10 The Impact of the Number of Warehouses on Inventories in a Distribution System

Bernhard Fleischmann, University of Augsburg, bernhard.fleischmann@wiwi.uni-augsburg.de

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 last 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, 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 presents a new analysis of this case and a relationship between the demand to be served by a warehouse and the necessary safety and cycle stock. As a consequence, the impact of the number of parallel warehouses in a distribution system on the total stock can be derived. This paper concerns a core subject of the International Workshop on Distribution Logistics (IWDL), founded by Jo van Nunen and the author in 1994 and continued in a series of more than 10 meetings, the last of which has been organized by Jo in s'Hertogenbosch in May 2009. The IWDL united international experts in location theory, network design, inventory theory, routing and warehousing. Thanks to Jo's strong orientation towards the practice of distribution logistics, the IWDL stimulated application oriented research also on interrelations between these fields, like network-routing, routing-inventory and network-inventory. This paper falls into the latter area.

10.1 Introduction and Problem Setting

The design of a distribution system usually starts from given plant locations, customer locations or areas with a certain demand and service requirements concerning 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 (Hagdorn and Van Nunen 1998, Fleischmann 1998). They define the distribution network which consists of plants, warehouses, transshipment points, customers, and arcs linking these nodes. 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.

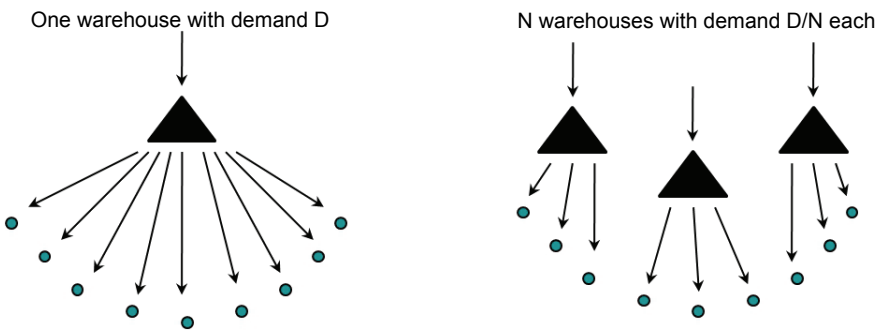

Figure 1. Distribution structures considered

Most manufacturers outsource the distribution task to a logistics service provider (LSP) who is then responsible for all operations. He usually owns the warehouses and transshipment locations and the information/communication system, but not the inventory. The main question remaining for the manufacturer is then, 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 stockpoints. In the last 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 in a 48h or 72h lead time through a stockless transportation network. This so-called "Euro Logistics" concept is practiced for instance by Bosch Power Tools, by OSRAM, and by the pencil producer Staedtler.

The primary motivation for reducing the number of stockpoints in a distribution system is the goal to cut down inventories. But what is the exact relationship between the number of parallel stock points in a distribution system and the total stock? That is the question that will be analyzed in this paper.

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

1. There are N parallel stock points supplying customers with a given demand of a single product.
2. The total demand per unit of time is stochastic with a fixed expectation D and does not depend on N. It is composed of many small uncorrelated customer demands.
3. There is a fixed allocation of the customers to the stock points so that the demand rate per stock point is $d = D/N$.
4. The fixed lead time L for replenishing the stock points is the same for all stock points and independent of N.
5. The replenishment of the stock points is done by trucks in full truck load (FTL) of the size Q, as long as Q does not exceed the average demand of a given maximal cycle time t_{max} (for instance one week).
6. Every stock point must satisfy a service level constraint in form of a fill rate β which is independent of N.

Some of these assumptions will be relaxed later on, such as the single product case (Assumption 1) and the equal allocation of demand (Assumption 3). Assumption 6 is criticized by Tempelmeier (2006, p. 154) because the distances from the warehouses to the customers increase with decreasing N and therefore the lead time for customer supply is likely to increase as well. But in a modern Euro Logistics system, as explained above, the service time for the customers is nearly independent of N.

