

Dynamic capacitated lot sizing with random demand and dynamic safety stocks

Stefan Helber^a, Florian Sahling^a and Katja Schimmelpfeng^b

^aDepartment of
Production Management
Leibniz University Hannover, Germany
stefan.helber@prod.uni-hannover.de, ++49 511 7625650

^bChair of
Accounting and Control
Brandenburg University of Technology, Germany

February 22, 2011, revised July 18 and November 14, 2011

accepted for publication in *Operations Research Spectrum*

Abstract

We present a stochastic version of the single-level, multi-product dynamic lot-sizing problem subject to a capacity constraint. A production schedule has to be determined for random demand so that expected costs are minimized and a constraint based on a new backlog-oriented δ -service-level measure is met. This leads to a non-linear model that is approximated by two different linear models. In the first approximation, a scenario approach based on random samples is used. In the second approximation model, the expected values of physical inventory and backlog as functions of the cumulated production are approximated by piecewise linear functions. Both models can be solved to determine efficient, robust and stable production schedules in the presence of uncertain and dynamic demand. They lead to dynamic safety stocks that are endogenously coordinated with the production quantities. A numerical analysis based on a set of (artificial) problem instances is used to evaluate the relative performance of the two different approximation approaches. We furthermore show under which conditions precise demand forecasts are particularly useful from a production-scheduling perspective.

Keywords: Stochastic demand, lot-sizing, service level, dynamic safety stocks

Acknowledgement: The authors thank the referees and the editor for their very helpful comments.

1 Introduction

In make-to-stock production environments, production planners often face an uncertain dynamic demand when deciding about the timing and sizing of production quantities. These decisions are often made for a finite planning horizon divided into discrete time periods, e.g., days or weeks. One principle approach to deal with this demand uncertainty is to determine and/or adjust planned production quantities whenever demand for the current period has actually been realized, i.e., the demand is known with certainty. The drawback of this approach is that it can induce substantial system nervousness and, in multi-level supply chains, lead to the notorious bullwhip effect. The other principle approach is to determine dynamic demand forecasts and to establish a production plan that is fixed or “frozen” for the next periods. This schedule is then executed irrespective of the demand realizations for those next periods. While the latter approach leads to more stable planning and production systems, it typically requires safety stocks in order to meet an uncertain customer demand. If demand is uncertain and dynamic, cost-minimizing planned safety stocks must be dynamic as well. If multiple different product types compete for scarce production capacity, their production quantities and safety stocks need to be coordinated. Furthermore, if setup times and/or setup costs lead to lot sizes covering the demand for multiple consecutive periods, these production lots substantially reduce the need to hold specific safety stocks. Both in the academic literature and in current management practice, these problems are addressed in an often inconsistent and unsatisfactory manner, see Fleischmann (2003). We hence treat these problems jointly and consistently in the context of a proposed Stochastic Capacitated Lot-Sizing Problem (SCLSP), which is a stochastic generalization of the established deterministic Capacitated Lot-Sizing Problem (CLSP).

In this SCLSP, we assume that unmet demand can be back-ordered and that the management uses a constraint on the resulting expected backlog, i.e., the cumulated open (or yet unmet) back-orders, to control the service that is offered to the customers. To this end, the newly developed δ -service-level measure is introduced and used. As both the expected physical inventory and the backlog in the SCLSP are non-linear functions of the production quantities, a non-linear model formulation is developed first. Since we are not aware of a way to solve this first model directly, we further introduce two approximations of the SCLSP by numerically tractable mixed-integer linear models. The first approximation is based on a sample of scenarios, each representing a single realization of the demand process. The second approximation operates with piecewise linear approximations of the expected values of physical inventory and backlog for a given cumulated production quantity. Both modeling approaches are solved using a specific Fix-and-Optimize algorithm as introduced in Helber and Sahling (2010) and Sahling (2010).

The outline of the paper is as follows: In section 2 we present the problem in more detail, analyze service-level measures and discuss the related literature. The model variants of the SCLSP are developed in sections 3 to 5. The Fix-and-Optimize algorithm used to solve both approximation models is outlined in section 6. Section 7 presents numerical results with respect to the accuracy of the approximation models and the relationship between demand forecasts, lot sizes and dynamic planned safety stocks. Section 8 concludes and outlines future research topics.

2 Problem statement, service-level measures and relevant literature

2.1 Capacitated lot sizing subject to dynamic and stochastic demand

We assume that a single production system or machine is required to produce K different products. This machine has a (regular) capacity b_t in each of the T discrete periods of the planning horizon. It can be extended by overtime at a cost oc per time unit. We do not consider the sequence of the products within any period. If a product k is produced during period t , i.e., with production quantity $q_{kt} > 0$, a setup time $ts_k \geq 0$ is required and a setup cost $sc_k \geq 0$ occurs. The processing time for a unit of product k is tb_k . The cost of holding one unit of physical inventory for one period is denoted as hc_k .

The demand of product k in period t is modeled as a random variable D_{kt} with a given probability distribution, given expected value $E[D_{kt}]$ and variance $\text{VAR}[D_{kt}]$. The demand for the product-period combination (k, t) is assumed to be independent from those for any other combination (\hat{k}, \hat{t}) with $k \neq \hat{k}$ and/or $t \neq \hat{t}$. Estimators of $E[D_{kt}]$ and variance $\text{VAR}[D_{kt}]$ are assumed to be provided by a forecasting system. In order to avoid production schedules that systematically create new backlog at the end of the planning horizon, we assume that for each product k , the total production $\sum_{t=1}^T q_{kt}$ over the entire planning horizon must at least be sufficient to meet the expected total demand $\sum_{t=1}^T E[D_{kt}]$.

If in any period t the cumulated (random) demand $\sum_{\tau=1}^t D_{k\tau}$ of product k exceeds the cumulated (deterministic) production $\sum_{\tau=1}^t q_{k\tau}$, the unmet demand is back-ordered and a positive value of the (random) backlog

$$BL_{kt} = \max \left(0, \sum_{\tau=1}^t (D_{k\tau} - q_{k\tau}) \right) \quad (1)$$

occurs. The opposite case results in a positive value of the (random) physical inventory:

$$YP_{kt} = \max \left(0, \sum_{\tau=1}^t (q_{k\tau} - D_{k\tau}) \right) \quad (2)$$

Due to the non-linear maximum functions in equations (1) and (2), both the expected physical inventory $E[YP_{kt}]$ and the expected backlog $E[BL_{kt}]$ are non-linear functions of the cumulated production in periods 1 to t .

Below we develop a new service-level constraint to limit the total expected backlog relative to the maximum possible total expected backlog. The overall objective is to minimize the expected costs due to setups, inventory and overtime.

If the demand was deterministic and we did not allow any backorders, this setting would describe the Capacitated Lot-Sizing problem (CLSP) which has already attracted a great deal of research, in particular as it is extended to multi-level production systems with multiple capacity constraints and/or to systems in which the setup state can be carried over from one period to the next.

In this paper, however, we assume that demand is stochastic. We still aim at a production plan that is capacity feasible. In addition, this production plan has to be robust in the sense that it systematically reflects demand uncertainty and leads to an a

priori given service level. After the formal presentation of the model we will show that it leads to dynamic safety stocks that are coordinated with the production quantities.

Note that with respect to the lot sizing aspect of the problem, we are working with a so-called “big-bucket” model as we assume in the (S)CLSP that multiple different products can be produced during a single period. We do not consider the sequence within those periods or setup carry-overs at period borders. If one assumes that (relatively stable) demand forecasts are provided for larger time periods, such as weeks, this appears to be a natural approach. However, in principle it would also be possible to formulate stochastic versions of established “small-bucket” models and hence to integrate sequencing decisions.

2.2 Service-level measures

In order to limit the (expected) backlog associated with a production plan subject to random demand, one can penalize it in the objective function using a backlog cost parameter. However, as such a parameter reflects essentially opportunity costs and is hence debatable, it is frequently proposed to limit backorders and/or backlog via a constraint on a suitably defined service-level measure. The following three different service-level measures are frequently discussed in the inventory management literature considering random demand, see Tempelmeier (2006, pp. 26-27):

- The α -service level represents the probability that a stockout occurs during a (production or procurement) cycle, or during a period in the rare case of a discrete time model. While this measure may be mathematically tractable and hence attractive from the mathematical point of view, it offers no insights as to how severe the stockout event is once it occurs. From a managerial perspective, the usefulness of this service-level measure seems to be limited.
- The β -service level (or “fill rate”) is the fraction of the demand per cycle that is met immediately, i.e., without backlogging. In a stationary environment these cycles are stochastically identical, so that their average length and hence the β -service-level can be computed. However, if production quantities $q_{kt} \geq 0$ are determined for each product-period-combination (k, t) , a β -service-level constraint per production cycle can only be enforced once the setup epochs are known. This makes it difficult to *simultaneously* determine production quantities. Furthermore, the β -service-level does not reflect the waiting time of the customer. Only the size of the backorder within a cycle is relevant, but not the time required to meet the back-ordered demand.
- Finally, the γ -service level (here for a specific period t) is defined as

$$\gamma_t = 1 - \frac{\text{expected backlog in period } t}{\text{expected demand in period } t}. \quad (3)$$

The attractive feature of this measure is that it reflects backlog and hence, to some extent the waiting time of the customers. It can be defined either (as above) for a specific period t , or as an average over the entire planning horizon. However, in a particular period the expected demand may be smaller than the expected backlog or even be zero. Therefore this measure can be negative or even undefined (if the expected demand is zero). From a management perspective, the γ -service level measure does not seem to offer a clear interpretation.

