A Practical Method for Joint Inventory Replenishment under Supplier Constraints

T.S. Winter[1*] & J.H. van Vuuren[1]

ARTICLE INFO

Article details
Submitted by authors 12 Jan 2025
Accepted for publication 10 Jul 2025
Available online 29 Aug 2025

Contact details
* Corresponding author
thomas14601@gmail.com

Author affiliations
1 Stellenbosch Unit for Operations Research in Engineering, Department of Industrial Engineering, Stellenbosch University, Stellenbosch, South Africa

ORCID® identifiers
T.S. Winter
https://orcid.org/0009-0005-4157-3141

J.H. van Vuuren
https://orcid.org/0000-0003-4757-5832

DOI
http://dx.doi.org//10.7166/36-2-3194

ABSTRACT

Retail distributors routinely encounter a variety of supplier constraints, including those related to the size and timing of inventory replenishment orders, such as minimum order quantities and minimum reorder intervals. In addition, constraints on the transportation of replenishment stock, such as minimum order values and container filling requirements, necessitate the joint replenishment of commodities in the case of multi-commodity inventories. These supplier constraints typically lead to larger order quantities and excess inventory, giving rise to the need for a joint replenishment model that is capable of mitigating the impact of increased inventory levels. In this paper, we propose a novel simulation-optimisation modelling approach in support of the complicated joint replenishment decisions faced by inventory managers in the presence of the aforementioned constraints. The modelling approach is capable of consolidating multi-product inventory replenishment orders, guided by minimum reorder intervals and minimum order quantities, subject to container filling and/or minimum order value constraints. It aims to optimise safety stock levels heuristically with a view to reducing the increased inventory costs associated with the particular supplier constraints imposed. We show in a real-world case study that the modelling approach is capable of leading to substantial cost savings during joint inventory replenishment decisions.

OPSOMMING

In hierdie artikel stel ons ’n nuwe simulasie-optimerings-modelleringsbenadering ter ondersteuning van die ingewikkelde gesamentlike aanvullingsbesluite voor wat voorraadbestuurders in die gesig staar. Die modelleringsbenadering is in staat om multi-produk voorraadaanvullingsbestellings te konsolideer, en word deur minimum herbestelintervalle en minimum bestelhoeveelhede gerig, onderhewig aan houervullings- en/of minimum bestelwaardebeperkings, en is daarop gemik om veiligheidsvoorraadvlakke heuristies te optimeer met die doel om die verhoogde voorraadkoste wat verband hou met spesifieke verskaffersbeperkings te verminder. Ons demonstreer in ’n werklike gevallestudie dat die benadering daartoe in staat is om tot aansienlike kostebesparings gedurende gesamentlike voorraad-aanvullingsbesluite te lei.

1. INTRODUCTION

The literature on the *joint replenishment problem* (JRP) in inventory management has long focused on the design of algorithms that guarantee optimal solutions to instances of the classical JRP [1]. Recently, however, there has been a shift towards the design of models that are able to address this problem in more realistic practical settings, where distributors are often confronted with the imposition of a variety of resource and supplier constraints on their inventory replenishment decisions [2].

Constraints imposed by suppliers are typically aimed at facilitating production scheduling and ensuring the economic viability of commodity production. These supplier constraints usually manifest themselves as *minimum reorder intervals* (MRIs) and *minimum order quantities* (MOQs). Uniquely defined for each supplier, MRIs determine the permissible frequency at which consolidated replenishment orders can be placed with the supplier in question. MOQs, on the other hand, specify the smallest permissible number of stock units in any inventory replenishment order for a specific product [3]. Based on forecast demand over an MRI (a period during which replenishment orders for a product cannot be placed), distributors have to decide on suitable replenishment order quantities while adhering to MOQs. This is complicated, because larger replenishment quantities than those that realise as actual demand over the MRI will lead to unnecessary excess capital being tied down in inventory and higher inventory holding costs, while smaller-than-demand quantities may result in stock-outs and thus lower service levels.

In practice, suppliers also impose constraints aimed at ensuring the economic viability of product transportation, which typically take the form of *minimum order values* (MOVs) or container filling requirements (more generally referred to in the literature as "the pursuit of full truckloads"). MOV specifications result in suppliers being willing to accept replenishment orders only if the goods in each order conform to an MOV expressed in monetary terms. Container filling requirements, on the other hand, dictate that replenishment orders may be placed at suppliers who use container shipments only if these orders are economically justifiable in the sense that shipment containers should be utilised effectively by filling them up to at least some fraction of their full capacity, as specified by the distributor. When retail distributors endeavour to make inventory management decisions subject to MOV and container filling requirements, they have to order enough replenishment stock units at a time so as to meet the MOV or fill the container to a justifiable utilisation level, or else order no stock units at all. This reduces the degree of management flexibility at the disposal of retail distributors when they respond to customer demand, since distributors cannot order only what they actually need when they need it, but are instead forced to order additional stock units [4]. Again this creates a difficulty, in that inventory managers have to determine which products, and how many of each, to include in consolidated product replenishment orders so as to mitigate the impact of the increase in inventory levels that results directly from MOV and container filling requirements.

The problems associated with the constraints mentioned above are typically addressed only in isolation in the literature. The constraint requiring full truckloads, for instance, has only been tackled in the literature without allowing for the possibilities of production constraints such as MRIs and MOQs. To ensure full truckloads, Miltenburg [5] proposed including products in a consolidated replenishment order on the basis of the expected time to their next replenishment order, thereby aiming to maximise the time between consecutive consolidated replenishment orders. Van Eijs [6] and Cachon [7] considered the full-truckload constraint in the context of a periodic review policy. In the former study, uncoordinated replenishment order quantities were determined and then increased in a manner that minimised expected ordering and holding costs until a full truckload was achieved. In the latter study, replenishment orders were restricted until the order quantities of all items that were scheduled for replenishment met minimum transportation requirements or exceeded transportation capacity, thereby allowing only for the reduction of replenishment order quantities. Kiesmuller [8] later introduced a dynamic order-up-to policy that allowed for an increase or decrease of replenishment order quantities and that aimed to minimise holding and backorder costs. Kiesmuller [9] subsequently addressed the full-truckload constraint in the context of the QS policy initially proposed by Renberg and Planche [10] and by proposing an allocation policy. None of these contributions accommodates replenishment order quantities based on MRI and MOQ constraints that suppliers impose. Moreover, these approaches fail to achieve the target service level without incurring additional costs. In some cases, the service level is allowed to increase, with efforts focused on minimising the additional costs ([6], [8]), while in others, the service level deteriorates while the costs remain unchanged ([7]).

The literature therefore currently lacks a model for generating joint replenishment order quantities based on suppliers' MRIs and MOQs while ensuring full truckloads. Such a model would effectively reduce holding costs, achieve the desired target service levels, and maximise the time between consecutive consolidated replenishment orders.

To this end, we propose a novel simulation-optimisation model in this paper. The simulation component of the model incorporates a practical strategy for placing joint replenishment orders subject to MRIs, MOQs, container filling, and/or MOV requirements, as specified by suppliers, by following the logical approach of Miltenburg [5]. The optimisation component of the model goes further, however, by providing a method for determining suitable safety stock levels linked to an overall target service level with a view to mitigate against increased inventory levels that result from such supplier constraints. This is done in the context of empirical demand patterns, and is aimed at reducing the cost impact of consolidated replenishment. Lead times are treated as (non-zero) fixed values, as specified by the relevant suppliers, and the model does not require any demand time series forecasting. Instead, demand proxy data are sampled from the historical distribution of past demand realised.

