



Stock control analytics: a data-driven approach to compute the fill rate considering undershoots

Eugenia Babiloni¹ · Ester Guijarro¹ · Juan R. Trapero²

Received: 29 September 2021 / Revised: 19 September 2022 / Accepted: 27 September 2022
© The Author(s) 2023

Abstract

One of the most frequently used inventory policies is the order-point, order-up-to-level (s, S) system. In this system, the inventory is continuously reviewed and a replenishment request is placed whenever the inventory position drops to or below the order point, s . The variable replenishment order quantity and the variable replenishment cycle characterize the system by the use of complex mathematical computations. Different methodological approaches diminish the mathematical complexity by neglecting the undershoots, i.e., the quantity that the inventory position is below the order point when it is reached. In this paper, we conceptually and empirically analyse the bias that neglecting the undershoots introduces into the estimation of the fill rate. After that, we suggest a new methodology developed under a data-driven perspective that uses a state-dependent parameter algorithm to correct such a bias. As a result, we propose two new methods, one parametric and the other nonparametric, to enhance the fill rate estimate. Both methods, named analytics fill rate methods, remove the bias that neglecting the undershoots introduces and are used to illustrate the practical implications of this hypothesis on the performance and design of the (s, S) system. This research is developed in a lost sales context with simulated stochastic and i.i.d. discrete demands as well as actual sales data.

Keywords Inventory · Fill rate · Lost sales · Undershoots · State-dependent parameter

✉ Ester Guijarro
esguitar@doe.upv.es

¹ Dept. Organización Empresas, CEGEA (Centre of Business Management Research), Universitat Politècnica de València, 46022 Valencia, Spain

² Dept. Administración de Empresas, Universidad de Castilla-La Mancha, Ciudad Real, Spain

1 Introduction

Inventory management plays a critical role in the success of most enterprises because of its direct implication for decision-making both at the operational and strategic levels. In the current markets, with more demanding customers and higher costs, effective inventory management systems are essential. Therefore, developing an intelligent inventory management system and transforming inventory data into valuable insights becomes a key factor in achieving competitive advantages (Stefanovic 2015; Zhou et al. 2017).

The core purpose of inventory management is to address two important key issues: (1) when to launch a replenishment order, and (2) the size of this replenishment order. The control parameters of inventory policies are devoted to determining these issues according to a specific objective, normally using a cost or a customer service criterion. Even though the cost approach often leads to less mathematical complexity, the use of the service criterion avoids the difficulty of accurately estimating the stockout cost due to the presence of intangible components such as the loss of customer goodwill and the diminishing of future business (Gutgutia and Jha 2018) opportunities. For that reason, this paper follows a customer service approach that introduces a control parameter known as the service level that becomes a constraint in determining the control parameters of the system (Silver et al. 2017; Escalona et al. 2019). The service level addressed in this paper is the fraction of the demand that is immediately satisfied from the stock without causing a shortage (Brown 1962), which is known as the fill rate. Despite its simple definition, its calculation hides complex mathematical nuances, especially when the system is managed by a continuous review system (the status of the inventory is known at any given moment) and unfilled demand is lost, as happens in highly competitive sectors such as retail, machinery spare parts, service sector, or e-commerce of fixed-date services (Gruen et al. 2002; Breugelmans et al. 2006; Diels and Wiebach 2011; Disney et al. 2021).

This paper focuses on the order-point, order-up-to-level (s, S) system due to its versatility and extended use in practice (Caplin and Leahy 2010; Bijvank et al. 2015). By definition, this system assumes continuous review and, in a stochastic demand context, both the replenishment cycle, i.e., time elapsed between two consecutive order deliveries, and the order quantity, are variable; therefore, simultaneous determination of the parameters is necessary to guarantee the mathematical optimality of the policy once a service objective is set. However, as Silver et al. (2017) point out, in practice, the most common approach consists of assuming that all the demand transactions are unit sized (see for example Silver 1970; Vincent 1983; Platt et al. 1997; Agrawal and Seshadri 2000; Axsater 2000; Larsen and Thorstenson 2014). This means that the inventory position always exactly reaches the order point, s (known also as reorder point or ROP), and therefore, the order quantity is constant and equal to the order-up-to-level, S , minus the order point, s , (i.e. $S - s$). However, in real situations, the inventory position may not exactly be at the ROP but at a certain amount below it, known as the undershoot at ROP. Neglecting undershoots greatly reduces the mathematical

complexity of the problem but can severely reduce the service level of the system (Schneider 1981), which leads to increased stockout costs, especially when managing lumpy or erratic demand items for which the unitary demand assumption is not appropriate (Cardós and Babiloni 2011a).

In the research literature, we find several studies that address the complexity of estimating undershoots. We find quite a number of works that calculate the undershoots based on the asymptotic results of renewal theory (Schneider 1981; Tijms and Groenevelt 1984; Baganha et al. 1999; Johansen and Hill 2000; Kouki et al. 2009). Other authors, in contrast, compute probability distribution functions to estimate the expected value of the undershoots and incorporate it into service measures (Cohen et al. 1988; Baganha et al. 1996; Gutierrez and Rivera 2021). However, in general, these studies approximate the value of the undershoots, which may have an impact on many aspects of inventory management. Regarding the fill rate estimations in service models, Schneider (1981), Tijms and Groenevelt (1984), Moors and Strijbosch (2002), Silver et al. (2009) and Silver et al. (2012) suggest different approximations of the fill rate that consider undershoots under the backordering assumption. In the lost sales context, only Cohen et al. (1988) present a fill rate method in a order-point, order-up-to-level (s, S) system; however in this approach, the undershoot estimations are calculated from an approximate distribution function that may bias the fill rate estimation.

Then, although the (s, S) system is widely used in practice, there is no optimal approach to establish the control parameters when the fill rate is the service constraint of the system and the unfilled demand is lost. The objective of this research is to propose a new methodology to obtain the fill rate based on a bias analysis and a correction of the undershoot assumption. The methodology follows a data-driven perspective that avoids any assumption regarding the undershoot distribution and is suitable for service models.

The implementation of business intelligence affects the agile efficiency of the supply chain (Kaur 2021). In general, the digitalization of companies offers a large amount and variety of data (Big Data) that transforms the way supply chain management works (Wang et al. 2016). In particular, Gandomi and Haider (2015) indicate the need to develop appropriate and efficient methods to leverage such data, and they review big data analytics areas such as text analytics or predictive analytics. In this sense, this work is an example of how to extract more value from inventory data and proposes use of a novel area called stock control analytics, where data-driven methods, which include machine learning and artificial intelligence techniques, can enhance the state-of-the art inventory management methods.

Although the nomenclature may be new, the use of data analytics tools to overcome the limitations associated with theoretical inventory models is not that new. For instance, Strijbosch et al. (2000) correct the bias introduced by using demand forecasts instead of statistical demand distribution first moments, which are unknown prior to the commencement of any inventory processes. That article proposes a heuristic method based on simulations, which substantially reduces the bias. Novel methods based on this data-driven perspective have been recently published to deal with the uncertainty of demand distributions. Huber et al. (2019) revisit the newsvendor problem subject to a target cycle service level by

employing machine learning and quantile regression approaches that circumvent the need to assume a statistical distribution for the demand. Regarding the uncertainty of statistical demand distribution, Trapero et al. (2019a) investigate parametric and nonparametric methods to enhance the safety stock estimation problem for a certain target cycle service level. Their results show that nonparametric approaches such as kernel density estimators provide a good performance in terms of cost and service for lower lead time values, and parametric approaches such as GARCH outperform the rest of the methods for higher lead time values. Trapero et al. (2019b) continue with this line of research by proposing a combination scheme of GARCH and Kernel whose weight parameters are determined by minimizing the tick-loss (newsvendor) linear asymmetric cost function. Additionally, regarding the performance of machine learning forecasting methods versus traditional counterparts, Spiliotis et al. (2020) compared both methods for daily SKU demand forecasting. The work highlights the improvements possible with machine learning methods.

