

Authors are encouraged to submit new papers to INFORMS journals by means of a style file template, which includes the journal title. However, use of a template does not certify that the paper has been accepted for publication in the named journal. INFORMS journal templates are for the exclusive purpose of submitting to an INFORMS journal and should not be used to distribute the papers in print or online or to submit the papers to another publication.

Asymptotic Optimality of Constant-Order Policies for Lost Sales Inventory Models with Large Lead Times

David A. Goldberg

H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology, dgoldberg9@isye.gatech.edu

Dmitriy A. Katz-Rogozhnikov

Mathematical Sciences Department, IBM Research, dkatzrog@us.ibm.com.

Yingdong Lu

Mathematical Sciences Department, IBM Research, yingdong@us.ibm.com.

Mayank Sharma

Cognitive Computing Department, IBM Research, mxsharma@us.ibm.com.

Mark S. Squillante

Mathematical Sciences Department, IBM Research, mss@us.ibm.com.

Lost sales inventory models with large lead times, which arise in many practical settings, are notoriously difficult to optimize due to the curse of dimensionality. In this paper we show that when lead times are large, a very simple constant-order policy, first studied by Reiman, performs nearly optimally. The main insight of our work is that when the lead time is very large, such a significant amount of randomness is injected into the system between when an order for more inventory is placed and when the order is received, that “being smart” algorithmically provides almost no benefit. Our main proof technique combines a novel coupling for suprema of random walks with arguments from queueing theory.

Key words: inventory; lost sales; approximation/heuristics; asymptotics; queueing theory; random walk; coupling

MSC2000 subject classification: Primary: 90B05; secondary: 90B22

OR/MS subject classification: Inventory/production, approximation/heuristics

1. Introduction. In this paper we consider a stochastic inventory control problem under the so-called single-item, periodic-review, lost-sales model with positive lead times and independent and identically distributed (i.i.d.) demand. This model is based on sales being lost whenever there is insufficient supply to fulfill demand, i.e., unfulfilled demand is lost rather than being carried over, or backlogged, to a later time. Furthermore, there is a constant delay of $L > 0$ periods (i.e., a single lead time) between when an order for additional inventory is placed and when that inventory is received. The problem then is to determine the best policy for a series of orders across a planning horizon comprised of a finite number of discrete time periods, with the goal of minimizing cost in expectation.

The cost structure of this model consists of a per-unit penalty for lost sales due to unfulfilled demand within each period and a per-unit cost for holding excess inventory within each period.

Unlike the corresponding backorder inventory control problem when unfulfilled demand is fully backlogged from period to period, where the optimal policy is well known to be an order-up-to policy, the optimal order policy for the lost-sales inventory model is not known in general, and in fact remains poorly understood (cf. Bijvank and Vis [4]).

Such periodic-review, lost-sales models have a long and rich history in the operations research, operations management and management science literature. Here we briefly review some of the most relevant literature, and refer the interested reader to the recent survey paper of Bijvank and Vis [4] for a more comprehensive exposition. This class of inventory models was first introduced by Bellman et al. [2]. Certain properties of the optimal policy were explored for the case of $L = 1$ by Karlin and Scarf [24] and by Yaspan [40], where it was shown that the order-up-to policy is not optimal for the lost-sales inventory model. Morton [31] extended this analysis to the case of general L . Other properties of the optimal policy, including various notions of convexity and monotonicity, were explored in Zipkin [44, 43] and Huh and Janakiraman [18]. With respect to computation of the optimal policy, the primary approach taken in the literature is dynamic programming, combined with various heuristics to speed up computations (cf. Morton [32], Zipkin [43]). However, since the state-space of any such dynamic program grows exponentially in the lead time, such computations become extremely challenging even for lead times less than ten (cf. Zipkin [43]). Namely, this family of techniques suffers from the curse of dimensionality as the lead time grows. Indeed, even for a lead time of four and geometrically distributed demand, Zipkin [43] reports that computing the optimal policy requires solving a dynamic program with 228,581 states. This is not surprising because the problem at hand and several closely related problems are known to be NP-complete (cf. Halman et al. [17]).

The difficulty of computing optimal policies for the lost-sales model has led to a considerable body of work on heuristics. The computational performance and properties of various algorithms, including order-up-to policies, have been analyzed by numerous authors (cf. Gaver [13], Morse [30], Morton [32], Yaspan [41], Pressman [35], Nahmias [34], Downs et al. [9], Johansen [22], Janakiraman and Roundy [21], Johansen and Thorstenson [23], Bijvank and Vis [3]). With respect to policies that have provable performance guarantees, the breakthrough work of Levi et al. [25] proved that a certain dual-balancing heuristic, inspired by previous results for other models (cf. Levi et al. [26, 27]), yields a policy whose cost is always within a factor of two of optimal. Huh et al. [19] show that in a certain scaling regime, in which the ratio of the lost-sales penalty to the holding cost asymptotically tends to infinity, an order-up-to policy is asymptotically optimal; and a similar result has been recently derived by Lu et al. [29]. Using a very different approach, Halman et al. provide an approximate dynamic programming algorithm that, combined with ideas from discrete convexity, yields a so-called fully polynomial-time approximation scheme for various related inventory control problems (cf. Halman et al. [16, 17]). These techniques were recently extended to lost-sales models with positive lead times (as considered in this paper) by Chen et al. [6], who provide a pseudo-polynomial-time additive approximation algorithm. Namely, under a suitable encoding scheme, an algorithm is presented that, for any $\epsilon > 0$, returns a policy whose performance differs additively from that of the optimal policy by at most ϵ , in time which is polynomial in ϵ^{-1} if the overall encoding length of the problem is held fixed while ϵ is varied, and otherwise is pseudo-polynomial in the overall encoding length (which grows with the lead time L); we refer the reader to Chen et al. [6] for details. In a follow-up study (Chen et al. [7]), the authors prove several interesting integrality results for these and related models.

The work closest to our own is that of Reiman [36], who studies a very simple policy for a certain continuous-review, lost-sales model with positive lead times and demand arriving as a Poisson process. In particular, the author analyzes an open-loop constant-order policy, which at time 0 selects an interval size τ and simply orders a single unit of inventory every τ time units. The author observes that this simple policy can be analyzed as a $D/M/1$ queue, and goes on to perform

an interesting asymptotic analysis, showing that for any fixed holding cost and lost-sales penalty, there exists a critical lead-time value L^* such that (s.t.): (i) for all lead times less than L^* , the best base-stock policy outperforms the best constant-order policy; and (ii) for all lead times greater than L^* , the best constant-order policy outperforms the best base-stock policy. The author makes no attempt to compare either policy to the true optimal policy, which he notes is unknown.

Of course, there is no a priori reason to believe that such a simple constant-order policy should be nearly optimal. However, numerical results from a recent study by Zipkin [43], in which the optimal policy is computed for a lost-sales model with i.i.d. demand and positive lead times (nearly identical to the model we consider, but with discounting), show that the constant-order policy (in which the same fixed constant is ordered in every time period) can perform surprisingly well. More precisely, in numerical experiments for a lead time of four, the constant-order policy always incurs an expected cost at most twice that incurred by the optimal policy; in 62.5% of the cases, the constant-order policy incurs a cost at most 1.33 times that incurred by the optimal policy; and in 38% of the cases, it incurs a cost at most 1.12 times that incurred by the optimal policy. This begs the question of how such a simple policy can perform so well on reasonable problem instances.

In the present paper we derive theoretical results that shed light on this and related phenomena. Specifically, we prove that, as the lead time grows (with the demand distribution, lost-sales penalty and holding cost remaining fixed), the best constant-order policy is in fact asymptotically optimal. We also establish explicit bounds on how large the lead time should be to ensure that the best constant-order policy incurs an expected cost of at most $1 + \epsilon$ times that incurred by the optimal policy. To the best of our knowledge, this is the first algorithm proven to be within $1 + \epsilon$ of optimal for lost-sales models when the lead time is large, and whose runtime does not grow with the lead time. The main insight of our work is that when the lead time is very large, such a significant amount of randomness is injected into the system between when an order for more inventory is placed and when the order is received, that “being smart” algorithmically provides almost no benefit. Our main proof technique combines a novel coupling for suprema of random walks with arguments from queueing theory. Since this simple policy succeeds exactly when known algorithms start running into trouble due to the curse of dimensionality, our results open the door for the creation of “hybrid” algorithms that use more elaborate forms of dynamic programming when the lead time is small, and gradually transition to less computationally intensive algorithms (with the constant-order policy at the extreme) as the lead time grows.

Outline of paper and overview of proof. The remainder of this paper, including the underlying proofs of our main results, is organized as follows. Section 2 formally defines the model of study and Section 3 states our main results (Theorem 1 and Corollary 1), namely the asymptotic optimality of the constant-order policy and associated explicit performance guarantees. In Section 4, we explicitly describe the dynamics and associated costs for a general policy over any consecutive L time periods, if one conditions on the pipeline-vector and inventory at the start of those L time periods, in terms of the maxima of various partial sums (Lemma 1). We then customize this result to the constant-order policy (Lemma 2), under which various associated expressions simplify considerably.

In Section 5, we develop lower bounds on the cost incurred by a particular optimal policy $\bar{\pi}$ described in Zipkin [44], which never orders more than a certain quantity that depends on the underlying costs and demand distribution, but not on the leadtime L nor planning horizon T . First, we formulate a lower bound on the cost incurred by $\bar{\pi}$ over any consecutive L time periods by supposing that $\bar{\pi}$ was able to choose the state of the system at the start of those L time periods to be as favorable as possible. This results in a “best-case” pipeline vector \mathbf{x}^* and inventory level \mathcal{I}^* , which can be described as the solution to an appropriate optimization problem (4), and then can be used together with Lemma 1 to derive a lower bound for the cost incurred by $\bar{\pi}$ over any consecutive L periods (Lemma 4), again in terms of the maxima of various partial sums.

