

# Optimality gap of constant-order policies decays exponentially in the lead time for lost sales models

**Linwei Xin**

Department of Industrial and Enterprise Systems Engineering, University of Illinois at Urbana-Champaign, Urbana, IL 61801

**David A. Goldberg**

School of Industrial & Systems Engineering, Georgia Institute of Technology, Atlanta, GA 30332

Inventory models with lost sales and large lead times have traditionally been considered intractable due to the curse of dimensionality. Recently, Goldberg and co-authors laid the foundations for a new approach to solving these models, by proving that as the lead time grows large, a simple constant-order policy is asymptotically optimal. However, the bounds proven there require the lead time to be very large before the constant-order policy becomes effective, in contrast to the good numerical performance demonstrated by Zipkin even for small lead time values. In this work, we prove that for the infinite-horizon variant of the same lost sales problem, the optimality gap of the same constant-order policy actually converges *exponentially fast* to zero, with the optimality gap decaying to zero at least as fast as the exponential rate of convergence of the expected waiting time in a related single-server queue to its steady-state value. We also derive simple and explicit bounds for the optimality gap, and demonstrate good numerical performance across a wide range of parameter values for the special case of exponentially distributed demand. Our main proof technique combines convexity arguments with ideas from queueing theory.

*Key words:* inventory, lost-sales, constant-order policy, lead time, convexity, queueing theory.

---

## 1. Introduction

It is well-known that there is a fundamental dichotomy in the theory of inventory models, depending on the fate of unmet demand. If unmet demand remains in the system and can be met at a later time, we say the system exhibits *backlogged demand*; if unmet demand is lost to the system, we say the system exhibits *lost sales*. Which of these assumptions is appropriate depends heavily on the application of interest. For example, in many retail applications demand can be met by a competing supplier, making lost sales a more appropriate assumption (cf. Bijvank and Vis (2011)).

A second important feature of many inventory models, intimately related to the above dichotomy, is that of positive lead times, i.e. settings in which there is a multi-period delay between when an order for more inventory is placed and when that order is received. This feature leads to an enlarged state-space (growing linearly with the lead time), to track all orders already placed but not yet received, i.e. the pipeline vector. It is a classical result that models with backlogged demand remain tractable even in the presence of positive lead times. Namely, it can be proven that a so-called base-stock (i.e. order-up-to) policy, based only on the total inventory position (i.e. sum of the current inventory and all orders in the pipeline vector), is optimal in this setting (cf. Scarf (1960), Iglehart (1963), Veinott (1966)). However, it is known that such simple policies are no longer optimal for models with lost sales and positive lead times (cf. Karlin and Scarf (1958)). *For over fifty years, inventory models with lost sales and positive lead times were generally considered intractable, as the primary solution method (dynamic programming) suffered from the curse of dimensionality as the lead time grew.* As noted in Bijvank and Vis (2011), this has led to many researchers using models with backlogging as approximations for settings in which a lost sales assumption is more appropriate, which may lead to very suboptimal policies.

Although the optimal policy for lost-sales models with positive lead times remains poorly understood, the model has been studied now for over fifty years, and we refer to Bijvank and Vis (2011) and Goldberg et al. (2015) for a comprehensive review. Due to the difficulty of computing the optimal policy, there has been considerable focus on understanding structural properties of an optimal policy, and analyzing heuristics. In particular, convexity results were obtained in Karlin and Scarf (1958), Morton (1969), and Zipkin (2008a). Janakiraman, Seshadri and Shanthikumar (2007) compared the optimal costs between the backlogged and lost sales systems with identical problem parameters, and Huh et al. (2009) further proved that the base-stock policy was asymptotically optimal as the lost-sales penalty became large compared to the holding cost, with similar results also derived in Lu, Squillante and Yao (2012). Levi, Janakiraman and Nagarajan (2008) proposed the family of so called dual-balancing policies, and proved that the cost incurred by such a policy

was always within a factor of 2 of optimal. Another recent line of research yields efficient approximation algorithms *for any fixed lead time* by carefully approximating the associated dynamic programs (cf. Halman et al. (2009), Halman, Orlin and Simchi-Levi (2012), Chen, Dawande and Janakiraman (2014)). *Despite this progress, the aforementioned work leaves open the problem of deriving efficient algorithms with arbitrarily small error, when the lead time is large.*

A very simple and natural policy, which will be the subject of our own investigations, is the so-called *constant-order policy*, which places the same order in every period, independent of the state of the system. Perhaps surprisingly, Reiman (2004) proved that for lost sales inventory models with positive lead times, sometimes the best constant-order policy outperforms the more sophisticated base-stock policy, and performed a detailed analysis under a certain asymptotic scaling. This phenomena was further illuminated by the computational study of Zipkin (2008b), which confirmed that in several scenarios the constant-order policy performed favorably. These observations were recently given a solid theoretical foundation by Goldberg et al. (2015), who proved that for lost sales inventory models with positive lead times, as the lead time grows with all other parameters remaining fixed, the best constant-order policy is in fact *asymptotically optimal*. This is quite surprising, as the policy is so simple, and performs nearly optimally exactly in the setting which had stumped researchers for over fifty years. However, the bounds proven there are impractical, requiring the lead time to be very large before the constant-order policy becomes nearly optimal, e.g. requiring a lead time which is  $\Omega(\epsilon^{-2})$  to ensure a  $(1 + \epsilon)$ -approximation guarantee, and involving a massive prefactor. The authors note that the numerical experiments of Zipkin (2008b) suggest that the constant-order policy performs quite well even for small lead times, and pose closing this gap (thus making the results practical) as an open problem. The authors also point out that if one could prove that the constant-order policy performs well for small to moderate lead times, this would open the door for the creation of practical hybrid algorithms, which solve large dynamic programs when the lead time is small, and transition to simpler algorithms for larger lead times.

## 1.1. Our contributions

In this work, we make significant progress towards solving this open problem and closing this gap. For the infinite-horizon variant of the finite-horizon problem considered by Goldberg et al. (2015), we prove that the optimality gap of the same constant-order policy actually converges *exponentially fast* to zero, i.e. we prove that a lead time which is  $O(\log(\epsilon^{-1}))$  suffices to ensure a  $(1 + \epsilon)$ -approximation guarantee. We demonstrate that the corresponding rate of exponential decay is at least as fast as the exponential rate of convergence of the expected waiting time in a related single-server queue to its steady-state value, which we prove to be monotone in the ratio of the lost-sales penalty to the holding cost. We also derive simple and explicit bounds for the optimality gap. For the special case of exponentially distributed demand, we further compute all expressions appearing in our bound in closed form, and numerically evaluate them, demonstrating good performance for a wide range of parameter values. Our main proof technique combines convexity arguments with ideas from queueing theory, and is simpler than the coupling argument of Goldberg et al. (2015).

## 1.2. Outline of paper

The rest of the paper is organized as follows. We formally define the inventory problem of interest in Section 2.1, and describe the constant-order policy in Section 2.2. We state our main results in Section 2.3, and prove our main results in Section 3. We summarize our results and propose directions for future research in Section 4. A technical appendix is provided in Section 5, which contains several proofs from throughout the paper, as well as an in-depth analysis of our bounds (both analytic and numeric) for the illustrative special case of exponentially distributed demand.

## 2. Main results

### 2.1. Model description and problem statement

In this section, we formally define our lost-sales inventory optimization problem. Let  $\{D_t, t \geq 1\}$  be a sequence of independent and identically distributed (i.i.d.) demand realizations, distributed as the non-negative random variable (r.v.)  $D$  with distribution  $\mathcal{D}$ , which we assume to have finite mean, and (to rule out certain trivial degenerate cases) to have strictly positive variance. Let  $L \geq 1$  be the deterministic lead time, and  $h, p (> 0)$  be the unit holding cost and lost-sales penalty

respectively. In addition, let  $I_t$  denote the on-hand inventory, and  $\mathbf{x}_t = (x_{1,t}, \dots, x_{L,t})$  denote the pipeline vector of orders placed but not yet delivered, at the beginning of time period  $t$ , where  $x_{i,t}$  is the order to be received in period  $i+t-1$ . The ordered sequence of events in period  $t$  is then as follows.

- A new amount of inventory  $x_{1,t}$  is delivered and added to the on-hand inventory.
- A new order is placed.
- The demand  $D_t$  is realized.
- Costs for period  $t$  are incurred, and the on-hand inventory and pipeline vector are updated.

