

Optimality of (s, S) Inventory Policies under Renewal Demand and General Cost Structures

Sandun Perera 

School of Management, The University of Michigan–Flint, Flint, Michigan 48502, USA, sperera@umich.edu

Ganesh Janakiraman*, Shun-Chen Niu

Naveen Jindal School of Management, The University of Texas at Dallas, Richardson, Texas 75080, USA,
ganesh@utdallas.edu, scnium@utdallas.edu

We study a single-stage, continuous-time inventory model where unit-sized demands arrive according to a renewal process and show that an (s, S) policy is optimal under minimal assumptions on the ordering/procurement and holding/backorder cost functions. To our knowledge, the derivation of almost all existing (s, S) -optimality results for stochastic inventory models assume that the ordering cost is composed of a fixed setup cost and a proportional variable cost; in contrast, our formulation allows virtually any reasonable ordering-cost structure. Thus, our paper demonstrates that (s, S) -optimality actually holds in an important, primitive stochastic setting for *all other* practically interesting ordering cost structures such as well-known quantity discount schemes (e.g., all-units, incremental and truckload), multiple setup costs, supplier-imposed size constraints (e.g., batch-ordering and minimum-order-quantity), arbitrary increasing and concave cost, as well as any variants of these. It is noteworthy that our proof only relies on elementary arguments.

Key words: stochastic inventory models; (s, S) -optimality; general ordering/procurement cost structures

History: Received: May 2017; Accepted: September 2017 by Qi Annabelle Feng, after 1 revision.

1. Introduction

Inventory management is at the core of supply chain optimization and Operations Management. In standard inventory models, it is typical to assume a fixed-plus-proportional ordering cost structure; that is, there is a fixed cost for ordering any strictly positive quantity together with an incremental cost proportional to this quantity. However, inventory models with more general ordering-cost structures are prevalent in practice. For example, discount schemes such as all-unit discounts, incremental discounts, truckload discounts are widely-used and studied in the literature (cf. Altintas et al. 2008, Benton and Park 1996, Chen 2009, Federguen and Lee 1990, Li et al. 2004, 2012, Zipkin 2000). Even more sophisticated cost structures, such as modified all-unit discounts (cf. Chan et al. 2002), generalized truckload discounts (cf. Li et al. 2004), multiple setup costs (cf. Alp et al. 2014, Lippman 1969, 1971), piecewise concave costs (cf. Koca et al. 2014, Tunc et al. 2016), and quantity-dependent fixed costs (cf. Caliskan-Demirag et al. 2012), are also abundant in applications. Surprisingly, the extant literature on the optimality of (s, S) policies—policies that raise the net-inventory level to S every time it decreases to s —in inventory models with *stochastic* demand is almost exclusively limited to the fixed-plus-proportional

ordering cost structure. Scarf (1960), Veinott (1966), and Zheng (1991) are classical examples, and Table 1 provides a short summary of the related literature. We refer the reader to the papers listed in the table, as well as to the books Beyer et al. (2010) and Porteus (2002) for thorough discussions.

Interestingly, in the arguably simpler setting of *deterministic* EOQ-type inventory models, the optimality of (s, S) policies has been established under weaker cost assumptions (see, e.g., Beyer and Sethi 1998, Lippman 1971, Perera et al. 2017). In particular, Perera et al. (2017) provides a characterization of when an (s, S) policy is optimal in the EOQ setting with general cost structures. In the same spirit as Perera et al. (2017), we establish in this paper the existence of an optimal (s, S) policy for a model with renewal demand under a minimal set of cost assumptions. To the best of our knowledge, this appears to be the first demonstration in a *primitive* stochastic setting of (s, S) -optimality under completely general cost structures. We do note however that in a recent study,¹ He et al. (2017) has shown that in a different stochastic setting where demand is governed by a Brownian motion, (s, S) -optimality holds when the fixed-plus-proportional ordering cost structure is relaxed to allow the setup cost to be a bounded and lower semi-continuous function of the order quantity.

Table 1 Summary of Papers on (s, S)-Optimality

Paper	Time horizon/ Review method	Ordering cost	Demand/Performance measure
Scarf (1960)	Finite/Periodic	F + P	IID/DC
Beckmann (1961)	Infinite/Continuous	F + P	Arbitrary inter- arrival and quantity distributions/DC
Veinott (1966)	Finite/Periodic	F + P	ID/DC
Constantinides and Richard (1978)	Infinite/Continuous	F + P	Diffusion/DC
Sulem (1986)	Infinite/Continuous	F + P	Diffusion/DC & AC
Zheng (1991)	Infinite/Periodic	F + P	IID/AC & DC
Bensoussan et al. (2005)	Infinite/Continuous	F + P	Compound Poisson with diffusion and exponentially distributed jump sizes/DC
Presman and Sethi (2006)	Infinite/Continuous	F + P	Compound Poisson with constant rate/ AC & DC
Benkherouf and Bensoussan (2009)	Infinite/Continuous	F + P	Compound Poisson with diffusion and arbitrarily distributed jump sizes/DC

Notes: F + P: Fixed plus proportional. IID: Independent and identically distributed with a common distribution. ID: Independently distributed with possibly different distributions. DC: Discounted cost. AC: Average cost.

More specifically, we study a single-stage, continuous-time inventory model where unit-sized demands arrive according to a renewal process and orders are allowed to be placed only at arrival epochs (see, e.g., Beckmann 1961, Sahin 1979, and references therein). The objective is to minimize the long-run expected average cost. We assume that the cost of ordering is finite for at least one positive order size and that the inventory holding/backorder cost is governed by a quasi-convex² function that approaches infinity when the amount of inventory/backorder approaches infinity. These assumptions are rather minimal and they cover any reasonable cost structure of practical interest, including the examples noted above as well as supplier-imposed size constraints such as batch-ordering (see, e.g., Chen 2000, Li et al. 2004) and minimum-order-quantity (see, e.g., Zhao and Katehakis 2006). Under our cost assumptions, we show that an optimal (s, S) policy always exists. Our proof is based on a lower-bounding approach that only involves elementary arguments. Thus, the primary merits of our work are the generality of the (s, S)-optimality result and the simplicity of our proof (to be summarized shortly).

It is important to note that when the ordering-cost function deviates from the fixed-plus-proportional form, an optimal (s, S) policy may not exist in the standard periodic-review inventory models with IID/ID stochastic demands (see, e.g., Alp et al. 2014,

Caliskan-Demirag et al. 2012, Chao and Zipkin 2008, Lu and Song 2014, Porteus 1990, and references therein). In particular, Porteus (1971, 1990) shows that a generalized (s, S) policy is optimal for periodic-review models with a fixed plus an increasing and strictly concave variable ordering cost and with an arbitrary (not necessarily unit-sized) demand distribution. In contrast, we show that (s, S)-optimality continues to prevail with almost no assumptions on the ordering/holding/shortage cost functions if (i) unit-sized demands arrive according to a renewal process and (ii) orders are placed at demand-arrival epochs.

Our assumption of unit-sized demands guarantees that the inventory trajectory never decreases more than one level at an arrival epoch; we shall refer to this feature as the *skip-free-to-the-left* (or simply *skip-free*) property. It will be shown later in the paper (see Lemma 7) that this is the primary driving force behind our strong conclusion. Not coincidentally, it is interesting to observe that the inventory trajectories in the deterministic-demand EOQ model of Perera et al. (2017) and in the Brownian-demand model of He et al. (2017) also satisfy the same property for a continuous state space. Hence, preserving (and exploiting) this property appears to be essential when one attempts to weaken the standard linear-cost requirements (except for the setup cost) in inventory models. While all of these three different demand assumptions share the skip-free property, we do note, however, that the conditions on the cost functions in He et al. (2017) are significantly stronger than those in Perera et al. (2017) and in the present paper. A more detailed discussion of this disparity will be provided in section 5 below.

We now briefly summarize our proof. There are four basic steps. First, we show that it is sufficient to consider policies in a slightly modified model that maintain the net-inventory level between a judiciously-chosen pair of lower and upper thresholds. This is a key step in our approach as it reduces the solution of our problem to one with a finite number of states and actions. Second, we show that the optimal cost in this stochastic, continuous time inventory problem is lower bounded by the optimal cost in a deterministic, discrete-time inventory problem in which a period is of length equal to the expected inter-arrival time in the original model. This lower-bounding problem is thus a deterministic dynamic program (DP). In our third step, we show that the assumption of unit-sized demands implies that any stationary policy in this DP is an (s, S) policy. In our final step, we invoke standard results for DPs with finite state and action spaces (see, e.g., Bertsekas 2001) to show that, under our cost assumptions, an optimal stationary policy,

now necessarily of the (s, S) type, exists in our lower-bounding DP, that this policy is feasible for the original problem, and that its cost equals the above lower bound on the optimal cost. Therefore, this (s, S) policy must be optimal for the original problem.