Previous articles show how data-driven approaches offer versatile solutions to address the limitations of traditional theoretical assumptions. In this work, following the same philosophy, we focus on the assumption of neglecting the undershoots on the (s, S) system when the demand is modelled by any discrete distributions for a lost sales case. To the best of the authors' knowledge, this particular topic has been overlooked in the literature. Basically, this research aims to determine the extent of the bias introduced by such an assumption and proposes a methodology to reduce it. The data-driven approach employed to cope with this problem is the state-dependent parameter (SDP) estimation technique. SDP estimation involves the nonparametric identification of the state dependency using recursive methods for time variable parameter estimation, which allow for rapid (state dependent) parametric change (Young et al. 2001). SDP estimation has been successfully applied in a supply chain context in Trapero et al. (2011), particularly to correct judgemental forecast biases. Once the nonparametric estimation between the state and the parameter is available, another advantage of the SDP procedure is that it reveals any potential nonlinear pattern between the state and the parameter and can thus be further parametrized, obtaining a fully self-contained model with just a few parameters (Young et al. 2001). Following this methodology, we propose two methods to estimate the fill rate, one parametric and another nonparametric; the methods remove the bias and outperform the initial approach. These methods are named analytics fill rate methods, by the authors, to distinguish them from the conventional formulas.

The rest of this paper is organized as follows. Section 2 describes the problem the paper addresses and states the main assumptions and the notations. Section 3 introduces the initial approach of the fill rate formula, which neglects the presence of undershoots. This section provides insight of the potential bias introduced by neglecting the undershoots. In Sect. 3, the experimental setup is also detailed. Section 4 explores the SDP approach to correct the bias associated with the initial fill rate estimation and proposes the new analytics fill rate methods. Section 5 assesses the methodology using a case study. Practical implications and a discussion of the results are presented in Sect. 6. Finally, the main conclusions are drawn in Sect. 7.

2 Problem formulation and description

2.1 System description, notation, and assumptions

This paper considers a single-echelon, single-item inventory system where the demand is stochastic. The system is modelled by a discrete time model where the demand arriving during a given interval of time comes from a discrete distribution. The stock is controlled following a continuous review order-point, order-up-to-level (s, S) system for the lost sales case. In this system, when the inventory position is at or below the ROP, a sufficient amount is ordered to raise the inventory position up to the order-up-to-level, S . The replenishment order is received L units of time after being launched. We consider that the replenishment cycle (also referred to as just, the cycle, further on) is the time between two consecutive deliveries and starts at order delivery. Figure 1 shows an example of the evolution of the on-hand stock and the inventory position when the system is not out of stock (a) and when it is out of stock (b).

The notations in Fig. 1 and in the rest of the paper are as follows:

- s = order point, ROP (units),
- S = order-up-to-level (units),
- L = lead time (time),
- τ = elapse of time from the beginning of the cycle to the ROP is reached (time),
- z_t = on-hand stock at t (units),
- d_t = demand at instant t (units),
- D_t = accumulated demand over t (units),
- $f_L(\cdot)$ = probability function of demand during L ,
- X^+ = maximum $[X, 0]$ for any expression X ,

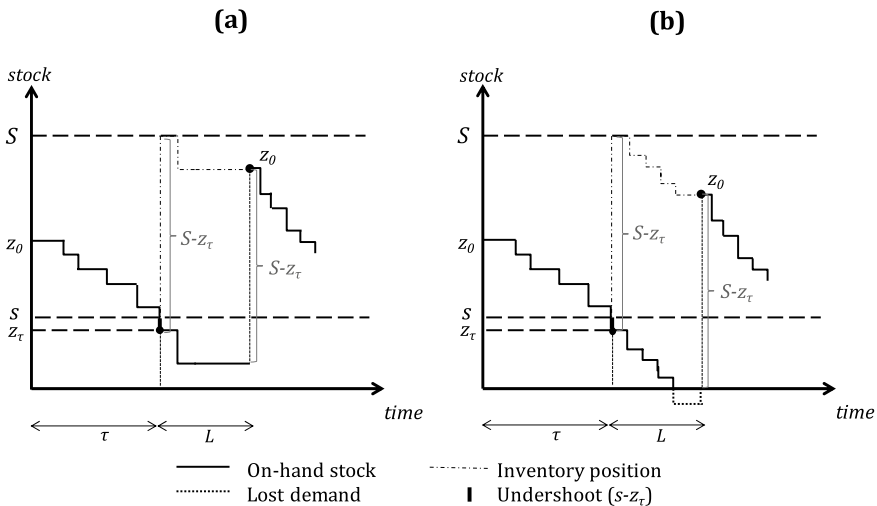


Fig. 1 Evolution of the stock in a (s, S) inventory policy and lost sales when the system is not out of stock (a) and when it is out of stock (b)

Q_t = order quantity at instant t (units),
 β_A = achieved fill rate,
 β_0 = initial fill rate,
 β_i = estimated fill rate for method i , $i = LR, SDP$.

The general assumptions of this paper are as follows: (1) the lead time, L , is constant and known; (2) there is never more than one outstanding order, which implies that the condition $s < S - s$ is always true; (3) the demand process is considered stationary and i.i.d., and is defined by any discrete distribution function; (4) unfilled demand is lost and; (5) z_t is the on-hand used to satisfy d_t . Note that assumption (3) is generally accepted in inventory research since, if it is not the case, the numerical difficulties are insurmountable (Schneider 1981), especially for the lost sales case (see, for example, Hadley and Whitin 1963; Cohen et al. 1988; Johansen 2001).

This research considers that time is discrete and is organized into a numerable and infinite succession of equally spaced instants, which is tantamount to assuming that each transaction is updated in a specific instant of time considering a sequence of discrete and infinitely small intervals. This procedure is indeed the way IT systems that are used to control continuous inventories work, i.e., they report every transaction that occurs in a specific instant of time.

2.2 Problem description

The (s, S) system implicitly states that both the replenishment cycle and the order quantity are variable. These two characteristics make it truly difficult to find a mathematical approach with the potential for a practical implementation.

First, the size of the replenishment order depends on the stock level when the ROP is reached, i.e., $Q = S - z$. In a stochastic context, to compute the expected value of z_τ , bearing in mind that the evolution of the on-hand stock can be modelled as an ergodic Markov chain, it is necessary to know the probability transition matrix of the on-hand stock between two consecutive replenishment cycles \overline{M} in such a way that $\lim_{n \rightarrow \infty} \overline{P(z)} \cdot \overline{M}^n = \overline{P(z_0)}$, where $\overline{P(z)}$ is an arbitrary vector and $\overline{P(z_0)}$ is the principal left eigenvector of \overline{M} , i.e., the probability vector of the on-hand stock levels at order delivery. This procedure is very time-consuming, especially for large values of S , since it determines the dimension of the transition matrices (Cardós et al. 2017). If the undershoots are negligible such that $z_\tau = s$, the order quantity is simply $Q = S - s$.

