An overview of inventory systems with several demand classes

Marcel J. Kleijn   Rommert Dekker

September 1998
Econometric Institute Report 9838/A

Abstract

In this chapter we discuss inventory systems where several demand classes may be distinguished. In particular, we focus on single-location inventory systems and we analyse the use of a so-called critical level policy. With this policy some inventory is reserved for high-priority demand. A number of practical examples where several demand classes naturally arise are presented, and the implications and modelling of the critical level policy in distribution systems are discussed. Finally, an overview of the literature on inventory systems with several demand classes is given.

Keywords. Inventory, demand classes, rationing, critical level

1 Introduction

Inventory systems often face customer demand for many different products. The demand characteristics may vary from product to product and therefore an inventory manager will generally apply a customised policy for every product. However, in most cases, all customer demand for a single product is handled in a uniform way. Although the order sizes may vary greatly and some orders can be handled in a different way than others, each unit demanded is considered equally important. In this chapter we will focus on the situation where this is not the case, i.e. demand for a single product may be classified into different levels of priority. In particular, we tackle the problem where some customers have a higher stockout cost and/or required minimum service level than others. Customer differentiation, i.e. distinguishing classes of customers and giving them different service, has not received much attention in inventory control theory. The topic does not appear in several reviews on the area (e.g. Veinott (1966), Chikán (1990), Lee & Nahmias (1993) and Porteus (1990)). All listed papers consider one type of customers only, and thus all demand is assumed to be equal. Also in well-known books on logistics and inventory control (e.g. Ballou (1992) and Silver, Pyke & Peterson (1998)) the situation of different demand classes is not mentioned. We think that considering multiple demand classes in inventory control is an interesting extension of existing theory which has many practical applications. In Section 2 we will list four examples of inventory systems where different demand classes with different stockout cost and/or required service levels arise naturally. Thereafter, a policy to efficiently handle different demand classes in inventory systems is introduced and its characteristics are discussed. The problem of determining the optimal policy parameters is discussed in Section 4. Finally, we present an overview of the existing literature in Section 5 and summarise the contents of this chapter in the last section.

2 Examples of multiple demand classes

In this section we will discuss real-life examples[1] where multiple demand classes for a single product arise naturally.

Example 1:

The first example deals with the inventory of so-called *rotables* in the airline industry. A rotable is a part of an aircraft that can be repaired after it breaks down. A major airline has founded an independent company to take care of the inventory of serviceable parts. This company now faces different types of demand for these serviceable parts. The most important are the requests of the major airline. There is a contractual agreement stating that in 95% of the times the company should supply a part within 24 hours. The company also has contractual agreements with other airlines, with similar service standards. Some airlines not having a contractual agreement with the company, the so-called 3rd parties, also request serviceable parts from time to time. In such a case, the company may decide to sell a part, to loan it, or to exchange it for an unserviceable (broken down) part. In all cases, the profit to the company will be different. The company wants to analyse the possibility of having some rules to decide whether or not to deliver a request from a 3rd party. So far, such decisions have been made based on the knowledge and experience of the inventory manager. An advantage of having a decision rule is that less experienced people can also do the job.

[1]We acknowledge Ruud Teunter and Mirjam Maatman for drawing our attention to Example 1.

Example 2:
The second example occurs in a two-echelon inventory system where the highest echelon (say warehouse) faces demand both from customers and from lower echelon stocking points (say retailers). Such a situation may arise if the break quantity rule is applied, and large orders at the retailers are routed to the warehouse (Kleijn & Dekker (1998)). A stockout for customer demand at the warehouse will induce a large stockout cost, whereas a stockout for a retailer's replenishment order merely causes a delay in the replenishment lead time, which usually yields a much lower cost. Therefore, customer demand would normally be considered more important than retailer demand. This example is illustrated in Figure 2.1, where the factory faces demand both from customers and the warehouse. Note that the example can be extended to general multi-echelon inventory systems. A similar example was mentioned by Cohen, Kleindorfer & Lee (1988). They described a multi-echelon system where the retailer could place normal replenishment orders and, in case of a stockout, emergency orders at the warehouse. The latter type of orders would receive a higher priority at the warehouse.