The second lower bound formulated in Section 5 (Lemma 5) represents the most critical step of the entire proof, where we reason as follows. Ultimately, we wish to show that the performance of an appropriate constant-order policy nearly matches the lower bound established in Lemma 4. To accomplish this, we first note that if we wanted to select a constant-order policy which came close to matching the aforementioned lower bound, a natural approach would be to select the constant-order policy that “best mimics” the pipeline vector \mathbf{x}^* , i.e., by defining $r^* \triangleq L^{-1} \sum_{i=1}^L x_i^*$ and considering the policy that orders r^* in every period (assuming $r^* < \mathbb{E}[D]$). Second, we note that if one compares the maxima appearing in the lower bound of Lemma 4, and the maxima appearing in the dynamics of the constant-order policy when r^* is ordered in every period (as described in Lemma 2), the associated expressions are markedly similar. A fundamental difficulty in precisely comparing these expressions is that the indices at which various maxima appearing in Lemma 4 attain their suprema may depend on fluctuations in the vector \mathbf{x}^* , while the maxima appearing in Lemma 2 do not have this property, as the associated pipeline vector is constant. To remedy this, we describe an explicit coupling between the maxima that appear in Lemma 4 and the maxima that appear in Lemma 2 when r^* is ordered in every period. Indeed, instead of computing the “true values” of the maxima appearing in Lemma 4, we derive bounds by examining the associated expressions not at the index at which they attain their supremum, but at a different index, namely the index at which a corresponding expression appearing in Lemma 2 attains its supremum. In this way, we are able to derive a second lower bound (Lemma 5), which is much closer to the corresponding expression appearing in Lemma 2, and highly amenable to analysis.

Then in Section 6 we use Lemma 5, along with a careful analysis of various quantities associated with certain random walks and their suprema (Lemma 6), to explicitly bound the performance gap between the aforementioned constant-order policy and $\bar{\pi}$, over any consecutive L periods (Theorem 2). In Section 7, we use Theorem 2 to complete the proof of our main results. Along these lines, we first prove that r^* is bounded away from $\mathbb{E}[D]$ (Lemma 7), as the bounds of Theorem 2 depend sensitively on this gap. The associated proof proceeds by contradiction in which it is shown that if r^* is “too close” to $\mathbb{E}[D]$, then a certain “sharpening” of the central limit theorem, known as Stein’s lemma, guarantees that too large of a holding cost is incurred. Finally, we note that every constant-order policy is dominated by a certain “best-possible” constant-order policy. Combining all of the above with a straightforward asymptotic analysis completes the proof.

Closing remarks and directions for future research are presented in Section 8. We also include a short technical appendix in Section 9.

2. Model description and problem statement. Let us consider a standard lost-sales inventory optimization problem. One is given as input the holding cost, lost-demand penalty, planning horizon, lead time, and demand distribution. The problem then is to determine the optimal ordering policy to control inventory in the so-called single-item, periodic-review, lost-sales model.

Specifically, we consider the following lost-sales model and associated optimization problem. Let h denote the per-unit holding cost and c the per-unit lost-demand penalty, where we assume $c, h > 0$. Time is slotted, with the planning horizon and lead time respectively comprised of T and L periods s.t. $T > L > 0$. The demand in each period t , denoted by D_t , is assumed to be i.i.d. and governed by a non-negative demand distribution \mathcal{D} with finite third moment. To prevent certain degenerate situations, we further assume that \mathcal{D} is not deterministic.

At the start of each time period t there is an amount of inventory I_t available. There is also an L -dimensional vector of pipeline inventory $\mathbf{x}_t = (x_{1,t}, x_{2,t}, \dots, x_{L,t})$ that captures the orders placed before period t but not yet received prior to period t . The system dynamics for time period t then proceed as follows. First, a new amount $x_{1,t}$ of goods is added to the inventory. Second, before seeing the demand of period t , an order for more inventory is placed. After placing this order, the pipeline inventory vector is updated in a manner analogous to that of a queue: the front entry $x_{1,t}$

is removed, all other entries move up one position (i.e., $x_{i,t+1} = x_{i+1,t}$ for $i = 1, 2, \dots, L - 1$), and the new order is appended at the end, becoming $x_{L,t+1}$.

The order placed at time t must be a function (albeit possibly a random function) only of the planning horizon T , the current time t , the inventory level I_t at the start of period t , the pipeline vector \mathbf{x}_t at the start of period t , and the model primitives L, h, c, \mathcal{D} . In particular, the ordering decision at time t *cannot* depend on the realizations of future demand. We call all such policies *admissible* policies, and denote the family of admissible policies by Π .

Next, a random demand D_t is realized from \mathcal{D} . At the end of time period t (but before the start of period $t + 1$), the inventory or lost-sales costs are incurred as follows. The amount of excess inventory at the end of period t is given by $I_{t+1} = (I_t + x_{1,t} - D_t)^+$, noting that D_t is independent of $I_t + x_{1,t}$. Conversely, the amount of lost demand (due to not having enough inventory on hand) in period t is denoted by $N_t \triangleq (I_t + x_{1,t} - D_t)^-$. Then, the holding cost incurred (for storing excess inventory) equals hI_{t+1} and the lost-sales penalty incurred (for lost demand) equals cN_t , noting that at most one of I_{t+1} and N_t is positive for any period t .

The goal of the inventory planner is to minimize the expected cost incurred over the entire planning horizon. Let $\mathbf{1}$ denote the vector of all ones and $\mathbf{0}$ the vector of all zeros, where the dimension is to be inferred from context. Let us suppose $I_0 = 0$ and $\mathbf{x}_0 = \mathbf{0}$, which should be assumed throughout as the given initial conditions, and let us further define

$$C_t \triangleq hI_{t+1} + cN_t = h(I_t + x_{1,t} - D_t)^+ + c(I_t + x_{1,t} - D_t)^-.$$

The planner then wishes to find the policy $\pi \in \Pi$ that minimizes $\mathbb{E}[\sum_{t=L+1}^{T+L} C_t]$, where the expectation is over the random demand and any random decisions taken by policy π , and where we suppose that ordering decisions are made only in periods $1, \dots, T$. Note that we do not penalize a policy for any costs incurred over the first L time periods, as this cost is completely determined by the initial pipeline vector and random demands. On the other hand, we do penalize a policy for costs incurred in periods $[T, T + L]$ (recalling $T > L$), as the cost incurred during these periods is completely determined by the ordering decisions made in periods $1, \dots, T$ and random demands. We note that such a convention is consistent with the previous literature; see, e.g., Zipkin [44]. As a notational convenience and without loss of generality (w.l.o.g.), we suppose that every policy orders 0 in periods $T + 1, \dots, T + L$, as these ordering decisions have no effect on the problem's cost. For a given policy π , let $\{N_t^\pi, C_t^\pi, I_t^\pi, \mathbf{x}_t^\pi; t = 1, \dots, T + L\}$ denote the associated random variables (r.v.s) when policy π is implemented (all constructed on the same probability space). The corresponding lost-sales inventory optimization problem is then given by

$$\inf_{\pi \in \Pi} \sum_{t=L+1}^{T+L} \mathbb{E}[C_t^\pi]. \tag{1}$$

3. Main results. Our main results establish that there exists a very simple constant-order policy which is asymptotically optimal as $L \rightarrow \infty$. This section formally states these results.

3.1. Additional definitions and notations. Let D denote a r.v. governed by \mathcal{D} . Note that if the same deterministic quantity $r < \mathbb{E}[D]$ is ordered in every period, then the inventory evolves exactly as the waiting time in a $GI/GI/1$ queue (initially empty) with interarrival time distribution \mathcal{D} and processing time distribution (the constant) r ; we refer the reader to Asmussen [1] for an excellent discussion of the dynamics and steady-state properties of the $GI/GI/1$ queue. Let I_∞^r denote a r.v. distributed as the steady-state waiting time in the corresponding $GI/D/1$ queue; namely, I_∞^r is distributed as $\sup_{k \geq 0} (kr - \sum_{i=1}^k D_i)$.

For two vectors \mathbf{x}, \mathbf{y} of equal dimension, we use the notation $\mathbf{x} \leq \mathbf{y}$ to denote component-wise domination, i.e., $x_i \leq y_i$ for all i . Define Q to be the $\frac{c}{c+h}$ quantile of the demand distribution, namely

$$Q \triangleq \inf\{s \in \mathcal{R}^+ : \mathbb{P}(D > s) \leq \frac{h}{c+h}\}.$$

We note that Q is the optimal inventory level for the corresponding single-stage newsvendor problem (cf. Zipkin [42]), i.e., for any policy π and any time t

$$\mathbb{E}[C_t^\pi] \geq g \triangleq h\mathbb{E}[(Q - D)^+] + c\mathbb{E}[(D - Q)^+].$$

Lastly, let σ denote the standard deviation of D , and $\zeta \triangleq \mathbb{E}[|D - \mathbb{E}[D]|^3] \sigma^{-3}$ denote the so-called skewness of D . We then define several functions that will be instrumental for our analysis:

$$\begin{aligned} m &\triangleq \left\lceil \left(26(3\zeta + c(h\sigma)^{-1}\mathbb{E}[D] + 1) \right)^2 \right\rceil, & z &\triangleq \arg \min_{v \geq 0} (h\mathbb{E}[I_\infty^v] - cv), \\ y_1(\epsilon) &\triangleq 2^{14} (h(\mathbb{E}^2[D] + \mathbb{E}[D^2])^3 g^{-1} + 1) (Q + 2^{\frac{3}{2}}\mathbb{E}[D] + 1) (\sigma^{-6} + 1) m^3 \epsilon^{-1}, \\ y_2(\epsilon) &\triangleq \left(12\mathbb{E}[D]cg^{-1}((2ch^{-1})^{\frac{1}{2}} + 3) \right)^2 \epsilon^{-2}, \quad \text{and} \quad y(\epsilon) \triangleq \max(y_1(\epsilon), y_2(\epsilon)). \end{aligned}$$

Although the above quantities are functions of c, h, \mathcal{D} , since there will be no ambiguity, we make this dependence implicit.