The purpose of this paper is to analyze the total inventory in the distribution system as a function of N. For that it is necessary to distinguish safety stock and cycle stock. The number of warehouses may also have an impact on the stock in transit, but this stock can be included in the cost of transportation. As no optimization of N is considered, the transportation cost is disregarded. It is well known that decreasing N implies lower transportation costs from the warehouses to the customers, whereas the warehouse replenishment costs only increase slightly, if at all, as long as the FTL mode can be kept (Chopra and Meindl, 2007, p.78f).

The rest of the paper is organized as follows: In the next Section, the usual answer that is found in literature to our question, the "Square-Root Law", and its limited justification are analyzed. Section 3 presents a continuous review model leading to a new relationship between the total inventory and the number of warehouses. Section 4 modifies the results for a more realistic periodic review model. Section 5 considers extensions for several products, Section 6 draws conclusions and points out open questions.

10.2 The Square-Root Law

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

$S^T(N)$ total safety stock

$C^T(N)$ total average cycle stock.

Due to d = D/N we calculate $S^T(N)$ and $C^T(N)$ as

$$S^T(N) = N \cdot S(D/N) \quad \text{and} \quad C^T(N) = N \cdot C(D/N) \,. \tag{1}$$

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

$$S(d) = A\sqrt{d} \quad \text{with some constant A}$$

and hence

$$S^T(N) = A'\sqrt{N} \quad \text{with a constant } A' = A\sqrt{D} \,.$$

It appears in many textbooks (e.g. Bretzke 2008, Chopra and Meindl 2007, Christopher 2005, Tempelmeier 2006) and is used in many articles on the design of distribution systems (e.g. Erlebacher and Meller 2000, Ozsen et al. 2008, Shen 2007, Snyder et al. 2007). The following justifications of the Square Root law in literature are often referred to:

The multilocation newsboy model of Eppen (1979):

Eppen considers N warehouses with normally distributed demand in a single period. In the special case of uncorrelated equal demands, the expected demand is d = D/N and the standard deviation $\sigma = \sigma_D/\sqrt{N} = \sigma_0\sqrt{d}$ where σ_D is the standard deviation of the total demand and $\sigma_0 = \sigma_D/\sqrt{D}$. It is well known that the optimal safety stock in a newsboy model is k σ with a safety factor

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

where h is the unit holding cost and p the penalty cost. Hence

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

satisfies the Square-Root Law. In addition, Eppen shows that this is also true for the total holding and penalty cost. Note that in case of a constant safety factor k, the Square-Root Law 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 variance is proportional to $1/\sqrt{d}$ and to $\sqrt{N}$.

However, this model is not compatible with the problem setting in Section 1: 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

$$R(k) = \int_k^\infty (x-k)\varphi(x)dx$$

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

α-service-level constraint:

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 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 2003, Chapter 11.4).

β-service-level constraint and EOQ

L.B. Schwarz (1981) derives the Square-Root Law for a continuous-review (r,q) policy with backorders where the replenishment quantity q is determined as an economic order quantity (EOQ) and a constant safety factor k. Then, with the standard deviation of the total lead-time demand σ_{LD}, the safety stock is $S(d) = k\,\sigma_L\sqrt{d}$, where $\sigma_L = \sigma_{LD}/\sqrt{D}$. But in order to yield a given fill rate β, k has to satisfy (Silver et al. 1998, p. 268)

$$k = R^{-1}\left(\frac{1}{\sigma_L\sqrt{d}}(1-\beta)\right). \tag{3}$$

This formula holds, if the expectation of a stockout immediately after replenishment can be neglected. Thus, the safety factor k depends on d and on q. The only situation where k is constant is when q is an EOQ, because then $q = a\sqrt{d}$ with some constant a . In this case the Square-Root Law holds not only for the safety stock S(d) and $S^T(N)$, but also for the cycle stock $C(d) = ½\, q = ½\, a\sqrt{d}$ and for $C^T(N)$.

However, when the warehouse replenishment is done by trucks, q is restricted by the truck capacity Q which is in most cases much smaller than the EOQ, as can be seen from realistic values of the following parameters:

Q : 12 tons (net weight of consumer goods)

v : value of the good in €/kg , $v \leq 100$

i : interest rate = 0.08

F : cost per trip, $F \geq 500$ € for a long-distance trip

t : demand per week in number of FTLs, i.e. $d = t\,Q$, $t \geq 2$.