In the context of dynamic lot sizing, a backlog-oriented measure seems to be desirable that reflects both the size of the backorders and the waiting time of the customers, but other than the γ -service-level measure, it should be well-defined and have a clear and transparent interpretation. For this reason, we now define the new δ -service-level measure conceptually as follows:

$$\delta = 1 - \frac{\text{total expected backlog}}{\text{total maximum expected backlog}}. \quad (4)$$

The maximum expected backlog in a period equals the expected cumulated demand including that of the considered period. It occurs if nothing is produced up to that period. Formally, the δ -service level for product k with random demand D_{kt} , deterministic production quantity q_{kt} , and random backlog $BL_{kt} = \max(0; \sum_{\tau=1}^t (D_{k\tau} - q_{k\tau}))$ is computed as

$$\delta_{kt} = 1 - \frac{\text{E}[BL_{kt}]}{\sum_{\tau=1}^t \text{E}[D_{k\tau}]} \quad (5)$$

for a single period t and averaged over T periods as

$$\delta_k = 1 - \frac{\sum_{t=1}^T \text{E}[BL_{kt}]}{\sum_{t=1}^T (T - t + 1) \text{E}[D_{kt}]} \quad (6)$$

The *average* value of the expected backlog of product k is $\frac{1}{T} \sum_{t=1}^T \text{E}[BL_{kt}]$. Furthermore, the average throughput is $\frac{1}{T} \sum_{t=1}^T \text{E}[D_{kt}]$ if we assume that long-run production meets long-run demand. By Little's Law we have for the *average* expected waiting time $\text{E}[W_k]$ of product k

$$\text{E}[W_k] = \frac{\frac{1}{T} \sum_{t=1}^T \text{E}[BL_{kt}]}{\frac{1}{T} \sum_{t=1}^T \text{E}[D_{kt}]} = \frac{\sum_{t=1}^T \text{E}[BL_{kt}]}{\sum_{t=1}^T \text{E}[D_{kt}]} \quad (7)$$

We can use equation (7) to eliminate the expected backlog from (6) to find the equivalent result

$$\delta_k = 1 - \frac{\sum_{t=1}^T \text{E}[D_{kt}]}{\sum_{t=1}^T (T - t + 1) \text{E}[D_{kt}]} \cdot \text{E}[W_k]. \quad (8)$$

which shows that the δ -service level is a linear function of the average expected waiting time $\text{E}[W_k]$. Note that this definition of the δ -service level of product k has a clear interpretation: It represents the expected percentage of the maximum possible demand-weighted waiting time that the customers of product k are protected against. If no demand is back-ordered, no backlog occurs and we obtain $\delta_k = 1$. If nothing is produced and delivered until period T , all the demand is back-ordered and the maximum possible backlog occurs. In this situation we get $\delta_k = 0$. All other cases lie somewhere within the interval $[0; 1]$. Based on this δ -service-level measure, we develop our optimization models in the next sections. Note that it is not possible to transform this δ -service level over the entire planning horizon into a β -service-level constraint per cycle as used by Tempelmeier and Herpers (2010) since this requires the additional knowledge of the setup pattern to compute their service measure.

2.3 Related work

Numerous publications address dynamic lot-sizing problems under a broad variety of differing assumptions. However, the overwhelming majority of these works assumes a deterministic setting as the recent reviews by Karimi et al. (2003), Jans and Degraeve (2008), Robinson et al. (2009) and Buschkühl et al. (2010) showed.

While the literature on discrete time, dynamic stochastic lot sizing appears to be relatively limited, a broad body of literature deals with inventory policies. These policies determine on the one hand (often stationary) safety stocks and on the other hand production (or order) quantities as demand is realized over time, see, e.g., the textbooks by Tempelmeier (2006), Zipkin (2000) and the discussion in the older review by Sox et al. (1999). A more recent survey of the “Stochastic Economic Lot Scheduling Problem” was given by Winands et al. (2011) who stated that “... the finite production capacity has to be *dynamically* distributed among the products in order to be reactive to the stochastic demands, processing and setup times...” (Winands et al. (2011, p. 1)). These “dynamic” and “reactive” elements are the defining features of a production *policy* serving as a general rule how to act in a future and yet uncertain situation. A typical assumption is that demands “... arrive according to stationary and mutually independent stochastic processes...” (Winands et al. (2011, p. 3)). Our work, however, aims at a specific schedule instead of a general policy for the case of both uncertain and non-stationary demand.

Much of the literature on stochastic lot sizing neither considers non-stationary demand distributions, capacity constraints over multiple products nor the attempt to determine production quantities in advance of the demand realization. The resulting inventory policies can easily lead to a nervous planning system and to the notorious bullwhip effect in multi-level systems. In addition, these approaches also do not reflect the popular practice to operate with “frozen” schedules (Zhao et al. (2001)) to achieve at least some degree of planning stability in the presence of dynamic and/or uncertain demand. Their deficiencies have led to some recent publications. Tunc et al. (2011) addressed the cost of using stationary safety stocks in the case of non-stationary demand and Kanet et al. (2010) emphasized the general positive effect of dynamic planned safety stocks. Absi and Kedad-Sidhoum (2009) presented models to determine production quantities for *given* values of maximum lost sales and minimum safety stock levels, but did not endogenously determine the latter values.

In a seminal paper for the case of dynamic uncertain demand, Bookbinder and Tan (1988) described three fundamental strategies for probabilistic lot-sizing problems subject to a service-level constraint. In the “static uncertainty” approach, both the timing and the size of production quantities (for a single product) are determined in advance of the demand realizations. The other extreme is the “dynamic uncertainty” approach which assumes that production decisions for a period are made when precise period demands are finally known. To strike a compromise between the extremes, Bookbinder and Tan (1988) proposed the “static-dynamic uncertainty” approach in which the production epochs are fixed beforehand, thus yielding a stable setup pattern, but production quantities are only determined once the demand has been realized. Customer service is reflected via the α -service level which limits the probability of a stock-out or back-order event.

In this paper, we use an adapted version of the “static uncertainty” approach as a benchmark to evaluate our own schedules. The basic idea of this approach as presented by Bookbinder and Tan is to transform a single-product dynamic lot-sizing problem subject to random demand and an period-specific α -service level constraint into a corresponding deterministic dynamic lot-sizing problem of the Wagner-Whitin type by adjusting the

demand data. In order to determine the adjusted demand data d_t for each period t , Bookbinder and Tan compute the required cumulated production in periods 1 to t for each period t that yields a desired α -service level in period t . The adjusted demand d_t for the deterministic dynamic lot-sizing problem of the Wagner-Whitin type is then simply this required cumulated production until period t minus the required cumulated production until period $t - 1$. They furthermore approximate the expected value of the physical inventory by the expected value of the net inventory position, which can be crude.

Several authors solved stochastic uncapacitated lot sizing problems based on scenario trees. Guan and Miller (2008) presented a dynamic programming approach which determines a solution for the stochastic uncapacitated lot sizing problem in polynomial time. Guan et al. (2006) adapted the well-known valid inequalities proposed by (Pochet and Wolsey 2006, p. 321) and showed that these are also valid for the case of stochastic demand. Based on these results, Di Summa and Wolsey (2008) presented simplified valid inequalities and reformulations. Brandimarte (2006) proposed a model formulation of the stochastic capacitated lot sizing problem that is also based on a scenario tree. The author reformulated this lot sizing problem as a simple-plant location problem and proposed a Fix-and-Relax heuristic.

Tempelmeier and Herpers (2010) have recently argued that it may be desirable to determine both the timing and the size of production orders based on forecasted demand and keep both fixed or “frozen” over the complete planning horizon, i.e., to follow the “static uncertainty approach”. This seems to be particularly important if production schedules have to be coordinated over multiple products due to multi-level production processes or capacity restrictions of the production system. Tempelmeier and Herpers (2010) studied such a stochastic version of the CLSP. They assumed in a specific β -service-level constraint, that during each production cycle only a maximum fraction of the demand of this particular cycle may be back-ordered. In their paper, a modified version of the ABC-heuristic by Maes and van Wassenhove (1986) was used to determine production quantities. In a subsequent paper, Tempelmeier (2011) presented a different solution approach based on column generation to solve the same problem. It significantly outperforms the adapted ABC-heuristic. The research presented in our paper is on the one hand closely related to those presented in Tempelmeier and Herpers (2010) and Tempelmeier (2011). Instead of limiting back-orders via a β -service-level constraint, we aim at limiting backlog and hence take the customer waiting time into account using the δ -service level. Furthermore, we impose a service-level constraint for the entire planning horizon instead of one for each single production cycle. On the other hand, the most closely related papers were presented by Martel et al. (1995) and Sox and Muckstadt (1997) who penalized backlog in the objective function via a penalty cost which did not give them a direct control on the backlog and motivated us to work with a backlog-oriented service level measure.

3 The non-linear stochastic capacitated lot-sizing problem (SCLSP) with a δ -service-level constraint

Based on the assumptions in section 2.1, the δ -service level introduced in section 2.2, and the notation in Table 1, we now state the SCLSP as follows:

Table 1: Notation used for the SCLSP model

Indices and index sets:

\mathcal{K}	set of products ($k \in \{1, \dots, K\}$)
\mathcal{T}	set of periods ($t \in \{1, \dots, T\}$)

Deterministic parameters:

b_t	available capacity in period t
δ_k	minimum δ -service level of product k
hc_k	holding cost of product k per unit and period
M	big number
oc	overtime cost per unit of overtime
sc_k	setup cost of product k
tb_k	production time per unit of product k
ts_k	setup time of product k

Random variables:

BL_{kt}	backlog of product k in period t
D_{kt}	demand of product k in period t
Y_{kt}	net inventory position of product k at the end of period t
YP_{kt}	physical inventory of product k at the end of period t

Decision variables:

o_t	overtime in period t
q_{kt}	production quantity (lot size) of product k in period t
x_{kt}	binary setup variable of product k in period t

SCLSP Model

$$\min Z = \sum_{k \in \mathcal{K}} \sum_{t \in \mathcal{T}} hc_k \cdot \mathbb{E}[YP_{kt}] + \sum_{k \in \mathcal{K}} \sum_{t \in \mathcal{T}} sc_k \cdot x_{kt} + \sum_{t \in \mathcal{T}} oc \cdot o_t \quad (9)$$

subject to:

$$\sum_{k \in \mathcal{K}} (ts_k \cdot x_{kt} + tb_k \cdot q_{kt}) \leq b_t + o_t, \quad \forall t \quad (10)$$

$$q_{kt} - M \cdot x_{kt} \leq 0, \quad \forall k, t \quad (11)$$

$$Y_{k,t-1} + q_{kt} - Y_{kt} = D_{kt}, \quad \forall k, t \quad (12)$$

$$YP_{kt} = \max(0, Y_{kt}), \quad \forall k, t \quad (13)$$

$$BL_{kt} = \max(0, -Y_{kt}), \quad \forall k, t \quad (14)$$

$$\sum_{t \in \mathcal{T}} q_{kt} \geq \sum_{t \in \mathcal{T}} \mathbb{E}[D_{kt}], \quad \forall k \quad (15)$$

$$\sum_{t \in \mathcal{T}} \mathbb{E}[BL_{kt}] \leq (1 - \delta_k) \sum_{t \in \mathcal{T}} (T - t + 1) \mathbb{E}[D_{kt}], \quad \forall k \quad (16)$$

$$q_{kt} \geq 0, \quad \forall k, t \quad (17)$$

$$o_t \geq 0, \quad \forall t \quad (18)$$

$$x_{kt} \in \{0, 1\}, \quad \forall k, t \quad (19)$$

The objective (9) is to minimize the total expected costs of physical inventory, setups, and overtime. Constraints (10) guarantee that the capacity needed for setups and production does not exceed the sum of regular and overtime capacity. Constraints (11) enforce setups in production periods. The inventory balance equations (12) relate the (random) end-of-period inventory position Y_{kt} to the random demand, the production quantity and the inventory position from the previous period. In the following equations, (13) and (14), the physical inventory as well as the backlog are determined. Constraints (15) make sure that at least the expected cumulated demand of product k is produced until the last period T . Furthermore, constraints (16) ensure that the backlog does not exceed the target δ -service level.