Note that the on-hand inventory is updated according to  $I_{t+1} = \max(0, I_t + x_{1,t} - D_t)$ , and the pipeline vector is updated such that (s.t.)  $x_{1,t}$  is removed,  $x_{i,t+1}$  is set equal to  $x_{i+1,t}$  for  $i \in [1, L-1]$ , and  $x_{L,t+1}$  is set equal to the new order placed. We now formalize the family of admissible policies  $\Pi$ , which will determine the new order placed, mirroring the discussion in Huh, Janakiraman and Nagarajan (2011). An admissible policy  $\pi$  consists of a sequence of measurable maps  $\{f_t^\pi, t \geq 1\}$ , where each  $f_t^\pi$  is a deterministic measurable function with domain  $\mathcal{R}^{+,L+1}$  and range  $\mathcal{R}^+$ . In that case, for a given policy  $\pi$ , the order placed in period  $t$  equals  $f_t^\pi(\mathbf{x}_t, I_t)$ ; and  $\Pi$  denotes the family of all such admissible policies  $\pi$ . We let  $C_t \triangleq h(I_t + x_{1,t} - D_t)^+ + p(I_t + x_{1,t} - D_t)^-$  denote the sum of the holding cost and lost-sales penalty in time period  $t$ , where  $x^+ \triangleq \max(x, 0)$ ,  $x^- \triangleq \max(-x, 0)$ . Let  $\mathbf{0}(\mathbf{1})$  denote the all-zeros (all-ones) vector, where the dimension is to be inferred from context. As we will be primarily interested in minimizing long-run average costs, and for simplicity, we suppose that the initial conditions are such that the initial inventory is 0, and the initial pipeline vector is empty, i.e.  $I_1 = 0, \mathbf{x}_1 = \mathbf{0}$ . For a given policy  $\pi \in \Pi$ , we let  $C_t^\pi$  denote the corresponding cost incurred in period  $t$ , and  $C(\pi) \triangleq \limsup_{T \rightarrow \infty} \frac{\sum_{t=1}^T \mathbb{E}[C_t^\pi]}{T}$  denote the long-run average cost incurred by  $\pi$ . Then the corresponding infinite-horizon lost-sales inventory optimization problem which we will consider (identical to the problem considered in Huh, Janakiraman and Nagarajan (2011)) is

$$\text{OPT}(L) \triangleq \inf_{\pi \in \Pi} C(\pi). \tag{1}$$

For a given policy  $\pi \in \Pi$ , let  $\mathbf{x}_t^\pi(I_t^\pi)$  denote a random vector (variable) distributed as the pipeline vector (inventory) at the start of period  $t$  under policy  $\pi$ . Further suppose that  $\{(\mathbf{x}_t^\pi, I_t^\pi), t \geq 1\}$  are all constructed on a common probability space, and have the appropriate joint distribution induced by the operation of  $\pi$  over time. Recall that a policy  $\pi \in \Pi$  is said to be *stationary* if  $f_t^\pi$  is independent of  $t$ , i.e.  $f_t^\pi = f_1^\pi$  for all  $t$ , in which case we denote this common function by  $f^\pi$ . In that case, it follows immediately from the results of Huh, Janakiraman and Nagarajan (2011) and Zipkin (2008a) that for all  $L \geq 1$ , there exists a stationary optimal policy  $\pi^{*,L}$  for Problem (1) whose induced inventory level and pipeline vector is uniformly bounded over time. More formally, for all  $L \geq 1$ , there exists a stationary policy  $\pi^{*,L} \in \arg \min_{\pi \in \Pi} C(\pi)$  and finite positive constant  $M_L$  s.t. with probability 1 (w.p.1), for all  $t \geq 1$ ,

$$\sum_{i=1}^L x_{i,t}^{\pi^{*,L}} + I_t^{\pi^{*,L}} < M_L. \quad (2)$$

As a notational convenience, we denote  $\mathbf{x}_t^{\pi^{*,L}}$  by  $\mathbf{x}_t^{*,L}$ , and  $I_t^{\pi^{*,L}}$  by  $I_t^{*,L}$ .

## 2.2. Constant-order policy

In this section, we formally define the constant-order policy, and characterize the best constant-order policy. As a notational convenience, let us define all empty sums to equal zero, let  $e$  denote Euler's number,  $\log(x)$  denote the natural logarithm of  $x$ ,  $\frac{1}{\infty}$  denote 0,  $\frac{1}{0}$  denote  $\infty$ ,  $\log(\infty)$  denote  $\infty$ , and  $\mathbb{I}(A)$  denote the indicator of the event  $A$ . For any  $r \in [0, \mathbb{E}[D]]$ , the constant-order policy  $\pi_r$  is the policy that places the constant order  $r$  in every period. It is well-known (cf. Asmussen (2003)) that for any such  $r \in [0, \mathbb{E}[D]]$ ,  $\{I_t^{\pi_r}, t \geq 1\}$  converges in distribution, and in expectation, to a limiting r.v. with finite mean, which we denote by  $I_\infty^r$ .  $I_\infty^r$  has the same distribution as the steady-state waiting time in the corresponding  $GI/GI/1$  queue with interarrival distribution  $\mathcal{D}$  and processing time distribution the constant  $r$ . For two r.v.s  $X, Y$ , let  $X \sim Y$  denote equivalence in distribution between  $X$  and  $Y$ . In that case, it is well-known (cf. Asmussen (2003)) that  $I_\infty^r \sim \sup_{j \geq 0} \left( jr - \sum_{i=1}^j D_i \right)$ , and that for any  $r \in [0, \mathbb{E}[D]]$

$$C(\pi_r) = h\mathbb{E} \left[ \sup_{j \geq 0} \left( jr - \sum_{i=1}^j D_i \right) \right] + p\mathbb{E}[D] - pr, \quad (3)$$

independent of the lead time  $L$ . As it is easily verified from (3) that  $C(\pi_r)$  is a convex function of  $r$  on  $[0, \mathbb{E}[D]]$ , to find the best possible constant-order policy, it suffices to select the  $r$  minimizing this one-dimensional convex function over the compact set  $[0, \mathbb{E}[D]]$ . We note that the existence of at least one such optimal  $r$  follows from the well-known properties of convex optimization over a compact set, and that the set of all such optimal solutions must be bounded away from  $\mathbb{E}[D]$ , since by assumption  $h > 0$ , and  $\lim_{r \uparrow \mathbb{E}[D]} \mathbb{E}[I_\infty^r] = \infty$  (since  $D$  has strictly positive variance, cf. Asmussen (2003)). Let  $r_\infty \in \arg \min_{0 \leq r \leq \mathbb{E}[D]} C(\pi_r)$  denote the infimum of this set of optimal constant-order levels, in which case the best constant-order policy will refer to  $\pi_{r_\infty}$ .

### 2.3. Main results

In this section, we present our main results, demonstrating that the optimality gap of the best constant-order policy decays exponentially in the lead time. For  $\theta \geq 0$ , let us define

$$\phi(\theta) \triangleq \exp(\theta r_\infty) \mathbb{E}[\exp(-\theta D)] \quad , \quad \gamma \triangleq \inf_{\theta \geq 0} \phi(\theta),$$

and  $\vartheta \in \arg \min_{\theta \geq 0} \phi(\theta)$  denote the supremum of the set of minimizers of  $\phi(\theta)$ , where we define  $\vartheta$  to equal  $\infty$  if the above infimum is not actually attained. Note that  $\phi(\theta)$  is a continuous and convex function of  $\theta$  on  $(0, \infty)$ , and right-continuous function of  $\theta$  at 0. In addition, it follows from Folland (1999) Theorem 2.27 that  $\phi(\theta)$  is right-differentiable at zero, with derivative equal to  $r_\infty - \mathbb{E}[D]$ . As  $r_\infty < \mathbb{E}[D]$ , we conclude from the definition of derivative and a straightforward contradiction argument that  $\vartheta > 0$  (i.e.  $\vartheta$  is strictly positive), and  $\gamma \in [0, 1)$  (i.e.  $\gamma$  is strictly less than 1). It follows from the celebrated Cramér's Theorem, and more generally the theory of large deviations, that up to exponential order  $\mathbb{P}(kr_\infty \geq \sum_{i=1}^k D_i)$  decays like  $\gamma^k$  as  $k \rightarrow \infty$  (cf. Deuschel and Stroock (1989)). Furthermore, as we will explore in detail later in the proof of our main result,  $\gamma$  corresponds (again up to exponential order) to the rate at which the expected waiting time in an initially empty single-server queue, with inter-arrival distribution  $D$  and processing time distribution (the constant)  $r_\infty$ , converges to its steady-state value (cf. Kingman (1962)). Let  $g \triangleq \inf_{x \in \mathcal{R}} \mathbb{E}[h(x - D)^+ + p(x - D)^-] > 0$ , and  $Q$  denote the  $\frac{p}{p+h}$  quantile of the demand distribution, where we note that  $Q$  is the optimal order quantity in the single-stage newsvendor problem and  $\text{OPT}(L) \geq g$  (cf. Zipkin (2000)). Then our main result is as follows.

**THEOREM 1 (Exponential convergence of constant-order policy to optimality).** *For all  $L \geq 1$ ,*

$$\frac{C(\pi_{r_\infty})}{OPT(L)} \leq 1 + h((1-\gamma)g)^{-1} \left( \mathbb{E}[D] - r_\infty + (e\vartheta(L+1))^{-1} \right) \gamma^{L+1}. \quad (4)$$

Our results prove that for the corresponding infinite-horizon problem, the optimality gap of the constant-order policy converges exponentially fast to zero. In particular, a lead time which is  $O(\log(\epsilon^{-1}))$  suffices to ensure a  $(1+\epsilon)$ -approximation guarantee. This contrasts with the bounds of Goldberg et al. (2015), which had an inverse polynomial dependence on  $\epsilon$ . Furthermore, our explicit bounds are much tighter than those of Goldberg et al. (2015). This takes a large step towards answering several open questions posed in Goldberg et al. (2015) with regards to deriving bounds tight enough to be useful in practice. We note that since our results only hold for the infinite-horizon problem, and will use critically certain relationships which only hold in this regime, our results are not directly comparable to those of Goldberg et al. (2015), whose bounds also hold for finite-horizon problems. Closing this gap, and proving tighter bounds for the finite-horizon problem, remains an interesting open question. To further illustrate the performance of our bounds, in the technical appendix Section 5 we explicitly compute (in closed form) the right-hand side of (4) for the special case that demand is exponentially distributed, and demonstrate good numerical performance for a wide range of parameter values.