Remarks.

- If $\arg \min_{v \geq 0} (h\mathbb{E}[I_\infty^v] - cv)$ is not unique, we define z to be the infimum of all such values. We will later show that z is the best constant possible if the same amount has to be ordered in every time period, in an appropriate sense. Note that $z \in [0, \mathbb{E}[D])$, since $\mathbb{E}[I_\infty^0] = 0$ and $\lim_{r \uparrow \mathbb{E}[D]} \mathbb{E}[I_\infty^r] = \infty$.

- The function $y(\epsilon)$ captures how large L should be so that our constant-order policy is within a $(1 + \epsilon)$ multiplicative factor of the optimal policy.

- Note that, for all sufficiently small ϵ , the term $y_2(\epsilon)$ will dominate $y(\epsilon)$ due to its quadratic dependence on ϵ^{-1} .

- We note that our assumptions that D has finite third moment and that the initial conditions satisfy $I_0 = 0$ and $\mathbf{x}_0 = \mathbf{0}$ are not strictly necessary. Indeed, in an earlier version of this work (i.e. Goldberg et al. [14]), similar results were presented under the assumption on D of only finite second moment, but for technical reasons requiring that D had unbounded support. Our assumptions in this paper, however, allow for a considerably simplified exposition. As such, we do not pursue the formalization of any further generalizations along these lines, instead leaving as open questions the minimal required assumptions on D for such an asymptotic optimality to hold, as well as precisely how non-existence of moments and/or non-empty initial conditions would change our explicit bounds.

3.2. Formal statement of results. For $r \in [0, \mathbb{E}[D])$, let π_r denote the policy that orders the random amount $I_\infty^r + r$ (i.e., the order amount is drawn from the distribution that governs the r.v. I_∞^r plus the constant r) in the first time period, and then orders the constant r in all subsequent time periods. We note that by ordering $I_\infty^r + r$ in the first time period, as opposed to r , the associated sequence of inventory levels becomes a stationary process, which considerably simplifies our analysis. With a slight abuse of notation, we still refer to π_r as a constant-order policy. Let $\text{OPT}(L, T)$ denote the optimal value of the lost-sales inventory optimization problem (1) for a given L and T . We then have our main theorem and an important corollary.

THEOREM 1. For all $\epsilon \in (0, 1)$, $L \geq y(\epsilon)$, and $T \geq (1 + 3\epsilon^{-1})L$,

$$\frac{\mathbb{E}[\sum_{t=L+1}^{T+L} C_t^{\pi_z}]}{OPT(L, T)} \leq 1 + \epsilon.$$

COROLLARY 1.

$$\lim_{L \rightarrow \infty} \limsup_{T \rightarrow \infty} \frac{\mathbb{E}[\sum_{t=L+1}^{T+L} C_t^{\pi_z}]}{OPT(L, T)} = 1.$$

In particular, the simple constant-order policy is asymptotically optimal as $L \rightarrow \infty$.

4. Inventory dynamics. In this section, we derive several expressions that explicitly describe the inventory dynamics for any policy over any consecutive L periods, and then customize these results to the constant-order policy for later use in our proofs.

4.1. General policy dynamics. We first explicitly characterize the cost incurred under any given policy during any consecutive L periods. Although such a characterization is generally well-known (see, e.g., Zipkin [44]), we include a proof for completeness. Let \mathcal{I} denote the initial inventory level and, for positive integers j, k , define $\delta_{j,k}$ to be 1 if $j = k$ and to be 0 otherwise.

LEMMA 1. For any policy $\pi \in \Pi$ and time $\tau \in [1, T + 1]$,

$$\begin{aligned} \mathbb{E} \left[\sum_{t=\tau}^{\tau+L-1} C_t^\pi \mid \mathbf{x}_\tau^\pi = \mathbf{x}, I_\tau^\pi = \mathcal{I} \right] &= h \sum_{k=1}^L \mathbb{E} \left[\max_{j=0, \dots, k} \left(\sum_{i=k+1-j}^k (x_i - D_{\tau+i-1}) + \delta_{j,k} \mathcal{I} \right) \right] \\ &\quad + c \left(\mathbb{E}[I_{\tau+L}^\pi \mid \mathbf{x}_\tau^\pi = \mathbf{x}, I_\tau^\pi = \mathcal{I}] - \mathcal{I} + L\mathbb{E}[D] - \sum_{i=1}^L x_i \right). \end{aligned}$$

Remark. Note that, instead of using the $\delta_{j,k}$ notation to indicate that the initial inventory is only accounted for in a single term appearing in the associated maximum, we could have, e.g., considered a transformed set of “inventory position” variables, with each associated variable corresponding to the sum of certain pipeline (and possibly inventory) variables, as was done in Zipkin [44]. However, since our subsequent analysis will rely sensitively on the indices at which various associated maxima are attained, and precisely which pipeline-vector components appear in the associated partial sums, as well as whether the associated initial inventory level \mathcal{I} is accounted for in the associated partial sum for which the maximum is attained (i.e., the maximum occurs at index k), we believe that the $\delta_{j,k}$ notation leads to greater clarity of exposition and use this notation throughout.

Proof: The result follows from a straightforward, generally well-known (see, e.g., Zipkin [44]), induction that for any $k \in [1, L]$:

$$I_{\tau+k}^\pi = \max_{j=0, \dots, k} \left(\sum_{i=k+1-j}^k (x_i - D_{\tau+i-1}) + \delta_{j,k} \mathcal{I} \right). \quad (2)$$

Note that for any times t_1, t_2 s.t. $\tau \leq t_1 < t_2 \leq \tau + L$, the net amount of demand that is met during $[t_1, t_2 - 1]$ equals $\sum_{t=t_1}^{t_2-1} D_t - \sum_{t=t_1}^{t_2-1} N_t^\pi$. It follows that

$$\sum_{t=t_1}^{t_2-1} N_t^\pi = I_{t_2}^\pi - I_{t_1}^\pi + \sum_{t=t_1}^{t_2-1} D_t - \sum_{t=t_1}^{t_2-1} x_{t-\tau+1}. \quad (3)$$

Combining (2) with (3) completes the proof. \square

4.2. Constant-order policy dynamics. We next customize Lemma 1 to the constant-order policy π_r . As previously noted, if the same deterministic quantity r is ordered in every period, the inventory evolves exactly as the waiting time in a $GI/GI/1$ queue with interarrival distribution \mathcal{D} and processing time distribution (the constant) r . Recall that for $r \in [0, \mathbb{E}[D]]$, I_∞^r denotes a r.v. distributed as the steady-state waiting time in the corresponding stable $GI/D/1$ queue, i.e., I_∞^r is distributed as $\sup_{k \geq 0} (kr - \sum_{i=1}^k D_i)$. We let i_∞^r denote the non-negative integer-valued r.v. corresponding to the largest index at which the associated random walk attains its supremum. It follows that $\{I_{L+k}^r; k \geq 1\}$ is a stationary sequence of r.v.s, with I_{L+k}^r distributed as I_∞^r for all $k \geq 1$.

Before explicitly showing how Lemma 1 simplifies when applied to π_r , it will be useful to introduce some additional notations. Let $I_{1,\infty}^r$ denote a particular replication of I_∞^r s.t. $I_{1,\infty}^r$ and $\{D_i; i \geq 1\}$ are mutually independent, where we note that $I_{1,\infty}^r$ will play the role of \mathcal{I} when Lemma 1 is applied to π_r . For $k \geq 0$, let us define

$$W_k^r \triangleq \max_{j=0,\dots,k} \left(jr - \sum_{i=1}^j D_i + \delta_{j,k} I_{1,\infty}^r \right)$$

and

$$i_k^r \triangleq \max \left\{ j^* : j^* \in [0, k], j^* r - \sum_{i=1}^{j^*} D_i + \delta_{j^*,k} I_{1,\infty}^r = W_k^r \right\}.$$

In words, i_k^r is the (largest) index at which the random walk W_k^r attains its maximum.

From Lemma 1 and the fact that $\{D_i, i \geq 1\}$ is a sequence of i.i.d. r.v.s, we obtain the following explicit characterization for the cost incurred by π_r over any L consecutive periods.

LEMMA 2. *For any $r \in [0, \mathbb{E}[D]]$ and $\tau \in [L+1, T+1]$,*

$$\mathbb{E} \left[\sum_{t=\tau}^{\tau+L-1} C_t^{\pi_r} \right] = h \sum_{k=1}^L \mathbb{E} \left[i_k^r r - \sum_{i=1}^{i_k^r} D_i + \delta_{i_k^r, k} I_{1,\infty}^r \right] + c(L\mathbb{E}[D] - Lr).$$

5. Lower bound on an optimal policy. We now derive in this section a lower bound on the cost incurred by an optimal policy during any consecutive L time periods. First, it will be convenient to review a result of Zipkin [44], which establishes an upper bound on the ordering quantities of a family of optimal policies for Problem (1).

LEMMA 3 (Zipkin [44]). *There exists an optimal policy $\bar{\pi}$ for Problem (1) that with probability (w.p.) 1 never orders more than Q , i.e., $\mathbf{x}_t^{\bar{\pi}} \leq Q\mathbf{1}$ for all $t \in [L+1, T]$.*

Remark. Note that Lemma 3 implicitly asserts the existence of at least one optimal policy for Problem (1). A priori, it could have been possible that no such policy existed, i.e., the optimal value was only attained in the limit by some sequence of policies. However, the dynamic programming formulation for Problem (1) provided in Zipkin [44], combined with the convexity and monotonicity results proven in the same paper, as well as the strict positivity of all underlying cost parameters and continuity of all relevant cost-to-go functions, indeed ensures the existence of at least one optimal policy; we refer the interested reader to Zipkin [44] for details.

We now construct a lower bound by computing the cost that $\bar{\pi}$ incurs over any L consecutive time periods, if the policy were able to choose the state of the system at the start of those L time periods to be as favorable as possible, subject only to the conditions imposed by Lemma 3 (which must hold w.p.1).

