THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

On Service Level Measures in Stochastic Inventory Control

Citation for published version:

Rossi, R 2013, 'On Service Level Measures in Stochastic Inventory Control', Paper presented at IFAC Conference on Manufacturing Modeling, Management and Control, MIM 2013, Saint Petersburg, Russian Federation, 19/06/13 - 21/06/13 pp. 1991-1996. <https://doi.org/10.3182/20130619-3-RU-3018.00295>

Digital Object Identifier (DOI):

[10.3182/20130619-3-RU-3018.00295](https://doi.org/10.3182/20130619-3-RU-3018.00295)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Early version, also known as pre-print

Publisher Rights Statement:

© Rossi, R. (2013). On Service Level Measures in Stochastic Inventory Control. 1991-1996. Paper presented at IFAC Conference on Manufacturing Modeling, Management and Control, MIM 2013, Saint Petersburg, Russian Federation. [10.3182/20130619-3-RU-3018.00295](https://doi.org/10.3182/20130619-3-RU-3018.00295)

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

On Service Level Measures in Stochastic Inventory Control

Roberto Rossi*

* *University of Edinburgh Business School, Edinburgh, EH8 9JS UK
(e-mail: roberto.rossi@ed.ac.uk).*

Abstract: We consider the issue of modeling service level measures in stochastic decision making via chance constraints. More specifically we focus on service level measures in production/inventory control under stochastic demand and α service level constraints, which are constraints enforcing a prescribed non-stockout probability for the system. We introduce multiple ways of expressing these chance constraints by using conditional probability. Then we demonstrate that, when these constraints are formulated by using expressions that do not involve a conditional probability, a base stock policy is optimal for this problem only under a number of assumptions. To demonstrate this, we discuss a number of examples for simple cases in which it is possible to find better policies and we also present some analytical results. In contrast, when our novel measure involving a conditional probability is used, a base stock policy is optimal under much less restrictive assumptions, although the cost performance of the system tends to deteriorate.

Keywords: inventory control; stochastic modeling; service level measures; non-stationary demand; control policies

1. INTRODUCTION

We consider problems of decision making under uncertainty. These problems typically comprise a set of decision and observation stages. At each stage, one or more decisions must be made in order to satisfy a number of deterministic or stochastic constraints, which are constraints involving decision and random variables, and in certain cases to maximize or minimize a given objective function. In the following discussion, we assume that randomized decisions at a given stage — i.e. randomized policies — are forbidden. Each stage, also comprises one or more random variables, for which a probability distribution is given. At a given stage, after a decision as been made, random variables are observed. The observed realizations may affect subsequent decisions and observations.

Several approaches to decision making under uncertainty exist in the literature. To gain insights into the issues we are going to discuss next, we choose to adopt a graphical depiction of the problem known as “decision trees”. Decision trees are commonly used in decision making under uncertainty and provide a convenient representation of the problem at hand. In a decision tree, we have decision nodes and chance forks. Branches originating from decision nodes ($\square$) represent decisions, each decision is usually indicated in text above the respective branch. Branches originating from a chance fork ($\circ$) represent realizations for the random variables. Associated with a branch originating from a chance fork there are two numbers, the value taken by the random variable associated with the chance fork — in our case the demand in a given period — and the probability of this event. Note that, in general, this probability is the conditional probability of the event given the realizations associated with parent chance forks in the tree. Consider

Fig. 1. The key elements of a decision tree

for example the decision tree in Fig. 1. It is possible to observe decisions (D_i) at different stages ($i = 1, 2$). It is also possible to observe that the probability of events at stage 2 are conditional probabilities that depend on random variable (r_i) realizations in the previous stages. The tree comprises a total of 4 scenarios. A scenario is a possible realization for all the random variables in the problem.

