

Log-concavity of compound distributions with applications in operational and actuarial models

F. G. Badía * A. Federgruen † and C. Sangüesa ‡

*Department of Statistical Methods and IUMA, University of Zaragoza, Zaragoza, 50018, SPAIN

† Graduate School of Business, Columbia University, NY, 10027, UNITED STATES

‡ Department of Statistical Methods and IUMA, University of Zaragoza, Zaragoza, 50009, SPAIN

July 16, 2019

Abstract

We establish that a random sum of independent and identically distributed (i.i.d.) random quantities has a log-concave cumulative distribution function (cdf) if (i) the random number of terms in the sum has a log-concave probability mass function (pmf) and (ii) the distribution of the i.i.d. terms has a non-increasing density function (when continuous) or a non-increasing pmf (when discrete). We illustrate the usefulness of this result using a standard actuarial risk model and a replacement model.

We apply this fundamental result to establish that a compound renewal process observed during a random time interval has a log-concave cdf if the observation time interval and the inter-renewal time distribution have log-concave densities, while the compounding distribution has a decreasing density or pmf. We use this second result to establish the optimality of a so-called (s,S) policy for various inventory models with a stock-out cost coefficient of dimension [\$/unit], significantly generalizing the conditions for the demand and leadtime processes, in conjunction with the cost structure in these models. We also identify implications of our results for various algorithmic approaches to compute optimal policy parameters.

2000 Mathematics subject classification: 62E10; 60E15

Key words and phrases: log-concavity, compound distribution, inventory model, replacement model, actuarial model

*E-mail: gbadia@unizar.es

†e-mail: af7@gsb.columbia.edu

‡e-mail: csangues@unizar.es

1 Introduction

In many stochastic operational and financial systems, the key determinant for optimal control strategies is given by the distribution of the aggregate number of events that occur over a given (random) time interval. We refer to this random variable as the Aggregate Event quantity AE. In many of these systems, the events are generated by a compound renewal process, or special cases thereof. In several contexts, specific shape properties of the Aggregate Event quantity play a critical role to establish optimality results or to guarantee the validity of various algorithmic approaches to compute optimal control policy parameters. Of particular interest among these shape properties is the log-concavity of the cumulative distribution function, or the stronger property of log-concavity of the density function or probability mass function.

Most commonly used families of distributions have log-concave density functions or pmfs. For continuous distributions this applies to all Normals, Uniforms, Laplace, Logistic and Extreme Value distributions, as well as all Gamma (r, λ) distributions with shape parameter $r \geq 1$, Beta (a, b) distributions, with $a, b \geq 1$, Weibull(a) distributions with shape parameter $a \geq 1$, among various others. Similarly, all discrete uniform, Poisson, binomial, hypergeometric and negative binomials with shape parameters $r \geq 1$ have a log-concave pmf.

As our main result, we establish that a random sum of i.i.d. random quantities has a log-concave cdf if (i) the random number of terms in the sum has a log-concave pmf and (ii) the distribution of the i.i.d. terms has a non-increasing density function (when continuous) or a non-increasing pmf (when discrete). To date, the log-concave shape of the cdf of a random sum of i.i.d. terms was established only when the distribution of the number of terms in the sum is Binomial or Poisson (see Cai and Wilmot [11]), or Geometric (see Badía and Sangüesa [7]), rather than a general distribution with a log-concave pmf.

We illustrate the usefulness of this first result in three settings: in a classical risk model, we show that the 'non-ruin probability' of an insurance pool is a log-concave function of its initial endowment; the result is very useful when allocating an insurance firm's aggregate cash position across its various insurance pools. Applying the duality between risk models and queueing models, the previous result also shows the log-concavity of the cdf of the virtual waiting time in the $M/G/1$ queueing model. A third application deals with reliability theory, more specifically, replacement models for systems with additive damage.

We apply our fundamental result to establish that an Aggregate Event Distribution generated by a compound renewal process has a log-concave cdf if the observation time interval and the inter-renewal time distribution have log-concave densities, while the compounding distribution has a decreasing density or pmf. This result was only known for compound Poisson processes and an observation time that is the convolution of exponentials, see Badía and Sangüesa [7, Proposition 5.2.]. To obtain the stronger result of the pdf or pmf of the Aggregate Event distribution being log-concave itself, it again suffices that the random number of terms has log-concave density, but far more restrictive assumptions are required in terms of the distribution of each term.

To illustrate the usefulness of our second major result, i.e. the log-concavity of Aggregate Event Distributions, consider, for example, the standard single item inventory model in which demands are generated by a compound renewal process. Replenishment orders may be placed at any of the demand epochs and are received a lead time later, where the lead time is usually random, but independent of the demand process during this lead time. Under backlogging of stockouts, the state of the system may be represented by a single state variable, the so-called inventory position, defined as the aggregate of the inventory level and all outstanding orders. (The inventory level is the net difference between inventory on hand and the backlog volume.) Of critical importance in this system is the so-called lead time demand, defined as the total amount demanded over the

course of a leadtime interval, started either at a demand epoch or at an arbitrary steady-state epoch.

The cost structure in this type of inventory model, typically includes variable and fixed order costs, as well as inventory carrying costs that are proportional or convexly increasing in the inventory level. In almost all of the literature, the cost consequences of stockouts are represented by a backlogging penalty which is incurred for every backlogged unit and every unit of time it is backlogged. Under this structure it is well known that the so-called instant cost rate function of the inventory position is convex, irrespective of the shape of the leadtime demand distribution. This condition guarantees the optimality of an (s, S) policy under the long-run average cost criterion, as shown in various settings by Scarf [50], Veinott [63, 64], Zheng [71] and Chen [13].

However, in practice, the cost consequences of stockouts are often represented by an (additional)

(i) one-time penalty for any demand unit that is backlogged, irrespective of the duration of the fulfillment delay, and/or

(ii) a cost that is proportional to the time backlogs are on the books, regardless of the volume of backlogs that accumulates.

These two alternative backorder cost structure were recognized in early classical inventory textbooks, such as Hadley and Whitin [30] and Peterson and Silver [40]. Moreover, the two cost structures are logically connected to the most popular service level constraints, the so-called fill-rate and ready-rate constraints, which often are used, to design the company's procurement strategy, in particular, when the actual (full) cost consequences of stockouts are difficult to estimate. The two cost structures arise, respectively, when these service level constraints are dualized.

Their importance notwithstanding, the two cost structures received scant attention because it was unclear what type of policy is optimal under these structures. A seminal breakthrough arose with Rosling [43] who showed, in several models, differentiated by the type of demand process and the set of review epochs, that the inventory cost rate function continues to be quasi-convex, and an (s, S) policy continues to be optimal if the leadtime demand distribution has a log-concave cdf. However, it remained unclear under what general conditions with respect to the primitives of the model, this log-concavity property can be guaranteed, other than for deterministic leadtimes and (compound) Poisson demand processes with a concave cdf for the customers' demand size distribution.

Based on the above main result of the paper, we are able to establish the optimality of an (s, S) policy whenever the leadtime distribution as well as the interdemand time distribution have an arbitrary log-concave density, while the cdf of the demand size distribution is concave. (Deterministic leadtimes or interdemand times are included as limiting cases). In contrast, to date, optimality of an (s, S) policy under general parameters could only be shown when leadtimes are deterministic or convolutions of exponentials, and interdemand times are exponential as well, see Rosling [43] and Badía and Sangüesa [7]. In the absence of a backlog cost per period, i.e. cost structure (ii), the requirement of the cdf of the demand size distribution being concave can be relaxed to a property, even weaker than log-concavity.

When the demand process is an ordinary renewal process, Perera, Janakiraman and Shun-Chen [39] recently established that an (s, S) policy is optimal under a fully general order cost function and a general quasi-convex instantaneous inventory and shortage cost function. However, this optimality result was only obtained when orders arrive instantaneously, or when the leadtime is deterministic and the demand process Poisson. Based on our results, we extend the optimality result to a far broader combination of leadtime and interdemand time distributions: it suffices for both distributions to have a log-concave density.

Beyond establishing the optimality of (s, S) policies, under a much broader and practically relevant set of cost structures, or dualization of service level constraints, our results also have major

algorithmic implications. For example, Federgruen and Zheng [23] and Zheng and Federgruen [72] establish the most efficient algorithms, known to date, to compute the optimal policy parameters s^* and S^* . (The former applies to the special case where the demand process is a renewal process, so that (s, S) policies reduce to so-called (r, q) policies. Here, an order of fixed size q is placed whenever the inventory position drops to the level r). Their algorithms merely require that the inventory cost rate function be quasi-convex. Thus, the same algorithms can be used in the presence of the alternative shortage cost structures, under the very same conditions mentioned above to guarantee optimality of an (s, S) policy, to begin with.

As a further example, consider the standard inventory model with traditional backlogging and zero fixed order costs. Under the long-run average cost criterion a simpler base-stock policy is optimal and the optimal base-stock level is given by a given fractile of the leadtime demand distribution. Often, it is difficult to compute this leadtime distribution even if its primitive components are known (i.e. the distribution of the leadtime, that of the interdemand times and the distribution of an individual customer demand). In other settings, even these primitive distributions are unknown, but historical data are available for the (aggregate) leadtime demand volumes. A natural approximation for the desired critical fractile S^* of the leadtime demand distribution is to draw a sample from the historical data and compute the same or a closely related fractile within the sample. This is an application of the well-known Sample Average Approximation (SAA) method. In the context of the simple newsvendor problem, Levi, Pekaris and Uichanco [33] derived a tight optimality bound for this heuristic when the demand distribution in this single period model has a log-concave density which is decreasing to the right of S^* . Our results establish that the same optimality bound continues to apply in our general infinite horizon model under the above conditions guaranteeing a log-concave density for the leadtime demand distribution.

When the true leadtime distribution fails to be available, an alternative approach is to employ robust optimization. Here, the base-stock level is selected to minimize long-run average costs under the worst possible leadtime demand distribution within a given class. We show how knowledge of the log-concavity property of the cdf of the leadtime demand distribution can be exploited to obtain a far more cost effective solution.

Other settings in which an Aggregate Event distribution plays a central role include insurance and queueing models. In the former, this applies to the total claims presented in an insurance portfolio during a deterministic or random coverage period when generated by a compound renewal process, as in the classical Cramer-Lundberg risk model. Recent treatments, for example Albrecher, Cheung and Thonhauser [1, 2] consider the total claims paid out during a random observation period, see also the recent text book by Asmussen and Albrecher [6].