In this context of the SCLSP, we can now formally introduce the concept of a dynamic safety stock ss_{kt} of product k in period t . Though the safety stock is not explicitly modeled in the SCLSP, it can be derived from its solution. Assume to this end that an arbitrary solution to the SCLSP is given. This solution does only not specify production quantities q_{kt} , but also production cycles, each starting with a setup period and ending with the last period preceding the next setup. We now define the safety stock ss_{kt} of product k in period t as the difference between the cumulated production quantities and the cumulated expected demand served by the lots up to period t :

$$ss_{kt} = \sum_{\tau=1}^t q_{k\tau} - \sum_{\tau=1}^{\min(T, \min(\theta | (\theta \geq t) \wedge (x_{k, \theta+1} = 1)))} E[D_{k\tau}] \quad \forall k, t \quad (20)$$

Note that the safety stock is constant within a production cycle, but can differ from cycle to cycle. Note furthermore that it can be both positive and negative. In our numerical analysis in section 7 we will use examples to show the impact of different problem parameters on these safety stocks.

Unfortunately, we are not aware of a method available to solve the SCLSP in the non-linear form presented above. For this reason we develop below two different models to approximate the SCLSP. Both can be solved using standard methods for mixed-integer programming.

As a benchmark, we compare the solutions of these approximation models in sections 4 and 5 to those obtained by a modified “static uncertainty approach” as proposed by Bookbinder and Tan (1988), see section 2. Following Bookbinder and Tan (BT), we have to transform the SCLSP into an corresponding deterministic model with specific deterministic demand. To this end, we determine for each product k and period t the necessary cumulated production of product k over periods 1 to t such that a given period-specific δ -service level (5) is met in period t . The deterministic demand in the corresponding BT model for a product in period t is then simply this required cumulated production over periods 1 to t minus the required cumulated production over periods 1 to $t - 1$ leading to the desired δ -service level in the respective periods. This demand data is then used in a deterministic CLSP. As a benchmark, we hence extend the approach by Bookbinder and Tan by

- replacing the period-specific α -service level constraint by a period-specific δ -service level constraint,
- considering a constrained resource and overtime as in constraints (10), and
- adding a constraint (15) on the total production.

For the resulting schedules we then determine the expected value of the backlog and the physical inventory under uncertain demand and finally the costs. As these schedules meet the required δ -service level for each single period, they also always meet the aggregate service level constraint (16). As a matter of fact, they tend to exceed the aggregate constraint which leads to schedules that are more costly than those that match but do not exceed the aggregate service level constraint. As this benchmark approach was inspired by Bookbinder and Tan, we call it CLSP-BT.

4 Linear approximation of the SCLSP via a scenario approach

4.1 Approximating expected values by sample averages using scenarios

The basic idea of the first approximation model is to consider a set of demand scenarios $s \in \mathcal{S}$. Each such scenario s is a random trajectory that represents one of several *equally likely* paths of realizations of the random demand variables D_{kt} for all k and t . In order to achieve a robust production plan that meets the required δ -service level, we optimize the production quantities over the entire set of scenarios $s \in \mathcal{S}$. Random variables D_{kt} , Y_{kt} , YP_{kt} , and BL_{kt} are therefore replaced by their respective scenario-specific realizations d_{kt}^s , y_{kt}^s , yp_{kt}^s , and bl_{kt}^s for $s \in \mathcal{S}$. Furthermore, expected values of these random variables in the SCLSP are approximated by sample averages as follows:

$$E[D_{kt}] \approx \frac{\sum_{s \in \mathcal{S}} d_{kt}^s}{|\mathcal{S}|}, \quad \forall k, t \quad (21)$$

$$E[YP_{kt}] \approx \frac{\sum_{s \in \mathcal{S}} yp_{kt}^s}{|\mathcal{S}|}, \quad \forall k, t \quad (22)$$

$$E[BL_{kt}] \approx \frac{\sum_{s \in \mathcal{S}} bl_{kt}^s}{|\mathcal{S}|}, \quad \forall k, t \quad (23)$$

Other applications of this sample average approximations can be found, e.g., in Bihlmaier et al. (2009), Helber and Henken (2010). A general overview of scenario analysis methods is presented, e.g., in Mißler-Behr (1993), while Freimer et al. (2010) discuss the impact of different sampling methods in linear models using sample average approximations.

4.2 Formulation of the SCLSP-SCN

The modified or additional notation of this second model is shown in Table 2. With this notation, the scenario approximation model SCLSP-SCN can be formulated as follows:

SCLSP-SCN Model

$$\min Z = \sum_{k \in \mathcal{K}} \sum_{t \in \mathcal{T}} hc_k \cdot \left(\frac{\sum_{s \in \mathcal{S}} yp_{kt}^s}{|\mathcal{S}|} \right) + \sum_{k \in \mathcal{K}} \sum_{t \in \mathcal{T}} sc_k \cdot x_{kt} + \sum_{t \in \mathcal{T}} oc \cdot o_t \quad (24)$$

subject to constraints (10), (11), (15), (17), (18), (19) and:

Table 2: Additional or modified notation used for the SCLSP-SCN model

<u>Indices and index sets:</u>	
\mathcal{S}	set of scenarios ($s \in \{1, \dots, S\}$)
<u>Scenario-specific parameters:</u>	
d_{kt}^s	demand of product k in period t and in scenario s
<u>Scenario-specific decision variables:</u>	
bl_{kt}^s	backlog of product k in period t and in scenario s
y_{kt}^s	net inventory of product k at the end of period t and in scenario s
yp_{kt}^s	physical inventory of product k at the end of period t and in scenario s

$$y_{k,t-1}^s + q_{kt} - y_{kt}^s = d_{kt}^s, \quad \forall k, t, s \quad (25)$$

$$yp_{k,t}^s \geq y_{k,t}^s, \quad \forall k, t, s \quad (26)$$

$$bl_{k,t}^s \geq -y_{k,t}^s, \quad \forall k, t, s \quad (27)$$

$$\sum_{t \in \mathcal{T}} \frac{\sum_{s \in \mathcal{S}} bl_{kt}^s}{|\mathcal{S}|} \leq (1 - \delta_k) \sum_{t \in \mathcal{T}} (T - t + 1) E[D_{kt}], \quad \forall k \quad (28)$$

$$yp_{kt}^s, bl_{kt}^s \geq 0, \quad \forall k, t, s \quad (29)$$

All the constraints of the SCLSP in section 3 that contain random variables (12, 13, and 14) are replaced by $|\mathcal{S}|$ systems of constraints (25)-(27) reflecting independent scenario-specific realizations of the random variables. Instead of the expected value of the physical inventory in the objective function (9) and the expected value of backlog in the δ -service-level restriction (16), we now use sample averages in the objective function (24) and the service-level constraint (28) as explained in section 4.1. Note that in order to minimize costs, both backlog and physical inventory will tend to be minimized. For this reason, it is possible to replace the non-linear maximum function in Equations (13) and (14) by the new constraints (26) and (27).

If the standard deviation of the demand D_{kt} equals 0 for all k and t , a single scenario represents the entire knowledge about the demand. As a result we obtain a version of the deterministic CLSP in which backlogs are allowed. These backlogs are restricted using the δ -service-level constraint. If we additionally assume the case of a desired δ -service level of 100%, the SCLSP-SCN reduces to be the well-known CLSP.

Note that this scenario approach offers a substantial degree of flexibility. In particular, it is possible to construct scenarios that consistently model probabilistic dependencies of the demands for different products or periods. This is a potentially very useful feature of this approximation model.

The realizations of the demand data d_{kt}^s for product k in period t can be determined independently for each single scenario s based on the distribution of the random variable D_{kt} . However, this so-called Simple Random Sampling (SRS) tends to lead to a substantial variance of any function computed based on such a sample. As a variance-reducing technique, Saliby (1990) proposed the so-called Descriptive Sampling (DS) which he showed to work extremely well for the Newsboy problem, a problem related to the SCLSP. Descriptive Sampling requires the ex-ante knowledge of the sample size $|\mathcal{S}|$, which is always exactly known in our context. The basic idea is to deterministically construct a set of

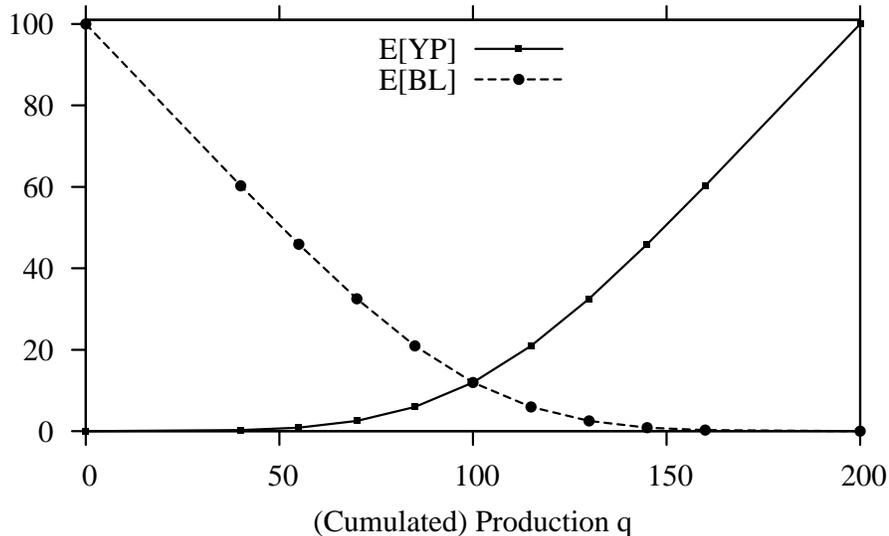


Figure 1: Linearized expected backlog and physical inventory over production

$|\mathcal{S}|$ values d_{kt}^s such that their “empirical distribution” follows the underlying theoretical distribution of the random variable D_{kt} as close as possible. Assume, e.g., that for a random variable Z with probability function $F_Z(z)$ a set of size 10 is required. Then the set elements are determined by inverting the probability function as $z_i = F_Z^{-1}((i - 0.5)/10)$ for $i = 1, \dots, 10$ and finally randomly “shuffled” to induce a random sequence. We implemented the SCLSP-SCN with both Simple Random Sampling and Descriptive Sampling to analyze the impact of the different sampling methods.

5 Piecewise Linear Approximation of the SCLSP

5.1 Approximating expected values via piecewise linear functions

It is possible to replace the non-linear functions of expected backlog and physical inventory in the SCLSP by piecewise linear functions. Such a piecewise linearization is shown in Figure 1 for the case of a single period with normally distributed demand, expected demand of 100 and variance of 900. (In the single-period case, the period production equals the *cumulated* production.) The graph shows that in order to linearize these functions, a point for the minimal (zero) production is required ($q = 0$) and one for the realistic maximum possible production (e.g., $q = 200$). As we do not have any a priori knowledge about the optimal cumulated production for any given period, the other supporting points of the linearization are concentrated as shown in Figure 1 in the area where the non-linearity of the expected backlog and physical inventory (over the production quantity) is strongest, i.e., around the expected demand, so that the overall deviation from the original non-linear functions is small.