**2.3.1. Impact of the ratio  $\frac{p}{h}$  on our bounds** In this section, we discuss the dependence of our demonstrated exponential rate of convergence  $\gamma$  on the ratio of the lost-sales penalty to the holding cost. In particular, we show that  $\gamma$  is non-decreasing in  $\frac{p}{h}$ . Note that by a simple scaling argument, for any fixed demand distribution  $\mathcal{D}$ ,  $r_\infty$  (and thus  $\gamma$ ) is a function of  $\frac{p}{h}$  only, as opposed to the particular values of  $p, h$ . To make this dependence explicit, let  $\gamma(\varrho)$  denote the value of  $\gamma$  when  $\frac{p}{h} = \varrho$  (with the dependence on  $\mathcal{D}$  implicit).

**LEMMA 1.** *For any fixed demand distribution  $\mathcal{D}$ ,  $\gamma(\varrho)$  is non-decreasing in  $\varrho$ .*

We include a proof of Lemma 1 in the technical appendix Section 5. This result suggests that the optimality gap of the constant-order policy may be larger when  $\frac{p}{h}$  is large. Interestingly, this is



exactly the regime in which Huh et al. (2009) proved that order-up-to policies are nearly optimal. More formally understanding this connection remains an interesting open question.

### 3. Proof of Theorem 1

In this section, we complete the proof of our main result Theorem 1. Let  $\delta_{i,j}$  equal 1 if  $i = j$  and 0 otherwise. Our proof begins by relating the “long-run behavior” of  $\pi^{*,L}$  to a certain constant-order policy. At a high level, we will combine Lindley’s recursion and convexity with the fact that the r.v.s corresponding to (appropriately defined stationary versions of) the different components of the pipeline vector (under  $\pi^{*,L}$ ) have the same mean, which will (approximately) coincide with the constant in our constant-order policy. However, this program immediately encounters a technical problem. Namely, the natural way to analyze the “long-run behavior” of  $\pi^{*,L}$  is through the steady-state distribution of the Markov chain induced by the policy  $\pi^{*,L}$ , i.e.  $\{(\mathbf{x}_t^{*,L}, I_t^{*,L}), t \geq 1\}$ . However, it is not obvious that this steady-state exists, i.e. that  $\{(\mathbf{x}_t^{*,L}, I_t^{*,L}), t \geq 1\}$  converges weakly to a limiting stationary vector as  $t \rightarrow \infty$ . To overcome this, we first observe that we will not actually need a random vector which is truly the steady-state of the aforementioned Markov chain (which in principle may not exist), but only need to demonstrate the existence of a random vector which has several properties that we would want such a steady-state (if it existed) to have. We now show the existence of such a random vector. We note that although closely related questions have been studied in the Markov decision process (MDP) literature (cf. Arapostathis et al. (1993)), and perturbative approaches similar to the approach we take in our own proof are in general well-known (cf. Filar (2007)), to the best of our knowledge the desired result does not follow directly from any results appearing in the literature. As such, we include a proof for completeness in the technical appendix Section 5.

**THEOREM 2.** *One may construct an  $L$ -dimensional random vector  $\chi^*$ , r.v.  $\mathcal{I}^*$ , and  $\{D_i, i \geq 1\}$  on a common probability space s.t. the following are true.*

(i)  $(\chi^*, \mathcal{I}^*)$  is w.p.1 non-negative, has finite mean, and is independent of  $\{D_i, i \geq 1\}$ .

(ii)  $\mathbb{E}[\mathcal{I}^*] = \mathbb{E} \left[ \max_{j=0, \dots, L} \left( \sum_{i=1}^j (\chi_{L+1-i}^* - D_{L+1-i}) + \delta_{j,L} \mathcal{I}^* \right) \right]$ .

(iii)  $\mathbb{E}[\chi_i^*] = \mathbb{E}[\chi_1^*] \leq \mathbb{E}[D]$  for all  $i$ .

(iv)  $h\mathbb{E}[\mathcal{I}^*] + p\mathbb{E}[D_1] - p\mathbb{E}[\chi_1^*] = OPT(L)$ .

We note that in the above,  $\chi_1^*$  corresponds intuitively to (a stationary version) of the first component of the pipeline vector under  $\pi^{*,L}$ , i.e. the amount of inventory received in any single period.  $\chi_1^*$  *does not* correspond to any notion of inventory position (i.e. some aggregate measure of the total inventory in the pipeline), nor does  $\chi_i^*$  for any  $i$ . Indeed, although certain previous papers on lost sales models sometimes perform their analysis on a transformed set of variables which do correspond to inventory positions (cf. Zipkin (2008a)), we *do not* use any such transformation here.

Next, we observe that the right-hand side of Theorem 2.(ii) is a jointly convex function of  $\chi^*$  and  $\mathcal{I}^*$ , which will allow us to apply the multi-variate Jensen's inequality (cf. Dudley (2002)). More formally, for fixed  $\mathbf{d} \in \mathbb{R}^L$  and  $\chi_1, \dots, \chi_L, I \in \mathcal{R}$ , let us define  $\mathfrak{f}_{\mathbf{d}}(\chi_1, \dots, \chi_L, I) \triangleq \max_{j=0, \dots, L} \left( \sum_{i=1}^j (\chi_{L+1-i} - d_{L+1-i}) + \delta_{j,L} I \right)$ .

**Observation 1** For each fixed  $\mathbf{d} \in \mathbb{R}^L$ ,  $\mathfrak{f}_{\mathbf{d}}(\chi_1, \dots, \chi_L, I)$  is a jointly convex function of  $(\chi_1, \dots, \chi_L, I)$  over  $\mathbb{R}^{L+1}$ . Combining with Theorem 2.(i) - (iii), the multi-variate Jensen's inequality, and the i.i.d. property of  $\{D_i, i \geq 1\}$ , we conclude that

$$\mathbb{E}[\mathcal{I}^*] \geq \mathbb{E} \left[ \max_{j=0, \dots, L} \left( j\mathbb{E}[\chi_1^*] - \sum_{i=1}^j D_i + \delta_{j,L} \mathbb{E}[\mathcal{I}^*] \right) \right].$$

We note that the observation of such a convexity in terms of the on-hand inventory and pipeline vector is not new. Indeed, the so called  $L$ -natural-convexity of the relevant cost-to-functions has been studied extensively (cf. Karlin and Scarf (1958), Morton (1969), Zipkin (2008a), Chen, Dawande and Janakiraman (2014)), and used to obtain both structural results and algorithms. In contrast, here we use convexity to relate the expected inventory under an optimal policy to the expected inventory under a particular constant-order policy, intuitively that which orders  $\mathbb{E}[\chi_1^*]$  in every period. Very similar ideas and arguments have appeared previously in the queueing theory literature, to demonstrate the extremality (with regards to expected waiting times) of certain queueing systems with constant service (or inter-arrival) times (cf. Humblet (1982), Hajek (1983)).

Before proceeding, let us define several additional notations. In particular, for  $r \in [0, \mathbb{E}[D]]$  and  $L \geq 1$ , let

$$I_L^r \triangleq \max_{j=0, \dots, L} \left( jr - \sum_{i=1}^j D_i \right), \quad C_L(r) \triangleq h\mathbb{E}[I_L^r] + p\mathbb{E}[D] - pr,$$

and  $r_L \in \arg \min_{0 \leq r \leq \mathbb{E}[D]} C_L(r)$  denote the infimum of the set of minimizers of  $C_L(r)$ . We note that  $I_L^r$  is distributed as the waiting time of the  $L$ -th customer in the corresponding  $GI/GI/1$  queue (initially empty) with interarrival distribution  $D$  and processing time  $r$ . We also note that for  $r \in [0, \mathbb{E}[D]]$ ,  $I_\infty^r$  is the weak limit, as  $L \rightarrow \infty$ , of  $I_L^r$ . Similarly,  $C(\pi_r) = \lim_{L \rightarrow \infty} C_L(r)$ , and  $C_L(r)$  is monotone increasing in  $L$ .

We now combine Theorem 2 with Observation 1 to bound the optimality gap of the constant-order policy.