$GI^x/G/1$ denotes the general single-server queueing model with batch arrivals and i.i.d. inter-arrival times. It is usually analyzed via the process embedded on service completion epochs. The net addition of customers in between a consecutive pair of service completion epochs is an Aggregate Event distribution, driven by a compound arrival renewal process.

We complete this Section with a brief review of the relevant literature. There is an extensive literature on comparison results for random sums of random variables. See Shaked and Shanthikumar [53, Chapter 1] for a survey of these results. In contrast, there are only a few papers establishing shape properties for such random sums of random variables.

Shanthikumar [54] showed that a geometrically distributed number of i.i.d. random variables with a Decreasing Failure Rate (DFR) distribution, is DFR itself. Yu [70, p. 588] recently provided an alternative proof of this result when the summands have a discrete distribution. Ross, Shanthikumar and Zhu [45] established IFR and DFR properties for random sums of exponential distributions.

As mentioned, Cai and Willmot [11] showed that such random sums have a log-concave cdf when

the number of terms has a Binomial or Poisson distribution, while the summands have a concave cdf. Badía and Sangüesa [7] established the same when the number of terms is geometrically distributed. In this paper, we establish the log-concavity of the cdf of these random sums of random variables, when the number of terms has a general log-concave pmf. This general sufficient condition includes, of course, the special distributions (Binomial, Poisson, geometric) for which this result was known, thus far.

For a random sum of random variables to have the stronger property of a log-concave density, several results were obtained by Ninh and Prekopa [37].

Rosling [43], Huh et al. [31] and Perera et al. [39] identify important implications of leadtime demand distribution having a log-concave pdf or cdf as discussed above and, in more detail, in Section 4.

The remainder of this paper is organized as follows: In Section 2 we review several shape properties and derive some preliminary results. Section 3 derives our general conditions under which a random sum of i.i.d. terms has a log-concave cdf. Section 4 develops the new optimality and algorithmic results for (s, S) policies in continuous time inventory models (the periodic review one is left as an appendix). Section 5 provides conditions under which a random sum of i.i.d. terms has the stronger (but still very general) property of having a log-concave density of pmf.

2 Preliminary results

In this paper, we use the log-concavity property, repeatedly, and this in various contexts: log-concave sequences, probability densities (pdf), probability mass functions (pmf) and cumulative distribution functions (cdf)

Definition 1. *A sequence of non-negative real numbers $(a(n))_{n=0,1,\dots}$ is said to be log-concave if and only if*

(a) *The sequence $(a(n))_{n=0,1,\dots}$ does not have internal zeroes i.e., $a(n) > 0$ and $a(n+k) > 0$ for $k > 0 \Rightarrow a(n+i) > 0$, $1 \leq i \leq k$.*

(b) *The sequence $(a(n))_{n=0,1,\dots}$ satisfies*

$$\log(a(n+1)) \geq \frac{1}{2} \log(a(n)) + \log(a(n+2)) \quad \text{for } n = 0, 1, 2, \dots \quad (1)$$

Remark 1. *Condition (a) is a technical condition which needs to be imposed in order to obtain the preservation of discrete log-concavity under convolutions (see Ninh and Prekopa [37, p.3018]).*

A random variable has a log-concave cdf $F(\cdot)$ if $\log F(\cdot)$ is concave on its entire domain. If the random variable has a pdf f , it is said to be log-concave, if $\log f(\cdot)$ is concave. If the random variable is discrete with support on a subset of the set of integers \mathbb{Z} , its cdf cannot be log-concave on its entire domain (because this cdf has discontinuity points). In this case, we will say that it has a log-concave cdf $F(\cdot)$ if $F(i)$, $i \in \mathbb{Z}$ is a log-concave sequence, and a log-concave pmf $f(\cdot)$ if $f(i)$, $i \in \mathbb{Z}$ is a log-concave sequence.

As mentioned in the Introduction, most commonly used families of continuous distributions have a log-concave pdf, see there, for a long list of examples. See also An [3] and Table 1 in Bagnoli and Bergstrom [8] for a more extensive list of such families. Even so, log-concavity of the pdf [pmf] implies log-concavity of the cdf $F(\cdot)$ and its complement $\bar{F} = 1 - F$. The latter is therefore an even weaker property than the former. The log-concavity property is also preserved under convolutions. We summarize:

Lemma 1.

- (a) If a random variable has a log-concave pdf or pmf, both its cdf $F(\cdot)$ and ccdf $\bar{F}(\cdot)$ are log-concave as well.
- (b) The convolution of two random variables with log-concave pdfs [cdf, ccdf] has a log-concave pdf [cdf, ccdf]. The same closure property applies when the random variables are discrete, with respect to the pmf, cdf and ccdf of the convolution.

Proof. See Proposition 1-1 in Rosling [43]. □

The following lemma will be useful.

Lemma 2. Let M be a non-negative integer-valued random variable.

(a) The following assertions are equivalent

- (i) M has a log-concave cdf.
- (ii) $P(M = n)P(M \leq n) \geq P(M = n + 1)P(M \leq n - 1)$, $n = 1, 2, \dots$
- (iii) $g(n) := \frac{P(M = n)}{P(M \leq n - 1)}$ is a decreasing function (in the interval for which the denominator is strictly positive).

(b) If M has a log-concave cdf, then the Reversed Hazard Rate $P(M = n)/P(M \leq n)$ is decreasing.

Proof. (a) When the cdf of M is log-concave

$$P(M \leq n)^2 \geq P(M \leq n - 1)P(M \leq n + 1), \quad n = 1, 2, \dots \quad (2)$$

This can be written as

$$(P(M \leq n - 1) + P(M = n))P(M \leq n) \geq P(M \leq n - 1)(P(M \leq n) + P(M = n + 1)), \quad n = 1, 2, \dots$$

which is equivalent to

$$P(M = n)P(M \leq n) \geq P(M = n + 1)P(M \leq n - 1), \quad n = 1, 2, \dots \quad (3)$$

Dividing both sides of the inequality by $P(M \leq n)P(M \leq n - 1)$ we get, as a third equivalent condition

$$\frac{P(M = n)}{P(M \leq n - 1)} \geq \frac{P(M = n + 1)}{P(M \leq n)}. \quad (4)$$

In other words, properties (i)-(iii) are equivalent. Note that the condition for no internal zeroes is satisfied as $P(M \leq n)$ is increasing.

(b) By part (a), log-concavity of the cdf M is equivalent to $g(n) := \frac{P(M = n)}{P(M \leq n - 1)}$ being decreasing. Note that the Reversed Hazard rate:

$$\frac{P(M = n)}{P(M \leq n)} = \frac{P(M = n)/P(M \leq n - 1)}{P(M \leq n)/P(M \leq n - 1)} = \frac{g(n)}{1 + g(n)},$$

which is an increasing function of $g(n)$. Thus, if $g(n)$ is decreasing, so is the Reversed Hazard Rate. □

Part (b) of Lemma 2 defines the Reversed Hazard Rate for discrete random variables. For a continuous random variable with pdf $f(\cdot)$, the Reversed hazard rate is defined as $f(\cdot)/F(\cdot)$. It is immediate that in the continuous case, log-concavity of the cdf is equivalent to the Reversed Hazard Rate being decreasing.

Remark 2. *If a distribution function F is log-concave, then it has to be continuous except, perhaps, at the left end-point of its support (see Sengupta and Nanda [52, p.427]).*

It is easily verified that log-concavity implies unimodality. Log-concave densities are also known as Polya frequencies of order 2 (PF_2), see Ross [44] and Porteus [42].

Associated with the Reversed Hazard Rate property, is a stochastic order with the same name.

Definition 2. (See Shaked and Shanthikumar [53, p. 37]). *Let X and Y be two random variables. X is said to be smaller than Y in the reversed hazard rate order ($X \leq_{rh} Y$) if*

$$\frac{P(Y \leq x)}{P(X \leq x)} \quad \text{is increasing in } x$$

(on the interval for which the denominator is strictly positive).

3 Random sums of i.i.d. terms: log-concavity of the cumulative distribution function

In this Section we derive our main structural results. We consider a random sum of i.i.d. random terms as follows:

$$S_M := \sum_{i=1}^M Z_i, \tag{5}$$

where $(Z_i)_{i=1,2,\dots}$ is a sequence of i.i.d. terms and M is an integer valued random variable which is independent of this sequence. We first provide a general set of conditions under which S_M has a log-concave cdf. We then apply our result to a special case where M represents the number of events generated during a random time interval under a pure or equilibrium renewal process.

We start our development assuming the random variables $(Z_i)_{i=1,2,\dots}$ are strictly positive and integer-valued.

Theorem 1. *Let $(Z_i)_{i=1,2,\dots}$ denote a sequence of i.i.d. positive and integer valued terms with a concave cdf, i.e. a decreasing pmf $(q(n))_{n=1,2,\dots}$. Let M be a non-negative and integer-valued random variable, with a log-concave pmf, which is independent of $(Z_i)_{i=1,2,\dots}$. Then S_M has a log-concave cdf.*

We first need the following Lemma

Lemma 3. *Under the conditions of Theorem 1, let*

$$G_r(n, x) = \sum_{k=r}^n P(S_k \leq n) x^k, \quad x \geq 0, \quad n = r, r+1, \dots \quad (G_r(n, x) = 0, r > n) \tag{6}$$

$G_r(n, x)$ is a log-concave sequence in n for every $x > 0$ and $r = 0, 1, \dots$

Proof. Let $x > 0$ be fixed and call $H_r(n) := G_r(n, x)$. We first derive the following recursive formula

$$H_r(n) = P(S_r \leq n)x^r + x \sum_{i=1}^{n-r} H_r(n-i)q(i), \quad n = r, r+1, \dots \quad (7)$$

Note that

$$\begin{aligned} H_r(n) &= P(S_r \leq n)x^r + \sum_{k=r+1}^n P(S_k \leq n)x^k \\ &= P(S_r \leq n)x^r + \sum_{k=r+1}^n \sum_{i=1}^{n-k+1} q(i)P(S_{k-1} \leq n-i)x^k \\ &= P(S_r \leq n)x^r + x \sum_{i=1}^{n-r} q(i) \sum_{k=r+1}^{n-i+1} P(S_{k-1} \leq n-i)x^{k-1} \\ &= P(S_r \leq n)x^r + x \sum_{i=1}^{n-r} q(i) \sum_{j=r}^{n-i} P(S_j \leq n-i)x^j = P(S_r \leq n)x^r + x \sum_{i=1}^{n-r} q(i)H_r(n-i), \end{aligned}$$

which verifies Eq. (7). Next, we will show that:

$$H_r(n+1) = x^r P(S_r = n+1) + (1+xq(1))H_r(n) + x \sum_{j=1}^{n-r} H_r(n-j)(q(j+1) - q(j)). \quad (8)$$