In this work we are concerned with the issue of “service level” measures in decision making under uncertainty. A common service level measure in stochastic decision making takes the form of a chance constraint. A chance constraint is a particular type of stochastic constraint that must be satisfied according to a prescribed probability. For instance, consider a stochastic constraint $D_1 - r_1 + D_2 - r_2 = 0$. Assume that r_1 and r_2 are independent and that the two values in their support are equally likely to occur. Under a chance constraint $\Pr\{D_1 - r_1 + D_2 - r_2 \geq 0\} \geq 0.25$ the assignment $D_1 = 2$, $(D_2|r_1 = 3, D_1) = 2$ and $(D_2|r_1 = 2, D_1) = 0$ is feasible; in fact $\Pr\{D_1 - r_1 + D_2 - r_2 \geq 0\}$ can be computed as $\Pr\{r_1 = 3\} \Pr\{r_2 = 1|r_1 = 3\} = 0.5 \cdot 0.5$ since r_1 and r_2 are independent. In this simple example, we have conditioned the probability to an event. However, in general, it is possible to condition on a random variable. That is, let E be an event, we may write $\Pr\{E|r_1\}$; this is a function which takes value $\Pr\{E|r_1 = i\}$ when $r_1 = i$. Enforcing a chance constraint such as $\Pr\{E|r_1\} \geq \alpha$ means making sure that this function does not take a value less than α for each value in the support of r_1 . In practical situations chance constraints such as $\Pr\{D_1 - r_1 + D_2 - r_2 \geq 0|r_1\} \geq \alpha$ may become relevant. Under this chance constraint, the above assignment would be infeasible, since when $r_1 = 3$, $\Pr\{D_1 - r_1 + D_2 - r_2 \geq 0|r_1 = 3\} = \Pr\{r_2 = 1|r_1 = 3\} = 0.5$ however, when $r_1 = 2$, $\Pr\{D_1 - r_1 + D_2 - r_2 \geq 0|r_1 = 2\} = 0$. In what follows we will discuss why these service level measures become particularly relevant in stochastic inventory control.

2. STOCHASTIC PRODUCTION/INVENTORY PROBLEM UNDER SERVICE LEVEL CONSTRAINTS

We consider the periodic review production/inventory problem under stochastic demand and service level constraints. In a periodic review system inventory is reviewed only at discrete points in time. We review inventory only at the beginning and at the end of a period. Orders can be placed only at the beginning of a period. Demand is a random variable d with known distribution. The delivery lead-time is constant and equal to L periods. Unmet demand is backordered and fulfilled as soon as a replenishment arrives. We consider a service level constraint enforcing a specified probability α of no-stockout per period — α service level. A holding cost of $\$h$ per period is paid for each unit carried in stock to the next period.

Example 1 We consider a planning horizon comprising 4 periods. In each period we observe a random demand that follows a discrete distribution. The probability mass functions for the demand in each period are the following.

$$\begin{aligned} \text{pmf}(d_1) &= \{18(0.5), 26(0.5)\} \\ \text{pmf}(d_2) &= \{52(0.5), 6(0.5)\} \\ \text{pmf}(d_3) &= \{9(0.5), 43(0.5)\} \\ \text{pmf}(d_4) &= \{20(0.5), 11(0.5)\}. \end{aligned}$$

Accordingly, in period 1 we may observe 2 values for the random demand, 18 and 26, each of which occurs with probability 0.5. The complete set of scenarios is presented in Table 1. The delivery lead time is set to 0; therefore orders placed at the beginning of a period are received immediately. The holding cost h is set to \$10 and it is charged on items in stock at the end of a period, after demand has been observed. The prescribed no stockout

Scenario	d_1	d_2	d_3	d_4	Probability
1	18	52	9	20	0.0625
2	18	52	9	11	0.0625
3	18	52	43	20	0.0625
4	18	52	43	11	0.0625
5	18	6	9	20	0.0625
6	18	6	9	11	0.0625
7	18	6	43	20	0.0625
8	18	6	43	11	0.0625
9	26	52	9	20	0.0625
10	26	52	9	11	0.0625
11	26	52	43	20	0.0625
12	26	52	43	11	0.0625
13	26	6	9	20	0.0625
14	26	6	9	11	0.0625
15	26	6	43	20	0.0625
16	26	6	43	11	0.0625

Table 1. Scenarios for Example 1

probability α is 0.85. The optimal solution to this problem, obtained by a trivial scenario based MILP model, can be represented by means of a decision tree. Let I_t denote the inventory level — i.e. on hand stock minus backorders — at the end of period t . If the event of interest is a no stockout in period t , i.e. $I_t \geq 0 \Leftrightarrow I_{t-1} + Q_t - d_t \geq 0$, then we may enforce either the constraint $\Pr\{I_{t-1} + Q_t \geq d_t\} \geq \alpha$ or the constraint $\Pr\{I_{t-1} + Q_t \geq d_t | I_{t-1}\} \geq \alpha$; see Bitran and Yanasse (1984). We present both the optimal solution (Fig. 2) under the service level measure

$$\Pr\{I_t \geq 0\} \geq \alpha \quad (1)$$

and the optimal solution (Fig. 3) under the service level

Fig. 2. The optimal policy for Example 1 under the service level measure in Eq. 1

measure

$$\Pr\{I_t \geq 0 | I_{t-1}\} \geq \alpha \quad (2)$$

Fig. 3. The optimal policy for Example 1 under the service level measure in Eq. 2