The remainder of the paper is organised as follows. An overview of the most relevant literature may be found in §2. This is followed in §3 by a detailed presentation of the approach adopted in our simulation-optimisation model. The model is then applied in §4 to a case study involving two distinct suppliers, illustrating its potential impact in terms of cost saving while adhering to target service levels at a specified level of statistical confidence. The paper closes in §5 with an evaluation of the effectiveness of the suggested modelling approach and with recommendations for potential future follow-up work.

2. RELATED WORK

This section contains a brief review of the most pertinent literature on simulation and other modelling techniques in the area of inventory management. The review opens with a discussion in §2.1 on popular multi-product inventory replenishment models and this is followed in §2.2 by an explanation of how simulation may be used to determine safety stock levels.

2.1. Multi-product inventory replenishment models

The recent literature on inventory management has been oriented more towards multi-product inventory replenishment models, according to which orders for products are not placed individually but are grouped together so as to save on ordering and transportation costs. The problem associated with this type of inventory management is called the JRP. This problem may assume one of two forms: One involves deterministic/dynamic demand, and the other involves stochastic demand. While we focus primarily on the stochastic JRP in this paper, the reader is referred to Khouja and Goyal [1], Peng *et al.* [2], and Bastos *et al.* [11] for a comprehensive review of the entire JRP literature.

The first major contribution to solving the JRP in a stochastic setting appeared in a paper by Balintfy [12], in which the (s, c, S) policy was introduced. According to this policy, a replenishment order is triggered whenever the inventory level of a product drops below s, and all other products with inventory levels below or equal to c are included in the replenishment order. Later, Atkins and Iyogun [13] demonstrated that their base-stock level policy often outperforms the (s, c, S) policy. More recently, Feng *et al.* [14] introduced the so-called (s, c, d, S) policy as an enhancement of the (s, c, S) policy. This improvement allows for products to be ordered up to levels lower than S, with d representing the lowest order level, thereby effectively reducing holding costs.

The concept of an aggregate inventory level was first introduced by Renberg and Planche [10] in respect of their (A, S) policy. According to this policy, all products associated with a particular supplier are ordered up to level S when the aggregate inventory level of all products drops below an optimal level A. This notion was refined by Viswanathan [15] in the form of the $P(s, S)$ policy. A more recent approach to solving the JRP, by incorporating the notion of an aggregate inventory level, is the (A, S, T) policy proposed by Ozkaya *et al.* [16]. According to this policy, an order is triggered to replenish all products to a level S when the aggregate inventory level of all products drops below A, or once a periodic review period T has elapsed. This approach is a combination of the (A, S) and $P(s, S)$ policies (introduced by Renberg and Planche [10] and by Viswanathan [15], respectively).

Determining the optimal order-up-to level S in a practical setting for each product is not always feasible when considering all products simultaneously. Moreover, the questions of which products to order and when to order them change significantly in the presence of container filling and MOV requirements, as more than one replenishment order for a particular product is often included in the overall order. In the literature, this supplier constraint extension of the JRP is referred to as "achieving full truckloads," initially considered in the work of Miltenburg [5], Van Eijs [6], and Cachon [7]. Later, Kiesmüller [9] showed how these supplier constraints can be incorporated into classical inventory replenishment policies, such as the (Q, S) policy. Lately, there has been a notable shift towards addressing the JRP in the presence of resource and/or supplier constraints [2]. This shift is aimed at enhancing the practicality and applicability of such models.

2.2. Determining safety stock levels

The reader is referred to Gonçalves *et al.* [17] for a comprehensive review of the literature for the period 1977 to 2019 on approaches that have been adopted to determine safety stock values. This section covers notable simulation and simulation-optimisation approaches for safety stock determination, along with more recent works in these fields.

The pioneering work on safety stock determination by means of simulation is due to Eilon and Elmaleh [18], Wemmerlöv and Whybark [19], and Callarman and Hamrin [20]. These authors aimed compared inventory management policies with a view to ascertaining which policy results in the lowest cost. When evaluating inventory management policies in respect of cost, it is necessary to maintain a constant service level in all cases. The safety stock required to achieve a constant service level was determined by analysing exchange curves in the case of Eilon and Elmaleh [18], by carrying out a search routine in the case of Wemmerlöv and Whybark [19], and by invoking the service level decision rule in the case of Callarman and Hamrin [20].

In order to address the difficulties associated with the typically large number of simulations required to generate accurate results, but also to avoid the complexities of a search routine or regression analysis, Gudum and Kok [21] proposed the *safety stock adjustment procedure*. This procedure is based on a netting approach that ensures the independence of the net requirement and replenishment processes from the safety stock under the assumption of a *time-phased order point* policy [22]. Recently Da Cruz *et al.* [23] and Zhou and Zhang [24] used respectively Monte Carlo simulation and a simulation-based Bayesian network to determine appropriate safety stocks.

In the realm of simulation-optimisation aimed at determining target safety stock levels, Bouslah *et al.* [25] determined the economic production quantity, optimal safety stock level, and economic sampling plan associated with the minimisation of total costs. They used a stochastic model based on the *response surface methodology* (RSM). Gansterer *et al.* [26] adopted a discrete-event simulation approach in conjunction with six optimisation methodologies to determine "good" values for lead times, safety stock, and lot sizes. OptQuest, RSM, and *variable neighbourhood search* formed the basis of these optimisation methodologies. More recently, Avci and Selim [27] proposed a *decomposition-based multi-objective differential evolution* (MODE/D) algorithm for determining safety stock values. This approach was later enhanced to incorporate supply and demand uncertainties in convergent supply chains [27]. Hosseini *et al.* [29] introduced a reinforcement learning-based simulation-optimisation model to determine, among other objectives, the optimal safety stock levels.

To the best of our knowledge, no research has been documented on the determination of safety stock values in the case where container filling and MOV requirements are imposed by suppliers, let alone on simulation models or simulation-optimisation models aimed at addressing this specific problem.

3. PROPOSED SIMULATION-OPTIMISATION MODEL

In this section, we present a simulation-optimisation approach for modelling joint replenishment orders in the presence of supplier constraints that aim to determine appropriate safety stock durations (when safety stock is expressed in respect of time rather than product quantity) associated with the products of a single retail supplier. The approach is aimed at achieving a specified target service level while taking into account container filling and MOV requirements. The section opens in §3.1 with a high-level overview of the proposed modelling approach, after which discussions follow in §3.2 on simulating daily sales (based on empirical distributions) and invoking the inverse-transform method. Our base method of modelling product replenishment orders independently from one another is described in §3.3, after which the focus of discussion turns in §3.4 to the consolidation of these replenishment orders. Our approach to modelling inventory conservation is explained briefly in §3.5, before the model components responsible for service level computation and confidence level establishment are described in §3.6. The section closes in §3.7 with a discussion of five heuristic approaches to determining recommended safety stock durations.

For a more detailed explanation of the simulation-optimisation model and its various components, the reader is directed to Winter and Van Vuuren [30] for a worked example that illustrates the functionality of the model. This example is based on information from a hypothetical supplier. Similar examples were used to validate our proposed simulation-optimisation model.