To this end, we use Eq. (7) to obtain

$$\begin{aligned} H_r(n+1) - xH_r(n)q(1) &= x^r P(S_r \leq n+1) + x \sum_{i=2}^{n+1-r} H_r(n+1-i)q(i) \\ &= x^r P(S_r \leq n+1) + x \sum_{i=2}^{n+1-r} H_r(n-(i-1))q(i-1) + x \sum_{i=2}^{n+1-r} H_r(n-(i-1))(q(i) - q(i-1)). \end{aligned}$$

Change the summation index to $j = i - 1$ and apply Eq. (7) again, to obtain:

$$H_r(n+1) = xH_r(n)q(1) + x^r P(S_r = n+1) + H_r(n) + x \sum_{j=1}^{n-r} H_r(n-j)(q(j+1) - q(j)),$$

which verifies Eq. (8). In order to show that H_r is log-concave, we will show that (recall Eq. (1))

$$\frac{H_r(n+1)}{H_r(n)} \text{ is a decreasing function, for } n = r, r+1, \dots \quad (9)$$

We will proceed by induction with respect to n . First, we check the case $n = r$. We have, applying Eq. (8) that

$$\begin{aligned} \frac{H_r(r+1)}{H_r(r)} &= \frac{x^r P(S_r = r+1)}{H_r(r)} + 1 + xq(1); \\ \frac{H_r(r+2)}{H_r(r+1)} &\leq \frac{x^r P(S_r = r+2)}{H_r(r+1)} + 1 + xq(1). \end{aligned}$$

Note that in the last inequality we disregarded the last term in Eq. (8) (it is negative). Therefore,

$$\frac{H_r(r+1)}{H_r(r)} - \frac{H_r(r+2)}{H_r(r+1)} \geq \frac{x^r P(S_r = r+1)}{H_r(r)} - \frac{x^r P(S_r = r+2)}{H_r(r+1)} \quad (10)$$

Thus, we have to prove that

$$\frac{x^r P(S_r = r+1)}{H_r(r)} - \frac{x^r P(S_r = r+2)}{H_r(r+1)} \geq 0 \quad (11)$$

This is a special case ($n = r$) of the following fact (which will be needed next)

$$D(n) := H_r(n+1)P(S_r = n+1) - H_r(n)P(S_r = n+2) \geq 0, \quad n \geq r \quad (12)$$

To show Eq.(12), note first that $H_r(n+1) \geq \sum_{k=r}^n P(S_k \leq n+1)x^k$. Therefore,

$$D(n) \geq \sum_{k=r}^n (P(S_k \leq n+1)P(S_r = n+1) - P(S_k \leq n)P(S_r = n+2))x^k.$$

We will show that all the terms inside the summation are non-negative. In fact, the k -th term has the same sign as:

$$\frac{P(S_r = n+1)}{P(S_k \leq n)} - \frac{P(S_r = n+2)}{P(S_k \leq n+1)}$$

This term can alternatively be written as

$$\frac{P(S_r = n+1)}{P(S_r \leq n)} \frac{P(S_r \leq n)}{P(S_k \leq n)} - \frac{P(S_r = n+2)}{P(S_r \leq n+1)} \frac{P(S_r \leq n+1)}{P(S_k \leq n+1)}$$

But notice that

$$\frac{P(S_r = n+1)}{P(S_r \leq n)} \geq \frac{P(S_r = n+2)}{P(S_r \leq n+1)} \quad \text{and} \quad \frac{P(S_r \leq n)}{P(S_k \leq n)} \geq \frac{P(S_r \leq n+1)}{P(S_k \leq n+1)}$$

The first inequality is due to the fact that that S_r , $r = 1, 2, \dots$ has a log-concave cdf and Lemma 2 (a). In fact, $S_1 = Z_1$ has a log-concave cdf since $(q(n))_{n=1,2,\dots}$ is a decreasing sequence, so that $P(Z_1 \leq n)$, $n = 1, 2, \dots$ is a concave and, a fortiori, also log-concave sequence. Moreover, the cdf of S_r , $r = 2, 3, \dots$ is also log-concave because this property is preserved by convolutions (see Lemma 1). The second inequality is due to the fact that $S_r \leq_{rh} S_k$, $r < k$ (see Shaked and Shanthikumar [53, Theorem 1.C.12, p. 47]), and therefore (recall Definition 2)

$$\frac{P(S_k \leq n)}{P(S_r \leq n)} \leq \frac{P(S_k \leq n+1)}{P(S_r \leq n+1)}, \quad k \geq r.$$

Thus, Eq. (12) has been verified, and so has Eq. (9) for $n = r$. Assume the monotonicity property in Eq. (9) holds up to $n - 1$. We have, applying Eq. (8), that

$$\begin{aligned} \frac{H_r(n+1)}{H_r(n)} &= \frac{x^r P(S_r = n+1)}{H_r(n)} + 1 + xq(1) + x \sum_{j=1}^{n-r} \frac{H_r(n-j)}{H_r(n)} (q(j+1) - q(j)); \\ \frac{H_r(n+2)}{H_r(n+1)} &= \frac{x^r P(S_r = n+2)}{H_r(n+1)} + 1 + xq(1) + x \sum_{j=1}^{n-r} \frac{H_r(n+1-j)}{H_r(n+1)} (q(j+1) - q(j)) \\ &\quad + x \frac{H_r(r)}{H_r(n+1)} (q(n+2-r) - q(n+1-r)) \\ &\leq \frac{x^r P(S_r = n+2)}{H_r(n+1)} + 1 + xq(1) + x \sum_{j=1}^{n-r} \frac{H_r(n+1-j)}{H_r(n+1)} (q(j+1) - q(j)), \end{aligned}$$

since the last term to the left of the inequality is non-positive, as $q(\cdot)$ is decreasing. Therefore,

$$\begin{aligned} \frac{H_r(n+1)}{H_r(n)} - \frac{H_r(n+2)}{H_r(n+1)} &\geq \left[\frac{x^r P(S_r = n+1)}{H_r(n)} - \frac{x^r P(S_r = n+2)}{H_r(n+1)} \right] + \\ + x \sum_{j=1}^{n-r} \left(\frac{H_r(n-j)}{H_r(n)} - \frac{H_r(n+1-j)}{H_r(n+1)} \right) &(q(j+1) - q(j)). \end{aligned} \quad (13)$$

Note that by the induction assumption and the fact that the sequence $(q(n))_{n=1,2,\dots}$ is decreasing, the last term to the right of Eq. (13) is non-negative, as both factors of the summand term are non-positive (for the first factor, we apply a well-known property of log-concave sequences, namely twice positivity (see Ninh and Prekopa [37, p. 3018]). The first summand term to the right of Eq. (13) is non-negative, as well, as follows by multiplying both sides in Eq. (12) by $x^{-r} H_r(n) H_r(n+1)$. This completes the proof. \square

Proof of Theorem 1 This is an immediate consequence of Wang and Yeh [67, Thm. 2.3]), in which it was proved that if the polynomials

$$G_r(n, x) = \sum_{k=r}^n P(S_k \leq n) x^k, \quad x \geq 0, \quad n = r, r+1, \dots \quad (G_r(n, x) = 0, r > n)$$

are log-concave in n for any $x > 0$ and all $r = 0, 1, 2, \dots$, then the sequence

$$a(n, k) := P(S_k \leq n), \quad k \leq n$$

preserves log-concavity, and therefore

$$\sum_{k=0}^n a(n, k) p(k), \quad n = 0, 1, \dots$$

is a log-concave sequence, whenever $(p(n))_{n=0,1,\dots}$ is log-concave. Thus, the assertion is an immediate consequence of this result, together with the fact that

$$P(S_M \leq n) = \sum_{k=0}^n P(S_k \leq n) p(k), \quad n = 0, 1, 2, \dots$$

and Lemma 3. \square

Using an approximation procedure, Theorem 1 can easily be extended to settings where the summands are i.i.d. absolutely continuous random variables, with a decreasing density function.

Corollary 1. *Let $(Z_i)_{i=1,2,\dots}$ be a sequence of i.i.d. non-negative, absolutely continuous random variables, all distributed like a random variable Z with decreasing density f (and thus, concave distribution function F). Let M be an integer-valued random variable independent of the previous sequence, with pmf $(p(n))_{n=0,1,\dots}$. Consider the random sum S_M . If M has a discrete log-concave pmf, then S_M has a log-concave cdf.*

Proof. Note first that

$$G(x) := P(S_M \leq x) = \sum_{k=0}^{\infty} P(S_k \leq x) p(k), \quad x \geq 0.$$

We approximate the continuous random variable Z with a sequence of discrete distributions. First of all we divide \mathbf{R}_+ into intervals of length h and consider the sequence $(Z_i^h)_{i=1,2,\dots}$, in which each realization of Z_i is replaced by the smallest multiple of h that is larger, (i.e. $Z_i^h = h(\lfloor Z_i/h \rfloor + 1)$, with $\lfloor x \rfloor$ the integer part of x). It is clear that if Z has a decreasing density, then the variables Z_i^h are strictly positive, with a decreasing discrete density. Then, $P(Z_1^h + \dots + Z_M^h \leq nh)$, $n = 0, 1, 2, \dots$ is a log-concave sequence, by Theorem 1. Also, it is clear that Z_i^h converges in distribution to Z_i , as $h \rightarrow 0$. Moreover, the distribution function of $S_M^h := Z_1^h + \dots + Z_M^h$ (say G_h) converges to G , as $h \rightarrow 0$. In fact, we have

$$\begin{aligned} G(x) &= \sum_{k=0}^{\infty} P(S_k \leq x)p(k) = \sum_{k=0}^{\infty} \lim_{h \rightarrow 0} P(S_k^h \leq x)p(k) \\ &= \lim_{h \rightarrow 0} \sum_{k=0}^{\infty} P(S_k^h \leq x)p(k) = \lim_{h \rightarrow 0} G_h(x) \end{aligned}$$

