OR Spectrum (2016) 38:899–920
DOI 10.1007/s00291-016-0442-2

REGULAR ARTICLE

The impact of the number of parallel warehouses on total inventory

Bernhard Fleischmann[1]

Received: 2 April 2015 / Accepted: 31 March 2016 / Published online: 16 April 2016

Abstract In the strategic design of a distribution system, the right number of stock points for the various products is an important question. In the past decade, a strong trend in the consumer goods industry led to centralizing the inventory in a single echelon consisting of a few parallel warehouses or even a single distribution center for a Europe-wide distribution system. Centralizing inventory is justified by the reduction in total stock which mostly overcompensates the increasing transportation cost. The effect of centralization is usually described by the "Square Root Law", stating that the total stock increases with the square root of the number of stock points. However, in the usual case where the warehouses are replenished in full truck loads and where a given fill rate has to be satisfied, the Square Root Law is not valid. This paper explores that case. It establishes functional relationships between the demand to be served by a warehouse and the necessary safety and cycle stock for various demand settings and control policies, using an approximation of the normal loss function and its inverse. As a consequence, the impact of the number of parallel warehouses on the total stock can be derived. The results can be used as tools in network design models.

Keywords Consumer goods distribution · Full truck load replenishment · Fill rate constraint · Normal loss function approximation

1 Introduction

The design of a distribution system usually starts from given plant locations, customer locations or areas with a certain demand, and service requirements concerning

✉ Bernhard Fleischmann
bernhard.fleischmann@wiwi.uni-augsburg.de

[1] Chair of Production and Supply Chain Management, Augsburg University, Universitätstraße 16, 86135 Augsburg, Germany

time and reliability of the deliveries. The decisions to be taken are the locations of warehouses and transshipment points, and the distribution paths from the plants to the customers (Chopra and Meindl 2007, Ch. 4.3; Fleischmann 1993). They define the distribution network which consists of plants, warehouses, transshipment points, customers. The objective is to minimize the cost of installing and maintaining the warehouses and transshipment points, including the information system, the cost of operations (transport, handling, and dispatching), and the cost of inventories.

Most manufacturers outsource the distribution task to a logistics service provider (LSP) who is then responsible for all operations. He usually owns the warehouses, transshipment locations, and the information/communication system, but not the inventory. The main question remaining for the manufacturer is then as to where to keep stocks in the network of the LSP, i.e. the decisions on the number of inventory echelons and the number and locations of stock points. In the past years, a strong trend in the consumer goods industry led to distribution systems with a single inventory echelon consisting of a single distribution center (DC) or a few parallel DCs. This is even true for Europe-wide distribution systems where the major part of the customers can be reached from a DC within a 48-h or 72-h lead time through a transportation network without further stock points. This so-called "Euro Logistics" concept is practiced or being implemented for instance by Bosch Power Tools with the International Distribution Center (IDC) in Worms (press release of July 24, 2007), by Gardena, a producer of garden tools, with a European Distribution Center in Ulm (Hartel 2009, Ch. 7.2.3), by the pencil producer Staedtler with the European Logistics Center in Nuremberg (press release of March 27, 2006), and by Viking, a producer of mowing machines, with a European Distribution center in Strasbourg (press release of March 1, 2011).

The motivation for centralizing the inventories within a distribution system is to reduce the cost of operating the warehouses and to reduce the overall inventory. On the other hand, reducing the number of warehouses increases the distances from the warehouses to the customers and hence the most expensive part of the transportation, the delivery of small customer orders (see e.g. Chopra and Meindl 2007, Ch. 4.2; Croxton and Zinn 2005). In the Euro Logistics concept, this effect is attenuated by bundling the deliveries within the LSP network. The distances for the replenishment from the suppliers to the warehouses decrease only slightly with the decreasing number of warehouses. Only if the inventory is centralized at a single location near to the main source of supply, for instance a factory, the replenishment cost can be cut down considerably.

The impact of the warehouse locations on the transportation costs can be easily expressed in a network design model, using the distances between the relevant locations and appropriate transport tariffs. The same is true for the warehouse operation costs. The impact on the inventory level, however, is less obvious and is, therefore, often disregarded or estimated only roughly. Shapiro and Wagner (2009) state a deficiency of inventory deployment decisions in network optimization models due to the incompatibility of mathematical programming models and probabilistic models for inventory planning.

This paper addresses the functional relationship between the average inventory level at a certain stock point and its throughput, i.e. the demand which is served from it, and, consequently, the relationship between the number of warehouses and the overall inventory. This question is not new, but has been investigated for about 40 years. Main results are summarized in the next section. Most authors differentiate between safety stock and cycle stock, the latter being caused by the replenishment orders or, equivalently, the transport lot sizes. In addition, the warehouse locations impact the average transit inventories, which depend on the travel times, hence on the travel distances, but not on the transport lot sizes. Therefore, costs of the transit inventories can be easily modeled, for every link of the network, as linear functions of the flow in the link (Fleischmann et al. 2015). That is why transit inventories will not be considered in the following.

This paper considers only a single echelon within a distribution system, which is typical for the most downstream stock-keeping echelon of the network of a consumer goods manufacturer. The interdependence with upstream inventories, for instance in a factory warehouse, is disregarded, in particular the effect of central control. Downstream in the supply chain, there may be additional retail warehouses. But usually the retailer operates with a very low stock level or with stockless transshipment points and requests a very high service level of the manufacturer, although from a total supply chain perspective the highest service level is required at the retail echelon. Multi-echelon inventory systems with local and central control are discussed by Tempelmeier (2011, Ch. D) and in the survey papers of Diks et al. (1996) and van Houtum et al. (1996).

The most popular model of the relationship between the number of warehouses and inventories is the Square Root Law. It states that the total inventory in N parallel warehouses is proportional to the square root of N. It appears in many textbooks (e.g. Bretzke 2010; Chopra and Meindl 2007, Ch. 11.4; Christopher 2005, p. 215; Tempelmeier 2011, Ch. C. 6) and is employed in many articles on the design of distribution systems (e.g. Croxton and Zinn 2005; Erlebacher and Meller 2000; Miranda and Garrido 2004; Ozsen et al. 2008; Shen 2007; Snyder et al. 2007), for the safety stock or for the total stock. Literature on the justification and limitations of the Square Root Law will be discussed in the next section.

However, a crucial limitation of the Square Root Law is the assumption, that the replenishment takes place in economic order quantities (EOQ), which is required not only for the cycle stock, but also for the safety stock in case of a fill rate constraint. However, the use of the EOQ is not questioned in literature. According to the author's experience with distribution network design in the consumer goods industry (Fleischmann 1993; Fleischmann et al. 2015), the EOQ assumption is absolutely unrealistic, when the transport is done by trucks, as usual in a national or continental distribution system. The reason is that the EOQ is typically much larger than the capacity of a truck. That is why transports of consumer goods between factories and warehouses nearly always run in the full truck load (FTL) mode.

This can be easily understood from the following typical data: The capacity of a truck with trailer or a semitrailer is usually restricted by the volume of 32 up to 34 pallets. The net weight of consumer goods per pallet is mostly less than 0.5 tons so that the capacity, expressed in weight, is up to 17 tons. The comparison between the EOQ and FTL is based on realistic ranges of the following parameters:

Q: truck capacity, $Q \leq 17$ tons
v: value of the good in €/kg, $v \leq 100$
i: interest rate, $i = 0.06$
F: cost per trip, $F \geq 500$ € for a long-distance trip
t: demand per week in number of FTLs, $t \geq 2$.

The EOQ is larger than Q if, for the FTL replenishment, the transport cost exceeds the holding cost per week, i.e. if $t\,F \geq \frac{1}{2} \cdot 17{,}000 \cdot v \cdot \frac{i}{52} \approx 9.8\,v$ or $v \leq 0.1\,t\,F$, which is satisfied for the above typical ranges of the parameters. Therefore, the EOQ assumption is not realistic. The replenishment quantities are FTLs. Exceptions may occur for very small warehouses, where an FTL would cause undesired long replenishment cycles. But in this case, the less than truck load (LTL) lot sizes are even more away from the EOQ.

