

Haubitz, Christiane B.; Thonemann, Ulrich W.

Article — Published Version

How to Change a Running System—Controlling the Transition to Optimized Spare Parts Inventory Policies

Production and Operations Management

Provided in Cooperation with:

John Wiley & Sons

Suggested Citation: Haubitz, Christiane B.; Thonemann, Ulrich W. (2021) : How to Change a Running System—Controlling the Transition to Optimized Spare Parts Inventory Policies, Production and Operations Management, ISSN 1937-5956, Wiley, Hoboken, NJ, Vol. 30, Iss. 5, pp. 1386-1405,
<https://doi.org/10.1111/poms.13327>

This Version is available at:

<https://hdl.handle.net/10419/233696>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<http://creativecommons.org/licenses/by/4.0/>

How to Change a Running System—Controlling the Transition to Optimized Spare Parts Inventory Policies

Christiane B. Haubitz* , Ulrich W. Thonemann 

Department of Supply Chain Management and Management Science, University of Cologne, Cologne, 50923, Germany,
christiane.haubitz@uni-koeln.de, ulrich.thonemann@uni-koeln.de

Inventory optimization approaches typically optimize steady-state performance, but do not consider the transition of an initial state to the optimized state. In this study, we address this transition. Our research is motivated by a company that implemented an improved inventory policy for its spare parts division. The improved policy suggested new base stock levels for the majority of the parts. For parts with increased base stock levels, inventory increases were realized after the part lead times, but for low-demand parts with decreased base stock levels, inventory reductions were slow. As a result, inventory cost increased over the first months after the new inventory policy had been introduced and exceeded the inventory budget substantially. To avoid such undesirable effects, base stock level changes must be phased in. We consider a multi-item spare parts inventory system, initially operating under an item approach inventory policy that achieves identical fill rates for all parts. Our approach addresses the transition to a superior system approach inventory policy that maximizes the system fill rate. We model the inventory transition as a finite-horizon optimization problem and apply column generation and a marginal analysis heuristic to determine transient base stock levels for all parts. Using data from the company that motivated our research, we illustrate how the transition can be controlled to quickly improve fill rates without exceeding the initial inventory budget.

Key words: inventory; spare parts; system approach; inventory transition; column generation

History: Received: December 2019; Accepted: October 2020 by M. Eric Johnson, after 2 revisions.

*Corresponding author.

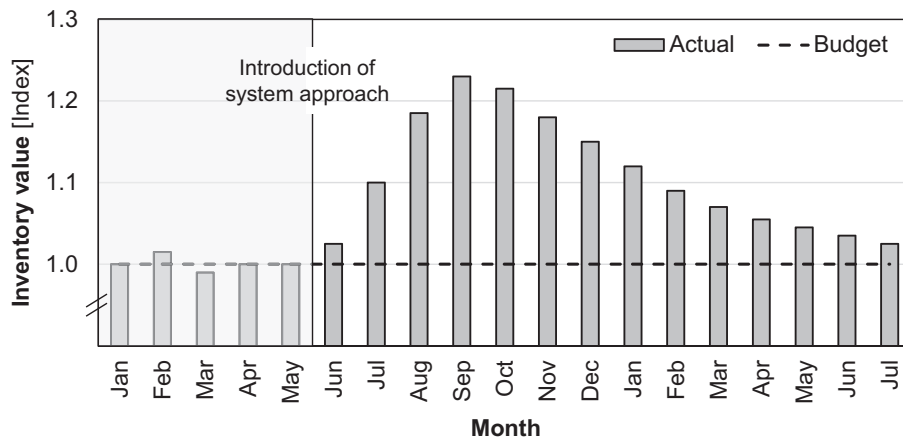
1. Introduction

The trigger of inventory optimization projects is often a suboptimal performance of the current inventory system. In spare parts management, for instance, companies can set base stock levels such that a certain fill rate is achieved by every individual part. This approach is referred to as *item approach*. Instead of optimizing each part individually, companies can reduce inventory or increase the system fill rate by considering all parts in the inventory system jointly when making decisions about base stock levels. This approach is referred to as *system approach* (Sherbrooke 2004). For a spare parts inventory system for high-end computer servers, for example, Thonemann et al. (2002) show improvements in inventory investment of up to 25% when applying a system approach instead of an item approach. For a spare parts inventory system at the Royal Netherlands Navy, Rustenburg et al. (2003) demonstrate an increase in spare parts availability by 34 percentage points, while simultaneously reducing the inventory investment by about 10%.

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

Implementing a system approach requires the adjustment of inventory control policy parameters. For example, base stock levels of inexpensive fast movers are increased and base stock levels of expensive slow movers are reduced. Overall inventory performance is improved once the inventory system has reached its new steady state. However, during the transition to the new steady state, system fill rate or system inventory holding cost can deteriorate temporarily, particularly if lead times are long or demand rates are low.

We experienced this in an inventory optimization project with the service division of a global business-to-business (B2B) equipment manufacturer. The company generates annual multi-billion euros turnover and operates in more than 50 countries. Its service division offers repair and maintenance services for the specialized and expensive equipment. The division wanted to improve its inventory performance but the inventory budget was restricted and we were asked to improve the system fill rate without exceeding the current budget. We suggested moving the inventory system from an item approach to a system approach and projected a long-run system fill rate increase of 12 percentage points while keeping the inventory holding cost constant. After pilot implementation of the new approach, the system fill rate

Figure 1 Inventory Index Development at the Service Division after Introducing the System Approach Base Stock Levels in May

started to improve, but the inventory holding cost increased by more than 20% within four months (Figure 1). This placed a major burden on the service organization's inventory budget and it took more than a year for the inventory holding cost to finally reach the projected value.

The observed increase in inventory cost is a severe issue, since complying with inventory budgets is of critical importance to companies. Companies establish inventory budgets because the capital they can invest in inventory is restricted (Silver et al. 2016, Yang et al. 2017). Requiring significantly more capital, even if only temporarily, can cause severe financial straits and even impacts the company's valuation. Moreover, the negative consequences of increased inventory cost during the transition period are not limited to the direct financial impact. Major changes in inventory policies like moving from an item approach to a system approach are usually initiated in the scope of large and costly inventory projects. If the inventory performance falls short of its targets and expectations, the pressure on accountable divisions can be substantial.

These observations motivated our research on controlling the transition of inventory systems when the inventory policy changes. We analyze the transient behavior of a spare parts inventory system that operates under a periodic review base stock policy. The goal of our approach is to phase-in changes in base stock levels over time while keeping inventory holding cost within a given budget during the transition.

We focus on analyzing the base stock level transition of an inventory system that moves from an item approach to a system approach, taking into account replenishment lead times and expected future demands. However, our model is not restricted to managing transitions of inventory systems starting from an item approach. It can be generally applied to arbitrary inventory systems that operate under

periodic review base stock policies and that require an adaption of the base stock levels to a system approach. Such adaptations might be triggered, for example, if lead times or prices are renegotiated with suppliers, new suppliers are introduced, new customers are acquired, or major customers terminate their service contracts (Çetinkaya and Parlar 2010).

We contribute to the inventory control literature by considering the transition of a multi-item inventory system from a current state to an optimized state, which is a topic that has not previously been addressed. We formulate the problem as a multi-period optimization problem and present two solution approaches that rely on column generation and on marginal analysis, respectively. For small and moderate size problems, for which we can determine upper bounds on the objective function values, both approaches generate solutions that are close to optimal. Large size problems can only be solved by marginal analysis and we use the approach in an extended case study that is based on data from a global equipment manufacturer.

This study is structured as follows. In section 2, we review the literature. In section 3, we present a model for optimizing the inventory system transition. In section 4, we introduce two solution approaches. In section 5, we present numerical results based on data from the service division of the global manufacturer that motivated our research. In section 6, we summarize our findings and conclude.

2. Literature Review

We review steady-state system approaches with modeling assumptions similar to our work, and non-steady-state inventory models for base stock level transitions. We also review two relevant solution techniques for system approach optimization problems. For a comprehensive overview of the literature on spare parts inventory systems, we refer to Sherbrooke

(2004), Muckstadt (2005), Basten and Van Houtum (2014), and Van Houtum and Kranenburg (2015).

System approaches for spare parts inventory systems have gained increasing attention since their introduction in the seminal paper by Sherbrooke (1968). Their benefits compared to simpler item approaches have been shown by Mitchell (1988), Thonemann et al. (2002), Rustenburg et al. (2003), and Sherbrooke (2004). Spare parts inventory systems that operate under a periodic review base stock policy have been analyzed by Mitchell (1988), Cohen et al. (1989), and Cohen et al. (1992). Budget-constrained, service maximizing system approaches have been analyzed by Rustenburg et al. (2000, 2003) and Sherbrooke (2004). Schwarz et al. (1985) maximize the system fill rate subject to a constraint on the system safety stock.

While most system approach models consider the initial supply of spare parts, there are only a few models that include the current state of a running inventory system. For such a system, Rustenburg et al. (2000) address a budget-constrained re-supply problem. They assume annual budgets for the total purchasing cost of new parts, which arrive after a fixed lead time. They develop operational re-supply strategies for spending the limited budget, but do not cover situations in which the optimal target base stock levels change.

Van Houtum and Kranenburg (2015 ch. 2.8) introduce a modified system approach for periodic base stock level optimization. They account for the current inventory position of each spare part by modifying the cost function of the standard initial supply model when updating the base stock levels. While their approach generates good solutions for the next planning period, it is myopic and does not aim to reach the standard system approach base stock levels. This is reasonable for inventory systems that already operate under a system approach, when changes in base stock levels are typically minor. Their approach also assumes steady-state behavior of the system in each period with quickly realized changes in the base stock levels. The authors note that it is necessary to consider the transient behavior of the inventory system if their assumptions do not hold, for example, if parts have long replenishment lead times. We address such situations with our research.

Several publications address the non-steady-state transition period when steady-state base stock levels change for single items. When base stock levels increase, the implication in a single-item setting is straightforward: new parts must be ordered and arrive after the replenishment lead time. When base stock levels decrease, inventory above the reduced base stock level must be considered. Pinçe and Dekker (2011) note that in such situations “timely adaptation of the base stock levels is crucial for optimal stock

control” (p. 83). They adapt the steady-state base stock policy for a single, low demand item with a fixed lead time and decreasing demand and propose a transition control policy that minimizes the total expected cost during the transition. Similar to our model, their policy is based on the current inventory position of the item, but it only allows for two different base stock levels and focuses on the timing of the switch. Our model allows for multiple adjustments in base stock levels during the transition period, which is hypothesized by Pinçe and Dekker (2011) to improve results. Teunter and Klein Haneveld (2002) develop an ordering policy for the end-of-life phase of service parts with stationary, Poisson-distributed demand. Assuming a higher replenishment price and risk of obsolescence in the final phase, they propose a sequence of decreasing base stock levels for a deterministic planning horizon of a single item. Çetinkaya and Parlar (2010) consider the introduction of a periodic review base stock policy without an explicit spare parts focus. In their single-item setting, the target inventory policy is similar to ours. However, the authors do not consider inventory systems with multiple items and do not include replenishment lead times. All of the reviewed non-steady-state models focus on the transition of a single item and are not directly applicable to the multi-item problem that we consider.

We apply two solution approaches to solve our optimization problem. The first approach, decomposition and column generation, exploits structural properties of multi-item spare parts systems. Instead of solving the original complex optimization problem, the problem is decomposed by spare part and the resulting single-item optimization problems are solved repeatedly. This method has recently received increasing attention in the multi-item spare parts inventory optimization literature (Alvarez et al. 2013, 2015, Arts 2017, Drent and Arts 2020, Kranenburg and Van Houtum 2007, Topan et al. 2010, Topan et al. 2017, Wong et al. 2007). Details on fundamentals and theoretical background of decomposition and column generation are provided by Dantzig and Wolfe (1960) and Desrosiers and Lübbecke (2005).