Such a linearization is required for each combination of product k and period t for both the expected physical inventory and the expected backlog as functions of the cumulated production up to period t . Let L denote the number of segments of the continuous and piecewise linear functions approximating the non-linear functions of expected inventory or backlog. In the example in Figure 2, the expected inventory is approximated by a piecewise linear function with four segments, i.e., $L = 4$. Each such segment l is characterized by

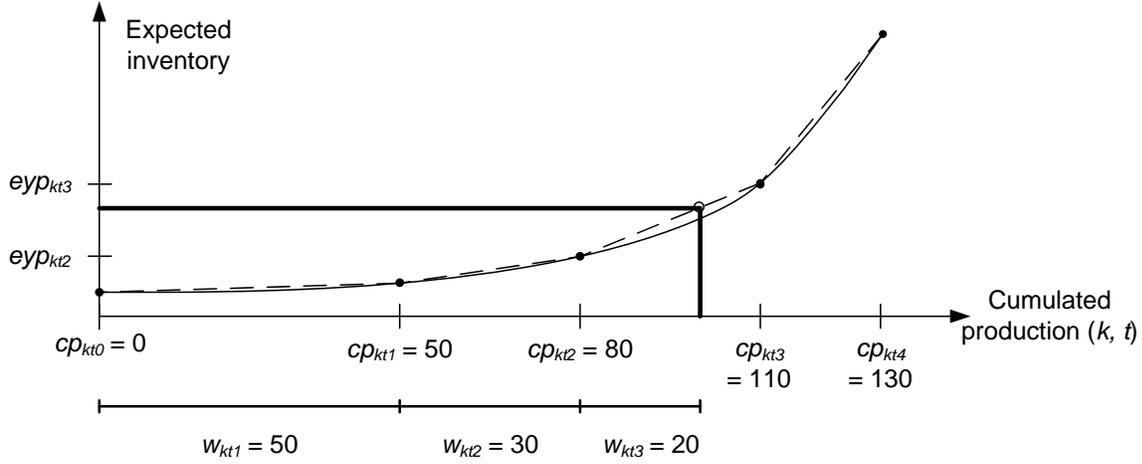


Figure 2: Linearization of the expected physical inventory function

an upper limit on the cumulated production cp_{ktl} as shown in Figure 2. Assume that l^* is the index of the segment that corresponds to the cumulated production of product k in period t , i.e.,

$$cp_{kt,l^*-1} \leq \sum_{\tau=1}^t q_{k\tau} \leq cp_{kt,l^*}. \quad (30)$$

In the example in Figure 2, we find $l^* = 3$ and $cp_{kt,l^*-1} = 80 \leq \sum_{\tau=1}^t q_{k\tau} = 100 \leq cp_{kt,l^*} = 110$. Now we define w_{ktl} as the part of the cumulated production quantity related to segment $l = 1, \dots, L$ of the linearization. Then the following two equations must hold:

$$\sum_{\tau=1}^t q_{k\tau} = \sum_{l=1}^{l^*} w_{ktl} \quad (31)$$

$$\sum_{l=l^*+1}^L w_{ktl} = 0 \quad (32)$$

The segments 1 to $l^* - 1$ are filled to their (maximum) “capacity” w_{ktl} and pieces $l^* + 1$ to L must be zero:

$$w_{ktl} = cp_{ktl} - cp_{kt,l-1}, \quad l = 1, \dots, l^* - 1 \quad (33)$$

$$w_{ktl} = \sum_{\tau=1}^t q_{k\tau} - cp_{kt,l-1}, \quad l = l^* \quad (34)$$

$$w_{ktl} = 0, \quad l = l^* + 1, \dots, L \quad (35)$$

In the example in Figure 2, we hence have $w_{kt1} = 50$, $w_{kt2} = 30$, $w_{kt3} = 20$, and $w_{kt4} = 0$.

Note that the part of the cumulated production quantity w_{ktl} that is assigned to segment l of the linearization must not exceed the difference between the limits of cumulated production related to segment l , i.e.,

$$w_{ktl} \leq cp_{ktl} - cp_{kt,l-1}, \quad \forall k, t, l \quad (36)$$

In order to determine the supporting points for the approximation of $E[BL_{kt}]$ and $E[YP_{kt}]$, it is necessary to determine for each value of cumulated production cp_{ktl} the corresponding value of the expected backlog

$$ebl_{ktl} = E \left[\max \left(0; \sum_{\tau=1}^t D_{k\tau} - cp_{ktl} \right) \right] \quad (37)$$

and the corresponding physical inventory:

$$eyp_{ktl} = E \left[\max \left(0; cp_{ktl} - \sum_{\tau=1}^t D_{k\tau} \right) \right] \quad (38)$$

We now consider the special case that the period demand D_{kt} is normally distributed with mean $E[D_{kt}]$ and variance $\text{VAR}[D_{kt}]$ estimated by a forecasting system. We furthermore assume that all demands are mutually independent, such that there is neither cross- nor autocorrelation. The cumulated demand $Z_{kt} = \sum_{\tau=1}^t D_{k\tau}$ has the mean $\mu_{Z_{kt}} = \sum_{\tau=1}^t E[D_{k\tau}]$. As the demand is period-wise independent, we furthermore determine the standard deviation $\sigma_{Z_{kt}} = \sqrt{\sum_{\tau=1}^t \text{VAR}[D_{k\tau}]}$. Let $\phi(x)$ and $\Phi(x)$ denote the density function and the probability function of a random variable X following a standard normal distribution. Then the expected backlog and also the physical inventory corresponding to a specific (cumulated) production cp_{ktl} can be computed via the “first-order loss function”, see Tempelmeier (2006, p. 292).

$$ebl_{ktl} = E[BL_{ktl}] = \sigma_{Z_{kt}} \cdot \left(\phi \left(\frac{cp_{ktl} - \mu_{Z_{kt}}}{\sigma_{Z_{kt}}} \right) - \frac{cp_{ktl} - \mu_{Z_{kt}}}{\sigma_{Z_{kt}}} \left[1 - \Phi \left(\frac{cp_{ktl} - \mu_{Z_{kt}}}{\sigma_{Z_{kt}}} \right) \right] \right) \quad (39)$$

$$eyp_{ktl} = E[YP_{ktl}] = \sigma_{Z_{kt}} \cdot \left(\phi \left(\frac{cp_{ktl} - \mu_{Z_{kt}}}{\sigma_{Z_{kt}}} \right) + \frac{cp_{ktl} - \mu_{Z_{kt}}}{\sigma_{Z_{kt}}} \left[1 - \Phi \left(-\frac{cp_{ktl} - \mu_{Z_{kt}}}{\sigma_{Z_{kt}}} \right) \right] \right) \quad (40)$$

Let ebl_{kt0} and eyp_{kt0} denote the respective values corresponding to zero production. Hence $L + 1$ different pairs (cp_{ktl}, ebl_{ktl}) and (cp_{ktl}, eyp_{ktl}) define the supporting points used to approximate the non-linear functions of expected backlog and inventory over (cumulated) production. Then the non-linear functions can be approximated as follows:

$$E[BL_{kt}] \approx ebl_{kt0} + \sum_{l=1}^L \frac{ebl_{ktl} - ebl_{kt,l-1}}{cp_{ktl} - cp_{kt,l-1}} \cdot w_{ktl} \quad (41)$$

$$E[YP_{kt}] \approx eyp_{kt0} + \sum_{l=1}^L \frac{eyp_{ktl} - eyp_{kt,l-1}}{cp_{ktl} - cp_{kt,l-1}} \cdot w_{ktl} \quad (42)$$

However, we have to ensure that the cumulated production in periods 1 to t is assigned to the correct pieces of the linearization as required in Equations (33) to (35) and shown in the example in Figure 2.

If the period demand D_{kt} follows a distribution type other than the normal distribution, the values of ebl_{ktl} and eyp_{ktl} in Equations (37) and (38) have to be determined based on the specific features of the respective demand distribution. The general principle of the piecewise linearization, however, still applies.

Table 3: Additional notation for SCLSP-PLA model

Indices and index sets:

\mathcal{L} set of linearization segments ($l \in \{1, \dots, L\}$)

Deterministic parameters:

abl_{ktl} expected backlog of product k in period t for segment l
 abl_{kt0} expected backlog of product k in period t for zero cumulated production
 cp_{ktl} cumulated production of product k in period t for segment l
 eyp_{ktl} expected physical inventory of product k in period t for segment l
 eyp_{kt0} expected physical inventory of product k in period t for zero cumulated production

Decision variables:

w_{ktl} cumulated production quantity of product k in periods 1 to t assigned to stage l of the linearization

5.2 Formulation of the SCLSP-PLA

With the additional notation in Table 3, the piecewise linear approximation SCLSP-PLA of the SCLSP can be defined as follows:

SCLSP-PLA Model

$$\begin{aligned} \min Z = & \sum_{k \in \mathcal{K}} \sum_{t \in \mathcal{T}} hc_k \cdot \left(eyp_{kt0} + \sum_{l=1}^L \frac{eyp_{ktl} - eyp_{kt,l-1}}{cp_{ktl} - cp_{kt,l-1}} \cdot w_{ktl} \right) \\ & + \sum_{k \in \mathcal{K}} \sum_{t \in \mathcal{T}} sc_k \cdot x_{kt} + \sum_{t \in \mathcal{T}} oc \cdot o_t \end{aligned} \quad (43)$$

subject to constraints (10), (11), (15), (17), (18), (19) and:

$$q_{kt} = \sum_{l=1}^L w_{ktl} - \sum_{l=1}^L w_{k,t-1,l}, \quad \forall k, t \quad (44)$$

$$w_{ktl} \leq cp_{ktl} - cp_{kt,l-1}, \quad \forall k, t, l \quad (45)$$

$$\begin{aligned} \sum_{t \in \mathcal{T}} \left(abl_{kt0} + \sum_{l=1}^L \frac{abl_{ktl} - abl_{kt,l-1}}{cp_{ktl} - cp_{kt,l-1}} \cdot w_{ktl} \right) \leq \\ (1 - \delta_k) \sum_{t \in \mathcal{T}} (T - t + 1) E[D_{kt}], \quad \forall k \end{aligned} \quad (46)$$

$$w_{ktl} \geq 0, \quad \forall k, t, l \quad (47)$$

In the objective function (43) and the δ -service-level constraint (46) the linearized versions of the expected values of physical inventory and backlog are used. The production quantity per period equals the difference of cumulated production of subsequent periods, see (44). Constraint (45) was already introduced as (36). Finally, the production quantity assigned to segment l of the linearization must be non-negative, see (47).

It should be noted that due to the direction of optimization and the structure of the service-level constraint, the conditions in Equations (33) to (35) are always met! I.e., it is not necessary to define additional binary variables to ensure that cumulated production quantities w_{ktl} are properly determined.

6 A Fix-and-Optimize algorithm to solve the two approximation models

As the deterministic CLSP is already \mathcal{NP} -hard (see Florian et al. (1980)), it tends to be difficult to solve the proposed approximations of the SCLSP to optimality in reasonable time. Therefore, we use an adapted version of the flexible Fix-and-Optimize heuristic proposed by Helber and Sahling (2010) as it leads to high-quality solutions for the multi-level version of the deterministic CLSP, see also Sahling et al. (2009) and Sahling (2010).