Figure 2.1: Illustration of Example 2.

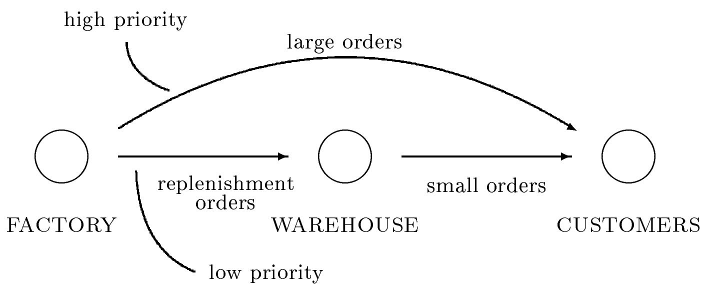

Example 3:
Consider again a two-echelon inventory system, consisting of a warehouse and a number of retailers. If the retailers are located in different countries, it may be desirable at the warehouse level to set different priorities for the retailers. For example, Belgian customers most surely have different needs and expectations than French, Finnish or African customers (Henaux & Semal (1998)).

Example 4:
The final example is related to the previous one, and occurs in spare parts inventory control where an item is used in several types of equipment of varying degrees of criticality. Recently, a case study was done on the inventory control of slow moving spare parts

in a large petrochemical plant (see Dekker, Kleijn & De Rooij (1998)). The management wanted to study the possibility of having equipment criticality determine stock levels. Some equipment in a plant may be critical, while others may be almost redundant. Equipment criticality was defined as the importance of equipment for sustaining production in a safe and efficient way. In the petrochemical plant one could distinguish between vital, essential and auxiliary equipment. Many similar parts were installed in equipment of different criticality and the cost of a stockout depends on the degree of criticality of the equipment in which the part was installed. Hence, the management wished to maintain different service levels for the same part. A related example is mentioned in Ha (1997b), where in an assemble-to-order system a common component is shared by several end-products which have different values to the firm. If, for example, a component is used in both a coffee maker and a television set, then in case of low inventory priority will be given to the television set because it will result in a higher profit to the firm.

Well-known related problems where inventory for a single product is limited and different demand classes are distinguished are found in the health care and airline/hotel industry. In health care, for example, the demand for suitable kidneys generally exceeds its availability and therefore the limited number of suitable kidneys needs to be rationed. In the United States, among those recommended for kidney transplantation there appears to be explicit rationing based on race, sex, age, health condition and income (Greenberg (1991), Held (1988)).
In the travel and leisure industry, which markets space such as seats on airline flights and rooms in hotels, the notion of multiple demand classes has long been recognised. Some customers are willing to pay more for a hotel room than others, and therefore it may be beneficial to refuse the request of a low-price customer in anticipation of a future request of a high-price customer. If customers from the different classes arrive sequentially (first the customers who are willing to pay less) then the optimal policy can be represented as a set of *protection levels*, i.e. the minimum number of rooms reserved for future classes (Robinson (1995)). For traditional inventory systems with several demand classes a similar policy, which we shall refer to as a *critical level policy*, is often used. In the next section the use of this policy is motivated and its implications for the inventory system are analysed.

3 The critical level policy

In the literature on inventory systems with multiple demand classes, the problem of stock rationing is discussed in two different settings. First, a number of authors consider a periodic review situation (see Section 5), where all demand in a period is observed before a rationing decision has to be made. In this case the rationing decision is easy: satisfy demand from the highest priority class first, then from the one-but-highest priority class, and so on. We will consider the continuous review situation, where a rationing decision has to be made at the moment demand occurs.