The EOQ is larger than Q if, for the FTL replenishment, the transport cost exceeds the holding cost per week, i.e. if $t\,F \geq ½ \cdot 12000 \cdot v \cdot i/52 \approx 10\,v$ or $v \leq t\,F/10$, which is satisfied for the above typical ranges of the parameters.

Therefore, the EOQ assumption is not realistic. The replenishment quantities are FTLs or, for very small warehouses, even smaller. This situation will be considered in the following section.

10.3 A continuous review model

As before, let the warehouses be replenished using a continuous-review (r, q) policy. The demand is composed of many very small units so that an "undershoot" under the recorder 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 a realistic model for that purpose. Instead, Assumption 5 (see Section 1) leads to the following relationship:

$$q = \begin{cases} Q, & \text{if } d \geq Q / t_{max} \quad \text{(FTL case)} \\ d\, t_{max}, & \text{if } d \leq Q / t_{max} \quad \text{(less than truckload, LTL case)} \end{cases} \tag{4}$$

Using (3) and $C(d) = ½\, q$ we get the stock of a single warehouse as a function of d:

FTL case:

$$S(d) = R^{-1}\left(\frac{Q}{\sigma_L\sqrt{d}}(1-\beta)\right)\sigma_L\sqrt{d} \quad \text{and} \quad C(d) = ½\, Q$$

LTL case

$$S(d) = R^{-1}\left(\frac{\sqrt{d}\, t_{max}}{\sigma_L}(1-\beta)\right)\sigma_L\sqrt{d} \quad \text{and} \quad C(d) = ½\, d\, t_{max}.$$

The total stock for N warehouses is than, using (1) and $\sigma_{LD} = \sigma_0\sqrt{LD}$:

FTL case ($N \leq D\, t_{max}/Q$)

$$S^T(N) = R^{-1}\left(\frac{Q\sqrt{N}}{\sigma_{LD}}(1-\beta)\right)\sigma_{LD}\sqrt{N} \quad \text{and} \quad C^T(N) = ½\, NQ.$$

LTL case ($N \geq D\, t_{max}/Q$)

$$S^T(N) = R^{-1}\left(\frac{D\, t_{max}}{\sigma_{LD}\sqrt{N}}(1-\beta)\right)\sigma_{LD}\sqrt{N} \quad \text{and} \quad C^T(N) = ½\, D\, t_{max}$$

Figure 2 shows, for a realistic data setting, the safety stock and cycle stock functions for a single warehouse and in total.

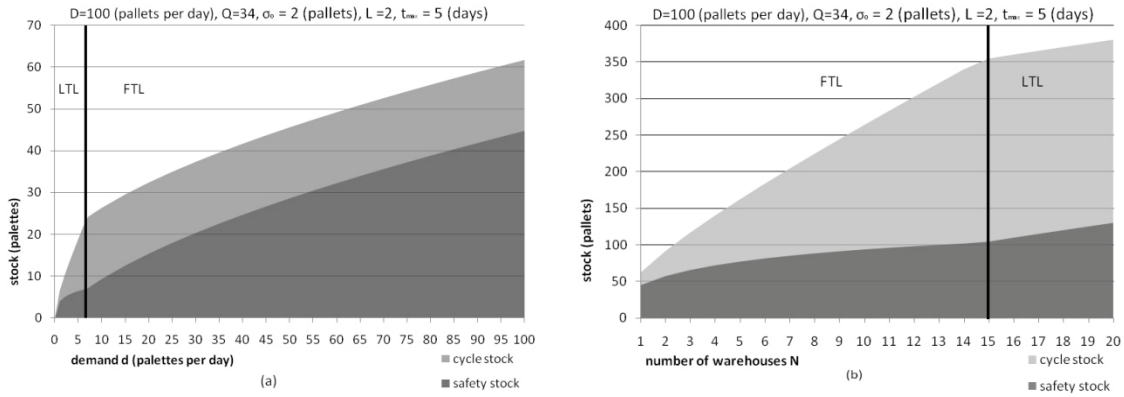

Figure 2. Safety stock and cycle stock for continuous review, (a) as a function of the demand d per warehouse, (b) in total for N warehouses

Surprisingly, in the FTL case, the total safety stock $S^T(N)$ is not a monotonic function of N because $R^{-1}(\ )$ is monotonically decreasing. An analysis of the function

$$G(x) = R^{-1}\left(c\sqrt{x}\right)\sqrt{x}$$

shows that it has a maximum at approximately $x_0 = 1/(6c)^2$ and that this maximum is rather flat: The values in the interval $[0.4\ x_0,\ 2x_0]$ differ from the maximum by less than 10%. Thus, $S^T(N)$ is nearly constant, as can be seen in Figure 2b.