In particular, let $(\mathbf{x}^*, \mathcal{I}^*)$ denote any solution to the optimization problem

$$\min_{\mathbf{x} \in [0, Q\mathbf{1}], \mathcal{I} \in \mathbb{R}^+} \mathbb{E} \left[\sum_{t=\tau}^{\tau+L-1} C_t^{\bar{\pi}} \mid \mathbf{x}_{\tau}^{\bar{\pi}} = \mathbf{x}, I_{\tau}^{\bar{\pi}} = \mathcal{I} \right], \quad (4)$$

where the existence of $(\mathbf{x}^*, \mathcal{I}^*)$ follows from the fact that $\mathbb{E}[\sum_{t=\tau}^{\tau+L-1} C_t^{\bar{\pi}} \mid \mathbf{x}_{\tau}^{\bar{\pi}} = \mathbf{x}, I_{\tau}^{\bar{\pi}} = \mathcal{I}]$ is continuous with respect to $(\mathbf{x}, \mathcal{I})$ and goes to infinity as \mathcal{I} goes to infinity, combined with a routine compactness argument.

Note that w.l.o.g. we can take $(\mathbf{x}^*, \mathcal{I}^*)$ to be independent of the particular value of τ , and thus a function of \mathcal{D}, c, h, L only. Further note that conditional on the event $\{\mathbf{x}_1^{\bar{\pi}} = \mathbf{x}^*, I_1^{\bar{\pi}} = \mathcal{I}^*\}$, the conditional joint distribution of $\{I_{t+1}^{\bar{\pi}}, N_t^{\bar{\pi}}, C_t^{\bar{\pi}}; t = 1, \dots, L\}$ does not depend on the particular policy choices of $\bar{\pi}$, and we denote these conditional r.v.s as $\{I_{t+1}^*, N_t^*, C_t^*; t = 1, \dots, L\}$ where $I_1^* = \mathcal{I}^*$. Let

$$V_k \triangleq \max_{j=0, \dots, k} \left(\sum_{i=k+1-j}^k x_i^* - \sum_{i=1}^j D_i + \delta_{j,k} \mathcal{I}^* \right)$$

and

$$v_k^* \triangleq \max \left\{ j^* : j^* \in [0, k], \sum_{i=k+1-j^*}^k x_i^* - \sum_{i=1}^{j^*} D_i + \delta_{j^*,k} \mathcal{I}^* = V_k \right\}.$$

Then, combining Lemma 1 with the fact that $\{D_i, i \geq 1\}$ is a sequence of i.i.d. r.v.s, we derive the following lower bound for $\bar{\pi}$.

LEMMA 4. *For any $\tau \in [1, T+1]$,*

$$\mathbb{E} \left[\sum_{t=\tau}^{\tau+L-1} C_t^{\bar{\pi}} \right] \geq h \sum_{k=1}^L \mathbb{E}[V_k] + c(\mathbb{E}[I_{L+1}^*] - \mathcal{I}^* + L\mathbb{E}[D] - \sum_{i=1}^L x_i^*).$$

Ultimately, we wish to show that the performance of an appropriate constant-order policy nearly matches the lower bound established in Lemma 4. To accomplish this, we begin by making several observations. First, note that if we wanted to select a constant-order policy which came close to matching the aforementioned lower bound, a natural approach would be to select the constant-order policy which “best mimics” the pipeline vector \mathbf{x}^* , i.e., by considering the policy π_{r^*} (assuming $r^* < \mathbb{E}[D]$) and recalling that $r^* = L^{-1} \sum_{i=1}^L x_i^*$. In this case, we are left with the task of showing that the cost incurred over any consecutive L periods by policy π_{r^*} , as detailed in Lemma 2, is “close” to the lower bound established in Lemma 4.

To this end, first note that $\sum_{i=1}^L x_i^* = Lr^*$. As such, when comparing the terms appearing in Lemma 2 (applied with constant r^*) to those appearing in Lemma 4, the primary difficulty lies in comparing $\sum_{k=1}^L \mathbb{E}[V_k]$ to $\sum_{k=1}^L \mathbb{E}[i_k^{r^*} r^* - \sum_{i=1}^{i_k^{r^*}} D_i]$. We will accomplish this through a particular coupling as follows. For $k \in [1, L]$, let us construct $V_k, W_k^{r^*}$, and $i_k^{r^*}$ on the same probability space, using the same sequence of demands $\{D_i; i = 1, \dots, L\}$ (independent of $I_{1,\infty}^*$). Since the maximum of several terms is at least any one of the terms (even if selected randomly in an arbitrary manner), for $k \in [1, L]$, it follows that w.p.1

$$V_k \geq \sum_{i=k+1-i_k^{r^*}}^k x_i^* - \sum_{i=1}^{i_k^{r^*}} D_i + \delta_{i_k^{r^*},k} \mathcal{I}^*.$$

Upon combining the above with the non-negativity of $\delta_{i_k^{r^*}, k} \mathcal{I}^*$, we conclude that

$$\sum_{k=1}^L \mathbb{E} \left[i_k^{r^*} r^* - \sum_{i=1}^{i_k^{r^*}} D_i \right] - \sum_{k=1}^L \mathbb{E}[V_k] \leq r^* \sum_{k=1}^L \mathbb{E}[i_k^{r^*}] - \sum_{k=1}^L \mathbb{E} \left[\sum_{i=k+1-i_k^{r^*}}^k x_i^* \right]. \quad (5)$$

Note that the intuition behind why the above coupling works is that the indices $j \in [0, k]$ for which $\sum_{i=1}^j D_i$ is exceptionally small are good candidates for both $i_k^{r^*}$ and v_k^* . From a purely technical perspective, the coupling is convenient for two fundamental reasons. First, it eliminates all terms of the form $\sum_{i=1}^Z D_i$, where Z is a random index that may depend in a complicated way on $\{D_i; i \geq 1\}$, and is not in general adapted to the filtration generated by $\{D_i; i \geq 1\}$. Second, as we shall formalize below, the terms $r^* \sum_{k=1}^L \mathbb{E}[i_k^{r^*}]$ and $\sum_{k=1}^L \mathbb{E}[\sum_{i=k+1-i_k^{r^*}}^k x_i^*]$ can each be well-approximated (in an appropriate sense) by $\mathbb{E}[i_\infty^{r^*}] \sum_{i=1}^L x_i^*$.

Combining Lemma 4 with (5), we conclude the following refined lower bound, which will be convenient for comparing the cost incurred by policy π_{r^*} to that incurred by $\bar{\pi}$.

LEMMA 5. *For any $\tau \in [1, T + 1]$, $\mathbb{E} \left[\sum_{t=\tau}^{\tau+L-1} C_t^{\bar{\pi}} \right]$ is at least*

$$\begin{aligned} & h \left(\sum_{k=1}^L \mathbb{E} \left[i_k^{r^*} r^* - \sum_{i=1}^{i_k^{r^*}} D_i \right] + \sum_{k=1}^L \mathbb{E} \left[\sum_{i=k+1-i_k^{r^*}}^k x_i^* \right] - r^* \sum_{k=1}^L \mathbb{E}[i_k^{r^*}] \right) \\ & + c \left(\mathbb{E}[I_{L+1}^*] - \mathcal{I}^* + L\mathbb{E}[D] - \sum_{i=1}^L x_i^* \right). \end{aligned}$$

6. Difference between constant-order policy and lower bound. In this section, we combine Lemmas 2 and 5 to show that the cost incurred by π_{r^*} over any L consecutive time periods nearly matches the lower bound established in Lemma 5. We accomplish this through a term-by-term comparison of the expressions appearing in Lemmas 2 and 5.

Remark. We note that conceptually, our approach is closely related to several results in the queueing literature which prove that for certain queueing systems, the arrival process in which all inter-arrival times are the same constant is asymptotically extremal with regards to mean waiting time (cf. Hajek [15], Humblet [20]). It is an interesting open question to further quantify the connection between our approach and the convexity-type arguments typically used to prove extremality in such queueing systems (cf. Humblet [20]), which could likely be used to provide an alternate proof of our main results.

For $r \in [0, \mathbb{E}[D]]$, let us define

$$\Theta_r \triangleq \frac{(\mathbb{E}[D] - r)^2}{4(\mathbb{E}^2[D] + \mathbb{E}[D^2])}. \quad (6)$$

Then the main result of this section is formally stated as follows.

THEOREM 2. *If $r^* < \mathbb{E}[D]$, then for any $\tau \in [L + 1, T + 1]$,*

$$\mathbb{E} \left[\sum_{t=\tau}^{\tau+L-1} C_t^{\pi_{r^*}} \right] - \mathbb{E} \left[\sum_{t=\tau}^{\tau+L-1} C_t^{\bar{\pi}} \right] \quad (7)$$

is at most

$$h(Q + 2^{\frac{3}{2}} \mathbb{E}[D]) \Theta_{r^*}^{-3} + c\mathcal{I}^*.$$

Before proceeding with the proof of Theorem 2, it will be convenient to derive several bounds for the distribution of i_k^r and i_∞^r , where we defer the associated proofs to the technical appendix in Section 9.

LEMMA 6. *For any $r \in [0, \mathbb{E}[D])$ and integers $j, k \geq 0$, i_k^r has the same distribution as $\min(k, i_\infty^r)$, namely $\mathbb{P}(i_k^r = j) = \mathbb{P}(\min(k, i_\infty^r) = j)$. Furthermore,*

$$\mathbb{P}(i_\infty^r \geq k) \leq \Theta_r^{-1}(1 - \Theta_r)^k, \quad \sum_{k=0}^{\infty} \sum_{j=k}^{\infty} \mathbb{P}(i_\infty^r \geq j) \leq \Theta_r^{-3}, \quad \mathbb{E}[(I_\infty^r)^2] \leq 2\Theta_r^{-3}\mathbb{E}[D].$$

Remark. We note that a more precise analysis of the quantities in Lemma 6 would be possible using the theory of ladder heights and epochs (cf. Asmussen [1]), especially the precise results for the relevant moments given in Sparre-Andersen [38, 39], Feller [10], Loulou [28] and the recent work by Nagaev [33]. However, since the increments of the random walks that we consider have a very special structure (i.e., they are absolutely bounded from above), as well as for the sake of simplicity, we do not pursue such an analysis here.

