittent Demand

The development of intermittent demand has been driven by industry use cases: forecasting, inventory planning, and smoothing. These different approaches may seem distinct but share the common trait of imposing an underlying model for the intermittent demand. For modeling intermittent inventory demand, which is the object of research in this paper, two families of models exist: parametric and nonparametric. Regardless of use case, determining the optimal intermittent demand model is still an open problem (Kourentzes et al., 2019).

For parametric models, a known distribution is chosen for the model and then fit to the data. The most widely used parametric model is Croston's method (Croston, 1972). Croston's method first divides intermittent demand into two constituent parts: the demand amount and the interdemand interval (the time between two periods of non-zero demand). The demand amount is assumed to follow a normal distribution and the interdemand arrival, a geometric distribution. Syntetos and Boylan (2005) provided a bias correction for Croston's method, and it is their Syntetos-Boylan Approximation (SBA) method which continues to receive substantial empirical support (Syntetos et al., 2016; Syntetos et al., 2015).

Nonparametric models do not assume an underlying distribution for the intermittent demand model. We further divide nonparametric models into three broad groups: temporal aggregation, bootstrapping, and machine learning. Temporal aggregation methods resolve the issue of zero demand observations by aggregating time series from higher (e.g., daily) to lower (e.g., weekly) frequencies. The aggregated time series can then be forecasted and disaggregated to return it to the original frequency. Temporal aggregation has received significant support following its initial concept in (Nikolopoulos et al., 2011; Spithourakis et al., 2012). Kourentzes et al. (2014) proposed a systematic approach that combines the result of multiple exponential smoothing forecasts done at different aggregated frequencies. Murray et al. (2018a) proposed the use of Croston's method to improve the aggregation. Bootstrapping methods sample observations to create a histogram that replaces a theoretical distribution, as very erratic demand may not comply with any standard theoretical distribution (Hasni et al., 2019b). The downside of bootstrapping techniques is the added complexity (Syntetos et al., 2015), especially when considering that difference in performance is small (Hasni et al., 2019c). Finally, the very broad machine learning category leverages recent innovations in machine learning to tackle intermittent demand. Methods used have included neural networks (Kourentzes, 2013; Lolli et al., 2017) and segmentation (Murray et al., 2018b). These methods are not widely applied, owing to their increased complexity, the mixed evidence of their superiority, and their narrower use cases (Mukhopadhyay et al., 2012; Teunter and Duncan, 2009).

In inventory planning under intermittent demand, the choice is either to model the demand as a known distribution or to use a nonparametric bootstrapping method (Zhou and Viswanathan, 2011). In the simple order-up-to-level case with a known demand distribution, the order-up-to-level can be determined by using the inverse of the cumulative distribution function of the model's chosen distribution (Syntetos et al., 2015). For nonparametric methods, the models produce a point forecast of the order-up-to-level for a targeted service level (Syntetos et al., 2015). When comparing these two approaches, Syntetos et al. (2015) recommended the simpler parametric models in terms of performance compared to the added complexity of nonparametric models. Sillanpää and Liesiö (2018) further recommends against the use of point forecasts over parametric distributions modeling. In more complex inventory management situations, full scale simulation of the inventory is done based on parametric or nonparametric point forecasts of the inventory demand which must then be filled (Eaves and Kingsman, 2004). Estimates for the safety stock can also be determined from the point forecast error uncertainty (Trapero et al., 2019).

For intermittent demand models used in VMI systems, Fu and Chien (2019) used a combination forecast of temporal aggregation and Croston's method, autoregressive integrated moving average, and recurrent neural networks. The VMI methodologies proposed in (Scala et al., 2013; Wu and Hsu, 2008) were agnostic as to the type of intermittent demand model. The lack of research on the intermittent demand for VMI is to be expected, as most VMI papers assume that the distribution is known. This is in line with a similar observation made by (Syntetos et al., 2015) who noted the lack of studies on intermittent demand for inventory management, of which VMI is a subproblem.

3 Experimental Design

The goal of this paper is to compare demand information strategies in a VMI arrangement under intermittent demand. To do this, we will simulate a VMI arrangement with point-of-sale telemetry demand and with historical deliveries demand, and then evaluate their respective performance. Thus, we need to describe the VMI arrangement, the simulation design and parameters, both sources of demand data, the demand forecasting model, and the performance measurements.

3.1 Vendor Managed Inventory

Properly describing the VMI arrangement is of utmost importance, as VMIs differ widely in both their industrial implementations and in the literature (Borade and Sweeney, 2015). The VMI arrangement presented in this paper is adapted from the one used by the industrial partner who provided the data required to simulate the VMI. The industrial partner is a large supplier of raw materials operating across the contiguous United States. The partner delivers to thousands of customers operating in a variety of industries, such as manufacturing, agriculture, food processing, and medical services. These different industrial sectors are known to have varied intermittent demand behaviors. The supplied product requires a dedicated storage container installed on the customer's site-of-operations. These containers can be equipped with telemetry that measure the level of the product.

The VMI arrangement is for a single supplier shipping to several customers. The supplier monitors the inventory of each customer and determines the replenishment policy (amount and date). Both the replenishment amount and date are forecasted based on the available demand information. While ensuring that no stock-out occurs, the replenishment is done as late as possible to maximise the useable onsite inventory and thus minimise the number of deliveries being made. The forecasted replenishment amount attempts to bring a customer back to its maximum level. This is defined as a non-periodic order-up-to-level policy. Under the details of the arrangement, the values of the safety stock level (*SSL*) and the maximum order-up-to-level are fixed for each customer. The maximum order-up-to-level is constrained by the physical system which contains the stock and negotiated amongst the supplier and the customer based on the expected product consumption over the duration of the arrangement. The *SSL* is determined through simulation optimization. The amount delivered will vary for each delivery. The product being delivered has a targeted lead time + review (*LT*), which the supplier attempts to respect.