3.1. Model framework

Figure 1 contains a schematic illustration of the simulation and optimisation components of the framework proposed in this paper. The simulation component represents our recommended approaches to modelling sales, determining order sizes, consolidating replenishments, and calculating service levels. These approaches can vary from company to company and also across sectors, depending on the inventory policy in place. In contrast, the optimisation part of the model is universal and can be applied to all companies, regardless of the inventory policy adopted. This component adjusts the safety stock durations associated with each product, based on one of five heuristic approaches. These adjustments are made in accordance with the service level calculated in the simulation component.

3.2. Modelling sales

Let $\mathcal{T} = \{1, \dots, T\}$ index the time intervals (typically days or weeks) during the planning horizon under consideration, and let $\mathcal{P}$ index the products in the basket of some supplier of the retail distributor under consideration.

As a simulation run progresses through each time period of the planning horizon, customers are modelled as collectively purchasing $S_p(t)$ units of product $p \in \mathcal{P}$ from the distributor during time period $t \in \mathcal{T}$. A Monte Carlo simulation approach is adopted to determine the value of $S_p(t)$, driven by historical sales patterns. These historical patterns take the form of a sales data set comprising past product-specific sales records, typically encompassing a wide range of sales quantities per time period, from zero (indicating no sales) to substantial volumes. The empirical distribution of the historical sales of product $p \in \mathcal{P}$ serves as the foundation for stochastically generating the sales quantity $S_p(t)$ associated with each time period $t \in \mathcal{T}$.

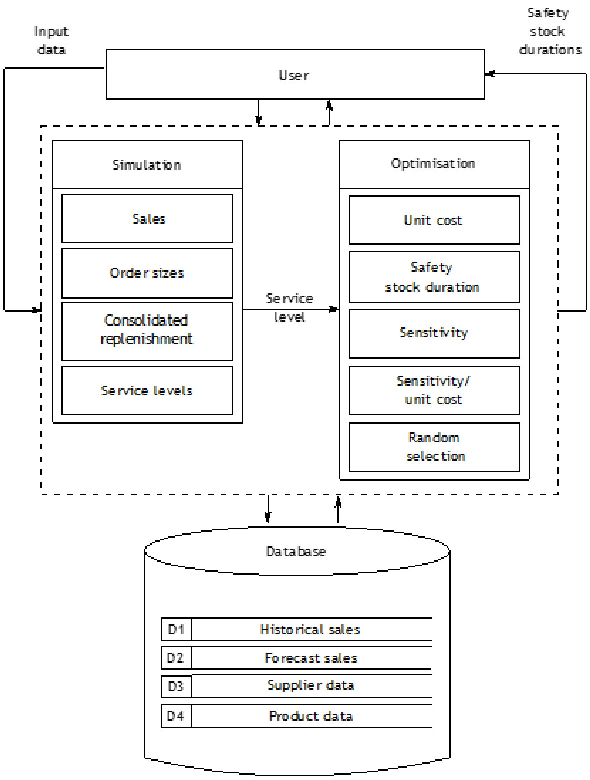

Figure 1: Simulation-optimisation modelling approach to inventory replenishment order consolidation

The procedure begins by categorising the historical sales quantities per time period into bins of approximately equal width. These bins span the entire non-outlier range of historical sales quantities per time period, starting from the lowest sales quantity (excluding zero) to the highest recorded sales quantity that is not considered to be an outlier. A dedicated bin is created to represent the occurrences of zero sales per time period, given the typical frequency of such occurrences. Another dedicated bin is created for outlier sales quantities per time period, extending up to infinity, in order to account for a small probability of sales quantities being realised in this range. A count of historical sales quantities in each bin is performed and subsequently normalised by division by the total historical sales count of product $p \in \mathcal{P}$ over the entire planning horizon. A well-known rule of thumb [31] is that bin widths should be chosen so that the sales count in each bin is at least five and so that there are no more than 30 bins. Cumulatively, these bins form an empirical *cumulative distribution function* (CDF) of the sales quantity of product $p \in \mathcal{P}$ per time period.

The inverse-transform method is then applied to generate sales quantities per time interval stochastically during the execution of a simulation run. This technique involves generating a random number, denoted by U, from a uniform distribution ranging between 0 and 1. The value of U is then mapped by the CDF to a corresponding bin. Thereafter, a random number is generated from a uniform distribution ranging between the confines of the bin thus identified. For the outlier bin, however, a random number is drawn from an exponential distribution within the bin's boundaries. This approach is adopted to increase the likelihood of selecting lower sales quantities in the outlier bin, as the exponential distribution tapers off quickly at large values. This helps to ensure that higher sales quantities in the outlier bin are not over-represented among the sales quantities generated.

The resulting value serves as the simulated sales quantity for a specific product during a particular time interval. This process is repeated to produce sales quantities for every time interval within the planning horizon, considering each distinct product $p \in \mathcal{P}$ separately.

3.3. Modelling order sizes

Each product $p \in \mathcal{P}$ is associated with the following three inventory-related variables: A total inventory level at the end of time interval $t \in \mathcal{T}$ denoted by $L_p(t)$; an inventory-on-hand value at the end of time interval t denoted by $I_p(t)$; and potentially the order quantity of a pending inventory replenishment order denoted by $P_p(t)$ (representing a product quantity ordered from the supplier) that has not yet been received by the end of time period $t \in \mathcal{T}$. These three variables satisfy the relationship

$$L_p(t) = I_p(t) + P_p(t), \quad p \in \mathcal{P}, t \in \mathcal{T}, \tag{1}$$

where $L_p(t) \geq 0$, $I_p(t) \geq 0$ and $P_p(t) \geq 0$.

The reorder point (ROP) of a product serves as a threshold level that triggers the need for an inventory replenishment order. If the lead time of the supplier is denoted by λ, this threshold level is calculated for each product $p \in \mathcal{P}$ at the end of time interval t as

$$r_p(t) = \sum_{\tau=t+1}^{t+\lambda+1+s_p} S_p^*(\tau), \quad p \in \mathcal{P}, t \in \mathcal{T}, \tag{2}$$

where $S_p^*(t)$ denotes the sales quantity of product $p \in \mathcal{P}$ anticipated for time period $t \in \mathcal{T}$, and s_p denotes the safety stock duration for product $p \in \mathcal{P}$ under the condition that $s_p > 0$ to ensure that buffer inventory is held in order to protect against inventory stock-outs as a result of supplier and customer sales uncertainty. This ensures that the replenishment order for product $p \in \mathcal{P}$ arrives when the lead time λ has elapsed. That is, when

$$I_p(t) = \sum_{\tau=t+\lambda+1}^{t+\lambda+1+s_p} S_p^*(\tau), \quad p \in \mathcal{P}, t \in \mathcal{T} \tag{3}$$

(i.e., when the inventory on hand equals the safety stock quantity). When the inventory level $L_p(t)$ of any product falls below its corresponding ROP value (i.e., when $L_p(t) - r_p(t) \leq 0$), an order is placed with the supplier to ensure an adequate stock level for product $p \in \mathcal{P}$ throughout. If a periodic review policy is in place for determining the inventory-on-hand value for each product $p \in \mathcal{P}$ at the end of every periodic review period, the relationship between $L_p(t)$ and $r_p(t)$ will only be determined after the periodic review period has elapsed, since the inventory-on-hand value is required to calculate the inventory level value as described in (1).