As in Perera et al. (2017), where a *different* set of lower-bounding arguments is developed for the EOQ setting, the proof of each of the steps outlined above is again quite elementary. As will be seen, with the exception of our use of the optimality of a deterministic, stationary policy in finite state, finite action, deterministic DPs, our proof is entirely based on first principles. While the classic discrete-time papers on (s, S) -optimality (Iglehart 1963a,b, Scarf 1960, Veinott 1966, Zheng 1991) are all elegant in their own terms, their analyses are considerably more involved and necessarily so (as the model setups are different). The corresponding (s, S) -optimality proofs in the literature for continuous-time models (Beckmann 1961, Benkherouf and Bensoussan 2009, Bensoussan et al. 2005, Constantinides and Richard 1978, Presman and Sethi 2006, Sulem 1986) involve even more sophisticated mathematical machineries such as quasi-variational inequalities. The simplicity of the analysis in our setting is therefore noteworthy.

The remainder of the paper is organized as follows. The model, assumptions, and our main result are given in section 2. The proof of the main result is presented in section 3. In section 4, we show that when the inter-arrival times are assumed to be exponential (i.e., demands follow a Poisson process), our (s, S) -optimality result can be extended to the case with a positive (constant) replenishment lead time. Finally, we provide some concluding remarks in section 5.

2. Model Formulation and the Main Result

Consider a single-product inventory model where demands of unit size arrive according to a renewal process. We assume for simplicity that a demand occurs at time 0; and that the successive inter-demand times are IID random variables with mean $1/\lambda$, where $\lambda > 0$.

Denote by \mathbb{Z} and \mathbb{N}^0 the set of integers and the set of non-negative integers, respectively; and let

- $A_n :=$ the arrival epoch of the n -th demand, with $A_0 = 0$;
- $\bar{A}_n := E[A_n]$, the expected time to the n -th arrival/demand epoch;
- $c(q) :=$ the non-negative cost for ordering q units, where $q \in \mathbb{N}^0$;
- $g(x) :=$ the non-negative holding/shortage cost rate when net-inventory is x , where $x \in \mathbb{Z}$; and

$I_0 :=$ the net-inventory level at time 0, right after demand arrival but prior to any order.

Then, we will make the following assumptions on the ordering and holding/shortage cost functions.

ASSUMPTION 1. *The function $c(\cdot)$, not necessarily finite-valued, satisfies $c(0) = 0$ and $c(q) < \infty$ for some $q \geq 1$.*

ASSUMPTION 2. *The function $g(\cdot)$, with $g(0) < \infty$, is non-increasing on $\{\dots, -2, -1, 0\}$ and non-decreasing on $\{0, 1, 2, \dots\}$ (i.e., $g(\cdot)$ is quasi-convex with $g(0) \leq g(x)$ for $x \in \mathbb{Z}$); moreover, $\lim_{x \rightarrow \pm\infty} g(x) = \infty$.*

These assumptions are minimal.³ In particular, the ordering-cost function $c(\cdot)$ need not be increasing, nor everywhere finite; thus, it indeed covers all of the cost schemes noted in section 1, including any possible constraints on order sizes. The function $g(\cdot)$ is also general, and it allows potential warehouse capacity constraints as well as backordering limits.

We next define the class of policies considered in this paper. In general, a policy is any rule for replenishing inventory. However, we shall limit attention to policies that place orders only at arrival epochs. With this assumption, let q_n^π be the size of the order placed by a policy π at the n -th arrival epoch; here, q_n^π is allowed to be zero, which simply indicates that a “genuine” order is not placed. We then claim that the sequence $\{q_n^\pi, n \geq 0\}$ fully characterizes the inventory trajectory under π . This follows immediately from the fact that the net-inventory level just after the placement of the n -th order is, independent of the arrival/demand process, given by $I_n^\pi := I_0 + \sum_{i=0}^n q_i^\pi - n$. Note in addition that, for $n \geq 1$, I_n^π depends on the “history” $\{q_i^\pi, 0 \leq i \leq n-1\}$ only via the “current” state $I_{n-1}^\pi - 1$ (or equivalently, via the sum $\sum_{i=0}^{n-1} q_i^\pi$); that is, the sequence $\{I_n^\pi, n \geq 0\}$, i.e., the *inventory trajectory* under π , is Markovian. Next, without loss of generality, we will also limit attention to policies that only involve finite ordering costs and finite holding/shortage cost rates. Thus, we shall define an admissible policy as follows:

DEFINITION 1. *Given I_0 , an admissible policy π is a non-negative deterministic sequence $\{q_n^\pi, n \geq 0\}$ with $c(q_n^\pi) < \infty$ and $g(I_n^\pi) < \infty$ for every $n \geq 0$.*

Now, for $0 \leq a < b$, let $C^\pi[a, b]$ denote the total cost (i.e., the sum of ordering, holding and backordering costs) incurred over the time interval $[a, b]$ for a given policy π . The performance of a policy π will be measured by its long-run expected average cost, which is defined⁴ as

$$f^\pi := \limsup_{T \rightarrow \infty} \frac{E[C^\pi[0, T]]}{T}. \quad (1)$$

Our objective is to minimize f^π over all admissible policies. Denote by Π the set of all admissible policies. Note that, in general, the set Π could be empty⁵ under Assumptions 1 and 2. A simple sufficient condition for Π to be non-empty and for the existence of a π in Π with $f^\pi < \infty$ will be provided shortly.⁶

For $s < S$, an (s, S) policy is defined as a policy that raises the net-inventory level to S every time it decreases to s . Note that it is implicit in this definition that an (s, S) policy may involve (if $I_0 < s$) an initialization phase⁷ that terminates at an arrival epoch where the net-inventory strictly up-crosses level s for the first time. Clearly, such an initialization phase need not be unique; thus, multiple instances of an (s, S) policy could exist. Unless an explicit clarification is necessary, we shall henceforth refer to any instance of an (s, S) policy simply as an/the (s, S) policy. Note that, in general, an arbitrary (s, S) policy need not be admissible for the given I_0 . Consider now any admissible (s, S) policy (assuming one exists). Since all cost parameters in an admissible policy are finite by definition, it follows immediately from standard renewal theory (see, e.g., Ross 1996, p. 133, Theorem 3.6.1) that if $s = x$ and $S = y$, then, the long-run expected average cost under this (x, y) policy is, independent of I_0 , given by

$$\alpha(x, y) := \frac{c(y - x) + \sum_{z=x+1}^y g(z)/\lambda}{(y - x)/\lambda}. \quad (2)$$

Our goal is to investigate whether or not an optimal (s, S) policy exists in Π . We next argue that the search for such a policy can be limited to (x, y) vectors satisfying the constraint $x < 0 \leq y$. To see this, consider, irrespective of the initial state I_0 , an arbitrary vector (x, y) of integers with $x < y$ and $\alpha(x, y) < \infty$. Observe that for any such vector, Assumption 2 implies that if $0 \leq x$, then $\alpha(x, y) \geq \alpha(-1, y - x - 1)$. Similarly, if $y < 0$, then $\alpha(x, y) \geq \alpha(x - y, 0)$. Clearly, both of the “relocated” vectors $(-1, y - x - 1)$ and $(x - y, 0)$ satisfy the above constraint.

As a final preparation for the formal statement of our main (s, S) -optimality result, we will next address, and then eliminate, the special case where the sequence $\{I_n^\pi, n \geq 0\}$ is non-increasing for every $\pi \in \Pi$. We will argue that this case, which is of little practical interest, has a trivial solution.

It is easily seen that the above case could occur under two scenarios: (i) $c(q) < \infty$ holds only for $q \leq 1$; and (ii) $c(q) < \infty$ for some $q \geq 2$ but the

“range/coverage” of the set $\{z \in \mathbb{Z} : g(z) < \infty\}$ is finite and is too narrow to admit/accommodate any order of size greater than one. For Scenario (i), Assumption 2 implies that the (s, S) policy of ordering one unit each time the net-inventory drops down to $\min\{-1, I_0\}$ is optimal.⁸ For Scenario (ii), the policy prescribed for Scenario (i) is again optimal if $c(1)$ is finite; otherwise, an admissible policy does not exist (i.e., the set Π is empty) for any I_0 . Interestingly, while these scenarios are pathological, an optimal (s, S) policy, which could be functionally dependent on the initial state I_0 , exists nonetheless.