LEMMA 2.  $OPT(L) \geq C_L(r_L)$ , and

$$C(\pi_{r_\infty}) - OPT(L) \leq h(\mathbb{E}[I_\infty^{r_\infty}] - \mathbb{E}[I_L^{r_\infty}]) + h(\mathbb{E}[I_\infty^{r_\infty}] - \mathbb{E}[I_L^{r_L}]) - p(r_\infty - r_L). \quad (5)$$

*Proof* Combining Observation 1 with the nonnegativity of  $\mathbb{E}[Z^*]$ , we conclude that  $\mathbb{E}[Z^*] \geq \mathbb{E}[I_L^{\mathbb{E}[\chi_1^*]}]$ . Thus by Theorem 2.(iv),  $OPT(L) \geq h\mathbb{E}[I_L^{\mathbb{E}[\chi_1^*]}] + p\mathbb{E}[D] - p\mathbb{E}[\chi_1^*]$ . Combining with Theorem 2.(iii) and the definition of  $r_L$ , we conclude that  $OPT(L) \geq C_L(r_L)$ . It then follows from (3) that

$$\begin{aligned} C(\pi_{r_\infty}) - OPT(L) &\leq (h\mathbb{E}[I_\infty^{r_\infty}] + p\mathbb{E}[D] - pr_\infty) - (h\mathbb{E}[I_L^{r_L}] + p\mathbb{E}[D] - pr_L) \\ &= h(\mathbb{E}[I_\infty^{r_\infty}] - \mathbb{E}[I_L^{r_\infty}]) + h(\mathbb{E}[I_\infty^{r_\infty}] - \mathbb{E}[I_L^{r_L}]) - p(r_\infty - r_L), \end{aligned}$$

completing the proof.  $\square$

We proceed by bounding the terms appearing in the right-hand side of (5) separately. We begin by recalling a classical result of Kingman (1962), which uses the celebrated Spitzer's identity to bound the difference between the expected waiting time of the  $L$ th job to arrive to a single-server queue (initially empty), and the steady-state expected waiting time. As this difference is exactly  $\mathbb{E}[I_\infty^r] - \mathbb{E}[I_L^r]$ , the result will allow us to bound the relevant term of (5). We state Kingman's results as customized to our own setting, notations, and assumptions.

LEMMA 3 (**Theorems 1 and 4 of Kingman (1962)**). For all  $r \in [0, \mathbb{E}[D]]$  and  $L \geq 1$ ,

$$\mathbb{E}[I_L^r] = \sum_{n=1}^L \frac{1}{n} \mathbb{E} \left[ \left( nr - \sum_{i=1}^n D_i \right)^+ \right].$$

If in addition  $r < \mathbb{E}[D]$ , then

$$\mathbb{E}[I_\infty^r] = \sum_{n=1}^{\infty} \frac{1}{n} \mathbb{E} \left[ \left( nr - \sum_{i=1}^n D_i \right)^+ \right].$$

Also,

$$\mathbb{E}[I_\infty^r] - \mathbb{E}[I_L^r] \leq ((1 - \gamma)e\vartheta(L + 1))^{-1} \gamma^{L+1}.$$

To bound the remaining term  $h(\mathbb{E}[I_L^{r_\infty}] - \mathbb{E}[I_L^{r_L}]) - p(r_\infty - r_L)$ , we begin by proving that  $r_\infty \leq r_L$  for all  $L$ . This makes sense at an intuitive level, since  $r_L$  is minimizing a function which “penalizes less” for carrying inventory. However, in spite of this clear intuition, the analysis is not entirely trivial, as one function dominating another does not necessarily imply a similar comparison of the appropriate minimizers, and we defer the proof to the technical appendix Section 5.

LEMMA 4.  $r_\infty \leq r_L$  for all  $L \geq 1$ .

Before proceeding, we also derive a certain critical inequality, which we will use to show that the term  $h(\mathbb{E}[I_L^{r_\infty}] - \mathbb{E}[I_L^{r_L}])$  and the term  $p(r_\infty - r_L)$  essentially “cancel out”. This inequality follows from the first-order optimality conditions of the convex optimization problem associated with  $r_\infty$ , but requires some care, as the relevant functions are potentially non-differentiable, and the desired statement in principle involves an interchange of expectation and differentiation. We again defer the proof to the technical appendix Section 5.

LEMMA 5.  $\sum_{n=1}^{\infty} \mathbb{P}(nr_\infty \geq \sum_{i=1}^n D_i) \geq \frac{p}{h}$ .

With Lemmas 3, 4, and 5 in hand, we now complete the proof of our main result, Theorem 1.

*Proof of Theorem 1* Recall that the remaining term on the right-hand side of (5) which we are yet to bound is

$$h(\mathbb{E}[I_L^{r_\infty}] - \mathbb{E}[I_L^{r_L}]) - p(r_\infty - r_L). \quad (6)$$

We first bound  $\mathbb{E}[I_L^{r_\infty}] - \mathbb{E}[I_L^{r_L}]$ , which by Lemma 3 equals

$$\sum_{n=1}^L \frac{1}{n} \mathbb{E} \left[ \left( nr_\infty - \sum_{i=1}^n D_i \right) \mathbb{I} \left( nr_\infty \geq \sum_{i=1}^n D_i \right) \right] - \sum_{n=1}^L \frac{1}{n} \mathbb{E} \left[ \left( nr_L - \sum_{i=1}^n D_i \right) \mathbb{I} \left( nr_L \geq \sum_{i=1}^n D_i \right) \right].$$

Combining with Lemma 4 (i.e. the fact that  $r_L \geq r_\infty$ ), which implies that

$$\left( nr_L - \sum_{i=1}^n D_i \right) \mathbb{I} \left( nr_L \geq \sum_{i=1}^n D_i \right) \geq \left( nr_L - \sum_{i=1}^n D_i \right) \mathbb{I} \left( nr_\infty \geq \sum_{i=1}^n D_i \right),$$

we conclude that  $\mathbb{E}[I_L^{r_\infty}] - \mathbb{E}[I_L^{r_L}]$  is at most

$$\sum_{n=1}^L \frac{1}{n} \mathbb{E} \left[ \left( nr_\infty - \sum_{i=1}^n D_i \right) \mathbb{I} \left( nr_\infty \geq \sum_{i=1}^n D_i \right) \right] - \sum_{n=1}^L \frac{1}{n} \mathbb{E} \left[ \left( nr_L - \sum_{i=1}^n D_i \right) \mathbb{I} \left( nr_\infty \geq \sum_{i=1}^n D_i \right) \right],$$

which itself equals  $-(r_L - r_\infty) \sum_{n=1}^L \mathbb{P}(nr_\infty \geq \sum_{i=1}^n D_i)$ . It follows that (6) is at most

$$(r_L - r_\infty) \left( p - h \sum_{n=1}^L \mathbb{P} \left( nr_\infty \geq \sum_{i=1}^n D_i \right) \right). \quad (7)$$

Note that Lemma 5 implies that

$$\sum_{n=L+1}^{\infty} \mathbb{P} \left( nr_\infty \geq \sum_{i=1}^n D_i \right) \geq \frac{p}{h} - \sum_{n=1}^L \mathbb{P} \left( nr_\infty \geq \sum_{i=1}^n D_i \right). \quad (8)$$

Combining (7) and (8), we conclude that (6) is at most  $(r_L - r_\infty) h \sum_{n=L+1}^{\infty} \mathbb{P}(nr_\infty \geq \sum_{i=1}^n D_i)$ .

It follows from the well-known Chernoff's inequality (cf. Deuschel and Stroock (1989)) that  $\mathbb{P}(nr_\infty \geq \sum_{i=1}^n D_i) \leq \gamma^n$ . By summing the associated geometric series, and combining with the fact that by definition  $r_L \leq \mathbb{E}[D]$ , we conclude that (6) is at most  $h(\mathbb{E}[D] - r_\infty)(1 - \gamma)^{-1} \gamma^{L+1}$ . Combining the above bound for (6) with Lemma 3, plugging into (5), and applying the fact that  $\text{OPT}(L) \geq g$ , completes the proof.  $\square$

## 4. Conclusion

In this paper, we proved that for lost sales models with large lead times, the optimality gap of the simple constant-order policy converges *exponentially fast* to zero as the lead time grows with the other problem parameters held fixed, and derived effective explicit bounds for this optimality gap. This takes a large step towards answering several open questions of Goldberg et al. (2015), who recently proved the asymptotic optimality of the constant-order policy in this setting, but

whose bounds on the rate of convergence were impractical. We also demonstrated that the corresponding rate of exponential decay is at least as fast as the exponential rate of convergence of the expected waiting time in a certain single-server queue to its steady-state value, which we proved to be monotone in the ratio of the lost-sales penalty to the holding cost. For the special case of exponentially distributed demand, we further computed all expressions appearing in our bound in closed form, and numerically evaluated these bounds, demonstrating good performance for a wide range of parameter values.