For a specific customer and a single replenishment, the equation to forecast the date of the replenishment $\widehat{t_r}$ and the amount of the replenishment $\widehat{Q_r}$ are follows.

$$\widehat{t_r} = \underset{t}{argmin}\left(\sum_{i=t_s}^{t} \widehat{D_\iota} \geq S_{t_s} - SSL\right) \tag{1}$$

For a starting date t_s, the forecasted date of replenishment $\widehat{t_r}$ is the smallest time at which the cumulated forecasted demand $\widehat{D_\iota}$ time series is greater than the difference between the starting stock S_{t_s} and the *SSL*, i.e., the first date our cumulated customer demand exceeds the safety reserves.

$$\widehat{Q_r} = \sum_{i=t_s}^{\widehat{t_r}} \widehat{D_\iota} \tag{2}$$

The forecasted replenishment amount is equal to the customer's demand from t_s to $\widehat{t_r}$ inclusively. We include the $\widehat{t_r}$ in the sum as some demand will be consumed on the day of delivery. Generally, the starting date t_s will be the day immediately following a delivery as after a delivery. The supplier will wish to know when to schedule the next delivery.

A unique benefit of VMI arrangements is that since the demand data is being continuously fed to the supplier, either through data collection or information sharing, the supplier can update the delivery forecast $\widehat{D_\iota}$ whenever new observations of the demand are acquired. Updating the demand forecast with new observations improves the accuracy of the replenishment date and amount as the true demand is now partly known. However, because the product has an *LT*, there is a set amount of time during which the supplier cannot update the forecast. The forecast is locked in for production, scheduling, and review.

For example, consider a product with an *LT* of 5 and a starting date t_s. An initial forecast for $\widehat{t_r}$ is 8. The supplier waits a day to collect that day's actual demand data and performs a new forecast. This second forecast for $\widehat{t_r}$ is 5. Both forecasts have the same starting date t_s. This second forecast is sent off for production and scheduling since our *LT* has been reached.

Two additional replenishment cases are treated separately. First, if the forecasted replenishment date following a delivery is less than the *LT*, we assume that the supplier will accelerate the production and perform the delivery anyway. Second, if due to a bad forecast of the demand, the level of stock ends up under the *SSL* or if there is a stock-out, an emergency delivery is done immediately the following day. Increasing the safety stock level can reduce the occurrence of these two events.

3.2 Simulation

Simulation is a common strategy when validating proposed methodologies in both VMI literature (Borade and Sweeney, 2015) and intermittent demand inventory management (Bartezzaghi et al., 1999; Eaves and Kingsman, 2004; Persson et al., 2017; Syntetos et al., 2015). Simulating a VMI arrangement entails determining optimal replenishments over a time frame for a targeted service level. Simulations are useful because they allow for retrospective analysis to determine a minimal optimal safety stock level (Eaves and Kingsman, 2004; Persson et al., 2017; Scala et al., 2013). The experimental nature of this paper will be to simulate our VMI arrangement and measure the performance under the demand information scenarios. Our strategy is iterative in nature. The VMI is simulated under a set number of parameters and re-evaluated iteratively until the optimal safety stock level and replenishments are found.

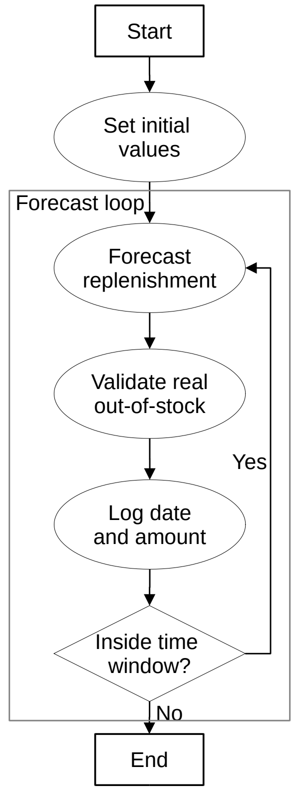

Fig. 1. Simulation flowchart

Before the simulation begins, we set the values of the product's *LT*, the customer's maximum order-up-to-level, and the safety stock level for the customer. The supplier will attempt to always keep the customer above the safety level. Then, a starting date for the simulation is chosen. The initial stock level at the start date is measured from the real telemetry data. From this date and stock level, we forecast the date and amount of the next replenishment following the forecasting process described in the previous section. While a replenishment date is forecasted, we also follow the real stock changes at the customer's site to see if an out-of-stock may occur before the forecasted replenishment is done. If a stock-out occurs, the event is logged, and an emergency delivery is sent out the next day to fill the customer to the maximum level. Once the delivery is done, the procedure is repeated using the replenishment date and amount as the starting date and stock amount for the next forecast. This process is repeated sequentially over the window of time dedicated for the simulation. These steps are shown visually on the flowchart in Fig. 1.

This process represents a single iteration of a simulation of the supply chain for a specific safety stock level. The service level of an iteration is 1 minus the ratio of out-of-stock deliveries over the total number of deliveries done in the simulation. A service level of 100% means there were no stock-outs. To determine the required safety stock for a targeted service level for a customer, we iterate the simulation over all possible safety stock amounts, i.e., [0-100%] of the maximum stock. To speed up the iterating process, we used a bisection algorithm with a stopping tolerance criterion of 2%.

Generally, VMI arrangements target a 100% service level. However, since our experimental design allowed for it, we simulate for service levels of 85%, 90%, 95%, and 100%. Targeted service levels in our simulations are always achieved at a minimum, i.e., we may target a 90% service level but achieve a 92%, since one more stock-out would bring us to an 89% service level, which is not allowed. Furthermore, since we included the *LT* in our replenishment forecasts, we can study its effect during our simulation by changing the *LT* value, assuming the supplier is capable of such a feat.