The above discussion suggests that a compatibility condition between the cost functions $c(\cdot)$ and $g(\cdot)$ is necessary to ensure that upward movements in inventory trajectories are permitted for at least some policies in Π . This is formalized in the following new assumption:

ASSUMPTION 3. *For some $q \geq 2$, $c(q)$ is finite. Moreover, let $q_L := \min\{q \geq 2 : c(q) < \infty\}$, $z_L := \inf\{z \leq 0 : g(z) < \infty\}$, and $z_H := \sup\{z \geq 0 : g(z) < \infty\}$; then, the inequality $z_H - z_L + 1 \geq q_L$ holds whenever both z_L and z_H are finite.*

It is easily seen that Assumption 3, which can be readily verified in practice, precisely excludes the special case discussed above.

Next, let $\alpha^* := \inf\{\alpha(x, y) : x < 0 \leq y\}$. Then, we claim that Assumption 3 implies that a vector (\tilde{x}, \tilde{y}) satisfying $\tilde{x} < 0 \leq \tilde{y}$ and $\alpha(\tilde{x}, \tilde{y}) < \infty$ exists; that is, we have $\alpha^* < \infty$. To see this, note that if z_H is finite, then $\alpha(z_H - q_L, z_H) < \infty$; otherwise, $\alpha(-1, q_L - 1) < \infty$. Now, as discussed in the paragraph below Equation (2), these vectors can be relocated (if necessary) to one that meets the requirement $\tilde{x} < 0 \leq \tilde{y}$ while achieving a possibly-lower expected average cost; this establishes the claim.

A further important consequence of Assumption 3 is as follows. Suppose both z_L and z_H are finite. Then, observe that whenever the net-inventory drops down to level $z_L - 1$, a series of orders of size q_L can be placed⁹ to take it up to level z_H . It follows that starting from any $I_0 \in [z_L - 1, z_H]$, i.e., from any admissible I_0 , the inventory trajectory can reach/access any level in the set $\{z_L, \dots, z_H\}$ at a finite expected cost. Observe on the other hand that if at least one of z_L and z_H is not finite, then this “accessibility” property also holds¹⁰ under the sole requirement that $c(q)$ is finite for some $q \geq 2$. These two observations together imply that for any admissible I_0 and any (x, y) vector satisfying $x < y$ and $\alpha(x, y) < \infty$, an admissible instance of the corresponding (x, y) policy *always* exists.

Finally, let $f^* := \inf_{\pi \in \Pi} f^\pi$; and we are ready for the statement of the main result of this paper:

THEOREM 1. *Under Assumptions 1–3, there exist x^* and y^* , with $x^* < 0 \leq y^*$, such that $\alpha(x^*, y^*) = \alpha^* = f^*$; that is, any admissible instance of the (x^*, y^*) policy is optimal over Π .*

The proof of Theorem 1 is given in the next section.

3. Proof of Theorem 1

We will begin by motivating and organizing the main steps in our argument. Detailed proofs are then given in an Appendix and two subsections.

Our model has a countably-infinite state space \mathbb{Z} . The first simplifying step in our proof is to reduce the necessary argument to that for a modified model in which only a finite subset of the state space needs to be considered. The modified model will be referred to as *Model U*; and it will have a pair of ordering and holding/shortage cost functions that are suitably designed so that the search for an optimal policy in that model can be limited to policies with uniformly-bounded inventory trajectories.

Specifically, let U_L and U_H be a given pair of integers satisfying $U_L \leq 0$, $U_H \geq 0$, $U_L \leq I_0 < U_H$, and $U_H - U_L \geq 2$. (The last specification is only intended to simplify notation below.) The holding/shortage cost rates for Model U will be defined as:

$$g_U(z) := \begin{cases} g(z), & \text{if } U_L \leq z \leq U_H, \\ \infty, & \text{otherwise.} \end{cases} \quad (3)$$

In contrast with the original model where $c(\cdot)$ is a function of order size only, the ordering-cost function for Model U will be state dependent. Suppose the net-inventory levels right before and right after the placement of an order at an arrival epoch are given by x and y , respectively; then, the cost for such an order is defined as:

$$c_U(x, y) := \begin{cases} c(y - x), & \text{if } U_L \leq x < y < U_H, \\ 0, & \text{otherwise.} \end{cases} \quad (4)$$

The set of admissible policies for Model U, to be denoted by Π^U , is defined in the same manner as that for the original model. Note that $\Pi^U \subset \Pi$ does not necessarily hold; this is because when y in Equation (4) equals, for example, U_H or U_L , the corresponding ordering cost in the original model could be infinite. In addition, recall that Assumption 3 ensures that Π is nonempty; our argument below will show that Π^U is nonempty as well. For a given $\pi \in \Pi^U$, denote by $C_U^\pi[a, b]$, with $0 \leq a < b$, the total cost in the interval $[a, b]$ for Model U; the corresponding long-run expected average cost associated with π , denoted by f_U^π , is defined similar to Equation (1).

The rationale behind the definitions in Equations (3) and (4) is as follows. The intent of Equation (3) is simply to ensure that any policy for Model U that takes the inventory trajectory outside the set $B := \{U_L, U_L + 1, \dots, U_H\}$ is excluded from consideration (i.e., is inadmissible). We next turn to Equation (4), which is more intricate. Observe first that the cost of ordering up to either U_H or U_L in Model U from any level below, including in particular $U_L - 1$, is zero; this implies that whenever necessary, the inventory or backorder level in Model U can be *maintained* at U_H or U_L , respectively, with zero ordering costs (by following a “just-in-time” policy). Observe further that whenever the inventory trajectory in Model U meanders within the set $B^* := \{U_L + 1, \dots, U_H - 1\}$, every order incurs the same cost as specified in the original model. Consider now a policy $\pi \in \Pi$ for the original model; then, we claim that, by exploiting the above-noted features of Equation (4) to “truncate” possible excursions of the inventory trajectory of π outside the set B to stay at the boundaries U_H or U_L , we could construct a corresponding policy $\hat{\pi} \in \Pi^U$ for Model U that always maintains the inventory trajectory within the set B with a possibly-lower ordering cost for every order. Moreover, this claim, in conjunction with Equation (3) and Assumption 2, implies that the cumulative holding/shortage cost under π in the original model will also be no less than that of $\hat{\pi}$ in Model U. Thus, we actually have an expanded claim that sums up the purpose of the definitions in both Equations (3) and (4); and this is stated in the following lemma:

LEMMA 1. *For every $\pi \in \Pi$, there exists a corresponding $\hat{\pi} \in \Pi^U$ satisfying $f^\pi \geq f_U^{\hat{\pi}}$.*

Given the above discussion, the formal proof of this lemma is straightforward; it will be given in Appendix A.

We now motivate the remaining steps in our proof. Clearly, Lemma 1 immediately implies that

$$\inf_{\pi \in \Pi} f^\pi \geq \inf_{\pi \in \Pi^U} f_U^\pi. \quad (5)$$

Recall that our aim is to reduce the solution of the original model to that of Model U. With that objective in mind, observe that Inequality (5) is of little use unless the magnitude of $U_H - U_L$ is sufficiently large so that at least one optimal policy in the original model, if one exists, meets the following requirements: (i) It is in Π^U ; (ii) it is also optimal for Model U; and (iii) it never takes the inventory trajectory outside the set B^* . Requirement (i) is obvious. Requirement (ii) is needed because Model U has a different cost structure. Finally, requirement (iii) is mandated because, according to Equation (4), such a

policy will necessarily have identical cumulative cost for both models.

We will actually establish the stronger result that there exists an (s, S) policy that satisfies requirements (i)–(iii). The strategy is to judiciously select a pair of U_L and U_H to ensure that the right-hand side of Inequality (5) is bounded from below by α^* and, furthermore, that α^* is attained by an (s, S) policy in Π^U with $U_L \leq s < S \leq U_H - 1$ (i.e., with its inventory trajectory confined to the set B^*).

Formally, we will prove the following pair of lemmas:

LEMMA 2. *There exist integers U_L and U_H such that an (s, S) policy in Π violating $U_L \leq s < S \leq U_H - 1$ is never optimal for the original model. Moreover, any such policy in Π^U is also not optimal for Model U.*

LEMMA 3. *Consider any U_L and U_H satisfying Lemma 2. Then, the inequality $\inf_{\pi \in \Pi^U} f_U^\pi \geq \alpha^*$ holds; moreover, there exists an (s, S) policy in Π^U whose long-run expected average cost exactly equals α^* .*

Lemma 2 implies that if an optimal (s, S) policy exists in the original model, then that policy is necessarily admissible in Model U with the same long-run expected average cost. Lemma 3 establishes that an (s, S) policy that attains α^* is optimal in Model U; in light of Inequality (5), this policy, namely an (x^*, y^*) policy in Π^U with $\alpha(x^*, y^*) = \alpha^*$, is then also optimal for the original model. Hence, Lemmas 2 and 3, together, imply Theorem 1.