This property contradicts the Square-Root Law and leads to the question: Why is there no 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 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, the reduction of the cycle stock, not of the safety stock. In our example, the cycle stock would decrease from 5/2 Q to ½ Q.

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 $\sqrt{N}$.

10.4 A periodic review model

In a continuous review model, as considered before it is assumed that the replenishments of a warehouse can occur at any time on a continuous time axis. It is more realistic to assume that the trucks arrive only once per day within a small time window. This situation can be modeled by a (r,nq) policy: Whenever the inventory position s is below the reorder point r, an order of nq ($n \geq 1$, integer) is placed at the end of the day such that s is raised into the interval [r, r + q). As before, we set $q = \min(Q, d\ t_{max})$. The lead time L is integer, the order put on a day x arrives and is available on day x+L, and the expected demand per day is d with the standard deviation $\sigma = \sigma_0\sqrt{d}$.

For the approximate calculation of the safety stock, two cases have to be distinguished:

Case d ≤ Q:

In this case, the average number of trucks arriving per day is at most 1. Then, the *undershoot* U of the inventory position under the reorder point, just before ordering, has approximately the expectation and variance (see Tempelmeier 2006, p. 76f)

$$E(U) = \tfrac{1}{2}\,(d + \sigma^2/d\,) = \tfrac{1}{2}\,(d + \sigma_0^2)$$

$$V(U) = \tfrac{1}{2}\,\sigma^2(1 - \tfrac{1}{2}\,\sigma^2/d^2) + d^2/12 = \tfrac{1}{2}\,\sigma_0^2\,(d - \tfrac{1}{2}\,\sigma_0^2) + d^2/12\ .$$

The safety stock has to cover the variance of the lead-time demand and of the undershoot and can therefore be calculated as

$$S(d)=R^{-1}\left(\frac{q}{\sigma_{LU}}(1-\beta)\right)\sigma_{LU} \quad \text{with} \quad \sigma_{LU}=\sqrt{\sigma_0^2\, L\, d + ½\, \sigma_0^2(d - ½\, \sigma_0^2) + d^2/12}\ . \tag{5}$$

The cycle stock is roughly C(d) = ½ q, neglecting the fact that sometimes two trucks may arrive together.

Case d > Q

The above safety stock calculation is only valid, if d is not too large (Tijms 1998, p.61). Therefore, we use another approximation for the case d > Q, i.e. if there is more than one truck arriving per day on average. Following Hadley and Whitin (1963, p. 248), in this case the *overshoot* O of the inventory position over r immediately after placing an order is uniformly distributed in the interval [0,Q), and the safety stock has to cover the variance of the demand in the lead time plus one day (the review period) and of the overshoot. Therefore

$$S(d)= R^{-1}\left(\frac{d}{\sigma_{LO}}(1-\beta)\right)\sigma_{LO} \quad \text{with} \quad \sigma_{LO}=\sqrt{\sigma_0^2(L+1)\, d + Q^2/12}\ . \tag{6}$$

The cycle stock is roughly C(d) = ½ d.

The total stock functions $S^T(N)$ and $C^T(N)$ for N warehouses can be calculated from S(d) and C(d) as before. Figure 3 shows the stock curves including the three cases:

LTL ($d < Q / t_{max}$ or $N > D\, t_{max} / Q$),

FTL1 with up to one truck per day ($d \le Q$ or $N \ge D/Q$), and

FTL2 with more than one truck per day ($d > Q$ or $N < D/Q$).