3.3 Data

The data available from the supplier's VMI arrangement are the demand time series of its customers as seen from two points of view. At a customer's site, the telemetry periodically measures the stock level. Ignoring the stock changes due to deliveries, we take the first difference and aggregate to the daily level. This is the telemetry demand time series and represents the customer's stock usage. On the supplier's side, a historical record of the deliveries sent to customers is maintained. These records are aggregated to the daily level and make the deliveries demand time series. To be concise when distinguishing between the results using these two demand time series, the telemetry demand time series will be referred to by *telemetry* and the historical delivery time series by *delivery*. These two sources of demand data are shown visually on Fig. 2.

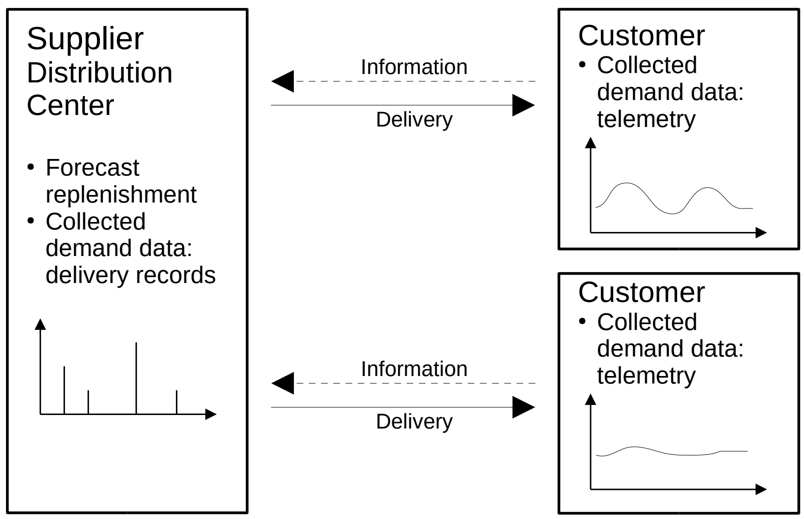

Fig. 2. Supply chain information overview

For both time series, the supplier provided observations for 921 customers from 2015-07-01 to 2016-12-31, i.e., the second half of 2015 and the full year of 2016. From this, we chose for the starting training data the second half of 2015 and for the simulation time window the year of 2016. The simulation begins on 2016-01-01, using as starting value the last stock value measured on 2015-12-31. There is a minimum of two observations of the demand in the second half of 2015 in order to train the forecasting model.

The descriptive statistics for the delivery and consumption time series are shown respectively in Table 1 and Table 2 for the three constituent parts of the Croston intermittent model: demand, interdemand interval, and demand per period. We note the interdemand interval of 1 in the telemetry up to the third quartile. Since the telemetry time series are daily observations of the consumed stock, this indicates that in roughly 75% of observations some amount of stock is used every day. This would initially lead us to believe that most telemetry time series are not intermittent, as they will most likely have a low number of 0 observations.

Table 1
Descriptive statistics of the historical deliveries demand time series

	Demand (units)		Interdemand interval (days)		Demand per period (units/day)	
	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Min.	43.36	23.22	9.19	15.94	3.37	3.64
1st Qu.	76.25	33.48	15.64	18.94	9.02	12.52
Median	93.15	39.57	20.13	22.87	13.10	18.94
3rd Qu.	108.55	44.56	25.08	26.45	19.93	29.52
Max.	157.79	91.58	39.01	35.77	67.63	91.01

Table 2
Descriptive statistics of the telemetry demand time series

	Demand (units)		Interdemand interval (days)		Demand per period (units/day)	
	Mean	Std. Dev.	Mean	Std. Dev.	Mean	Std. Dev.
Min.	2.85	7.08	1	0	2.84	7.08
1st Qu.	7.93	14.35	1	0	7.93	14.34
Median	11.31	18.33	1	0	11.30	18.33
3rd Qu.	15.52	24.95	1.00	0.03	15.52	24.95
Max.	71.60	117.80	1.48	2.28	71.11	117.18

SBC-KH-SES intermittent time series categorization (Petropoulos and Kourentzes, 2015) of the telemetry and delivery time series is shown in Table 3. Time series are categorized based on the coefficient of variation of the demand and the interdemand interval. The table also specifies the optimal model for forecasting between Croston, SBA, and Simple Exponential Smoothing (SES). Time series categorized as Croston or SBA are intermittent. Time series categorized as SES are not intermittent. 100% of delivery time series in our dataset are classified as intermittent compared to only 24% of telemetry time series.

Table 3
SBC-KH-SES intermittent time series categorization of the delivery and telemetry time series

Category	Delivery	Telemetry
Croston	7	66
SBA	914	153
SES	0	702

The collection of point-of-sale data using telemetry does not entirely resolve the issue of working with intermittent data under VMI. However, this intermittent behavior is due almost entirely to variations in the demand amount, as the interdemand interval is 1 in most cases. VMI is proposed as a solution under the belief that the demand collected at or shared by the customer is not intermittent. However, this is not the case if that customer's own demand is also intermittent. Furthermore, there always remains the case that if the data is unavailable, the supplier falls back to using historical delivery records, which are more likely to be intermittent (Murray et al., 2018a).

3.4 Forecasting Model

Core to the VMI framework is a demand forecast. It is from this demand forecast that the replenishment amount and date are determined (eq. 1-2). When choosing the demand forecasting model, we were guided by two considerations. First, the forecasting model must be capable of accommodating both intermittent time series and standard time series. As shown in Table 3, the telemetry and deliveries time series possess different degrees of intermittency. Fortunately, most intermittent time series forecasting models can model non-intermittent data and have been tested on a wide range of intermittent and non-intermittent data (e.g., M-competition) (Kourentzes et al., 2014; Spithourakis et al., 2012).