We now complete the proof of Theorem 2.

Proof of Theorem 2: Suppose $r^* < \mathbb{E}[D]$. Combining Lemmas 2 and 5 with the definition of r^* and the non-negativity of all relevant terms, and then simplifying, allows us to conclude that (7) is at most

$$h\left(r^* \sum_{k=1}^L \mathbb{E}[i_k^{r^*}] - \sum_{k=1}^L \mathbb{E}\left[\sum_{i=k+1-i_k^{r^*}}^k x_i^*\right] + \sum_{k=1}^L \mathbb{E}[\delta_{i_k^{r^*}, k} I_{1, \infty}^{r^*}]\right) + c\mathcal{I}^*. \quad (8)$$

We proceed by bounding each term appearing in (8), beginning with

$$r^* \sum_{k=1}^L \mathbb{E}[i_k^{r^*}] - \sum_{k=1}^L \mathbb{E}\left[\sum_{i=k+1-i_k^{r^*}}^k x_i^*\right]. \quad (9)$$

First, it will be convenient to generalize our notation $\delta_{j,k}$ as follows. For an integer j and set S , define $\delta_{j,S}$ to be 1 if $j \in S$, and 0 otherwise. Observe that

$$\sum_{i=k+1-i_k^{r^*}}^k x_i^* = \sum_{i=1}^k x_i^* \delta_{i, [k+1-i_k^{r^*}, k]} = \sum_{i=1}^k x_i^* \delta_{i_k^{r^*}, [k+1-i, \infty)},$$

and thus by interchanging the order of summation and applying Lemma 6, we obtain

$$\begin{aligned} \sum_{k=1}^L \mathbb{E}\left[\sum_{i=k+1-i_k^{r^*}}^k x_i^*\right] &= \sum_{i=1}^L x_i^* \sum_{k=i}^L \mathbb{E}[\delta_{i_k^{r^*}, [k+1-i, \infty)}] \\ &= \sum_{i=1}^L x_i^* \sum_{k=i}^L \mathbb{P}(i_k^{r^*} \geq k+1-i) \\ &= \sum_{i=1}^L x_i^* \sum_{k=1}^{L+1-i} \mathbb{P}(i_{k+i-1}^{r^*} \geq k) \\ &= \sum_{i=1}^L x_i^* \sum_{k=1}^{L+1-i} \mathbb{P}(i_\infty^{r^*} \geq k). \end{aligned} \quad (10)$$

Applying the definition of r^* together with Lemma 3 and the fact that $\mathbb{E}[i_k^{r^*}] \leq \mathbb{E}[i_\infty^{r^*}]$ by Lemma 6, yields that (9) is at most

$$\begin{aligned} r^* \sum_{k=1}^L \mathbb{E}[i_\infty^{r^*}] - \sum_{i=1}^L x_i^* \sum_{k=1}^{L+1-i} \mathbb{P}(i_\infty^{r^*} \geq k) &= \sum_{i=1}^L x_i^* \sum_{k=1}^{\infty} \mathbb{P}(i_\infty^{r^*} \geq k) - \sum_{i=1}^L x_i^* \sum_{k=1}^{L+1-i} \mathbb{P}(i_\infty^{r^*} \geq k) \\ &= \sum_{i=1}^L x_i^* \sum_{k=L+2-i}^{\infty} \mathbb{P}(i_\infty^{r^*} \geq k) \\ &\leq Q \sum_{i=1}^L \sum_{k=L+2-i}^{\infty} \mathbb{P}(i_\infty^{r^*} \geq k) \\ &\leq Q \sum_{i=0}^{\infty} \sum_{k=i}^{\infty} \mathbb{P}(i_\infty^{r^*} \geq k), \end{aligned}$$

where the final inequality follows from a straightforward reindexing. Upon combining this with Lemma 6, we conclude that (9) is at most

$$Q\Theta_{r^*}^{-3}. \quad (11)$$

Next we turn to bound

$$\sum_{k=1}^L \mathbb{E}[\delta_{i_k^{r^*}, k} I_{1, \infty}^{r^*}]. \quad (12)$$

Applying the Cauchy-Schwartz inequality and Lemma 6, yields

$$\begin{aligned} \sum_{k=1}^L \mathbb{E}[\delta_{i_k^{r^*}, k} I_{1, \infty}^{r^*}] &\leq \sum_{k=1}^L \mathbb{E}^{\frac{1}{2}}[\delta_{i_k^{r^*}, k}] \mathbb{E}^{\frac{1}{2}}[(I_\infty^{r^*})^2] \\ &= \sum_{k=1}^L \mathbb{P}^{\frac{1}{2}}(i_\infty^{r^*} \geq k) \mathbb{E}^{\frac{1}{2}}[(I_\infty^{r^*})^2] \\ &\leq \sum_{k=1}^{\infty} (\Theta_{r^*}^{-1}(1 - \Theta_{r^*})^k)^{\frac{1}{2}} (2\Theta_{r^*}^{-3} \mathbb{E}[D])^{\frac{1}{2}} \\ &= 2^{\frac{1}{2}} \mathbb{E}[D] \Theta_{r^*}^{-2} \frac{(1 - \Theta_{r^*})^{\frac{1}{2}}}{1 - (1 - \Theta_{r^*})^{\frac{1}{2}}} \leq 2^{\frac{3}{2}} \mathbb{E}[D] \Theta_{r^*}^{-3}, \end{aligned} \quad (13)$$

where the final inequality follows from multiplying and dividing by $1 + (1 - \Theta_{r^*})^{\frac{1}{2}}$ and from noting that $(1 - \Theta_{r^*})^{\frac{1}{2}}(1 + (1 - \Theta_{r^*})^{\frac{1}{2}}) \leq 2$.

Finally, using both (11) to bound (9) and (13) to bound (12) in (8), completes the proof. \square

7. Proof of main result. We now complete the proof of our main result, namely Theorem 1, by combining Theorem 2 with several additional bounds. In light of Theorem 2, the primary difficulty which remains is proving that r^* is bounded away from $\mathbb{E}[D]$ as $L \rightarrow \infty$. Recall that $z = \arg \min_{v \geq 0} (h\mathbb{E}[I_\infty^v] - cv)$. The final step will be to show that π_{r^*} is itself dominated by the policy π_z , which we will prove to be the “best-possible” constant-order policy. Combining these results with a few straightforward asymptotic arguments will complete the proof.

7.1. Bounding r^* away from $\mathbb{E}[D]$. Let us begin by proving that r^* is bounded away from $\mathbb{E}[D]$ as $L \rightarrow \infty$. In particular, we will prove the following result, recalling that $m = \lceil (26(3\zeta + c(h\sigma)^{-1}\mathbb{E}[D] + 1))^2 \rceil$.

LEMMA 7. For all $L \geq 8(\sigma^{-1}Q + 1)m^{\frac{3}{2}}$, we have that $\mathbb{E}[D] - r^* \geq \frac{1}{2}\sigma m^{-\frac{1}{2}}$.

Although at first glance one might think such a result to be straightforward, the fact that in principle the components of \mathbf{x}^* could vary considerably over the L periods creates difficulties here. Indeed, we will actually argue indirectly as follows. Roughly speaking, we argue that for any time t , if in the m periods leading up to and including t (i.e., periods $t - m + 1, \dots, t$) the corresponding pipeline vector components (i.e., $x_{t-m+1}^*, \dots, x_t^*$) were much larger (on average) than $\mathbb{E}[D]$ (i.e., $\sum_{k=t-m+1}^t x_k^* - m\mathbb{E}[D]$ is too large), then the expected value of the inventory at the end of period t , I_{t+1}^* , will be very large.

More precisely, defining $\gamma_t \triangleq \sum_{k=t-m+1}^t x_k^*$, we note that $\mathbb{E}[I_{t+1}^*] \geq \mathbb{E}[\max(0, \gamma_t - \sum_{k=t-m+1}^t D_k)]$. We then apply a type of central limit theorem scaling, combined with certain explicit bounds on the rate of convergence in the central limit theorem (i.e., Stein’s method), to argue that if $\gamma_t - m\mathbb{E}[D]$ is large, then $\mathbb{E}[I_{t+1}^*]$ is large. It will follow that $\gamma_t - m\mathbb{E}[D]$ cannot be large for too many values of t . A key insight here is that because of the maximum within the expectation, we get a strong lower bound even when $\gamma_t - m\mathbb{E}[D] = 0$, which in turn allows us to show that $\gamma_t - m\mathbb{E}[D]$ should typically be significantly less than zero. We then argue that $\sum_{t=m}^L \gamma_t$ is sufficiently close to $m \sum_{k=1}^L x_k^*$, and combine the above observations to conclude that $\mathbb{E}[D] - r^*$ must be bounded away from zero.

Proof of Lemma 7: It follows from (2) and non-negativity that, for all $t \in [m, L]$, w.p.1

$$I_t^* + x_t^* - D_t \geq \gamma_t - \sum_{k=t-m+1}^t D_k.$$

Thus, for all $t \in [m, L]$, $\mathbb{E}[\max(0, I_t^* + x_t^* - D_t)]$ is at least

$$\begin{aligned} & \mathbb{E} \left[\max \left(0, \gamma_t - \sum_{k=t-m+1}^t D_k \right) \right] \\ &= \sigma m^{\frac{1}{2}} \mathbb{E} \left[\max \left(0, \frac{\sum_{k=t-m+1}^t (\mathbb{E}[D] - D_k)}{\sigma m^{\frac{1}{2}}} + \frac{\gamma_t - m\mathbb{E}[D]}{\sigma m^{\frac{1}{2}}} \right) \right]. \end{aligned} \quad (14)$$

Let N denote a standard normal r.v. We now show that (14) is well-approximated by

$$\sigma m^{\frac{1}{2}} \mathbb{E} \left[\max \left(0, N + \frac{\gamma_t - m\mathbb{E}[D]}{\sigma m^{\frac{1}{2}}} \right) \right],$$

using known results on the rate of convergence in the central limit theorem. Such results are typically derived via Stein’s method, and we refer the interested reader to Chen and Shao [5] for details. Specifically, the following explicit bound on the rate of convergence in the central limit theorem is generally well known.