Our second solution approach, marginal analysis, is a well-known solution technique for optimizing spare parts inventory systems that operate under a system approach (Sherbrooke 2004, Van Houtum and Kranenburg 2015). The general principle is to compare the cost of increasing a spare part’s base stock level by one unit with its impact on the system’s performance (e.g., the expected backorders or the system fill rate). In each iteration, the base stock level of the part with the highest performance improvement per monetary unit spent is increased. The procedure is repeated successively until the system target is reached. Marginal analysis is efficient, provides good results, and

is easy to implement in practice. It has been applied to various spare parts inventory optimization problems, for example, by Wong et al. (2005, 2007) and Topan et al. (2017). For a finite-horizon periodic review setting, Caggiano et al. (2006) apply heuristics that are based on marginal analysis to make operational repair and inventory allocation decisions. We use this solution approach to determine target base stock levels for multiple items and multiple periods based on the expected system state.

3. Problem Formulation and Mathematical Model

We consider a multi-item spare parts inventory system that operates under a periodic review base stock policy (e.g., Cohen et al. 1992, Sherbrooke 2004). This inventory control policy is suited when spare parts are ordered periodically and is applied in many real-world spare parts inventory systems (Cavalieri et al. 2008, Tiemessen et al. 2013, Wang 2012). It is also applied at the company that motivated our research. We denote the set of spare parts i in the inventory system by I . The part-specific replenishment lead time is $l_i \in \mathbb{N}$ and the unit holding cost is $c_i \in \mathbb{R}^+$. Demand per period, D_i , is Poisson-distributed with demand rate $\lambda_i \in \mathbb{R}^+$ and the aggregated mean demand per period is $\Lambda = \sum_{i \in I} \lambda_i$. We denote the part demand over p consecutive periods by D_i^p . Excess demand is backlogged. We use the fill rate as service level measure, that is, the fraction of demand that is fulfilled from on-hand inventory (Axsäter 2015, Cachon and Terwiesch 2006).

We will next discuss the model of the steady-state inventory system and define the initial state of the inventory system under the item approach, as well as the target state of the inventory system under the system approach. Then, we will build on the steady-state model to optimize the inventory transition from the item approach to the system approach.

3.1. The Inventory System in Steady State

We denote the base stock level of part $i \in I$ by s_i . $s = (s_1, s_2, \dots, s_{|I|})$ is the vector of base stock levels for the parts in the system. The sequence of events in a period is as follows: First, the current inventory positions of the parts are observed. Then, orders are placed to increase the inventory positions to the base stock levels s_i . Next, the orders that were placed l_i periods ago are received. Demands D_i are satisfied and excess demands are backlogged. Inventory holding costs are determined based on the on-hand inventories at the end of a period.

The fill rate of part i under a periodic review base stock policy with base stock level s_i is (Cachon and Terwiesch 2006)

$$\beta_i(s_i) = 1 - \frac{\sum_{d=s_i+1}^{\infty} (d - s_i) \cdot (P(D_i^{l_i+1} = d) - P(D_i^{l_i} = d))}{E[D_i]} \quad (1)$$

The demand-weighted system fill rate is

$$\beta(s) = \frac{\sum_{i \in I} \lambda_i \cdot \beta_i(s_i)}{\Lambda}$$

The expected on-hand inventory of part i at the end of a period is

$$oh_i(s_i) = \sum_{d=0}^{s_i} P(D_i^{l_i+1} = d) \cdot (s_i - d) \quad (2)$$

and the expected system inventory holding cost is

$$c(s) = \sum_{i \in I} c_i \cdot oh_i(s_i).$$

3.1.1 The Inventory System under the Item Approach. Consider an inventory system that is operated under the item approach. The base stock levels $s^{item} = (s_1^{item}, \dots, s_{|I|}^{item})$ are chosen such that all parts achieve the same fill rate (Thonemann et al. 2002). We denote the resulting expected system inventory holding cost by $c(s^{item})$ and the system fill rate by $\hat{\beta} = \beta(s^{item})$.

3.1.2 The Inventory System under the System Approach. The base stock levels that maximize the system fill rate $\beta(s)$, given an expected inventory cost budget b , can be determined by solving

$$\begin{aligned} (P^S) \quad & \max \quad \beta(s) \\ & s.t. \quad c(s) \leq b \\ & \quad \beta_i(s_i) \geq \beta_{min} \quad \forall i \in I \\ & \quad s_i \in \mathbb{N}_0 \quad \forall i \in I. \end{aligned}$$

The first constraint ensures that the inventory budget is not exceeded. The second constraint imposes a minimum fill rate $\beta_{min} \in [0, 1)$ on each part. Companies impose minimum fill rates on parts to avoid that some expensive parts with low demand rates receive base stock levels of zero. This can be the optimal solution, but it can also be undesirable as it implies that failures of critical parts result in system downtimes that are equal to the part lead time. To avoid excessively long downtimes and to avoid that some parts are generally unavailable, we include the minimum fill rate. With $\beta_{min} = 0$, the problem corresponds to the system approach without minimum fill rates.

We solve problem (P^S) by marginal analysis (Sherbrooke 2004, Van Houtum and Kranenburg

2015). The obtained solution is efficient if condition $s_i \geq \frac{\lambda_i}{\ln(l_i+1) - \ln(l_i)} - 1$ holds for all $i \in I$ (proof contained in Appendix A). This condition holds in all our numerical experiments; however, in situations where this is not the case, there might be solutions that achieve a higher system fill rate at lower cost. Note that an efficient solution is not necessarily optimal for optimization problems with discrete decision variables. If the final increase in the base stock level determined by the algorithm does not result in a fully utilized budget, a solution that increases the base stock level of a part with a smaller marginal increase $\beta(s)/\Delta c(s)$ can be superior. However, in situations like the one we consider, with a large number of spare parts and unit holding costs that are very small compared to the budget, the suboptimality is negligible (Sherbrooke 2004).

3.2 The Inventory Transition

The transition from the item approach to the system approach is realized during a finite planning horizon T with periods indexed by $t=0, \dots, T$, $T > l^{max}$ with $l^{max} = \max_{i \in I} l_i$. The base stock levels in period $t=0$ are those of the item approach. After period $t=T$, the transition is completed and the base stock levels are those of the system approach. The objective of the inventory transition is to maximize the average system fill rate during the transition without exceeding the inventory budget b in any period. The budget b is determined by the expected system inventory holding cost under the item approach, $b = c(s^{item})$. While any $T > l^{max}$ is feasible for the transition model, the planning horizon T should be chosen long enough for the on-hand inventories oh_i , $i \in I$, to reach the system approach steady-state levels. Otherwise, the budget might be exceeded in period $T+1$. We further discuss the appropriate length of the planning horizon in section 5.2.3.

We model the transition as a sequence of base stock levels \bar{s}_i , $i \in I$, during the transition periods $0, \dots, T$. For notational convenience, \bar{s}_i also includes base stock levels for the l_i periods before $t=0$. The sequence $\bar{s}_i = (s_i^{-l_i}, s_i^{-l_i+1}, \dots, s_i^0, \dots, s_i^T)$ has the following properties: Up to $t=0$, the base stock levels are those of the item approach: $s_i^t = s_i^{item} \forall t \leq 0$. After period $T - l_i$, the base stock levels s_i^t are set to the system approach base stock levels s_i^{system} : $s_i^t = s_i^{system} \forall t > T - l_i$. We assume \bar{s}_i to be either monotone increasing ($s_i^{item} < s_i^{system}$), monotone decreasing ($s_i^{item} > s_i^{system}$) or constant ($s_i^{item} = s_i^{system}$).

The base stock levels s_i^t , $t \leq 0 \cup t > T - l_i$ are input parameters and $s_i^1, s_i^2, \dots, s_i^{T-l_i}$ are decision variables of the transition model. $\beta_i^t(\bar{s}_i)$ and $oh_i^t(\bar{s}_i)$ denote

the transient fill rate and the transient on-hand inventory of part i in period t , respectively. Due to the non-stationarity of the base stock levels during the transition, we cannot directly apply Equations (1) and (2) for computing the transient fill rates and the transient on-hand inventories. Depending on the direction of the base stock level transition for different parts, we distinguish between three cases.

If $s_i^{item} = s_i^{system}$, then there is no change in the base stock level of part i and we can directly apply Equations (1) and (2) for the transient fill rate and the transient on-hand inventory in period t , $\beta_i^t(\bar{s}_i) = \beta_i(s_i^{system})$ and $oh_i^t(\bar{s}_i) = oh_i(s_i^{system})$.

If $s_i^{item} < s_i^{system}$, then the base stock levels of part i are monotone increasing and the transient fill rate and corresponding on-hand inventory in period t only depend on the transient base stock level of period $t - l_i$. Therefore, we can again use Equations (1) and (2), $\beta_i^t(\bar{s}_i) = \beta_i(s_i^{t-l_i})$ and $oh_i^t(\bar{s}_i) = oh_i(s_i^{t-l_i})$.

We can apply the steady-state equations for increasing and constant base stock levels because of the characteristics of the base stock policy. Under the steady-state base stock policy, the on-hand inventory depends on the base stock level s_i and the demand over lead time + 1 periods. An order of part i arriving in period t was initiated in period $t - l_i$. The order size in period $t - l_i$ is the difference between the base stock level $s_i^{t-l_i}$ and the inventory position in period $t - l_i$. All orders that are part of the inventory position in period $t - l_i$ will arrive until period t and whatever happened before period $t - l_i$ will be compensated with the order initiated in period $t - l_i$. In period $t - l_i$, the inventory position of part i is either below $s_i^{t-l_i}$ and we order the difference to $s_i^{t-l_i}$, or it is equal to $s_i^{t-l_i}$ and we order nothing. We cannot observe an inventory position greater than $s_i^{t-l_i}$ as $s_i^{t'} \leq s_i^t \forall t' < t$. Thus, the transient on-hand inventory in period t only depends on the base stock level of period $t - l_i$. The transient on-hand inventory in period t does further not depend on any decisions made in periods after $t - l_i$.

If $s_i^{item} > s_i^{system}$, then the base stock levels of part i are monotone decreasing during the transition. As the on-hand inventory of a part only decreases with demand, we lose the characteristic that the transient fill rate and the transient on-hand inventory in period t only depend on the base stock level in period $t - l_i$. For example, let us consider a situation with a positive difference between the base stock levels of two periods t' and t'' with $t' < t''$, $s_i^{t'} - s_i^{t''} > 0$. If the total demand during the time span between periods

t' and t'' , $D_i^{t''-t'}$, is smaller than the difference of the base stock levels, the inventory position of part i in period t'' is greater than $s_i^{t''}$, which then also influences the on-hand inventory of later periods. The probability that this situation occurs increases with lower demand rates. Therefore, considering this case is particularly important in the spare parts context as demand rates are often low.

The transient fill rate and the transient on-hand inventory at the end of period t depend on the inventory position in period $t - l_i$ immediately after an order is placed. In the standard base stock policy, this is the base stock level s_i but for parts with decreasing base stock levels, the inventory position $ip_i^{t-l_i}$ can be between s_i^{item} and $s_i^{t-l_i}$. For each potential inventory position $ip_i^{t-l_i} = k, k \in \{s_i^{t-l_i}, \dots, s_i^{item}\}$, we can compute the resulting fill rate $\beta_i(k)$ and the expected on-hand inventory $oh_i(k)$ using Equations (1) and (2). The overall transient fill rate $\beta_i^t(\bar{s}_i)$ and the transient on-hand inventory $oh_i^t(\bar{s}_i)$ in period t are the averages of those values, weighted with the probabilities $P(ip_i^t = k)$. We, therefore, must derive those probabilities.