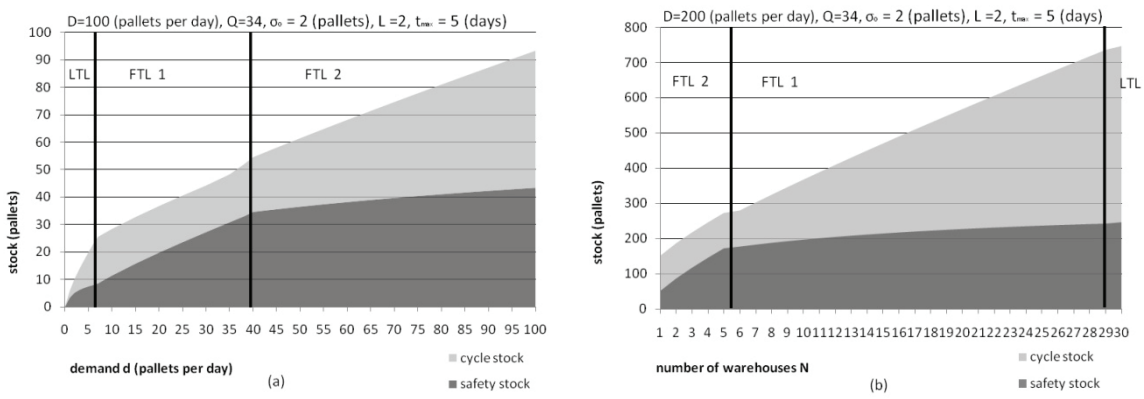

Figure 3. Safety stock and cycle stock for periodic review, (a) as a function of the demand d per warehouse, (b) in total for N warehouses

Since the standard deviations σ_{LU} and σ_{LO} increase with d, approximately like $\sqrt{d}$, the stock curves for the first two cases show the same behavior as in the continuous review model. In the FTL2 case, the average replenishment quantity is d, and therefore the stock curves behave similarly as in the LTL case: The total safety stock increases nearly linearly with N, whereas the total cycle stock is constant and equals ½ D.

The approximations (5) and (6) have been validated in a simulation study. For various date settings, the (r,nq) policy has been simulated with

$$r = \begin{cases} S(d) + d\,L + ½\,(d + \sigma_0^2) & \text{for } d \le Q \\ S(d) + d\,(L + 1) - ½\,Q & \text{for } d > Q \end{cases}$$

where S(d) has been determined by (5) and (6), respectively. In runs over 1000 days with 10 repetitions, the observed fill rate showed no significant deviation from the target fill rate. The same was true for the observed average stock before a replenishment, compared with S(d), and for the observed undershoot and overshoot, compared with the theoretical values. Only the observed average

stock was slightly higher than S(d) + C(d), because the approximation C(d) = ½ min(q,d) underestimates the cycle stock.

10.5 Extensions

Several products

In a distribution warehouse, usually a large range of products is stored. For products coming from the same source in the same truck, a joint replenishment strategy is required. However, in case of stationary demand and FTL replenishment, it is optimal to include every product in every truck load (see Fleischmann 2000).

Consider the jointly replenished products i = 1,...,m with the total demand rate D_i and $D = \Sigma_i D_i$. The product mix $a_i = D_i / D$ is assumed to be the same in every demand area so that the demand at one of N warehouses is $d_i = D_i / N$ or $d_i = a_i d$. Then, the relationship between the average replenishment quantity q_i of product i and the demand d_i is the same as stated for q and d in the previous sections:

$$
\begin{array}{lll}
q_i = d_i\, t_{max}\,, & \text{if } d \le Q / t_{max} & \text{(LTL)} \\
\quad a_i\, Q\,, & \text{if } Q / t_{max} \le d \le Q & \text{(FTL1)} \\
\quad d_i\,, & \text{if } d > Q & \text{(FTL2)}
\end{array}
$$

Therefore, all stock functions derived from Sections 3 and 4 remain valid for the single products.

Unequal demand of the warehouses

The calculation of the total stock from the stock functions S(d) and C(d) according to (1) can be easily extended to the case of unequal demands of the warehouses j = 1, …, N : let d_j the demand of warehouse j and $\Sigma_j d_j = D$. Then the total safety and cycle stock is

$$S^T(N) = \sum_{j=1}^{N} S(d_j),\ \ C^T(N) = \sum_{j=1}^{N} C(d_j)$$

where S(d) and C(d) can be determined by one of the models of Sections 3 or 4.