The remainder of this section will consist of two subsections, devoted to the proofs of Lemma 2 and Lemma 3, respectively.

3.1. Proof of Lemma 2

We will consider four possible scenarios: (i) Both z_L and z_H are finite; (ii) both z_L and z_H are not finite; (iii) z_L is finite but $z_H = \infty$; and (iv) z_H is finite but $z_L = -\infty$. For each scenario, we will provide suitable choices for U_L and U_H . The argument is actually the same for both the original model and Model U. We will therefore focus only on the original model.

Scenario (i) is easy. The settings $U_L = \min\{z_L - 1, I_0\}$ and $U_H = z_H + 1$ would satisfy the stipulations in Lemma 2, as the long-run expected average cost of any (x, y) vector with either $x \leq U_L - 1$ or $y \geq U_H$ is infinite.

Next, consider Scenario (ii), where $g(z)$ is finite for all $z \in \mathbb{Z}$. Recall from section 2 that, under Assumptions 3 and 2, there exists a vector (\tilde{x}, \tilde{y}) with $\tilde{x} < 0 \leq \tilde{y}$ and $\tilde{\alpha} := \alpha(\tilde{x}, \tilde{y}) < \infty$. Our choices for U_L and U_H will be tied to the value $\tilde{\alpha}$, and they are defined as follows. Since $\lim_{x \rightarrow \pm\infty} g(x) = \infty$ from

Assumption 2, we can pick integers u_L and u_H that satisfy $u_L < \min\{I_0, \tilde{x}\}$, $u_H > \max\{I_0, \tilde{y}\}$, $g(u_L) > 2\tilde{\alpha}$, and $g(u_H) > 2\tilde{\alpha}$. Let $\Delta := u_H - u_L$; then, we define $U_H := u_H + \Delta$ and $U_L := u_L - \Delta$. Note that the inequalities $U_L < u_L < \tilde{x} < 0 \leq \tilde{y} < u_H < U_H$ hold. These definitions are illustrated in Figure 1.

We now need to show that the above pair of U_L and U_H meets the requirements in Lemma 2. Consider any (x, y) vector with $\alpha(x, y) < \infty$ that violates $U_L \leq x < y \leq U_H - 1$. Then, at least one of the following holds: (a) $x \leq U_L - 1$ and (b) $y \geq U_H$. We claim that for both cases, we must have $\alpha(x, y) > \tilde{\alpha}$ and hence the corresponding (x, y) policy cannot be optimal.

To prove the claim, observe that from Equation (2), we have

$$\alpha(x, y) = \frac{c(y-x) + \sum_{z=x+1}^y g(z)/\lambda}{(y-x)/\lambda} \geq \frac{\sum_{z=x+1}^y g(z)/\lambda}{(y-x)/\lambda} \quad (6)$$


$$= \frac{\sum_{z=x+1}^y g(z)}{y-x},$$

where the inequality is due to the non-negativity of $c(\cdot)$. Now, observe further that whenever z is outside the set $\{u_L+1, \dots, u_H-1\}$, which has cardinality $u_H - u_L - 1$ or $\Delta - 1$, we have $g(z) > \min\{g(u_L), g(u_H)\} > 2\tilde{\alpha}$. It follows that if $g(z) > 2\tilde{\alpha}$ holds for a least half (i.e., no less than a majority) of the z 's in $\{x+1, \dots, y\}$, a set with cardinality $y - x$, then Inequality (6) implies that

$$\alpha(x, y) > \frac{2\tilde{\alpha}(y-x)/2}{y-x} = \tilde{\alpha}. \quad (7)$$

Consider now Case (a) with a given x satisfying $x \leq U_L - 1$. In this case, the possible ranges for y

Figure 1 Selection of u_L, u_H, U_L and U_H



are: $x < y \leq u_L$, $u_L < y < u_H$, or $y \geq u_H$. Note that our definitions of Δ and U_L imply that the cardinality of the set $\{U_L, \dots, u_L\}$ is given by $u_L - U_L + 1 = \Delta + 1$; and that this count is greater than $\Delta - 1$, the cardinality of the set $\{u_L+1, \dots, u_H-1\}$. It is then easily seen from Figure 1 that, for any corresponding choice of y with $y > x$, no less than a majority of the values in the set $\{x+1, \dots, y\}$ lie outside the set $\{u_L+1, \dots, u_H-1\}$. Similarly, a symmetric analysis for Case (b) with a given y satisfying $y \geq u_H$ will yield this same conclusion. Hence, our claim above is a consequence of Inequality (7). This completes the argument for Scenario (ii).

Our choices for U_L and U_H in the remaining two scenarios are similar to those in the first two. Consider Scenario (iii). Let $U_L = \min\{z_L-1, I_0\}$ on the negative side. For the positive side, let u_H be as defined in Scenario (ii) but now let $\Delta := u_H - U_L$. Again, set $U_H := u_H + \Delta$. Then, it should be clear that these choices would meet the requirements in Lemma 2. Finally, Scenario (iv) can be handled in a symmetric manner; and this completes the proof.

3.2. Proof of Lemma 3

Our proof will be based on a further reduction of the solution of Model U to that of a corresponding deterministic model in which every inter-demand time has a constant duration of size $1/\lambda$. We shall refer to this new model as *Model D*. Following our original notation, the n -th demand in Model D, for $n \geq 0$, will then occur precisely at epoch \bar{A}_n . The cost functions (3) and (4) and the set of admissible policies Π^U in Model U will be shared with Model D.

For $0 \leq a < b$, let $C_D^\pi[a, b]$ denote the total costs incurred in Model D over the time interval $[a, b]$ for a given policy $\pi \in \Pi^U$. For Model D, our objective will be to minimize

$$f_D^\pi := \liminf_{T \rightarrow \infty} \frac{C_D^\pi[0, T]}{T} \tag{8}$$

over all $\pi \in \Pi^U$. The reason behind the adoption of \liminf in Equation (8) will become clear later in the proof (see Inequality (17) below).

Paralleling the spirit of Lemma 1, the key step in our argument is the following lemma:

LEMMA 4. *For every $\pi \in \Pi^U$, we have $f_U^\pi \geq f_D^\pi$.*

This lemma implies that the optimal solution of Model D would serve as a lower bound for that of Model U. We will first establish several preliminary results before proving Lemma 4.

For $n \geq 1$, let $X_n := A_n - A_{n-1}$. Recall from Equations (3) and (4) that, for any $\pi \in \Pi^U$ and all $i \geq 0$,

$$C_U^\pi[A_i, A_{i+1}] = \begin{cases} c(q_i^\pi) + g(I_i^\pi)X_{i+1}, & \text{if } U_L < I_i^\pi < U_H, \\ g(I_i^\pi)X_{i+1}, & \text{if } I_i^\pi = U_L \text{ or } I_i^\pi = U_H. \end{cases} \tag{9}$$

Note that the admissibility of π implies that the case for $I_i^\pi = U_L$ or $I_i^\pi = U_H$ in Equation (9) is applicable only when $g(U_L)$ or $g(U_H)$ is finite (cf. the first paragraph in section 3.1). Next, the finiteness of the state space of Model U implies that

$$c_M := \sup_{\pi \in \Pi^U} \sup_{i \geq 0} c(q_i^\pi) < \infty$$

and

$$g_M := \sup_{\pi \in \Pi^U} \sup_{i \geq 0} g(I_i^\pi) < \infty.$$

It then follows from Equation (9) that for any $\pi \in \Pi^U$ and all $i \geq 0$, we have

$$C_U^\pi[A_i, A_{i+1}] \leq c_M + g_M X_{i+1}. \tag{10}$$

For $T > 0$, let $N(T) := \max\{i \geq 0: A_i \leq T\}$. Then, since all costs are non-negative, we have

$$\frac{E[C_U^\pi[0, A_{N(T)}]]}{T} \leq \frac{E[C_U^\pi[0, T]]}{T} \leq \frac{E[C_U^\pi[0, A_{N(T)+1}]]}{T}. \tag{11}$$

Note that, by definition,

$$\begin{aligned} & \frac{E[C_U^\pi[0, A_{N(T)+1}]]}{T} - \frac{E[C_U^\pi[0, A_{N(T)}]]}{T} \\ &= \frac{E[C_U^\pi[A_{N(T)}, A_{N(T)+1}]]}{T}, \end{aligned} \tag{12}$$

and that, in light of Inequality (10),

$$E[C_U^\pi[A_{N(T)}, A_{N(T)+1}]]$$

is bounded from above by $c_M + g_M E[X_{N(T)+1}]$. Hence, a standard result from renewal theory implies that the right-hand side of Equation (12) converges to 0 as $T \rightarrow \infty$; see, e.g., part (ii) of Theorem 3.6.1 in Ross (1996, p. 134) (let $E[R_{N(t)+1}]$ there be $E[X_{N(t)+1}]$). This convergence, together with Inequality (11), then yields the following lemma:

LEMMA 5. *For every $\pi \in \Pi^U$,*

$$f_U^\pi = \limsup_{T \rightarrow \infty} \frac{E[C_U^\pi[0, A_{N(T)+1}]]}{T}. \tag{13}$$

In the next lemma, we relate $E[C_U^\pi[0, A_{N(T)+1}]$, the numerator in Equation (13), to the expectation of a corresponding cumulative cost in Model D that is terminated just before the demand epoch with the random index $N(T) + 1$ from Model U.