This work leaves many interesting directions for future research. First, it would be interesting to investigate the tightness of our exponential bound, e.g. to determine whether our exponential rate captures the true exponential rate of convergence of the optimality gap of the constant-order policy. Although one can come up with pathological examples for which this is not true, e.g. discrete demand distributions with probability at least  $\frac{p}{p+h}$  at 0 (for which  $Q = 0$ ,  $\pi_0$  is optimal amongst all policies for all  $L \geq 0$ , yet  $\gamma > 0$ ), we conjecture that under mild assumptions  $\gamma$  indeed captures the true rate of convergence of the optimality gap. Second, it is an open challenge to analyze the performance of more sophisticated policies for lost sales inventory models, e.g. affine policies, order-up-policies, and dual-balancing policies. Another interesting question involves the formal construction and analysis of “hybrid” algorithms, which e.g. solve large dynamic programs when  $L$  is small and transition to using simpler policies when  $L$  is large, or use base-stock policies when  $\frac{p}{h}$  is large (relative to  $L$ ) and a constant-order policy when  $\frac{p}{h}$  is small. The analysis of such policies would likely require a more precise understanding of the complexity of computing  $r_\infty$  to any given accuracy, and a deeper understanding of how to apply the techniques of Halman, Orlin and Simchi-Levi (2012) and Chen, Dawande and Janakiraman (2014) to infinite-horizon problems. Third, it would be interesting to prove that a similar phenomena occurs (albeit possibly with a different simple policy) in more general inventory settings, e.g. models with non-i.i.d. demand, fixed ordering costs, integrality constraints, multiple suppliers and/or products, and network structure. On a final note, our results and methodology (combined with that of Goldberg et al. (2015)) provide

a fundamentally new approach to lost sales inventory models with positive lead times. We believe that our approach, combined with other recent developments in inventory theory (e.g. the efficient solution of related dynamic programs), represents a considerable step towards making these models solvable in practice. Such progress may ultimately help to free researchers from having to use backlogged demand inventory models as approximations to lost-sales inventory models, even when such an approximation is not appropriate, which has been recognized as a major problem in the inventory theory literature (cf. Bijvank and Vis (2011)).

## 5. Appendix

### 5.1. Example: exponentially distributed demand

In this section, for the special case of exponentially distributed demand, we further compute all expressions appearing in our bound in closed form, and numerically evaluate them, demonstrating good performance for a wide range of parameter values. Thus suppose demand is exponentially distributed with rate  $\lambda$ , i.e. mean  $\lambda^{-1}$ . In this case, it is well-known that for  $r \in [0, \mathbb{E}[D]]$ ,  $\mathbb{E}[I_\infty^r]$  is the expected steady-state waiting time in a corresponding  $M/D/1$  queue, and equals  $\frac{r^2\lambda}{2(1-r\lambda)}$  (cf. Haigh (2013)). For  $p, h > 0$ , let  $\tau_{p,h} \triangleq \sqrt{\frac{h}{2p+h}}$ , and  $\gamma_{p,h} \triangleq (1 - \tau_{p,h}) \exp(\tau_{p,h})$ , where it may be easily demonstrated that  $\gamma_{p,h} \in (0, 1)$ . In that case, when demand is exponentially distributed, Theorem 1 is equivalent to the following bound.

**COROLLARY 1 (Case of exponentially distributed demand).** *Suppose  $D$  is exponentially distributed with rate  $\lambda$ . Then  $C(\pi_{r_\infty}) = \lambda^{-1}(\sqrt{h(2p+h)} - h)$ , and for all  $L \geq 1$ ,*

$$\frac{C(\pi_{r_\infty})}{OPT(L)} \leq 1 + \left( \tau_{p,h} + (\tau_{p,h}^{-1} - 1)(e(L+1))^{-1} \right) \left( (1 - \gamma_{p,h}) \log(1 + ph^{-1}) \right)^{-1} \gamma_{p,h}^{L+1}. \quad (9)$$

*Proof* Suppose  $D$  is exponentially distributed with rate  $\lambda$ . It is well-known that in this case, for all  $r \in [0, \lambda^{-1}]$ ,  $\mathbb{E}[I_\infty^r] = \frac{r^2\lambda}{2(1-r\lambda)}$  (cf. Haigh (2013)). It follows from (3) that for all  $r \in [0, \lambda^{-1}]$ ,  $C(\pi_r) = h \frac{r^2\lambda}{2(1-r\lambda)} + p\lambda^{-1} - pr$ , and  $\frac{d}{dr}C(\pi_r) = \frac{h}{2}((\lambda r - 1)^{-2} - 1) - p$  strictly increases from  $-p$  to  $\infty$  on  $[0, \lambda^{-1}]$ . It follows that  $r_\infty$  must be the unique solution to the equation  $\frac{d}{dr}C(\pi_r) = 0$  on  $[0, \lambda^{-1}]$ , and hence by a straightforward calculation that  $r_\infty = \lambda^{-1}(1 - \tau_{p,h})$ , and  $C(\pi_{r_\infty}) = \lambda^{-1}(\sqrt{h(2p+h)} - h)$ . As  $\mathbb{E}[\exp(-\theta D)] = \lambda(\lambda + \theta)^{-1}$  for all  $\theta \geq 0$ , we conclude that  $\phi(\theta) = \exp(\lambda^{-1}(1 - \tau_{p,h})\theta)\lambda(\lambda + \theta)^{-1}$ .

It then follows from another straightforward calculation that  $\vartheta$  equals the unique solution to  $\frac{d}{d\theta}\phi(\theta) = 0$  on  $[0, \infty)$ ,  $\vartheta = \tau_{p,h}\lambda(1 - \tau_{p,h})^{-1}$ , and  $\gamma = \gamma_{p,h} = (1 - \tau_{p,h})\exp(\tau_{p,h})$ . It is easily calculated that  $Q = \lambda^{-1}\log(1 + ph^{-1})$ , and

$$g = h \int_0^Q (Q - x)\lambda \exp(-\lambda x) dx + p \int_Q^\infty (x - Q)\lambda \exp(-\lambda x) dx = h\lambda^{-1} \log(1 + ph^{-1}).$$

Combining the above with another straightforward calculation completes the proof.  $\square$

Note that (9) does not depend on  $\lambda$ , which follows from the scaling properties of the exponential distribution. We now numerically evaluate (9) under different lost-demand penalty and lead time scenarios, with the holding cost fixed to 1, and present the results in Table 1.

Evaluation of (9)	L=1	L=4	L=10	L=20	L=30	L=50	L=70	L=100
p=1/4	2.13	1.08	1.00	1.00	1.00	1.00	1.00	1.00
p=1	3.36	1.89	1.15	1.01	1.00	1.00	1.00	1.00
p=4	6.42	3.99	2.62	1.72	1.34	1.08	1.02	1.00
p=9	12.26	6.77	4.43	3.12	2.45	1.73	1.38	1.15
p=39	62.26	27.60	14.86	9.62	7.62	5.75	4.75	3.81
p=99	204.5	85.21	41.77	24.43	18.20	12.92	10.49	8.49

**Table 1** When  $h = 1$ , values of (9) under different  $p$  and  $L$ .

We note that consistent with Lemma 1, when  $\frac{p}{h}$  is small (i.e. less than or equal to 1), our bounds demonstrate an excellent performance by the constant-order policy even for lead times as small as 10. When  $\frac{p}{h}$  is moderate (i.e. less than or equal to 9), our bounds demonstrate a similarly good performance for lead times on the order of 70. Even when  $\frac{p}{h}$  is very large, our bounds still imply non-trivial performance guarantees, e.g. the constant-order policy is always within a factor of 4 of optimal when  $p = 39$  and  $L = 100$ . Note that combining (9) with our explicit evaluation of  $C(\pi_{r_\infty})$  yields tight bounds on  $\text{OPT}(L)$  whenever (9) is close to 1. For example, for  $p = \lambda = 1$ , our bounds imply that  $\text{OPT}(10) \in [.63, .73]$ .



### 5.2. Proof of Lemma 1

*Proof of Lemma 1* Suppose  $\frac{p_1}{h_1} < \frac{p_2}{h_2}$ , and let  $r_\infty^i \in \arg \min_{0 \leq r \leq \mathbb{E}[D]} (h_i \mathbb{E}[I_\infty^r] + p_i \mathbb{E}[D] - p_i r)$ ,  $i = 1, 2$ . From the respective optimality of  $r_\infty^1, r_\infty^2$ , we conclude that

$$\begin{aligned} \mathbb{E}[I_\infty^1] + \frac{p_1}{h_1} \mathbb{E}[D] - \frac{p_1}{h_1} r_\infty^1 &\leq \mathbb{E}[I_\infty^2] + \frac{p_1}{h_1} \mathbb{E}[D] - \frac{p_1}{h_1} r_\infty^2, \\ \mathbb{E}[I_\infty^2] + \frac{p_2}{h_2} \mathbb{E}[D] - \frac{p_2}{h_2} r_\infty^2 &\leq \mathbb{E}[I_\infty^1] + \frac{p_2}{h_2} \mathbb{E}[D] - \frac{p_2}{h_2} r_\infty^1. \end{aligned}$$

Summing these two inequalities together implies  $\left(\frac{p_2}{h_2} - \frac{p_1}{h_1}\right)(r_\infty^2 - r_\infty^1) \geq 0$ . It follows that  $r_\infty^2 \geq r_\infty^1$ , and for all  $\theta \geq 0$ ,  $\exp(\theta r_\infty^2) \mathbb{E}[\exp(-\theta D)] \geq \exp(\theta r_\infty^1) \mathbb{E}[\exp(-\theta D)]$ . Combining with the definition of  $\gamma$  completes the proof.  $\square$