At the beginning of the transition, in period 0, the system is in steady-state under the base stock policy with base stock levels s_i^{item} and, thus, $P(ip_i^0 = k) = 1$ for $k = s_i^{item}$ and $P(ip_i^0 = k) = 0$ for $k \neq s_i^{item}$. For every period $t > 0$, we describe the probabilities recursively with

$$P(ip_i^t = k) = \begin{cases} \sum_{j=k}^{s_i^{item}} P(ip_i^{t-1} = j) \cdot P(D_i = j - k) & k > s_i^t \\ \sum_{j=k}^{s_i^{item}} P(ip_i^{t-1} = j) \cdot P(D_i \geq j - k) & k = s_i^t \\ 0 & k < s_i^t. \end{cases}$$

With these probabilities, the transient fill rate and the transient on-hand inventory in period t for parts with decreasing base stock level are defined by

$$\beta_i^t(\bar{s}_i) = \sum_{j=s_i^{t-l_i}}^{s_i^{item}} P(ip_i^{t-l_i} = j) \cdot \beta_i(j)$$

and

$$oh_i^t(\bar{s}_i) = \sum_{j=s_i^{t-l_i}}^{s_i^{item}} P(ip_i^{t-l_i} = j) \cdot oh_i(j).$$

Thus, we can compute the transient fill rates and the transient on-hand inventories for every part and every period of the inventory transition by

$$\beta_i^t(\bar{s}_i) = \begin{cases} \beta_i(s_i^{t-l_i}) & s_i^{item} < s_i^{system} \\ \beta_i(s_i^{system}) & s_i^{item} = s_i^{system} \\ \sum_{j=s_i^{t-l_i}}^{s_i^{item}} P(ip_i^{t-l_i} = j) \cdot \beta_i(j) & s_i^{item} > s_i^{system} \end{cases}$$

and

$$oh_i^t(\bar{s}_i) = \begin{cases} oh_i(s_i^{t-l_i}) & s_i^{item} < s_i^{system} \\ oh_i(s_i^{system}) & s_i^{item} = s_i^{system} \\ \sum_{j=s_i^{t-l_i}}^{s_i^{item}} P(ip_i^{t-l_i} = j) \cdot oh_i(j) & s_i^{item} > s_i^{system}. \end{cases}$$

Based on these results, we formulate the inventory transition as a finite-horizon optimization problem:

$$(P^T) \quad \max \frac{1}{T} \cdot \sum_{t=1}^T \frac{\sum_{i \in I} \lambda_i \cdot \beta_i^t(\bar{s}_i)}{\Lambda} \quad (3)$$

$$s.t. \sum_{i \in I} c_i \cdot oh_i^t(\bar{s}_i) \leq b \quad \forall t = 1, \dots, T \quad (4)$$

$$s_i^t = s_i^{item} \quad \forall i \in I, \forall t \leq 0 \quad (5)$$

$$s_i^t = s_i^{system} \quad \forall i \in I, \forall t > T - l_i \quad (6)$$

$$\bar{s}_i \text{ is monotone} \quad \forall i \in I \quad (7)$$

$$s_i^t \in \mathbb{N}_0 \quad \forall i \in I, t = -l_i, \dots, T. \quad (8)$$

(P^T) builds on the optimization problem (P^S) for the steady-state system approach. The objective is to maximize the average system fill rate during the transition. The Budget Constraints (4) ensure that the expected system inventory holding cost does not exceed the budget in any period of the transition. Constraints (5)–(8) formalize the required properties on the base stock level sequences \bar{s}_i . Similar to the stationary model, the model of the inventory transition is a non-linear integer programming problem. The solution of the model is an optimal sequence of base stock levels for every part over the planning horizon T .

4. Solution Approaches

We provide two solution approaches for problem (P^T) . The first solution approach is based on decomposition and column generation. It generates upper bounds for problem (P^T) and can also be used to find a well-performing, feasible solution. The second solution approach is a heuristic that is based on marginal analysis.

4.1. Column Generation Approach

The optimization problem (P^T) is non-linear and on-hand inventories depend on base stock levels of previous periods. To efficiently solve this problem, we use decomposition and column generation (Dantzig and Wolfe 1960, Desrosiers and Lübbecke 2005). Column generation is an efficient technique to solve large-scale linear problems. This technique alternates between solving a master problem with a restricted set of decision variables and generating new decision variables (or as the name implies columns) that might improve the solution.

We first reformulate the problem (P^T) as an integer program, for which we later relax the integrality constraints. We refer to it as the master problem (MP). The set \bar{S}_i contains all feasible base stock level sequences \bar{s}_i for part i . For each element $n \in \bar{S}_i$, we introduce a decision variable $x_i^n \in \{0, 1\}$ that indicates whether a sequence is selected ($x_i^n = 1$) or not ($x_i^n = 0$). We add a constraint to ensure that exactly one policy is selected for each part i . The transient fill rate $\beta_i^{t,n}$ and the transient on-hand inventory $oh_i^{t,n}$ of the corresponding sequence n in period t are input parameters to the master problem. We obtain:

$$(MP) \quad \max \frac{1}{T} \cdot \sum_{t=1}^T \frac{\sum_{i \in I} \sum_{n \in \bar{S}_i} \lambda_i \cdot \beta_i^{t,n} \cdot x_i^n}{\Lambda} \quad (9)$$

$$s.t. \quad \sum_{i \in I} \sum_{n \in \bar{S}_i} c_i \cdot oh_i^{t,n} \cdot x_i^n \leq b \quad \forall t = 1, \dots, T \quad (10)$$

$$\sum_{n \in \bar{S}_i} x_i^n = 1 \quad \forall i \in I \quad (11)$$

$$x_i^n \in \{0, 1\} \quad \forall i \in I, n \in \bar{S}_i. \quad (12)$$

For every part, the number of potential sequences $|\bar{S}_i|$ and, therefore, the number of decision variables is finite. However, the number quickly increases with planning horizon T and absolute difference between initial and target base stock levels, $|s_i^{system} - s_i^{item}|$. $|\bar{S}_i|$ can be computed as

$$|\bar{S}_i| = f(s_i^{item}, s_i^{system}, T, l_i) = \begin{pmatrix} |s_i^{system} - s_i^{item}| + T - l_i \\ T - l_i \end{pmatrix}. \quad (13)$$

The number of constraints of (MP) increases linearly in the number of parts and the length of the planning horizon.

In principle, it would be possible to determine \bar{S}_i and all corresponding fill rates $\beta_i^{t,n}$ and on-hand inventories $oh_i^{t,n}$ upfront, and then solve (MP) optimally. However, $|\bar{S}_i|$ can be large and the problem becomes numerically intractable even for moderate problem sizes. This specifically holds for cases with

large differences between the item approach base stock levels and the system approach base stock levels, that is, cases for which we can potentially gain the biggest impact by optimizing the transition.

To overcome this computational challenge, we relax the integrality constraints on x_i^n and apply column generation to a restricted master problem (RMP) that contains only a subset of possible sequences $\bar{S}'_i \subseteq \bar{S}_i$ (see Appendix B). The solution of the final (RMP) after column generation provides an upper bound to (MP) and thus to the original problem (P^T). If in this solution all x_i^n are integer, this upper bound is the optimal solution to (P^T) and the base stock level sequences can be applied directly. Otherwise, we must still find a well-performing feasible integer solution to (MP).

We achieve this by solving the final (RMP) after column generation as an integer program, thus with enforced Integrality Constraints (12), and denote it by CG (column generation) approach. Such an approach has recently been applied by Alvarez et al. (2013, 2015), Arts (2017), and Drent and Arts (2020). Alvarez et al. (2013, 2015) showed that it performs better than using the fractional solution of (RMP) as a starting point for local search as done by Kranenburg and Van Houtum (2007, 2008).

4.2. Marginal Analysis Approach

Our second solution approach applies marginal analysis to all periods of the planning horizon. In every period, it considers the projected state of the inventory system, based on decisions of previous periods, different lead times of the parts, projected part demand, and restrictions on the base stock level sequences. We refer to this approach as marginal analysis approach (MA approach).

We perform a two-stage algorithm on the planning horizon T to determine all transient base stock levels s_i^t , $t = 1, \dots, T - l_i$. With the vector $\bar{s} = (\bar{s}_1, \bar{s}_2, \dots, \bar{s}_{|I|})$ that contains the base stock level sequences \bar{s}_i of all parts i , the system fill rate in period t is $\beta^t(\bar{s}) = \frac{\sum_{i \in I} \lambda_i \cdot \beta_i^t(\bar{s}_i)}{\Lambda}$ and the expected system inventory holding cost in t is $c^t(\bar{s}) = \sum_{i \in I} c_i \cdot oh_i^t(\bar{s}_i)$. To initialize the algorithm, we set $s_i^t = s_i^{item} \forall t \leq 0$ and $s_i^t = \min\{s_i^{item}, s_i^{system}\} \forall t > 0$ and compute the corresponding values for $\beta^t(\bar{s})$ and $c^t(\bar{s})$, $t = 1, \dots, T$.

The **first stage** of our algorithm only adapts the transient base stock levels for parts with increasing base stock levels. At the beginning of each period $t = 1, \dots, T$ we determine which parts are applicable for a base stock level increase. Because of the transient behavior of the system, $\beta^t(\bar{s})$ and $c^t(\bar{s})$ are not affected by the base stock levels of the current period t but by the base stock levels of period $t - l_i$. We therefore check for which parts the base stock levels $s_i^{t-l_i}$ can be

increased without violating the requirements for the base stock level sequence ($t > l_i$ and $s_i^{t-l_i} < s_i^{system}$). Then we apply marginal analysis to this subset of parts, starting with the current base stock levels $s_i^{t-l_i}$. We incrementally increase the base stock level $s_i^{t-l_i}$ with the highest relative improvement in system fill rate compared to the impact on expected system inventory holding cost. A part is removed from the marginal analysis whenever an increase by one unit would violate the budget in period t or the requirements for the base stock level sequence. When the base stock level of no part can be further increased, the marginal analysis finishes. For all parts with increasing base stock levels, the base stock levels of periods greater than $t - l_i$ are set to the value of $s_i^{t-l_i}$ and $\beta^\tau(\bar{s})$ and $c^\tau(\bar{s})$ of later periods $\tau > t$ are updated. In the last period T , all base stock levels $s_i^{\tau-l_i}$, $\tau > T - l_i$ are set to the system approach base stock levels.

The **second stage** only adapts the transient base stock levels for parts with *decreasing* base stock levels. For this subset of parts, the second stage follows the process of the first stage with two adaptations. First, we do not only check the impact of a base stock level increase on the budget of period t , but also on the budget of all succeeding periods $\tau = t, \dots, T$. This is necessary as for parts with decreasing base stock levels, the expected inventory holding cost in period τ depends on the decision in period t (and not only on the decision of period $\tau - l_i$). Second, we do not change any base stock levels of periods greater than $t - l_i$ after the marginal analysis in period t .

The detailed marginal analysis approach is provided in pseudocode in the Online Appendix.