It should be noted that I_{t-1} is a random variable functionally dependent on our past ordering decisions and on past demand realizations. Therefore we can write $I_{t-1} = f(d_{t-1}, \dots, d_1, Q_{t-1}, \dots, Q_1)$. Since I_{t-1} is given, this means that all demand realizations and decisions taken in periods $1, \dots, t-1$ are given. Intuitively, Eq. 2 simply requires that, for each possible subtree at stage t in the decision tree, a sufficient quantity Q_t is ordered so that the conditional probability of meeting demand over the next period given all demand realizations and decisions taken in periods $1, \dots, t-1$ is greater than α . Since we only consider holding cost, this Q should be the smallest possible one. Adopting such a measure therefore greatly simplifies the analysis of the problem, in fact the optimal policy takes a well-known form called “base stock” policy. To describe this policy, we must introduce the notion of inventory position. The inventory position comprises items in stocks minus backorders plus incoming orders not yet received. In a base stock policy, an order is placed as soon as the inventory position drops below the “base stock level”. In general, this example shows that the optimal policy under the service measure in Eq. 1 is not a base stock policy. Instead, for the measure in Eq. 2, a simple forward analysis of the decision tree immediately leads to the optimal policy, which in this example takes the form of a base stock policy with base stock levels 26, 52, 43, 20 in period 1, 2, 3, and 4, respectively. The expected total cost is clearly higher under the measure in Eq. 2. However, in Section 4 we will discuss why such a measure may better reflect contractual requirements and industrial practices.

3. STATIONARY STOCHASTIC DEMAND

We consider the continuous review production/inventory problem under stationary stochastic demand and service level constraints. In a continuous review production/inventory problem inventory is monitored continuously and orders can be placed at each time instant. Demand — measured in units per period — is a random variable d with known distribution. The delivery lead-time is constant and equal to L periods. Unmet demand is back-ordered and fulfilled as soon as a replenishment arrives. We consider a service level constraint enforcing a specified probability α of no-stockout over the replenishment lead time — α service level. A holding cost of $\$h$ per period is paid for each unit carried in stock.

In the literature, the α service level is generally defined informally as the “no-stockout probability over the replenishment lead time” or “no-stockout probability per period”, for the continuous or periodic review case, respectively. Similarly to what discussed in the previous section, we now aim to formalize the service level definition for the continuous review case in mathematical terms.

We denote the inventory position at time t as I_t^p . Assuming that demand is a stationary stochastic process and that the current time is t , we may express the service level constraint by using chance constraint $\Pr\{I_{t+L} \geq 0\} \geq \alpha$, where I_t is a random variable representing the inventory level, i.e. items in stocks minus backorders, at time t . Let d_L denote the demand distribution over the lead-time. By exploiting the fact that the inventory position tracks incoming orders, we can rewrite our service level constraint

$$\Pr\{I_t^p - d_L + Q_t \geq 0\} \geq \alpha \quad (3)$$

Alternatively, as discussed in the previous section, one may adopt the following chance constraint to express the service level constraint

$$\Pr\{I_t^p - d_L + Q_t \geq 0 | I_t\} \geq \alpha \quad (4)$$

When the service level is formulated as in Eq. 4, since we condition the event $I_t^p - d_L + Q_t \geq 0$ on I_t , I_t^p becomes a scalar value and the optimal order quantity can be immediately obtained by simply inverting the cumulative distribution function of the demand over lead-time

$$Q_t = \min\{Q | I_t^p + Q \geq \text{cdf}_{d_L}^{-1}(\alpha)\}$$

where $\text{cdf}_{d_L}^{-1}(\alpha)$ denotes the inverse cumulative distribution of d_L . This shows that, under the cost structure discussed above and this service level measure, a base stock policy with base stock level

$$S = \min\{s | s \geq \text{cdf}_{d_L}^{-1}(\alpha)\} \quad (5)$$

becomes optimal regardless of the nature of the demand distribution, i.e. continuous or discrete, as long as randomized policies are forbidden. This optimal policy is directly related to the classical optimality proof which exploits the connection with a pure cost oriented formulation of the problem via the so-called critical fractile solution (see e.g. van Houtum and Zijm (2000)).