THEOREM 3 (Chen and Shao [5]). Suppose that $F : \mathcal{R} \rightarrow \mathcal{R}$ is any Lipschitz-continuous function with Lipschitz constant at most unity, i.e., for all $x, y \in \mathcal{R}$, $|F(x) - F(y)| \leq |x - y|$. Suppose that $\{X_i; i \geq 1\}$ is any sequence of i.i.d. r.v.s s.t. $\mathbb{E}[X_1] = 0$, $\mathbb{E}[X_1^2] = 1$, and $\mathbb{E}[|X_1|^3] < \infty$. Then, for all $n \geq 1$,

$$\left| \mathbb{E} \left[F \left(n^{-\frac{1}{2}} \sum_{i=1}^n X_i \right) \right] - \mathbb{E}[F(N)] \right| \leq 3n^{-\frac{1}{2}} \mathbb{E}[|X_1|^3].$$

Letting $F_t(x) \triangleq \max \left(0, x + \frac{\gamma_t - m\mathbb{E}[D]}{\sigma m^{\frac{1}{2}}} \right)$, it follows from Theorem 3 and (14) that, for all $t \in [m, L]$,

$$\mathbb{E}[\max(0, I_t^* + x_t^* - D_t)] \geq \sigma m^{\frac{1}{2}} \left(\mathbb{E} \left[\max \left(0, N + \frac{\gamma_t - m\mathbb{E}[D]}{\sigma m^{\frac{1}{2}}} \right) \right] - 3m^{-\frac{1}{2}}\zeta \right),$$

and thus

$$\sum_{t=1}^L \mathbb{E}[C_t^*] \geq h\sigma m^{\frac{1}{2}} \sum_{t=m}^L \mathbb{E} \left[\max \left(0, N + \frac{\gamma_t - m\mathbb{E}[D]}{\sigma m^{\frac{1}{2}}} \right) \right] - 3h\sigma\zeta L. \quad (15)$$

Note that $\psi(y) \triangleq \mathbb{E}[\max(0, N + y)]$ is a convex function of y . It then follows from Jensen's inequality that

$$\sum_{t=m}^L (L - m + 1)^{-1} \mathbb{E} \left[\max \left(0, N + \frac{\gamma_t - m\mathbb{E}[D]}{\sigma m^{\frac{1}{2}}} \right) \right]$$

is at least

$$\mathbb{E} \left[\max \left(0, N + (L - m + 1)^{-1} \sum_{t=m}^L \frac{\gamma_t - m\mathbb{E}[D]}{\sigma m^{\frac{1}{2}}} \right) \right]. \quad (16)$$

Since $(\mathbf{0}, 0)$ is a feasible solution to Problem (4), with value $cL\mathbb{E}[D]$, the optimality of $(\mathbf{x}^*, \mathcal{I}^*)$ implies

$$\sum_{t=1}^L \mathbb{E}[C_t^*] \leq Lc\mathbb{E}[D]. \quad (17)$$

Combining (15), (16), and (17) yields

$$\mathbb{E} \left[\max \left(0, N + (L - m + 1)^{-1} \sum_{t=m}^L \frac{\gamma_t - m\mathbb{E}[D]}{\sigma m^{\frac{1}{2}}} \right) \right] \leq \frac{L}{L - m + 1} (3\zeta + c(h\sigma)^{-1}\mathbb{E}[D])m^{-\frac{1}{2}}. \quad (18)$$

We now relate $\sum_{t=m}^L \gamma_t$ to $\sum_{t=1}^L x_t^*$ by proving that

$$\sum_{t=m}^L \gamma_t = \sum_{t=m}^L \sum_{k=t-m+1}^t x_k^* \geq m \sum_{k=1}^L x_k^* - 2m^2Q. \quad (19)$$

Indeed, it follows from a straightforward counting argument that for all $t \in [m, L - m]$, x_t^* appears exactly m times in the double sum $\sum_{t=m}^L \sum_{k=t-m+1}^t x_k^*$, and thus

$$\sum_{t=m}^L \gamma_t \geq m \sum_{t=m}^{L-m} x_t^*.$$

Moreover, since $x_t^* \leq Q$ for all t , we conclude that

$$m \sum_{t=m}^{L-m} x_t^* \geq m \sum_{k=1}^L x_k^* - 2m^2Q.$$

In combination, these results yield (19).

Next, upon combining (18) and (19) with the monotonicity of ψ , and simplifying all relevant expressions, we obtain

$$\mathbb{E} \left[\max \left(0, N + \frac{m^{\frac{1}{2}}}{\sigma} \left(\frac{L}{L - m + 1} r^* - \mathbb{E}[D] - \frac{2mQ}{L - m + 1} \right) \right) \right] \leq \frac{L}{L - m + 1} (3\zeta + c(h\sigma)^{-1}\mathbb{E}[D])m^{-\frac{1}{2}}. \quad (20)$$

Noting that $L \geq 2m$ implies $\frac{L}{L - m + 1} \leq 2$ and $L \geq 8\sigma^{-1}m^{\frac{3}{2}}Q$ implies $\frac{2mQ}{L} \leq \frac{1}{4}\sigma m^{-\frac{1}{2}}$, we devise from the monotonicity of ψ that, for all $L \geq \max(2m, 8\sigma^{-1}m^{\frac{3}{2}}Q)$,

$$\mathbb{E} \left[\max \left(0, N + \frac{m^{\frac{1}{2}}}{\sigma} \left(r^* - \mathbb{E}[D] - \frac{1}{2}\sigma m^{-\frac{1}{2}} \right) \right) \right] \leq 2(3\zeta + c(h\sigma)^{-1}\mathbb{E}[D])m^{-\frac{1}{2}}. \quad (21)$$

It is easily verified that $(\mathbb{E}[\max(0, N - 1)])^{-1} \leq 13$, and thus from definitions and basic algebra

$$m \geq \left(2(3\zeta + c(h\sigma)^{-1}\mathbb{E}[D])(\mathbb{E}[\max(0, N - 1)])^{-1} \right)^2.$$

We conclude that the right-hand-side of (21) is at most $\mathbb{E}[\max(0, N - 1)]$, in which case the monotonicity of ψ implies

$$\frac{m^{\frac{1}{2}}}{\sigma} \left(r^* - \mathbb{E}[D] - \frac{1}{2}\sigma m^{-\frac{1}{2}} \right) \leq -1,$$

namely $\mathbb{E}[D] - r^* \geq \frac{1}{2}\sigma m^{-\frac{1}{2}}$. Combining the above with some straightforward algebra completes the proof. \square

We end this subsection by combining Theorem 2 and Lemma 7 to bound the difference between the constant-order policy and an optimal policy.

COROLLARY 2. *For all $L \geq 8(\sigma^{-1}Q + 1)m^{\frac{3}{2}}$ and any $\tau \in [L + 1, T + 1]$, (7) is at most*

$$2^{12}h(Q + 2^{\frac{3}{2}}\mathbb{E}[D])(\mathbb{E}^2[D] + \mathbb{E}[D^2])^3\sigma^{-6}m^3 + c(\lceil(2ch^{-1}L)^{\frac{1}{2}}\rceil + 2)\mathbb{E}[D].$$

Proof: It follows from Lemma 7 that

$$h(Q + 2^{\frac{3}{2}}\mathbb{E}[D])\Theta_{r^*}^{-3} \leq 2^{12}h(Q + 2^{\frac{3}{2}}\mathbb{E}[D])(\mathbb{E}^2[D] + \mathbb{E}[D^2])^3\sigma^{-6}m^3.$$

Thus, by Theorem 2, to complete the proof it suffices to demonstrate that $\mathcal{I}^* \leq (\lceil(2ch^{-1}L)^{\frac{1}{2}}\rceil + 2)\mathbb{E}[D]$. Indeed, suppose for contradiction that $\mathcal{I}^* > (\lceil(2ch^{-1}L)^{\frac{1}{2}}\rceil + 2)\mathbb{E}[D]$. For all $k \in [1, \lceil(2ch^{-1}L)^{\frac{1}{2}}\rceil + 2]$, we have

$$\mathbb{E}[I_{1+k}^*] \geq \mathcal{I}^* - k\mathbb{E}[D] > \mathbb{E}[D](\lceil(2ch^{-1}L)^{\frac{1}{2}}\rceil + 2 - k).$$

The resulting holding costs ensure that $\sum_{t=1}^L \mathbb{E}[C_t^*]$ is strictly greater than

$$h\mathbb{E}[D] \sum_{k=1}^{\lceil(2ch^{-1}L)^{\frac{1}{2}}\rceil} k \geq c\mathbb{E}[D]L.$$