We split the approach into two stages to mitigate the negative effects of its myopicity. The algorithm optimizes the base stock levels of the different periods successively. As a result, the transition to new base stock levels could be prolonged unnecessarily without the two stages. For parts with decreasing base stock levels, the myopicity of the approach could keep the base stock levels s_i^t high for too many periods, delaying the increase of parts with $s_i^{item} < s_i^{system}$. The base stock levels could stay high because the effects on the expected system inventory holding cost are small enough to still fulfill the budget constraint in the *current* period t . However, this decision also leads to budget consumption in *succeeding* periods. This can prevent the buildup of budget gaps that would allow for the base stock level increase of expensive parts with $s_i^{item} < s_i^{system}$.

The two-stage approach addresses this issue and ensures that base stock levels converge quickly to the target base stock levels, since it prioritizes the increase in base stock levels for parts with $s_i^{item} < s_i^{system}$ over the

delayed decrease in base stock levels for parts with $s_i^{item} > s_i^{system}$. The second stage ensures that base stock levels for parts with $s_i^{item} > s_i^{system}$ are only kept high if there is spare budget for the base stock level above the system base stock level in *all* later periods.

To guarantee the adherence to the budget constraint during the transition, only the first stage of the MA approach is crucial. Depending on the characteristics of the inventory system, the second stage can additionally increase the transient system fill rate during the transition, for example, if parts with increasing base stock levels have long lead times and parts with decreasing base stock levels have short lead times and high demands. For most inventory systems in the numerical study, however, the impact of the second stage is limited and most benefits are realized by applying the first stage.

The objective of the MA approach differs from the objective of the original problem (P^T), since the MA approach is myopic and optimizes the system fill rates in the periods $t = 1, \dots, T$ sequentially instead of considering the average system fill rate over the planning horizon holistically. However, the constraints are identical. Therefore, every solution obtained by the MA approach is feasible and can be used as the starting solution of the CG approach of section 4.1. We show the effectiveness of our MA approach in the numerical study. Benchmarking it to the upper bound obtained from column generation, we observe average gaps below 0.1% and a maximum gap below 1.2% across all tested instances. This close-to-optimal performance together with the short runtimes makes the MA approach particularly valuable for controlling inventory transitions of large systems.

5. Numerical Study

The numerical study presented in this section consists of two parts. First, we compare the performance of the CG approach and the MA approach for small to medium-sized inventory systems with medium-length planning horizons (section 5.1). Second, we analyze the value of optimizing the inventory transition in a case study based on company data (section 5.2). After applying the MA approach to a real-world inventory system in section 5.2.1, we analyze the impact of changes in system characteristics in an extensive sensitivity analysis (section 5.2.2) and discuss managerial implications (section 5.2.3). We implemented all approaches in C++ and used Gurobi to solve the linear and integer programs. We conducted the computations on a Windows 10 64-bit system with 16 GB memory and two Intel Xeon 2.30 GHz processors.

Table 1 Parameter Values for the Performance Evaluation

Parameter	Values
(a) Number of parts in I	20, 35, 50
(b) System fill rate $\hat{\rho}$	75%, 85%, 95%
(c) Demand rates $\lambda_i \sim \text{lognorm}(\mu_d, \sigma_d^2)$	$(\mu_d, \sigma_d^2) \in \{(0.28, 0.87), (0.55, 1.75), (0.83, 2.63)\}$
(d) Holding costs $c_i \sim \text{lognorm}(\mu_c, \sigma_c^2)$	$(\mu_c, \sigma_c^2) \in \{(1, 3), (1, 6), (1, 9)\}$
(e) Lead times $l_j \sim u(\hat{a}, \hat{b})$	$\hat{a} \in \{1\}; \hat{b} \in \{1, 2, 3, 4\}$
(f) Planning horizon T	10, 12, 14

5.1. Performance Evaluation of the Solution Approaches

We benchmark the performance of the CG approach and the MA approach against the upper bound for small to medium-sized inventory systems. The inventory systems are characterized by the number of parts, the initial system fill rate under the item approach, and the distributions of demand rates, unit holding costs, and lead times. Moreover, we vary the planning horizon of the transition. To obtain robust results, we randomly generate inventory systems with different parameter values. To get realistic scenarios, we base the parameter ranges on data from the company that motivated our research. Table 1 contains the parameter values that we cover in a full factorial analysis. We test the solution approaches for inventory systems with 20, 35, and 50 parts. Demand rates λ_i are drawn from a lognormal distribution. This distribution type best describes the spare parts data from our industry partner, which reveals a highly skewed distribution of demand rates: most of the parts in the inventory system have very low demand rates, but a few parts have high rates. Our observations are in line with the literature, which observed that demand rates in spare parts systems with many independent parts follow lognormal distributions (Wright 1991, 1992). We also draw the spare parts’ unit holding costs from a lognormal distribution. This distribution describes the situation of many spare parts inventory systems very well: only a small fraction of parts contributes to the majority of the system inventory holding cost. By fitting the lognormal distribution to the spare parts set, we obtain the distribution parameters $\mu_d = 0.55$ and $\sigma_d^2 = 1.75$ for the demand rates. The holding costs are best described by $\sigma_c^2 = 6$ and we normalize μ_c to 1 without loss of generality. To show the robustness of the computational results for inventory systems that run under different conditions, we vary μ_d and σ_d^2 by $\pm 50\%$, keeping the σ^2 -to- μ -ratio constant. We also vary σ_c^2 by $\pm 50\%$. The demand and holding cost distributions are truncated at $F(\cdot) = 0.99$ to avoid extreme outliers and provide realistic inventory system representations.

We analyze four lead time settings. In the first setting, all parts have a lead time of 1 period. In the other settings, lead times are drawn from a discrete uniform

distribution between 1 and an upper bound of 2, 3, and 4. We solve all instances with a minimum part fill rate of 50% and initial system fill rates of 75%, 85%, and 95%. The planning horizons are set to 10, 12, and 14 periods, which are reasonable time spans for the considered lead times. The combination of parameter values results in $3^5 \times 4 = 972$ settings. We randomly generate 20 instances per setting, leading to a total of 19440 instances. For all instances, we determine item approach and system approach base stock levels and then optimize the inventory transition with both solution approaches.

We assess the performance of the two approaches by evaluating the solutions and the runtimes for each instance. The solution quality is measured by comparing the objective value of the respective solution approach (z_{CG}^* and z_{MA}^*) to the upper bound (z_{UB}^*) that is provided by applying column generation to the restricted master problem (RMP). We compute the relative gap to the upper bound with

$$Gap_{UB-x} = \frac{z_{UB}^* - z_x^*}{z_{UB}^*}, x \in \{CG, MA\}.$$

Table 2 summarizes the computational results for the 19440 tested instances. We report averages and maximum values for the gaps to the upper bound and runtimes with respect to the different parameter values (Table 2a–f) and present results across all instances at the end (Table 2g). The gaps to the upper bound are reported in percent. Across all instances (Table 2g), we observe that both solution approaches perform well with an average gap to the upper bound of 0.025% (CG approach) and 0.058% (MA approach). The maximum gaps of all 19440 instances are 0.644% and 1.163%, respectively. The gap of the CG approach to the upper bound is always smaller than or equal to the gap of the MA approach to the upper bound, as the base stock level sequences from the MA approach act as the starting solution for the CG approach. On average, the CG approach is able to close approximately half of the gap of the MA approach.

Analyzing the influence of different parameter values, we observe that the gap to the upper bound decreases with an increasing number of parts (Table 2a), thus with increasing problem size. This is

Table 2 Summary of the Performance of the CG Approach and MA Approach

Parameter value	Gap to upper bound (%)				Runtime (sec)			
	CG		MA		CG		MA	
	avg	max	avg	max	avg	max	avg	max
(a) Number of parts in \mathcal{I}								
20	0.036	0.644	0.085	1.163	12.47	1,478	0.002	0.013
35	0.023	0.352	0.053	0.616	41.61	19,115	0.002	0.006
50	0.016	0.262	0.037	0.426	83.71	26,122	0.002	0.011
(b) Item approach fill rate $\hat{\beta}$								
75%	0.039	0.644	0.096	1.163	51.31	26,122	0.002	0.009
85%	0.025	0.390	0.057	0.746	50.13	19,115	0.002	0.011
95%	0.011	0.217	0.021	0.296	36.32	5,905	0.002	0.013
(c) Demand parameter (μ_d, σ_d^2)								
(0.28, 0.87)	0.030	0.644	0.074	1.163	18.59	387	0.002	0.006
(0.55, 1.75)	0.026	0.584	0.058	0.875	28.38	5,101	0.002	0.013
(0.83, 2.63)	0.019	0.535	0.042	0.815	90.82	26,122	0.002	0.011
(d) Holding costs parameter σ_c^2								
3	0.031	0.644	0.067	0.758	50.19	26,122	0.002	0.008
6	0.024	0.609	0.056	0.829	44.85	24,121	0.002	0.013
9	0.020	0.593	0.050	1.163	42.72	11,919	0.002	0.011
(e) Lead time upper bound \hat{b}								
1	0.024	0.609	0.067	1.163	13.15	2,147	0.002	0.013
2	0.024	0.593	0.062	1.130	37.46	2,135	0.002	0.006
3	0.024	0.500	0.053	0.792	47.17	5,388	0.002	0.006
4	0.025	0.644	0.051	0.815	85.90	26,122	0.002	0.011
(f) Planning horizon T								
10	0.025	0.644	0.050	1.163	22.19	3,053	0.002	0.011
12	0.025	0.490	0.060	0.829	42.96	19,115	0.002	0.006
14	0.024	0.584	0.064	0.753	72.61	26,122	0.002	0.013
(g) Total	0.025	0.644	0.058	1.163	45.92	26,122	0.002	0.013


particularly convenient since instances of real-world size, like the ones we investigate in section 5.2, are large. The gaps also decrease with an increasing initial system fill rate $\hat{\beta}$ (Table 2b). With regard to the characteristics of the parts in the inventory system, the average gaps to the upper bound decrease with increasing μ_d and σ_d^2 (Table 2c) and σ_c^2 (Table 2d). We do not observe a strong effect of the lead time upper bound \hat{b} for the CG approach and a decreasing gap with increasing \hat{b} for the MA approach (Table 2e). An increasing planning horizon T (Table 2f) has no strong influence on the gap of the CG approach to the upper bound. For the MA approach, we observe a slight increase in the average gap with an increasing planning horizon but the maximum gaps are decreasing.

While both solution approaches perform well with respect to gaps to the upper bound, the differences in runtimes are substantial. The MA approach provides a solution within milliseconds for all instances (maximum runtime of 13 milliseconds). The runtimes of the CG approach vary strongly. Across all instances, the average CG runtime is 45.92 seconds. However, the runtimes of single instances reveal large differences between 0.02 seconds (not explicitly shown in Table 2) and 26,122 seconds (7.25 hours). We observe this variation because the runtimes of the CG approach

depend on the specific size of every instance, that is, the number of potential decision variables.

The problem size is driven by three elements: the number of parts in the system, the length of the planning horizon, and the difference between the item approach and the system approach base stock levels. The increase in average and maximum runtimes for the first two elements can be observed in Table 2a–f, respectively. An increasing number of parts and an increasing planning horizon result in longer runtimes. The third element, the difference in the base stock levels, drives the large spread of runtimes across all instances. It also drives the increasing runtimes for increasing demand parameters (Table 2c). In systems with low demand rates, more parts have minimum base stock levels of 1 under both the item approach and the system approach, thus fewer differences in the base stock levels and smaller problem sizes. Systems that transition from a lower system fill rate $\hat{\beta}$ (Table 2b) are more difficult to solve, resulting in longer runtimes. Regarding the variation in unit holding costs (Table 2d), runtimes of the CG approach decrease with increasing σ_c^2 . The higher σ_c^2 , the bigger the differences in unit holding costs within a system, which leads to clearer decisions when optimizing the transition. Runtimes of the CG approach are also

Figure 2 Average Gap of the MA Approach to the Upper Bound Depending on Planning Horizon and Number of Parts in the System



increasing with the increasing upper bound on lead time (Table 2e).