### 5.3. Proof of Theorem 2

Before providing a formal proof, we first provide an intuitive overview, noting that our proof is very similar to many proofs appearing in the literature in which one proves that a “small perturbation” of a given MDP has additional nice properties (cf. Filar (2007)). We proceed by constructing a sequence of random vectors, one for each sufficiently small  $\epsilon > 0$ , where the random vector ultimately chosen to satisfy the requirements of the theorem will be an appropriate subsequential weak limit (as  $\epsilon \downarrow 0$ ) of these vectors. Given such an  $\epsilon > 0$ , we will pick a sufficiently large time  $T_\epsilon$  s.t. the expected performance of  $\pi^{*,L}$  up to time  $T_\epsilon$  is “close” to  $\text{OPT}(L)$ . We then construct a “modified Markov chain”, which behaves very-much like  $\{(\mathbf{x}_t^{*,L}, I_t^{*,L}), t \geq 1\}$ , but has an extra “time-accounting” dimension (i.e. it will be  $(L + 2)$ -dimensional instead of  $(L + 1)$ -dimensional). This extra dimension will allow the chain (if initialized in a special restart state) to mimic  $\pi^{*,L}$  (started from an empty pipeline and zero inventory) for exactly  $T_\epsilon$  periods, then mimic a policy which orders nothing until the dimensions corresponding to the pipeline and inventory re-enter the  $(\mathbf{0}, 0)$  state, then enter a dummy state (in which it continues to mimic a policy which orders nothing) and stay for a geometrically distributed number of periods, and finally re-enter the special restart state and repeat the process. This will yield a regenerative process, which will have finite expected regeneration time “close” to  $T_\epsilon$  since the pipeline vector and inventory level are bounded by  $M_L$

under  $\pi^{*,L}$ , and which will be aperiodic due to the geometrically distributed number of visits to the dummy state. We will thus be able to prove, using standard results from the theory of regenerative processes, that this process has a steady-state distribution. Furthermore, our construction will ensure that the corresponding random vector has approximately (made precise in terms of  $\epsilon$ ) the desired features laid out in the statement of the theorem. We will then prove that this sequence of random vectors itself has a subsequential weak limit (along a subsequence of  $\epsilon$ 's converging to 0), which will satisfy the conditions of the theorem.

*Proof of Theorem 2* Recall that  $\text{OPT}(L) \geq g > 0$ . Similarly, as  $C(\pi_0) = p$ ,  $\text{OPT}(L) \leq p$ . Given any  $\epsilon \in (0, g)$ , we now construct a random vector  $(\chi_\infty^\epsilon, \mathcal{I}_\infty^\epsilon)$  which satisfies (i) - (iii) (after replacing  $\chi^*$  and  $\mathcal{I}^*$  by  $\chi_\infty^\epsilon$  and  $\mathcal{I}_\infty^\epsilon$  in all relevant statements and expressions), and for which

$$|h\mathbb{E}[\mathcal{I}_\infty^\epsilon] + p\mathbb{E}[D_1] - p\mathbb{E}[\chi_{1,\infty}^\epsilon] - \text{OPT}(L)| < \epsilon, \quad (10)$$

i.e. (iv) is approximately satisfied. It follows from the definition of limsup that there exists  $T_\epsilon > 2(L + \frac{M_L}{\mathbb{E}[D]} + 2)(hM_L + (\mathbb{E}[D] + 1)p)\epsilon^{-1}$  s.t.  $(\text{OPT}(L) - \frac{\epsilon}{2})T_\epsilon < \sum_{t=1}^{T_\epsilon} \mathbb{E}[C_t^{\pi^{*,L}}] < (\text{OPT}(L) + \frac{\epsilon}{2})T_\epsilon$ . We now construct an  $(L+2)$ -dimensional non-negative discrete-time Markov process  $\{\mathbf{Y}_t^\epsilon, t \geq 1\} = \{(\chi_t^\epsilon, \mathcal{I}_t^\epsilon, \tau_t^\epsilon), t \geq 1\}$ , where  $\chi_t^\epsilon = (\chi_{1,t}^\epsilon, \dots, \chi_{L,t}^\epsilon)$  is an  $L$ -dimensional random vector, and  $\mathcal{I}_t^\epsilon$  and  $\tau_t^\epsilon$  are both 1-dimensional r.v.s. Let  $\{B_t, t \geq 1\}$  denote an i.i.d. sequence of Bernoulli r.v.s, each of which equals 1 w.p.  $\frac{1}{2}$  and 0 w.p.  $\frac{1}{2}$ . Then  $\{\mathbf{Y}_t^\epsilon, t \geq 1\}$  evolves as follows.  $\chi_1^\epsilon = \mathbf{0}, \mathcal{I}_1^\epsilon = 0, \tau_1^\epsilon = 1$ . For  $t \geq 1$ , the chain evolves as follows.  $\chi_{i,t+1}^\epsilon = \chi_{i+1,t}^\epsilon$  for  $i \in [1, L-1]$ , and  $\mathcal{I}_{t+1}^\epsilon = (\mathcal{I}_t^\epsilon + \chi_{1,t}^\epsilon - D_t)^+$ . If  $\tau_t^\epsilon \in [1, T_\epsilon)$ , then  $\chi_{L,t+1}^\epsilon = f^{\pi^{*,L}}(\chi_t^\epsilon, \mathcal{I}_t^\epsilon)$  and  $\tau_{t+1}^\epsilon = \tau_t^\epsilon + 1$ . If  $\tau_t^\epsilon = T_\epsilon$  and  $(\chi_t^\epsilon, \mathcal{I}_t^\epsilon) \neq (\mathbf{0}, 0)$ , then  $\chi_{L,t+1}^\epsilon = 0$  and  $\tau_{t+1}^\epsilon = T_\epsilon$ . If  $\tau_t^\epsilon = T_\epsilon$  and  $(\chi_t^\epsilon, \mathcal{I}_t^\epsilon) = (\mathbf{0}, 0)$ , then  $\chi_{L,t+1}^\epsilon = 0$  and  $\tau_{t+1}^\epsilon = 0$ . If  $\tau_t^\epsilon = 0$ , then  $\chi_{L,t+1}^\epsilon = 0$  and  $\tau_{t+1}^\epsilon = B_t$ . One may easily verify the following properties of  $\{\mathbf{Y}_t^\epsilon, t \geq 1\}$ . Let  $D'_1$  denote another r.v. distributed as  $\mathcal{D}$  (independent of  $\{D_t, t \geq 1\}$ ),  $z(x) \triangleq \mathbb{E}[h(x - D'_1)^+ + p(x - D'_1)^-]$ , and  $\hat{T}_\epsilon$  denote a r.v. distributed as the time between the chain's initial and second visit to  $(\mathbf{0}, 0, 1)$ .

(a) It follows from a straightforward induction (essentially Lindley's recursion, cf. Goldberg et al.

$$(2015)) \text{ that for all } t \geq 1, \mathcal{I}_{t+L}^\epsilon \sim \max_{j=0, \dots, L} \left( \sum_{i=1}^j (\chi_{L+1-i,t}^\epsilon - D_{t+L-i}) + \delta_{j,L} \mathcal{I}_t^\epsilon \right).$$

- (b) It follows directly from the Markov chain dynamics that for all  $t \geq 1$ ,  $\mathbb{E}[\mathcal{I}_{t+1}^\epsilon] = \mathbb{E}[(\mathcal{I}_t^\epsilon + \chi_{1,t}^\epsilon - D_t)^+]$ , and thus  $\mathbb{E}[(\mathcal{I}_t^\epsilon + \chi_{1,t}^\epsilon - D_t)^-] = \mathbb{E}[\mathcal{I}_{t+1}^\epsilon] - \mathbb{E}[\mathcal{I}_t^\epsilon] + \mathbb{E}[D_t] - \mathbb{E}[\chi_{1,t}^\epsilon]$ ;
- (c) It follows directly from the Markov chain dynamics that for all  $t \geq 1$ ,  $\chi_{i,t+1}^\epsilon \sim \chi_{i+1,t}^\epsilon$  for  $i \in [1, L-1]$ .
- (d) Conditional on the event  $\{\mathbf{Y}_t^\epsilon = (\mathbf{0}, 0, 1)\}$ , the joint distribution of  $\{\mathbf{Y}_i^\epsilon, i \in [t, t + T_\epsilon - 1]\}$  is identical to that of  $\{(\mathbf{x}_i^{*,L}, I_i^{*,L}, i), i \in [1, T_\epsilon]\}$ .
- (e) Conditional on the event  $\{\tau_t^\epsilon = T_\epsilon\}$ , the expected number of time steps until the Markov chain next enters the state  $(\mathbf{0}, 0, 0)$  is at most  $L + \frac{M_L}{\mathbb{E}[D]}$ . We note that this follows from (2), combined with a straightforward application of Wald's identity and the fact that ordering nothing for  $L$  periods clears the pipeline vector.
- (f) Upon entering state  $(\mathbf{0}, 0, 0)$ , the Markov chain remains in that state for a geometrically distributed number of time steps (with mean 2), and then transitions to state  $(\mathbf{0}, 0, 1)$ .
- (g) W.p.1,  $\hat{T}_\epsilon \geq T_\epsilon$ , and  $\mathbb{E}[\hat{T}_\epsilon] - T_\epsilon \leq L + \frac{M_L}{\mathbb{E}[D]} + 2$ .
- (h)  $0 \leq \mathbb{E}[\sum_{t=1}^{\hat{T}_\epsilon} z(\chi_{1,t}^\epsilon + \mathcal{I}_t^\epsilon)] - \mathbb{E}[\sum_{t=1}^{T_\epsilon} z(\chi_{1,t}^\epsilon + \mathcal{I}_t^\epsilon)] \leq (L + \frac{M_L}{\mathbb{E}[D]} + 2)(hM_L + p\mathbb{E}[D])$ .