This paper makes the following contributions: It investigates the relationship between the number of warehouses and the total inventory for the first time for the realistic situation of replenishment by trucks and a fill rate constraint. It emphasizes the limitation of the Square Root Law by the EOQ assumption, which has been disregarded in literature so far. It analyzes the total inventory as a function of the number of warehouses for various demand distributions (normal and Gamma distribution) and inventory control policies (continuous and periodic review) and explores the effect of demand correlation, using an approximation of the normal loss integral and its inverse. On the whole, it provides novel functional relationships between the inventory level at a stock point and its throughput, which can be used as modules in network design.

The remainder of the paper is organized as follows: The next section defines the problem setting and discusses the justifications and limitations of the Square Root Law. In Sect. 3 a basic model with normally distributed demand and continuous review is used for determining the stock functions. Three realistic data sets are presented which are used throughout the paper. Section 4 analyzes the effect of demand correlation. Sections 5 and 6 consider the case of Gamma distributed demand and a periodic review model. Section 7 draws the conclusions. The Appendix presents an approximation of the normal loss integral and proves some properties of it which are used in the paper for the analysis of the stock functions.

2 The square root law

The following situation is considered (cf. Fig. 1):

1. There are N parallel stock points supplying customers with a given demand of a single product.
2. The total demand per day is stochastic with expectation D and standard deviation σ_D and does not depend on N.
3. There is a fixed allocation of the customers to the stock points so that the expected daily demand per stock point is $d = \frac{D}{N}$.
4. The total demand is composed of many small uncorrelated customer demand so that the daily demand per stock point has the standard deviation $\sigma_0\sqrt{d}$ where $\sigma_0 = \frac{\sigma_D}{\sqrt{D}}$.

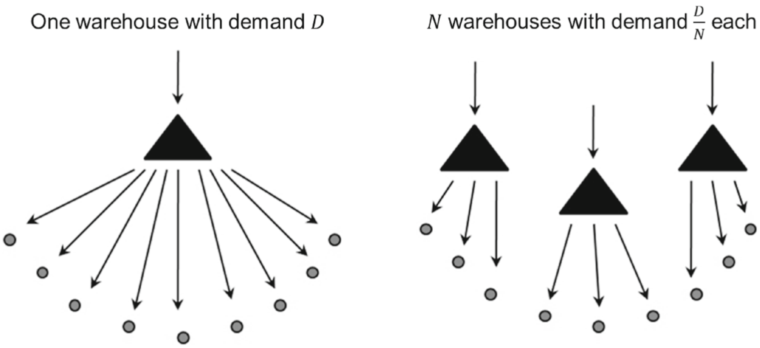

Fig. 1 Distribution structures considered

5. The fixed lead time for replenishing the stock points is L days for all stock points and independent of N.
6. The replenishment of the stock points is done by trucks in FTL of the size Q, if Q does not exceed the average demand of a given maximal cycle time $t_{\max}$ (for instance 1 week), otherwise in the LTL quantity $d\,t_{\max}$.
7. Every stock point must satisfy a service level constraint in form of a fill rate β which is independent of N. Unsatisfied demand is backordered.

The critical Assumption 4 about independent demand will be relaxed within Sect. 4. Assumption 7 is criticized by Tempelmeier (2011, p. 159) because the distances from the warehouses to the customers increase with decreasing N and, therefore, the lead time toward the customers is likely to increase as well. Therefore, a higher service level at the warehouse is required to keep a constant service level for the customers. But in a modern Euro Logistics system, as explained above, the lead time for the customers is nearly independent of the transport time to the customers, which is only a part of the lead time.

Given the probability distribution of the demand, a replenishment strategy, and a service level constraint, we define the following stock functions:

For a single warehouse with demand rate d:

$S(d)$ necessary safety stock
$C(d)$ average cycle stock

and for the distribution system with N stock points:

$S^{\mathrm{T}}(N)$ total safety stock
$C^{\mathrm{T}}(N)$ total average cycle stock.

Due to $d = \frac{D}{N}$ we have

$$S^{\mathrm{T}}(N) = N \cdot S\left(\frac{D}{N}\right) \quad \text{and} \quad C^{\mathrm{T}}(N) = N \cdot C\left(\frac{D}{N}\right). \tag{1}$$

In the following, the structure of the functions $S(d)$, $C(d)$, $S^{\mathrm{T}}(N)$, and $C^{\mathrm{T}}(N)$ will be analyzed for various situations. A wide-spread assumption about this function is the *Square Root Law*, which states

$$S^{\mathrm{T}}(N) = A\sqrt{N} \quad \text{with some constant } A$$

or, equivalently,

$$S(d) = A'\sqrt{d} \quad \text{with a constant } A' = \frac{A}{\sqrt{D}}.$$

The following justifications of the Square Root law in literature are often referred to. Let Φ and φ denote the cdf and the density function of the standard normal distribution.

The multilocation newsboy problem Eppen (1979) considers N warehouses with normally distributed demand in a single period. The special case of uncorrelated equal demands satisfies the Assumptions 1 to 3 above. The optimal safety stock in a newsboy model is $k\,\sigma$ with a safety factor

$$k = \Phi^{-1}\left(\frac{p}{h+p}\right),$$

where h is the unit holding cost, p the penalty cost and $\sigma = \sigma_0\sqrt{d}$ the standard deviation of the demand. Hence

$$S(d) = k\,\sigma_0\sqrt{d} \tag{2}$$

satisfies the Square Root Law. In addition, Eppen (1979) shows that this is also true for the total holding and penalty cost and he derives cost functions for the more general case of correlated demands with given covariances.

However, this model is not compatible with the problem setting: First, the single-period model does not include warehouse replenishments and cycle stock. Second, the fill rate β is not fixed: With the normal loss integral

$$R(k) = \int_k^\infty (x-k)\,\varphi(x)\,\mathrm{d}x \tag{3}$$

the fill rate $\beta = 1 - R(k)\frac{\sigma}{d} = 1 - R(k)\frac{\sigma_0}{\sqrt{d}}$ increases with increasing d.

α service level constraint and fixed safety factor In case of a constant safety factor k, the Square Root Law for the safety stock is simply based on the "Risk Pooling Effect" for uncorrelated demands, i.e. the fact that the standard deviation is proportional to $\sqrt{d}$ and the coefficient of variation (CV) is proportional to $\frac{1}{\sqrt{d}}$ and to $\sqrt{N}$.

Stulman (1987) modifies Eppen's model by introducing an α service-level constraint. Then the safety factor k in (2) becomes $k = \Phi^{-1}(\alpha)$. As it is again independent of d, the Square Root Law is satisfied. This remains true in a multi-period model with a continuous-review (r, q) policy (Chopra and Meindl 2007, Ch. 11.4). There, σ refers to the lead time demand, and the safety stock is $S(d) = k\,\sigma_0\sqrt{L\,d}$. Many authors justify the Square Root Law by a safety factor, which is assumed to be constant (e.g. Croxton and Zinn 2005; Evers and Beier 1993, 1998; Evers 1995; Maister 1976; Schwarz 1981).

EOQ If the replenishment quantity q in a (r, q) policy is determined as an economic order quantity (EOQ), then $q = z\sqrt{d}$ with some constant z depending on the cost factors. In this case the Square Root Law holds also for the cycle stock $C(d) = \frac{1}{2}q$ and consequently for $C^{\mathrm{T}}(N)$. This relationship is widely used in literature. Evers (1995) compares the EOQ with an order up to policy where q is a fixed multiple of the lead time demand. In this case, the total cycle stock is constant, independent of the centralization. This accords with the LTL policy considered in this paper (see Assumption 6).