Second, calculation speed is a practical concern since time is limited. As described in 3.1, the demand forecast can be updated when new demand data is received. In our case study, the telemetry demand data was collected at the end of the day. The forecast is repeated multiple times from the start date until the product's *LT* is reached. Thus, the number of forecasts being performed scales with the number of customers and the number of days in the simulation time frame. This comes out to ~325,000 point forecasts (921 x 365) for each bisection iteration of a supply chain simulation.

The parametric SBA Croston's method was chosen for the forecasting method as it offers fast computation and the most empirical validation (Gardner, 2006; Syntetos et al., 2016; Syntetos and Boylan, 2005; Syntetos et al., 2015). The SBA method is presented in equations (3-5).

$$\hat{Z}_t = \alpha_z z_t + (1 - \alpha_z)\hat{z}_{t-1} \tag{3}$$

$$\hat{X}_t = \alpha_x x_t + (1 - \alpha_x)\hat{x}_{t-1} \tag{4}$$

$\hat{Z}_t$ and $\hat{X}_t$ are the SES forecasts of the non-zero demand amounts and the intervals between non-zero demand periods, respectively.

$$\hat{y}_t = \left(1 - \frac{\alpha_x}{2}\right)\hat{Z}_t/\hat{X}_t \tag{5}$$

To determine the point forecast $\hat{y}_t$, the ratio of the demand amount over the interval is performed with a bias correction.

Both smoothing parameters α_x and α_z are set to the commonly used values of 0.05 (Syntetos and Boylan, 2005). Although it has also been recommended to use different values for the smoothing parameters (Schultz, 1987) or to optimise these values on training data (Kourentzes, 2014), we found that considering the very low number of observations in the training dataset for each customer (minimum of 2 observations in the training data), either option could lead to overfitting the parameters. Furthermore, setting the values of Croston's method makes it quicker to compute compared to more sophisticated alternatives.

Croston's method can be viewed as a more general case of simple exponential smoothing. When the interdemand interval $\hat{X}_t$ is equal to 1, i.e., the time series is non-intermittent, Croston's method reduces itself to simple exponential smoothing. This allows for Croston's method to model both intermittent and non-intermittent time series.

A further computational advantage of Croston's method is that the point forecast $\hat{y}_t$ is only updated when a non-zero demand observation occurs. This means that the forecast can be "valid" for many days if no new demand is observed. In the VMI with delivery data strategy, this means we only need to compute the replenishment once, regardless of *LT*, as no new information is collected until the forecasted replenishment delivery is done which can be used to update the forecast.

3.5 Performance Measurement

The effects of forecast errors on inventory management have been studied for more general use cases (Sanders and Graman, 2009). Under intermittent demand, accuracy measures are known for being misleading (Teunter and Duncan, 2009; Wallström and Segerstedt, 2010). Therefore, as prescribed in other studies on intermittent demand stock control accuracy (Kourentzes et al., 2020; Syntetos et al., 2010), the performance is evaluated only on the stock.

When evaluating the impact of an inventory strategy on stock, what needs to be shown is the relationship between the costs of implementing the strategy versus the inventory outcome. Trade-off curves offer a visual understanding of the trade-off between total inventory investment and the outcome (Gardner, 1990, 2006). Plotting multiple trade-off curves on the same graph makes it possible to compare the performance of different scenarios.

The result of the VMI simulation is the optimal replenishments and the safety stock level required to achieve a targeted service level for each customer. The cost of the VMI strategy contains 2 aspects: the deliveries and the inventory. Inventory costs are measured as the total stock which must remain on site for all customers. Delivery costs are the number of deliveries. Trade-off curves are presented separately for these two costs. On each trade-off curve, we superimpose the result for each simulation scenario.

Another useful measurement of the inventory performance is the exact safety stock level required to achieve a targeted service level (Eaves and Kingsman, 2004). This level is a fraction of the maximum stock allowed on a customer's site. In our VMI simulations, the safety stock level is calculated for each customer. To present these results, we will show the distribution of the safety level under each information scenario for targeted service levels using a box plot.

4 Empirical Evaluation and Discussion

The VMI arrangement was simulated for targeted service levels of 85%, 90%, 95%, and 100% under both telemetry and delivery demand information with a product *LT* of 3 and 7. As described in 3.4, a shorter *LT* has no effect when the demand data is the historical deliveries with Croston's forecasting method; it is thus omitted from the results. A shorter lead time can be understood as allowing the supplier to gain more demand knowledge before performing a replenishment. The question these results answer is the impact of demand information strategies on the VMI costs as explored through data acquisition strategies and the *LT*. Under each data scenario, the simulation returns the replenishment dates and amounts, the service interruptions for service levels of less than 100%, and the safety stock level required to achieve the service level.

In the results figures, the results for different data strategies are labelled as follows: the demand data collected through a customer's point-of-sale telemetry system is labelled as *telemetry*, and the demand data collected by the supplier's aggregated delivery records is labelled as *delivery*. The *LT* is indicated to the right of the data label.

Firstly, we were successful in determining replenishment and safety stock levels under intermittent delivery information for all our customers. This shows the possibility of operating a VMI using solely a supplier's historical deliveries without information sharing or telemetry data collection. It is possible even while ensuring a 100% service level. This is important, as VMIs, like any supply chain arrangement, are imperfect (Colicchia et al., 2019). Information may be missing or specific partners may be unwilling to provide all the required information (Kembro and Näslund, 2014; Wang et al., 2014). Although specialized forecasting models have been suggested to perform point forecasts of the demand in situations where data is missing, their use has not been evaluated empirically (Murray et al., 2018a). Our results offer such an empirical validation under a complex inventory management arrangement like VMI. Even without information sharing or collection, intermittent demand forecasting makes it possible and even worthwhile to engage in VMI.