In the Fix-and-Optimize heuristic, a sequence of subproblems is solved to (sub-)optimality in an iterative fashion. In such a subproblem, the optimal values for a small number of binary setup variables are determined while the remaining binary setup variables are fixed to a given setup pattern from the solution of a previously solved subproblem. Furthermore, all real-valued decision variables are optimized in each subproblem as well. Due to the reduced number of binary setup variables, the time effort to solve the subproblem is moderate.

To apply the Fix-and-Optimize heuristic on the proposed approximations, we initially define the set $\mathcal{KT} = \{1, \dots, K\} \times \{1, \dots, T\}$ of all product-period combinations (k, t) . Subsequently, the subset $\mathcal{KT}^{fix} \subseteq \mathcal{KT}$ can be defined. This subset \mathcal{KT}^{fix} contains those product-period combinations whose respective binary setup variables are fixed to a given setup pattern \bar{x}_{kt} from the solution of a previous subproblem. Given the subset \mathcal{KT}^{fix} , a subproblem can be defined by adding the following constraints to the SCLSP-SCN and the SCLSP-PLA as well as the benchmark CLSP-BT

$$x_{kt} = \bar{x}_{kt}, \quad \forall (k, t) \in \mathcal{KT}^{fix}. \quad (48)$$

These additional constraints (48) assign a fixed setup pattern \bar{x}_{kt} to those binary setup variables whose product-period combination (k, t) belongs to the set \mathcal{KT}^{fix} . Hence, the optimization of the binary setup variables is limited to those product-period combinations in the set $\mathcal{KT}^{opt} = \mathcal{KT} \setminus \mathcal{KT}^{fix}$. We typically have $|\mathcal{KT}^{opt}| \ll |\mathcal{KT}^{fix}|$. As an initial solution, we start similarly to Helber and Sahling (2010) with a lot-for-lot setup pattern, i.e. all binary setup variables x_{kt} are fixed to 1.

We use a product-oriented decomposition so that each subproblem corresponds to a single product. All binary setup variables of the respective product are optimized within the subproblem, while the others are fixed. Starting with the first product all products are considered once per iteration in the increasing sequence of the product index. The Fix-and-Optimize heuristic stops when a local optimum is reached, i.e., the last iteration yields no further improvement.

7 Numerical results

7.1 Accuracy of the different approximations

7.1.1 Description of the test instances

For our numerical investigation, we defined 1,296 (artificial) test instances by systematically varying different parameters, e.g., the number of products and the number of periods. Table 4 gives an overview of the parameter values; details are presented in Appendix A.

Table 4: Parameters of the test instances

Number of products	K	=	{5, 10, 20}
Number of periods	T	=	{5, 10, 20}
Inter-period coefficient of variation of expected demand	VC^{ip}	=	{0.2, 0.3}
Demand coefficient of variation	VC^d	=	{0.1, 0.3}
Time between orders	TBO	=	{1, 2, 4}
Utilization of resource due to processing	$Util$	=	{0.6, 0.75}
Setup time as fraction of period processing time	ts^{rel}	=	{0.0, 0.25}
Service-level goal	δ	=	{0.8, 0.9, 0.95}
Number of scenarios (SCLSP-SCN only)	S	=	{10, 30, 50}

In order to create demand data that is both dynamic and stochastic, we first specified average demands $E[D_k]$ for 20 different products k as reported in Table 11 of the Appendix. To create *dynamic* time series of expected demand $E[D_{kt}]$ per period t around the mean $E[D_k]$, we then drew (pseudo) random numbers from normal distributions with mean $E[D_k]$ and standard deviation $VC^{ip} \cdot E[D_k]$. In the case of the inter-period coefficient of demand $VC^{ip} = 0.2$ this led to the dynamic time series with moderate time variability reported in Table 12. The more volatile time series for $VC^{ip} = 0.3$ is given in Table 13.

It is important to distinguish this inter-period variability of the *expected* demand $E[D_{kt}]$ over the different periods $t = 1, \dots, T$ from the randomness of the demand D_{kt} for a given (k, t) combination. In our models in sections 3 to 5.2, we require that the probability distribution of D_{kt} is given for each (k, t) combination. For our numerical study, we now assume that all random demands D_{kt} follow a normal distribution with mean $E[D_{kt}]$. We furthermore assume that the standard deviation of demand is time-invariant, proportional to the average expected demand

$$\overline{E[D_k]} = \frac{\sum_{t=1}^T E[D_{kt}]}{T} \quad (49)$$

and proportional to the demand coefficient of variation VC^d , i.e.,

$$\sigma_{kt} = \sigma_k = \overline{E[D_k]} \cdot VC^d. \quad (50)$$

We used this approach to model the usage of a properly designed demand forecasting system yielding point forecasts of the expected demand $E[D_{kt}]$ and (as the most simple case) normally distributed forecasting errors with zero mean and time-invariant standard deviation, i.e., homoscedasticity of the forecasting residuals. Figure 3 shows examples of dynamic stochastic demand time series with time-invariant variance. The case with $VC^d = 0.1$ ($VC^d = 0.3$) models the use of a more (less) precise forecasting system, respectively.

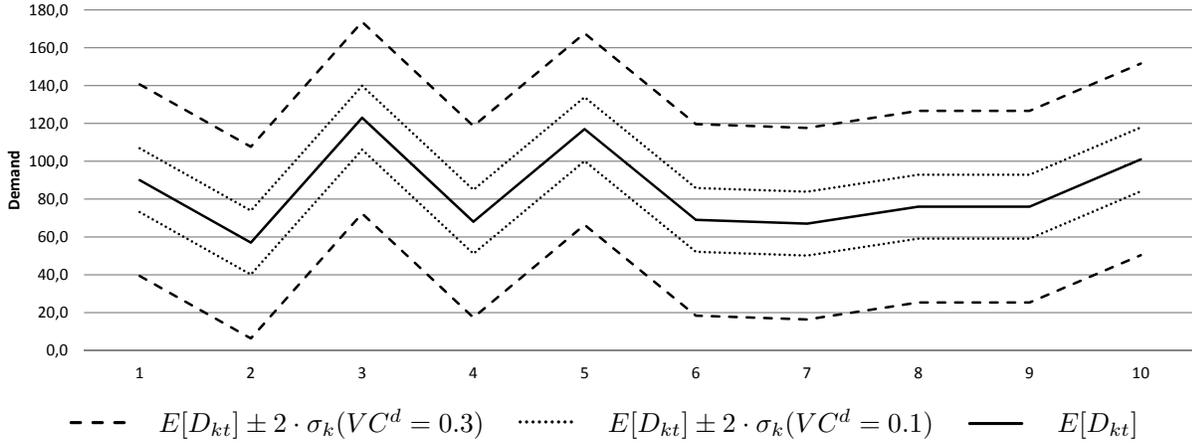


Figure 3: Demand of a product with lower and upper 95% intervals for smaller ($VC^d = 0.1$) and higher ($VC^d = 0.3$) demand variability

To determine the scenario-specific demand d_{kt}^s of product k in period t and scenario s in the SCLSP-SCN in section 4, realizations were drawn from a normal distribution with expected value $E[D_{kt}]$ and standard deviation σ_k as specified in Equation (50) using either SRS or DS, see section 4.

7.1.2 Analysis of the numerical results

In our numerical experiments, we used an Intel Xeon CPU with 2.93 GHz and 8 GB of RAM. The Fix-and-Optimize heuristic was implemented in the 64-bit version of GAMS 23.6. Each subproblem was solved by CPLEX 12.2 using two parallel threads. Because of the potentially large number of real-valued decision variables in both approximation models (number of scenarios in the SCLSP-SCN and linearization segments in the SCLSP-PLA), each subproblem optimization was aborted as soon as an integrality gap of 0.5% of the solution was reached or a time limit of 30 seconds was exceeded. If CPLEX did not find a mathematically feasible solution to the subproblem within this time limit, the time limit was extended to 300 seconds to find a solution. In the case of the SCLSP-PLA, we operated with 10 linearization segments which are specific for each product-period combination (k, t) and distributed as shown in Figure 1 to relate expected physical inventory and backlog to *cumulated* production.

In Table 5, we illustrate the average results over all 1,296 test instances for the scenario approach SCLSP-SCN based on Simple Random Sampling (SRS) and Descriptive Sampling (DS), the piecewise linear approximation model SCLSP-PLA and, as a benchmark, the modified Bookbinder and Tan approach CLSP-BT as described at the end of section 3. In the first row we address the service level (SL) and report the percentage of problem instances for which the required δ -service level is met for all products in this instance. Remember that for *given* production quantities, the expected backlog and hence also the δ -service level can easily be determined, either exactly (e.g., in our case for normally distributed demand) or by approximating any theoretical demand distribution by an empirical distribution. In the next two rows, we report the fraction of problem instances for which all products in this instance missed the required service level by at most 1 or 2 percentage points, respectively. E.g., if the SCLSP-SCN based on Simple Random Sampling (SRS) was solved using 10 scenarios, 61.5% of the problem instances had solutions which missed the required δ -service level for all products by at most 2 per-

centage points, whereas this holds true for 100 % of the problem instances if Descriptive Sampling is used. The results indicate that DS leads to a substantial variance reduction, here with respect to the aggregate service level.

Note that the SCLSP-PLA as well as the CLSP-BT always met the service-level target for all products, while the SCLSP-SCN almost always failed to meet the target exactly due to the variance of the estimated δ -service level. In the next two rows we compare only those solutions to test instances for which all products met the required δ -service-level target and first report the percentage of best solutions “BSol” determined by the respective model. The percentage of best solutions found by the SCLSP-PLA is exorbitantly high with 90.43% compared to values between 0.08% and 2.78% for the SCLSP-SCN and 4.86 % for the CLSP-BT. This is due to the fact that we only consider those solutions which meet the required δ -service level for *all* products in the instance.

The average deviation “DevBSol” from the best known solution is almost always very small, except for the benchmark CLSP-BT which delivers a very high deviation. The reason for this very high deviation is that a constraint on the δ -service level for each period is imposed in the CLSP-BT, see equation (5). The SCLSP, however, operates with an aggregate service level constraint only. The period-specific constraint can lead to a substantial increase of the number of setups and hence the setup costs. The final three rows “SLP” to “SLP-2%” report over all instances the average of the percentage of the number of products for which the δ -service level was met or missed by at most 1% or 2%, respectively. On average, 80.44% of the products in the solutions of the SCLSP-SCN: SRS for 10 scenarios either met the required δ -service level or missed it by at most 1 percentage point. The comparison of between SRS and DS shows again the strong variance reduction due to DS. With respect to the fraction of products that missed the target service level by at most 1 percentage point, DS with only 10 scenarios led to a fraction of 98.46% and hence even outperformed SRS with 50 scenarios and a fraction of only 95.07%. The overall numerical results indicate that the SCLSP-PLA is clearly superior to the SCLSP-SCN, but the latter still delivers solutions that can be considered to be robust from a practical point of view.