β service level constraint The fill rate or β service level, which is more significant for the customer service than the α service level, is a popular service measure in inventory control research and in practice. For a continuous-review (r, q) policy and a given fill rate β, the safety factor is

$$k = R^{-1}\left(\frac{q}{\sigma_0\sqrt{L\,d}}(1-\beta)\right). \tag{4}$$

This formula holds, if the expected shortage immediately after replenishment can be neglected, as it is the case for a high fill rate β (Tempelmeier 2011, Ch. C. 1.1). Thus, the safety factor k depends on d and on q. The only situation where k is constant is when $\frac{q}{\sqrt{d}}$ is constant, as is the case for the EOQ. This relevance of the EOQ for the Square Root Law of the safety stock has not been considered in literature so far. Unfortunately, as explained in the Introduction, the EOQ is not appropriate for the warehouse replenishment by trucks. While the β service level is often considered in multi-echelon models (see Diks et al. 1996), it has not been used in models of centralization of inventory. The only exception, to our knowledge, is Tempelmeier (2011, Ch. C. 6), who also stresses the necessity of a constant safety factor for the Square Root Law. But he considers the safety factor as a function of β only, disregarding its dependence on d and q.

Correlated demand Obviously, the Square Root Law gets lost if the demand at different locations is correlated. Models with explicit covariances (Eppen 1979; Schwarz 1981) can only be used for comparing a decentralized system that has N fixed stock points with a centralized system with $N = 1$. Nevertheless, Evers and Beier (1993, 1998) and Evers (1995) consider the effect of consolidating N stock points with given covariances into M new stock points $(1 \leq M < N)$. But they assume tacitly that there is perfect correlation *within* the demand of each of the original stock points i with demand D_i and variance V_i so that any partial demand of D_i, say $W\,D_i$ with $0 < W < 1$, has the variance $W^2\,V_i$. Thus, the original N stock points hold an artificial special position, which explains the strange result, that in case of no correlation *between* the D_i the safety stock is always reduced by the factor $\frac{1}{\sqrt{N}}$, independent of M. This effect is not due to centralization, but due to the reallocation of demand, $\frac{1}{M}$ of the demand of every original location to each of the new locations. This would even work for $M = N$.

Another way to generalize the Square Root Law is suggested by Ballou (2005) and by Shapiro and Wagner (2009) who construct stock functions of the form $S(d) = \eta\, d^{\theta}$ from empirical data, where the exponent $\theta \geq 0.5$ can take demand correlation into account. In a similar way, correlated demand will be considered in Sect. 4 using the model $\sigma = \eta\, d^{\theta}$ for the standard deviation σ.

3 A continuous review model

In this section, stock functions are derived for normal distribution of the demand, a continuous-review (r, q) policy, and a given fill rate β, based on the assumptions in Sect. 2. For the ease of writing we use the notations $\sigma_L = \sigma_0 \sqrt{L}$, and $\sigma_{LD} = \sigma_0 \sqrt{L\,D}$. As the demand is composed of many very small units, an "undershoot" under the reorder point r can be neglected. When a variation of the demand rate d is considered, it is important to know how the replenishment quantity q varies with d. As explained before, the EOQ is not adequate for that purpose. Instead, Assumption 6 leads to the following relationship:

$$q = \begin{cases} Q, & \text{if } d \geq \frac{Q}{t_{\max}} \text{ (FTL case)} \\ d\, t_{\max}, & \text{if } d < \frac{Q}{t_{\max}} \text{ (LTL case)} \end{cases} \tag{5}$$

Using (4) and $C(d) = \frac{1}{2} q$ we get the stock of a single warehouse as a function of d:

FTL case

$$S(d) = R^{-1}\left(\frac{Q}{\sigma_L \sqrt{d}}(1-\beta)\right)\sigma_L \sqrt{d} \quad \text{and} \quad C(d) = \frac{1}{2} Q \tag{6}$$

LTL case

$$S(d) = R^{-1}\left(\frac{\sqrt{d}\, t_{\max}}{\sigma_L}(1-\beta)\right)\sigma_L \sqrt{d} \quad \text{and} \quad C(d) = \frac{1}{2} d\, t_{\max}. \tag{7}$$

The total stock for N warehouses is then, using (1):

FTL case

$$S^{\mathrm{T}}(N) = R^{-1}\left(\frac{Q\sqrt{N}}{\sigma_{LD}}(1-\beta)\right)\sigma_{LD}\sqrt{N} \quad \text{and} \quad C^{\mathrm{T}}(N) = \frac{1}{2} N Q \tag{8}$$

LTL case

$$S^{\mathrm{T}}(N) = R^{-1}\left(\frac{D\, t_{\max}}{\sigma_{LD}\sqrt{N}}(1-\beta)\right)\sigma_{LD}\sqrt{N} \quad \text{and} \quad C^{\mathrm{T}}(N) = \frac{1}{2} D\, t_{\max}. \tag{9}$$

The above safety stock functions are valid for the expected average net inventory $I = I^{+} - I^{-}$ with the expected average stock on hand I^{+} and the expected average

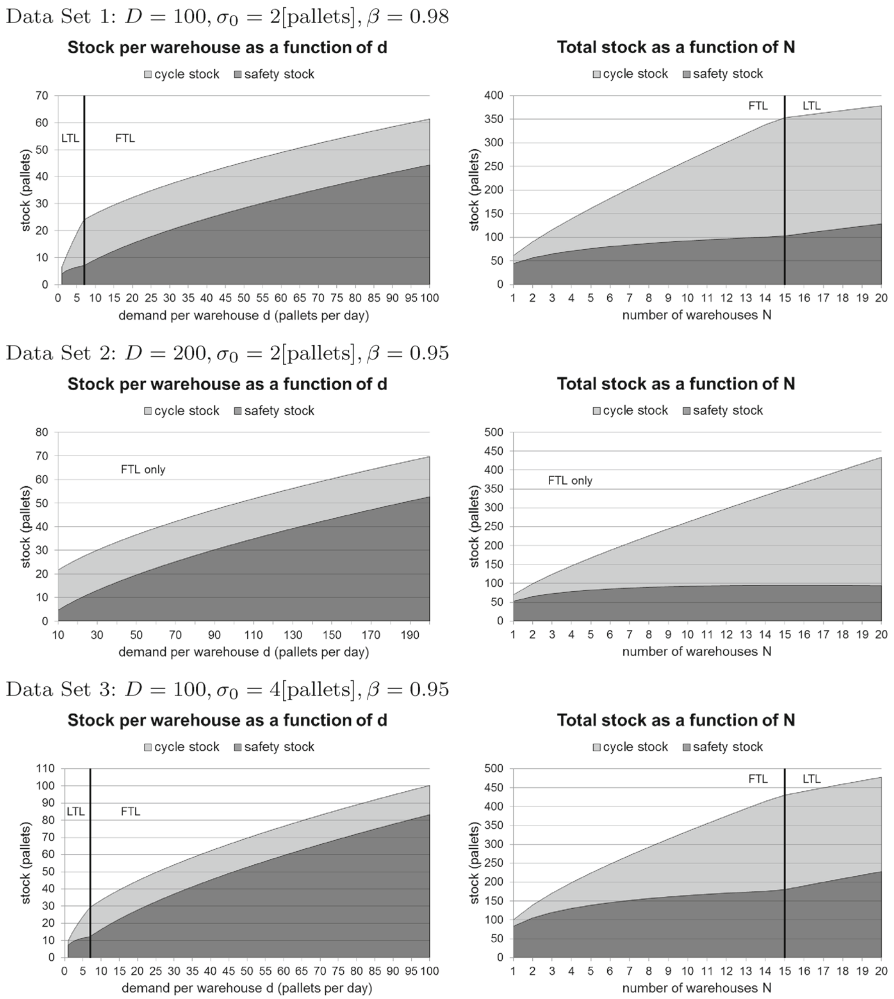

Fig. 2 Safety stock and cycle stock for continuous review and normally distributed demand. Common data: $Q = 34$ [pallets], $L = 2$, $t_{\max} = 5$ [days]

shortage I^-. I^- can be approximated by $\frac{1}{2}(I_0^- + I_1^-)$ with I_0^-, I_1^- being the expected shortage at the beginning and at the end of a replenishment cycle, respectively. As the safety stock is determined such that $I_1^- - I_0^- = (1-\beta)q$, it follows $I^- \approx \frac{1}{2}(1-\beta)q + I_0^-$, where I_0^- is small compared to I^- and can be neglected in most cases. Thus, for given data q and β, the difference between the average net inventory and the stock on hand is nearly constant and, for a high service level β, small compared to $S(d)$. In the following, in accordance with most textbooks on inventory management, the correction of the safety stock by I^- is forgone and all results represent net inventories, giving a clearer view of the structure of the stock functions.