The results of the full factorial analysis demonstrate that applying column generation becomes impracticable for large inventory systems and long planning horizons. However, with the MA approach, we can obtain close-to-optimal solutions at short runtimes. To get further indications on how this performance might scale to larger inventory systems with longer planning horizons, we relax the full factorial setup. We fix the system parameters at the values for which we observed the maximum gaps of the MA approach to the upper bound ($\hat{\beta} = 75\%$, $(\mu_d, \sigma_d^2) = (0.28, 0.87)$, $(\mu_c, \sigma_c^2) = (1, 9)$, $\hat{b} = 1$). Figure 2 shows the development of the average gap for planning horizons up to $T = 40$ and inventory systems with up to 100 parts. Even for the parameter values that led to the largest gaps in the full factorial analysis, the gaps to the upper bound remain low.

This further strengthens the positive indications for utilizing the MA approach for transitions that are impractical to solve with column generation. We will apply this solution approach to solve the real-world inventory systems in the next section. Although we cannot provide upper bounds for the large systems, the results in this section suggest that we can expect the good performance of the MA approach also for larger planning problems.

5.2. The Value of Controlling the Transition

We benchmark the *controlled transition* obtained from our solution approach against the default solution of directly applying the system approach base stock levels, which we will refer to as the *uncontrolled transition*. We first analyze the transition of a large-scale real-world inventory system (section 5.2.1). We demonstrate the resulting cost savings of controlling the transition and the impact on the system fill rate compared to the uncontrolled transition. To better understand the effects of the inventory system characteristics on the results, we then perform an extensive

sensitivity analysis (section 5.2.2). We conclude with discussing managerial implications of the results (section 5.2.3). As we solve large systems with many spare parts, long replenishment lead times and long planning horizons, all inventory transitions in this section are solved with the MA approach.

5.2.1 Real-World Application. We analyze the value of controlling the inventory transition based on data from the global manufacturer that motivated our research. The inventory system contains 3191 spare parts. Replenishment orders can be initiated on a daily basis (corresponding to a review period of one day) and arrive after lead times between 1 and 373 days. The demand rates vary between 0.5 and 1,524 units per year. The annual unit holding costs range between below 0.10 euros and above 30,000 euros per spare part, with an average of 416.80 euros.

Currently, the base stock levels of the system are determined with the item approach and the initial system fill rate is 85%. This results in annual system inventory holding cost of approximately 2.2 million euros. With a required minimum part fill rate of 50%, we predict an increase in the system fill rate of 13 percentage points with similar holding cost when moving the inventory system from an item approach to a system approach.

Figure 3 shows the development of the annual system inventory holding cost and the system fill rate for the uncontrolled and the controlled transition. We observe that directly implementing the system approach base stock levels (the uncontrolled transition) results in violations of the allowed holding cost budget of up to 15%, 100 days after the new base stock levels are introduced. On average, the budget is exceeded by 6.3% (142,000 euros) during the first year and by 1.8% (40,000 euros) during the second year of the transition.

The controlled transition does not exceed the budget. The savings in holding cost go along with minor losses in the system fill rate of 0.7% on average during the first year and 0.1% during the second year, with a maximum of 1.8% on day 100. Those losses are the consequence of deliberately delaying the increase of base stock levels for certain parts to stay within the budget constraints in the controlled transition. This leads to a delayed increase in fill rates for the corresponding parts, and a lower system fill rate in the controlled transition than in the uncontrolled transition. However, the differences in system fill rate are marginal compared to the savings in system inventory holding cost. Similar to the uncontrolled transition, the system fill rate steeply increases at the beginning of the transition and the largest share of system fill rate improvement is realized quickly.

Figure 3 Comparison of the Controlled and the Uncontrolled Transition - System Inventory Holding Cost (left) and System Fill Rate (right)




Figure 4 shows the base stock level transition and the resulting on-hand inventory development for three exemplary parts. The first part (Part A) has a high demand over lead time, resulting in a higher system approach base stock level than under the item approach. Since it has relatively high unit holding cost, the increase of base stock levels is delayed and the system approach base stock level is reached relatively late in the transition. The changes in base stock levels are reflected in inventory after the lead time (90 days). Prolonging the base stock level increase of expensive parts is reasonable since all considered parts are of high criticality and any available part provides similar value to the company’s B2B customers. Considering the benefit per monetary unit, it is therefore more beneficial to increase the base stock level of less expensive parts with a similar demand over lead time earlier. This can be observed for Part B. It has a similar demand over lead time but lower unit holding cost than Part A. Therefore, the base stock levels are increased earlier during the transition. The third part (Part C) has a rather low demand over lead time and relatively high unit holding cost. This results in a lower system approach base stock level than under the item approach. The system approach base stock level is set directly in period 1, however, it takes time

until the projected steady-state on-hand inventory value is reached.

5.2.2 Sensitivity Analysis. In the real-world application analyzed in the previous section, we observe significant savings in system inventory holding cost and only marginal losses in system fill rate when controlling the transition. In this section, we investigate how the characteristics of the inventory system (demand rates, unit holding costs, lead times, and initial system fill rates under the item approach) influence the impact on system holding cost and system fill rate during the transition.

We must find appropriate representations of the real-world inventory system’s characteristics that can be varied subsequently. Therefore, we describe the inventory system by the distributions of demand rates, unit holding costs, and lead times. We fitted various distributions to the parameter value realizations and found that the demand rates and the unit holding costs for different spare parts are well represented by lognormal distributions ($\mu_d = -4.50$ and $\sigma_d^2 = 1.60$ for the demand rate distribution and $\mu_c = 0.87$ and $\sigma_c^2 = 2.30$ for the unit holding cost distribution; parameters obtained by maximum likelihood estimation). The lead times are well represented by a

Figure 4 Transition of Base Stock Levels and Expected On-Hand Inventories for Three Exemplary Parts





Figure 5 Comparison of the Transitions of the Real and the Fitted Inventory System - System Inventory Holding Cost (left) and System Fill Rate (right)

zero-truncated negative binomial distribution ($r = 3.03$ and $p = 0.05$). Details on the distribution selection and the fitting of parameters are provided in the Online Appendix.

We substitute the original values for the part characteristics with the theoretical values obtained from the distributions. For every part in the original dataset, we first determine the percentile ranks of its demand rate λ_i , unit holding cost c_i , and lead time l_i (denoted by ρ_i^λ , ρ_i^c , and ρ_i^l). Then, we substitute the empirical demand rate λ_i , unit holding cost c_i , and lead time l_i with the theoretical demand rate $\hat{\lambda}(\rho_i^\lambda)$, unit cost $\hat{c}(\rho_i^c)$, and $\hat{l}(\rho_i^l)$ obtained from the quantile functions of the fitted distributions at the corresponding percentiles. For example, a part at the 75th demand percentile, 32th unit cost percentile, and 10th lead time percentile in the original dataset, will be at the same percentiles of the dataset generated from the theoretical distributions. This procedure generates an inventory system that is close to the original dataset and, therefore, has a similar transition as shown in Figure 5. Without loss of generality, we normalize the budget to 1.

To analyze in which settings controlling the inventory transition is particularly beneficial, we vary the characteristics of the inventory system by modifying the corresponding distribution parameters. For the demand rates of the system's spare parts, we increase and decrease μ_d and σ_d^2 by 25% and 50%. For the unit holding costs, we analogously vary σ_c^2 (since μ_c has no influence on the relative transition). Moreover, we vary the mean of the lead time distribution (denoted by μ_l) and the variance (denoted by σ_l^2) by 25% and 50%. The initial system fill rate $\hat{\beta}$ is varied between 75%, 85%, and 95%.

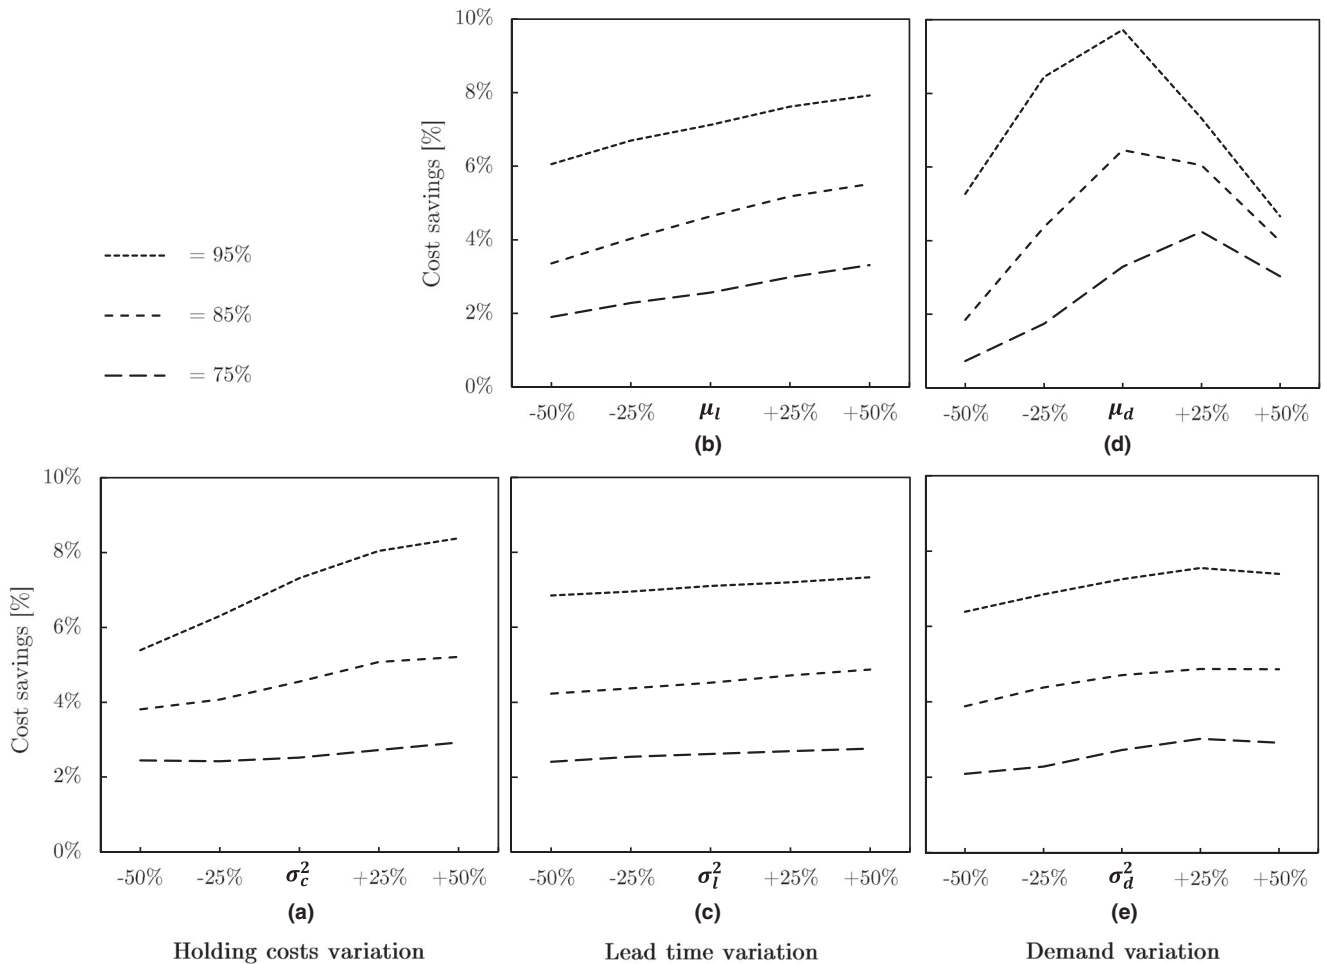
We analyze the transition for all different parameter combinations. Figures 6a–e show the aggregated effect of the different parameter values on the savings

in system inventory holding cost during the first year of the controlled transition, compared to the uncontrolled transition. Across all systems, we observe the highest savings in system inventory holding cost for systems with high initial system fill rates $\hat{\beta}$. The higher the initial system fill rate, the higher the initial item approach base stock levels. The system approach reduces those base stock levels for expensive parts. With the resulting freed-up budget, base stock levels of inexpensive parts are increased. As a result, we observe a larger number of base stock level changes with an increasing magnitude for higher initial system fill rates $\hat{\beta}$. This leads to a more expensive uncontrolled transition. For a similar reason, the spread in unit holding costs influences the cost savings (Figure 6a). A higher spread in unit holding costs leads to larger differences between the item approach and system approach base stock levels and the larger the differences, the higher the temporary increase in system inventory holding cost during the uncontrolled transition.