Now suppose that the product subset

$$\mathcal{C}(t) \subseteq \mathcal{P}, \quad t \in \mathcal{T} \tag{4}$$

indexes the assortment of distinct products for which inventory has to be replenished during time interval $t \in \mathcal{T}$ when the inventories of these products are considered completely independently of one another (i.e., if no inventory replenishment order consolidation takes place). Also, let $R_p(t)$ denote the number of units of product $p \in \mathcal{P}$ that are to be included in the replenishment order during time period $t \in \mathcal{T}$, still under the assumption that products are considered independently of one another. Then

$$R_p(t) \begin{cases} = 0 & \text{if } p \notin \mathcal{C}(t), \\ > 0 & \text{if } p \in \mathcal{C}(t) \end{cases} \tag{5}$$

for any $p \in \mathcal{P}$ and $t \in \mathcal{T}$. The order replenishment size $R_p(t)$ is determined by invoking the concept of an MRI. Uniquely defined for each supplier, the MRI determines the permissible frequency at which consolidated replenishment orders can be placed at the supplier in question. For instance, an MRI of 30 days indicates that at least 30 days should elapse between successive orders placed at the particular supplier. Therefore, the order size has to be chosen carefully to ensure that the available inventory can adequately facilitate sales over this MRI, at a specified level of statistical confidence.

Given the MRI, denoted here by μ, and the lead time of the supplier (both assumed to be constant), the replenishment order size may be computed as

$$R_p(t) = r_p(t) - L_p(t) + \sum_{\tau=t+\lambda+1}^{t+\lambda+1+\mu} S_p^*(\tau), \quad p \in \mathcal{P}, t \in \mathcal{T}. \tag{6}$$

In addition, each product is assigned an MOQ and an *increase lot size* (ILS) value, denoted by E_p and e_p, respectively. The condition $R_p(t) \geq E_p$ is imposed to ensure that the order size meets or exceeds the MOQ. Moreover, any order quantity exceeding the MOQ necessarily has to be a multiple of the ILS. This leads to the relationship

$$R_p(t) = \begin{cases} E_p & \text{if } R_p(t) \leq E_p, \\ E_p + ne_p & \text{if } R_p(t) > E_p \end{cases}$$

for any $p \in \mathcal{P}$ and $t \in \mathcal{T}$, where n is the smallest positive integer for which $E_p + ne_p \geq R_p(t)$.

3.4. Modelling inventory replenishment order consolidation

If the inventory level of a single product in the supplier's basket falls below its ROP, orders are likely necessitated for other products as well, given the intricacies of container filling and MOV considerations (irrespective of whether their inventory levels have fallen below their corresponding ROP values), giving rise to the opportunity for a joint replenishment order. The product replenishment timings and quantities are consolidated according to a heuristic adjustment process during which the timings of product replenishments are brought forward with a view to improve on the supplier's utilisation of delivery containers or to meet the required MOV imposed by the supplier. When a product $p \in \mathcal{P}$ is included in a consolidated replenishment order, its adjusted replenishment order size in the consolidated replenishment order is denoted by $R'_p(t)$. The replenishment product index sets $\mathcal{C}(t)$ in (4) are also updated to reflect replenishment order consolidations, and these updated sets are denoted by $\mathcal{C}'(t)$ for all $t \in \mathcal{T}$.

The determination of which product replenishments are to be included in a joint replenishment order involves carrying out an inventory transactions simulation run for each product in the supplier's basket in isolation from other products, as described in §3.3. The orders placed for each product during these simulation runs are compiled in tabular form, organised chronologically based on ROPs (with the earliest ROP taking precedence, followed by subsequent ROPs in sequence, as suggested by Miltenburg [5]). Separate product replenishment orders are considered one by one, in the chronological order specified by this table of ROPs, for inclusion in a larger consolidated replenishment order. If a product replenishment order does not violate any of the constraints associated with container filling imposed by the supplier, then it is included in the consolidated order. Otherwise, the next product replenishment order in the sequence is considered for inclusion, and so on.

Consider, for example, an inventory of three products supplied by the same supplier, each of unit volume and indexed by the set $\mathcal{P} = \{1,2,3\}$, over a planning horizon of ten time intervals, indexed by the set $\mathcal{T} = \{1, \ldots, 10\}$. Suppose the volume capacity of the delivery containers used by the supplier of these products is 600 volumetric units, and that the target minimum utilisation of the containers is 90% of their volumetric capacity (that is, the contents of each container should have a volume of between $0.9 \times 600 = 540$ and 600 units). Finally, suppose the appropriate inventory replenishment time intervals for these products, when viewed independently of one another, are specified by the sets $C(1) = \{1,2,3\}$, $C(2) = \emptyset$, $C(3) = \{2\}$, $C(4) = \emptyset$, $C(5) = \{1,3\}$, $C(6) = \emptyset$, $C(7) = \emptyset$, $C(8) = \{1,2\}$, $C(9) = \emptyset$, and $C(10) = \{1,3\}$, adopting the same notation as introduced in the previous section, while the order quantities corresponding to these replenishment timings are as shown in Table 1.

Table 1: Replenishment timings and corresponding order quantities for three hypothetical products over a ten-period planning horizon

$R_1(1)$ = 230	$R_2(1)$ = 160	$R_3(1)$ = 50
$R_1(2)$ = 0	$R_2(2)$ = 0	$R_3(2)$ = 0
$R_1(3)$ = 0	$R_2(3)$ = 110	$R_3(3)$ = 0
$R_1(4)$ = 0	$R_2(4)$ = 0	$R_3(4)$ = 0
$R_1(5)$ = 320	$R_2(5)$ = 0	$R_3(5)$ = 35
$R_1(6)$ = 0	$R_2(6)$ = 0	$R_3(6)$ = 0
$R_1(7)$ = 0	$R_2(7)$ = 0	$R_3(7)$ = 0
$R_1(8)$ = 260	$R_2(8)$ = 190	$R_3(8)$ = 0
$R_1(9)$ = 0	$R_2(9)$ = 0	$R_3(9)$ = 0
$R_1(10)$ = 200	$R_2(10)$ = 0	$R_3(10)$ = 170

Consolidation of the product replenishment orders of the three products during time interval 1 is permissible according to the upper bound on the volume of the replenishment inventory in a delivery container, because $R_1(1) + R_2(1) + R_3(1) = 230 + 160 + 50 = 440 \leq 600$ volumetric units. This consolidation is not sufficient, however, in reference to the 90% target lower bound on the volume of the replenishment inventory in a delivery container, because $R_1(1) + R_2(1) + R_3(1) = 440 \ngeq 540$ volumetric units. The product replenishment orders for product 2 during time interval 3 and for product 3 during time interval 5 may also be included in the joint replenishment order to yield a total replenishment inventory volume of 585 volumetric units during time interval 1, which is indeed between the lower bound of 540 and the upper bound of 600 volumetric units. This replenishment consolidation order corresponds to the adjusted order quantities $R_1'(1) = 230$, $R_2'(1) = 160 + 110 = 270$ and $R_3'(1) = 50 + 35 = 85$ (using the same colour shadings as in Table 1). In a similar manner, the product order quantity $R_1(5) = 320$ induces another consolidated replenishment order for time interval 5 corresponding to the adjusted order quantities $R_1'(5) = 320 + 260 = 580$, $R_2'(5) = 0$ and $R_3'(5) = 0$. This yields a total replenishment inventory volume of 580 volumetric units during time interval 5, which is again between the lower bound of 540 and the upper bound of 600 volumetric units. Finally, the product order quantity $R_2(8) = 190$ induces another consolidated replenishment order for time interval 8 corresponding to the adjusted order quantities $R_1'(8) = 200$, $R_2'(8) = 190$ and $R_3'(8) = 170$. This yields a total replenishment inventory volume of 560 volumetric units during time interval 8, which is yet again between the lower bound of 540 and the upper bound of 600 volumetric units. It follows that $\mathcal{C}'(1) = \{1,2,3\}$, $\mathcal{C}'(5) = \{1\}$ and $\mathcal{C}'(8) = \{1,2,3\}$. (Note that $R_p'(t) = 0$ and $\mathcal{C}'(t) = \emptyset$ for all $p \in \mathcal{P}$ and all $t \in \{2,3,4,6,7,9,10\}$.)