Notice that the interchange of limit and summation is justified due to the Bounded Convergence Theorem, as we are dealing with distribution functions. Now consider \hat{G}_h , where $\hat{G}_h(x) = 0, 0 \leq x < h$ and, for $x \geq h$, $\log(\hat{G}_h(x))$ is obtained by linear interpolations between consecutive values of $\{\log G_h(nh)\}$ (and henceforth $\log(\hat{G}_h(x))$ is a concave function for $x \geq h$). We have, by construction, that \hat{G}_h is a log-concave distribution function. Also, by construction, $\{\hat{G}_h\}_{h>0}$ converges to G , just like $\{G_h\}_{h>0}$. Thus, G is log-concave, as the limit of log-concave functions is log-concave, hence obtaining the claimed result. \square

Remark 3. *As mentioned in Section 2, Theorem 1 and Corollary 1 extend the log-concavity property of the cdf of the random sum of i.i.d summands to a random number of terms M that has a log-concave pmf. Until now, the result was known only for very special classes of distributions, in particular when the random number of terms M has a Poisson, Binomial or Geometric distribution.*

Remark 4. *Theorem 1 and Corollary 1 cannot be generalized to settings where the random number of terms M has the weaker property of having a log-concave cdf (rather than pmf). As a counterexample, consider summands distributed like Z with a two-point non-increasing pmf $P(Z = 1) = P(Z = 2) = 0.5$. By Lemma 2 (a), if S_M has a log-concave cdf, then*

$$\begin{aligned} P(S_M = 1)P(S_M \leq 1) &\geq P(S_M = 2)P(S_M = 0) \Leftrightarrow \\ \frac{p(1)}{2} \left(p(0) + \frac{p(1)}{2} \right) &\geq \left(\frac{p(1)}{2} + \frac{p(2)}{4} \right) p(0) \Leftrightarrow \\ p(1)^2 &\geq p(0)p(2). \end{aligned}$$

However, it is well known that even strictly log-convex pmfs generate a cdf which is log-concave. An example, assume M is a Negative Binomial random variable with shape parameter $\alpha < 1$, i.e. $p(n) = \frac{\Gamma(\alpha + n)}{\Gamma(n + 1)\Gamma(\alpha)} p^\alpha (1 - p)^n$, $n = 0, 1, 2, \dots$. Because of the strict log-convexity of its pmf, $p(1)^2 < p(0)p(2)$.

Remark 5. *The previous remark exhibits that log-concavity of the pmf of M is essential to obtain that S_M has a log-concave cdf for an arbitrary concave summand. However, for specific classes of summand distributions, weaker properties for M suffice. An example is Ross et al. [45] who in Theorem 7.1. showed that with exponentially distributed summands, M may be DFR, i.e. may have a decreasing failure rate (Distributions with a log-concave pmf have an increasing failure rate, see e.g. Corollary 2 in Bagnoli and Bergstrom [8]).*

An application in risk theory: Log-concavity of non-ruin probability in the classical risk model. Consider the classical risk model (see, for instance, Asmussen and Albrecher [6]). Claims are presented to an insurance pool, in accordance with a Poisson process $(N(t), t \geq 0)$ with rate λ . The associated claim amounts $(X_i)_{i=1,2,\dots}$ are independent and identically distributed positive random variables with finite mean, which are also independent of $(N(t), t \geq 0)$. Suppose that the initial capital of the insurance pool is $U(0) := u \geq 0$ and that premia are earned at a constant rate c . The probability of eventual ruin $\psi(u)$ is the probability that the pool ever runs out of money, i.e.

$$\psi(u) = P \left(\inf_{t \geq 0} \left(u + ct - \sum_{i=0}^{N(t)} X_i \right) < 0 \right), \quad u \geq 0. \quad (14)$$

Assume, without loss of generality, that the pool, in the long-run, is profitable in expectation, i.e. $p := \frac{\lambda EX_1}{c} < 1$. Under this condition, it is well-known (see, for instance, Asmussen and Albrecher [6]) that

$$1 - \Psi(u) = P(L \leq u), \quad u \geq 0,$$

where L can be written as

$$L = \sum_{i=0}^M L_i, \quad (15)$$

in which M has a geometric distribution with success probability p , and $(L_i)_{i=1,2,\dots}$ are identically distributed positive random variables (called the record lows), independent on M , and such that L_i has the equilibrium distribution of X_i , i.e.,

$$P(L_i \leq x) = \frac{1}{EX_i} \int_0^x P(X_i > s) ds, \quad x \geq 0. \quad (16)$$

As a consequence of Corollary 1, we have the following:

Corollary 2. *In the above standard risk model, the 'non-ruin' probability $1 - \Psi(u)$ is a log-concave function of the pool's initial cash balance u .*

Proof. Geometric random variables have a log-concave pmf, see e.g. Bagnoli and Bergstrom [8]. Observe also that the record lows have a concave distribution function, as their density is $f(x) = \frac{1}{EX_i} P(X_i > x)$, $x \geq 0$ which is obviously a decreasing function. Thus, the result is an immediate consequence of Corollary 1. \square

Remark 6. *One application where the log-concavity of the 'non-ruin' probability is very useful, is when an insurance company underwrites several, say N , independent insurance pools each subject to a risk model of the above type and with initial endowments u_i , $i = 1, \dots, N$, and non-ruin probabilities $1 - \Psi_i(u_i)$, $i = 1, \dots, N$. Assume the company wants to allocate its initial cash balance U among the pools so as to maximize the likelihood that all pools remain solvent, i.e. that no pool runs out of money. Replacing this objective, equivalently, with its logarithm, results, in view of Corollary 1, in a separable concave objective in the pools' initial endowments u_i , $i = 1, \dots, N$ subject to the single budget constraint $\sum_{i=1}^N u_i = U$, the unique optimal solution of which can be found very efficiently, see e.g. Zipkin [73] and the references therein.*

Remark 7. *Due to the well-known duality between risk and queueing models (see, for instance, Asmussen and Albrecher [6]) we obtain, in the same manner as Corollary 2 that the stationary workload (virtual waiting time) in a $M/G/1$ queue has a log-concave cdf.*

A classical replacement model A production system is subject to a sequence of random shocks generated by a renewal process with interarrival times having a log-concave pdf. Each shock causes a random amount of damage Z . The sequence of consecutive damage quantities $(Z_i)_{i=1,2,\dots}$ is i.i.d. with a common cdf $F(\cdot)$. The system may fail at any shock epoch, and the survival likelihood is a function of the cumulative damage incurred. Thus, if up to time t the cumulative damage incurred equals $X(t) = x$, and a shock occurs adding a new damage quantity z , then the survival probability equals $r(x+z)$, with $r(\cdot)$ a non-increasing function. K denotes the replacement cost at failure and $C(x) < K$ when the system is replaced before failure, with $C(\cdot)$ an increasing function and x the cumulative damage incurred. While the system is working, it incurs an income at a rate $I(x)$, again a function of the cumulative damage x .

The above is a standard replacement model. For the special case of a Poisson shock process, it was first analyzed by Zuckerman [76], generalizing Esary, Marshall and Proschan [19] and Taylor [60] who considered further special cases where a unit of damage is incurred at each shock and where $C(\cdot)$ and $I(\cdot)$ are constant functions, respectively. Zuckerman showed that the long-run average profit is maximized when the system is replaced as soon as the cumulative damage exceeds a critical threshold ξ^* (or up to failure, whichever occurs first), i.e. at a stopping time T given by

$$T := \min\{\inf\{t \geq 0 : X(t) \geq \xi^*\}, \delta\}, \quad (17)$$

where δ denotes the system's failure time, see (2.4) in Zuckermann.

The same class of threshold policies may be used, effectively, when the shock process is a general compound renewal process, although for general distributions, it is difficult to compute the distribution of $X(t)$, $t > 0$, and hence of T . However, by Theorem 4.5. in Badía and Sangüesa [7] we deduce that $N(t)$, the number of shocks up to time t has a log-concave pmf, as long as the interrenewal times have a log-concave pdf. Thus, assuming that F is concave, it follows from Corollary 1 that $X(t)$ has a log-concave cdf. This result maybe used, for example, to fit a log-concave distribution to empirical data. To illustrate this, fix $t \geq 0$. Assume that $X(t)$ has a discrete distribution, and let $\pi_k = -\log P(X(t) \leq k)$. Assume we know the mean and standard deviation of $X(t)$, and for a limited set of integer values I , we have empirical estimates $\{\pi_k^0 =, k \in I\}$ for $\{\pi_k =, k \in I\}$. The fitting problem can then be formulated as:

$$\min \sum_{k \in I} (\pi_k - \pi_k^0)^2 \quad (18)$$

$$(19)$$

$$\text{s.t.} \quad \sum_{k \in I} (1 - e^{-\pi_k}) = \mu \quad (\text{first moment}) \quad (20)$$

$$2 \sum_{k \in I} k(1 - e^{-\pi_k}) = \mu^2 + \sigma^2 \quad (\text{second moment}) \quad (21)$$

$$\pi_{k+1} \leq \frac{1}{2}(\pi_k + \pi_{k+2}), \quad k = 0, 1, \dots \quad (\text{log-concavity of cdf}) \quad (22)$$

$$0 \leq \pi_{k+1} \leq \pi_k, \quad k = 0, 1, \dots \quad (\text{monotonicity of cdf}) \quad (23)$$

Instead of the L_2 -norm, a different norm of $\{\pi_k - \pi_k^0, k \in I\}$ may be optimized in Eq. (19).

The maximization problem has a separable convex objective function and all but two constraints are linear. Klincewicz and Whitt [32] used a similar approach to derive a lower and an upper bound for the mean steady state queue length in the $GI/M/1$ queueing model, assuming the pmf of the inter-arrival time distribution is log-concave. In their numerical study, the authors showed that the addition of the log-concavity shape property reduces the gap between the upper and lower bound to just a few percentage points, resulting in truly robust estimates.

The next result, which generalizes the above-mentioned Theorem 4.5. in Badía and Sangüesa [7] will be used, repeatedly, in the next section.

Corollary 3. (a) Let $\{N(t) : t \geq 0\}$ be a renewal process. Let $\{N^*(t) : t \geq 0\}$ denote the associated equilibrium renewal process. Let L be a random variable independent of the previous processes $\{N(t) : t \geq 0\}$ and $\{N^*(t) : t \geq 0\}$, with a log-concave pmf. Define $M = N(L)$ and $M^* = N^*(L)$. Let $(Z_i)_{i=1,2,\dots}$ be a sequence of i.i.d. summands with a concave cdf. Then,

$$S_M = \sum_{i=1}^M Z_i \quad \text{and} \quad S_{M^*} = \sum_{i=1}^{M^*} Z_i \quad \text{have a log-concave cdf.}$$

(b) In the specific case where $Z_i = 1$, $S_M = M$ and $S_{M^*} = M^*$ have a log-concave pmf.