10.6 Conclusions

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

1. The total safety stock that is required in a distribution system with N parallel warehouses to satisfy a given fill rate, crucially depends on the relationship between the replenishment size q and the demand d per warehouse. The common Square-Root Law of the safety stock is only valid if q is determined as an EOQ.
2. In the more realistic FTL situation, three cases have to be distinguished: FTL2 with more than one truck per day on average, FTL1 with up to one truck per day and LTL with $q = d\, t_{max}$. In all three cases, the total safety stock $S^T(N)$ is far from satisfying the Square-Root-Law.
3. The safety stock $S^T(N)$ and the cycle stock $C^T(N)$ have to be considered together if the impact of reducing the number N of warehouses, for instance by centralization, has to be estimated. $S^T(N)$ and $C^T(N)$ behave complementarily in the three cases: Either $S^T(N)$ increases stronger than $\sqrt{N}$ and $C^T(N)$ is constant (cases LTL and FTL2) or $S^T(N)$ is nearly constant and $C^T(N)$ increases linearly (case FTL1).
4. Roughly, over all three cases, $S^T(N) + C^T(N)$ behaves like a + b N where a is greater than the minimal cycle stock ½ D.

An important open question is the extension of this analysis on a multi-echelon inventory system: What is the effect of the variation of the number of stock points on a certain echelon?

References

Bretzke, W.-R. (2008). Logistische Netzwerke. Berlin: Springer.

Chopra, S., & Meindl, P. (2007). Supply Chain Management: Strategy, Planning and Operation. Upper Saddle River: Pearson Education.

Christopher, M. (2005). Logistics and Supply Chain Management: Creating Value-Adding Networks. Harlow: Prentice Hall Financial Times.

Eppen, G. (1979). Effects of Centralization on Expected Costs in a Multi-Location Newsboy Problem, Management Science, 25, 498-501.

Erlebacher, S. J., & Meller, R. D. (2000). The Interaction of Location and Inventory in Designing Distribution Systems, IIE Transactions, 32, 155-166.

Fleischmann, B. (1998). Design of Freight Traffic Networks. In Fleischmann, B., van Nunen, J. A. E. E., Speranza, M. G., & Stähly, P. (Eds.), Advances in Distribution Logistics (pp. 55-81). Berlin: Springer.

Fleischmann, B. (2000). Transport and Inventory Planning with Discrete Shipment Times. In Speranza, M. G., & Stähly, P. (Eds.), New Trends in Distribution Logistics (pp. 159-178). Berlin: Springer.

Hadley, G., & Whitin, T. M. (1963). Analysis of Inventory Systems. Englewood Cliffs: Prentice-Hall.

Hagdorn-van der Meijden, L. & van Nunen, J.A.E.E. (1994). Strategic Decision Making for Logistics Network Design. In Fleischmann, B., van Nunen, J.A.E.E., Speranza, M.G., & Stähly, P. (Eds.), Advances in Distribution Logistics (pp. 83-109). Berlin: Springer.

Ozsen, L., Coullard, C. R., Daskin, M. S. (2008). Capacitated Warehouse Location Model with Risk Pooling. Naval Research Logistics, 55(4), 295-312.

Shen, Z.-J. M. (2007). Integrated Supply Chain Design Models: A Survey and Future Research Directions, Journal of Industrial and Management Optimization, 3(1), 1-27.

Schwarz, L. B. (1981). Physical Distribution: The Analysis of Inventory and Location. AIIE Transactions, 13(2), 138-150.

Silver, E.A, Pyke, D. F., & Peterson, R. (1998). Inventory Management and Production Planning and Scheduling. New York: Wiley.

Snyder, L. V., Daskin, M. S. & Teo, C.-P. (2007). The Stochastic Location Model with Risk Pooling. European Journal of Operational Research, 179 (3), 1221-1238.

Stulman, A. (1987). Benefits of Centralized Stocking for the Multi-Centre Newsboy Problem with First Come, First Served Allocation. Journal of the Operational Research Society, 38(9), 827-832.

Tempelmeier, H. (2006). Inventory Management in Supply Networks: Problems, Models, Solutions. Norderstedt: Books on Demand.

Tijms, H. C. (1998). Stochastic models: An Algorithmic Approach. Chichester: Wiley.