Combining this with (17) completes the proof. \square

7.2. Proof of Theorem 1. With Corollary 2 in hand, we now proceed with the proof of our main result, i.e., Theorem 1.

Proof of Theorem 1: Suppose $T \geq L$, $\epsilon \in (0, 1)$, and $L \geq 8(\sigma^{-1}Q + 1)m^{\frac{3}{2}}$. It then follows from Lemma 7 that $r^* < \mathbb{E}[D]$. Note that

$$\frac{\sum_{t=L+1}^{T+L} \mathbb{E}[C_t^{\pi_{r^*}}]}{\sum_{t=L+1}^{T+L} \mathbb{E}[C_t^{\pi}]} \tag{22}$$

equals

$$\frac{\sum_{k=1}^{\lfloor \frac{T}{L} \rfloor} \sum_{t=kL+1}^{(k+1)L} \mathbb{E}[C_t^{\pi_{r^*}}] + \sum_{t=(\lfloor \frac{T}{L} \rfloor + 1)L+1}^{T+L} \mathbb{E}[C_t^{\pi_{r^*}}]}{\sum_{k=1}^{\lfloor \frac{T}{L} \rfloor} \sum_{t=kL+1}^{(k+1)L} \mathbb{E}[C_t^{\pi}] + \sum_{t=(\lfloor \frac{T}{L} \rfloor + 1)L+1}^{T+L} \mathbb{E}[C_t^{\pi}]}$$

As the policy π_{r^*} is stationary and yields a stationary sequence of inventory and ordering levels, it follows that $\mathbb{E}[C_t^{\pi_{r^*}}] = \mathbb{E}[C_{L+1}^{\pi_{r^*}}]$ for all $t \geq L + 1$. Further note that $\mathbb{E}[C_t^{\pi}] \geq g$ for all $t \geq L + 1$,

and thus $\sum_{t=kL+1}^{(k+1)L} \mathbb{E}[C_t^{\pi}] \geq Lg$ for all $k \in [1, \lfloor \frac{T}{L} \rfloor]$. Combining the above with Corollary 2, and the non-negativity of all relevant terms, we conclude that (22) is at most

$$\frac{\lfloor \frac{T}{L} \rfloor}{\lfloor \frac{T}{L} \rfloor} \left(1 + \left(2^{12} h(Q + 2^{\frac{3}{2}} \mathbb{E}[D]) (\mathbb{E}^2[D] + \mathbb{E}[D^2])^3 \sigma^{-6} m^3 + c(\lceil (2ch^{-1}L)^{\frac{1}{2}} \rceil + 2) \mathbb{E}[D] \right) (gL)^{-1} \right). \quad (23)$$

Next, we note that $L \geq y_1(\epsilon)$ implies

$$2^{12} h(Q + 2^{\frac{3}{2}} \mathbb{E}[D]) (\mathbb{E}^2[D] + \mathbb{E}[D^2])^3 \sigma^{-6} m^3 (gL)^{-1} \leq \frac{\epsilon}{4},$$

$$L \geq y_2(\epsilon) \quad \text{implies} \quad c(\lceil (2ch^{-1}L)^{\frac{1}{2}} \rceil + 2) \mathbb{E}[D] (gL)^{-1} \leq \frac{\epsilon}{12},$$

and

$$T \geq \left(1 + \frac{3}{\epsilon} \right) L \quad \text{implies} \quad \frac{\lfloor \frac{T}{L} \rfloor}{\lfloor \frac{T}{L} \rfloor} \leq 1 + \frac{\epsilon}{3}.$$

Combining the above with the facts that, for all $\epsilon \in (0, 1)$, $(1 + \frac{\epsilon}{3})^2 \leq 1 + \epsilon$ and $y_1(\epsilon) \geq 8(\sigma^{-1}Q + 1)m^{\frac{3}{2}}$ then completes the proof that the stated performance guarantees are attained by the policy π_{r^*} , for any L, T satisfying the conditions of Theorem 1.

The final step is to prove that the same guarantees extend to π_2 . Indeed, it follows from stationarity that, for any $r \in [0, \mathbb{E}[D]]$ and $t \geq L + 1$,

$$\mathbb{E}[C_t^{\pi_r}] = h\mathbb{E}[(I_\infty^r + r - D)^+] + c\mathbb{E}[(I_\infty^r + r - D)^-]. \quad (24)$$

However, since I_∞^r is distributed as $(I_\infty^r + r - D)^+$, we conclude

$$\mathbb{E}[(I_\infty^r + r - D)^+] = \mathbb{E}[I_\infty^r].$$

Furthermore, since

$$\mathbb{E}[(I_\infty^r + r - D)^+] - \mathbb{E}[(I_\infty^r + r - D)^-] = \mathbb{E}[I_\infty^r] + r - \mathbb{E}[D],$$

it follows that

$$\mathbb{E}[(I_\infty^r + r - D)^-] = \mathbb{E}[D] - r.$$

In combination with (24), we have that for any $r \in [0, \mathbb{E}[D]]$ and $t \geq L + 1$

$$\mathbb{E}[C_t^{\pi_r}] = h\mathbb{E}[I_\infty^r] + c(\mathbb{E}[D] - r). \quad (25)$$

The desired result then follows from the fact that z is a minimizer of (25), completing the proof. \square

8. Conclusion. In this paper, we considered the single-item, periodic-review, lost-sales model with positive lead times and i.i.d. demand, for which the optimal policy is poorly understood and computationally intractable. We proved that, as the lead time grows (with the demand distribution, lost-sales penalty, and holding cost remaining fixed), a simple, open-loop constant-order policy is in fact asymptotically optimal. We also established explicit bounds on how large the lead time should be to ensure that the best constant-order policy incurs an expected cost of at most $1 + \epsilon$ times that incurred by the optimal policy. To the best of our knowledge, this is the first algorithm proven to be within $1 + \epsilon$ of optimal for lost-sales models when the lead time is large, while maintaining a runtime that does not grow with the lead time. Our main proof technique involved a novel coupling for suprema of random walks, and may be useful in other settings.

This work leaves many interesting directions for future research. We suspect that our explicit bounds are not tight, and a more precise analysis of the constant-order policy would further help to explain the good performance of the algorithm for lead times as small as four, as reported by Zipkin [43]. Since lost sales models commonly arise in practice, an interesting challenge is to combine the core ideas of our analysis with known results from dynamic programming to derive and analyze practical “hybrid” algorithms, which use more elaborate forms of dynamic programming when the lead time is small and gradually transition to less computationally intensive algorithms (with the constant-order policy at the extreme) as the lead time grows. It would also be interesting to prove that a similar phenomenon occurs for other policies, as well as in other inventory models. In particular, it is an interesting open question whether other simple (but perhaps slightly more sophisticated) policies, such as the order-up-to policy considered by Huh et al. [19] and the cost-balancing policy considered by Levi et al. [26], exhibit a similar asymptotic optimality as the lead time grows. On a related note, the fact that $y(\epsilon)$ contains terms of the form ch^{-1} suggests that the constant-order policy may require larger lead times to approach near-optimality when the ratio of lost-sales-penalty to holding-cost is large. Since this is exactly the regime in which the order-up-to policy of Huh et al. [19] provably works well, it is an interesting open question to try and combine these (and perhaps other) algorithms to yield tighter performance guarantees over a larger range of parameters.

Philosophically, our main results and insights fall under the broad heading of “long-range independence / decay of correlations” phenomena, in which so much uncertainty is introduced into a model that even very sophisticated algorithms cannot perform significantly better than very simple algorithms. Such ideas have led to significant progress on fundamental models in other fields (cf. Gamarnik et al. [11], Sly [37], Gamarnik and Sudan [12], Deshpande and Montanari [8]), and may prove useful in other operations management problems.

9. Appendix.

9.1. Proof of Lemma 6.

In this appendix, we provide the proof of Lemma 6. *Proof of Lemma 6:* We first prove that i_k^r has the same distribution as $\min(k, i_\infty^r)$. Let $\{D'_i; i \geq 1\}$ be an additional sequence of i.i.d. realizations from \mathcal{D} , mutually independent from $\{D_i; i \geq 1\}$. Then, for any $k \geq 1$, we can construct $I_{1,\infty}^r, W_k^r, \{D_i; i \geq 1\}, \{D'_i; i \geq 1\}$ on the same probability space s.t. $I_{1,\infty}^r = \max_{j \geq 0} (jr - \sum_{i=1}^j D'_i)$. It is easy to see that we only need to show

$$\mathbb{P}(i_k^r = j) = \mathbb{P}(i_\infty^r = j), \quad j = 0, \dots, k-1.$$

By definition, we have $\mathbb{P}[i_k^r = j] = \mathbb{P}[\mathcal{I}_1 \cap \mathcal{I}_2 \cap \mathcal{I}_3]$ where

$$\begin{aligned} \mathcal{I}_1 &= \left\{ \ell r - \sum_{i=1}^{\ell} D_i \leq jr - \sum_{i=1}^j D_i, \quad \forall \ell \leq j \right\}, \\ \mathcal{I}_2 &= \left\{ \ell r - \sum_{i=1}^{\ell} D_i < jr - \sum_{i=1}^j D_i, \quad \ell = j+1, \dots, k-1 \right\}, \\ \mathcal{I}_3 &= \left\{ jr - \sum_{i=1}^j D_i > kr - \sum_{i=1}^k D_i + I_{1,\infty}^r \right\}. \end{aligned}$$

Note that

$$\begin{aligned} \mathcal{I}_3 &= \left\{ jr - \sum_{i=1}^j D_i > kr - \sum_{i=1}^k D_i + \max_{\ell \geq 0} \left(\ell r - \sum_{i=1}^{\ell} D'_i \right) \right\} \\ &= \left\{ jr - \sum_{i=1}^j D_i > \max_{\ell \geq 0} \left[(k+\ell)r - \left(\sum_{i=1}^k D_i + \sum_{i=1}^{\ell} D'_i \right) \right] \right\}. \end{aligned}$$

It therefore follows, since $\{D_i; i \geq 1\}$ and $\{D'_i; i \geq 1\}$ are mutually independent i.i.d. sequences with common distribution \mathcal{D} , that

$$\mathbb{P}[i_k^r = j] = \mathbb{P}[\mathcal{I}_1 \cap \mathcal{I}_2 \cap \mathcal{I}'_3],$$

where

$$\mathcal{I}'_3 = \left\{ jr - \sum_{i=1}^j D_i > \max_{\ell \geq k} \left[\ell r - \sum_{i=1}^{\ell} D_i \right] \right\}.$$

Noting that this is the definition of $\mathbb{P}(i_\infty^r = j)$ completes the first part of the proof.