Proof. In view of Corollary 1, it suffices to show that both M and M^* have a log-concave pmf. The latter follows immediately from Theorem 4.5 and Remark 4.6 in Badía and Sangüesa [7]. \square

4 Implications for inventory models

In this Section we establish the optimality of simply structured policies in various standard inventory models, significantly generalizing the conditions for the demand and leadtime processes, in conjunction with the cost structure in these models. We also identify implications of our results for various algorithmic approaches to compute optimal policy parameters.

4.1 Continuous review Inventory Model with Compound Renewal demand Process

Consider the following standard inventory model. Assume that demands are generated by a compound renewal process. The model includes the frequently used periodic review model with i.i.d. demands, as a special case, selecting a deterministic interdemand time (We treat this special case in the Appendix).

Orders may be placed at any demand epoch, but not between demand epochs. When the interarrival times of the demands are exponential, that is, when the demand process is Poisson, this restriction is without loss of generality, due to the memoryless property of the exponential distribution. For general interdemand time distributions, this restriction is, potentially, with some loss of optimality in as much as it is standard in all Semi-Markov Decision Problems (SMDP), in general, and all inventory models with compound renewal demands, in particular.

An order arrives after a leadtime L which is assumed to be constant, or more generally characterized by a stochastic process $\{L(t) : t \geq 0\}$ with $L(t)$ the leadtime experienced by an order placed at time t . We assume that the demand process is exogenous, that is, independent of the demand process, and sequential, that is, $t + L(t) \leq t' + L(t')$ for all $t < t'$, with probability one. Under sequential leadtime processes, orders do not cross. Let L denote a random variable distributed like the steady-state distribution of the process $\{L(t) : t \geq 0\}$. We refer to Zipkin [74] for an extensive discussion of such processes and their applications.

Versions of this model go back to Beckman [9], Tijms [61], Federgruen and Schechner [21], Sahin [46, 47], Zipkin [74] and Chen and Zheng [14]; for recent treatments, see Rosling [43], Chen and Simchi-Levy [16], Shi, Katehakis and Melamed [59], Shi et al. [58], Federgruen and Wang [22] and Badía and Sangüesa [7]. For the case of Poisson demand processes, the model goes to back to the seminal textbook by Hadley and Whitin [30]; see, also, Archibald and Silver [4].

The cost structure consists of order costs and inventory and shortage costs. As far as the former are concerned, the cost of an order has a fixed and a variable component, where the latter grows proportionally with the size of the order. (See, however, Section 4.2 for more general order cost structures). Let y_t denote the inventory position at time t , defined as the inventory level plus the current and any outstanding orders. Under backlogging, and since orders do not cross under an exogenous, sequential leadtime process, I_{t+L} , the inventory level at the end of period $t + L$ is given by:

$$I_{t+L} = y_t - D[t, t + L], \quad (24)$$

with $D[t, t + L]$ the aggregate demand in the time interval $[t, t + L]$. Thus, as long as all inventory and shortage costs are determined by inventory levels, the expectation of the costs associated with I_{t+L} can be specified as a function $G(y_t)$, and hence, accounted for at time t .

It is well known that an (s, S) policy acting on the inventory position, optimizes long-run average or discounted costs, as long as the function $G(\cdot)$ is quasi-convex (See Federgruen and Wang [22] for a rigorous extension from periodic review to continuous review settings).

For infinite horizon models, this optimality result is due to Zheng [71], after Veinott [64] had obtained it for periodic review, finite horizon models. Recently, Feinberg and Liang [25] refined Zheng's optimality result to allow for continuous demand distributions and they established convergence properties of optimal discounted thresholds for discounted problems to the corresponding thresholds for average cost problems. These results are based on their earlier paper, Feinberg and Liang [24].

Most standard inventory models, naturally assume that inventory holding costs are proportional (or convexly increasing) with the inventory levels. However, the consequences of shortages are more varied and multi-faceted. In some settings, there are explicit out-of-pocket costs, for example a one time penalty $b > 0$ for any demand unit that is backlogged. This cost component arises with considerable and increasing frequency: many retailers in the U.S. impose such a penalty on their suppliers; for example, in 2017, Target instituted a 5% penalty for any demand unit delivered, beyond its specific 2-day delivery window and Walmart initiated a similar 3% penalty, see Vigil [66]. Some suppliers sell exclusively to one of the national retailers (such as Walmart and Target). Even if a supplier sells to multiple retail organizations, it typically maintains dedicated inventories for any national retail organization in its customer base.

In addition to these direct cost consequences, one often considers service level constraints, precisely because estimating the costs associated with backlogging and loss of goodwill is, typically, very difficult.

The most frequently used service level measures, both in the literature and practice, are the fill-rate and the ready-rate, the percentage of demand units that are fulfilled in time and the percentage of time without any shortage, respectively, see e.g. Chain Store Age [12] and Özer and Xiang [38].

Muckstadt [36] and Cohen, Deshpande and Wang [18] describe various examples from military systems or the semi-conductor industry where inventory management is driven by the above two service constraints. Lee, Billington and Carter [34] document the same for the inventory management practices at Hewlett-Packard: 'for sites that face end-customers there is usually a target service level (fill-rate) desired by management'. See Bertsimas and Paschalidis [10], Chen and Krass [15] and Wei, Jasin and Xin [68] for additional discussion of the disadvantages of cost-based models.

It is prohibitively difficult to solve multi-period inventory problem with such service level constraints. The standard approach is, therefore to dualize the constraints, adding a one time delay penalty b for any backlogged demand unit, and a penalty π for every time unit with shortages,

respectively.

When backlogging costs are accrued at a rate proportional to the backlog size, i.e. if only the cost structure (i) prevails it is easy to verify that the function $G(\cdot)$ is convex, and certainly, quasi-convex, under arbitrary leadtime and demand size distributions.

However, quasi-convexity of the inventory cost function $G(\cdot)$ is much harder to establish under the above alternative and practically important cost structures, or under the above service level constraints, even after dualization. Thus, while these cost structures and service level constraints were prominently discussed in the early textbooks on inventory theory, e.g. Hadley and Whitin [30] and Peterson and Silver [40], they were largely abandoned in view of their perceived intractability. However, Rosling [43] and Huh et al. [31] provided seminal results for lead-time and demand conditions under which the above cost structures can be handled and optimal policies are of a simple structure. The results in this paper allow us to generalize these conditions considerably, see Theorem 2 (a) and Remark 8, below.

At any time epoch with an inventory position of y , let $G(y)$ denote the aggregate expected inventory and shortage cost incurred one leadtime L later. Let $(Z_i)_{i=1,2,\dots}$ denote the sequence of demand sizes at consecutive demand epochs with common cdf Ψ and mean μ . Let Ψ_r denote the r -fold convolution of the cdf Ψ , $r \geq 1$. If M is random, Ψ_M similarly denotes the cdf of the random sum $Z_1 + \dots + Z_M$. Let $(X_i)_{i=1,2,\dots}$ denote the sequence of interdemand times, again with a common distribution with mean $1/\lambda$. Thus, $D := \lambda\mu$ denotes the long-run expected demand per time unit. We denote by $\{N(t) : t \geq 0\}$ the (renewal) process of demands epochs and by $\{\tilde{N}(t) : t \geq 0\}$ its associated equilibrium process.

Apart from the alternative afore-mentioned backlogging cost structures, we will assume, as usual, that each unit of inventory (backlog) incurs a cost of h (p) per unit of time.

Rosling [43] shows that $G(y)$, the total expected inventory and shortage cost rates assigned to any time epoch with an inventory position of y units is

$$G(y) = h\mathcal{I}(y) + pB_p(y) + \pi B_\pi(y) + bB_b(y), \quad (25)$$

where

$$\mathcal{I}(y) = [y - E[L]D] + \int_y^\infty [1 - \Psi_{\tilde{N}(L)}(u)]du, \quad (26)$$

$$B_p(y) = \int_y^\infty [1 - \Psi_{\tilde{N}(L)}(u)]du, \quad (27)$$

$$B_\pi(y) = \int_y^\infty [\Psi_{N(L)}(u) - \Psi_{N(L)+1}(u)]du. \quad (28)$$

$$B_b(y) = 1 - \Psi_{\tilde{N}(L)}(y). \quad (29)$$

To show the optimality of an (s, S) policy, it is sufficient to show that this inventory cost rate function $G(\cdot)$ is quasi-convex. The following represents our main application to inventory models.

Theorem 2. *Assume the demand process is compound renewal, and that both the interdemand times and leadtime variables have a log-concave density, with the leadtime process independent of the demand process.*

- (a) *An (s, S) policy is optimal under fully general h , p , π and b parameters if the demand sizes $(Z_i)_{i=1,2,\dots}$ are i.i.d. with a concave cdf.*
- (b) *See Proposition 5.2(b) in Badía and Sangüesa [7]. If $b = 0$, an (s, S) policy is optimal under fully general h , p and π parameters, if the demand sizes $(Z_i)_{i=1,2,\dots}$ are i.i.d., with a common*

cdf that is MCR (See Remark 10 for the definition of this property, which is weaker than the cdf being log-concave).

Proof. Part (a): The optimality of an (s, S) policy is guaranteed if G is quasi-convex. This result can be shown by observing that the inventory control problem embedded on demand epochs is a Semi-Markov Decision Problem (SMDP) with identical, i.e., state invariant transition times between decision epochs. With a classical data-transformation, this SMDP is equivalent to a standard periodic review model in which an (s, S) policy is optimal, see e.g. Tijms [62, chapter 7]).

To prove quasi-convexity of $G(\cdot)$ in this general model, with $b \geq 0$ we need, appart from the sufficient conditions for Theorem 2(b) to hold, the log-concavity of the cdf $\Psi_{\tilde{N}(L)}$ (see Rosling [43, Proposition 2-2] for more details). The latter was shown in Corollary 3 with the assumed properties for the leadtimes, interdemand times and demand sizes. \square

Remark 8. *To date, optimality of an (s, S) policy under compound renewal demands and general h, p, π and b parameters was only shown when leadtimes are deterministic or convolutions of exponentials, and interdemand times that are exponential as well, see Rosling [43, Proposition 2-2]) and Badía and Sangüesa [7, Proposition 1.3].*

4.2 The case of (simple) renewal demand processes

In this Subsection, we consider the special case where the demand process is renewal i.e. each arriving customer requests a single unit of our item.