Examples Figure 2 shows, for three realistic data settings, the safety stock and cycle stock functions for a single warehouse and in total. The quantity unit is a pallet, as

Table 1 Maximum of $S^{\mathrm{T}}(N)$ for FTL

Data set	c	N_0	max $S^{\mathrm{T}}(N)$	Indifference area
1	0.024	48	117	[19, 96] (outside FTL area)
2	0.0425	15	95	[6, 30]
3	0.030	31	188	[12, 62] (FTL up to $N = 14$)

usual in the distribution business. The full truck load of $Q = 34$ pallets corresponds to the capacity of a truck with trailer or a semitrailer. The replenishment lead time is $L = 2$ days, as is typical in a national or European distribution system. The maximal cycle time is $t_{\max} = 5$ days, i.e. any warehouse must be replenished at least once a week. The number of warehouses is varied in the range $N = 1, \ldots, 20$. The data sets differ in the following parameters:

The service level is $\beta = 98$ % (Data Set 1) or $\beta = 95$ % (Data Sets 2 and 3).

The total demand is $D = 100$ or $D = 200$ pallets per day, that is about 3 or 6 full loads per day, respectively. In the first case (Data Sets 1 and 3), FTL replenishment is possible up to $N = 14$ warehouses. For $N = 15$, the load is $q = 33.3$ pallets, it goes down to $q = 25$ pallets for $N = 20$. In the second case (Data Set 2), only FTL occurs.

The standard deviation per pallet and day is $\sigma_0 = 2$ or 4. In the first case (Data Sets 1 and 2), the CV of the lead time demand of 200 is 0.14 (such as for $D = 100$, $N = 1$), but increases to 0.63 for $N = 20$, which is already critical for the assumption of normal distribution. In the second case (Data Set 3), the CV of the lead time demand is 0.28 for $N = 1$, but 1.26 for $N = 20$, which strongly contradicts the normal distribution. Nevertheless, in this section and in Sect. 6, the stock functions are based on this model, which permits an easy explanation of the observed stock curves. In Sect. 5, the more realistic Gamma distribution is used instead, but it will turn out that this does not change the general structure of the stock functions.

Surprisingly, in the FTL case, the total safety stock $S^{\mathrm{T}}(N)$ increases rather slowly and may even decrease as for Data Set 2, as Fig. 2 and the following numerical analysis show:

Using the approximations and properties A.2 and A.3 in the Appendix, $S^{\mathrm{T}}(N) \approx H(c\sqrt{N})\sqrt{N}\sigma_{LD}$ with $c = \frac{Q(1-\beta)}{\sigma_{LD}}$ attains the maximum at $N_0 \approx \frac{1}{(6c)^2}$ with the value $S^{\mathrm{T}}(N_0) \approx \frac{\sigma_{LD}}{10c}$. Within the indifference area $[0.4N_0, 2N_0]$, $S^{\mathrm{T}}(N) \geq 0.9 S^{\mathrm{T}}(N_0)$ holds. Table 1 specifies the results, which accord well with the safety stock functions in Fig. 2. The maximum itself can only be observed for $S^{\mathrm{T}}(N)$ in Data Set 2, whereas the other cases show the slowly increasing safety stock left to the maximum.

This trend of $S^{\mathrm{T}}(N)$ contradicts the Square Root Law and leads to the following question: Why is there not a stronger risk pooling effect? This can be explained in a simple example: Suppose, the total demand per week is 5 Q. Thus, if $N = 5$, every warehouse gets one replenishment per week. In case of centralization ($N = 1$), the single warehouse receives one truck every day. It is well known that for a (r, q) policy, reducing the cycle time requires a higher safety stock. This effect partly or totally compensates the risk pooling effect. If all five trucks were sent together on the same day once a week, risk pooling would become effective, but this would cause unreasonable and unnecessary cycle stock. Nevertheless, centralization yields a big advantage, namely the reduction of the cycle stock, not of the safety stock. In our example, the cycle stock would decrease from $\frac{5}{2}\,Q$ to $\frac{1}{2}\,Q$.

The functions $S(d)$ for FTL and $S^{\mathrm{T}}(N)$ for LTL are of the form $h_2(N)$ (see Appendix A.4) and, therefore, increase stronger than $\sqrt{N}$ and weaker than linearly, but in all cases nearly linearly. Thus, in the LTL case, the role of the two types of stock is exchanged: The cycle stock remains constant, whereas the safety stock increases monotonically and stronger than $\sqrt{N}$.

4 Correlated demand

The assumption that the demands at the different stock points are independent may be unrealistic in many cases. In particular, the local demands in a distribution system may be influenced by global factors entailing positive correlation. In this section, we analyze the effect of demand correlation on the centralization of inventories.

As the focus of this paper is the general relationship between the number of stock points and the total inventory, we do not consider explicit covariances. Instead the following model is used which generalizes the variance V of the demand as a function of the expected value d as $V(d) = V_0\, d^{2\theta}$, where the exponent $\theta = 0.5$ is valid for independent demand, $\theta > 0.5$ or $\theta < 0.5$ for positive or negative correlation, respectively, and V_0 is a constant. A relationship between θ and the correlation coefficient ρ can be established as follows: Let X_1 and X_2 be the demand of two stock points, both with expectation d and variance σ^2, and with the correlation coefficient ρ. Then the variance of $X_1 + X_2$ must equal $V(2d)$ or $2\sigma^2(1+\rho) = 2^{2\theta}\,\sigma^2$, which leads to $\theta = \frac{1}{2}(1+\frac{\ln(1+\rho)}{\ln 2})$. We assume $-1 < \rho < 1$, i.e. the trivial cases of perfect correlation are not considered.

In order to show the effect of correlation for the Data Sets 1 to 3 compared with the case of independent demand, V_0 is fixed to the variance of the lead time demand for $N = 20$, the highest number of stock points considered, when there is no correlation:

$$V_0 = \frac{\sigma_0^2\, L\, D}{20},$$

so that the lead-time demand of each of N stock points with daily demand $\frac{D}{N}$ has the standard deviation

$$\sigma = \sqrt{V_0}\left(\frac{20}{N}\right)^{\theta} = \sigma_{LD} N^{-\theta} \quad \text{with } \sigma_{LD} = \sqrt{V_0}\, 20^{\theta}.$$

Then, the safety stock $S^{\mathrm{T}}(20)$ is equal for all ρ, including $\rho = 0$, but $S^{\mathrm{T}}(N)$ depends on ρ for $N < 20$. Note that this variable transformation does not affect any results for $\rho = 0$ or correspondingly $\theta = 0.5$. ρ can be interpreted as the correlation either between the lead-time demands per stock point or between their daily demands if

Table 2 Maximum of $S^{\mathrm{T}}(N)$ for various correlation coefficients ρ (Data Set 2)

ρ	θ	σ_{LD}	N_0	$S^{\mathrm{T}}(N_0)$
−0.1	0.4240	31.86	24.5	94.6
0	0.5	40	15.5	94.9
0.1	0.5688	49.15	10.3	100.6
0.3	0.6893	70.52	4.4	123.0
0.5	0.7925	96.07	1.3	163.4

there is no correlation between different days. The cycle stock is not affected by the correlation.