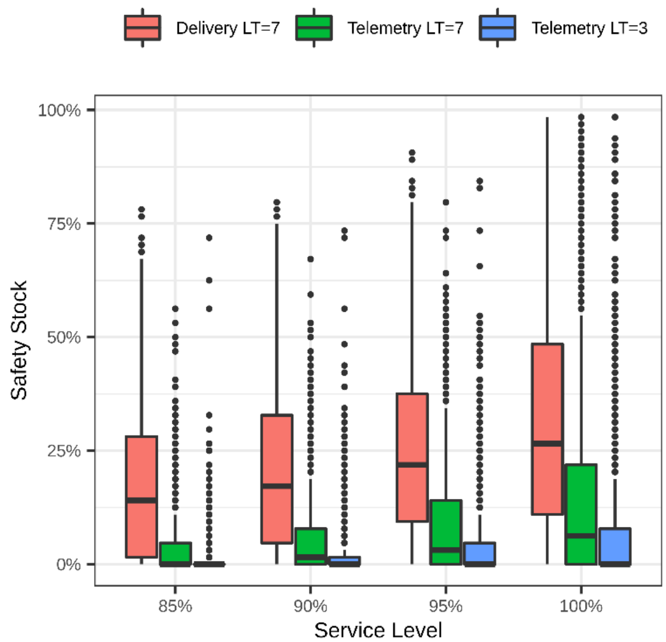

Fig. 3. Exact safety stock level distribution

Fig. 3 presents the distribution of the safety stock level as a percentage of the maximum allowed stock on site for a targeted service level. For a service level of 100% (most common in VMI), the mean safety stock is halved from 32% to 17% when going from delivery data to telemetry data. Lowering the product's lead time further halves the mean safety stock from 17% to 9%. These ratios are roughly maintained for the other service level values. However, in absolute terms, the improvement is greatest for higher service levels, since the mean safety stock is smaller for smaller values of service levels.

Looking at outliers in all three simulations reveals similarities in both their number and their required safety stock. For a given service level, the furthest outlier is similar in each configuration. This behavior is more obvious for the outliers at a service level 1 where certain customers' demand is intermittent enough to require nearly 100% safety stock to ensure the service level under all information scenarios. The number of outliers also remains consistent for each information scenario across the different service level. For example, there are only a couple of outliers for the delivery data at each service level. These two observations indicate that this is an issue directly with the demand of some specific customers. Their demands are so lumpy, that regardless of the data, they are unforecastable. Individual cost benefit of outlying customers may indicate to the supplier that a VMI arrangement is unprofitable with them. The supplier may then decide to return to a more standard arrangement or ask for more information sharing from those customers.

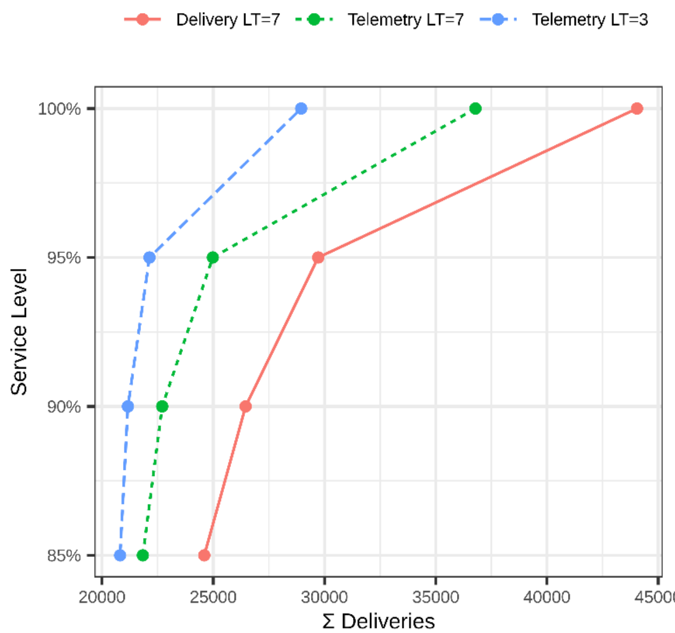

Fig. 4. Number of deliveries vs. service level

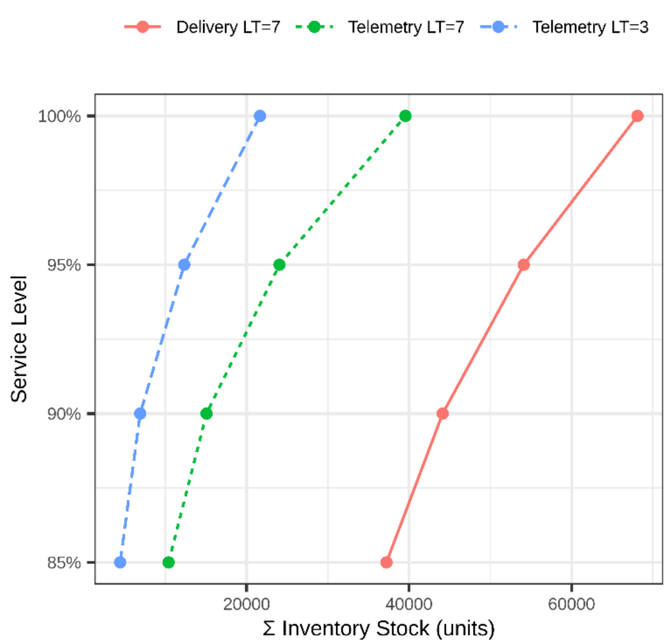

Fig. 5. Inventory vs. service level

Fig. 4 and Fig. 5 present the trade-off curves for the inventory stock and the number of deliveries versus the service level. For both performance measurements at all service levels, telemetry data is better than delivery data and lower lead times are better. For a service level of 100%, telemetry data offers an improvement of 16% in the number of deliveries and of 43% in inventory stock over delivery data. Lowering the lead time from 7 to 3 offers a further improvement of 21% in the number of deliveries and of 45% in inventory stock. There are greater improvements in inventory stock than in the number of deliveries for all service levels between the information strategies. Intuitively, there will always be a minimum number of deliveries for customers even under perfect information. However, the inventory stock nears 0 with perfect information, i.e., no safety level is necessary with perfect forecasts. Telemetry data lowers the number of deliveries and inventory stock since there are fewer deliveries of greater amounts.