4.2.1 Optimality of (s, S) policies under general order cost functions

Instead of restricting ourselves to fixed-plus-linear order costs, as we have done thus far, assume the order cost function is fully general and possibly infinitely valued on some part of the positive half line, (as long as it is finite for some positive order quantity). We thus include the possibility of capacity limits, all-unit discounts, incremental unit discounts, truckload discounts, general convex or concave order costs, batch ordering or minimal order quantities. In addition, let inventory and shortage costs be incurred at a rate which is given by a general quasi-convex function $g(\cdot)$, this includes the tree types of shortage cost functions, covered above, as special cases. Charging, one more time, at any time t with an inventory position of y units the expected value of these costs, one leadtime later, we get an inventory position based cost function

$$G(y) = Eg(y - \tilde{N}(L)), \quad (30)$$

since demand sizes are equal to one.

The above general model was recently studied by Perera et al. [39] who showed that an (s, S) policy is optimal, either when there is no leadtime, or when the leadtime is deterministic and the demand process Poisson. We extend this result to allow for stochastic leadtimes and a general renewal demand process, merely requiring that the leadtime and interdemand distribution belong to the broad category of distributions that have a log-concave density.

Theorem 3. *Assume the demand process is a renewal process. Consider a fully general order cost function $c(\cdot)$ which may be infinitely valued on part of the positive half line. Let $g(\cdot)$, the instantaneous inventory and cost order function be a general quasi-convex function. An (s, S) policy, acting on the inventory position is optimal, provided both the leadtime and interdemand time distribution have a log-concave density.*

Proof. By the proof of Theorem 1 in Perera et al. [39], it suffices to show that the function $G(\cdot)$ is quasi-convex. By Lemma 5.4. in Huh et al. [31] this is guaranteed as long as $\tilde{N}(L)$ has a log-concave pmf. The result follows under the assumed log-concavity properties, as mentioned in the proof of Corollary 3, itself the direct result of Theorem 4.5. and Remark 4.6. in Badía and Sangüesa [7]. \square

The astounding generality of the above optimality result is due to the fact that under a renewal demand process, the inventory position process is "skip-free", i.e. when decreasing, it decreases in unit steps, only.

As mentioned in the Introduction, in a skip-free system, an (s, S) policy reduces to a (r, q) policy: an order of fixed size q is placed whenever the inventory position decreases to the level r .

In a periodic review setting, or under compound renewal processes, the system fails to be skip-free. Hence, it is well known that an (s, S) policy fails to be optimal under general order cost functions, see e.g. Porteus [41, 42], or the case where a capacity limit prevails in conjunction with a fixed-plus-linear order cost, see Shaoxiang and Lambrecht [56], Shaoxiang [55] and Federgruen, Liu and Li [20].

4.2.2 Unknown leadtime demand distribution: a data driven approach

Levi et al. [33] consider the single period newsvendor model where the demand distribution is unknown, but a sample of realizations from this distribution is available. In the newsvendor problem, an average cost $h > 0$ is incurred for every unit that remains unsold, and an underage cost $b > 0$ for every unit of demand that cannot be fulfilled because of an insufficient supply. It is elementary that expected costs are minimized by the $b/(b+h)$ -the fractile of the demand distribution, i.e.

$$S^* = \inf\{S : F(S) \geq \frac{b}{b+h}\}.$$

When the true cdf $F(\cdot)$ is unknown it is natural to replace $F(\cdot)$ by the empirical distribution $\hat{F}_N(\cdot)$ derived from a given sample of size N that is drawn from this distribution. In addition, one may want to select a slightly larger fractile from this empirical distribution. Thus, define, for all $\alpha > 0$:

$$\hat{S}_N^\alpha = \inf\left\{S : \hat{F}_N(S) \geq \frac{b}{b+h} + \frac{1}{2} \frac{\alpha}{b+h}\right\}.$$

Note, \hat{S}_N^α is a random variable since it depends on the specific sample of observations encountered. In the absence of any information about the shape of the distribution, the best known bound for the optimality gap of any of the empirical fractiles \hat{S}_N^α is due to Levi, Roundy and Shmoys [35]. However, when the distribution is known to have a log-concave density, Levi et al. [33, Theorem 4]) recently derived the following tight probabilistic bound for the likelihood that a specific $b/(b+h)$ -related fractile of the empirical distribution is ϵ -optimal. Recall that if a distribution has a log-concave density, it is unimodal, say with mode μ .

Proposition 1. (See Levi et al. [33, Theorem 4]). *Assume f is log-concave and $S^* \geq \mu$. Fix $\epsilon > 0$ and let $\alpha = \sqrt{2\epsilon b h \min(b, h)/(b+h)} + O(\epsilon)$. Then \hat{S}_N^α is ϵ -optimal with probability at least $1 - 2U^*(\epsilon)$ where $U^*(\epsilon) \sim \exp\left(-\frac{1}{4}N\epsilon \frac{\min(b, h)}{b+h}\right)$, as $\epsilon \rightarrow 0$.*

A numerical study in Levi et al. [33] shows that, for a given value of $\epsilon > 0$, and associated confidence level, the general Levi et al. [35] bound requires a sample size approximately 1000 times that required by the bound in Proposition 1, for various families of log-concave densities.

While the oldest and most basic of stochastic inventory models, the newsvendor model is of limited practical usage, since dealing with a single sales period and a single procurement opportunity. Consider, now, the class of infinite horizon models with a renewal demand process, and ordering opportunities at any one of the demand epochs, as considered in this Section. Moreover, consider the important special cost structure, when ordering, holding and backlogging costs are linear, without any fixed cost components. It is well known that a base-stock policy with order-up-to-level S optimizes long-run average costs, under fully general leadtime demand distributions. (A base-stock policy is an (s, S) policy with $s = S$.) Moreover, the optimal choice of the base-stock level S minimizes

$$Eh(S - LD)^+ + Eb(LD - S)^+, \quad (31)$$

with LD the steady state leadtime demand distribution. This is analogous to the objective function $Eh(S - D)^+ + Eb(S - D)^+$ in the newsvendor problem. Therefore, if the exact cdf of the leadtime demands distributions is unknown, but a sample of N i.i.d. historical observations from this distribution is available, it is again possible to obtain an ϵ -optimal approximation for S^* via \hat{S}_N^α , now with $\hat{F}_N(\cdot)$ denoting the empirical sample distribution of the leadtime demand. The following probabilistic bound follows immediately from Corollary 3 (b):

Theorem 4. *Consider an infinite horizon inventory model with a renewal demand process and linear ordinary, holding and backlogging costs resulting in the long-run average cost objective Eq. (31). Assume that the interdemand time as well as the leadtime distribution has a log-concave density. Fix $\epsilon > 0$ and let $\alpha = \sqrt{2ebh \min(b, h)/(b + h)} + O(\epsilon)$. Then \hat{S}_N^α is ϵ -optimal with probability at least $1 - 2U^*(\epsilon)$ where $U^*(\epsilon) \equiv \exp\left(-\frac{1}{4}N\epsilon\frac{\min(b, h)}{b+h}\right)$, as $\epsilon \rightarrow 0$.*

As a final remark, Theorem 4 continues to apply under periodic review models or compound renewal demand processes, as long as the leadtime demand distribution can be either (i) shown to, or (ii) assumed to have a log-concave pdf; see the next Section for a few sufficient conditions.

4.3 Unknown leadtime demand distribution: a robust optimization approach

We return to the general setting of Subsection 4.1 where the demand process is a general compound renewal process and inventory can be controlled at any demand epoch: as mentioned, the periodic-review, periodic-demand model may be viewed as a specific case with deterministic interdemand times. The long-run average cost under a base-stock policy with base-stock level S , and standard inventory carrying and shortage costs is of the form

$$\varphi(S) = E\Phi(S, LD) = Eh(S - LD)^+ + Eb(LD - S)^+, \quad (32)$$

as in Eq. (31). Here, $h > 0$ and $b > 0$ are given constants. Alternatively, they may be convex functions and Eq. (32) has to be adapted to $Eh((S - LD)^+) + Eb((LD - S)^+)$. Recall that in the absence of fixed order costs, a base-stock policy is, in fact, optimal. Moreover, decision makers often restrict themselves to base-stock policies, even if this is suboptimal.

For a known distribution of the leadtime demand LD , the problem of finding the optimal base-stock level S^* thus reduces to minimizing the convex single variable function $\varphi(\cdot)$. However, as mentioned in the previous Section, the (full) leadtime demand distribution is often unknown.

In the previous subsection, we showed how a data-driven empirical distribution may be used as a substitute. An alternative approach is to restrict leadtime distributions to a given class \mathcal{LD} and to solve the min-max problem

$$\min_S \max_{LD \in \mathcal{LD}} E\Phi(S, LD). \quad (33)$$

This represents the robust optimization approach. In the single stage newsvendor problem, this approach was initiated by Scarf [49] and extended by Gallego and Moon [28, 29] and Gallego [26, 27].

Recently, this approach has been extended to finite and infinite horizon inventory models with repeated procurement opportunities; see, for example, See and Sim [51] and Choi, Ruszczyński and Zhao [17]. Typically, the class of distributions \mathcal{LD} is specified by fixing the value of a given number of moments of the distribution. For example, when specifying the first two moments of the distribution, it follows from Proposition 6.10 in Shapiro, Dentcheva and Ruszczyński [57] that the leadtime distribution which achieves the maximum in Eq. (33) has support on at most three points. A far more reasonable but still robust bound can be obtained when these moment conditions are complemented with (or replaced by) the specification that the leadtime demand distribution has a log-concave cdf, a property for which Corollary 3 provides sufficient conditions. Focusing on the case where the leadtime demand distribution is discrete with known mean μ and standard deviation σ , the min-max problem in Eq. (33) can then be formulated as the mathematical program Eqs. (20)-(22), so that the objective function Eq. (19) is replaced by

$$\min_S \max_{\pi_k} \left\{ h(S - \mu) + (h + b) \sum_{k=S}^{\infty} (1 - e^{-\pi_k}) \right\}. \quad (34)$$

5 Log-concavity of the density of random sums

In certain settings, it is important to establish that the pdf or pmf of $S_M = \sum_{i=1}^M Z_i$ is log-concave, a stronger property than log-concavity of the cdf, even though this property is satisfied under many of the most commonly used classes of distributions. As an example, establishing this property greatly facilitates maximum likelihood estimation, see the recent survey paper by Samworth [48]. In this Section, we review the known sufficient conditions for the stronger log-concavity property. They are unfortunately considerably more restrictive than those of Theorems 1 and Corollary 1.