However, as we will see, if the service level is formulated as in Eq. 3, a base stock policy is not optimal, in general, when demand follows a discrete distribution; furthermore, if demand follows a continuous distribution, a base stock policy is only optimal if specific conditions are met.

3.1 Discrete demand distribution

When demand distribution is discrete, it may not be possible to find base stock level that guarantees exactly a service level α as defined in Eq. 3.

Example 2 We consider a Poisson demand with rate $\lambda = 3$ units/period. The lead time for an order is $L = 2$ periods. Holding cost is $h = \$4$ per unit per period. We enforce an α service level with $\alpha = 0.7$. It immediately follows that the demand over lead-time follows a Poisson distribution with rate λL and that the optimal base stock level is $S = 7$, since this is the minimum base stock level for which the cumulative distribution of a poisson with rate λL exceeds $\alpha = 0.7$ — the actual service level associated with this base stock level is 0.743. The expected total cost per period of the optimal base stock policy is \$ 9.14.

If the service level is defined as in Eq. 3, when demand is discrete and there exists no base stock level that satisfies Eq. 5, a common strategy is to pick the minimum base stock level that guarantees a service level larger than α . However, in such a case it is easy to show that a base stock policy is not optimal. In fact, it is possible to devise a better control policy by allowing the base stock level to depend on the demand realizations in previous periods. For instance, for the problem described in Example 2, we may adopt a policy which orders up to 7 if demand in the past L periods has been less or equal to 7, and that orders up to 6 if demand has been greater than 7. This policy ensures a no-stockout probability per replenishment cycle of 0.708, the expected total cost per period is 8.82. In fact, it is easy to see that also this policy is not optimal and that if our aim is to guarantee a given no-stockout probability over the replenishment lead time then there may be a better “past demand over lead time dependent policy” which guarantees a service level of exactly 0.7 at minimum cost. However, as stated below finding such a policy is NP-hard.

Theorem 1. Under a discrete demand distribution finding a policy that guarantees exactly a given α service level is NP-hard.

Proof. Reduction from 0-1 Knapsack (omitted).

3.2 Continuous demand distribution

If demand follows a continuous distribution, a base stock policy is optimal only if specific conditions are met.

Theorem 2. A base stock policy is optimal for the production/inventory problem with continuously distributed stationary stochastic demand under α service level constraints formulated as in Eq. 3 if

$$\alpha \frac{d \text{cdf}_{d_L}^{-1}(\alpha)}{d\alpha}$$

is increasing in α .

Proof. Consider the following function

$$b(S) = \int_0^S (S - i) \text{pdf}_{d_L}(i) di$$

which computes the expected buffer stock associated with a base stock level S . To prove that a base stock policy

is optimal we must show that a demand or stock level dependent policy cannot produce a better cost. This can be shown by considering the following function

$$h(\alpha) = \int_0^{\text{cdf}_{d_L}^{-1}(\alpha)} (\text{cdf}_{d_L}^{-1}(\alpha) - i) \text{pdf}_{d_L}(i) di$$

that represents the expected buffer stock as a function of the service level α ; and by showing that such a function is is convex in α .

We apply Leibniz integral rule to compute the derivative for the above function w.r.t. α and thus obtain

$$\begin{aligned} \frac{dh(\alpha)}{d\alpha} &= \int_0^{\text{cdf}_{d_L}^{-1}(\alpha)} \frac{\partial \alpha}{\partial} (\text{cdf}_{d_L}^{-1}(\alpha) - i) \text{pdf}_{d_L}(i) di \\ \frac{dh(\alpha)}{d\alpha} &= \int_0^{\text{cdf}_{d_L}^{-1}(\alpha)} \left(\frac{1}{\text{pdf}_{d_L}(\text{cdf}_{d_L}^{-1}(\alpha))} \right) \text{pdf}_{d_L}(i) di \\ \frac{dh(\alpha)}{d\alpha} &= \left(\frac{1}{\text{pdf}_{d_L}(\text{cdf}_{d_L}^{-1}(\alpha))} \right) \int_0^{\text{cdf}_{d_L}^{-1}(\alpha)} \text{pdf}_{d_L}(i) di \\ \frac{dh(\alpha)}{d\alpha} &= \left(\frac{1}{\text{pdf}_{d_L}(\text{cdf}_{d_L}^{-1}(\alpha))} \right) \text{cdf}_{d_L}(\text{cdf}_{d_L}^{-1}(\alpha)) \\ \frac{dh(\alpha)}{d\alpha} &= \frac{\alpha}{\text{pdf}_{d_L}(\text{cdf}_{d_L}^{-1}(\alpha))} \\ \frac{dh(\alpha)}{d\alpha} &= \alpha \frac{d \text{cdf}_{d_L}^{-1}(\alpha)}{d\alpha} \end{aligned}$$

Note that to prove that this derivative is increasing we may want to verify that the second derivative is positive.