Table 5: Overall numerical results for all test instances

	SCLSP-SCN: SRS			SCLSP-SCN: DS			SCLSP- PLA	CLSP- BT
	$ S = 10$	$ S = 30$	$ S = 50$	$ S = 10$	$ S = 30$	$ S = 50$		
SL [%]	0.77	4.71	1.08	1.77	1.47	0.77	100.00	100.00
SL-1% [%]	35.26	64.74	72.53	87.65	100.00	99.85	100.00	100.00
SL-2% [%]	61.50	88.12	98.15	100.00	100.00	100.00	100.00	100.00
BSol [%]	0.08	2.78	0.46	0.85	0.39	0.15	90.43	4.86
DevBSol [%]	2.51	0.58	1.78	1.23	1.09	0.76	1.24	105.72
SLP [%]	49.24	60.10	48.56	47.31	48.86	42.92	100.00	100.00
SLP-1% [%]	80.44	93.28	95.07	98.46	100.00	99.99	100.00	100.00
SLP-2% [%]	92.92	98.46	99.88	100.00	100.00	100.00	100.00	100.00

As the DS generates much better results than the SRS, we only report further results for DS which explains why percentages in the following tables do not necessarily add up to 100%. In Table 6, a detailed analysis with respect to the number of products and periods is given. An entry “***” indicates that a quantity could not be computed as not a single instance met the service level constraint for all products. For the SCLSP-SCN, the solution quality tends to be higher for smaller test instances. In the case of the SCLSP-SCN, the fraction of solutions which met the required δ -service level for *all* products decreases with the number of products and periods. Only for at most 3.7% of

Table 6: Numerical results for all test instances based on the number of products and periods

		$K = 5$	$K = 10$	$K = 20$	$T = 5$	$T = 10$	$T = 20$
SL [%]	DS, $ S = 10$	3.70	1.39	0.23	2.55	1.16	1.62
	DS, $ S = 30$	2.55	1.62	0.23	2.55	0.69	1.16
	DS, $ S = 50$	2.08	0.23	0.00	1.39	0.00	0.93
	SCLSP-PLA	100.00	100.00	100.00	100.00	100.00	100.00
	CLSP-BT	100.00	100.00	100.00	100.00	100.00	100.00
SL-1% [%]	DS, $ S = 10$	92.82	92.59	77.55	74.54	90.74	97.69
	DS, $ S = 30$	100.00	100.00	100.00	100.00	100.00	100.00
	DS, $ S = 50$	100.00	100.00	99.54	99.54	100.00	100.00
	SCLSP-PLA	100.00	100.00	100.00	100.00	100.00	100.00
	CLSP-BT	100.00	100.00	100.00	100.00	100.00	100.00
BSol [%]	DS, $ S = 10$	1.62	0.69	0.23	1.62	0.69	0.23
	DS, $ S = 30$	0.69	0.23	0.23	0.93	0.23	0.00
	DS, $ S = 50$	0.46	0.00	***	0.00	***	0.46
	SCLSP-PLA	85.19	91.44	94.68	83.80	94.44	93.06
	CLSP-BT	4.86	5.09	4.63	6.25	4.17	4.17
DevBSol [%]	DS, $ S = 10$	1.32	1.18	0.00	0.50	1.21	2.38
	DS, $ S = 30$	1.13	1.18	0.00	1.03	0.58	1.55
	DS, $ S = 50$	0.70	1.25	***	0.66	***	0.90
	SCLSP-PLA	1.30	1.21	1.20	1.95	1.28	0.47
	CLSP-BT	103.20	104.73	109.23	52.75	92.56	171.86
SLP [%]	DS, $ S = 10$	53.80	46.99	41.13	48.45	47.27	46.20
	DS, $ S = 30$	45.32	56.69	44.57	54.02	47.40	45.17
	DS, $ S = 50$	43.15	43.63	41.99	47.07	38.37	43.33
	SCLSP-PLA	100.00	100.00	100.00	100.00	100.00	100.00
	CLSP-BT	100.00	100.00	100.00	100.00	100.00	100.00
SLP-1% [%]	DS, $ S = 10$	98.56	98.84	97.97	96.23	99.27	99.88
	DS, $ S = 30$	100.00	100.00	100.00	100.00	100.00	100.00
	DS, $ S = 50$	100.00	100.00	99.98	99.98	100.00	100.00
	SCLSP-PLA	100.00	100.00	100.00	100.00	100.00	100.00
	CLSP-BT	100.00	100.00	100.00	100.00	100.00	100.00

the test instances with $K = 5$, a solution that is feasible for all products has been found. This portion decreases to 0.23% for test instances with 20 products. The influence of the number of periods is less distinct. From the perspective of individual products, it is interesting that for at least 40% of the products the required service-level has been reached. If it has not been reached, it has usually been missed by at most one percentage point. Again the huge average deviation “DevBSol” of the CLSP-BT is eye-catching, in particular for larger number of periods.

In Table 7, numerical results are presented with respect to the variability of demand (reflecting the forecast error) and the target δ -service level. While the SCLSP-PLA again performs extremely well under any conditions, the solution quality of the SCLSP-SCN seems to deteriorate as the demand variability or the target service-level increase. During our numerical investigation, we realized that the influence of the remaining parameters on the quality of the solutions is limited.

Table 8 reports the average run time of the Fix-and-Optimize algorithm for the different model variants. The solution times tend to increase with the number of products, periods, and—in case of the SCLSP-SCN—with the number of scenarios. Furthermore, the piecewise linear approximation model SCLSP-PLA appears to outperform the SCLSP-SCN from the computational perspective as well. For example, a test instance with 5 products and 20 periods can be solved using the SCLSP-PLA in 16 seconds on average. By contrast, the solution of the SCLSP-SCN with 50 scenarios and DS requires on average 185 seconds. The very fast solutions times of the CLSP-BT are attended by the costly results discussed earlier.

Based on this analysis, the reader might be tempted to conclude that the SCLSP-SCN is essentially useless. However, it should be noted that the scenario approach offers some kind of flexibility that does not exist in the other approximation model SCLSP-PLA: It is possible to model within a scenario s of the SCLSP-SCN a probabilistic dependency of the demand for two different product-period combinations (k, t) and (\hat{k}, \hat{t}) , respectively. In our problem specification in section 2.1, we explicitly ruled this possibility out, but it might be relevant in a practical setting and we do not see a way to incorporate it in the SCLSP-PLA.

7.2 Impact of forecast (in-)accuracy on costs and planned dynamic safety stocks

From the planner’s perspective, the randomness of demand is reflected in the errors of the forecasts from the forecasting system. A large standard deviation σ_k of the demand of product k leads to a large standard error of the demand forecast. For this reason, we now study the impact of demand variability (as the source of forecasting inaccuracy) on costs and safety stocks.

To this end, we consider from Table 4 the six cases with $K = 5$, $T = 10$, $VC^{ip} = 0.3$, $Util = 0.75$, $ts^{rel} = 0.25$, and now a target δ -service level of either 95% or 99%. Each test instance was solved to optimality using the SCLSP-PLA. The total costs of each solution are given for each combination of TBO and VC^d in Tables 9 and 10. In both tables, we report the relative cost increase due to the increased demand variability (or forecast inaccuracy).

As the setup costs (and hence the TBO) increase, the total costs also increase. They furthermore rise with the target δ -service level or the demand variability. The relative growth of the costs due to a higher demand variability is much stronger for low setup costs ($TBO = 1$) than for high setup costs ($TBO = 4$). The phenomenon is particularly

Table 7: Solution quality for all test instances based on VC^d and the given δ -service level

		$VC^d = 0.1$	$VC^d = 0.3$	$\delta = 0.80$	$\delta = 0.90$	$\delta = 0.95$
SL [%]	DS, $ S = 10$	2.62	0.93	4.86	0.46	0.00
	DS, $ S = 30$	2.16	0.77	4.17	0.23	0.00
	DS, $ S = 50$	1.23	0.31	1.62	0.69	0.00
	SCLSP-PLA	100.00	100.00	100.00	100.00	100.00
	CLSP-BT	100.00	100.00	100.00	100.00	100.00
SL-1% [%]	DS, $ S = 10$	100.00	75.31	94.91	85.42	82.64
	DS, $ S = 30$	100.00	100.00	100.00	100.00	100.00
	DS, $ S = 50$	100.00	99.69	100.00	100.00	99.54
	SCLSP-PLA	100.00	100.00	100.00	100.00	100.00
	CLSP-BT	100.00	100.00	100.00	100.00	100.00
BSol [%]	DS, $ S = 10$	1.08	0.62	2.55	0.00	***
	DS, $ S = 30$	0.62	0.15	0.93	0.23	***
	DS, $ S = 50$	0.15	0.15	0.46	0.00	***
	SCLSP-PLA	93.98	86.88	88.43	97.45	85.42
	CLSP-BT	0.00	9.72	0.00	0.00	14.58
DevBSol [%]	DS, $ S = 10$	0.88	2.22	1.28	0.68	***
	DS, $ S = 30$	0.84	1.81	1.15	0.00	***
	DS, $ S = 50$	0.83	0.49	0.85	0.55	***
	SCLSP-PLA	0.07	2.40	0.21	0.02	3.48
	CLSP-BT	121.56	89.88	190.96	81.74	44.47
SLP [%]	DS, $ S = 10$	51.47	43.15	57.34	44.16	40.43
	DS, $ S = 30$	49.72	48.00	55.69	46.10	44.79
	DS, $ S = 50$	44.39	41.46	50.34	40.93	37.51
	SCLSP-PLA	100.00	100.00	100.00	100.00	100.00
	CLSP-BT	100.00	100.00	100.00	100.00	100.00
SLP-1% [%]	DS, $ S = 10$	100.00	96.92	99.54	98.10	97.74
	DS, $ S = 30$	100.00	100.00	100.00	100.00	100.00
	DS, $ S = 50$	100.00	99.98	100.00	100.00	99.98
	SCLSP-PLA	100.00	100.00	100.00	100.00	100.00
	CLSP-BT	100.00	100.00	100.00	100.00	100.00

Table 8: Average solution time (in CPU seconds) for all test instances

		$T = 5$	$T = 10$	$T = 20$
$K = 5$	DS, $ S = 10$ SCN: SRS	1,08	4,32	19,83
	DS, $ S = 30$ SCN: SRS	3,64	21,00	105,66
	DS, $ S = 50$ SCN: SRS	7,58	46,00	262,73
	DS, $ S = 10$ SCN: DS	1,07	4,47	19,41
	DS, $ S = 30$ SCN: DS	3,64	20,00	103,44
	DS, $ S = 50$ SCN: DS	7,34	44,01	246,88
	SCLSP-PLA	0,51	1,90	16,05
	CLSP-BT	0,20	0,34	0,61
$K = 10$	DS, $ S = 10$ SCN: SRS	13,82	81,18	575,25
	DS, $ S = 30$ SCN: SRS	3,77	17,71	112,13
	DS, $ S = 50$ SCN: SRS	30,10	186,80	1381,87
	DS, $ S = 10$ SCN: DS	3,74	18,17	102,06
	DS, $ S = 30$ SCN: DS	14,70	83,80	577,03
	DS, $ S = 50$ SCN: DS	30,92	190,51	1401,26
	SCLSP-PLA	1,64	6,61	85,95
	CLSP-BT	0,57	0,87	2,41
$K = 20$	DS, $ S = 10$ SCN: SRS	15,67	70,80	390,30
	DS, $ S = 30$ SCN: SRS	71,93	367,86	2156,08
	DS, $ S = 50$ SCN: SRS	174,84	859,91	5338,12
	DS, $ S = 10$ SCN: DS	16,72	69,55	372,68
	DS, $ S = 30$ SCN: DS	76,07	388,38	2265,90
	DS, $ S = 50$ SCN: DS	185,06	850,47	4910,22
	SCLSP-PLA	5,04	24,60	310,96
	CLSP-BT	1,20	3,29	5,90