To cover the simple case where $P(M \geq 1) = 1$, the distribution of the summands $(Z_i)_{i=1,2,\dots}$ must have a log-concave density, itself. But, this property is clearly not sufficient: If Z is a deterministic constant different from one, the pmf of S_M has internal zeroes assuming $P(M = 1) > 0$ and $P(M = 2) > 0$. But a pmf with internal zeroes cannot be log-concave; see Definition 1 and Remark 1.

Furthermore, all known conditions for S_M to have a log-concave cdf include the requirement that the summands $(Z_i)_{i=1,2,\dots}$ have a decreasing pdf or pmf. Since log-concavity of the pdf (or pmf) of S_M is a stronger property, the same monotonicity property for the density of the Z -variable appears to be required. However, the following example shows that it fails to be sufficient.

Let M take on the values 1 and 2, so that $p(2) = 1 - p(1)$. Similarly, let Z take on the values 1 and 2 with equal probability 0.5. Then

$$P(S_M = 2) = \frac{p(1)}{2} + \frac{p(2)}{4}, \quad P(S_M = 3) = \frac{p(2)}{2} \quad \text{and} \quad P(S_M = 4) = \frac{p(2)}{4},$$

and therefore

$$P(S_M = 3)^2 - P(S_M = 2)P(S_M = 4) = \frac{p(2)^2}{4} - \frac{p(1)p(2)}{8} - \frac{p(2)^2}{16} = \frac{3p(2)^2}{16} - \frac{p(1)p(2)}{8},$$

which is nonnegative if and only if $3p(2)/2 \geq p(1) = 1 - p(2)$, hence if and only if $p(1) \leq 3/5$. Thus if $p(1) > 3/5$, the pmf of S_M is not log-concave.

Similar difficulties arise when Z has a continuous distribution, with support on an interval $[a, b]$. Let M take on the values 1 and 2 only. Then, S_M has support on $[a, b] \cup [2a, 2b]$. Thus, if $b < 2a$, the support of S_M is the union of two distinct intervals, which precludes log-concavity of the pdf by its very definition.

To our knowledge, the following are the only known sufficient conditions:

Proposition 2.

- (a) If M has a log-concave pmf on the set of non-negative integers and the summands $(Z_i)_{i=1,2,\dots}$ have an exponential or Erlang(2)-distribution, then the absolutely continuous part of S_M has a log-concave pdf.
- (b) If M has a log-concave pmf on the set of non-negative integers and the summands $(Z_i)_{i=1,2,\dots}$ have a geometric or Negative Binomial (2)-distribution, then the pmf of S_M $(p(n))_{n=0,1,\dots}$ is a log-concave sequence for $n \geq 1$.
- (c) If M has a Poisson distribution with mean λ and the summands $(Z_i)_{i=1,2,\dots}$ have a log-concave pmf $(q(n))_{n=0,1,\dots}$, then S_M has a log-concave pmf, provided that $\lambda q(1)^2 \geq 2q(2)$.

Proof. Parts (a) and (b) follow from Nihn and Prekopa [37, Theorems 2.2, 3.2 and 5.2]), and part (c) by Yu [69]. □

Acknowledgments

This work has been supported by the Spanish research project MTM2015-63978 (MINECO/FEDER). The first and third authors acknowledge the support of DGA S11 and E64, respectively.

References

- [1] Albrecher, H., Cheung, E.C.K. & Thonhauser, S. (2011). Randomized observation periods for the compound Poisson risk model dividends *Astin Bulletin* **41**, no. 2.
- [2] Albrecher, H., Cheung, E.C.K., & Thonhauser, S. (2013). Randomized observation periods for the compound Poisson risk model: the discounted penalty function. *Scandinavian Actuarial Journal* **6**, 424–452.
- [3] An, M.Y. (1998). Log-concavity versus log-convexity: A complete characterization. *Journal of Economic Theory* **80**: 350–396.
- [4] Archibald, B. & Silver, E. (1978). (s,S) policies under continuous review and discrete compound Poisson demands *Management Science* **24**: 899–908.
- [5] Arrow, K., Karlin, S. & Scarf, H. (1958). *Studies in the mathematical theory of inventory and production*. Stanford University Press: California.
- [6] Asmussen, S. & Albrecher, H. (2010). *Ruin Probabilities*. Hackensack, NJ: World Scientific Publishing.
- [7] Badía, F.G. & Sangüesa, C. (2015). Inventory models with nonlinear shortage costs and stochastic lead times; Applications of shape properties of randomly stopped counting processes. *Naval Research Logistics* **62**, 345–356.

- [8] Bagnoli, M. & Bergstrom, T. (2005). Log-concave probability and its applications. *Economic Theory* **26**: 445–469.
- [9] Beckmann, M. (1961). An inventory model for arbitrary interval and quantity distribution of demand *Management Science* **8**: 35–57.
- [10] Bertsimas, D & Paschalidis, I.Ch. (2001). Probabilistic service level guarantees in a Make-to-Stock manufacturing system *Operations Research* **49**, 119-133.
- [11] Cai, J. & Willmot, G.E. (2005). Monotonicity and aging properties of random sums. *Statistics and Probability Letters* **73**: 381–392.
- [12] Chain Store Age (2002). Inventory mangement 2002: data, detail and discipline: Supermarkets. *December I.P. 17A*.
- [13] Chen, F. (2000). Optimal policies for multi-echelon inventory problems with batch ordering *Operations Research* **48**, 376-389.
- [14] Chen, F. & Zheng, Y.S. (1993). Inventory models with general backorder costs *European Journal of Operations Research* **65**, 175-186.
- [15] Chen, F.Y. & Krass, Y.D. (2001). Inventory models with minimal service level constraints *European Journal of Operations Research* **134**, 120-140.
- [16] Chen, X. & Simchi-Levi, D. (2006) Coordinating inventory control and pricing strategies: the continuous review model. *Operations Research Letters* **34**, 323-332.
- [17] Choi, S., Ruszczyński, A. & Zhao, Y. (2011) A multiproduct risk-averse newswendor with law-invariant coherent measures of risk. *Operations Research* **59**, 346-364.
- [18] Cohen, M.A., Deshpande, V. & Wang, Y. (2003). Managing supply chains with differentiated service requirements in *Practise of Supply Chain Management: Where Theory and Applications Converge*. Harrison, T. Lee, H. & Neale, J. (eds). Chapter 11, Boston: Kluwer.
- [19] Esary, J.D., Marshall, A.W. & Proschan, F. (1973) Shock models and wear processes. *Annals of Probabability* **1**, 627-649.
- [20] Federgruen, A. Liu, Z. & Li, L. Synthesis and Generalization of Structural Results in Inventory Management: A Generalized Convexity Property *to appear in Mathematics of Operations Research*
- [21] Federgruen, A. & Schechner, Z. (1983). Cost formulas for continuous review inventory models with fixed delivery lags *Operations Research* **31**, 957-965.
- [22] Federgruen, A. & Wang, M. (2015). Inventory models with shelf age and delay-dependent inventory costs *Operations Research* **63**, 155–186.
- [23] Federgruen, A. & Zheng, Y. (1992). An efficient algorithm for computing an optimal (r,Q) policy. *Operations Research* **40**, 808-813.
- [24] Feinberg, E.A. & Liang, Y. (2017a). Structure of optimal policies to periodic review inventory models with convex costs and backorders for all values of discount factors *Annals of Operations Research* <https://doi.org/10.1007/s10479-017-2548-6>

- [25] Feinberg, E.A. & Liang, Y. (2017b). On the optimality equation for average cost Markov decision processes and its validity for inventory control *Annals of Operations Research* <https://doi.org/10.1007/s10479-017-2561-9>
- [26] Gallego, G. (1998). New bounds and heuristics for (Q, r) policies. *Management Science* **63**, 701–715.
- [27] Gallego, G. (2001). Minimax analysis for finite-horizon inventory models *IIE Transactions* **33**, 861–874.
- [28] Gallego, G. & Moon, I. (1993). The distribution free newsboy problem: review and extensions *Journal of the Operations Research Society* **44**, 825–834.
- [29] Gallego, G. & Moon, I. (1994). Distribution free procedures for some inventory models *Journal of the Operations Research Society* **45**, 651–658.
- [30] Hadley, G. & Whitin, T. (1963). *Analysis of Inventory Systems..* Prentice-Hall, Englewood Cliffs.
- [31] Huh, W.T., Jankiraman, G., Muharremoglu, A. & Anshul, S. (2011). Inventory systems with a generalized cost model *Operations Research* **4**, 1040-1047.
- [32] Klineciewicz, J.G. & Whitt, W. (1984). On approximations for queues, II: shape constraints *AT & T Bell Laboratories Technical Journal* **63**, 139–161.
- [33] Levi, R. Perakis, G. & Uichanco, J. (2015). The data-driven newsvendor problem: New bounds and insights. *Operations Research* **63**, 1294-1306.
- [34] Lee, H.L., Billington, C. & Carter, B. (1993). Hewlett-Packard gains control of inventory for service through design for localization *Interfaces* **23**, 1-11.
- [35] Levi, R. Roundy, R. & Shmoys, D.B. (2015). Provably near-optimal sampling based policies for stochastic inventory control models. *Mathematics of Operations Research* **32**, 821-839.
- [36] Muckstadt, J.A. (2000). *Analysis and Algorithms for Service Part Supply Chains*. New York: Springer.
- [37] Ninh, A. & Prekopa, A. (2013). Log-concavity of compound distributions with applications in stochastic optimization *Discrete Applied Mathematics* **161**, 3017–3027.
- [38] Özer, O. & Xiong, H. (2008). Stock position and performance estimation for distribution systems with service constraints *IIE Transactions* **40**, 1141-1157.
- [39] Perera, S., Janakiraman, G. & Shun-Chen, N. (2018). Optimality of (s,S) Inventory Policies under Renewal Demand and General Cost Structure *Production and Operations Management* **27**, 368-383.
- [40] Peterson, R.E. & Silver, E.A. (2007). *Decision Systems for Inventory Control and Production Management*. New York: Wiley.
- [41] Porteus, E.L. (1971). On the optimality of generalized (s,S) policies *Management Science* **17**, 411-426.