Combining (d) - (f) with the basic definitions associated with the theory of regenerative processes (here we refer the interested reader to Asmussen (2003) for an excellent overview), we conclude that  $\{\mathbf{Y}_t^\epsilon, t \geq 1\}$  is a discrete-time aperiodic regenerative process, with regeneration points coinciding with visits to  $(\mathbf{0}, 0, 1)$ . It thus follows from standard results in the theory of regenerative processes (cf. Asmussen (2003)) that  $\{\mathbf{Y}_t^\epsilon, t \geq 1\}$  converges weakly (as  $t \rightarrow \infty$ ) to a limiting random vector  $\mathbf{Y}_\infty^\epsilon = (\chi_\infty^\epsilon, \mathcal{I}_\infty^\epsilon, \tau_\infty^\epsilon)$ . That  $(\chi_\infty^\epsilon, \mathcal{I}_\infty^\epsilon)$  satisfies (i) - (iii) then follows from (a) - (c). We now prove that  $(\chi_\infty^\epsilon, \mathcal{I}_\infty^\epsilon)$  also satisfies (10). Again applying the standard theory of regenerative processes (cf. Asmussen (2003)), we conclude that  $\mathbb{E}[z(\chi_{1,\infty}^\epsilon + \mathcal{I}_\infty^\epsilon)] = \frac{\mathbb{E}[\sum_{t=1}^{\hat{T}_\epsilon} z(\chi_{1,t}^\epsilon + \mathcal{I}_t^\epsilon)]}{\mathbb{E}[\hat{T}_\epsilon]}$ . Combining with (a) - (h), the definition of  $T_\epsilon$ , and some straightforward algebra completes the proof of (10), and we omit the details. To complete the proof of the theorem, we observe that the sequence of probability measures corresponding to  $\{(\chi_\infty^{\frac{1}{n}}, \mathcal{I}_\infty^{\frac{1}{n}}), n \geq \lceil g^{-1} \rceil\}$  is tight since all associated random vectors have all components bounded in absolute value w.p.1 by  $M_L$ , where we refer the interested reader to

Billingsley (2013) for a review of tightness and related notions. Thus this sequence of measures has at least one subsequence which converges weakly to some subsequential weak limit  $(\chi^*, \mathcal{I}^*)$ , which by the already proven properties of  $(\chi_\infty^{\frac{1}{n}}, \mathcal{I}_\infty^{\frac{1}{n}})$  and definition of weak limit will satisfy all required conditions of the theorem.  $\square$ .

#### 5.4. Proof of Lemma 4

*Proof of Lemma 4* Suppose for contradiction that there exists  $L \in [1, \infty)$  s.t.  $r_L < r_\infty$ . Note that in this case, both  $r_L, r_\infty < \mathbb{E}[D]$ , and thus by Lemma 3 both  $\mathbb{E}[I_\infty^{r_\infty}], \mathbb{E}[I_\infty^{r_L}] < \infty$ . From definitions and the associated respective optimality of  $r_L, r_\infty$ , we conclude that

$$\begin{aligned} h\mathbb{E}[I_\infty^{r_\infty}] + p\mathbb{E}[D] - pr_\infty &\leq h\mathbb{E}[I_\infty^{r_L}] + p\mathbb{E}[D] - pr_L, \\ h\mathbb{E}[I_L^{r_L}] + p\mathbb{E}[D] - pr_L &\leq h\mathbb{E}[I_L^{r_\infty}] + p\mathbb{E}[D] - pr_\infty. \end{aligned}$$

Summing these two inequalities together implies that

$$\mathbb{E}[I_\infty^{r_\infty}] + \mathbb{E}[I_L^{r_L}] \leq \mathbb{E}[I_\infty^{r_L}] + \mathbb{E}[I_L^{r_\infty}],$$

which, by Lemma 3, is equivalent to

$$\sum_{n=L+1}^{\infty} \frac{1}{n} \mathbb{E} \left[ \left( nr_\infty - \sum_{i=1}^n D_i \right)^+ \right] \leq \sum_{n=L+1}^{\infty} \frac{1}{n} \mathbb{E} \left[ \left( nr_L - \sum_{i=1}^n D_i \right)^+ \right] < \infty. \quad (11)$$

Here we note that the desired result intuitively follows from (11) and the monotonicity of the relevant functions, i.e. the fact that  $x > y$  implies  $\mathbb{E} \left[ \left( nx - \sum_{i=1}^n D_i \right)^+ \right] \geq \mathbb{E} \left[ \left( ny - \sum_{i=1}^n D_i \right)^+ \right]$ . However, we must rule out certain subtle problems that could potentially arise from the function  $\mathbb{E} \left[ \left( nr - \sum_{i=1}^n D_i \right)^+ \right]$  not being *strictly* monotonic in  $r$ , and proceed as follows. Definitions, non-negativity, and the fact that  $r_L < r_\infty$ , together imply that

$$\begin{aligned} \left( nr_\infty - \sum_{i=1}^n D_i \right)^+ &= \mathbb{I} \left( nr_\infty \geq \sum_{i=1}^n D_i \right) \left( nr_\infty - \sum_{i=1}^n D_i \right) \\ &\geq \mathbb{I} \left( nr_L \geq \sum_{i=1}^n D_i \right) \left( nr_\infty - \sum_{i=1}^n D_i \right). \end{aligned}$$

Combining with (11), we conclude that

$$\sum_{n=L+1}^{\infty} \frac{1}{n} \mathbb{E} \left[ \mathbb{I} \left( nr_L \geq \sum_{i=1}^n D_i \right) \left( nr_\infty - \sum_{i=1}^n D_i \right) \right] \leq \sum_{n=L+1}^{\infty} \frac{1}{n} \mathbb{E} \left[ \mathbb{I} \left( nr_L \geq \sum_{i=1}^n D_i \right) \left( nr_L - \sum_{i=1}^n D_i \right) \right] < \infty.$$

It follows that

$$\sum_{n=L+1}^{\infty} \frac{1}{n} \mathbb{E} \left[ \mathbb{I} \left( nr_L \geq \sum_{i=1}^n D_i \right) (nr_{\infty} - nr_L) \right] \leq 0.$$

However, since by assumption  $nr_{\infty} - nr_L > 0$ , it follows from non-negativity that

$$\sum_{n=L+1}^{\infty} \mathbb{E} \left[ \mathbb{I} \left( nr_L \geq \sum_{i=1}^n D_i \right) \right] = 0,$$

and thus  $\mathbb{P}(nr_L \geq \sum_{i=1}^n D_i) = 0$  for all  $n \geq L + 1$ . Further noting that  $\mathbb{P}(nr_L \geq \sum_{i=1}^n D_i) \geq \mathbb{P}^n(r_L \geq D_1)$ , we conclude that  $\mathbb{P}(r_L \geq D_1) = 0$ . It follows that  $\mathbb{E}[I_{\infty}^{r_L}] = \mathbb{E}[I_L^{r_L}] = 0$ , and thus by (3),  $C(\pi_{r_L}) = C_L(r_L)$ . Combining with Lemma 2, which implies that  $C_L(r_L) \leq C(\pi_{r_{\infty}})$ , and the optimality of  $r_{\infty}$ , we conclude that  $r_L \in \arg \min_{0 \leq r \leq \mathbb{E}[D]} C(\pi_r)$ . However, as  $r_{\infty}$  is by definition the infimum of  $\arg \min_{0 \leq r \leq \mathbb{E}[D]} C(\pi_r)$ , the fact that  $r_L < r_{\infty}$  thus yields a contradiction, completing the proof.  $\square$

### 5.5. Proof of Lemma 5

*Proof of Lemma 5* Since  $r_{\infty} < \mathbb{E}[D]$ , there exists  $\delta > 0$  s.t.  $r_{\infty} + \epsilon < \mathbb{E}[D]$  for all  $\epsilon \in [0, \delta]$ . Let us fix any such  $\epsilon > 0$ . The definition and associated optimality of  $r_{\infty}$  implies that  $C(\pi_{r_{\infty}}) \leq C(\pi_{r_{\infty} + \epsilon})$ . Combining with Lemma 3 and (3), we conclude that

$$h \sum_{n=1}^{\infty} \frac{1}{n} \mathbb{E} \left[ \mathbb{I} \left( nr_{\infty} \geq \sum_{i=1}^n D_i \right) \left( nr_{\infty} - \sum_{i=1}^n D_i \right) \right]$$

is at most

$$h \sum_{n=1}^{\infty} \frac{1}{n} \mathbb{E} \left[ \mathbb{I} \left( n(r_{\infty} + \epsilon) \geq \sum_{i=1}^n D_i \right) \left( n(r_{\infty} + \epsilon) - \sum_{i=1}^n D_i \right) \right] - p\epsilon.$$