Now the safety stock functions (8) and (9) can be generalized:

$$\text{for the FTL case:}\quad S^{\mathrm{T}}(N) = R^{-1}\left(\frac{Q\,N^{\theta}}{\sigma_{LD}}(1-\beta)\right)\sigma_{LD}N^{1-\theta}\quad\text{and}$$

$$\text{for the LTL case:}\quad S^{\mathrm{T}}(N) = R^{-1}\left(\frac{D\,t_{\max}}{\sigma_{LD}N^{1-\theta}}(1-\beta)\right)\sigma_{LD}N^{1-\theta}.$$

In this comparison, we also include the Square Root Law (SRL) for the safety stock and for the total stock, using the same variable transformation:

$$\text{for the SRL:}\quad S^{\mathrm{T}}(N) = \sqrt{\frac{N}{20}}\,S^{\mathrm{T}}(20)\quad\text{and}$$

$$S^{\mathrm{T}}(N) + C^{\mathrm{T}}(N) = \sqrt{\frac{N}{20}}\,(S^{\mathrm{T}}(20) + C^{\mathrm{T}}(20)).$$

Figure 3 shows the safety stock $S^{\mathrm{T}}(N)$ and the total stock $S^{\mathrm{T}}(N) + C^{\mathrm{T}}(N)$ for various $\rho \geq 0$, $\rho = -0.3$, and for the SRL. It is well known that the risk pooling effect is weakened by positive correlation and strengthened by negative correlation. Thus, the necessary safety stock after centralization increases with increasing ρ. As explained in the previous section, the risk pooling effect is also diminished, even for $\rho = 0$, by the fact that increasing demand entails shorter replenishment cycle times. Therefore, in the FTL case, $S^{\mathrm{T}}(N)$ turned out not to be monotonic, but to attain a maximum. This remains true in case of correlation, as shown in the Appendix A.2, using the approximation $S^{\mathrm{T}}(N) \approx H(c\,N^{\theta})N^{1-\theta}\sigma_{LD}$ with $c = \frac{Q}{\sigma_{LD}}(1-\beta)$. The maximum, which can be determined analytically, is larger and attained at smaller N if ρ increases.

Table 2 specifies these values, according with Fig. 3, for Data Set 2 as an example. For $\rho > 0.54$ the maximum is attained for $N < 1$ so that $S^{\mathrm{T}}(N)$ decreases monotonically in the relevant area $1 \leq N \leq 20$. For $\rho > 0.35$, $S^{\mathrm{T}}(N)$ is not concave in this area, but exhibits an inflection point, for instance at $N \approx 6$ for $\rho = 0.5$. For $\rho < -0.06$ the maximum is attained for $N > 20$, so that $S^{\mathrm{T}}(N)$ is monotonically increasing and concave, as shown in the Appendix A.2.

Thus, in case of positive correlation, the total safety stock may even increase significantly by centralization. Nevertheless, centralization still remains advantageous

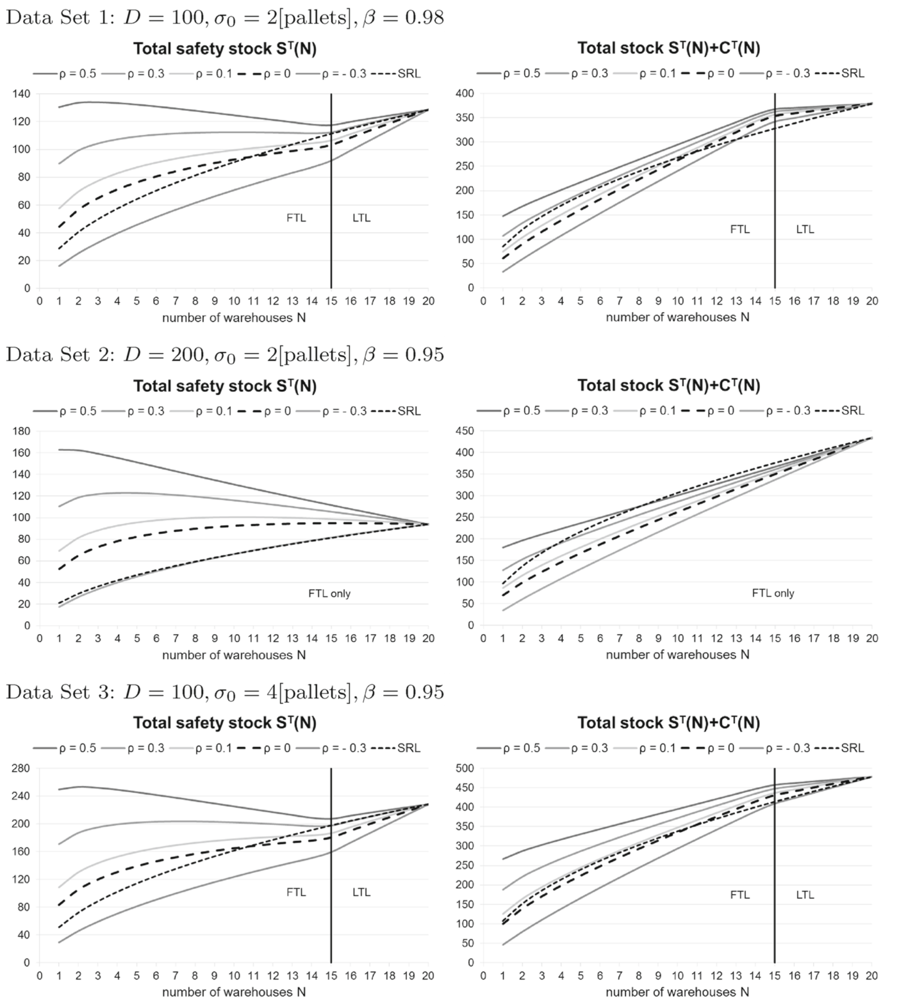

Fig. 3 Total safety stock and total stock for correlated demand and for the SRL. Common data: $Q = 34$ [pallets], $L = 2$, $t_{\max} = 5$ [days]

due to the reduced cycle stock, as can be seen from the total stock curves in Fig. 3. Negative correlation, which is of less importance for demands at different locations, leads to a monotonic increase of $S^T(N)$ and the total stock curve exhibits a still more distinct linear trend than for independent demand.

The SRL clearly overestimates the centralization effect for the safety stock in the FTL area, as it exhibits, for decreasing N, a stronger decrease of $S^T(N)$ than all cases with $\rho \geq 0$. However, for the total stock, where the centralization effect consists mainly in the linear decrease of the cycle stock, the SRL underestimates this effect compared with $\rho = 0$. In the LTL area, the comparison of the SRL with the other cases shows the opposite relation.

5 Gamma-distributed demand

As stated in Sect. 3, the assumption of normally distributed demand is questionable in case of higher variance and smaller demand, as it occurs in particular in Data Set 3. In this section, the lead-time demand is considered to be Gamma-distributed instead and again uncorrelated. Given the demand d per unit of time and the standard deviation σ_0 per unit and per unit of time as before, the parameters of the Gamma distribution are chosen such that the mean $\mu = L\,d$ and the variance $\sigma^2 = \sigma_0^2\,L\,d$ of the lead-time demand remain unchanged. Then, the parameters are (see Tempelmeier 2011, p. 341) the shape parameter $p = (\frac{\mu}{\sigma})^2 = \frac{L\,d}{\sigma_0^2}$ and the scale parameter $\lambda = \frac{\mu}{\sigma^2} = \frac{1}{\sigma_0^2}$.

Note that $\frac{p}{\lambda} = L\,d$. Let $F_G(r, p, \lambda)$ denote the cdf of the corresponding Gamma distribution. For a given reorder point r, the loss function for the Gamma distribution

$$R_G(r) = L\,d\,[1 - F_G(r, p+1, \lambda)] - r\,[1 - F_G(r, p, \lambda)]$$

represents the expected shortage at the end of a replenishment cycle. The expected backorder per replenishment cycle equals the difference between the shortages at the end and at the beginning of the cycle. The latter term, $R_G(r+q)$, cannot be neglected here, as it amounts to up to 30 % of $R_G(r)$ for Data Set 3. The reorder point is determined by the equation

$$R_G(r) - R_G(r+q) = (1-\beta)q$$

which can be easily solved by the Excel solver, using the reorder points from the normal distribution as starting values. The safety stock is the expected inventory at the end of the replenishment cycle, $S(d) = r - L\,d$, and the cycle stock is $C(d) = \frac{1}{2}q$.

Figure 4 shows the resulting stock functions. The comparison with Fig. 2 shows that the safety stock, as expected, is somewhat higher than for normally distributed demand, the difference being larger for smaller demand and for higher service level. For Data Set 1, the relative difference of $S(d)$ ranges from +33.6 % ($d = 5$) to +8.7 % ($d = 100$) for Data Set 3 from +30 % ($d = 5$) to +2.5 % ($d = 100$). But in all cases, the structure of the stock functions, as observed and analyzed in Sect. 3 for the normal distribution, remains the same. In particular, in the FTL case, the total safety stock $S^{\mathrm{T}}(N)$ increases only slowly or is nearly constant, such as for Data Set 2. In the LTL case, the nearly linear increase of $S^{\mathrm{T}}(N)$ is steeper than for the normal distribution, due to the larger difference for the smaller demands.