There are many different ways of rationing inventory among different demand classes with varying stockout cost and/or service requirements. Perhaps the easiest way is the use of separate stockpiles for each demand class. In this way it is very easy to assign a different service level for each class. Also the implementation in practice is extremely easy. However, an important drawback of this method is the loss of economies of scale because no advantage is taken from the so-called *portfolio effect* (Eppen & Schrage (1981)). It is well-known that splitting up the customer demand process will lead to higher inventory cost due to an increasing variability of demand. Nevertheless, this simple policy may outperform a policy where all demand classes are satisfied from a single stockpile. In this case, the highest required service level among the different classes will determine the total stock needed and thus the inventory cost. Although all demand is centralised and the demand variability is reduced, the service level may be unnecessarily high for many demand classes which may lead to a higher cost than in the separate stockpiles situation.

The critical level policy is in essence a mixture of the above two extreme policies. It reserves part of the stock for high-priority demand. In a system with n demand classes and unit demand the policy operates as follows: demand from class j is satisfied from stock on hand if the inventory level exceeds the so-called critical level associated with class j. It is assumed that demand for the highest priority class is satisfied from stock on hand whenever possible, so we have $n-1$ critical levels. If the demand classes are arranged such that class 1 represents the highest priority class, then the set of critical levels will be non-decreasing. If customers can order more than one unit at a time, there are different ways to operate the critical level policy. It may happen that prior to a demand from class j the inventory level exceeds the critical level for this class, but after issueing the customer order the inventory level will drop below the critical level. In that case, the inventory manager needs to decide whether to accept or refuse the order, or maybe to use deliver it partially. In a backorder environment, the critical policy also causes an operational problem with respect to the allocation of incoming replenishment orders. Clearly, whenever upon arrival of a replenishment order there is a backorder for a highest priority customer, it is optimal to use the incoming order to satisfy this backorder. Also, when the stock level exceeds the largest critical level (i.e. the critical level of the lowest priority class), one should first use (if necessary) the incoming order to satisfy any outstanding backorders of the other classes, in order of priority, and then to replenish inventory. However, whenever upon arrival of a replenishment order there are both outstanding orders for one of the other demand classes *and* the inventory level does not exceed the largest critical level, we need to make an allocation decision. Either we satisfy a backorder or we increase the inventory level. In an inventory system with two demand classes, this implies that there are basically two ways of allocating an incoming order in this situation:

1. satisfy backorders for low-priority customers

2. replenish inventory reserved for high-priority customers

For exponentially distributed lead times, Ha (1997a) proved that method 2 is optimal. However, in general this method has the disadvantage that the average backorder length for a low-priority customer may become too large. If method 1 is applied, then this backorder length is limited, at the expense of a lower service level for high-priority customers. In a lost sales environment, all the above complications are not relevant, because there are no backorders. Whenever a customer demand is not satisfied from stock it is lost, so any incoming replenishment order should be used to increase the inventory level.

The critical level policy has a number of implications for the inventory cost and service levels. In general, the service level for the lowest priority customers will decrease and the service level for the highest priority demand class will increase. The effect on the average holding cost depends on the inventory policy used. For example, by applying a critical level policy in the context of a lot-for-lot policy with lost sales the average inventory holding cost will increase, whereas in a lost sales (s, Q) model the inventory holding cost will be reduced (see Dekker, Hill & Kleijn (1997) and Melchiors, Dekker & Kleijn (1998)).