Second, the estimation of the total demand and concretely D is not straightforward and depends on the expected values of the stock levels and on the unknown number of periods until the ROP is reached, τ . To delve into this problem, please refer to Cardós and Babiloni (2011b). As a result, the replenishment cycle is variable, which, in a stochastic demand context, hinders the calculation of the demand distribution from the order delivery until the ROP is reached. Nevertheless, if the undershoots are neglected, the demand consumed from the beginning of the cycle until the order point is reached is the difference between the on-hand stock at the beginning of the cycle and the ROP.

For these reasons, in practice, “the values of the control parameters are set in a rather arbitrary fashion” (Silver et al. 2017), and the common derivation of the (s, S) system consists of assuming that all the demand transactions are unit sized or, stated another way, that undershoots at ROP are not possible.

However, what does neglecting undershoots mean? The answer is related to the expected value of the stockouts (also known as expected shortages) during the replenishment cycle. The assumption that the replenishment order is launched at the precise moment the ROP is reached implies that the on-hand stock at launch may be greater than it actually is, since $z_\tau \leq s$, in such a way that there is a higher probability of stockouts than expected. In the following sections, we describe the bias that neglecting undershoots introduces into the estimation of the fill rate as a consequence of underestimating the stockout probability of the system; we then propose a data-driven approach that eliminates this bias. The results are validated using a simulation since there is not an exact method able to compute the fill rate in the context of our research.

3 The bias of the initial approach of the fill rate

3.1 Conceptual analysis of the bias

The fill rate is commonly defined as the fraction of demand that is immediately fulfilled from the on-hand stock. However, this simple-looking definition hides many technical details and nuances that are often overlooked in the literature. In fact, the fill rate has been simplified through the traditional approximation which computes the fill rate in terms of units short, i.e., as the complement of the ratio between the expected shortage per replenishment cycle (ESPRC) and the total expected demand per replenishment cycle (EDPRC) (Bijvank and Vis 2011, 2012; Guijarro et al. 2012; Babiloni and Guijarro 2020), i.e.,

$$\beta = 1 - \frac{ESPRC}{EDPRC}$$

Although the definition of the fill rate does not introduce any assumption regarding the inventory system, as pointed out before, one possible approach to avoid the computational complexity that undershoots introduce is to assume that they are negligible at ROP. This assumption implies that, (1) the order quantity is constant and, (2) the system will only be out of stock during the lead time. Figure 2 shows the evolution of the stock levels when the undershoots are neglected, where both features can be easily appreciated.

When the stock exactly reaches the ROP, the system is only out of stock if $D_L > s$, and therefore, the stockouts only take place during the lead time. Hence, the expected shortage per replenishment cycle is straightforwardly computed as $ESPRC = \sum_{i=s+1}^{\infty} (i - s) \cdot f_L(i)$.

To compute the expected total demand per replenishment cycle (EDPRC), we sum the accumulated demand from the beginning of the cycle until the ROP

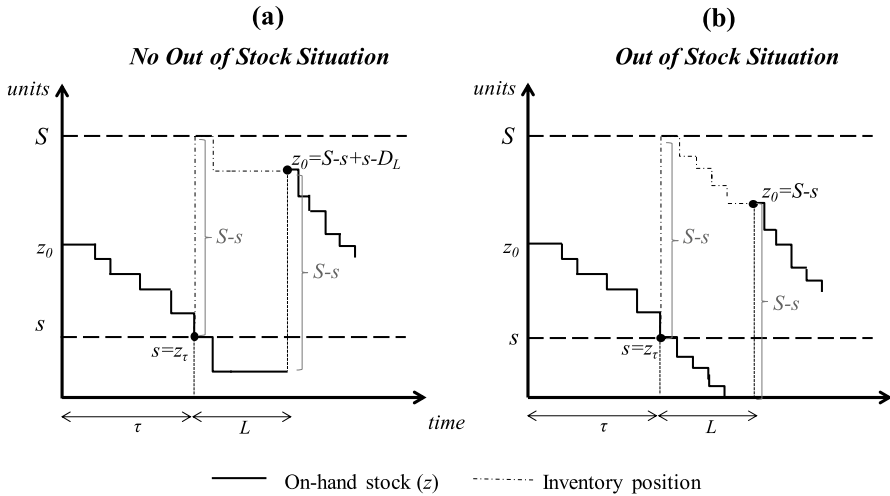


Fig. 2 Evolution of the stock in a (s, S) inventory policy when the undershoots are neglected

is reached, D , and the accumulated demand over the lead time, D_L . The latter is simply computed when the demand distribution is known. However, to compute D_τ , we need to know the on-hand stock balance at the beginning of the cycle, i.e., at order delivery, which, in a lost sales context, is obtained by $z_0 = S - s + E[s - D_L]^+$. Therefore, taking into account that D_τ must guarantee that the ROP is exactly reached, then $D_\tau = z_0 - s = S - 2s + E[s - D_L]^+$.

Consequently, the expression to estimate the fill rate when the undershoots are neglected (named initial fill rate, β_0 , further on) for the lost sales case that applies to any discrete distribution is:

$$\beta_0 = 1 - \frac{\sum_{i=s+1}^{\infty} (i - s) \cdot f_L(i)}{S - 2s + \sum_{j=0}^s (s - j) \cdot f_L(j) + \sum_{k=0}^{\infty} k \cdot f_L(k)} \tag{1}$$

It is clear that conceptually neglecting undershoots leads to a systematic overestimation of the fill rate and, therefore, introduces a bias in its estimation, which can be clearly appreciated in the following example. If the stock at t is one unit above the ROP, and demand at $t + 1$ is 2 units, the ROP is reached, and a replenishment order is placed. The real stock at launch is $s - 1$ units. Therefore, the on-hand stock remaining on the shelves available to meet the demand during the lead time is one unit less than what is assumed if the undershoot is neglected. Thus, the initial fill rate is higher than the real rate, and thus, the stockouts are larger than expected. This simple example shows why neglecting undershoots introduces a significant bias in the estimation of the fill rate.

3.2 Empirical analysis of the bias

To analyse the bias of the initial fill rate, we designed an experiment based on simulations as follows (see Fig. 3): (1) we combine a set of input parameters that define the inventory policy and the simulated demand; (2) for each combination of the input data, we simulate a (s, S) inventory system using the Monte Carlo method and calculate the achieved fill rate; (3) for each input data combination, we also compute the initial fill rate using expression (1). Note that inventory simulation is an adequate approach given the complexities of the system (Chinello et al. 2020).