6 A periodic review model

In a continuous review model, as considered before, replenishments of the warehouse can occur at any time on a continuous time axis. A more realistic assumption is that the trucks arrive only once per day within a small time window. This situation can be modeled by a $(r, n\,q)$ policy: The inventory is reviewed at the beginning of every day. When the inventory position s is less than or equal to the reorder point r, an order of $n\,q$ ($n \geq 1$, integer) is placed such that s is raised into the interval $(r, r+q]$. As before,

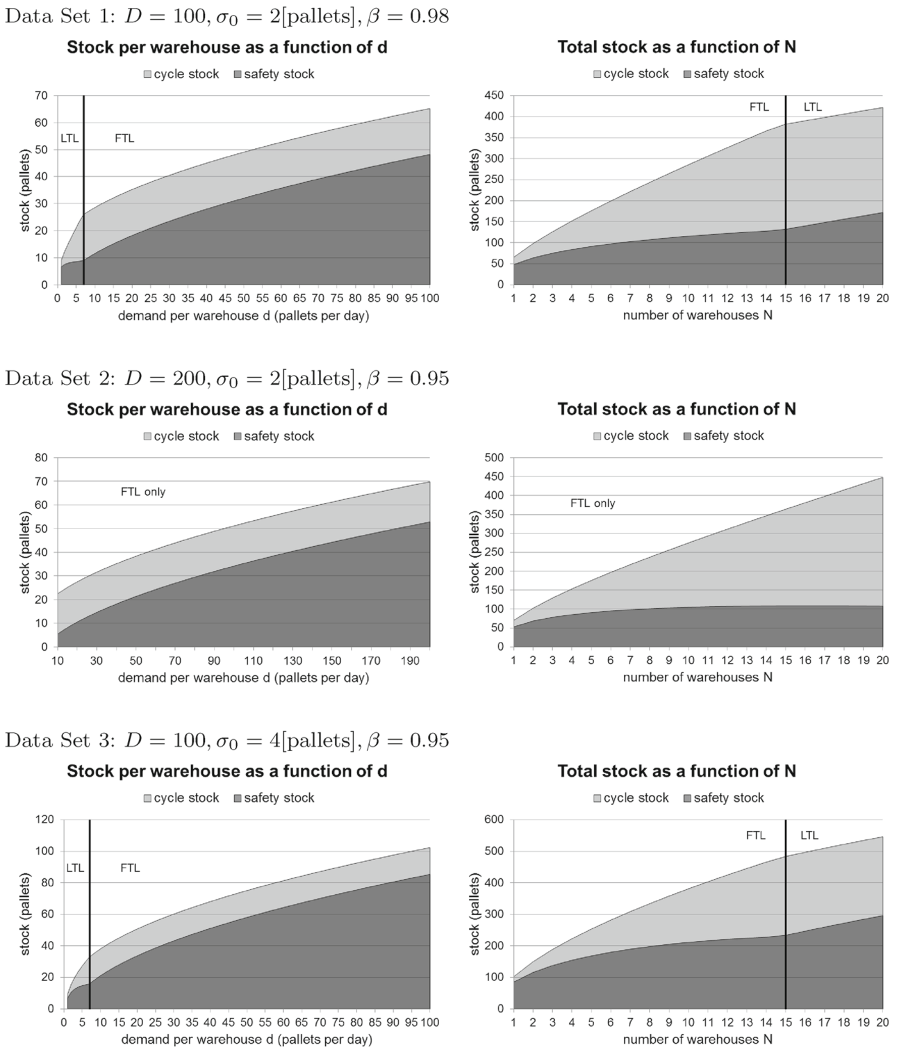

Fig. 4 Safety stock and cycle stock for continuous review and gamma distributed demand. Common data: $Q = 34$ [pallets], $L = 2$, $t_{\max} = 5$ [days]

we set $q = \min(Q, d\, t_{\max})$. The lead time $L \geq 0$ is integer; the order placed on day t arrives and is available on day $t + L$. The demand per day is normally distributed with $\mu = d$ and $\sigma = \sigma_0\sqrt{d}$.

The analysis of the $(r, n\, q)$ policy is based on the fact that the inventory position at the beginning of any day, after a potential order placement, is $r + y$, where the overshoot y is uniformly distributed within $(0, q]$. Hence, the expected net inventory at the beginning of any day, after a potential replenishment, is $I_0 = r + \frac{1}{2}q - L\, d$ and at the end of the day $I_1 = r + \frac{1}{2}q - (L + 1)\, d$. Hadley and Whitin (1963, Ch. 5-5) show that for normally distributed demand the expected shortage at the beginning and at the end of the day are

$$I_0^- = \frac{1}{q}[u(r, L) - u(r+q, L)] \quad \text{and} \quad I_1^- = \frac{1}{q}[u(r, L+1) - u(r+q, L+1)],$$

where

$$u(v, t) = \frac{1}{2}\sigma_0^2\, t\, d[(1+x^2)(1-\Phi(x)) - x\varphi(x)] \quad \text{and} \quad x = x(v, t) = \frac{v - t\,d}{\sigma_0\sqrt{t\,d}}.$$

Thus, the reorder point r is determined by the equation

$$\frac{1}{d\,q}[u(r, L+1) - u(r+q, L+1) - u(r, L) + u(r+q, L)] = \beta. \tag{10}$$

The expected average net inventory is then

$$I = \frac{1}{2}(I_0 + I_1) = r - (L+1)d + \frac{1}{2}(d+q).$$

How to split it into safety stock and cycle stock is not obvious for the $(r, n\,q)$ policy because the usual definition of the safety stock, the expected net inventory at the end of a replenishment cycle, does not appear appropriate. Even for deterministic demand, the stock at the end of the replenishment cycle does not vanish, due to the discrete time axis and the discrete replenishment quantities $n\,q$. Therefore, the total stock is split, somewhat arbitrarily, into

$$S(d) = r - (L+1)d \quad \text{and} \quad C(d) = \frac{1}{2}(d+q).$$

Note that in the LTL case, $C^{\mathrm{T}}(N) = \frac{1}{2}(1 + t_{\max})\, D$ is constant like for continuous review. Figure 5 shows the stock functions for the three data sets.

Alternatively, the reorder point r can be calculated by the following approximation: r must cover the demand during the lead time and one review period minus the overshoot y, i.e. a variable z with mean $\mu_z = (L+1)d - \frac{1}{2}q$ and standard deviation $\sigma_z = \sqrt{\sigma_0^2(L+1)d + \frac{q^2}{12}}$. With the safety factor $k = R^{-1}(\frac{(1-\beta)d}{\sigma_z})$, the approximate reorder point is

$$r' = k\,\sigma_z + (L+1)d - \frac{1}{2}q. \tag{11}$$

The approximation considers z as normally distributed, disregarding its true distribution, a convolution of a normal and a uniform distribution. Moreover, the shortage at the beginning of a day is neglected. Nevertheless, the approximation (11) is very close to the exact reorder point defined by (10). In all cases of the Data Sets 1 to 3,