LEMMA 6. *For every $\pi \in \Pi^U$,*

$$E[C_U^\pi[0, A_{N(T)+1}]] = E[C_D^\pi[0, \bar{A}_{N(T)+1}]]. \tag{14}$$

PROOF. The argument is a simple adaptation of the proof of Wald’s equation (see, e.g., Theorem 3.3.2 in Ross 1996, p. 105). For $i \geq 1$, let Y_i be the indicator function of the event $\{A_{i-1} \leq T\}$. Then, we have, for any $\pi \in \Pi^U$,

$$\begin{aligned} E[C_U^\pi[0, A_{N(T)+1}]] &= E\left[\sum_{i=1}^{N(T)+1} C_U^\pi[A_{i-1}, A_i]\right] \\ &= E\left[\sum_{i=1}^{\infty} C_U^\pi[A_{i-1}, A_i] Y_i\right] \\ &= \sum_{i=1}^{\infty} E[C_U^\pi[A_{i-1}, A_i] Y_i] \\ &= \sum_{i=1}^{\infty} E[C_U^\pi[A_{i-1}, A_i]] E[Y_i] \\ &= \sum_{i=1}^{\infty} C_D^\pi[\bar{A}_{i-1}, \bar{A}_i] E[Y_i]. \end{aligned} \quad (15)$$

The 3rd equality above is due to monotone convergence (since all costs are non-negative); the 4th is due to the independence of $C_U^\pi[A_{i-1}, A_i]$ and Y_i , as well as the bound (10) which ensures the finiteness of $E[C_U^\pi[A_{i-1}, A_i]]$ for all $i \geq 1$; and the 5th is due to the fact that $E[C_U^\pi[A_{i-1}, A_i]]$ depends on X_i only through its mean (cf. Equation (9)). Clearly, the same argument shows that $E[C_D^\pi[0, \bar{A}_{N(T)+1}]]$ can also be written as Equation (15); and this completes the proof. \square

We are now ready to prove Lemma 4.

PROOF OF LEMMA 4. From Lemma 5 and Lemma 6, we have, for any $\pi \in \Pi^U$,

$$\begin{aligned} f_U^\pi &= \limsup_{T \rightarrow \infty} \frac{E[C_U^\pi[0, A_{N(T)+1}]]}{T} \\ &\geq \liminf_{T \rightarrow \infty} \frac{E[C_U^\pi[0, A_{N(T)+1}]]}{T} \\ &= \liminf_{T \rightarrow \infty} \frac{E[C_D^\pi[0, \bar{A}_{N(T)+1}]]}{T}. \end{aligned} \quad (16)$$

We will complete the proof by showing that the last limit above is bounded from below by f_D^π .

Since the costs are non-negative, the well-known Fatou’s lemma implies that

$$\liminf_{T \rightarrow \infty} E\left[\frac{C_D^\pi[0, \bar{A}_{N(T)+1}]}{T}\right] \geq E\left[\liminf_{T \rightarrow \infty} \frac{C_D^\pi[0, \bar{A}_{N(T)+1}]}{T}\right]. \quad (17)$$

Observe that

$$\liminf_{T \rightarrow \infty} \frac{C_D^\pi[0, \bar{A}_{N(T)+1}]}{T} = \liminf_{T \rightarrow \infty} \frac{C_D^\pi[0, \bar{A}_{N(T)+1}]}{N(T)+1} \frac{N(T)+1}{T}.$$

Since

$$\liminf_{T \rightarrow \infty} \frac{C_D^\pi[0, \bar{A}_{N(T)+1}]}{N(T)+1} = \liminf_{n \rightarrow \infty} \frac{C_D^\pi[0, \bar{A}_n]}{n}$$

and

$$\lim_{T \rightarrow \infty} \frac{N(T)+1}{T} = \lambda$$

with probability 1 (cf. Proposition 3.3.1 of Ross 1996, p. 102), we obtain (noting that $n/\lambda = \bar{A}_n$)

$$\liminf_{T \rightarrow \infty} \frac{C_D^\pi[0, \bar{A}_{N(T)+1}]}{T} = \liminf_{n \rightarrow \infty} \frac{C_D^\pi[0, \bar{A}_n]}{\bar{A}_n}, \quad (18)$$

which is a constant.

Now observe that, for any $T > 0$, we have

$$C_D^\pi[0, \bar{A}_{m(T)}] \leq C_D^\pi[0, T] \leq C_D^\pi[0, \bar{A}_{m(T)+1}],$$

where $m(T)$ denotes the index that satisfies $\bar{A}_{m(T)} \leq T < \bar{A}_{m(T)+1}$. It follows that

$$\begin{aligned} \liminf_{T \rightarrow \infty} \frac{C_D^\pi[0, \bar{A}_{m(T)}]}{\bar{A}_{m(T)}} \frac{\bar{A}_{m(T)}}{T} &\leq \liminf_{T \rightarrow \infty} \frac{C_D^\pi[0, T]}{T} \\ &\leq \liminf_{T \rightarrow \infty} \frac{C_D^\pi[0, \bar{A}_{m(T)+1}]}{\bar{A}_{m(T)+1}} \frac{\bar{A}_{m(T)+1}}{T}. \end{aligned} \quad (19)$$

Clearly, both sides of Inequality (19) converge to the right-hand side of Equation (18). Therefore,

$$f_D^\pi = \liminf_{n \rightarrow \infty} \frac{C_D^\pi[0, \bar{A}_n]}{\bar{A}_n}. \quad (20)$$

Finally, relations (16)–(18) and (20) together yield that $f_U^\pi \geq f_D^\pi$, completing the proof. \square

Our next step is to consider Model D and show that $\inf_{\pi \in \Pi^U} f_D^\pi \geq \alpha^*$. Note that Equation (20) implies that the solution of Model D is equivalent to that of a discrete-time deterministic dynamic program with finite state and action spaces. Hence, we will employ basic concepts from dynamic programming.¹¹

Observe that under the assumption that demands are unit sized, the inventory trajectory dictated by any $\pi \in \Pi^U$ never decreases more than one level at an arrival epoch; that is, it has the so-called “skip-free-to-the-left” property. Together with the fact that Model D has a finite state space, this property implies the following important result.

LEMMA 7. *Every stationary policy in Π^U is an (s, S) policy.*

PROOF. Let $\pi \in \Pi^U$ be a stationary policy in Model D. Recall from section 3.1 that $U_L \leq I_0 < U_H$. Given

this I_0 , suppose the first demand epoch at which a positive order is placed by π has index i_1 . Let x_{i_1} and y_{i_1} , respectively, be the net-inventory levels just before and just after the placement of this order. Clearly, we have $U_L - 1 \leq x_{i_1} \leq I_0$. We will consider two possible scenarios for y_{i_1} : (a) $y_{i_1} \leq I_0 + 1$ or (b) $I_0 + 1 < y_{i_1} \leq U_H$.

For Scenario (a), observe that the skip-free-to-the-left property and the stationarity of π together imply that, beyond epoch \bar{A}_{i_1} , the next positive order under π would occur only when net-inventory decreases to x_{i_1} ; and furthermore, the order size would again equal $y_{i_1} - x_{i_1}$. By iterating this argument, we see that π is an (x_{i_1}, y_{i_1}) policy.

For Scenario (b), we will restart the entire argument at epoch \bar{A}_{i_1+1} by pretending that $y_{i_1} - 1$, which is strictly greater than I_0 , is the “original I_0 .” Again, we will then either conclude that π is an (s, S) policy or advance to the first subsequent demand epoch at which a positive order under π takes the net-inventory level strictly above y_{i_1} . In the latter case, we simply repeat the same argument.