Before proving the remainder of the lemma, it will be useful to establish that

$$\mathbb{P}(i_\infty^r = k) \leq (1 - \Theta_r)^k. \quad (26)$$

By definition,

$$\mathbb{P}(i_\infty^r = k) \leq \mathbb{P}\left(\sum_{i=1}^k (r - D_i) \geq 0\right).$$

Applying a Chernoff bound, we find that for any $\theta > 0$

$$\mathbb{P}(i_\infty^r = k) \leq \mathbb{E}^k[\exp(\theta(r - D))]$$

where

$$\begin{aligned} \mathbb{E}[\exp(\theta(r - D))] &= \exp(\theta r) \mathbb{E}[\exp(-\theta D)] \\ &\leq \exp(\theta r) \mathbb{E}[(1 + \theta D)^{-1}] && \text{(since } \exp(v) \geq 1 + v\text{)} \\ &\leq \mathbb{E}\left[\frac{1 + \theta r + \theta^2 r^2}{1 + \theta D}\right], && \text{for all } \theta \in (0, r^{-1}], \end{aligned}$$

the final inequality following from a simple Taylor-series expansion. However, w.p.1, we have

$$\begin{aligned} \frac{1 + \theta r + \theta^2 r^2}{1 + \theta D} &= 1 + \theta(r - D) + \frac{\theta^2}{1 + \theta D}(r^2 - D(r - D)) \\ &\leq 1 + \theta(r - D) + \theta^2(r^2 + D^2), \end{aligned}$$

and thus

$$\mathbb{E}[\exp(\theta(r - D))] \leq 1 - \theta(\mathbb{E}[D] - r) + \theta^2(r^2 + \mathbb{E}[D^2]).$$

Observing that

$$\begin{aligned} \frac{\mathbb{E}[D] - r}{2(r^2 + \mathbb{E}[D^2])} &\leq \frac{\mathbb{E}[D]}{2\mathbb{E}[D^2]} \\ &\leq \frac{\mathbb{E}[D]}{2\mathbb{E}^2[D]} = \frac{1}{2\mathbb{E}[D]} < r^{-1}, \end{aligned}$$

we may take

$$\theta = \theta^* \triangleq \frac{\mathbb{E}[D] - r}{2(r^2 + \mathbb{E}[D^2])}$$

to conclude

$$\begin{aligned} \mathbb{E}[\exp(\theta^*(r - D))] &\leq 1 - \frac{(\mathbb{E}[D] - r)^2}{4(r^2 + \mathbb{E}[D^2])} \\ &\leq 1 - \frac{(\mathbb{E}[D] - r)^2}{4(\mathbb{E}^2[D] + \mathbb{E}[D^2])} = 1 - \Theta_r, \end{aligned}$$

where the final inequality follows from the fact that $r^2 \leq \mathbb{E}^2[D]$. Combining the above completes the proof of (26).

With (26) in hand, the lemma follows directly from the basic manipulation of a few geometric series and the fact that $I_\infty^r \leq r i_\infty^r \leq \mathbb{E}[D] i_\infty^r$ w.p.1, the details of which we omit. \square

Acknowledgments. The authors thank Maury Bramson, Jim Dai, Bruce Hajek, Retsef Levi, Marty Reiman, Alan Scheller-Wolf, and Paul Zipkin for several stimulating discussions. We also thank the anonymous referees and the editors for their constructive feedback on an earlier version of this work. The first author also thanks the IBM T.J. Watson Research Center, for providing a great research environment.

References

- [1] Asmussen, S. 2003. *Applied Probability and Queues*. Second ed. Springer.
- [2] Bellman, R., I. Glicksberg, O. Gross. 1955. On the optimal inventory equation. *Management Science* **2** 83–104.
- [3] Bijvank, M., I. Vis. 2012. Lost-sales inventory systems with a service level criterion. *European Journal of Operational Research* **220**(3) 610–618.
- [4] Bijvank, M., I.F.A. Vis. 2011. Lost-sales inventory theory: A review. *European Journal of Operations Research* **215** 1–13.
- [5] Chen, L., Q. Shao. 2005. Stein’s method for normal approximation. *An introduction to Stein’s method* **4** 1–59.
- [6] 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.
- [7] Chen, W., M. Dawande, G. Janakiraman. 2014. Integrality in stochastic inventory models. *Production and Operations Management* .
- [8] Deshpande, Y., A. Montanari. 2015. Finding hidden cliques of size $\sqrt{\frac{n}{e}}$ in nearly linear time. *Foundations of Computational Mathematics* 1–60.
- [9] Downs, B., R. Metters, J. Semple. 2001. Managing inventory with multiple products, lags in delivery, resource constraints and lost sales: A mathematical programming approach. *Management Science* **47**(3) 464–479.
- [10] Feller, W. 1971. *An Introduction to Probability Theory and Its Applications*, vol. II. Second ed. John Wiley and Sons.
- [11] Gamarnik, D., D.A. Goldberg, T. Weber. 2013. Correlation decay in random decision networks. *Mathematics of Operations Research* **39**(2) 229–261.
- [12] Gamarnik, D., M. Sudan. 2014. Limits of local algorithms over sparse random graphs. *Proceedings of the 5th conference on Innovations in theoretical computer science*. ACM, 369–376.
- [13] Gaver, D.P. 1959. On base-stock level inventory control. *Operations Research* **7** 689–703.
- [14] Goldberg, D., D. Katz-Rogozhnikov, Y. Lu, M. Sharma, M. Squillante. 2012. Asymptotic optimality of constant-order policies for lost sales inventory models with large lead times. *arXiv preprint arXiv:1211.4063v1* .
- [15] Hajek, B. 1983. The proof of a folk theorem on queuing delay with applications to routing in networks. *Journal of the ACM (JACM)* **30**(4) 834–851.
- [16] 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.
- [17] Halman, N., J.B. 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.
- [18] Huh, T., G. Janakiraman. 2010. On the optimal policy structure in serial inventory systems with lost sales. *Operations Research* **58**(2) 486–491.
- [19] Huh, T., G. Janakiraman, J. Muckstadt, P. Rusmevichientong. 2009. Asymptotic optimality of order-up-to policies in lost sales inventory systems. *Management Science* **55**(3) 404–420.
- [20] Humblet, P. 1982. Determinism minimizes waiting time in queues. *MIT Technical report* .
- [21] Janakiraman, G., R. Roundy. 2004. Lost-sales problems with stochastic lead times: Convexity results for base-stock policies. *Operations Research* **52**(5) 795–803.

- [22] Johansen, S.G. 2001. Pure and modified base-stock policies for the lost sales inventory system with negligible set-up costs and constant lead times. *International Journal of Production Economics* **71** 391–399.
- [23] Johansen, S.G., A. Thorstenson. 2008. Pure and restricted base-stock policies for the lost-sales inventory system with periodic review and constant lead times. *Technical report, Aarhus University* .
- [24] Karlin, S., H. Scarf. 1958. Inventory models of the Arrow-Harris-Marschak type with time lag. K.J. Arrow, S. Karlin, H. Scarf, eds., *Studies in the Mathematical Theory of Inventory and Production*. Stanford University Press, 155–178.
- [25] 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.
- [26] Levi, R., M. Pa’l, R. Roundy, D. Shmoys. 2008. Approximation algorithms for stochastic inventory control models. *Mathematics of Operations Research* **32**(2) 284–302.
- [27] Levi, R., R. Roundy, D. Shmoys, V. Truong. 2008. Approximation algorithms for capacitated stochastic inventory control models. *Operations Research* **56**(5) 1184–1199.
- [28] Loulou, R. 1978. An explicit upper bound for the mean busy period in a GI/G/1 queue. *Journal of Applied Probability* **15**(2) 452–455.
- [29] Lu, Y., M. Squillante, D. Yao. 2015. Matching supply and demand in production-inventory systems: Asymptotics and optimization. *arXiv preprint arXiv:1501.07141* .
- [30] Morse, P.M. 1959. Solutions of a class of discrete-time inventory problems. *Operations Research* **7** 67–68.
- [31] 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.
- [32] Morton, T. E. 1971. The near-myopic nature of the lagged-proportional-cost inventory problem with lost sales. *Operations Research* **19**(7) 1708–1716.
- [33] Nagaev, S. V. 2010. Exact expressions for the moments of ladder heights. *Sibirsk. Mat. Zh.* **51**(4) 848–870. (in Russian). English translation in *Siberian Mathematical Journal*, **51** (2010), no. 4, 675–695.
- [34] Nahmias, S. 1979. Simple approximations for a variety of dynamic leadtime lost-sales inventory models. *Operations Research* **27**(5) 904–924.
- [35] Pressman, I. 1977. An order-level-scheduling-period system with lost sales. *Management science* **17** 1328–1335.
- [36] Reiman, M. 2004. A new and simple policy for a continuous review lost-sales inventory model. Unpublished manuscript.
- [37] Sly, A. 2010. Computational transition at the uniqueness threshold. *Foundations of Computer Science (FOCS), 2010 51st Annual IEEE Symposium on*. IEEE, 287–296.
- [38] Sparre-Andersen, E. 1953. On the fluctuations of sums of random variables. *Math. Scand.* **1** 263–285.
- [39] Sparre-Andersen, E. 1954. On the fluctuations of sums of random variables. *Math. Scand.* **2** 195–223.
- [40] Yaspan, A. 1961. An inclusive solution to the inventory problem. *Operations Research* **9** 371–382.
- [41] Yaspan, A. 1972. Fixed-stockout-probability order quantities with lost sales and time lag. *Operations Research* **20** 903–904.
- [42] Zipkin, P. 2000. Fundamentals of inventory management. *McGraw-Hill, New York* .
- [43] Zipkin, P. 2008. Old and new methods for lost-sales inventory systems. *Operations Research* **56**(5) 1256–1263.
- [44] Zipkin, P. 2008. On the structure of lost-sales inventory models. *Operations Research* **56**(4) 937–944.