The inventory stock amount can be translated directly into an inventory cost by multiplying the amount by the product's price. In our context, there were only 2 different products being delivered of a comparable price. As such, the transformation from stock to cost is linear. Determining the cost for deliveries is much more complicated. A full analysis of the delivery costs would include, but not be limited to, routing and truck sizing (Borade and Sweeney, 2015). Unfortunately, this was not done as the required data was not available to us. Regardless, we can safely claim that lowering the number of deliveries will lower the total delivery costs. Even without directly calculating the costs of the deliveries and inventory stock, we can still offer the following recommendations.

Since the cost reduction is greater for inventory stock than deliveries, that is to say, the major cost improvements under a VMI are the inventory, we differ in the conclusion of (Ru et al., 2018) that savings are more important if customer inventory holdings costs are low. Greater savings will be had by a VMI if a customer's inventory holding costs are high. This is in continuation with the more standard recommendation in the literature (Kim, 2008). In this specific case study, our industrial partner's stock holdings costs are low, as specialized standardized containers are required to store the product. Thus, the deciding factor for managers when doing a cost analysis is most likely to be industry specific. Perhaps this is what led to the opposing recommendations between (Ru et al., 2018) and (Kim, 2008).

Performance improvements between the different scenarios is greater for higher service levels. Increased improvements at higher service levels is similar to previous inventory management strategy comparisons (Syntetos et al., 2015). However, there is one case in our results for which this is not respected. The difference in inventory stocks between delivery and telemetry data at a *LT* of 7 is consistent for all service levels. This has the managerial implication that acquiring improved data is more valuable at higher service levels when deliveries are the main cost driver, and valuable at all service levels when the inventory is the main cost driver.

We note the sharper increase in the number of deliveries required to achieve a 100% service level compared to the inventory stock. The managerial implication would be that more delivery savings are to be had if the service level can be lowered even slightly from 100%. In our case, 95% may be as low as a single stock-out during the whole year. A strategy under VMI may be to negotiate a policy with a stock-out penalty that still results in a cost saving when weighed against the reduced delivery and stock costs.

The safety stock level increased with the service level. We can also deduce that the number of deliveries increases as a function of the safety stock level (simply flip the variables of Fig. 4). This implies that lowering the safety stock level can reduce the number of deliveries, as there will be fewer deliveries but of each delivery will be larger. Thus, it is possible to reduce the safety stock level while still maintaining the same service level by increasing the maximum order-up-to-level at a customer's site. For example, doubling the maximum order-up-to-level reduces the safety stock level by half while the stock investment is the same. This strategy may help reduce the frequency of deliveries to this specific customer, but the deliveries will be larger. Furthermore, the cost of increasing the maximum order-up-to-level must also be considered. In our context, this would require installing a second container or replacing it with a larger one. This strategy would be more useful for customers with high safety stock levels that can be identified using distributions like the ones shown on Fig. 3.

5 Conclusion

Good information is critical to produce good forecasts. When forecasts have large impacts on inventory policy, the need for good information only increases. Technological or strategic decisions on how to acquire information then becomes paramount. This paper simulated a VMI arrangement under different information strategies. Only a single forecasting model was used to isolate and compare the influence of the information strategy. Regardless of the chosen information strategy, intermittent time series were present and challenging. Using an SBA Croston forecasting model to determine the replenishment amounts and dates, demand data from a customer's point-of-sale lowered the number of deliveries by 16% and the inventory stock by 43% for a targeted service level of 100% compared to using demand data from a supplier's delivery records. For lower service levels, lesser but consistent improvements between information strategies is maintained. Lowering product lead time to acquire more up-to-date demand information also reduces the number of deliveries and inventory stock.

The described VMI framework in this paper allows for practitioners to implement a VMI style arrangement even without any information gathering. This is sometimes necessary when the required information is not made available to supply chain members—an inherent risk of any information sharing arrangement (Colicchia et al., 2019). This is an enhancement to previous methods that have focused on purely improving intermittent demand forecasts under missing information (Murray et al., 2018a).

In the presented case, the performance improvements were larger in terms of inventory stock compared to the number of deliveries. We would thus recommend that suppliers will see more benefit for supply chains with higher inventory costs. This recommendation is in line with those given by (Kim, 2008). Furthermore, better information offers significantly more improvements for higher targeted service levels.

Regardless of the chosen information strategy, intermittent demand remains an important challenge that industries will continue to face (Nikolopoulos, 2021). Our work reinforces the use of Croton's model for inventory management under intermittent demand in terms of performance and for practical considerations. We concur with the recommended use of parametric models both for computational speed and their robustness (Syntetos et al., 2015). Our work also reinforces the growing consensus on the value of information sharing in supply chain management (Cao and Zhang, 2011; Jung et al., 2005; Zhou et al., 2017).