Finally, since the net-inventory level for Model D is bounded from above by U_H , Scenario (b) above can only occur a finite number of times. Hence, π must be an (s, S) policy. \square

Recall from the proof of Lemma 2 that if α^* , which was defined in the original model, is attained by an (x, y) vector, then we must have $U_L \leq x < y \leq U_H - 1$. The fact that the total number of (x, y) vectors satisfying this requirement is finite then implies that, indeed, there exists a vector (x^*, y^*) satisfying $U_L \leq x^* < 0 \leq y^* \leq U_H - 1$ and $\alpha(x^*, y^*) = \alpha^*$. Next, for the given I_0 , recall further that Assumption 3 implies that a finite number of positive orders can be prescribed, whenever $I_0 < x^*$, to take the net-inventory beyond level x^* at a finite cost. Therefore, at least one instance of the corresponding (x^*, y^*) policy is admissible in Model D; our next lemma strengthens this result to a partial converse of Lemma 7.

LEMMA 8. *A stationary instance of the (x^*, y^*) policy exists in Π^U .*

PROOF. Obviously, we only need to address the case with $I_0 < x^*$. The four scenarios in the proof of Lemma 2 are again relevant in our argument. We will first handle Scenarios (ii) and (iii). For these scenarios, we have $g(U_H) < \infty$. Note that, since $I_0 < x^* < 0 \leq U_H$, the difference $U_H - I_0$ is positive. Hence, the placement of an order of size $U_H - I_0$, which incurs zero cost (see Equation (4)), will immediately take the net-inventory beyond level x^* . This clearly then yields a stationary instance of the (x^*, y^*) policy in Π^U .

We next consider the remaining two scenarios where $g(U_H) = \infty$ and hence it is inadmissible to order up to U_H . To overcome this complication, the key idea is to construct a stationary ordering sequence (i.e., a sequence of order quantities that are specified by a deterministic function of the net-inventory levels just before ordering) that would also take the net-inventory level beyond x^* but now no greater than either 0 or z_H . For this purpose, it will be expedient to define and work with a set of “ladder steps”: Let $L_0 := 0$; and define $L_j := L_{j-1} - (q_L - 1)$ for $j \geq 1$. Note that $q_L - 1 \geq 1$ holds by definition (see Assumption 3); hence, the L_j 's are distinct.

We will begin with the simplest case with $L_1 \leq I_0 < L_0$. Observe that if $z_L \leq L_1$, then we can do nothing until the net-inventory level drops down to $L_1 - 1$ and then place an order of size q_L to take it back up to L_0 ; since $x^* < L_0$, we are done with the construction. Suppose on the other hand that $z_L > L_1$; this implies that we must have $I_0 \geq z_L - 1$. We could then either order q_L units immediately (if $I_0 = z_L - 1$) or wait (if $I_0 > z_L - 1$) until net-inventory drops down to level $z_L - 1$ and then order q_L units; in either case, we will end up with a net-inventory level of $z_L - 1 + q_L$. Note that Assumption 3 ensures that $z_L - 1 + q_L \leq z_H$ holds; hence, these actions are admissible. Since $x^* < L_0 < 1 \leq z_L - 1 + q_L$, we are again done. We will next move on to other cases, where I_0 is lower than L_1 .

Let $J := \sup \{j \geq 0: L_j \geq z_L\}$; note that J is not necessarily finite. Suppose $L_k \leq I_0 < L_{k-1}$ for some $J \leq k \leq 2$ (the case with $k = 1$ has already been dealt with above). Then, we will let net-inventory decrease to level $L_k - 1$ and then place an order of size q_L to take it back up to level L_{k-1} . If $L_{k-1} > x^*$, then we are done. Otherwise, we can simply repeat similar actions until the net-inventory level after the placement of a size- q_L order exceeds x^* for the first time; this, then, completes the construction. Note that the inventory trajectory ascends by consecutively stepping up to the “next L_j ” until a sufficient height is reached.

Finally, suppose J is finite; then, it is easily seen that the only remaining possibility satisfies the conditions $J \geq 2$ and $L_{J+1} \leq z_L - 1 \leq I_0 < L_J$. For this case, we will order q_L units immediately. If the resulting net-inventory level (which overshoots L_J but not L_{J-1}) exceeds x^* , we are done; otherwise, we continue as per the construction in the last paragraph, lowering the net-inventory level to either $L_{J-1} - 1$ (if $I_0 = L_J - 1$) or $L_J - 1$ (if $I_0 < L_J - 1$) first and then move further up the “ladder.” This completes the proof. \square

Lemma 8 and Lemma 7 now together imply that a stationary instance of the (x^*, y^*) policy is optimal

within the class of all *stationary* policies in Π^U . In our final step, we will further strengthen this result to *within all* policies in Π^U by invoking standard results from dynamic programming.

LEMMA 9. *A stationary instance of the (x^*, y^*) policy is optimal within Π^U .*

PROOF. When Model D is treated as a DP, the standard setup is to have an arbitrary initial state¹² (i.e., not limited to I_0). The proof of Lemma 8 actually shows that for *any* initial state, a stationary instance of the (x^*, y^*) policy exists in Π^U . Hence, the optimal average cost over the class of stationary policies in Π^U is equal to α^* for all initial states. It then follows from Proposition 4.2.4 in Bertsekas (2001, p. 203) that a solution to the standard “average-cost optimality equations” exists. Proposition 4.2.1 in the same reference (cf. p. 199) then implies that an optimal stationary policy exists in Π^U ; and hence, any stationary instance of the (x^*, y^*) policy in Π^U is optimal for Model D. \square

Lemma 9 shows that any admissible instance of the (x^*, y^*) policy (stationary or not) is in fact optimal for Model D. It then follows from Lemma 4 that $\inf_{\pi \in \Pi^U} f_U^\pi \geq \inf_{\pi \in \Pi^U} f_D^\pi \geq \alpha^*$, where α^* is attained by any instance of the (x^*, y^*) policy in Π^U . This completes the proof of Lemma 3, and of Theorem 1 as well.

4. Optimality of (s, S) Policies with Constant Lead Times

Our original inventory model in section 2 assumes that the supply is instantaneous, i.e., the lead time is zero. In this section, we show that when the inter-arrival times are assumed to be exponential (i.e., demands arrive according to a Poisson process), our (s, S)-optimality result can be extended to cover the case with a positive constant replenishment lead time τ .

In this new setting, denote by $\eta^\pi(t)$ the net-inventory level at time t under a policy π ; and let $\hat{\eta}^\pi(t)$ be the corresponding inventory position, defined as the sum of $\eta^\pi(t)$ and the total amount of inventory in the pipeline at time t . An (s, S) policy is now defined to be one that raises the *inventory position* to S every time it decreases to s .

We use the convention that $\eta^\pi(t)$ is a right-continuous step function with jumps only at demand arrival epochs and order-delivery epochs (which occur after a constant time delay τ following the demand epochs). That is, with $\eta^\pi(A_0-) := I_0$, we have

$$\eta^\pi(A_n) = \eta^\pi(A_n-) - 1 \quad \text{for all } n \geq 1;$$

and

$$\eta^\pi(A_n + \tau) = \eta^\pi((A_n + \tau)-) + q_n^\pi \quad \text{for all } n \in \mathbb{N}^0.$$

Similarly, $\hat{\eta}^\pi(t)$ is a right-continuous step function where jumps occur only at demand arrival epochs. That is, with $\hat{\eta}^\pi(A_0-) := I_0$ and $\hat{\eta}^\pi(A_0) = \hat{\eta}^\pi(A_0-) + q_0^\pi$, we have

$$\hat{\eta}^\pi(A_n) = \hat{\eta}^\pi(A_n-) - 1 + q_n^\pi \quad \text{for all } n \geq 1.$$

It is well known that the two processes $\{\eta^\pi(t) : t \geq 0\}$ and $\{\hat{\eta}^\pi(t) : t \geq 0\}$ are related through the identity (see, e.g., section 6.2 in Zipkin 2000)

$$\eta^\pi(t) = \hat{\eta}^\pi(t - \tau) - \delta(t - \tau, t] \quad \text{for all } t \geq 0,$$

where $\delta(t - \tau, t]$ is the total demand in $(t - \tau, t]$.