If $h(\alpha)$ is convex, we proceed by reductio ad absurdum. Assume that there exist a demand or stock dependent policy in which I different demand or stock dependent order quantities Q_i guarantee a service level of exactly α . Furthermore, let S_i be the inventory position — comprising items in stocks minus backorders plus incoming orders not yet received — immediately after an order Q_i is placed. Let $\Pr\{S_i\}$ be the probability of placing an order that will produce an inventory position S_i . In other words, we are looking for I values for S_i such that

$$\sum_{i=1}^I \Pr\{S_i\} \text{cdf}_{d_L}(S_i) = \alpha$$

The cost of this policy is

$$c(S) = \frac{h}{L} \left[E[d_L]/2 + \sum_{i=1}^I \Pr\{S_i\} b(S_i) \right]$$

where $E[\cdot]$ denotes the expected value and $\text{pdf}_{d_L}(\cdot)$ denotes the probability density function of the demand over the lead-time. However, if $h(\alpha)$ is convex, by applying Jensen's inequality, any of such ordering policies providing a service level α must incur a higher holding cost than a base stock policy with base stock level S computed as in Eq. 5. The extension to the continuous case in which I may be infinite is immediate. $\square$

Theorem 3. If demand over leadtime follows an exponential distribution with parameter λ , $h(\alpha)$ is convex.

Proof.

$$\frac{d^2 h(\alpha)}{d^2 \alpha} = \frac{d}{d\alpha} \alpha \frac{d \text{cdf}_{d_L}^{-1}(\alpha)}{d\alpha} = \frac{1}{(1 - \alpha)\lambda} + \frac{\alpha}{(1 - \alpha)^2 \lambda}$$

is positive.

Example 3 We consider a continuous review system in which the demand follows an Exponential distribution with parameter $\lambda = 1/3$ and thus expected value $1/\lambda = 3$ units/period. The lead time for an order is $L = 1$ period. Holding cost is $h = \$4$ per unit per period. We enforce an α service level with $\alpha = 0.7$. It immediately follows that the optimal base stock level is $S = 3.61$, since this is the base stock level for which the cumulative distribution of the demand over lead time is exactly equal to $\alpha = 0.7$. The expected total cost per period of the optimal base stock policy is \$ 7.01 per period.

We now provide an example in which $h(\alpha)$ is not convex and it is possible to find a policy that beats a base stock policy for a given service level.

Example 4 We consider a continuous review system in which demand follows a Beta distribution with parameter $\alpha = \beta = 0.2$ and thus expected value $\alpha/(\alpha + \beta) = 0.5$ units/period. The lead time for an order is $L = 1$ periods. Holding cost is $h = \$4$ per unit per period. We enforce an α service level with $\alpha = 0.7$. It immediately follows that the optimal base stock level is $S = 0.94$, since this is the base stock level for which the cumulative distribution of the demand over lead time is exactly equal to $\alpha = 0.7$. The expected total cost per period of the optimal base stock policy is \$ 4.70 per period. However, we now consider a policy that orders up to $S_1 = 0.9999$ if demand over lead time has been lower than 0.5 and that orders up to $S_2 = 0.4999$ if demand over lead time has been greater than 0.5. This policy ensures a service level of 0.7082, that is slightly higher than the prescribed one. However, the expected total cost per period of this policy is \$ 4.59.

Finally, it is worth mentioning that under a demand that follows a continuous distribution and that satisfies Theorem 2, the service level measures in Eq. 3 and Eq. 4 provide the same cost performance.

4. DISCUSSION

We introduced two possible strategies to capture α service level constraints in stochastic production/inventory control. The associated expressions, were presented in Eq. 3 and Eq. 4. Note that these expressions are easily extended to the case in which demand is non stationary. In this section we aim to discuss when one or the other measure is appropriate in practical settings.