In general, a simple critical level policy is not optimal. An optimal policy would incorporate knowledge about the remaining lead time. For example, if it is known that a replenishment order will arrive within a small amount of time and the inventory level is below the critical level, it may not be optimal to refuse a demand of a low-priority customer. The probability that a high-priority demand will occur before the replenishment order arrives is negligible, and thus the stockout cost for this low-priority customer will not be offset. In a continuous review setting, the optimality of a simple (lead time independent) critical level policy can only be proved for exponentially distributed lead times (Ha (1997a, 1997b)). However, instead of concentrating on sophisticated policies, we will discuss the simple critical level policy, because it is easy for practitioners to understand and this facilitates the practical implementation. We feel that this is a necessary requirement for any policy which is to be used in practice. In reality, an inventory manager can check the remaining lead time and he/she may decide to overrule the refusal of a low-priority customer demand. If the inventory manager can really improve the performance of the inventory policy by overruling the critical level policy, then he/she should be careful when setting a desired service level. If the required service level is e.g. 95% for a certain demand class, it may be more efficient to optimise the policy parameters for a slightly lower required service level, and let the inventory manager increase the service to its desired level by (occasionally) overruling the inventory policy.

4 Determining an optimal critical level policy

In this section we briefly discuss the issue of determining an optimal critical level policy. If the critical level policy is applied within the framework of an existing inventory policy, then this policy is extended by a set of critical levels. As mentioned before, we have $n-1$ critical levels if n is the number of demand classes, because demand from the highest priority class is always satisfied if possible. Consider, for example, an (s, Q) inventory model, where a replenishment order of size Q is placed whenever the inventory position reaches the reorder level s. With a critical level policy, this model is extended to a $(\boldsymbol{c}, s, Q)$ inventory model, where $\boldsymbol{c} := (c_1, \ldots, c_{n-1})$ denotes the set of critical levels. Demand from class j is satisfied from stock on hand if the inventory level exceeds the critical level c_{j-1} for this class. Given the set of critical levels and the other policy parameters, one needs to calculate the average inventory holding and shortage cost and/or the service level for both demand classes.

A possible way to analyse an inventory system with multiple demand classes and a critical level policy is to model it as an inventory system with state-dependent demand. Given the critical levels one can determine the demand process for each level of inventory. Another approach to determine the operating characteristics of such an inventory system is to model the so-called *hitting time* of every critical level. Given these hitting times, one can calculate (or approximate) the expected holding and shortage cost and/or service levels. Observe that the backorder case is much more difficult to analyse than the lost sales situation. The increased complexity arises because there may be both outstanding backorders (for low-priority customers) and a positive inventory level (stock reserved for high-priority customers).

After having derived an expression for the average cost and/or service level, one needs to determine the optimal policy parameters. To determine the optimal critical levels we may use enumeration over all relevant values, if the other parameter values are given. The set of relevant values for the critical levels depends on the inventory policy which is used. For example, for the $(\boldsymbol{c}, s, Q)$ policy described above, it is not worthwhile to consider critical levels exceeding $s + Q$. The introduction of a critical level policy may complicate the search for optimal values of the other parameters. For positive critical levels the demand process depends on the inventory level which may affect e.g. convexity properties. Thus, one should be cautious when applying standard optimisation procedures for the other parameters. Nevertheless, using the optimal values of the parameters for the situation where no critical level policy is applied, and determining for these parameter values the best critical levels, it is always possible to find a critical level policy which is at least as good as the old policy. Observe that a system without applying a critical level policy is equivalent to a system with a critical level policy where $c_j = -\infty$ for all demand classes j. In a lost sales environment it suffices to set all critical levels at zero.

The general approach to find an optimal critical level policy would be to develop new (approximate) expressions for the average inventory cost and service levels, and derive bounds and/or convexity results for the policy parameters. So far in the literature on inventory models with multiple demand classes, the focus has been on the derivation of (approximate) expressions for the inventory cost and service levels. The problem of optimising the policy parameters has been addressed only by few authors. In the next section an overview of the existing literature on multiple demand classes and the critical level policy is presented.