3.5. Modelling inventory conservation

The dynamics of a product's inventory are subject to changes that are influenced (positively) both by incoming consolidated inventory replenishments delivered by its supplier and (negatively) by product sales. Therefore, the inventory of product $p \in \mathcal{P}$ on hand at the end of time interval $t \in \mathcal{T}$ satisfies the conservation equation

$$I_p(t) = I_p(t-1) + R'_p(t-\lambda) - S_p(t), \quad p \in \mathcal{P}, t \in \mathcal{T}. \tag{7}$$

Here $I_p(0)$ denotes the inventory of product $p \in \mathcal{P}$ on hand at the start of the planning horizon, while a value of I_p or R'_p evaluated at a non-positive time index value represents respectively the inventory on hand and the consolidated replenishment order of product $p \in \mathcal{P}$ during some time interval of the previous planning horizon.

The simulation model is subject to two constraints linked to container filling during each consolidated inventory replenishment. First, sets of volume constraints of the form

$$\alpha V_{max} \leq \sum_{p \in \mathcal{C}'_p(t)} R'_p(t) v_p \leq V_{max}, \quad p \in \mathcal{P}, t \in \mathcal{T} \tag{8}$$

are imposed, where v_p denotes the volume (in m^3) of a single unit of product $p \in \mathcal{P}$, α is a (unitless) fraction between zero and one (non-inclusive), and V_{max} is the volume capacity (in m^3) of a delivery container used by the supplier in question. This constraint ensures that between $100\alpha\%$ and 100% of the capacity of each container of the supplier used to deliver replenishment inventory is, in fact, utilised (as illustrated in §3.4).

Concurrently, a weight constraint set also comes into play, according to which an upper bound is placed on the total inventory mass that may be transported in each container. This constraint takes the form

$$\sum_{p \in \mathcal{C}'_p(t)} R'_p(t) m_p \leq M_{max}, \quad p \in \mathcal{P}, t \in \mathcal{T}, \tag{9}$$

where m_p denotes the mass (in kg) of a single unit of product $p \in \mathcal{P}$ and M_{max} is the content mass capacity (in kg) of a delivery container used by the supplier under consideration. The use of local suppliers typically introduces an MOV criterion, necessitating the value of the entire consolidated inventory replenishment order to meet or exceed this threshold. This may be expressed mathematically as

$$\sum_{p \in \mathcal{C}'_p(t)} R'_p(t) c_p \geq A_{max}, \quad p \in \mathcal{P}, t \in \mathcal{T}, \tag{10}$$

where c_p denotes the cost of a single unit of product $p \in \mathcal{P}$ and A_{max} is the MOV required by the specific supplier in question.

While the constraints in (9) and (10) were not explicitly considered in the numerical example of §3.4, order consolidation, in fact, takes into account not only the volume constraints in (8), but also the weight constraint in (9) and, if required, the MOV constraint in (10). If any of these constraints are violated by the potential inclusion of a product replenishment order in a consolidated replenishment order, that product replenishment order is not included in the consolidated order. Instead, the next potential product replenishment order in a table similar in form to Table 1 is considered for inclusion in the consolidated order, and so on.

3.6. Service and confidence levels

The overall service level ζ_j associated with a simulation run j is determined by performing a count of the purchases that could be fulfilled from inventory on hand over the entire planning horizon during that simulation run. A successful customer purchase of product $p \in \mathcal{P}$ during time interval $t \in \mathcal{T}$ transpires when the inventory on hand for the desired product at the time of the sale is equal to or exceeds the quantity purchased ($i.e.$ if $I_p(t-1) \geq S_p(t)$). Conversely, an unsuccessful purchase occurs when $I_p(t-1) < S_p(t)$. It therefore follows that

$$\zeta_j = \frac{\text{\# of successful purchases during run } j}{\text{total \# of purchases during } j}. \quad (11)$$

In order to determine the appropriate number of simulation runs Ω^* for a simulation replication required to obtain an estimate of the true overall service level ζ^* with sufficient statistical confidence, an initial set of Ω = 50 (say) simulation runs may be performed. After these initial runs, the mean overall service level estimate can be computed as

$$\zeta^* = \frac{\sum_{j=1}^{\Omega} \zeta_j}{\Omega} \quad (12)$$

with a variance of

$$\mathrm{Var}(\zeta^*) = \frac{\sum_{j=1}^{\Omega} \left(\zeta_j - \zeta^*\right)^2}{\Omega - 1} \quad (13)$$

over these Ω = 50 simulation runs. At a user-specified statistical significance of α, the confidence interval half-width for the service level estimate ζ^* may then be calculated from the cumulative Student t-distribution as

$$h(\Omega) = t_{\Omega-1;1-\frac{\alpha}{2}} \sqrt{\frac{\mathrm{Var}(\zeta^*)}{\Omega}}. \quad (14)$$

This initial half-width $h(\Omega)$ can now be used to determine the number of simulation runs Ω^* required to obtain a user-desired half-width h^* associated with a simulation replication, as

$$\Omega^* = \left\lceil \Omega \left(\frac{h(\Omega)}{h^*} \right)^2 \right\rceil. \quad (15)$$

If this number of simulation runs were to be performed, then a confidence level of 100(1 − α)% may be associated with the product safety stock durations recommended by the simulation replication, yielding an overall service level of ζ^* with a confidence interval half-width h^*.

3.7. Safety stock calculation

Given the complexity of the problem of estimating individual product safety stock durations s_p for all $p \in \mathcal{P}$ that will lead (with sufficient statistical confidence) to a particular, pre-specified overall service level $\bar{\zeta}$, a heuristic approach is adopted to vary the product safety stock durations iteratively between successive replications of Ω^* simulation runs. After having completed all Ω^* simulation runs of such a replication, the product safety stock duration s_p for a specific product $p \in \mathcal{P}$ is reduced by one day if $\zeta^* > \bar{\zeta}$.

Several heuristic approaches may be considered for selecting the product for which the safety stock duration s_p is to be reduced, such as:

- **Unit cost:** Altering the safety stock duration of either the most expensive or the least expensive product.
- **Safety stock duration:** Modifying the safety stock duration of the product with the largest safety stock duration or the product with the smallest current safety stock duration.
- **Sensitivity:** Adjusting the safety stock duration of the product that precipitates the most frequent orders (i.e., it reaches its ROP first) or the product with the fewest such occurrences.
- **Sensitivity/unit cost:** Refining the safety stock duration of either the product with the largest sensitivity-to-unit cost ratio or the product with the smallest value for this ratio. This ratio is computed by dividing the sensitivity associated with each product $p \in \mathcal{P}$ by the unit cost associated with each product $p \in \mathcal{P}$.

- **Random selection:** After each simulation replication, one of the four aforementioned heuristics is randomly selected according to a uniform distribution, and is executed to reduce the safety stock duration of the particular product.