With regard to the input data, an extensive range of values for s , S and L is selected to provide realistic values of the fill rate (from 0.5 to 0.99). Regarding the demand, we select a set of parameters (r, θ) for the negative binomial distribution based on the consideration of four demand categories suggested by Syntetos et al. (2005): smooth, lumpy, intermittent and erratic, where r is the number of successes and θ is the probability of success. Moreover, the negative binomial distribution reproduces some characteristics typically seen in demand data; for instance, it is positive and asymmetric, which models potential marketing campaigns. Tables 1 and 2 present the set of data used in the experiment, which considers every feasible combination of these values per factor, that leads to 1360 different cases (excluding cases where $s \geq S - s$ to ensure that, as we assume, there is never more than one




Fig. 3 Experimental design

Table 1 Inventory input data

Reorder point	$s = 2, 3, 5, 6$
Order-up-to-level	$S = 5, 7, 12, 15$
Lead time	$L = 1, 2, 3, 4$

Table 2 Demand input data

Negative binomial (r, θ)	
Smooth	$(3, 0.4); (3, 0.5); (3, 0.6); (3.5, 0.4); (3.5, 0.5); (3.5, 0.6); (4, 0.4); (4, 0.5); (4, 0.6)$
Erratic	$(1, 0.1); (1.5, 0.1); (1.5, 0.3); (2, 0.1); (2, 0.3)$
Lumpy	$(0.1, 0.1); (0.1, 0.3); (0.25, 0.1); (0.25, 0.3); (0.5, 0.1); (0.5, 0.3); (1, 0.1); (1, 0.3)$
Intermittent	$(0.1, 0.5); (0.1, 0.8); (0.1, 0.9); (0.25, 0.5); (0.5, 0.5); (0.5, 0.8); (0.5, 0.9); (0.75, 0.5); (0.75, 0.8); (0.75, 0.9); (1, 0.8); (1, 0.9)$

outstanding order). The inventory system simulation follows the diagram (Fig. 4), and for each input data combination, the fill rate achieved (β_A) is obtained as the average fraction of the complement of the unfulfilled demand over the total demand in every replenishment cycle when considering 20,000 consecutive periods, i.e.:

$$\beta_A = \frac{1}{N} \sum_{n=1}^N \left(1 - \frac{\text{unfilled demand}_n}{\text{total demand}_n} \right) \tag{2}$$

where N indicates the total number of replenishment cycles. To assure the consistency of the results, Monte Carlo simulations undertake thirty replications of each case using the average of these replications as the final achieved fill rate.

Figure 5 compares the achieved (β_A) and the initial (β_0) fill rates plotted as red circles and illustrates the bias introduced in the calculation of the initial fill rate due to neglecting the undershoots. The larger the distance between the red circles and the blue line, the larger the bias. Note that the initial fill rate always overestimates the achieved fill rate, as explained in Sect. 3.1. What is empirically found is that the bias size is not constant, and it seems to increase over the central part of the figure.

Therefore, the problem is to find the bias between β_0 and β_A , which appears to have a nonlinear pattern.

4 Analytics fill rate methods

To model the bias shown in the previous section, we assume that the data are cross-sectional, i.e., the relationship between the achieved fill rate and the initial fill rate does not depend on time. The initial number of experiments was 1360. We focused on fill rates greater than 50%, filtering the number of experiments to 1011. The training and test sets are divided by random sampling, where 70% of the data (708 experiments) are used for estimation purposes and the rest of the data (303 experiments) are used for validation.

Figure 5 shows that the bias between the fill rates is nonlinear and is greater for fill rates between approximately 0.55 and 0.9, where the maximum bias is reached at approximately 0.6 and 0.75. To model that bias, this work adopts an SDP approach.

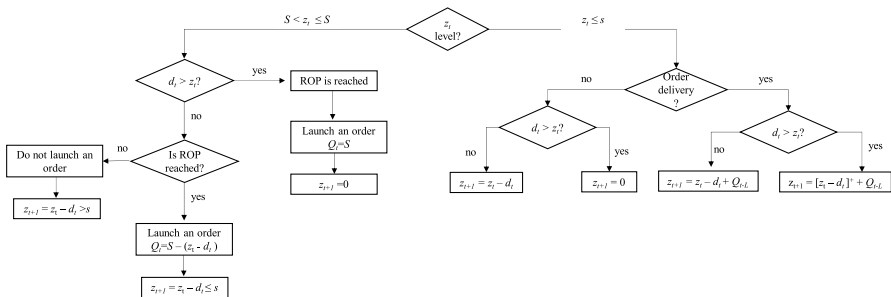


Fig. 4 Diagram of the simulation




Fig. 5 Achieved fill rate versus initial fill rate (red circles)

The approach requires that we define the state and the parameter that results in a variance of the state. An initial model could be as follows:

$$\beta_{SDP} = \alpha_1(\beta_0) \cdot \beta_0 \tag{3}$$

where α_1 is the state that somehow depends on the parameter β_0 . Note that we do not exactly know the relationship between the state (α_1) and the parameter (β_0), and it will be the SDP algorithm, using the fill rate simulated data, that will provide a curve to give us some insights about the pattern of such a dependency. In this sense, β_{SDP} approximates β_A .

4.1 A nonparametric approach: the SDP algorithm

The SDP approach can be seen as an extension of the time variable parameter estimation (TVP) (Young 2012), where the unknown parameters are slowly variable with time. Typically, TVP estimations are represented by a two-dimensional state (level and slope) to model its stochastic behaviour by means of a generalized random walk (GRW), which is a generic model to unify the different versions of the random walk as an integrated random walk (IRW), random walk (RW) and smoothed random walk (SRW). More information about the use of TVPs in time series analysis and forecasting can be found in Young et al. (2001). The SDP approach employs a similar procedure but the main difference is that the parameters evolve stochastically with respect to another variable instead of time. Such parameters are denoted as state-dependent parameters (SDPs). This can be achieved by “sorting” the data in a nontemporal order. If the new ordering provides a smoother and less rapid variation in the SDP, a GRW will be able

to describe its evolution in the transformed observation space. Utilizing a fixed interval smoothing (FIS) estimation, the estimated SDP can be “unsorted” to the original order.

In our particular case, the SDP α_1 can be sorted with respect to β_0 . This new organization of data is indexed by k , and thus, the stochastic dynamics of the SDP can be expressed as an integrated random walk defined in a state space framework as follows:

$$\begin{pmatrix} \alpha_1(k+1) \\ \alpha_1^*(k+1) \end{pmatrix} = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} \alpha_1(k) \\ \alpha_1^*(k) \end{pmatrix} + \begin{pmatrix} 0 \\ \eta^*(k) \end{pmatrix} \tag{4}$$

where α_1 is the level and α_1^* is the slope. Small variations with respect to the new sorting are introduced with the random Gaussian noise $\eta^*(k)$ with zero mean and the constant variance σ_α^2 . Such a variance is often referred to as a *hyperparameter* to differentiate it from the states that are the main object of the estimation analysis. The observation equation needed to complete the state space representation is:

$$\beta_{SDP} = \alpha_1(\beta_0) \cdot \beta_0 + \epsilon \tag{5}$$

The last term ϵ is the typical error component that assumes a Gaussian distribution with zero mean and constant variance σ^2 .

Fortunately, these routines are already implemented in the CAPTAIN toolbox of MATLAB available in <http://captaintoolbox.co.uk>. See “Appendix” for more information about the recursive filtering/smoothing algorithms used. Figure 6 shows the estimated SDP $\alpha_1(\beta_0)$ sorted with respect to the parameter β_0 for the training set. This figure shows the nonlinear pattern of the SDP coherently with the nonlinear bias found between β_A and β_0 in Fig. 6.