5 Literature review

Veinott (1965) was the first to consider the problem of several demand classes in inventory systems. He analysed a periodic review inventory model with n demand classes and zero lead time, and introduced the concept of a critical level policy. Topkis (1968) proved the optimality of this policy both for the case of backordering and for the case of lost sales. He made the analysis easier by breaking down the period until the next ordering opportunity into a finite number of subintervals. In any given interval the optimal rationing policy is such that demand from a given class is satisfied from existing stock as long as there remains no unsatisfied demand from a higher class and the stock level does not drop below a certain critical level for that class. The critical levels are generally decreasing with the remaining time until the next ordering opportunity. Independent of Topkis (1968), Evans (1968) and Kaplan (1969) derived essentially the same results, but for two demand classes. A single period inventory model where demand occurs at the end of a period is presented by Nahmias & Demmy (1981) for two demand classes. This work was later generalised to multiple demand classes by Moon & Kang (1998). Nahmias & Demmy generalised their results to a multi-period model, with zero lead times and an (s, S) inventory policy, with policy parameters satisfying $0 < c < s < S$. Atkins & Katircioğlu (1995) analysed a periodic review inventory system with several demand classes, backordering and a fixed lead time, where for each class a minimum service level was required. For this model they presented a heuristic rationing policy. Cohen, Kleindorfer & Lee (1988) also considered the problem of two demand classes, in the setting of a periodic review (s, S) policy with lost sales. However, they did not use a critical level policy. At the end of every period the inventory is issued with priority such that stock is used to satisfy high-priority demand first, followed by low-priority demand.

The first contribution considering multiple demand classes in a continuous review inventory model was made by Nahmias & Demmy (1981). They analysed an (s, Q) inventory model with two demand classes, Poisson demand, backordering, a fixed lead time and a critical level policy, under the assumption that there is at most one outstanding order. This assumption implies that whenever a replenishment order is triggered, the net inventory and the inventory position are identical. Their main contribution was the derivation of approximate expressions for the fill rates. In their analysis they used the notion of the *hitting time* of the critical level, i.e. the time that the inventory level reaches the critical level. Conditioning on this hitting time, it is possible to derive approximate expressions for the cost and service levels. Dekker, Kleijn & De Rooij (1998) considered a lot-for-lot inventory model with the same characteristics, but without the assumption of at most one outstanding order. They discussed a case study on the inventory control of slow moving spare parts in a large petrochemical plant, where parts were installed in equipment of different criticality. Their main result was the derivation of (approximate) expressions for the fill rates for both demand classes. The results of Nahmias & Demmy (1981) were generalised by Moon & Kang (1998). They considered an (s, Q) model with compound Poisson demand, and derived (approximate) expressions for the fill rates of the two demand classes. The model of Nahmias & Demmy (1981) is analysed in a lost sales context by Melchiors, Dekker & Kleijn (1998).

Ha (1997a) discussed a lot-for-lot model with two demand classes, backordering and exponentially distributed lead times, and showed that this model can be formulated as a queueing model. He showed that in this setting a critical level policy is optimal, with the critical level decreasing in the number of backorders of the low-priority class. Moreover, he proved that it is optimal to increase the stock level when upon the arrival of a replenishment order there are backorders for low-priority customers and the inventory level is below the critical level.

A critical level policy for two demand classes where the critical level depends on the remaining time until the next stock replenishment was discussed by Teunter & Klein Haneveld (1996). A so-called *remaining time* policy is characterised by a set of critical stocking times $L_1, L_2, \ldots$; if the remaining time until the next replenishment is at most L_1 no items are reserved for high-priority customers, if the remaining time is between L_1 and $L_1 + L_2$ then one item should be reserved, and so on. They first analyse a model which is the continuous equivalent of the periodic review models by Evans (1968) and Kaplan (1969). Teunter & Klein Haneveld also presented a continuous review (s, Q) model with nonnegative deterministic lead times. Under the assumption that an arriving replenishment order is large enough to satisfy all outstanding backorders for high-priority customers, they derived a method to find (near) optimal critical stocking times. They showed that such a remaining time policy outperforms a simple critical level policy where all critical levels are stationary.