Adopting standard inventory-theoretic convention (see Zipkin 2000), we now define $C_\tau^\pi[a, b]$ as the sum of the ordering costs incurred in the time interval $[a, b)$ and the holding/shortage costs incurred in the time interval $[a + \tau, b + \tau)$. Our new problem (with a lead time of $\tau > 0$) is then to minimize

$$f_\tau^\pi := \limsup_{T \rightarrow \infty} \frac{E[C_\tau^\pi[0, T])}{T}$$

over all $\pi \in \Pi$. This problem is mathematically identical to the inventory problem of section 2 (with no lead time), i.e., the problem defined by Equation (1), except for the following:

- (I) For all $n \in \mathbb{N}^0$, the inventory positions $\hat{\eta}^\pi(A_n)$ and $\hat{\eta}^\pi(A_n-)$, respectively, take the roles played by I_n^π and $I_{n-1}^\pi - 1$ (with $I_{-1}^\pi := I_0$) for the original model in section 2.
- (II) The holding and shortage cost rate at any instant is now given by the function $g_\tau(\hat{\eta})$ if $\hat{\eta}$ is the inventory position at that instant, where

$$g_\tau(\hat{\eta}) := E[g(\hat{\eta} - Q)]$$

and Q is a Poisson random variable with mean $\lambda\tau$.

Next, we will revisit Assumption 2 in section 2. That assumption stipulated that the function $g(\cdot)$ needs to satisfy the following properties: (a) $\lim_{y \rightarrow \pm\infty} g(y) = \infty$, and (b) $g(\cdot)$ is quasi-convex with zero as a minimizer (i.e., $g(0) \leq g(y)$ for all $y \in \mathbb{Z}$). (Note that $g(0) < \infty$ is not required for our analysis; see Footnote 3.) However, it is easy to verify that the main result of section 2, namely Theorem 1, holds when (b) is replaced by the weaker assumption that $g(\cdot)$ is quasi-convex with a minimizer z^* , for some $z^* \in \mathbb{Z}$. This fact, together with (II) above, implies that to establish the optimality of an (s, S) policy for the case with a positive lead time, we now need to show that (a) $\lim_{\hat{\eta} \rightarrow \pm\infty} g_\tau(\hat{\eta}) = \infty$, and (b) $g_\tau(\cdot)$ is a

quasi-convex function with a finite minimizer ζ^* ; the following lemma verifies these two conditions.

LEMMA 10. *The function $g_\tau(\cdot)$ satisfies the following properties: (a) $\lim_{\hat{\eta} \rightarrow \pm\infty} g_\tau(\hat{\eta}) = \infty$; and (b) $g_\tau(\cdot)$ is a quasi-convex function with a finite minimizer ζ^* .*

PROOF. (a) Observe that $E[g(\hat{\eta} - Q)] \geq \mathbb{P}(Q = 0)g(\hat{\eta})$. Then, we have

$$\begin{aligned} \lim_{\hat{\eta} \rightarrow \infty} g_\tau(\hat{\eta}) &= \lim_{\hat{\eta} \rightarrow \infty} E[g(\hat{\eta} - Q)] \\ &\geq \mathbb{P}(Q = 0) \lim_{\hat{\eta} \rightarrow \infty} g(\hat{\eta}) \\ &= \infty, \end{aligned}$$

where the last equality is due to Assumption 2. A similar argument yields $\lim_{\hat{\eta} \rightarrow -\infty} g_\tau(\hat{\eta}) = \infty$.

(b) The proof is based on Keilson and Gerber (1971). In particular, their statement S_3 and the proof of Theorem 3(a) (in their paper) imply the following: If l is a quasi-convex function and Q is a Poisson random variable, then $E[l(x - Q)]$ is a quasi-convex function of x . (See also the discussion on page 1008 in Rosling (2002) and Lemma 5.4 in Huh et al. (2011).) Part (b) then follows immediately as $g(\cdot)$ is quasi-convex and Q is a Poisson random variable. \square

5. Concluding Remarks

This section is organized as follows. In section 5.1, we complement the opening paragraphs of section 1 with a further discussion of the relevant history of the (s, S) -optimality problem in inventory theory; this sets the stage for an ensuing summary of the contribution of our paper relative to prior work. In section 5.2, we comment on recent parallel developments under the alternative stochastic framework where demand is governed by a Brownian motion. We conclude in section 5.3 with suggestions for future research.

5.1. Motivation and Contribution of the Paper

Ever since the influential paper of Scarf (1960), the question of “when and why an (s, S) policy is optimal” has been a central research concern in inventory theory. The model examined by Scarf assumes periodic review. Demands are IID and all costs are linear, except for a fixed setup cost for ordering. It is very interesting to note that Scarf himself has wondered about the answer to the above question in his setting. This was mentioned in the Introduction of Porteus (1971) with the following quote from Scarf (1963):

This type of cost function has appeared very frequently in the literature of inventory theory not necessarily because of its realism, but because it

provides one of the few examples of cost functions with a decreasing average cost for which the analysis of inventory policies is relatively easy.

Perhaps motivated by its simplicity, the same cost assumptions were also made in the classical EOQ model of Harris (1913). Perera et al. (2017) have recently revisited Harris’s model. Their work was precisely motivated by the same when/why question above; specifically, see the opening paragraph of that paper. To investigate the answer to this question, Perera et al. (2017) worked with an extended EOQ framework where the forms of the cost functions are left unspecified. While this apparently is the simplest and cleanest setup, it is interesting to note that only a few isolated (s, S) -optimality results under cost assumptions weaker than those in Scarf (1960) exist in the literature. Moreover, even under the standard linear-cost assumptions, simple proofs of (s, S) -optimality are not readily available. Nevertheless, Perera et al. (2017) established (in their Corollary 1) a simple necessary and sufficient condition¹³ for (s, S) -optimality under their general, but *deterministic*, framework. Their proof is constructive and elementary.

To push the above important theme further, it is then natural to ask: Can the results in Perera et al. (2017) be *fully* extrapolated to an appropriate setting with stochastic demands. Clearly, stochastic demands would add much greater realism, which is needed to support applications. Three possible settings immediately emerge as potential candidates. The first is the standard periodic-review model of Scarf (1960), where successive demands are IID random variables; the second is the renewal-demand model in this paper; and the third is the Brownian-demand model initiated in a pioneering paper by Bather (1966). We will discuss the first two in this subsection and defer our comments on the third to the next subsection.

As noted in the quoted passage above, Scarf himself pointed out that it is unlikely for (s, S) -optimality to prevail in his setting when cost functions are allowed to assume more general forms. Indeed, Porteus (1971) showed that when the ordering-cost function is concave and increasing, an ordinary (s, S) policy need not be optimal. This result appears to be in direct “conflict” with that in Perera et al. (2017); and hence it could be puzzling at first sight. However, observe that the inventory trajectory in the EOQ setting does not have any downward jumps; that is, it is skip-free. In contrast, observe further that this property typically does not hold in Scarf’s setting (unless the demands are Bernoulli¹⁴). This suggests that, for (s, S) -optimality to prevail, it is essential to preserve the skip-free property. Hence, the periodic-review model is not amenable for a stochastic extension in

the spirit of Perera et al. (2017). (It of course continues to be a very reasonable framework for inventory analysis.)

The second candidate which is that demands arrive according to a renewal process is adopted in our paper; we also assume that demands are unit-sized, which is now seen as *necessary* to preserve the skip-free property in models with an integer-valued state space. Our choice of this demand model is also motivated by the fact that the deterministic demand process (in the EOQ model) is commonly taken in the stochastic-models (especially queueing) literature as the fluid limit of a properly-scaled sequence of renewal processes (with increasing arrival rates and correspondingly decreasing jump sizes). In other words, these two demand models are “natural” counterparts of one another. Indeed, with the renewal-demand assumption, we are able to establish that (s, S) -optimality continues to prevail with virtually no assumptions on the cost functions; and this fully extends the strong conclusion in Perera et al. (2017) to an important, primitive stochastic setting.

It is noteworthy that while the essential “insight” of exploiting the skip-free property is common to Perera et al. (2017) and the present paper, the constructions of the proper proofs are still rather non-trivial in both cases. The proof in Perera et al. (2017) relied on path-wise cost dominance of inventory trajectories as well as on carefully lower-bounding the cumulative total costs based on a notion of ordering cycles. Our proof here relied instead on a novel reduction of the problem to one with uniformly-bounded inventory trajectories. This reduction allowed us to constructively exploit the skip-free property (see Lemma 7) so as to further link the problem to standard discrete dynamic programming results. Thus, the methods of proof are distinctly different. In particular, we note that the path-wise cost comparisons in Lemma 1 of Perera et al. (2017), which implicitly depend on the skip-free property, do not have an obvious counterpart in the stochastic setting here even after our reduction of the problem into one for the deterministic Model D (*cf.* Footnote 11).