Fig. 6 SDP $\alpha_1(\beta_0)$ sorted with respect to the parameter β_0

Figure 7 depicts the SDP estimation of the fill rate (β_{SDP}) together with the initial and achieved fill rates for the training set. That figure clearly shows how the SDP approach is capable of considerably reducing the bias. Essentially, the fill rate estimates provided by β_{SDP} estimate the mean of β_A , correcting the systematic deviation found for its initial counterpart β_0 .

4.2 Parameterization based on the SDP graph

The SDP approach provides a nonparametric estimation of the relationship of $\alpha_1(\beta_0)$ based on a graph, as shown in Fig. 6. Furthermore, this approach can be complemented by parameterizing such a curve and proposing another parametric alternative. Figure 8 shows two polynomial (quadratic and cubic) functions that could be used to fit the resulting nonlinear SDP $\alpha_1(\beta_0)$ curve.

Considering that, the following model can be proposed:

$$\beta_{LR} = \theta_0 + (\theta_1 + \theta_2 \cdot \beta_0 + \theta_3 \cdot \beta_0^2 + \theta_4 \cdot \beta_0^3) \cdot \beta_0 + \epsilon \tag{6}$$

where θ_i for $i=0, 1, 2, 3,4$ are constants that can be estimated by means of a linear regression and ϵ is the typical error component that assumes a Gaussian distribution with a zero mean and a constant variance. Again, β_{LR} approximates β_A .

It is important to note that, despite the nonlinearities involved in the initial fill rate bias, thanks to the SDP identification, we have arrived at a nonlinear model that is linear with respect to the parameters θ_i . Note that such parameterization makes the model fully self-contained and clearly reveals the nature of the graphically identified nonlinearity with just a few parameters.




Fig. 7 Fill rate estimations provided by the initial (β_0) and the SDP (β_{SDP}) approaches for the training set




Fig. 8 SDP $\alpha_1(\beta_0)$ modelled with quadratic (dotted line) and cubic (dashed line) polynomials

Table 3 summarizes the estimation results for the quadratic and cubic polynomials. The table shows that all parameters are statistically significant for a significance level of 1%. Although both approaches perform well, the cubic polynomial achieves a smaller RMSE in the estimation and a slightly better goodness of fit (R^2).

To validate that the parameters can be used to correct the bias with data that have not been used in the training set, we used those estimated parameters in the data test set. Figure 9 shows the fill rate estimates provided by the nonparametric SDP (β_{SDP}) and the parametric cubic polynomial model (β_{LR}) compared to the achieved and initial fill rates. That figure shows that the performance provided by β_{SDP} and β_{LR} is very similar, corroborating the utility of both models.

Table 3 OLS estimation results for the parameters in (6)

Parameter	Value (std dev.)	
	Cubic	Quadratic
θ_0	7.6 (1.9)*	- 1.93 (0.3)*
θ_1	- 40.9 (10.1)*	9.9 (1.2)*
θ_2	86.4 (19.8)*	- 13.4 (1.5)*
θ_3	- 79.4 (17)*	6.4 (0.6)*
θ_4	27.3 (5.4)*	-
R^2	0.967	0.965
RMSE	0.0299	0.0304

*Significant at the 1% confidence level




Fig. 9 Fill rate estimations provided by the initial formula (β_0), SDP approach (β_{SDP}) and parametric model (β_{LR}) for the test set

4.3 Error assessment

We can define the fill rate error as:

$$e_{i,j} = \beta_{i,j} - \beta_{A,j}$$

where β_{ij} for $i=1, 2, 3$ is the fill rate estimated by the initial equation, the nonparametric SDP and the parametric linear regression, respectively. The index $j=1, 2, \dots, J$ is the number of fill rate observations for the entire test set J . The performance of the different fill rate estimates can be summarized by the mean error (ME) and the root mean squared error (RMSE), respectively, such that:

$$ME_i = \frac{\sum_{j=1}^J e_{i,j}}{J} \tag{7}$$

$$RMSE_i = \sqrt{\frac{\sum_{j=1}^J e_{i,j}^2}{J}} \tag{8}$$

Note that the ME measures the error bias, whereas the RMSE focuses on the error size.

Figure 10 shows the ME and RMSE obtained for the different fill rate estimation techniques. First, the high positive value of ME for β_0 quantifies the overestimation produced by the expression (1). Note that such a formula also displays a high error variability as measured by the RMSE. This figure also clearly shows how the parametric approach (β_{LR}) and the SDP (β_{SDP}) alternatives outperform the initial (β_0) in




Fig. 10 ME (left panel) and RMSE (right panel) obtained for the different fill rate estimates with respect to the achieved fill rate for the test set

terms of the bias (ME) and error variability size (RMSE). On the other hand, the differences found between β_{SDP} and β_{LR} are minimal.

5 Case study

To verify the performance of the proposed methods for a real dataset, sales data previously used by Barrow and Kourentzes (2016) were utilized as the demand estimates. This dataset belongs to a major UK fast-moving consumer goods manufacturer specializing in the production of household and personal care products. In total, 229 products at the SKU level were available, with 173 weekly observations per product. According to Barrow and Kourentzes (2016), there are no seasonal SKUs, and only a minority (21%) exhibit a slight trend. To calculate the achieved fill rate for real data, we introduced the demand data of every SKU in our simulation, indicating the values of S , s and L for each product, and following the procedure summarized in Fig. 4, we calculated the fill rate with expression (2).

Proceeding as in the previous simulations, 70% of the data are held as a training set, and 30% are held as a test set. Figure 11 shows the initial (β_0) and achieved (β_A) fill rates for the training set of the case study. As previously occurred with the simulated demands, there is a nonlinear positive bias of β_0 with respect to β_A . That bias is higher around fill rates of 0.8 and 0.9.

Figure 12 depicts the results obtained in the test set, where β_{SDP} and β_{LR} have been reestimated using the training dataset. Note that β_{LR} was computed by means of a quadratic polynomial function. In general, the bias of β_{SDP} and β_{LR} was significantly




Fig. 11 Relationship between the initial fill rate (β_0) and achieved fill rate (β_A) for the training set of the case study data




Fig. 12 Fill rate estimations provided by the initial formula (β_0), SDP approach (β_{SDP}) and parametric model (β_{LR}) for the test set of the case study dataset




Fig. 13 ME (left panel) and RMSE (right panel) obtained for the different fill rate estimates with respect to the achieved fill rate for the test set of the case study dataset

reduced. The bias reduction is quantified in Fig. 13. Again, the ME and RMSE of β_0 with regard to the proposed alternatives are larger.

6 Discussion and practical implications

The previous sections showed that analytics fill rate methods successfully remove the bias associated with the initial fill rate formula. In this section, we will illustrate the managerial implications and benefits of using them in practice. Essentially, we will describe two possible applications: (1) a system performance measurement and (2) a policy design. In both examples, we will use the parametric alternative (β_{LR}), since it provides similar results when compared with the nonparametric counterpart (β_{SDP}).

6.1 System performance measurement

For companies, one possible application is to measure the performance of a determined (s , S) pair of parameters in terms of the fill rate. Consider a manager that sets a fixed value of S for physical restrictions, and s is defined judgementally. The manager would like to know beforehand what the resulting fill rate for those chosen parameters will be, given past sales as demand estimates and a given lead time. By using the initial formula, they would obtain β_0 . However, as previously demonstrated, such a formula overestimates the real fill rate, and the system will be less

protected than the manager expects with an unexpected increase in stockout costs. Thus, it should be corrected by means of β_{LR} .