Ha (1997b) considered a single-item, make-to-stock production system with n demand classes, lost sales, Poisson demand and exponential production times. He modelled the system as an $M/M/1/S$ queueing system and proved that a lot-for-lot production policy and a critical level rationing policy is optimal. Moreover, the optimal policy is stationary. For two demand classes he presented expressions for the expected inventory level and the stockout probabilities. To determine the optimal policy he used an exhaustive search, and he used the assumption that the average cost is unimodal in the order-up-to level. Dekker, Hill & Kleijn (1997) analysed a similar system, with n demand classes, lost sales, Poisson demand and general distributed lead times. They modelled this system as an $M/M/S/S$ queueing system to derive expressions for the average cost and service levels. It was shown by Nguyen (1991) that for exponential distributed lead times, a critical level policy is optimal in such a queueing system. Dekker, Hill & Kleijn (1997) have derived efficient algorithms to determine the optimal critical level, order-up-to level policy, both for systems with and without service level restrictions. Moreover, they presented a fast heuristic approach for the model without service level restrictions. In this model, the different demand classes are characterised by different unit lost sales costs.

The only contribution assuming deterministic demand was recently made by Moon & Kang (1998). They considered a single period model with n demand classes, lost sales and continuous review, and introduced the notion of *rationing trigger times*. Instead of having critical stock levels, there are critical times after which demand from certain classes will no longer be satisfied.

To conclude this section, we categorised the literature based on the following characteristics: periodic or continuous review, 2 or n demand classes. This categorisation is presented in Table 5.1. It is interesting to note that all contributions considering n demand classes and continuous review assume lost sales.

Table 5.1: Categorisation of literature on multiple demand classes.

periodic review	
2 classes	Evans (1968)
	Kaplan (1969)
	Nahmias & Demmy (1981)
	Cohen, Kleindorfer & Lee (1988)
n classes	Veinott (1965)
	Topkis (1968)
	Moon & Kang (1998)
	Atkins & Katircioğlu (1995)
continuous review	
2 classes	Nahmias & Demmy (1981)
	Dekker, Kleijn & De Rooij (1998)
	Ha (1997a)
	Teunter & Klein Haneveld (1996)
	Melchiors, Dekker & Kleijn (1998)
n classes	Ha (1997b)
	Moon & Kang (1998) (*deterministic demand*)
	Dekker, Hill & Kleijn (1997)

6 Summary and conclusions

In this chapter we discussed inventory systems where multiple classes of demand may be distinguished. A number of practical examples where multiple demand classes naturally arise were presented. We also introduced the critical level policy which reserves part of the stock for high-priority customers. In Section 4 the problem of determining an optimal critical level policy was discussed. Finally, we presented an overview of the existing literature on this subject.

In the recent literature, some authors have addressed the problem of multiple demand classes, but the main focus has been on the determination of the inventory cost and service level, given a critical level policy. Due to its analytical complexity, the optimisation of the critical level policy has not been given a lot of attention. In order to facilitate the implementation of this policy in practice, we must focus more attention on deriving easy and fast methods to determine (near) optimal policy parameters.

References

Atkins, D./Katircioğlu, K.K. (1995): Managing inventory for multiple customers requiring different levels of service. Working Paper 94-MSC-015, University of British Columbia, Vancouver, B.C.

Ballou, R.H. (1992): Business Logistics Management, 3rd edition. (Prentice Hall) Englewood Cliffs, NJ.

Chikán, A. (1990): Inventory Models. (Akadémiai Kiadó) Budapest.

Cohen, M.A./Kleindorfer, P.R./Lee, H.L. (1988): Service constrained (s, S) inventory systems with priority demand classes and lost sales, in: Management Science, 34:482–499.