- [42] Porteus, E.L. (2002). *Foundations of Stochastic Inventory Theory*. Stanford, CA: Stanford University Press
- [43] Rosling, K. (2002). Inventory cost rate functions with nonlinear shortage costs *Operations Research* **50**, 797-809. (Appendix with the proofs available at <http://www.informs.org/content/download/16442/184983/file/Rosling.pdf>).
- [44] Ross, S.M. (1983). *Stochastic Processes*. New York: Wiley.
- [45] Ross, S.M., Shanthikumar, J.G. & Zhu, Z. (2005). On increasing-failure-rate random variables. *Journal of Applied Probability* **42**, 797–809.
- [46] Sahin, I. (1979). On the stationary analysis of continuous review (s,S) inventory systems with constant lead times. *Operations Research* **27**, 717-729.
- [47] Sahin, I. (1983). On the continuous-review (s,S) inventory model under compound renewal demand and random lead times. *Journal of Applied Probability* **1**, 213-219.
- [48] Samworth, R.J. (2018). Recent Progress in Log-Concave Density Estimation. *Statistical Science* **4**, 493–509.
- [49] Scarf, H. (1958). A min-max solution of an inventory problem *Mathematical methods in the social sciences* K. A. Arrow, ed. Stanford University Press, 201-209.
- [50] Scarf, H. (1960). The optimality of (s,S) policies in the dynamic inventory problem. *Studies in the Mathematical Theory of Inventory and Production* K. A. Arrow, S. Karlin, P. Suppes, eds. Stanford University Press.
- [51] See, C.T. & Sim, M. (2010). Robust approximation to multiperiod inventory management. *Operations Research* **58**, 583-594.
- [52] Sengupta, D. & Nanda, A.K. (1999). Log-concave and concave distributions in reliability. *Naval Research and Logistics* **46**, 419-433.
- [53] Shaked, M. & Shanthikumar, J.G. (2007). *Stochastic Orders*. New York: Springer.
- [54] Shanthikumar, J.G. (1988). DFR property of first-passage times and its preservation under geometric compounding. *Annals of Probability* **16**, 397–406.
- [55] Shaoxiang, C. (2004). The infinite horizon periodic review problem with setup costs and capacity constraints: a partial characterization of the optimal policy. *Operations Research* **52**, 409–421.
- [56] Shaoxiang, C. & Lambrecht, M. (1996). $X - y$ band and modified (s,S) policy. *Operations Research* **44**, 1013–1019.
- [57] Shapiro, A.D. Dentcheva, D. & Ruszczyński, A. (2009). *Lectures on Stochastic Programming: Modeling and Theory*. Philadelphia: SIAM.
- [58] Shi, J., Katehakis, M. & Melamed, B. (2013). Martingale methods for pricing inventory penalties under continuous replenishment and compound renewal demands. *Annals of Operations Research* **208**, 593–612.

- [59] Shi, J., Katehakis, M., Melamed, B. & Xia, Y. (2014). Production-inventory systems with lost sales and compound Poisson demands. *Operations Research* **62**, 1048–1063.
- [60] Taylor, H.M. (1975). Optimal replacement under additive damage and other failure models *Naval Research and Logistics* **22**, 1-18.
- [61] Tijms, H. (1972). *Analysis of (s,S) inventory models..* Mathematical Centre Tracts, Amsterdam: Mathematisch Centrum.
- [62] Tijms, H. (1986). *Stochastic Modelling and Analysis: A Computational Approach.* New York: Wiley.
- [63] Veinott, A. Jr (1965). The optimal inventory policy for batch ordering. *Operations Research* **13**, 424-432.
- [64] Veinott, A. Jr (1966a). On the optimality of (s,S) inventory policies: new conditions and a new proof *SIAM Journal of Applied Mathematics* **14**, 1067-1083.
- [65] Veinott, A. Jr (1966b). The status of mathematical inventory theory *Management Science* **12**, 745-777.
- [66] Vigil, R. (2017) Suppliers, Don't Get Burned by Tight Delivery Windows & Steep Penalties at Big-Box Retailers (Available at <https://www.loaddelivered.com/blog/suppliers-dont-get-burned-by-tight-delivery-windows-steep-penalties-at-big-box-retailers>).
- [67] Wang, Y. & Yeh, Y.N. (2006). Log-concavity and LC-positivity *Journal of Combinatorial Theory. Series A* **114**, 195–210.
- [68] Wei, L., Jasin, S. & Xin, L. (2017). Asymptotic optimality of Order-Up-to-Control for stochastic inventory systems with sequential probabilistic service level constraints *University of Michigan, working paper*.
- [69] Yu, Y. (2009). On the entropy of compound distributions on non-negative integers *IEEE Transactions on Information Theory* **55**, 3645-3650.
- [70] Yu, Y. (2011). Concave renewal functions do not imply DFR interrenewal times. *Journal of Applied Probability* **2**, 583-588.
- [71] Zheng, Y. (1991). A simple proof of the optimality of (s,S) policies in infinite-horizon inventory systems *Journal of Applied Probability* **28**, 802-806.
- [72] Zheng, Y. & Federgruen A (1991). Finding optimal (s,S) policies is about as simple as evaluating a single policy *Operations Research* **39**, 654-665.
- [73] Zipkin, P. (1980). Simple ranking methods for allocation of one resource *Management Science* **26**, 34-43.
- [74] Zipkin, P. (1986). Stochastic leadtimes in continuous-time inventory models *Naval Research and Logistics* **33**, 763-774.
- [75] Zipkin, P. (2000). *Foundations on Inventory Management.* Boston: McGraw-Hill.
- [76] Zuckerman, D. (1977). Replacement models under additive damage *Naval Research and Logistics* **24**, 549-558.

Appendix: Periodic review and Periodic Demand model

In this Appendix, we discuss an important and special case of the model developed in Section 4.1. This most frequently studied and used single-item inventory model with stochastic demands was introduced by the seminal paper of Scarf [50], itself a generalization of the pioneering model in Arrow, Karlin and Scarf [5]; see also Veinott [65]. The planning horizon is divided into discrete time periods. At the start of each period, the inventory level is assessed and an order may be planned with an external or internal supplier. Demands in consecutive periods are i.i.d. with a given continuous cdf $F(\cdot)$ and probability density function $f(\cdot)$ (The case where the demand size distribution is discrete or mixed, can be handled with minor adaptations). Orders arrive after an integer valued leadtime of L periods. The leadtime may be deterministic or stochastic; in the latter case, we assume again that the leadtimes are generated by a so-called sequential exogenous process. Recall that under such processes consecutive orders are guaranteed not to cross. Stockouts are fully backlogged.

The model arises as a special case of the general model of Section 4.1, with deterministic interdemand times. In terms of the cost structure, assume each unit of inventory (backlog) incurs a cost of h (p) per unit of time. In addition, there is a fixed cost π for each unit of demand which is backlogged, irrespective of how long the unit is backlogged, and a fixed cost b for each period that ends with a shortage, irrespective of the size of the backlog.

Thus, let $(Z_i)_{i=1,2,\dots}$ denote the sequence of i.i.d. demands with common cdf Ψ and mean μ .

In addition, let:

$\mathcal{I}(y)$ = the expected inventory level at the end of some period $(t + L)$ if the starting inventory position (after ordering), at the beginning of period t , equals y units.

$B_p(y)$ = the expected backlog size at the end of some period $(t + L)$, if the starting inventory position (after ordering), at the beginning of period t , equals y units.

$B_\pi(y)$ = the expected new backlogs incurred at the end of some period $(t + L)$ if the starting inventory position (after ordering), at the beginning of period t , equals y units.

$B_b(y)$ = the expected likelihood of a shortage at the end of some period $(t + L)$ if the starting inventory position (after ordering), at the beginning of period t , equals y units.

Rosling [43] derived the following expressions for these quantities:

$$\mathcal{I}(y) = [y - (E[L] + 1)\mu] + \int_y^\infty [1 - \Psi_{L+1}(u)]du, \quad (35)$$

$$B_p(y) = \int_y^\infty [1 - \Psi_{L+1}(u)]du, \quad (36)$$

$$B_\pi(y) = \int_y^\infty [\Psi_L(u) - \Psi_{L+1}(u)]du. \quad (37)$$

$$B_b(y) = 1 - \Psi_{L+1}(y). \quad (38)$$

while

$$G(y) = h\mathcal{I}(y) + pB_p(y) + \pi B_\pi(y) + bB_b(y), \quad (39)$$

Theorem 5. (Optimality of (s,S) policies in the periodic review, periodic demand model) *An (s,S) policy is optimal under the cost structure Eq. (39) if,*

- (a) *(See Proposition 2-1(i) Rosling [43]) L is deterministic and the demand size variable Z has a log-concave cdf Ψ*
- (b) *L is stochastic with a log-concave discrete pmf, while the demand size variable Z has a concave cdf (i.e. decreasing density function)*

Proof. Part (b): it follows from Theorem 1 that both $\Psi_L(\cdot)$ and $\Psi_{L+1}(\cdot)$ are log-concave. The remainder of the proof is identical to that of Proposition 2-1 in Rosling [43]. \square

Remark 9. *Under stochastic leadtimes, part (b) of Theorem 5 extends the most general optimality condition, known to date. Badía and Sangüesa [7] established optimality when the lead time distribution is a convolution of geometric, binomial or Poisson distributions (all of which have a log-concave pmf, as mentioned above).*

Remark 10. *For part (a) of Theorem 5, Rosling identifies some, even weaker, conditions than simple log-concavity of the cdf $\Psi(\cdot)$. In the special case where $b = 0$, Theorem 1.2. (b) in Badía and Sangüesa [7] shows $G(\cdot)$ is quasiconvex, hence an optimal (s,S) policy exists, if the leadtime has a log-concave pmf and $\Psi_n(x)/\Psi_{n+1}(x)$ is non-increasing for all $n = 1, 2, \dots$. This property is referred to as the Monotone Convolution Ratio (MCR) property for the cdf $\Psi(\cdot)$, see Ross [44] and Chen and Zheng [14]. Proposition 1-2(ii) in Rosling [43] shows that this property is even weaker than log-concavity of the cdf $\Psi(\cdot)$.*