The scope of this paper was limited to the comparison of demand information in intermittent supply chain forecasting. We did not tackle the challenge of comparing different forecasting models. Performing the simulation with different forecasting models would have allowed us to compare them, but multiple papers have already compared intermittent demand forecasting models (Hasni et al., 2019a; Mukhopadhyay et al., 2012; Syntetos et al., 2015) and our computing time was limited. The forecasting model remains of practical importance as a more accurate forecast of the reorder point means improved delivery and inventory performance for the same price of the demand data collection strategy.

Collecting improved data on demand either through data sharing or technological means between a customer and a supplier is also known to reduce noise in the data caused by the Bullwhip effect (Gang et al., 2017; Jeong and Hong, 2019). The link between the Bullwhip effect and intermittent demand has been discussed (Murray et al., 2018a). Our research only considers this effect as another source of noise which is present in a supplier's delivery records demand data that must be overcome with better forecasting.

Future work in VMI under intermittent demand may look at extending the idea proposed in (Kourentzes et al., 2020) to optimise the forecasting model directly on the VMI results. We believe that this would also require a simulation framework for the VMI arrangement. The framework proposed here can be used as a starting point.

Finally, we are left pondering the issue of whether or not changes to the supply chain arrangement may also affect the intermittent behavior of the demand. Implementing a new replenishment policy may cause changes in the observed demand which would then change the replenishment policy. This type of feedback interaction could be analyzed from a systems theory angle.

Acknowledgments

The authors would like to acknowledge our industrial partner and the National Science and Engineering Research Council of Canada (NSERC) for funding under grant RDCPJ 492021-15, RGPIN-2019-04723. We also thank the members of the Data Intelligence Laboratory of Polytechnique Montreal for their constructive comments that helped the authors to make this research possible.

References

Achabal, D.D., McIntyre, S.H., Smith, S.A., Kalyanam, K., 2000. A decision support system for vendor managed inventory. Journal of Retailing 76, 430-454.

Ali, M.M., Babai, M.Z., Boylan, J.E., Syntetos, A.A., 2017. Supply chain forecasting when information is not shared. European Journal of Operational Research 260, 984-994.

Angulo, A., Nachtmann, H., Waller, M.A., 2004. Supply chain information sharing in a vendor managed inventory partnership. Journal of business logistics 25, 101-120.

Bartezzaghi, E., Verganti, R., Zotteri, G., 1999. A simulation framework for forecasting uncertain lumpy demand. International Journal of Production Economics 59, 499-510.

Borade, A.B., Sweeney, E., 2015. Decision support system for vendor managed inventory supply chain: a case study. International Journal of Production Research 53, 4789-4818.

Cao, M., Zhang, Q., 2011. Supply chain collaboration: Impact on collaborative advantage and firm performance. Journal of Operations Management 29, 163-180.

Colicchia, C., Creazza, A., Noè, C., Strozzi, F., 2019. Information sharing in supply chains: a review of risks and opportunities using the systematic literature network analysis (SLNA). Supply Chain Manag 24.

Croston, J.D., 1972. Forecasting and Stock Control for Intermittent Demands. Operational Research Quarterly (1970-1977) 23, 289-303.

Eaves, A.H., Kingsman, B.G., 2004. Forecasting for the ordering and stock-holding of spare parts. Journal of the Operational Research Society 55, 431-437.

Fu, W., Chien, C.-F., 2019. UNISON data-driven intermittent demand forecast framework to empower supply chain resilience and an empirical study in electronics distribution. Computers & Industrial Engineering 135, 940-949.

Gang, L., Gang, Y., Shouyang, W., Hong, Y., 2017. Bullwhip and anti-bullwhip effects in a supply chain. International Journal of Production Research 55, 5423-5434.

Gardner, E.S., 1990. Evaluating forecast performance in an inventory control system. Management Science 36, 490-499.

Gardner, E.S., 2006. Exponential smoothing: The state of the art—Part II. International journal of forecasting 22, 637-666.

Hasni, M., Aguir, M.S., Babai, M.Z., Jemai, Z., 2019a. On the performance of adjusted bootstrapping methods for intermittent demand forecasting. International Journal of Production Economics 216, 145-153.

Hasni, M., Aguir, M.S., Babai, M.Z., Jemai, Z., 2019b. Spare parts demand forecasting: a review on bootstrapping methods. International Journal of Production Research 57, 4791-4804.

Hasni, M., Babai, M.Z., Aguir, M.S., Jemai, Z., 2019c. An investigation on bootstrapping forecasting methods for intermittent demands. International Journal of Production Economics 209, 20-29.

Holweg, M., Disney, S., Holmström, J., Småros, J., 2005. Supply Chain Collaboration:: Making Sense of the Strategy Continuum. European Management Journal 23, 170-181.

Jeong, K., Hong, J.-D., 2019. The impact of information sharing on bullwhip effect reduction in a supply chain. Journal of Intelligent Manufacturing 30, 1739-1751.

Johnston, F.R., Boylan, J.E., Shale, E.A., 2003. An examination of the size of orders from customers, their characterisation and the implications for inventory control of slow moving items. Journal of the Operational Research Society 54, 833-837.

Jung, S., Chang, T., Sim, E., Park, J., 2005. Vendor Managed Inventory and Its Effect in the Supply Chain, in: Baik, D.-K. (Ed.), Systems Modeling and Simulation: Theory and Applications. Springer Berlin Heidelberg, pp. 545-552.

Kembro, J., Näslund, D., 2014. Information sharing in supply chains, myth or reality? A critical analysis of empirical literature. International Journal of Physical Distribution & Logistics Management 44, 179-200.

Kim, H.-S., 2008. Research note—revisiting "retailer-vs. vendor-managed inventory and brand competition". Management Science 54, 623-626.

Kourentzes, N., 2013. Intermittent demand forecasts with neural networks. International Journal of Production Economics 143, 198-206.