The distribution of lead times also affects the cost savings. Longer lead times (Figure 6b) positively correlate with the savings due to their influence on the demand over lead time, which is the main driver for base stock levels. A higher demand over lead time yields higher initial base stock levels, and therefore more budget flexibility. The variance of the lead times (Figure 6c) has only little effect on the cost savings. We observe a weak positive influence of a higher spread in lead times, which is mainly driven by the long lead time for some parts and the truncation of the lead time distribution at 1.

The impact of demand parameters on the cost savings is ambiguous. Figure 6d shows the influence of the demand parameter μ_d . We observe two competing effects that result in the inverted u-shape of the cost functions. On the one hand, a higher value for μ_d

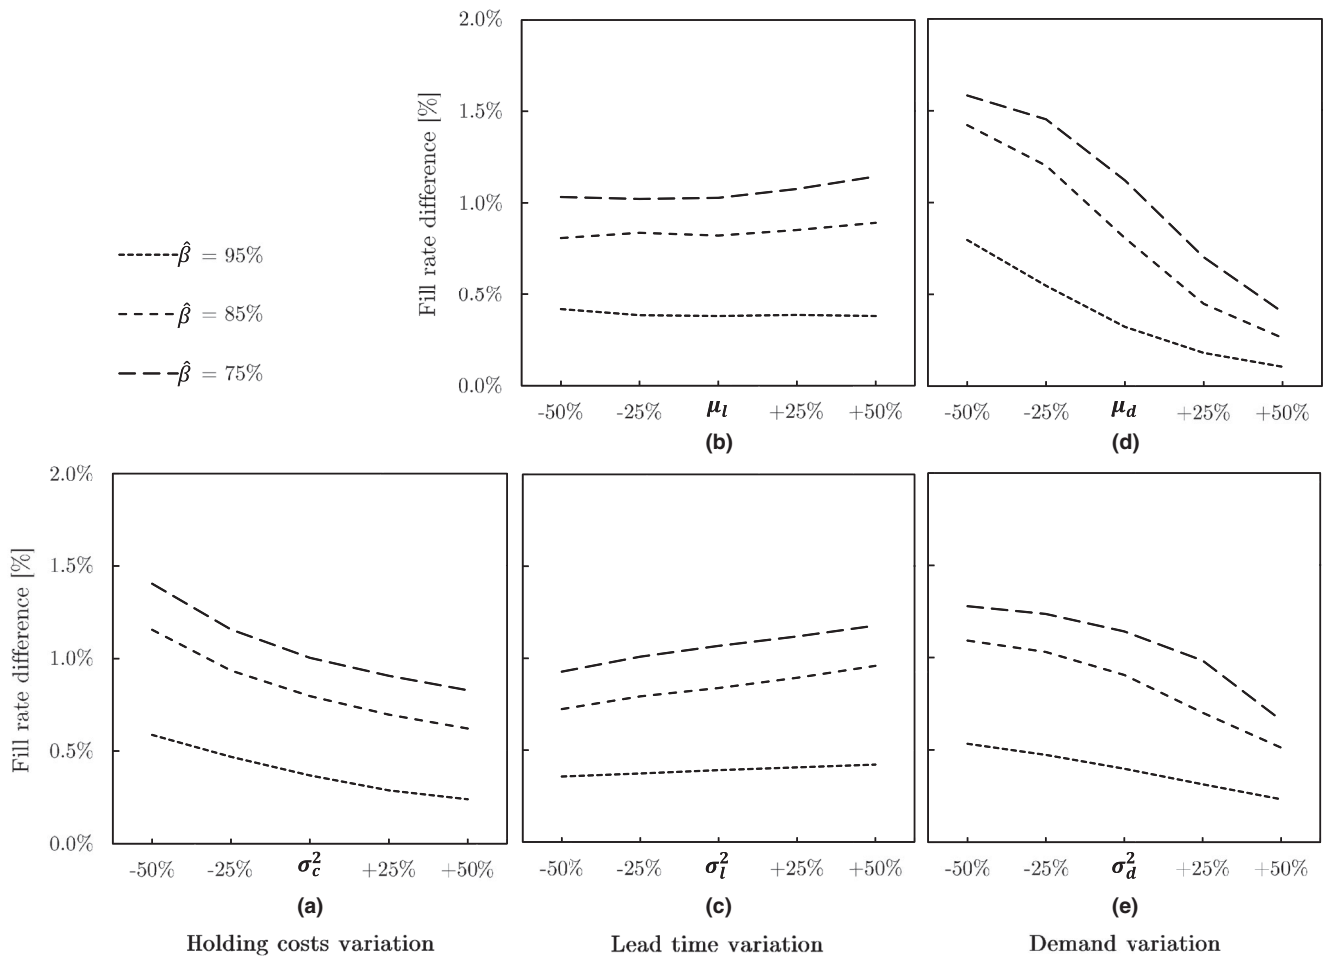
Figure 6 Average Savings in System Inventory Holding Cost When Controlling the Transition Depending on the Inventory System Characteristics



results in higher mean demand rates and faster consumption of excess stock for decreasing base stock levels. This leads to a faster transition and less cost savings when controlling the transition (first effect). On the other hand, base stock levels are bounded from below by 1 due to the minimum fill rate requirement for all parts. With lower demand rates, more parts will be planned with a base stock level of 1, for both the item approach and the system approach. As a result, there are fewer differences in the base stock levels of the two approaches and the impact of controlling the transition declines (second effect). The initial system fill rate $\hat{\beta}$ influences the turning point at which the relative dominance of one effect over the other changes. With a higher initial fill rate, fewer parts are planned with a base stock level of 1; thus, the first effect dominates the second effect for smaller values of μ_d compared to systems with lower initial fill rates. The spread σ_d^2 in demands (Figure 6e) has only limited effect on the cost savings.

We next analyze the impact of the inventory system characteristics on the relative difference in

average system fill rate between the controlled and uncontrolled transition during the first year (Figure 7a–e). Across all system parameter values, the differences in average system fill rate are small and rarely exceed 1.5%. This underlines the value of controlling the inventory transition since it allows to save a significant amount of money while not significantly affecting the system fill rate. We also observe that the higher the initial system fill rate $\hat{\beta}$, the lower the difference in average system fill rate when controlling the transition. The spread in unit holding costs has a relatively strong influence on the difference in average system fill rate. For systems with a smaller spread, the relative difference is higher than for systems with a higher spread (Figure 7a). The lead time characteristics of the inventory system have only very limited influence on the difference in average system fill rate (Figure 7b and c). Regarding the demand parameters, the relative differences in fill rate are decreasing with increasing μ_d and σ_d^2 .

Figure 7 Relative Differences in Average System Fill Rate When Controlling the Transition Depending on the Inventory System Characteristics

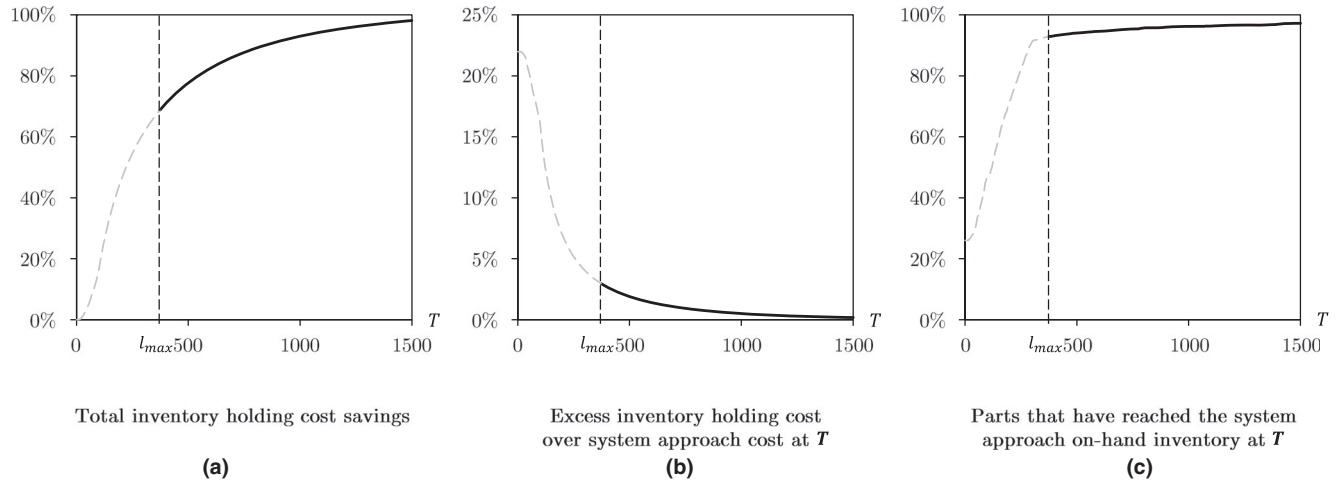
The maximum periodic cost savings and the maximum fill rate difference during the transition show similar patterns for different system characteristics as the discussed cost savings and differences in average system fill rate. Please see the Online Appendix for the corresponding graphs.

In summary, the sensitivity analysis reveals that the value of controlling the inventory transition is particularly sensitive to the initial system fill rate and the spread in unit holding costs. We can achieve the largest benefits in terms of system holding cost savings and small loss in average system fill rate by controlling the transition of systems that operate under a high system fill rate at the beginning of the transition. Also, the larger the spread in unit holding costs, the larger the benefits. Longer lead times and a higher spread in lead times increase the cost saving potential as well. The impact of demand parameters is ambiguous and needs further analysis for each particular situation.

5.2.3 Discussion of Managerial Implications. The results of the case study and sensitivity analysis

provide important managerial insights. They demonstrate that companies can realize large savings in inventory holding cost with moderate losses in system fill rate by controlling the transition as opposed to adapting an optimized solution for all parts at a single point in time. The magnitudes of the savings and losses depend on the inventory system characteristics, such as the initial system fill rate, the spread in unit holding costs, and the demand rate distribution. The savings and losses evolve over time and it can take long until the optimal inventory levels have been reached by all parts of the system. Thus, the required time for reaching the optimal solution for all parts will typically exceed reasonable lengths of the planning period. This raises the issue on the appropriate length of planning horizon T . Our solution approaches only require it to be longer than the maximum part lead time. However, a planning horizon T that is too short might not allow for a sufficient number of parts to reach their projected steady-state inventory levels during the transition. Our approach

Figure 8 Effect of Planning Horizon T in Days on Performance Metrics



ensures that the expected system inventory holding cost is within the budget until period T , but not necessarily beyond period T . Consequently, the inventory holding cost might be above budget in periods $t > T$.