In order to prevent the safety stock duration of a product from dropping below zero when applying any of these heuristics, thereby rendering the solution infeasible, a threshold factor δ_p specific to each product is incorporated. If the safety stock duration of the product under consideration drops below the threshold level $\delta_p s_p$, the next most suitable product, as determined by the specific heuristic, is selected and adjusted accordingly. A value of $\delta = 0.5$ may, for instance, be used for all $p \in \mathcal{P}$. This value may, however, be adjusted on the basis of company specifics, such as minimum safety stock duration constraints.

The entire simulation process terminates when, on completion of all Ω^* runs in a simulation replication, it holds that $\zeta^* \geq \bar{\zeta}$ (at the desired level of statistical confidence) and further lowering any product's safety stock duration results in a situation in which $\zeta^* < \bar{\zeta}$ (at the desired level of statistical confidence).

4. ILLUSTRATIVE CASE STUDY

The simulation-optimisation modelling approach of the previous section was applied to a practical case study involving real supply chain data, as described in §4.1. A description is also provided in §4.2 of the *status quo* at the industry partner in question in respect of decisions pertaining to the case study suppliers. The results returned by the simulation-optimisation model are presented and discussed in §4.3.

4.1. The case study data set

The case study in this section is based on two distinct suppliers, referred to here as Suppliers A and B, of a South African retail distributor, together with the products associated with these suppliers. The suppliers were deliberately selected to include an instance of a supplier imposing container filling constraints and an instance of a supplier imposing MOV requirements. Moreover, the suppliers vary in the size of their product portfolios, having basket sizes of ten products and fifty products, respectively.

Each supplier's product inventory includes several product-specific attributes, all of which are available in Excel file format[1]. These attributes are product IDs, current safety stock durations, MOQs, ILS values, product unit volumes, product unit costs, product unit weights, and the stock-on-hand values on 30/09/2023. Information is also included about replenishment orders that had already been placed, but were yet to arrive during the planning horizon of the case study. These orders were scheduled to arrive during the execution of the simulation-optimisation model.

Forecast sales data provided to the authors pertained to the sales quantities for each product in the baskets of the two suppliers in question over the following 365 days, for the period 01/10/2023 to 30/09/2024. This horizon aligns with the time period over which the simulation-optimisation model of §3 would be executed. The historical sales data set contains information about the sales history of each product from three years prior to the case study start date (that is, for the 1 095 days in the period 01/10/2020 to 30/09/2023).Table 2 contains a summary of the supplier-specific data considered during the case study.

Table 2: Supplier-specific parameters used to evaluate the performance of the simulation-optimisation model of §3 in the context of the case study

Supplier	A	B
Product basket size	10	50
Constraint	MOV	Container filling
MOV	R 100 000	–
Lead time	42 days	105 days
Target service level	97.0%	96.8%
Desired confidence level	98%	98%
Planning horizon	120 days	365 days
Review period	7 days	7 days
MRI	14 days	14 days

[1] Available at: https://github.com/InventoryManager101/Supporting-Material.

4.2. The current situation

Data pertaining to the overall service levels actually achieved by the industry partner for the period 01/10/2020 to 30/09/2023, as influenced by the effects of MOV or container filling requirements, were unfortunately not available. In order to estimate these as-is service levels for the two suppliers under consideration (before applying the simulation-optimisation model of §3 to generate any stock duration recommendations), the simulation part of the model described in §3 was invoked with the existing safety stock durations adopted by the industry partner at the time of writing this paper. (These stock durations are also available in Excel file format[1].) The resulting service level estimates, together with estimates of the corresponding average daily inventory values, may be found in Table 3 for the products of each of the two suppliers under consideration. A confidence level of 98% is attached to the estimates in the table.

Table 3: Estimates of the existing service levels achieved and the average daily inventory values for the products of the two case study suppliers, as determined with 98% confidence by the simulation part of the model of §3, when adopting the actual safety stock durations applied by the industry partner at the time of writing

Supplier	**A**
Service level achieved	97.3%
Average daily inventory value	R 621 112.44
Supplier	**B**
Service level achieved	97.5%
Average daily inventory value	R7 251 292.56

When comparing the estimates in Table 3 with the data in Table 2, it may be confirmed that the container filling and MOV requirements led to increases in the overall service level associated with each supplier. The safety stock durations applied by the industry partner were estimated to yield overall service levels of 97.3% and 97.5%, respectively, for the products of the two suppliers. As expected, these levels are higher than the target service levels for these suppliers, which are 97.0% and 96.8%, respectively.

4.3. Model application to the case study data

The results returned by the simulation-optimisation model of §3 on achieving the target service levels of $\bar{\zeta} = 97.0\%$ and $\bar{\zeta} = 96.8\%$, respectively, are summarised in Table 4. The corresponding safety stock durations recommended for each product by each of the heuristics may be found in Appendix A.

Table 4: Performance summary of each heuristic optimisation approach to achieving the target service level and reducing the average daily inventory costs for two suppliers

Supplier A	Unit cost	Safety stock duration
98% confidence interval for service level achieved	97.01% ± 0.11%	97.11% ± 0.12%
Number of simulation-optimisation replications	5	13
Computation time (hh:mm)	01:54	04:33
Average daily inventory value	R 615 572.83	R 613 339.38
Average daily inventory value reduction	R 5 539.61	R 7 773.05
Tied-up capital reduction	0.89%	1.25%

Supplier A	Sensitivity	Sensitivity/unit cost	Random selection
98% confidence interval for service level achieved	97.07% ± 0.14%	97.11% ± 0.13%	97.04% ± 0.16%
Number of simulation-optimisation replications	7	6	15
Computation time (hh:mm)	02:26	02:07	05:03
Average daily inventory value	R 613 633.29	R 613 790.53	R 612 637.94
Average daily inventory value reduction	R 7 479.15	R 7 321.91	R 8 474.49
Tied-up capital reduction	1.20%	1.18%	1.36%

Supplier B	Unit cost	Safety stock duration
98% confidence interval for service level achieved	96.93% ± 0.12%	96.85% ± 0.13%
Number of simulation-optimisation replications	113	466
Computation time (hh:mm)	31:41	132:02
Average daily inventory value	R 7 184 457.05	R 6 827 594.37
Average daily inventory value reduction	R 66 828.91	R 423 691.59
Tied-up capital reduction	0.92%	5.84%

Supplier B	Sensitivity	Sensitivity/unit cost	Random selection
98% confidence interval for service level achieved	96.87% ± 0.09%	96.85% ± 0.09%	96.83% ± 0.08%
Number of simulation-optimisation replications	111	168	196
Computation time (hh:mm)	30:16	32:34	33:51
Average daily inventory value	R 7 220 793.45	R 7 036 873.46	R 7 128 379.10
Average daily inventory value reduction	R 30 492.69	R 214 412.49	R 122 906.85
Tied-up capital reduction	0.42%	2.96%	1.69%

For Supplier A, all five heuristic approaches successfully achieved the target service level and reduced the average inventory value with a high degree of statistical confidence. Among these approaches, the random selection heuristic performed the best in terms of reducing the average daily inventory value, but was computationally the least efficient. After having conducted fourteen simulation replications over a period of five hours and three minutes, this approach managed to achieve a reduction in the average daily inventory value of $R\,8\,474.49$. This translates to a total reduction of $R\,1\,016\,938.80$ over the entire planning horizon of $T = 120$ days. While such a 1.36% reduction in inventory value may appear modest, it should be noted that distributors often deal with more than a thousand suppliers [32], making the inventory cost savings substantial.