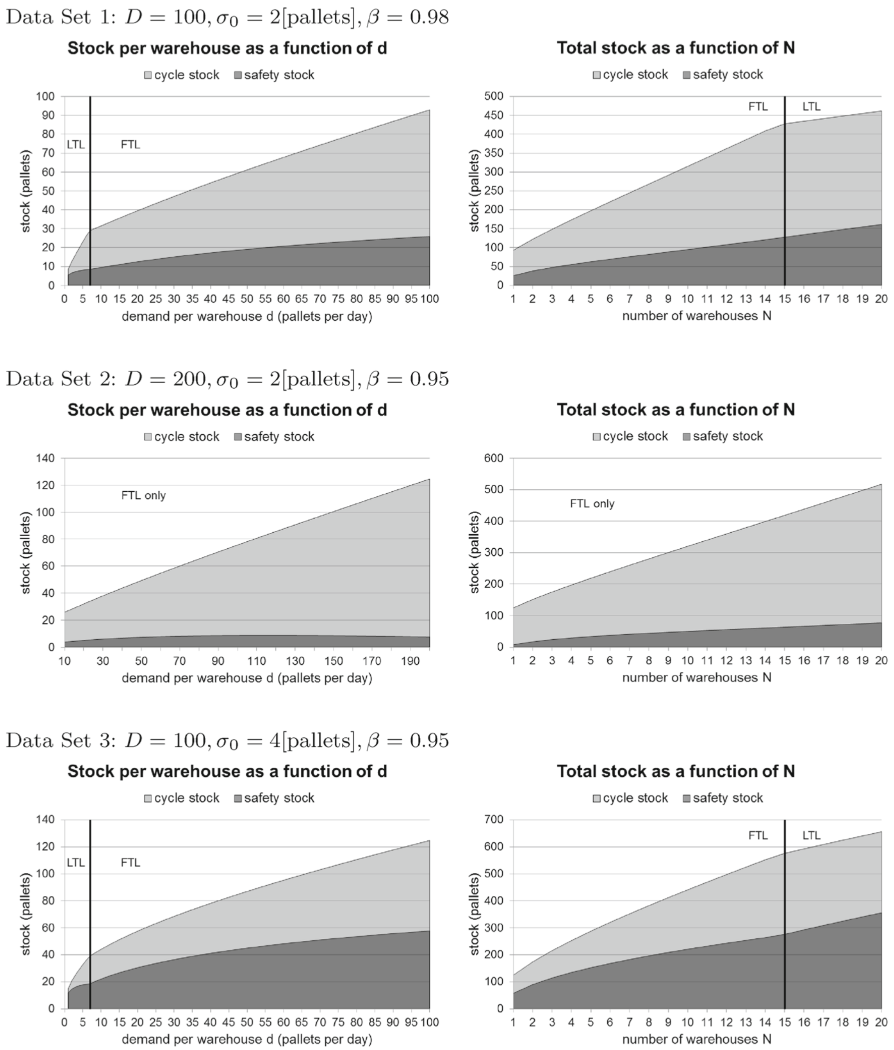

Fig. 5 Safety stock and cycle stock for periodic review and normally distributed demand. Common data: $Q = 34$ [pallets], $L = 2$, $t_{\max} = 5$ [days]

$|r' - r| < 1$ was observed, where r ranged from 33 to 608. Thus, the approximation

$$S^{\mathrm{T}}(N) + C^{\mathrm{T}}(N) = H\left(\frac{(1-\beta)D}{N\sqrt{\sigma_0^2(L+1)\frac{D}{N} + \frac{q^2}{12}}}\right) N\sqrt{\sigma_0^2(L+1)\frac{D}{N} + \frac{q^2}{12}} + \frac{1}{2}D \qquad (12)$$

can be used for analyzing the stock function.

The trend of the total stock (12) in the FTL case depends on the range of N. For small N the first term under the square root dominates so that the total stock proceeds like $h_2(N)$ (see Appendix A.4) with a concave, nearly linear increase. For large N the second term dominates so that the total stock proceeds like $h_3(N)$ (see Appendix A.5) with a convex increase. In the numerical analysis, the total stock indeed exhibits an inflection point at $N = 10$ for Data Set 1 and at $N = 15$ for Data Set 2, whereas it is concave in the whole FTL range for Data Set 3.

On the whole, the total stock function $S^{\mathrm{T}}(N) + C^{\mathrm{T}}(N)$ exhibits a similar pattern as for continuous review, but of course on a higher level.

7 Conclusions

The following conclusions can be drawn from the analysis presented in this paper:

1. The total safety stock that ensures a given fill rate in a system with N parallel warehouses crucially depends on the relationship between the replenishment size q and the demand d per warehouse. The Square Root Law for the safety stock is only valid if q is proportional to $\sqrt{d}$, as is the case for the EOQ. But the EOQ is usually much higher than an FTL and consequently the optimal replenishment quantity is an FTL. For small demands, if an FTL would cause too long cycle times, the replenishment size has to be reduced to a smaller LTL quantity proportional to d.
2. The total safety stock $S^{\mathrm{T}}(N)$ and the total cycle stock $C^{\mathrm{T}}(N)$ have to be considered together when estimating the impact of reducing the number N of warehouses. For continuous review (r, q) control, $S^{\mathrm{T}}(N)$ and $C^{\mathrm{T}}(N)$ behave complementarily in the two cases: In the FTL case $S^{\mathrm{T}}(N)$ is nearly constant and $C^{\mathrm{T}}(N)$ increases linearly, in the LTL case $S^{\mathrm{T}}(N)$ increases stronger than $\sqrt{N}$ and $C^{\mathrm{T}}(N)$ is constant. The total stock $S^{\mathrm{T}}(N) + C^{\mathrm{T}}(N)$ exhibits a rather linear trend. These result which have been derived analytically in case of normally distributed demand can as well be observed for Gamma distributed demand. Also for periodic review (r, nq) control the total stock behaves similarly.
3. In case of positive demand correlation, centralization may even increase the total safety stock significantly. Nevertheless, centralization still remains advantageous due to the reduced cycle stock.
4. The Square Root Law, in the FTL area, overestimates the effect of centralization for the safety stock and underestimates it for the total stock because the latter effect consists mainly in the linear decrease of the cycle stock.

The analysis can be easily extended to the case of unequal demand per warehouse, using the stock function $S(d)$. In the multi-product case, the FTL quantity has to be allocated to the products by means of an appropriate rule. The extension to multi-echelon distribution systems requires further research, combining concepts of multi-echelon inventory control with those of centralization.

Appendix

The analysis of the stock functions $S^{\mathrm{T}}(N)$ is based on the following approximation to the normal loss integral $G(k) \approx R(k)$ which allows an analytic inversion $H(x) = G^{-1}(x) \approx R^{-1}(x)$: Let

$$G(k) = \frac{1}{\sqrt{2\pi}} \exp(-ak^2 - bk) \quad \text{with parameters } a, b > 0. \tag{13}$$

The equation $x = G(k)$ or $k = G^{-1}(x)$ is equivalent to the quadratic equation in k $ak^2 + bk + \ln(\sqrt{2\pi}x) = 0$ with the solution

$$k = \frac{1}{2a}\left(-b + \sqrt{b^2 - 4a(\ln\sqrt{2\pi} + \ln x)}\right).$$

Hence the inverse to $G(k)$ is

$$H(x) = G^{-1}(x) = -A + \sqrt{B - \frac{1}{a}\ln x} \quad \text{with } A = \frac{b}{2a}, B = A^2 - \frac{1}{a}\ln\sqrt{2\pi}. \tag{14}$$

The parameters a, b are fitted so that, for any given $k \in [0, 4]$, $k' = H(R(k))$ coincides with k as well as possible. This is done by minimizing the sum of squares $(\frac{k'}{k} - 1)^2$ over $k = 0.03, 0.1, 0.2, \ldots, 3.9, 4.0$ resulting in

$$a = 0.36121504, \quad b = 1.22377537, \quad A = 1.69397068, \quad B = 0.32551597 \tag{15}$$

The resulting accuracy of the approximation is $G(0) = R(0)$ due to the definition (13) and $-0.0016 \le G(k) - R(k) \le 0.001$ for $0 \le k \le 4$, and $-0.0171 \le H(x) - R^{-1}(x) \le 0.00227$ for $R(4) \le x \le R(0)$.

The function $H(x)$ has the following properties:

A.1 $H(x) > 0$

Proof For $0 < x < \frac{1}{\sqrt{2\pi}}$ we have, using (14), $H(x) = -A + \sqrt{B - \frac{1}{a}\ln x} > -A + \sqrt{B + \frac{1}{a}\ln\sqrt{2\pi}} = -A + \sqrt{A^2} = 0$.

A.2

The function $h_1(x) = H(cx^\theta)x^{1-\theta}$ with constants $c > 0$ and $0 < \theta < 1$ is positive for $0 < cx^\theta < \frac{1}{\sqrt{2\pi}}$ and attains a maximum. It is concave if $\theta \le 0.5$.