First, we consider the case in which demand is stationary. If demand is truly stationary, as it may be the case for low-cost consumables (so called type “C” items), and it has been accurately estimated over a long time span, then we argue that it may make sense for management to adopt the service level in Eq. 3. In other words, management may want to “exploit” the long run properties of the stationary stochastic process to keep buffer stock as low as possible. However, in several practical cases a stationary demand process is only used for modeling convenience and stationarity does not reflect the actual nature of the demand. This is typically the case for inventory systems controlled under a “rolling horizon” strategy and for which forecast updates are periodically released. Most commercial inventory control packages cannot capture the complexity associated with modeling a non stationary demand

pattern. Therefore, in these packages, demand is typically approximated as a stationary process. An optimal control plan is then derived under this assumption, but only the most contingent ordering decision is implemented; then, periodically, the demand distribution is modified to reflect forecast updates received from the forecasting unit. As soon as forecast updates are received, a new stationary demand process is derived, which reflects the updated characteristics of the demand, and the stationary inventory control model is re-optimized in order to obtain a new plan. After obtaining a new plan, once more, only the most contingent ordering decision of this plan is implemented. This approach is known as rolling horizon approach with forecast updates. If this strategy is in place, we argue that adopting the service level measure in Eq. 3 makes little sense. In fact, this service measure tries to exploit long run “frequentist” characteristics of the demand to reduce buffer stock. Since we are constantly updating our belief on the demand process in a Bayesian sense, we should not rely on past belief, i.e. the distribution of the demand in past periods, to determine the service level associated with a given decision. The service level measure in Eq. 4 is thus more appropriate, since this measure decouples the computation of the service level in a given period from the demand distribution in previous periods. Similar considerations apply when demand process is non stationary. Also in this case the service level measure in Eq. 4 appears to be a more sensible choice. As an example, we may consider the production planning associated with the launch of a new product. A decision maker of course should never rely, for example, on the probability distribution associated with low demand scenarios during the first release week to lower the production quantity set at the beginning of the second week, after a very high demand over the first week has been observed. In fact, the event “demand over the first week” is a “one off” event which will never repeat, therefore the order quantity at the beginning of week two should not be affected by the probability distribution of this past event. The service level measure in Eq. 4 ensures this. However, there are situations in which adopting the service level measure in Eq. 3 may make sense under a non-stationary demand process subject to little or no forecast updates. Consider the case of a supermarket selling a consumable that presents a fairly well-known non-stationary weekly trend that repeats every week. For instance, this product may be delivered every day and have a very high demand on Monday, and then a demand that progressively decrease and becomes zero on Sunday, since the shop is closed. In this case, the non-stationary demand distribution actually reflects “frequentist” characteristics of the demand process, which may be exploited to adjust the order quantity on Tuesday, Wednesday etc. so that it accounts for low demand scenarios on previous weekdays.

5. RELATED WORKS

Service level measures are gaining momentum in inventory literature. Recently, Louly et al. (2012, 2008) discussed component supply planning in assembly systems under service level constraints; Louly and Dolgui (2012) presented a service level oriented approach to the problem of MRP offsetting for assembly systems with random component procurement times.

Despite the fact that service level measures are covered in most of inventory control textbooks, literature on mathematical formulations of service level measures in inventory control is quite sparse and service measure definitions are often presented informally. For instance in Silver et al. (1998) p. 245, the authors define the cycle or α service level as “the fraction of cycles in which a stockout does not occur”, where a “stockout” denotes an occasion in which the on-hand stock drops to zero level. For an (s, Q) policy, at p. 258 the authors phrase this measure as

$$\Pr\{\text{demand over lead time} \geq \text{reorder point } s\} \quad (6)$$

and compute the optimal reorder point s by inverting the cumulative distribution of the demand over lead time, this hints to the fact that they implicitly adopt the service measure in Eq. 4. In fact, under the measure in Eq. 3, for some distributions (e.g. Poisson) this strategy may result suboptimal. Also in Axsäter (2006), pp. 94–97, the authors seem to implicitly adopt the service measure in Eq. 4.

Sethi and Cheng (1995) investigate the optimality of (s, S) policies for inventory systems with markovian demand. In Section 5.2 they discuss models under service level constraints. The interesting aspect here is to observe that although the measure adopted is similar to the one presented in Eq. 6, the accompanying text clarifies that the authors actually refer to the service measure in Eq. 4, in fact the text clarifies that “given the demand state in a certain period” the service level measure can be converted into a simple deterministic equivalent expression by inverting the conditional density function of the demand.