Although all five heuristic approaches also successfully achieved the target service level and reduced the average inventory value with a high degree of statistical confidence for Supplier B, it is evident that the required computation time escalated significantly with the increase in the size of a supplier's product basket. The safety stock duration heuristic outperformed all other heuristics in respect of inventory cost savings, but performed the worst in respect of computational efficiency. This heuristic yielded a service level estimate of $\zeta^* = 96.85\%$ after 466 simulation replications and a total simulation run time of 132 hours. It nevertheless managed to lower the average daily inventory cost for products in the supplier's basket by $R\,423\,691.59$, as shown in Table 4, which is equivalent to a total inventory cost reduction of $R\,154\,647\,430.40$ (5.84%) over the entire planning horizon of $T = 365$ days.

These results underscore the value of a simulation-optimisation model in suggesting target safety stock durations aimed at mitigating the effects of container loading and MOV. In both supplier scenarios, a trade-off has to be struck between improved results in terms of inventory cost reduction and computational efficiency. Importantly, once the simulation model has been executed, the results obtained remain valid for the entirety of the planning horizon. For instance, in the case of Supplier B, the simulation-optimisation model only has to be run once a year (the planning horizon). Therefore, a simulation runtime of five days is negligible in the context of the significant cost savings that are achieved.

5. CONCLUSION

In this paper, we offered inventory managers the opportunity to apply a simulation-optimisation model aimed at enhancing the process of consolidating their replenishment orders in a manner that caters for the container-filling and/or MOV requirements imposed on them by suppliers. The results reported in §4 would seem to indicate that the simulation-optimisation model is capable of reducing safety stock durations so as to achieve a specified target service level and, in doing so, significantly decreasing the inventory holding costs. For the two case study suppliers considered, reductions of 1.36% and 5.84%, respectively, were achieved. Five heuristic safety stock adjustment approaches were suggested, although the selection of a suitable heuristic approach depends on the company's specific needs (e.g., requiring a trade-off between inventory value reduction and computational efficiency).

Possible variations in supplier lead times may be addressed in follow-up future work, as this variability poses difficulties for retail distributors in respect of predicting the arrival of replenishment inventory. Moreover, the case where distributors permit backorders may also be considered. Finally, the implementation of a natural selection-inspired heuristic that chooses the best-performing heuristic stochastically, based on its propensity to reduce inventory costs after each simulation replication, might also be a valuable contribution, and is anticipated to outperform the heuristics proposed in this paper. Given the computational expense associated with executing the simulation model Ω^* times for each heuristic per replication, however, this particular type of heuristic was not considered.

REFERENCES

[1] **Khouja, M., & Goyal, S.** 2008. A review of the joint replenishment problem literature: 1989–2005. *European Journal of Operational Research,* 186(1), pp 1–16. https://doi.org/10.1016/j.ejor.2007.03.007

[2] **Peng, L., Wang, L., & Wang, S.** 2022. A review of the joint replenishment problem from 2006 to 2022. *Management System Engineering*, 1(9). https://doi.org/10.1007/s44176-022-00010-3

[3] **Zhu, H., Liu, X., & Chen, Y. F.** 2015. Effective inventory control policies with a minimum order quantity and batch ordering. *International Journal of Production Economics,* 168, pp 21–30. https://doi.org/10.1016/j.ijpe.2015.06.008

[4] **Chow, P. S., Choi, T. M., & Cheng, T. C. E.** 2012. Impacts of minimum order quantity on a quick response supply chain. *IEEE Transactions on Systems, Man, and Cybernetics — Part A: Systems and Humans,* 42(4), pp 868–879. https://doi.org/10.1109/TSMCA.2012.2183351

[5] **Miltenburg, G. J.** 1985. Allocating a replenishment order among a family of items. *IEE Transactions,* 17(3), pp 261–267. https://doi.org/10.1080/07408178508975301

[6] **Van Eijs, M. J. G.** 1994. Multi-item inventory systems with joint ordering and transportation decisions. *International Journal of Production Economics,* 35(1–3), pp 285–292. https://doi.org/10.1016/0925-5273(94)90094-9

[7] **Cachon, G.** 2001. Managing a retailer's shelf space, inventory, and transportation. *Manufacturing & Service Operations Management,* 3(3), pp 211–229. https://doi.org/10.1287/msom.3.3.211.9893

[8] **Kiesmüller, G. P.** 2009. A multi-item periodic replenishment policy with full truckloads. *International Journal of Production Economics,* 118(1), pp 275–281. https://doi.org/10.1016/j.ijpe.2008.08.024

[9] **Kiesmüller, G. P.** 2010. Multi-item inventory control with full truckloads: A comparison of aggregate and individual order triggering. *European Journal of Operational Research*, 200(1), pp 54–62. https://doi.org/10.1016/j.ejor.2008.12.008

[10] **Renberg, B., & Planche, R.** 1967. Un modèle pour la gestion simultanée des n articles d'un stock. *Revue Francaise d'Informatique et de Recherche Operationnelle*, 6, pp 47–59.

[11] **Bastos, L. S. L., Mendes, M. L., Nunes, D. R. L., Melo, A. C. S., & Carneiro, M. P.** 2017. A systematic literature review on the joint replenishment problem solutions: 2006-2015. *Production,* 27. https://doi.org/10.1590/0103-6513.222916

[12] **Balintfy, J. L.** 1964. On a basic class of multi-item inventory problems. *Management Science*, 10(2), pp 287–297. https://www.jstor.org/stable/2627299

[13] **Atkins, D. R., & Iyogun, P. O.** 1988. Periodic versus 'can-order' policies for coordinated multi-item inventory systems. *Management Science*, 34(6), pp 791–796. https://www.jstor.org/stable/2632132

[14] **Feng, H., Wu, Q., Muthuraman, K., & Deshpande, V.** 2015. Replenishment policies for multi-product stochastic inventory systems with correlated demand and joint-replenishment costs. *Production and Operations Management,* 24(4), pp 647–664. https://doi.org/10.1111/poms.12290

[15] **Viswanathan, S.** 1997. Periodic review (s, S) policies for joint replenishment inventory systems. *Management Science,* 43(10), pp 1447–1454. https://doi.org/10.1287/mnsc.43.10.1447

[16] **Ozkaya, B. Y., Gurler, U., & Berk, E.** 2006. The stochastic joint replenishment problem: A new policy, analysis, and insights. *Naval Research Logistics,* 53, pp 525–546. https://doi.org/10.1002/nav.20147

[17] **Gonçalves, J. N. C., Carvalho, M. S., & Cortez, P.** 2020. Operations research models and methods for safety stock determination: A review. *Operations Research Perspectives,* 7, 100164. https://doi.org/10.1016/j.orp.2020.100164

[18] **Eilon, S., & Elmaleh, J.** 1968. An evaluation of alternative inventory control policies. *International Journal of Production Research,* 7(1), pp 1–14. https://doi.org/10.1080/00207546808929792

[19] **Wemmerlöv, U., & Whybark, D. C. 1983.** Lot-sizing under uncertainty in a rolling schedule environment. *International Journal of Production Research,* 22(3), pp 467–484. https://doi.org/10.1080/00207548408942467