Proof $h_1(x) = (-A + w)x^{1-\theta}$ with $w = \sqrt{B - \frac{1}{a}\ln c - \frac{\theta}{a}\ln x} > A$ due to A.1 and $\frac{\mathrm{d}}{\mathrm{d}x} w = -\frac{\theta}{2axw}$. The derivatives are

$$h_1'(x) = \frac{1-\theta}{x^\theta}(-A + w) - \frac{\theta}{2ax^\theta w} = \frac{1}{x^\theta}\left((1-\theta)(-A+w) - \frac{\theta}{2aw}\right)$$

and

$$\begin{aligned} h_1''(x) &= -\frac{\theta}{x^{1+\theta}}\left((1-\theta)(-A+w) - \frac{\theta}{2aw}\right) + \frac{1}{x^\theta}\left(\frac{-\theta(1-\theta)}{2axw} - \frac{\theta^2}{4a^2xw^3}\right) \\ &= -\frac{\theta}{x^{1+\theta}}\left((1-\theta)(-A+w) + \frac{1-2\theta}{2aw} + \frac{\theta}{4a^2w^3}\right). \end{aligned} \tag{16}$$

Therefore, $h_1''(x) < 0$ for $\theta \leq 0.5$ and hence the concavity of $h_1(x)$.
$h_1'(x) = 0$ holds if $(1-\theta)(-A+w) - \frac{\theta}{2aw} = 0$ or $w^2 - Aw - \frac{\theta}{2a(1-\theta)} = 0$.

From the solution of the last equation, $w_0 = \frac{1}{2}(A + \sqrt{A^2 + \frac{2\theta}{a(1-\theta)}})$, the solution of $h_1'(x) = 0$ is obtained, using the definition of w, as $x_0 = [\frac{1}{c}\exp(a(B - w_0^2))]^{\frac{1}{\theta}}$.

This is a global maximum for $0 < \theta \leq 0.5$. For $0.5 < \theta < 1$ it is at least a local maximum because for $w = w_0$

$$(1-\theta)(-A+w) + \frac{1-2\theta}{2aw} > (1-\theta)(-A+w) - \frac{\theta}{2aw} = 0$$

and hence, according to (16), $h_1''(x_0) < 0$.

A.3

For $\theta = 0.5$, $h_1(x)$ attains the maximum approximately at $x_0 \approx \frac{1}{(6c)^2}$ with the value $h(x_0) \approx \frac{1}{10c}$. In the *indifference area* $[0.4x_0, 2x_0]$, the value of $h_1(x)$ is less than 10 % below the maximum.

Proof For $\theta = 0.5$, we have, using the parameter values (15), $w_0 = \frac{1}{2}(A + \sqrt{A^2 + \frac{2}{a}}) = 2.2967$ and $x_0 = [\frac{1}{c}\exp(a(B - 5.2747))]^2 = (\frac{0.1673}{c})^2$, and $h_1(x_0) = \frac{0.1673}{c}H(0.1673) = \frac{0.1008}{c}$.

Let $y_1 = c\sqrt{0.4x_0} = \sqrt{0.4}\cdot 0.1673 = 0.1058$, $y_2 = c\sqrt{2x_0} = \sqrt{2}\cdot 0.1673 = 0.2366$. Then, $h_1(0.4x_0) = H(y_1)\frac{y_1}{c} = \frac{0.09142}{c} = 0.907h_1(x_0)$ and $h_1(2x_0) = H(y_2)\frac{y_2}{c} = \frac{0.09073}{c} = 0.9h_1(x_o)$. Due to the concavity of $h_1(x)$, this proves the property of the indifference area.

A.4

The function $h_2(x) = H(\frac{c}{\sqrt{x}})\sqrt{x}$ is positive and concave for $x > 2\pi c^2$ and increases asymptotically stronger than $\sqrt{x}$, but weaker than linearly.

Proof $h_2(x) = (-A + w)\sqrt{x}$ with $w = \sqrt{B - \frac{1}{a}\ln c + \frac{1}{2a}\ln x}$. $h_2''(x) < 0$ is proved analogously to $h_1''(x) < 0$ in A.2. For $x \to \infty$, $\frac{h_2(x)}{\sqrt{x}} = -A + w$ tends to infinity, but $\frac{h_2(x)}{x} = \frac{1}{\sqrt{x}}(-A + w)$ tends to zero.

A.5

The function $h_3(x) = H(\frac{c}{x})x$ is positive and convex for $x > c\sqrt{2\pi}$ and increases asymptotically stronger than linearly.

Proof $h_3'(x) = -A + w + \frac{1}{2aw}$ and $h_3''(x) = \frac{1}{2axw}(1 - \frac{1}{2aw^2}) > 0$, because $2aw^2 > 2a\,A^2 = 2.07$. This proves the convexity of $h_3(x)$. The asymptotic trend is obvious.

References

Ballou RH (2005) Expressing inventory control policy in the turnover curve. J Bus Logist 26(2):143–164

Bretzke WR (2010) Logistische Netzwerke, 2nd edn. Springer, Berlin

Chopra S, Meindl P (2007) Supply chain management: strategy, planning and operation. Pearson Education, Upper Saddle River

Christopher M (2005) Logistics and supply chain management: creating value-adding networks, 3rd edn. Prentice Hall Financial Times, Harlow

Croxton KL, Zinn W (2005) Inventory considerations in network design. J Bus Logist 26(1):149–168

Diks EB, de Kok AG, Lagodimos AG (1996) Multi-echelon systems: a service measure perspective. Eur J Oper Res 95(2):241–263

Eppen G (1979) Effects of centralization on expected costs in a multi-location newsboy problem. Manag Sci 25:498–501

Erlebacher SJ, Meller RD (2000) The interaction of location and inventory in designing distribution systems. IIE Trans 32:155–166

Evers PT (1995) Expanding the square root law: an analysis of both safety and cycle stocks. Logist Transp Rev 31(1):1–20

Evers PT, Beier FJ (1993) The portfolio effect and multiple consolidation points: a critical assessment of the square root law. J Bus Logist 14(2):109–125

Evers PT, Beier FJ (1998) Operational aspects of inventory consolidation decision making. J Bus Logist 19(1):173–198

Fleischmann B (1993) Designing distribution systems with transport economies of scale. Eur J Oper Res 70:31–42

Fleischmann B, Kopfer H, Sürie C (2015) Transport planning for procurement and distribution. In: Stadtler H, Kilger C, Meyr H (eds) Supply chain management and advanced planning. Springer, Berlin, pp 225–240

Hadley G, Whitin TM (1963) Analysis of inventory systems. Prentice-Hall, Englewood Cliffs

Hartel DH (2009) Consulting und Projektmanagement in Industrieunternehmen. Oldenbourg, München

Maister DH (1976) Centralisation of inventories and the “Square Root Law”. Int J Phys Distrib 6(3):124–134

Miranda PA, Garrido RA (2004) Incorporating inventory control decisions into a strategic network design model with stochastic demand. Transp Res Part E: Logist Transp Rev 40(3):183–207

Ozsen L, Coullard CR, Daskin MS (2008) Capacitated warehouse location model with risk pooling. Nav Res Logist 55(4):295–312

Schwarz LB (1981) Physical distribution: the analysis of inventory and location. AIIE Trans 13(2):138–150
Shapiro JF, Wagner SN (2009) Strategic inventory optimization. J Bus Logist 30(2):161–173
Shen ZJM (2007) Integrated supply chain design models: a survey and future research directions. J Ind Manag Optim 3(1):1–27
Snyder LV, Daskin MS, Teo CP (2007) The stochastic location model with risk pooling. Eur J Oper Res 179(3):1221–1238
Stulman A (1987) Benefits of centralized stocking for the multi-centre newsboy problem with first come, first served allocation. J Oper Res Soc 38(9):827–832
Tempelmeier H (2011) Inventory management in supply networks: problems, models, solutions, 2nd edn. Books on Demand, Norderstedt
van Houtum GJ, Inderfurth K, Zijm WHM (1996) Materials coordination in stochastic multi-echelon systems. Eur J Oper Res 95(1):1–23