An example of that application is shown in Fig. 14. Given the sales data from SKU 126, $S=1000$ and $L=2$, the figure shows that for different values of s , the β_0 overestimates the fill rate achieved β_A . In contrast, β_{LR} follows β_A in a closer fashion.

6.2 Parameter design

Another application of the proposed methodology is to design an inventory policy with the fill rate as a constraint. In other words, given a target fill rate, which values of s and S should be chosen? In this case, we have two options. First, if we do not apply the bias correction, then $\beta_0 = \beta_{target}$, and by means of Eq. (1), s can be obtained, denoted by s_0 . Alternatively, to avoid the bias associated with β_0 , we set $\beta_{LR} = \beta_{target}$, and we can obtain β'_0 from expression (9). Note that, unlike Eq. (6), in the expression (9), β_{LR} acts as an input, and β'_0 is the output. Proceeding in the manner shown in Sect. 4, θ'_i for $i=0, 1, 2, 3,4$ should be reestimated with the training dataset. Note that the parameters in (6) and (9) are not the same.

$$\beta'_0 = \theta'_0 + (\theta'_1 + \theta'_2 \cdot \beta_{LR} + \theta'_3 \cdot \beta_{LR}^2 + \theta'_4 \cdot \beta_{LR}^3) \cdot \beta_{LR} + \epsilon \tag{9}$$

With the new value of β'_0 together with Eq. (1), we can obtain the value of s , which is denoted by s_{LR} . Next, these values of s_0 and s_{LR} are fed into the simulation to obtain $\beta_A(s_0, S)$ and $\beta_A(s_{LR}, S)$, whose values are subject to a comparison with respect to the initially defined β_{target} .




Fig. 14 Fill rate calculations ($\beta_0, \beta_{LR}, \beta_A$) for different values of s for SKU 126, $L=2, S=1000$

Table 4 ROP computed without considering bias (s_{LR}) and considering bias (s_0) for several SKUs of the real dataset

	L	β_{target}	Policy			Fill rate achieved	
			S	s_0	s_{LR}	$\beta_A(s_0, S)$	$\beta_A(s_{LR}, S)$
SKU126	2	0.75	1000	104	300	0.72	0.76
SKU126	2	0.85	1100	283	481	0.72	0.85
SKU100	1	0.90	5000	244	1155	0.84	0.88
SKU100	1	0.95	4600	860	2227	0.87	0.95
SKU51	2	0.70	1675	498	774	0.66	0.71
SKU51	2	0.65	1500	417	695	0.62	0.65

Table 4 shows various examples for three different SKUs of the real dataset. For instance, the first row of the table is an example for SKU 126, given $L=2$, $S=1000$ and $\beta_{target} = 0.75$. The calculation of the reorder points is $s_0=104$ and $s_{LR}=300$, depending on whether the bias is corrected. These ROPs are fed as inputs into the simulation, and the fill rates achieved are 0.72 and 0.76, respectively. Since the fill rate achieved by $\beta_A(s_{LR}, S)$ is closer to β_{target} , it corroborates the usefulness of the proposed method to more effectively design the inventory policy, which means a reduction in stockout costs. Similar conclusions can be drawn regarding the rest of the SKUs in the table.

7 Summary and conclusions

This paper focuses on the versatile order-point, order-up-to-level (s, S) system in the lost sales context. The control procedure of this system consists of continuously reviewing the status of the inventory in order to know exactly when the inventory position is at or below the ROP in order to launch a replenishment order to raise the inventory position to the order-up-to-level, S . Despite this system being widely used in practice, it is normally implemented assuming that the undershoots at ROP are negligible, i.e., assuming the inventory position always exactly reaches the ROP, which is only possible if demand transactions are unit sized. However, if demand is not unitary, the assumption regarding neglecting the undershoots leads to a bias in the performance of the system. This paper shows both conceptually and experimentally that the initial formula to compute the fill rate is biased. It can provide an overestimate of almost 7% in the estimation of the fill rate and lead to wrong decisions that directly affect the performance and costs of the inventory system.

This paper proposes a new methodology from a data-driven perspective that uses a nonparametric SDP algorithm to model the bias and subsequently suggests two analytics fill rate methods that outperform the initial approach. Note that, despite the unknown complex nonlinearities involved in the initial fill rate bias, the SDP has revealed the nature of such a nonlinearity, which has been subsequently parameterized by a parametric nonlinear model, linear in the parameters, that allows the use of linear regression. The importance of the analytics fill rate methods is that both are unbiased and easily implemented in practice.

Although the SDP approach can be seen as a data-driven tool, it is important to distinguish it from other black-box data-driven tools typically associated with machine learning alternatives. In this work, the SDP approach was employed to learn the nature of the nonlinearities involved in the initial fill rate bias by neglecting the undershoots in the (s, S) system. Once those nonlinearities were identified by means of a curve, the expert model could develop a parametric solution taking advantage of the solid theory that underpins parametric theory. In this case, a simple linear regression was enough to model a complex nonlinearity. Note that this philosophy follows the Data Based Mechanistic approach proposed by Young (2012), where the author explains that “a model should not just explain the time series data well, but it should also provide a mechanistic description of the system under investigation”.

Practitioners interested in applying the methodology proposed should proceed as follows. First, they should compute the potential bias between their theoretical fill rate expression used and their achieved fill rate; if that bias is similar to the one presented here in the simulations and case study, then they can directly utilize the parametric model (β_{LR}) . If the bias found does not follow a similar nonlinear pattern, then the previous SDP identification step should be carried out. Note that the improvements obtained per SKU might mean a significant reduction in the stockout costs when applied to a large portfolio of SKUs.

Further research should analyse the effect on replacing the statistical demand distributions in the fill rate initial formula with the probabilistic forecasts of real demands. In addition, further research is also needed to extend these promising results to other systems.

Appendix

This appendix shows the recursive filtering/smoothing algorithms that are usually employed to reflect the TVP evolution (Young 2011) and how they have been extended and modified for SDP estimation. More details about this methodology can be found in Young (2011) and Young et al. (2001).

An overall SS model can be defined by these two equations:

$$\begin{aligned} \text{Observation equation : } & y_t = \mathbf{H}_t \mathbf{x}_t + e_t \\ \text{State equation : } & \mathbf{x}_t = \mathbf{F} \mathbf{x}_{t-1} + \mathbf{G} \boldsymbol{\eta}_t \end{aligned} \tag{10}$$

where \mathbf{x}_t is known as the state vector. In our case, the observation equation is defined in (5) and the state equation in (4). Then, $\mathbf{x}_t = [\alpha_{1,t}, \alpha_{1,t}^*]^T$. The remaining matrices and vectors are defined as follows:

$$\mathbf{F} = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}, \mathbf{G} = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}, \mathbf{H}_t = (\beta_c \quad 0), \boldsymbol{\eta}_t = (0 \quad \eta_t^*)^T \tag{11}$$

The white noise inputs, $\boldsymbol{\eta}_t$, are assumed to be independent of the observation noise, σ^2 , and have a covariance matrix \mathbf{Q} . These noise inputs determine the stochastic behaviour of the states.

The SDP estimation can be divided into two steps. First, forward and backwards recursive algorithms are employed. Second, a backfitting algorithm is applied with the sorted data with respect to the dependent variable. Note that both the backwards pass smoothing and backfitting algorithms can be applied only if it is not necessary to work in real time.