Kourentzes, N., 2014. On intermittent demand model optimisation and selection. International Journal of Production Economics 156, 180-190.

Kourentzes, N., Barrow, D., Petropoulos, F., 2019. Another look at forecast selection and combination: Evidence from forecast pooling. International Journal of Production Economics 209, 226-235.

Kourentzes, N., Petropoulos, F., Trapero, J.R., 2014. Improving forecasting by estimating time series structural components across multiple frequencies. International Journal of Forecasting 30, 291-302.

Kourentzes, N., Trapero, J.R., Barrow, D.K., 2020. Optimising forecasting models for inventory planning. International Journal of Production Economics 225, 107597.

Lolli, F., Gamberini, R., Regattieri, A., Balugani, E., Gatos, T., Gucci, S., 2017. Single-hidden layer neural networks for forecasting intermittent demand. International Journal of Production Economics 183, 116-128.

Mukhopadhyay, S., Solis, A.O., Gutierrez, R.S., 2012. The Accuracy of Non-traditional versus Traditional Methods of Forecasting Lumpy Demand. Journal of Forecasting 31, 721-735.

Murray, P.W., Agard, B., Barajas, M.A., 2018a. ASACT - Data preparation for forecasting: A method to substitute transaction data for unavailable product consumption data. International Journal of Production Economics 203, 264-275.

Murray, P.W., Agard, B., Barajas, M.A., 2018b. Forecast of individual customer's demand from a large and noisy dataset. Computers & Industrial Engineering 118, 33-34.

Nikolopoulos, K., 2021. We need to talk about intermittent demand forecasting. European Journal of Operational Research 291, 549-559.

Nikolopoulos, K., Syntetos, A.A., Boylan, J.E., Petropoulos, F., Assimakopoulos, V., 2011. An aggregate-disaggregate intermittent demand approach (ADIDA) to forecasting: an empirical proposition and analysis. Journal of the Operational Research Society 62, 544-554.

Persson, F., Axelsson, M., Edlund, F., Lanshed, C., Lindstrom, A., Persson, F., 2017. Using simulation to determine the safety stock level for intermittent demand, in: Chan, V., Dambrogio, A., Zacharewicz, G., Mustafee, N. (Eds.), 2017 Winter Simulation Conference. Ieee, New York, pp. 3768-3779.

Petropoulos, F., Kourentzes, N., 2015. Forecast combinations for intermittent demand. Journal of the Operational Research Society 66, 914-924.

Ru, J., Shi, R., Zhang, J., 2018. When Does A Supply Chain Member Benefit from Vendor-Managed Inventory? Production and operations management 27, 807-821.

Sanders, N.R., Graman, G.A., 2009. Quantifying costs of forecast errors: A case study of the warehouse environment. Omega 37, 116-125.

Scala, N.M., Rajgopal, J., Needy, K.L., 2013. A base stock inventory management system for intermittent spare parts. Military Operations Research 18, 63-77.

Schultz, C.R., 1987. Forecasting and inventory control for sporadic demand under periodic review. Journal of the Operational Research Society 38, 453-458.

Sillanpää, V., Liesiö, J., 2018. Forecasting replenishment orders in retail: value of modelling low and intermittent consumer demand with distributions. International Journal of Production Research 56, 4168-4185.

Spithourakis, G.P., Petropoulos, F., Nikolopoulos, K., Assimakopoulos, V., 2012. A systemic view of the ADIDA framework. IMA Journal of Management Mathematics, dps031.

Syntetos, A.A., Babai, Z., Boylan, J.E., Kolassa, S., Nikolopoulos, K., 2016. Supply chain forecasting: Theory, practice, their gap and the future. European Journal of Operational Research 252, 1-26.

Syntetos, A.A., Boylan, J.E., 2005. The accuracy of intermittent demand estimates. International Journal of Forecasting 21, 303-314.

Syntetos, A.A., Nikolopoulos, K., Boylan, J.E., 2010. Judging the judges through accuracy-implication metrics: The case of inventory forecasting. International Journal of Forecasting 26, 134-143.

Syntetos, A.A., Zied Babai, M., Gardner, E.S., 2015. Forecasting intermittent inventory demands: simple parametric methods vs. bootstrapping. J Bus Res 68, 1746-1752.

Teunter, R.H., Duncan, L., 2009. Forecasting intermittent demand: a comparative study. Journal of the Operational Research Society 60, 321-329.

Trapero, J.R., Cardós, M., Kourentzes, N., 2019. Quantile forecast optimal combination to enhance safety stock estimation. International Journal of Forecasting 35, 239-250.

Vigtil, A., 2007. Information exchange in vendor managed inventory. International Journal of Physical Distribution & Logistics Management 37, 131-147.

Wallström, P., Segerstedt, A., 2010. Evaluation of forecasting error measurements and techniques for intermittent demand. International Journal of Production Economics 128, 625-636.

Wang, Z., Ye, F., Tan, K.H., 2014. Effects of managerial ties and trust on supply chain information sharing and supplier opportunism. International Journal of Production Research 52, 7046.

Willemain, T.R., Smart, C.N., Shockor, J.H., DeSautels, P.A., 1994. Forecasting intermittent demand in manufacturing: a comparative evaluation of Croston's method. International Journal of Forecasting 10, 529-538.

Wu, M.-C., Hsu, Y.-K., 2008. Design of BOM configuration for reducing spare parts logistic costs. Expert Systems with Applications 34, 2417-2423.

Zhou, C., Viswanathan, S., 2011. Comparison of a new bootstrapping method with parametric approaches for safety stock determination in service parts inventory systems. International Journal of Production Economics 133, 481-485.

Zhou, M., Dan, B., Ma, S., Zhang, X., 2017. Supply chain coordination with information sharing: The informational advantage of GPOs. European Journal of Operational Research 256, 785-802.