Combining with the fact that

$$\mathbb{I} \left( nr_{\infty} \geq \sum_{i=1}^n D_i \right) \left( nr_{\infty} - \sum_{i=1}^n D_i \right) \geq \mathbb{I} \left( n(r_{\infty} + \epsilon) \geq \sum_{i=1}^n D_i \right) \left( nr_{\infty} - \sum_{i=1}^n D_i \right),$$

it follows that

$$h \sum_{n=1}^{\infty} \frac{1}{n} \mathbb{E} \left[ \mathbb{I} \left( n(r_{\infty} + \epsilon) \geq \sum_{i=1}^n D_i \right) \left( nr_{\infty} - \sum_{i=1}^n D_i \right) \right]$$

is at most

$$h \sum_{n=1}^{\infty} \frac{1}{n} \mathbb{E} \left[ \mathbb{I} \left( n(r_{\infty} + \epsilon) \geq \sum_{i=1}^n D_i \right) \left( n(r_{\infty} + \epsilon) - \sum_{i=1}^n D_i \right) \right] - p\epsilon.$$

Equivalently (as all relevant sums are finite)

$$\sum_{n=1}^{\infty} \mathbb{P} \left( n(r_{\infty} + \epsilon) \geq \sum_{i=1}^n D_i \right) \geq \frac{p}{h}.$$

As this holds for all sufficiently small  $\epsilon$ , the only remaining step is to demonstrate validity at  $\epsilon = 0$ .

By monotonicity, for each fixed  $n$  and all  $\epsilon \in [0, \delta]$ ,

$$\mathbb{P} \left( n(r_{\infty} + \epsilon) \geq \sum_{i=1}^n D_i \right) \leq \mathbb{P} \left( n(r_{\infty} + \delta) \geq \sum_{i=1}^n D_i \right).$$

Furthermore, since  $r_{\infty} + \delta < \mathbb{E}[D]$ , for any fixed  $\nu > 0$ , there exists  $M_{\nu} < \infty$  (depending only on  $\nu, \mathcal{D}, r_{\infty}, \delta$ ) s.t.

$$\sum_{n=M_{\nu}+1}^{\infty} \mathbb{P} \left( n(r_{\infty} + \delta) \geq \sum_{i=1}^n D_i \right) \leq \nu.$$

Indeed, the above follows from a standard argument (the details of which we omit) in which each term is bounded using Chernoff's inequality, and the terms are summed as an infinite series (cf. Deuschel and Stroock (1989)). Combining the above, we conclude that for all  $\nu > 0$ , and  $\epsilon \in (0, \delta]$ ,

$$\sum_{n=1}^{M_{\nu}} \mathbb{P} \left( n(r_{\infty} + \epsilon) \geq \sum_{i=1}^n D_i \right) \geq \frac{p}{h} - \nu.$$

As  $\mathbb{P}(nx \geq \sum_{i=1}^n D_i)$  is a right-continuous function of  $x$  (by the right-continuity of cumulative distribution functions), it follows that  $\sum_{n=1}^{M_{\nu}} \mathbb{P}(nx \geq \sum_{i=1}^n D_i)$  is similarly right-continuous in  $x$ .

Right-continuity at  $\epsilon = 0$  follows, and we conclude that

$$\sum_{n=1}^{M_{\nu}} \mathbb{P} \left( nr_{\infty} \geq \sum_{i=1}^n D_i \right) \geq \frac{p}{h} - \nu.$$

As this holds for all  $\nu$ , letting  $\nu \downarrow 0$  completes the proof.  $\square$

## Acknowledgments

The authors thank Maury Bramson, Jim Dai, Bruce Hajek, Nir Halman, Tim Huh, Ganesh Janakiraman, Dmitriy A. Katz-Rogozhnikov, Retsef Levi, Yingdong Lu, Marty Reiman, Alan Scheller-Wolf, Mayank Sharma, Mark Squillante, Kai Wang, Jiawei Zhang, and Paul Zipkin for several stimulating discussions. David A. Goldberg also gratefully acknowledges support from NSF grant no. 1453929.

## References

- Arapostathis, A., V. Borkar, E. Fernandez-Gaucherand, M. Ghosh, S. Marcus. 1993. Discrete-time controlled Markov processes with average cost criterion: a survey. *SIAM Journal on Control and Optimization*. **31**(2) 282-344.
- Asmussen, S. 2003. *Applied probability and queues*. Second ed., Springer.
- Bijvank, M., I.F.A. Vis. 2011. Lost-sales inventory theory: A review. *European Journal of Operations Research*. **215** 1-13.
- Billingsley, P. 2013. *Convergence of probability measures*. John Wiley and Sons.
- Chen, W., M. Dawande, G. Janakiraman. 2014. Fixed-Dimensional Stochastic Dynamic Programs: An Approximation Scheme and an Inventory Application. *Operations Research*. **62**(1) 81-103.
- Deuschel, J-D., D. W. Stroock. 1989. *Large deviations*. Academic Press, Boston.
- Dudley, R. 2002. *Real Analysis and Probability*. Cambridge University Press, Cambridge, UK.
- Filar, J. 2007. *Controlled Markov chains, graphs and Hamiltonicity*. Now Publishers Inc.
- Folland, G. B. 1999. *Real Analysis: Modern Techniques and Their Applications*. Second Edition, Wiley.
- Goldberg, D.A., D. A. Katz-Rogozhnikov, Y. Lu, M. Sharma, M. S. Squillante. 2015. Asymptotic optimality of constant-order policies for lost sales inventory models with large lead times. Accepted to *Mathematics of Operations Research*.
- Haigh, J. 2013. *Probability Models*. Springer, London.
- Hajek, B. 1983. The proof of a folk theorem on queueing delay with applications to routing in networks. *Journal of the ACM*. **30**(4) 834-851.
- Halman, N., D. Klabjan, M. Mostagir, J. Orlin, D. Simchi-Levi. 2009. A Fully Polynomial Time Approximation Scheme for Single-Item Inventory Control with Discrete Demand. *Mathematics of Operations Research*. **34** 674-685.
- Halman, N., J. Orlin, D. Simchi-Levi. 2012. Approximating the nonlinear newsvendor and single-item stochastic lot-sizing problems when data is given by an oracle. *Operations Research*. **60**(2) 429-446.

- Huh, W.T., G. Janakiraman, J.A. Muckstadt, P. Rusmevichientong. 2009. Asymptotic optimality of order-up-to policies in lost sales inventory systems. *Management Science*. **55**(3) 404-420.
- Huh, W.T., G. Janakiraman, M. Nagarajan. 2011. Average cost single-stage inventory models: an analysis using a vanishing discount approach. *Operations Research*. **59**(1) 143-155.
- Humblet, P. A. 1982. Determinism minimizes waiting time in queues. Technical report. *MIT Laboratory for Information and Decision Systems*. LIDS-P-1207.
- Iglehart, D. L. 1963. Optimality of (s, S) policies in the infinite horizon dynamic inventory problem. *Management Science*. **9**(2) 259-267.
- Janakiraman, G., S. Seshadri, G. Shanthikumar. 2007. A comparison of the optimal costs of two canonical inventory systems. *Operations Research*. **55**(5) 866-875.
- Karlin, S., H. Scarf. 1958. Inventory models of the Arrow-Harris-Marschak type with time lag. K. Arrow, S. Karlin, H. Scarf (Eds.), *Studies in the Mathematical Theory of Inventory and Production*, Stanford University Press, Stanford, CA, 155-178.
- Kingman, J.F.C. 1962. Some inequalities for the queue GI/G/1. *Biometrika* **49** 315-324.
- Levi, R., G. Janakiraman, M. Nagarajan. 2008. A 2-approximation algorithm for stochastic inventory control models with lost-sales. *Mathematics of Operations Research*. **33**(2) 351-374.
- Lu, Y., M. Squillante, D. Yao. 2012. Matching Supply and Demand – Asymptotics and Insights. Preprint.
- Morton, T. E. 1969. Bounds on the solution of the lagged optimal inventory equation with no demand backlogging and proportional costs. *SIAM Review*. **11**(4) 572-596.
- Reiman, M.I. 2004. A new simple policy for a continuous review lost-sales inventory model. Unpublished manuscript.
- Scarf, H. 1960. The Optimality of (s, S) Policies in the Dynamic Inventory Problem. *Mathematical Methods in the Social Sciences*. Stanford University Press. 196-202.
- Veinott, Jr., A. F. 1966. On the optimality of (s, S) inventory policies: new conditions and a new proof. *SIAM Journal on Applied Mathematics*. **14** 1067-1083.
- Zipkin, Paul. 2000. Foundations of inventory management. Vol. 2. New York: McGraw-Hill.



Zipkin, Paul. 2008a. On the structure of lost-sales inventory models. *Operations Research*. **56**(4) 937-944.

Zipkin, Paul. 2008b. Old and new methods for lost-sales inventory systems. *Operations Research*. **56**(5) 1256-1263.