Forward pass recursive LS equations

Prediction:

$$\begin{aligned} \hat{x}_{t|t-1} &= F\hat{x}_{t-1} \\ P_{t|t-1} &= FP_{t-1}F^T + GQ_rG^T \end{aligned} \tag{12}$$

Correction:

$$\begin{aligned} \hat{x}_t &= \hat{x}_{t|t-1} + P_{t|t-1}H_t^T [1 + H_tP_{t|t-1}H_t^T]^{-1} (y_t - H_t\hat{x}_{t|t-1}) \\ P_t &= P_{t|t-1} - P_{t|t-1}H_t^T [1 + H_tP_{t|t-1}H_t^T]^{-1} H_tP_{t|t-1} \end{aligned} \tag{13}$$

Backwards pass smoothing equations

$$\begin{aligned} \hat{x}_{t|N} &= F^{-1}[\hat{x}_{t+1|N} + GQ_rG^TL_t] \\ L_t &= [I - P_{t+1}H_{t+1}^T H_{t+1}]^T [F^TL_{t+1} - H_{t+1}^T (y_{t+1} - H_{t+1}\hat{x}_{t+1})] \\ P_{t|N} &= P_t + P_tF^T P_{t+1|N}^{-1} [P_{t+1|N} - P_{t+1|t}] P_{t+1|t}^{-1} FP_t \\ L_t &= 0 \end{aligned} \tag{14}$$

Note that Q_r and P_t are normalized with respect to the observation equation noise (σ^2), such as:

$$Q_r = \frac{Q}{\sigma^2}; P_t = \frac{P_t^*}{\sigma^2} \tag{15}$$

P_t^* is the error covariance matrix related to the state estimates. The parameters inside that NVR matrix are estimated by a maximum likelihood prior to applying the recursive algorithms (Young et al. 2001).

Backfitting algorithm for SDP models.

1. Assume that FIS estimation has yielded an initial TVP estimate of $\hat{\alpha}_{1,t|N}^0$
2. Iterate $k = 1, 2, \dots, k_c$
 - a. sort both β_A and β_c according to the ascending order of β_A .
 - b. obtain an FIS estimate $\hat{\alpha}_{1,t|N}^k$ in the relationship $\beta_a = \alpha_1(\beta_a) \cdot \beta_c$.

3. Continue Step 2 until iteration k_c , that is, when the SDP (which is a series of length N) remains approximately constant.

Funding Open Access funding provided thanks to the CRUE-CSIC agreement with Springer Nature. This work was supported by the European Regional Development Fund and Junta de Comunidades de Castilla-La Mancha (JCCM/FEDER, UE) under the project with reference SBPLY/19/180501/000151 and by the Vicerrectorado de Investigación y Política Científica from UCLM through the research group fund program (PREDILAB; [2021-GRIN-31210]). Funding for open access charge: CRUE-Universitat Politècnica de València.

Declarations

Conflict of interest Authors confirm that no conflict nor financial interest exist and that the Financial Disclosure is correct. We also confirm that the author list and affiliations are correct.

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

References

- Agrawal V, Seshadri S (2000) Distribution free bounds for service constrained (Q, r) inventory systems. *Nav Res Logist NRL* 47:635–656. [https://doi.org/10.1002/1520-6750\(200012\)47:8%3c635::AID-NAV3%3e3.0.CO;2-C](https://doi.org/10.1002/1520-6750(200012)47:8%3c635::AID-NAV3%3e3.0.CO;2-C)
- Axsater S (2000) *Inventory control*. Kuwer Academic Publishers, New York
- Babiloni E, Guijarro E (2020) Fill rate: from its definition to its calculation for the continuous (s, Q) inventory system with discrete demands and lost sales. *Cent Eur J Oper Res* 28:35–43. <https://doi.org/10.1007/s10100-018-0546-7>
- Baganha MP, Pyke DF, Ferrer G (1996) The undershoot of the reorder point: tests of an approximation. *Int J Prod Econ* 45:311–320. [https://doi.org/10.1016/0925-5273\(96\)00004-7](https://doi.org/10.1016/0925-5273(96)00004-7)
- Baganha MP, Ferrer G, Pyke DF (1999) The residual life of the renewal process: a simple algorithm. *Nav Res Logist NRL* 46:435–443. [https://doi.org/10.1002/\(SICI\)1520-6750\(199906\)46:4%3c435::AID-NAV6%3e3.0.CO;2-I](https://doi.org/10.1002/(SICI)1520-6750(199906)46:4%3c435::AID-NAV6%3e3.0.CO;2-I)
- Barrow DK, Kourentzes N (2016) Distributions of forecasting errors of forecast combinations: implications for inventory management. *Int J Prod Econ* 177:24–33. <https://doi.org/10.1016/j.ijpe.2016.03.017>
- Bijvank M, Vis IFA (2011) Lost-sales inventory theory: a review. *Eur J Oper Res* 215:1–13
- Bijvank M, Vis IFA (2012) Lost-sales inventory systems with a service level criterion. *Eur J Oper Res* 220:610–618. <https://doi.org/10.1016/j.ejor.2012.02.013>
- Bijvank M, Bhulai S, Huh WT (2015) Parametric replenishment policies for inventory systems with lost sales and fixed order cost. *Eur J Oper Res* 241:381–390
- Breugelmans E, Campo K, Gijsbrechts E (2006) Opportunities for active stock-out management in online stores: the impact of the stock-out policy on online stock-out reactions. *J Retail* 82:215–228. <https://doi.org/10.1016/j.jretai.2006.05.001>
- Brown RG (1962) *Smoothing forecasting and prediction of discrete time series*. Prentice-Hall, Hoboken