Table 9: Total costs of test instances ($\delta = 0.95$)

	$VC^d = 0.1$	$VC^d = 0.3$	Rel. Increase
$TBO = 1$	1806.47	2969.09	64.36 %
$TBO = 2$	5066.85	6007.12	18.56 %
$TBO = 4$	13008.63	14265.49	9.66 %

Table 10: Total costs of test instances ($\delta = 0.99$)

	$VC^d = 0.1$	$VC^d = 0.3$	Rel. Increase
$TBO = 1$	2758.99	12343.53	347.39 %
$TBO = 2$	7027.14	17134.86	143.84 %
$TBO = 4$	19548.03	33527.46	71.51 %

dramatic if a very high service level is demanded ($\delta = 99\%$). This result shows that as setup costs and setup times are reduced in an attempt to achieve “lean production”, it becomes c.p. more important to have an accurate forecasting system, i.e., to reduce the perceived demand variability.

We finally address the topic of dynamic safety stocks by studying Product 4 (out of the five products) in more detail. The expected demand per period as well as the lower and upper limits of the 95% demand intervals are depicted in Figure 3.

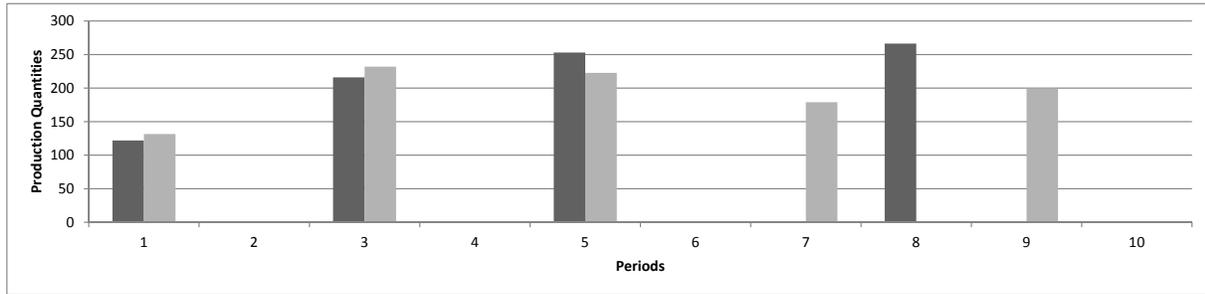
In Figure 4 we present for a target δ -service level of 99% the production quantities, the expected physical inventory, the expected backlog and the safety stocks as defined in Equation (20) for the case of a $TBO = 4$. The figures show that the overall expected physical inventory and the safety stock is substantially higher for the case with high demand variability ($VC^d = 0.3$) than for the case with low demand variability ($VC^d = 0.1$). The safety stocks are clearly dynamic and differ from lot to lot. They can also be negative, as in Periods 1 and 2, so that they lead to a substantial planned backlog in Period 2.

8 Conclusions, managerial insights, and outlook

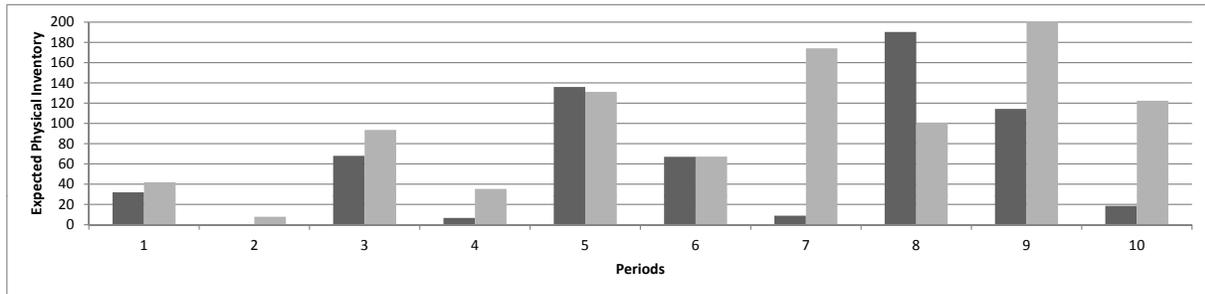
In this paper, we have proposed and evaluated different models to coordinate production quantities and safety stocks subject to random demand and a capacity-constrained production system. The solution of the models leads to robust production plans that contain dynamic safety stocks. We have furthermore presented a well-defined backlog-oriented δ -service level measure that overcomes the deficiencies of the γ -service level measure. It offers a clear interpretation and can in addition be directly related to the average waiting time of the demands. It can also be easily incorporated into dynamic and stochastic production planning models like the SCLSP. A Fix-and-Optimize heuristic was used to solve the two numerically tractable approximation models for the SCLSP as well as a model variant with period-specific service level constraints based on ideas suggested by Bookbinder and Tan (1988). Our numerical study shows that the model variant based on a piecewise linear approximation of the non-linear functions of expected backlog and expected physical inventory performs particularly well. The other modeling approach based on a scenario technique turns out to be less accurate, but more flexible with respect to probabilistic dependencies of the demands within a single scenario. To determine demand realizations for the scenario approach, Descriptive Sampling as proposed by Saliby (1990) clearly outperformed a Simple Random Sampling approach and showed a surprisingly high degree of accuracy even for small numbers of scenarios.

One important managerial insight can be derived from the comparison of the numerical results of the CLSP-model inspired by Bookbinder and Tan to enforce a period-specific δ -service level measure as opposed to an aggregate multi-period δ -service level measure in our approximation models: Imposing a period-specific service constraint can be much more expensive than imposing only an aggregate service level constraint and tolerating occasional backlogs. If the management is mainly interested in robust schedules that minimize system nervousness and bullwhip effects while meeting aggregate service level constraints, the models presented in this paper turn out to be very useful to determine such schedules. Our numerical examples have furthermore shown the potentially strong impact of demand variability (or forecast inaccuracy) on safety stocks and the costs of the schedule. We have in addition emphasized the dynamic nature of safety stocks in the case of dynamic demand and production quantities as well as the value of precise demand forecasts for production systems operating with small production lots and short production cycles. Our models can finally also be used to quantify the benefit from an improved forecasting accuracy, which is in our eyes an important indirect managerial aspect of our work.

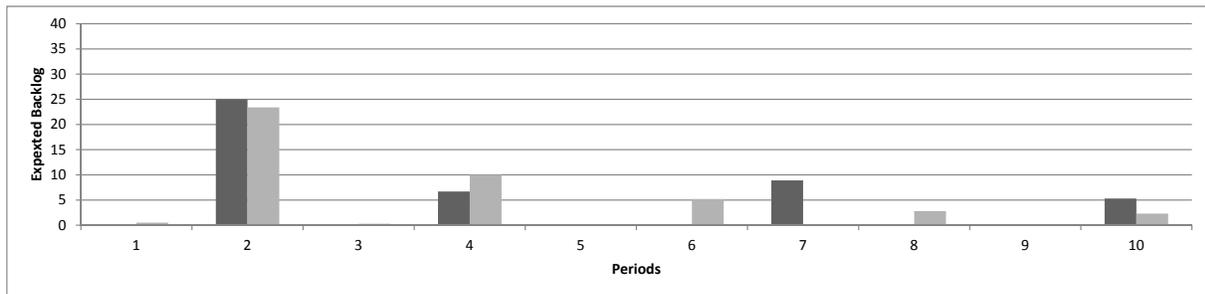
Future research should address the case of multi-level supply chains and/or additional randomness due to random production capacities and the use of the models within a rolling horizon context.



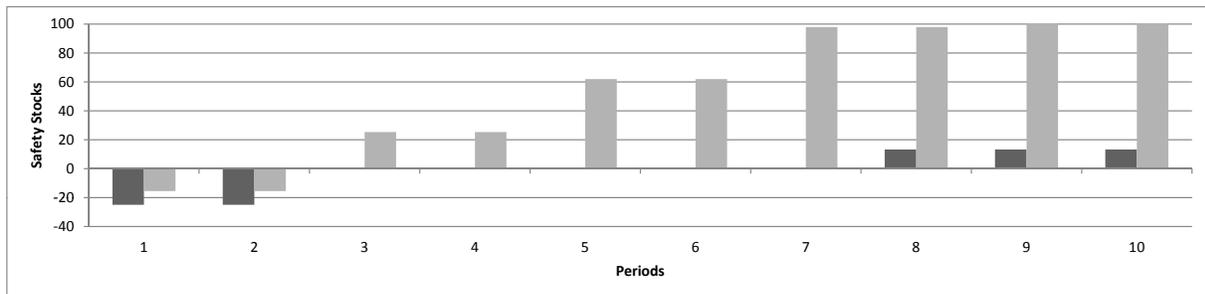
(a) Production Quantities



(b) Expected Physical Inventory



(c) Expected Backlog



(d) Safety Stocks

■ $VC^d = 0.1$ ■ $VC^d = 0.3$

Figure 4: $TBO = 4$ and $\delta = 0.99$

A Test sets

Table 11 reports the mean $E[D_k]$ of the demand. For each product k and period t , the expected dynamic demand $E[D_{kt}]$ is drawn from a normal distribution with mean $E[D_k]$ and standard deviation $VC^{ip} \cdot E[D_k]$. The drawn demand series are given in Table 12 for $VC^{ip} = 0.2$ and in Table 13 for $VC^{ip} = 0.3$.