[20] **Callarman, T. E., & Hamrin, R. S.** 1984. A comparison of dynamic lot sizing rules for use in a single stage MRP system with demand uncertainty. *International Journal of Operations & Production Management,* 4(2), pp 39–48. https://doi.org/10.1108/eb054713
[21] **Gudum, C. K., & Kok, T. G. D.** 2002. *A safety stock adjustment procedure to enable target service levels in simulation of generic inventory systems* (Tech. Rep.). Copenhagen: Copenhagen Business School.
[22] **Ptak, C. A., & Smith, C.** 2011. *Orlicky's material requirements planning*, 3rd ed. New York, NY: McGraw Hill.
[23] **Da Cruz, J. A., De Salles-Netto, L. L., & Schenekemberg, C. M.** 2024. An integrated production planning and inventory management problem for a perishable product: Optimisation and Monte Carlo simulation as a tool for planning in scenarios with uncertain demands. *Transactions in Operations Research*, 32, pp 263-303. https://doi.org/10.1007/s11750-024-00667-x
[24] **Zhou, J., & Zhang, S.** 2024. A simulation-based Bayesian network approach to the joint decision of the international transportation mode and safety inventory policy. *Transportation Research Record*, 2678(10), pp 26-44, https://doi.org/10.1177/03611981241231798
[25] **Bouslah, B., Gharbi, A., & Pellerin, R.** 2014. Joint production and quality control of unreliable batch manufacturing systems with rectifying inspection. *International Journal of Production Research,* 52(14), pp 4103–4117. https://doi.org/10.1080/00207543.2012.746481
[26] **Gansterer, M., Almeder, C., & Hartl, R. F.** 2014. Simulation-based optimisation methods for setting production planning parameters. *International Journal of Production Economics,* 151, pp 206–213. https://doi.org/10.1016/j.ijpe.2013.10.016
[27] **Avci, M. G., & Selim, H.** 2017. A multi-objective, simulation-based optimisation framework for supply chains with premium freights. *Expert Systems with Applications*, 67, pp 95–106. https://doi.org/10.1016/j.eswa.2016.09.034
[28] **Avci, M. G., & Selim, H.** 2018. A multi-objective simulation-based optimisation approach for inventory replenishment problem with premium freights in convergent supply chains. *Omega,* 80, pp 153–165. https://doi.org/10.1016/j.omega.2017.08.016
[29] **Hosseini, M., Madathil, S. C., & Khasawneh, M. T.** 2025. Reinforcement learning-based simulation optimisation for an integrated manufacturing-warehouse system: A two-stage approach. *Expert Systems with Applications,* 290, 128259. https://doi.org/10.1016/j.eswa.2025.128259
[30] **Winter, T. S., & Van Vuuren, J. H.** 2024. Supporting material for the current paper titled "A practical method for joint replenishment under supplier constraints." URL: https://github.com/InventoryManager101/Supporting-Material.
[31] **Winston, W. L.** 2003. *Operations research: Applications and algorithms*, 4th ed. Belmont, CA: Thomson Brooks/Cole.
[32] **Lützeler, A.** 2023. alutzeler@finstock.co.za (Supply chain consultant)

APPENDIX A: CASE STUDY RESULTS DATA

The safety stock durations for each product of supplier A, as determined by each heuristic optimisation approach of §3.7 for the case study of §4, are presented in Table 5.

Table 5: The heuristically optimised safety stock duration (in days) for each product of Supplier A

Product	Unit cost	Safety stock duration	Sensitivity	Sensitivity/Unit cost	Random selection
A1	22.59958	22.59958	22.59958	22.59958	22.59958
A2	20.93797	20.93797	14.93797	15.93797	15.93797
A3	21.46769	21.46769	21.46769	21.46769	1.46769
A4	22.2618	22.2618	22.26180	22.26180	22.26180
A5	24.76964	21.76964	24.76964	24.76964	22.76964
A6	24.21737	22.21737	24.21737	24.21737	22.21737
A7	25.06924	22.06924	25.06924	25.06924	23.06924
A8	22.69226	22.69226	22.69226	22.69226	22.69226
A9	24.25586	22.25586	24.25586	24.25586	22.25586
A10	20.0128	22.0128	24.01280	24.01280	23.01280

The safety stock durations for each product of supplier B, as determined by each heuristic optimisation approach of §3.7 for the case study of §4, are presented in Table 6.

Table 6: The heuristically optimised safety stock duration (in days) for each product of Supplier B

Product	Unit cost	Safety stock duration	Sensitivity	Sensitivity/Unit cost	Random selection
B1	56.06042	38.06042	56.06042	54.06042	53.06042
B2	55.36302	38.36302	55.36302	55.36302	52.36302
B3	42.00224	39.00224	42.00224	42.00224	42.00224
B4	67.90894	38.90894	67.90894	61.90894	52.90894
B5	42.0	39.0	42.0	42.0	42.0
B6	42.0	39.0	42.0	21.0	42.0
B7	42.00446	39.00446	42.00446	42.00446	42.00446
B8	42.00104	39.00104	42.00104	42.00104	42.00104
B9	42.0005	39.0005	42.0005	42.0005	42.0005
B10	42.0	39.0	42.0	42.0	42.0005
B11	53.81816	38.81816	53.81816	26.81816	26.81816
B12	42.0	39.0	42.0	21.0	42.0
B13	42.0	39.0	42.0	42.0	42.0
B14	42.0	39.0	42.0	42.0	42.0
B15	51.15323	38.15323	51.15323	51.15323	51.15323
B16	42.00316	39.00316	42.00316	42.00316	42.00316
B17	42.00255	39.00255	42.00255	42.00255	42.00255
B18	42.0	39.0	42.0	42.0	42.0
B19	42.0	39.0	42.0	39.0	42.0
B20	42.0	39.0	42.0	42.0	42.0
B21	56.97281	38.97281	56.97281	46.97281	52.97281
B22	42.00031	39.00031	42.00031	42.00031	42.00031
B23	42.00082	39.00082	41.00082	24.00082	41.00082
B24	42.0	39.0	42.0	42.0	42.0
B25	52.77863	38.77863	50.77863	52.77863	51.77863
B26	54.29855	38.29855	26.29855	54.29855	35.29855
B27	55.58772	38.58772	55.58772	55.58772	52.58772
B28	24.95105	38.95105	27.95105	49.95105	28.95105
B29	52.07068	38.07068	25.07068	52.07068	49.07068
B30	20.00511	38.00511	42.00511	42.00511	42.00511
B31	62.57973	38.57973	62.57973	62.57973	52.57973
B32	24.43924	38.43924	50.43924	50.43924	24.43924
B33	25.0	39.0	42.0	42.0	42.0
B34	52.30093	38.30093	52.30093	52.30093	52.30093
B35	50.50995	38.50995	50.50995	50.50995	50.50995
B36	52.8961	38.8961	52.8961	52.8961	52.8961
B37	58.18753	38.18753	28.18753	28.18753	28.18753
B38	57.21404	38.21404	57.21404	57.21404	53.21404
B39	58.06119	38.06119	58.06119	58.06119	53.06119
B40	57.61293	38.61293	57.61293	28.61293	43.61293
B41	53.84096	38.84096	53.84096	53.84096	52.84096
B42	48.24211	38.24211	48.24211	48.24211	48.24211
B43	42.0	39.0	42.0	42.0	42.0
B44	42.00185	39.00185	42.00185	42.00185	42.00185
B45	20.00308	39.00308	42.00308	42.00308	42.00308
B46	42.00513	38.00513	42.00513	42.00513	42.00513
B47	58.189	38.189	58.189	58.189	53.189
B48	42.00488	39.00488	42.00488	42.00488	42.00488
B49	42.00534	38.00534	42.00534	42.00534	42.00534
B50	42.00076	39.00076	42.00076	42.00076	42.00076