- Caplin A, Leahy J (2010) Economic theory and the world of practice: a celebration of the (S, s) model. *J Econ Perspect* 24:183–202. <https://doi.org/10.1257/jep.24.1.183>
- Cardós M, Babiloni E (2011a) Exact and approximate calculation of the cycle service level in periodic review inventory policies. *Int J Prod Econ* 131:63–68. <https://doi.org/10.1016/j.ijpe.2010.05.012>
- Cardós M, Babiloni E (2011b) Exact and approximated calculation of the cycle service level in a continuous review policy. *Int J Prod Econ* 133:251–255. <https://doi.org/10.1016/j.ijpe.2010.09.017>
- Cardós M, Guijarro E, Babiloni E (2017) On the estimation of on-hand stocks for base-stock policies and lost sales systems and its impact on service measures. *Int J Prod Res* 55(16):4680–4694. <https://doi.org/10.1080/00207543.2017.1279759>
- Chinello E, Lee Herbert-Hansen ZN, Khalid W (2020) Assessment of the impact of inventory optimization drivers in a multi-echelon supply chain: case of a toy manufacturer. *Comput Ind Eng* 141:106232. <https://doi.org/10.1016/j.cie.2019.106232>
- Cohen MA, Kleindorfer PR, Lee HL (1988) Service constrained (s, S) inventory systems with priority demand classes and lost sales. *Manag Sci* 34:482–499. <https://doi.org/10.1287/mnsc.34.4.482>
- Diels JL, Wiebach N (2011) Customer reactions in out-of-stock situations: Do promotion-induced phantom positions alleviate the similarity substitution hypothesis? SFB 649 discussion paper
- Disney SM, Ponte B, Wang X (2021) Exploring the nonlinear dynamics of the lost-sales order-up-to policy. *Int J Prod Res* 59:5809–5830. <https://doi.org/10.1080/00207543.2020.1790687>
- Escalona P, Angulo A, Weston J, Stegmaier R, Kauak I (2019) On the effect of two popular service-level measures on the design of a critical level policy for fast-moving items. *Comput Oper Res* 107:107–126. <https://doi.org/10.1016/j.cor.2019.03.011>
- Gandomi A, Haider M (2015) Beyond the hype: big data concepts, methods, and analytics. *Int J Inf Manag* 35:137–144. <https://doi.org/10.1016/j.ijinfomgt.2014.10.007>
- Gruen TW, Corsten D, Bharadwaj S (2002) Retail out-of-stocks: a worldwide examination of extent causes, rates and consumer responses. Grocery Manufacturers of America, Washington
- Guijarro E, Cardós M, Babiloni E (2012) On the exact calculation of the fill rate in a periodic review inventory policy under discrete demand patterns. *Eur J Oper Res* 218:442–447. <https://doi.org/10.1016/j.ejor.2011.11.025>
- Gutgutia A, Jha JK (2018) A closed-form solution for the distribution free continuous review integrated inventory model. *Oper Res* 18:159–186. <https://doi.org/10.1007/s12351-016-0258-5>
- Gutierrez M, Rivera FA (2021) Undershoot and order quantity probability distributions in periodic review, reorder point, order-up-to-level inventory systems with continuous demand. *Appl Math Model* 91:791–814
- Hadley G, Whitin T (1963) Analysis of inventory systems. Prentice-Hall, Englewood Cliffs
- Huber J, Müller S, Fleischmann M, Stuckenschmidt H (2019) A data-driven newsvendor problem: from data to decision. *Eur J Oper Res* 278:904–915. <https://doi.org/10.1016/j.ejor.2019.04.043>
- Johansen SG (2001) Pure and modified base-stock policies for the lost sales inventory system with negligible set-up costs and constant lead times. *Int J Prod Econ* 71:391–399
- Johansen SG, Hill RM (2000) The (r, Q) control of a periodic-review inventory system with continuous demand and lost sales. *Int J Prod Econ* 68:279–286. [https://doi.org/10.1016/S0925-5273\(00\)00051-7](https://doi.org/10.1016/S0925-5273(00)00051-7)
- Kaur K (2021) Business intelligence on supply chain responsiveness and agile performance: empirical evidence from Malaysian logistics industry. *Int J Supply Chain Manag* 6:31–63. <https://doi.org/10.47604/ijscm.1351>
- Kouki C, Sahin E, Jemai Z, Dallery Y (2009) Consideration of the undershoot in an (r, Q) inventory model for perishables. *IFAC Proc Vol* 42:1573–1578. <https://doi.org/10.3182/20090603-3-RU-2001.0154>
- Larsen C, Thorstenson A (2014) The order and volume fill rates in inventory control systems. *Int J Prod Econ* 147:13–19. <https://doi.org/10.1016/j.ijpe.2012.07.021>
- Moors JJA, Strijbosch LWG (2002) Exact fill rates for (R, s, S) inventory control with gamma distributed demand. *J Oper Res Soc* 53:1268–1274. <https://doi.org/10.1057/palgrave.jors.2601441>
- Platt DE, Robinson LW, Freund RB (1997) Tractable (Q, R) heuristic models for constrained service levels. *Manag Sci* 43:951–965. <https://doi.org/10.1287/mnsc.43.7.951>
- Schneider H (1981) Effect of service-levels on order-points or order-levels in inventory models. *Int J Prod Res* 19:615–631. <https://doi.org/10.1080/00207548108956694>
- Silver EA (1970) A modified formula for calculating customer service under continuous inventory review. *AIIE Trans* 2:241–245

- Silver EA, Naseraldin H, Bischak DP (2009) Determining the reorder point and order-up-to-level in a periodic review system so as to achieve a desired fill rate and a desired average time between replenishments. *J Opr Res Soc* 60:1244–1253
- Silver EA, Bischak DP, de Kok T (2012) Determining the reorder point and order-up-to level to satisfy two constraints in a periodic review system under negative binomial demand. *J Oper Res Soc* 63:941–949. <https://doi.org/10.1057/jors.2011.102>
- Silver EA, Pyke DF, Peterson R (2017) Inventory and production management in supply chains, 4th edn. Taylor and Francis Group, Boca Raton
- Spiliotis E, Makridakis S, Semenoglou A-A, Assimakopoulos V (2020) Comparison of statistical and machine learning methods for daily SKU demand forecasting. *Oper Res Int J*. <https://doi.org/10.1007/s12351-020-00605-2>
- Stefanovic N (2015) Collaborative predictive business intelligence model for spare parts inventory replenishment. *Comput Sci Inf Syst* 12:911–930. <https://doi.org/10.2298/CSIS141101034S>
- Strijbosch LWG, Heuts RMJ, van der Schoot EHM (2000) A combined forecast—inventory control procedure for spare parts. *J Oper Res Soc* 51:1184–1192. <https://doi.org/10.1057/palgrave.jors.2601013>
- Syntetos AA, Boylan JE, Croston JD (2005) On the categorization of demand patterns. *J Oper Res Soc* 56:495–503
- Tijms HC, Groenevelt H (1984) Simple approximations for the reorder point in periodic and continuous review (S, S) inventory systems with service level constraints. *Eur J Oper Res* 17:175–190
- Trapero JR, Fildes R, Davydenko A (2011) Nonlinear identification of judgmental forecasts effects at SKU level. *J Forecast* 30:490–508. <https://doi.org/10.1002/for.1184>
- Trapero JR, Cardós M, Kourentzes N (2019a) Quantile forecast optimal combination to enhance safety stock estimation. *Int J Forecast* 35:239–250. <https://doi.org/10.1016/j.ijforecast.2018.05.009>
- Trapero JR, Cardós M, Kourentzes N (2019b) Empirical safety stock estimation based on kernel and GARCH models. *Omega* 84:199–211. <https://doi.org/10.1016/j.omega.2018.05.004>
- Vincent P (1983) Practical methods for accurate fill rates. *INFOR Inf Syst Oper Res* 21:109–120. <https://doi.org/10.1080/03155986.1983.11731890>
- Wang G, Gunasekaran A, Ngai EWT, Papadopoulos T (2016) Big data analytics in logistics and supply chain management: certain investigations for research and applications. *Int J Prod Econ* 176:98–110. <https://doi.org/10.1016/j.ijpe.2016.03.014>
- Young PC (2012) Recursive estimation and time-series analysis: an introduction. Springer, Berlin
- Young PC, McKenna P, Bruun J (2001) Identification of non-linear stochastic systems by state dependent parameter estimation. *Int J Control* 74:1837–1857. <https://doi.org/10.1080/00207170110089824>
- Young PC (2011) Recursive estimation and time-series analysis: an introduction for the student and practitioner. Springer Science & Business Media
- Zhou X, Deng X, Deng Y, Mahadevan S (2017) Dependence assessment in human reliability analysis based on D numbers and AHP. *Nucl Eng Des* 313:243–252. <https://doi.org/10.1016/j.nucengdes.2016.12.001>

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