Figure 8 presents three metrics for the inventory system of section 5.2.1 as a function of the planning horizon T . Figure 8a depicts how the realized cumulative holding cost savings when controlling the transition depend on T . Compared to the total holding cost savings that were possible with an infinite planning horizon (corresponding to cost savings of 100%), we observe that the largest savings are realized at the beginning of the transition. The marginal savings decrease in the planning horizon. Figure 8b shows the excess inventory holding cost above the system approach cost for parts with decreasing base stock levels. This cost corresponds to the expected system inventory holding cost above the budget b for periods after $T + 1$. The excess inventory holding cost decreases with longer planning horizons T and is moderately above target after one year. However, requiring it to completely diminish would result in very long planning horizons T . Figure 8c shows the expected percentage of parts that have reached the system approach on-hand inventory at T . The vast majority of parts reach their system approach levels within a year.

The length of an appropriate planning horizon depends on the goals of the company and the characteristics of the inventory system. Companies can, for example, set a maximum percentage $x^+ \geq 0$ of the expected excess inventory holding cost above budget in period $T + 1$:

$$c^{T+1}(\bar{s}) \leq (1 + x^+) \cdot b. \quad (14)$$

To determine an appropriate planning horizon T , Equation (14) could be included in Problem (P^T) with

T as a decision variable. However, this would make (P^T) difficult to solve. Instead, companies can utilize Equation (14) and analyze the expected system inventory holding cost in period $T + 1$ of the uncontrolled transition, with \bar{s}^u denoting the corresponding vector of base stock level sequences. Since $c^{T+1}(\bar{s}^u)$ is monotone decreasing in T for $T > l^{max}$, they can determine the smallest $T > l^{max}$ for which Equation (14), $\bar{s} = \bar{s}^u$, is fulfilled. If base stock level decreases are delayed in the controlled transition, the system inventory holding cost in period $T + 1$ can exceed $c^{T+1}(\bar{s}^u)$. In our numerical study, this is negligible. For the case study of section 5.2.1, a percentage x^+ of 5%, 1%, and 0.5% results in planning horizons of 373, 720, and 1001 days, respectively.

To reduce the inventory above budget after the transition, companies can decide on whether to extend the planning horizon T or to remove excess inventory from stock at T . For our application, stock removal is rather undesirable since the company holds durable, specialized spare parts without an external market value. Therefore, removing inventory from stock is expensive. This particularly holds for the parts with lower base stock levels under the system approach than under the item approach. Due to the nature of the system approach, those parts will be the expensive ones in the spare parts portfolio.

6. Conclusion

In this study, we have analyzed how the transition of inventory systems can be controlled when the inventory policy changes from an item approach to a system approach. We have shown that an uncontrolled transition results in severe violations of operational constraints, which can jeopardize the successful implementation of the system approach in practice.

We have formulated an optimization problem that enables a smooth adaptation of base stock levels over time with the objective to maximize the average system fill rate without exceeding the inventory holding cost budget during the transition. We have presented two approaches to solve the problem. The first approach is based on decomposition and column generation and provides an upper bound and a well-performing feasible solution to the maximization problem. However, it requires long runtimes for large-scale inventory systems and long transition periods. The second solution approach is a heuristic that is based on marginal analysis. It is able to solve the transition of large-scale inventory systems with very short runtimes. We have demonstrated the performance of both approaches in an extensive numerical study.

Our research provides important managerial implications. First, it raises awareness for the unique challenges companies have to face when optimizing an existing inventory system. In such situations, it is not only important to establish a new target state of the system but it is also important to consider the transition period. Second, we provide insights on why the inventory system performance might fall short of the target performance during a transition period and how severe this deterioration can be. Even if managers decide against transition period control and cope with the resulting disturbances instead, our research helps to understand the implications and to avoid wrong interpretations and actions. It might also reduce the stress put on responsible units within a company. Third, we provide a solution approach that enables companies to keep inventory cost within a budget during the transition period. We also provide guidance on how to choose an appropriate length for the transition period. The suggested controlled transition is easy to apply in practice as the recommended base stock levels during the transition period can be integrated in existing operational routines and ERP systems.

To maximize the average system fill rate without exceeding the inventory holding cost budget, we allow the base stock levels to change gradually during the transition period. If a gradual adaptation is not desired, our approach could be simplified to a single-adjustment policy per item. Pinçe et al. (2015) have demonstrated that such a policy performs well for the base stock level transition of a single item in continuous-time. Going forward, it would be interesting to analyze how such a policy performs in our multi-item setting.

Even though we have focused on the introduction of the system approach in an inventory system that has been operated under the item approach before, our approach is suited for any adaptations of base stock levels in an inventory system that should

eventually operate under a system approach. This makes our approach applicable for a broad range of situations, for example, if new information regarding demand or reliability of parts is available, unit holding costs change, parts are substituted, or new supply sources are explored. If resulting changes in base stock levels are not quickly reflected in inventory, controlling the transition is important and valuable.

Appendix Characteristics of (P^S)

For the fill rate $\beta_i(s_i)$ and Poisson-distributed demand it holds:

PROPERTY 1. *The fill rate $\beta_i(s_i)$ is concave increasing for $s_i \geq \frac{\lambda_i}{\ln(l_i+1) - \ln(l_i)} - 1$.*

For the expected inventory holding cost $c_i \cdot oh_i(s_i)$ it holds:

PROPERTY 2. *The expected inventory holding cost $c_i \cdot oh_i(s_i)$ is convex increasing in s_i for $s_i \geq 0$.*

From Property 1 and 2 and Fox (1966 sect. 8), it follows that Problem (P^S) can be solved efficiently by marginal analysis for $s_i \geq \frac{\lambda_i}{\ln(l_i+1) - \ln(l_i)} - 1$. In the following, we provide proofs for Property 1 and 2.

PROOF OF PROPERTY 1. *The fill rate $\beta_i(s_i)$ is concave increasing in s_i if the first-order difference $\Delta\beta_i(s_i) = \beta_i(s_i + 1) - \beta_i(s_i) \geq 0$ and the second-order difference $\Delta^2\beta_i(s_i) = \Delta\beta_i(s_i + 1) - \Delta\beta_i(s_i) \leq 0$.*

Let us assume that

$$\frac{\lambda_i}{\ln(l_i + 1) - \ln(l_i)} - 1 \leq s_i$$

which is equivalent to

$$(s_i + 1) \cdot \ln(\lambda_i \cdot l_i) \leq -\lambda_i + (s_i + 1) \cdot \ln(\lambda_i \cdot (l_i + 1)).$$

With e^x being a monotone increasing function in $x \in \mathbb{R}$, $e^{-\lambda_i \cdot l_i} > 0$ and $(s_i + 1)! > 0$, this is equivalent to

$$e^{-\lambda_i \cdot l_i} \cdot \frac{(\lambda_i \cdot l_i)^{s_i+1}}{(s_i + 1)!} \leq e^{-\lambda_i \cdot (l_i+1)} \cdot \frac{(\lambda_i \cdot (l_i + 1))^{s_i+1}}{(s_i + 1)!}. \quad (A1)$$

Since D_i is Poisson-distributed with demand rate λ_i , $D_i^{l_i}$ and $D_i^{l_i+1}$ are Poisson-distributed with $\lambda_i \cdot l_i$ and $\lambda_i \cdot (l_i + 1)$ and Equation (A1) can be rewritten as

$$P(D_i^{l_i} = s_i + 1) \leq P(D_i^{l_i+1} = s_i + 1). \quad (A2)$$

From Equation (A2), it follows that $\Delta\beta_i(s_i) = E[D_i]^{-1} \cdot \sum_{d=s_i+1}^{\infty} (P(D_i^{l_i+1} = d) - P(D_i^{l_i} = d)) \geq 0$ and

$\Delta^2 \beta_i(s_i) = E[D_i]^{-1} \cdot (P(D_i^{l_i} = s_i + 1) - P(D_i^{l_i+1} = s_i + 1)) \leq 0$
 if $s_i \geq \frac{\lambda_i}{m(l_i+1) - m(l_i)} - 1$. Therefore, $\beta_i(s_i)$ is concave increasing in this domain.

PROOF OF PROPERTY 2. From

$$\Delta oh_i(s_i) = oh_i(s_i + 1) - oh_i(s_i) = \sum_{d=0}^{s_i} P(D_i^{l_i+1} = d) \geq 0,$$

it follows that $oh_i(s_i)$ is increasing for all $s_i \geq 0$. Furthermore,

$\Delta^2 oh_i(s_i) = \Delta oh_i(s_i + 1) - \Delta oh_i(s_i) = P(D_i^{l_i+1} = s_i + 1) \geq 0$,
 which shows that $oh_i(s_i)$ is convex in s_i for all $s_i \geq 0$. Therefore, with $c_i \in \mathbb{R}^+$, $c_i \cdot oh_i(s_i)$ is convex increasing for all $s_i \geq 0$.

Appendix B: Column Generation

The solution of the master problem (MP) is an optimal solution to the original problem (P^T). However, solving (MP) optimally becomes numerically intractable even for moderate problem sizes. To overcome this computational challenge, we relax the integrality constraints on x_i^n and apply column generation which provides an upper bound to (MP). We present this approach in the following.

B.1. Computing an Upper Bound

Our master problem (MP) contains all possible base stock level sequences per part as decision variables. We introduce a restricted master problem (RMP) that contains only a subset of possible sequences $\bar{S}_i \subseteq \bar{S}_i$, for example the naïve sequences of applying the minimum of the item approach and system approach base stock levels up to period $T - l_i$ and the system approach base stock levels starting from period $T - l_i + 1$. The (RMP) is smaller and easier to solve. Moreover, we relax the Integrality Constraints (12) to generate a linear program.

After solving the initial (RMP), we search for sequences that can be added to the sets \bar{S}_i to improve the solution. For parts with $s_i^{item} = s_i^{system}$, there is exactly one feasible sequence. No additional sequences can be added and we will disregard those parts in the following. Let y_t be the dual variables of the Budget Constraints (10) and z_i the dual variables of Constraints (11) that ensure one policy per part.

$$(SUB_i) \quad \max \quad \frac{\lambda_i}{\Lambda} \cdot \sum_{t=1}^T \beta_i^t(\bar{s}_i) - c_i \cdot \sum_{t=1}^T oh_i^t(\bar{s}_i) \cdot y_t - z_i \quad (A3)$$

$$\begin{aligned}
 \text{s.t.} \quad & s_i^t = s_i^{item} \quad \forall t \leq 0 \\
 & s_i^t = s_i^{system} \quad \forall t > T - l_i \\
 & \bar{s}_i \text{ is monotone} \\
 & s_i^t \in \mathbb{N}_0 \quad \forall t = 1, \dots, T.
 \end{aligned}$$

If (SUB_i) has a solution with a positive value of the Objective Function (A3), adding the obtained sequence to \bar{S}_i improves the solution of (RMP). We solve the subproblem for every part i and add the sequence with the largest positive objective value to the corresponding set \bar{S}_i .

We repeatedly solve (RMP) with the updated set \bar{S}_i , record the new dual variables and add new sequences obtained from the subproblems until no subproblem with a positive objective value remains (i.e., no sequence can be added to further improve (RMP)). The optimal solution of the final (RMP) is also an optimal solution to (MP) with relaxed integrality constraints. Therefore, the solution is an upper bound to the original problem.