For a problem instance from Table 4 with K products and T periods, the first K rows

Table 11: Series of expected demand $E[D_k]$

k	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
$E[D_k]$	67	135	105	80	79	72	85	136	56	100	150	108	150	126	63	114	66	75	117	93

Table 12: Series of expected demand $E[D_{kt}]$ for $VC^{ip} = 0.2$

$k \setminus t$	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
1	74	64	77	42	70	51	83	48	89	80	77	67	41	66	69	89	85	78	80	71
2	140	99	128	120	111	83	109	140	99	77	94	141	164	95	133	131	152	120	114	90
3	88	105	95	78	122	156	88	78	147	146	93	101	111	103	79	102	95	77	105	102
4	67	82	95	113	82	77	75	60	91	85	65	71	83	73	68	111	77	47	97	77
5	93	69	80	79	103	100	80	60	96	61	68	98	67	85	87	77	52	72	58	52
6	79	70	87	52	93	96	73	77	81	93	86	66	76	83	61	50	74	71	70	81
7	103	104	90	73	75	94	94	91	76	106	95	107	100	115	71	87	74	52	78	71
8	144	151	126	135	158	175	97	100	104	119	130	196	115	134	202	97	141	100	71	143
9	56	61	45	55	45	43	24	57	40	58	66	58	58	52	58	50	71	64	54	51
10	107	127	118	97	106	123	97	85	94	85	91	87	67	115	137	109	91	108	137	60
11	131	172	163	136	118	192	204	142	188	167	123	166	147	129	174	75	148	132	148	122
12	84	90	122	124	65	118	94	145	85	130	102	71	109	102	89	71	109	73	84	127
13	137	152	110	136	129	138	123	128	63	193	140	147	159	129	91	182	106	177	96	179
14	128	139	125	130	84	105	131	82	139	106	136	85	131	121	116	98	163	102	100	138
15	52	99	50	44	45	55	63	51	78	77	55	73	65	62	59	48	69	43	67	81
16	85	135	93	130	111	89	112	177	103	159	77	121	113	105	110	112	137	129	102	94
17	68	76	52	66	71	71	54	62	70	55	55	66	92	53	59	56	54	66	59	82
18	85	66	109	96	59	66	76	68	65	90	80	84	90	83	54	73	70	86	74	91
19	92	115	104	98	120	91	156	119	69	70	122	80	72	74	137	92	124	140	84	107
20	93	94	114	98	66	94	107	83	109	99	94	93	81	77	90	85	92	95	91	128

Table 13: Series of expected demand $E[D_{kt}]$ for $VC^{ip} = 0.3$

$k \setminus t$	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
1	48	76	69	76	68	58	57	69	56	70	56	104	83	84	73	62	48	96	65	68
2	80	119	77	145	68	98	135	183	120	63	163	166	121	164	178	104	159	114	141	119
3	113	59	104	63	100	143	106	81	104	158	158	28	112	52	85	120	185	117	159	69
4	90	57	123	68	117	69	67	76	76	101	109	135	58	73	77	82	111	43	89	54
5	47	66	99	91	55	92	69	88	73	76	96	80	72	74	51	70	87	50	104	81
6	101	58	95	109	28	78	67	107	71	37	46	64	66	34	30	79	102	100	40	51
7	103	61	91	70	101	74	77	98	106	80	108	80	64	48	75	69	44	87	113	72
8	211	192	199	191	113	136	113	125	183	100	118	118	113	194	142	117	140	91	97	148
9	39	41	61	61	34	69	58	55	47	28	52	27	79	47	39	94	67	56	48	70
10	98	149	76	105	153	133	43	114	96	220	83	94	99	127	113	23	90	86	113	74
11	72	108	146	222	146	126	204	142	187	226	147	220	115	227	103	200	206	164	208	153
12	117	136	108	149	97	126	87	119	93	161	128	113	45	102	150	56	132	89	105	81
13	101	140	146	207	80	121	157	106	164	152	187	123	108	133	79	192	198	191	152	79
14	108	174	125	118	130	113	188	121	125	156	156	199	152	91	95	124	139	166	118	132
15	79	76	67	50	67	61	72	52	36	58	29	53	55	12	63	51	60	44	36	46
16	102	156	112	137	97	131	64	118	152	46	83	109	92	139	103	143	168	153	42	96
17	72	37	91	47	92	75	81	46	75	78	49	53	57	97	79	57	56	54	75	43
18	66	84	85	75	31	78	75	81	57	89	82	58	90	100	59	91	53	109	88	37
19	133	89	150	134	182	123	131	75	110	140	129	114	134	156	108	192	115	238	148	108
20	78	87	57	51	119	94	108	99	129	84	23	160	136	108	94	118	74	107	101	69

and T columns were taken from Tables 12 and 13, respectively.

The holding cost hc_k and the processing times tb_k are all assumed to be 1 for all products k . The overtime cost oc is equal to 100 for one unit of overtime. The setup cost sc_k is derived based from the average expected demand $\overline{E[D_k]}$ of product k

$$\overline{E[D_k]} = \frac{\sum_{t=1}^T E[D_{kt}]}{T} \quad \forall k \quad (51)$$

and the time-between-orders TBO

$$sc_k = \frac{\overline{E[D_k]} \cdot TBO^2 \cdot hc_k}{2} \quad \forall k. \quad (52)$$

from the solution of the standard Economic Order Quantity model.

The setup time ts_k of product k depends on the given parameter ts^{rel} and the average expected processing time $\overline{E[D_k]} \cdot tb_k$ for the average period demand of product k

$$ts_k = ts^{rel} \cdot \overline{E[D_k]} \cdot tb_k \quad \forall k. \quad (53)$$

To create dynamic capacity parameters b_t , we divided the expected capacity requirement for processing by the utilization $Util$ due to processing:

$$b_t = \frac{\sum_{k=1}^K tb_k \cdot E[D_{kt}]}{Util} \quad \forall t. \quad (54)$$

Note that for instances with setup times, the actual utilization can be substantially higher, in particular in cases with low TBOs and hence many setups.

References

- Absi, N. and S. Kedad-Sidhoum (2009). The multi-item capacitated lot-sizing problem with safety stocks and demand shortage costs. *Computers & Operations Research* 36(11), 2926–2936.
- Bihlmaier, R., A. Koberstein, and R. Obst (2009). Modeling and optimizing of strategic and tactical production planning in the automotive industry under uncertainty. *OR Spectrum* 31(2), 311–336.
- Bookbinder, J. H. and J. Y. Tan (1988). Strategies for the probabilistic lot-sizing problem with service-level constraints. *Management Science* 34(9), 1096–1108.
- Brandimarte, P. (2006). Multi-item capacitated lot-sizing with demand uncertainty. *International Journal of Production Research* 44(15), 2997–3022.
- Buschkühl, L., F. Sahling, S. Helber, and H. Tempelmeier (2010). Dynamic capacitated lot-sizing problems: a classification and review of solution approaches. *OR Spectrum* 32(2), 231–261.
- Di Summa, M. and L. A. Wolsey (2008). Lot-sizing on a tree. *Operations Research Letters* 36, 7–13.
- Fleischmann, B. (2003). Bestandsmanagement zwischen Zero Stock und Inventory Control. *OR News* (19), 22–27.
- Florian, M., J. K. Lenstra, and A. H. G. R. Kan (1980). Deterministic production planning: Algorithms and complexity. *Management Science* 26, 669–679.
- Freimer, M. B., J. T. Linderoth, and D. J. Thomas (2010). The impact of sampling methods on bias and variance in stochastic linear programs. *Computational Optimization and Applications*, DOI 10.1007/s10589-010-9322-x.
- Guan, Y., S. Ahmed, G. L. Nemhauser, and A. J. Miller (2006). A branch-and-cut algorithm for the stochastic uncapacitated lot-sizing problem. *Mathematical Programming, Ser. A* 105, 55–84.
- Guan, Y. and A. J. Miller (2008). Polynomial-time algorithms for stochastic uncapacitated lot-sizing problems. *Operations Research* 56(5), 1172–1183.
- Helber, S. and K. Henken (2010). Profit-oriented shift scheduling of inbound contact centers with skills-based routing, impatient customers, and retrials. *OR Spectrum* 32(1), 109–134.
- Helber, S. and F. Sahling (2010). A fix-and-optimize approach for the multi-level capacitated lot sizing problem. *International Journal of Production Economics* 123(2), 247–256.
- Jans, R. and Z. Degraeve (2008). Modeling industrial lot sizing problems: a review. *International Journal of Production Research* 46(6), 1619–1643.
- Kanet, J. J., M. F. Gorman, and M. Stoesslein (2010). Dynamic planned safety stocks in supply networks. *International Journal of Production Research* 48(22), 6859–6880.
- Karimi, B., S. M. T. Fatemi Ghomi, and J. M. Wilson (2003). The capacitated lot sizing problem: a review of models and algorithms. *Omega* 31(5), 365–378.

- Maes, J. and L. N. van Wassenhove (1986). A simple heuristic for the multi item single level capacitated lotsizing problem. *Operations Research Letters* 4(6), 265–273.
- Martel, A., M. Diaby, and F. Boctor (1995). Multiple items procurement under stochastic nonstationary demands. *European Journal of Operational Research* 87(1), 74–92.
- Mißler-Behr, M. (1993). *Methoden der Szenario-Analyse*. Wiesbaden: Deutscher Universitätsverlag.
- Pochet, Y. and L. A. Wolsey (2006). *Production planning by mixed integer programming*. New York: Springer.
- Robinson, P., A. Narayanan, and F. Sahin (2009). Coordinated deterministic dynamic demand lot-sizing problem: A review of models and algorithms. *Omega* 37(1), 3–15.
- Sahling, F. (2010). *Mehrstufige Losgrößenplanung bei Kapazitätsrestriktionen*. Gabler Research: Produktion und Logistik. Wiesbaden: Gabler.
- Sahling, F., L. Buschkühl, H. Tempelmeier, and S. Helber (2009). Solving a multi-level capacitated lot sizing problem with multi-period setup carry-over via a fix-and-optimize heuristic. *Computers & Operations Research* 36(9), 2546–2553.
- Saliby, E. (1990). Descriptive sampling: A better approach to monte carlo simulation. *Journal of the Operational Research Society* 41, 1133–1142.
- Sox, C. R., P. L. Jackson, A. Bowman, and J. A. Muckstadt (1999). A review of the stochastic lot scheduling problem. *International Journal of Production Economics* 62(3), 181–200.
- Sox, C. R. and J. A. Muckstadt (1997). Optimization-based planning for the stochastic lot-scheduling problem. *IIE Transactions* 29(5), 349–357.
- Tempelmeier, H. (2006). *Inventory management in supply networks: Problems, models, solutions*. Norderstedt: Books on Demand.
- Tempelmeier, H. (2011). A column generation heuristic for dynamic capacitated lot sizing with random demand under a fill rate constraint. *Omega* 39(6), 627–633.
- Tempelmeier, H. and S. Herpers (2010). ABC_β - a heuristic for dynamic capacitated lot sizing with random demand under a fillrate constraint. *International Journal of Production Research* 48(17), 5181–5193.
- Tunc, H., O. A. Kilic, S. A. Tarim, and B. Eksioglu (2011). The cost of using stationary inventory policies when demand is non-stationary. *Omega* 39(4), 410–415.
- Winands, E. M. M., I. J. B. F. Adan, and G. J. van Houtum (2011). The stochastic economic lot scheduling problem: A survey. *European Journal of Operational Research* 210, 1–9.
- Zhao, X., J. Xie, and Q. Jiang (2001). Lot-sizing rule and freezing the master production schedule under capacity constraint and deterministic demand. *Production and Operations Management* 10(1), 45–67.
- Zipkin, P. H. (2000). *Foundations of inventory management*. Boston: McGraw-Hill.