Dekker, R./Hill, R.M./Kleijn, M.J. (1997): On the $(S-1, S)$ lost sales inventory model with priority demand classes. Report EI9743/A, Erasmus University Rotterdam, The Netherlands.

Dekker, R./Kleijn, M.J./de Rooij, P.J. (1998): A spare parts stocking system based on equipment criticality, to appear in: International Journal of Production Economics.

Eppen, G./Schrage, L. (1981): Centralized ordering policies in a multi–warehouse system with lead times and random demand, in: Multilevel production/inventory control systems: Theory and practice, L.B. Schwarz (editor), North–Holland, Amsterdam.

Evans, R.V. (1968): Sales and restocking policies in a single item inventory system, in: Management Science, 14:463–472.

Greenberg, W. (1991): Competition, Regulation and Rationing in Health Care. (Health Administration Press) Ann Arbor, Michigan.

Ha, A.Y. (1997a): Stock-rationing policy for a make-to-stock production system with two priority classes and backordering, in: Naval Research Logistics, 44:457–472.

Ha, A.Y. (1997b): Inventory rationing in a make-to-stock production system with several demand classes and lost sales, in: Management Science, 43:1093–1103.

Held, P.J. (1988): Access to kidney transplantation, in: Archives of Internal Medicine, 148:2594–2600.

Henaux, C./Semal, P. (1998): Delivery service: Expectiation, performances and costs for a distributor, in: Advances in Distribution Logistics, B. Fleischmann, Jo A.E.E. van Nunen, M. Grazia Speranza, and P. Stähly (eds.), Springer-Verlag, Berlin.

Kaplan, A. (1969): Stock rationing, in: Management Science, 15:260–267.

Kleijn, M.J./Dekker, R. (1998): Using break quantities for tactical optimisation in multi-stage distribution systems, in: Advances in Distribution Logistics, B. Fleischmann, Jo A.E.E. van Nunen, M. Grazia Speranza, and P. Stähly (eds.), Springer-Verlag, Berlin.

Lee, H.L./Nahmias, S. (1993): Single-product, single-location models, in: Handbooks in OR & MS, vol. 4, S.C. Graves, A.H.G. Rinnooy Kan, and P. Zipkin (eds), Elsevier Science Publishers, North-Holland.

Melchiors, P./Dekker, R./Kleijn, M.J. (1998): Inventory rationing in an (s, Q) inventory model with lost sales and two demand classes. Report EI9837/A, Econometric Institute, Erasmus University Rotterdam, The Netherlands.

Moon, I./Kang, S. (1998): Rationing policies for some inventory systems, in: Journal of the Operational Research Society, 49:509–518.

Nguyen, V. (1991): On the optimality of trunk reservation in overflow processes, in: Probability in the Engineering and Informational Sciences, 5:369–390.

Porteus, E.L. (1990): Stochastic inventory theory, in: Handbooks in OR & MS, vol. 2, D.P. Heyman and M.J. Sobel (eds), Elsevier Science Publishers, North-Holland.

Robinson, L.W. (1994): Optimal and approximate control policies for airline booking with sequential nonmonotonic fare classes, in: Operations Research, 43:252–263.

Silver, E.A./Pyke, D.F./Peterson, R. (1998): Inventory Management and Production Planning and Scheduling, 3rd edition. (John Wiley & Sons), New York.

Teunter, R.H./Klein Haneveld, W.K. (1996): Reserving spare parts for critical demand. Research Report, Graduate School/Research Institute System, Organisations and Management (SOM), University of Groningen, The Netherlands.

Topkis, D.M. (1968): Optimal ordering and rationing policies in a nonstationary dynamic inventory model with n demand classes, in: Management Science, 15:160–176.

Veinott, A.F. Jr. (1965): Optimal policy in a dynamic, single product, nonstationary inventory model with several demand classes, in: Operations Research, 13:761–778.

Veinott, A.F. Jr. (1966): The status of mathematical inventory theory, in: Management Science, 12:745–777.