B.2. Generating New Sequences by Solving the Subproblem

To apply column generation, it is crucial to solve the subproblem (SUB_i) efficiently. Generally, every sequence with a positive objective value of (SUB_i) can improve the restricted master problem (RMP). However, an optimal solution of (SUB_i) ensures that the most promising sequence is added. As shown in section 3.2, transient fill rates and transient on-hand inventories are calculated differently for parts with increasing and decreasing base stock levels. While for parts with increasing base stock levels, the values in period t only depend on the base stock level in period $t - l_i$, for decreasing base stock levels the whole sequence of base stock levels from period 0 to $t - l_i$ needs to be considered.

We exploit the characteristics of the transition for increasing base stock levels to solve the subproblem with an integer optimization problem:

$$(SUB_i^{incr}) \quad \max \quad \frac{\lambda_i}{\Lambda} \sum_{t=1}^T \sum_{s=s_i^{item}}^{s_i^{system}} \beta_i(s) \cdot \chi_s^{t-l_i} - c_i \cdot \sum_{t=1}^T y_t \sum_{s=s_i^{item}}^{s_i^{system}} oh_i(s) \cdot \chi_s^{t-l_i} - z_i \quad (A4)$$

$$\text{s.t.} \quad \sigma^t = s_i^{item} \quad \forall t \leq 0 \quad (A5)$$

$$\sigma^t = s_i^{system} \quad \forall t > T - l_i \quad (A6)$$

$$\sigma^{t-1} \leq \sigma^t \quad \forall t = 1, \dots, T \quad (A7)$$

$$\sum_{s=s_i^{item}}^{s_i^{system}} s \cdot \chi_s^t = \sigma^t \quad \forall t = -l_i, \dots, T \quad (A8)$$

$$\sum_{s=s_i^{item}}^{s_i^{system}} \chi_s^t = 1 \quad \forall t = -l_i, \dots, T \quad (A9)$$

$$\sigma^t \in \mathbb{N}_0 \quad \forall t = -l_i, \dots, T \quad (\text{A10})$$

$$\chi_s^t \in \{0, 1\} \quad \forall t = -l_i, \dots, T. \quad (\text{A11})$$

Problem (SUB_i^{incr}) has two sets of decision variables, σ^t with $t = (-l_i, \dots, 0, \dots, T)$ for the base stock level in each period and corresponding binary variables χ_s^t with $t = (-l_i, \dots, 0, \dots, T)$ and $s \in (s_i^{item}, \dots, s_i^{system})$, formalized in Constraints (A10) and (A11). Constraints (A5) and (A6) ensure that a feasible sequence starts with the item approach base stock level s_i^{item} and ends with the system approach base stock level s_i^{system} . Constraints (A7) guarantee the monotonicity of the sequences. Constraints (A8) connect the base stock level variables σ^t with the corresponding binaries χ_s^t and Constraints (A9) ensure one base stock level per period. With this set of constraints we guarantee that newly generated sequences fulfill the requirements outlined in section 3.2. The objective function is a reformulation of Objective Function (A3) with $\beta_i^t(\bar{s}_i) = \sum_{s=s_i^{item}}^{s_i^{system}} \beta_i(s) \cdot \chi_s^{t-l_i}$ and $oh_i^t(\bar{s}_i) = \sum_{s=s_i^{item}}^{s_i^{system}} oh_i(s) \cdot \chi_s^{t-l_i}$.

For parts with decreasing base stock levels, we cannot apply (SUB_i^{incr}) because the transient fill rates and transient on-hand inventories in period t do not only depend on the base stock level $s_i^{t-l_i}$ but on the entire sequence up to period $t-l_i$. We therefore solve (SUB_i) for decreasing base stock levels with enumeration of all feasible sequences that are not in \bar{S}_i' . We note that if s_i^{item} is much larger than s_i^{system} , this method is inefficient, especially for large T (see Equation 13). However, caused by the nature of the system approach, the differences between the initial and target base stock levels are typically much smaller for parts with decreasing base stock levels than for parts with increasing base stock levels. Ceteris paribus, parts with decreasing base stock levels are more expensive than parts with increasing base stock levels. Therefore, small base stock level decreases of the expensive parts allow for larger base stock level increases of the less expensive parts. In our numerical study, solving the subproblem for decreasing base stock levels is not a bottleneck.

References

- Alvarez, E. M., M. C. Van der Heijden, W. H. M. Zijm. 2013. The selective use of emergency shipments for service-contract differentiation. *Int. J. Prod. Econ.* **143**(2): 518–526.
- Alvarez, E. M., M. C. Van der Heijden, W. H. M. Zijm. 2015. Service differentiation in spare parts supply through dedicated stocks. *Ann. Oper. Res.* **231**(1): 283–303.
- Arts, J. 2017. A multi-item approach to repairable stocking and expediting in a fluctuating demand environment. *Eur. J. Oper. Res.* **256**(1): 102–115.
- Axsäter, S. 2015. *Inventory Control*, Vol. **225**. Springer, Cham.
- Basten, R. J. I., G.-J. Van Houtum. 2014. System-oriented inventory models for spare parts. *Surv. Oper. Res. Manag. Sci.* **19**(1): 34–55.
- Cachon, G., C. Terwiesch. 2006. *Matching Supply with Demand*, McGraw-Hill Publishing, New York.
- Caggiano, K. E., J. A. Muckstadt, J. A. Rappold. 2006. Integrated real-time capacity and inventory allocation for repairable service parts in a two-echelon supply system. *Manuf. Serv. Oper. Manag.* **8**(3): 292–319.
- Cavaliere, S., M. Garetti, M. Macchi, R. Pinto. 2008. A decision-making framework for managing maintenance spare parts. *Prod. Plann. Contr.* **19**(4): 379–396.
- Çetinkaya, S., M. Parlar. 2010. A one-time excess inventory disposal decision under a stationary base-stock policy. *Stochast. Anal. App.* **28**(3): 540–557.
- Cohen, M. A., P. R. Kleindorfer, H. L. Lee. 1989. Near-optimal service constrained stocking policies for spare parts. *Oper. Res.* **37**(1): 104–117.
- Cohen, M. A., P. R. Kleindorfer, H. L. Lee, D. F. Pyke. 1992. Multi-item service constrained (s, S) policies for spare parts logistics systems. *Naval Res. Logist.* **39**(4): 561–577.
- Dantzig, G. B., P. Wolfe. 1960. Decomposition principle for linear programs. *Oper. Res.* **8**(1): 101–111.
- Desrosiers, J., M. E. Lübbecke. 2005. A primer in column generation. G. Desaulniers, J. Desrosiers, M. M. Solomon, eds. *Column Generation*. Springer, Boston, 1–32.
- Drent, M., J. Arts. 2020. Expediting in two-echelon spare parts inventory systems. *Manuf. Serv. Oper. Manag.* Forthcoming.
- Fox, B. 1966. Discrete optimization via marginal analysis. *Management Sci.* **13**(3): 210–216.
- Kranenburg, A. A., G.-J. Van Houtum. 2007. Effect of commonality on spare parts provisioning costs for capital goods. *Int. J. Prod. Econ.* **108**(1–2): 221–227.
- Kranenburg, A. A., G.-J. Van Houtum. 2008. Service differentiation in spare parts inventory management. *J. Oper. Res. Soc.* **59**(7): 946–955.
- Mitchell, J. C. 1988. Multi-item inventory systems with a service objective. *Oper. Res.* **36**(5): 747–755.
- Muckstadt, J. A. 2005. *Analysis and Algorithms for Service Parts Supply Chains*, Springer Science & Business Media, Berlin.
- Pınar, Ç., R. Dekker. 2011. An inventory model for slow moving items subject to obsolescence. *Eur. J. Oper. Res.* **213**(1): 83–95.
- Pınar, Ç., J. B. G. Frenk, R. Dekker. 2015. The role of contract expirations in service parts management. *Prod. Oper. Manag.* **24**(10): 1580–1597.
- Rustenburg, W. D., G.-J. Van Houtum, W. H. M. Zijm. 2000. Spare parts management for technical systems: resupply of spare parts under limited budgets. *IIE Trans.* **32**(10): 1013–1026.
- Rustenburg, W. D., G.-J. Van Houtum, W. H. M. Zijm. 2003. Exact and approximate analysis of multiechelon, multi-indenture spare parts systems with commonality. J. G. Shantikumar, D. D. Yao, W. H. M. Zijm, eds. *Stochastic Modeling and Optimization of Manufacturing Systems and Supply Chains*. Kluwer Academic Publishers, Boston, 143–176.
- Schwarz, L. B., B. L. Deuermeyer, R. D. Badinelli. 1985. Fill-rate optimization in a one-warehouse n-identical retailer distribution system. *Management Sci.* **31**(4): 488–498.
- Sherbrooke, C. C. 1968. METRIC: A multi-echelon technique for recoverable item control. *Oper. Res.* **16**(1): 122–141.
- Sherbrooke, C. C. 2004. *Optimal Inventory Modeling of Systems: Multi-Echelon Techniques*, Kluwer Academic Publishers, Boston.
- Silver, E. A., D. F. Pyke, D. J. Thomas. 2016. *Inventory and Production Management in Supply Chains*, CRC Press, Boca Raton.

- Teunter, R. H., W. K. Klein Haneveld. 2002. Inventory control of service parts in the final phase. *Eur. J. Oper. Res.* **137**(3): 497–511.
- Thonemann, U. W., A. O. Brown, W. H. Hausman. 2002. Easy quantification of improved spare parts inventory policies. *Management Sci.* **48**(9): 1213–1225.
- Tiemessen, H. G., M. Fleischmann, G.-J. van Houtum, J. A. van Nunen, E. Pratsini. 2013. Dynamic demand fulfillment in spare parts networks with multiple customer classes. *Eur. J. Oper. Res.* **228**(2): 367–380.
- Topan, E., Z. P. Bayindir, T. Tan. 2010. An exact solution procedure for multi-item two-echelon spare parts inventory control problem with batch ordering in the central warehouse. *Oper. Res. Lett.* **38**(5): 454–461.
- Topan, E., Z. P. Bayindir, T. Tan. 2017. Heuristics for multi-item two-echelon spare parts inventory control subject to aggregate and individual service measures. *Eur. J. Oper. Res.* **256**(1): 126–138.
- Van Houtum, G.-J., B. Kranenburg, 2015. *Spare Parts Inventory Control under System Availability Constraints*, Vol. 227. Springer, New York.
- Wang, W. 2012. A stochastic model for joint spare parts inventory and planned maintenance optimisation. *Eur. J. Oper. Res.* **216**(1): 127–139.
- Wong, H., G.-J. Van Houtum, D. Cattrysse, D. Van Oudheusden. 2005. Simple, efficient heuristics for multi-item multi-location spare parts systems with lateral transshipments and waiting time constraints. *J. Oper. Res. Soc.* **56**(12): 1419–1430.
- Wong, H., B. Kranenburg, G.-J. Van Houtum, D. Cattrysse. 2007. Efficient heuristics for two-echelon spare parts inventory systems with an aggregate mean waiting time constraint per local warehouse. *OR Spectrum* **29**(4): 699–722.
- Wright, G. G. 1991. An empirical analysis and stochastic modelling of aggregate demand behaviour in a spare parts inventory system. PhD thesis, City University London.
- Wright, G. G. 1992. Optimising net margins and stock position in spare parts inventory echelons. *Omega* **20**(4): 503–511.
- Yang, L., H. Li, J. F. Campbell, D. C. Sweeney. 2017. Integrated multi-period dynamic inventory classification and control. *Int. J. Prod. Econ.* **189**: 86–96.

Supporting Information

Additional supporting information may be found online in the Supporting Information section at the end of the article.

Online Appendix: How to Change a Running System—Controlling the Transition to Optimized Spare Parts Inventory Policies.