A substantial literature exists on stochastic inventory control under non stationary demand and service level constraints when the so-called “static-dynamic uncertainty” policy is adopted. In this policy, inventory is controlled over a finite horizon comprising N periods. Ordering decisions are fixed at the beginning of the planning horizon, while decision on the actual order quantity is postponed until the very last moment. Bookbinder and Tan (1988) propose a two-stage heuristic for this problem, the first stage fixes all the order points, then based on this order schedule, the associated order-up-to-levels are determined in a separate fashion. Tarim and Kingsman (2004) propose a deterministic equivalent MIP model that can determine all these decisions at once thus producing an optimal plan. It is interesting to note that the α service level constraint is formulated in both these works as $\Pr\{\text{end of period inventory level} \geq 0\} \geq \alpha$, which corresponds to the measure in Eq. 3. However, the modeling strategy adopted in both works reveals that the actual measure adopted is the one in Eq. 4. A solution method to address the same problem under the measure in Eq. 3 is discussed in Rossi et al. (2008). An extension of this problem to the case in which a stochastic supplier lead time is considered is discussed in Rossi et al. (2010).

A discussion on the relation between cost and service level models in inventory system is presented in van Houtum and Zijm (2000). The authors point out that the majority of inventory model adopt penalty cost formulations instead of service level measures for dealing with stockout events; and that this is generally justified by a general belief that there exist one-to-one relationships between cost models and service level models, that is between the choice of

penalty cost on one hand and that of certain service level on the other. However, precise formulations and proofs for these one to one relationships are known only for few simple inventory systems. The authors then present a number of such relationships for a number of inventory models. However their discussion only holds under a demand whose cumulative distribution function is strictly increasing and for which it is possible to find an order up to levels that guarantees exactly a prescribed service level. Our discussion shows that the production/inventory control problem, both in the continuous and periodic review cases, becomes considerably more difficult when demand follows a discrete distribution. Furthermore, theorem 2 shows that there are also special classes of continuous distributions for which a base stock policy is not optimal under an α service level.

REFERENCES

- Axsäter, S. (2006). *Inventory Control*. International Series in Operations Research & Management Science. Springer.
- Bitran, G.R. and Yanasse, H.H. (1984). Deterministic approximations to stochastic production problems. *Operations Research*, 32(5), 999–1018.
- Bookbinder, J.H. and Tan, J.Y. (1988). Strategies for the probabilistic lot-sizing problem with service-level constraints. *Management Science*, 34(9), pp. 1096–1108.
- Louly, M.A., Dolgui, A., and Ahmari, A.A. (2012). Optimal MRP offsetting for assembly systems with stochastic lead times: POQ policy and service level constraint. *J. Intell. Manuf.*, 23(6), 2485–2495.
- Louly, M.A. and Dolgui, A. (2012). A note on analytic calculation of planned lead times for assembly systems under POQ policy and service level constraint. *International Journal of Production Economics*, 140(2), 778–781.
- Louly, M.A., Dolgui, A., and Hnaien, F. (2008). Supply planning for single-level assembly system with stochastic component delivery times and service-level constraint. *International Journal of Production Economics*, 115(1), 236–247.
- Rossi, R., Tarim, S.A., Hnich, B., and Prestwich, S. (2008). A global chance-constraint for stochastic inventory systems under service level constraints. *Constraints*, 13(4), 490–517.
- Rossi, R., Tarim, S.A., Hnich, B., and Prestwich, S. (2010). Computing the non-stationary replenishment cycle inventory policy under stochastic supplier lead-times. *International Journal of Production Economics*, 127(1), 180–189.
- Sethi, S.P. and Cheng, F. (1995). Optimality of (s, S) policies in inventory models with markovian demand. *Operations Research*, 45, 331–339.
- Silver, E., Pyke, D., and Peterson, R. (1998). *Inventory Management and Production Planning and Scheduling*. John Wiley & Sons.
- Tarim, S. and Kingsman, B.G. (2004). The stochastic dynamic production/inventory lot-sizing problem with service-level constraints. *International Journal of Production Economics*, 88(1), 105 – 119.
- van Houtum, G. and Zijm, W. (2000). On the relationship between cost and service models for general inventory systems. *Statistica Neerlandica*, 54(2), 127–147.