We conclude this subsection with a few remarks on the practicality of our demand model. The assumption of IID inter-demand times is reasonable when the pool of potential buyers in the entire population is large. It is particularly applicable in the now-prevalent platform of online retailing where buyers are mostly individual consumers. It is also a fairly standard assumption in the inventory literature. Beckmann (1961) is a notable early example; and other references include: Finch (1961), Rubalskiy (1972a,b), Sivazlian (1974), Tijms (1972), Sahin (1979, 1983), and Federgruen and Schechner (1983). The special case of a Poisson process is popular; see chapter 6 of Zipkin

(2000). Price-dependent Poisson demands have also been assumed in revenue-management settings;¹⁵ see, for example, Gallego and van Ryzin (1994). Finally, we believe the assumption of unit-sized demands is quite reasonable for durable goods.

5.2. Comparison with the Brownian-Demand Model

The assumption that demand is driven by a Brownian motion is a direct stochastic generalization of the deterministic demand in the EOQ model. This is because when the variance parameter σ in a Brownian motion is set to zero, the resulting cumulative demand process is a deterministic line with a slope given by the arrival/drift rate. However, it is important to note that when σ is strictly positive, the *cumulative* demand process in this setting could potentially *decrease* by any amount for any given time interval of positive duration, and furthermore that such decrements are completely independent of orders placed on the “actual supplier” of a given product. In other words, the Brownian-demand assumption implicitly features a second “hidden supplier” who is *constantly and randomly* injecting inventory in the background, independent of the actions of the actual supplier. Hence, it is arguable that the Brownian motion is a suitable choice for modeling continuous demand processes in an inventory-control setting. Indeed, it is explicitly noted in Bather (1966, p. 539) that:

... , we may hope that our model will lead to useful results provided that, in the final analysis, the policy determined is such that restocking occurs relatively infrequently. Alternatively, the present model can be regarded as a natural generalization of the deterministic inventory with fixed demand μ per unit time, where we represent only the variances of random demands.

Thus, Bather is fully aware of the difficulties with interpreting his optimality result in a realistic *inventory-control* setting; and this is due to the fact that the Brownian motion could serve as a reasonable approximation of a deterministic cumulative demand process *only* when σ is extremely small relative to the demand rate μ .

In a recent paper, He et al. (2017) (also see Yao et al. (2015, 2017) for related results) worked with Brownian demand and extended Bather’s result by showing that (s, S) -optimality prevails under a host of weaker assumptions on the cost structure; specifically, see Conditions (S1)–(S4) and Conditions (H1)–(H5) in that paper. Loosely speaking, their cost assumptions are as follows: (a) The ordering-cost function has a setup component and a linear component, where the setup component is allowed to be a fairly general bounded

lower-semicontinuous function of the order quantity; and (b) the holding/shortage cost function is strictly convex, reasonably smooth and polynomially bounded.

Note that the Brownian-demand model also preserves the skip-free property. In light of our discussion in section 5.1, it is then tempting to conjecture that (s, S) -optimality might again prevail under conditions similar to those in Perera et al. (2017). However, the required conditions in He et al. (2017) are easily seen to be much stronger. As a quick example, the simple and yet practical fixed-plus-concave ordering cost function discussed in Porteus (1971) is not readily covered by their assumptions; furthermore, any constraints on inventory level (e.g., the practical scenario of having a concave holding-cost function up to a given warehouse capacity) and/or on order quantity (e.g., the common vendor policies of batch ordering, minimum-order-quantity, or maximum-order-quantity) are also not covered. The origin of this significant discrepancy, which exists regardless of whether or not σ is close to 0 (i.e., whether the Brownian motion can serve as a reasonable approximation to deterministic demand), is not apparent. One possible reason is that the existence of an extra hidden supplier is not consistent with the fundamental physics of the deterministic EOQ setting. Another is that the current tool set in the Brownian-control area is not yet sufficient for a further weakening of their assumptions to the same extent as those in Perera et al. (2017) or the present paper; if this is the case, it would be of interest to see any new progress in the development of mathematical tools for the control of Brownian motion beyond what has been accomplished in He et al. (2017).

5.3. Future Work

As discussed in sections 1 and 5.1, almost all (s, S) -optimality results in the extant literature require the standard linear-cost assumption. Hence, it has been the norm for many authors to simply *assume* that an (s, S) policy is optimal when the cost functions are more general. However, even under this simplifying assumption, which is now fully supported theoretically for the deterministic-demand setting (Perera et al. 2017) and for the renewal-demand setting (the present paper), the ensuing task of identifying an optimal policy is still by no means obvious for complicated cost structures.

For the renewal-demand setting, we note that the proof of our Lemma 2 provides a procedure for the selection of an explicit pair of U_L and U_H for a given $g(\cdot)$. That is, the search for an optimal (s, S) vector has actually been reduced to one over a finite set of s and S values. Whether or not further simplifications within this finite set exist will depend, in addition, on the specific structures of the ordering-cost function $c(\cdot)$ as

well as the average-cost function $\alpha(\cdot, \cdot)$ defined in Equation (2); and this needs to be examined on a case-by-case basis.

As an example, consider the case of having a positive minimum order quantity q_m , where the cost of ordering q units is infinite when $0 < q < q_m$ but is equal to $K + vq$ (with positive constants K and v) for $q \geq q_m$. For this scenario, it appears that the algorithm discussed in Federgruen and Zheng (1992) could be adapted to compute the optimal values of s and S . The details, however, are somewhat involved even for this simple example. Therefore, the computation of an optimal (s, S) vector for specific, interesting cost structures is a useful area for future work.

To facilitate analysis, we have adopted the assumption that orders can be placed only at arrival epochs. To the best of our knowledge, all papers with a renewal demand cited in the last paragraph of section 5.1 also make the same assumption. This is a helpful theoretical simplification because knowledge of the net-inventory level at an *arbitrary* time epoch is, in general, not sufficient to determine the stochastic law that governs the future inventory trajectory. Fortunately, this assumption is readily justifiable in practice because it would greatly simplify the logistics of inventory management. (Of course, the assumption of periodic-view is also rather practical.) In general, a formulation that includes the elapsed time since the last arrival as part of the state definition might be needed to answer the question of whether or not this assumption can be made without loss of optimality (within the larger class of policies that allow restocking at any epoch). The resulting mathematical complexity is a substantial challenge. For the Poisson case, which has the memoryless property, we conjecture that there is no loss of optimality by limiting attention to arrival epochs only. A formal proof even for this much simpler case, however, remains open.

Acknowledgments

The authors sincerely thank the departmental editor Qi Annabelle Feng, the senior editor and two anonymous reviewers for their valuable comments and suggestions, and also for the prompt review process.

Appendix A: Proof of Lemma 1

Recall that the given initial state I_0 is assumed to be in the set $\{U_L, \dots, U_H - 1\}$. Clearly, if $g(U_L) = \infty$ and/or $g(U_H) = \infty$, then the inventory trajectory of an admissible policy in the original model will never down-cross level $U_L + 1$ and/or up-cross level $U_H - 1$. It is therefore sufficient to consider the case where both $g(U_L)$ and $g(U_H)$ are finite.

Let π be a policy in Π . If the inventory trajectory under π always stays within set B , then π is also in Π^U .

We will simply let $\hat{\pi} = \pi$. Now, the only difference between the original model and Model U is that whenever π places an order that takes the net-inventory up to either level U_L (necessarily from level $U_L - 1$) or level U_H , the ordering cost in Model U will be zero by definition (cf. Equation (4)). Since such ordering costs are never greater than the corresponding ones in the original model, it is immediate that $f^\pi \geq f_U^\pi$ holds.

Next, suppose that the inventory trajectory under π has at least one excursion outside the set B . Such an excursion necessarily begins with either (i) an up-crossing of level U_H or (ii) a down-crossing of level U_L . Case (i) could occur whenever the net-inventory is no higher than U_H and a positive order is placed at the ensuing arrival epoch; and Case (ii) occurs whenever the net-inventory drops down to level $U_L - 1$ due to a demand arrival but a positive order is not placed.

With initial state I_0 , let $A_j, j \geq 0$, be the epoch at which the inventory trajectory first exits set B under π . Let $\hat{\pi}$ be identical to π prior to this epoch.¹⁶ For both cases above, we will show how to construct $\hat{\pi}$ from A_j onward to meet the requirement of the lemma. Our prescription will be for the first exit only. Further exits, if any, are handled in the same manner.

Suppose the exit at A_j is due to an up-crossing of level U_H ; that is, Case (i) applies. Let x_j^π be the net-inventory right before the placement of the order at A_j ; then, we have $U_L - 1 \leq x_j^\pi < U_H$ and $x_j^\pi + q_j^\pi > U_H$. At A_j , we will let $\hat{\pi}$ place an order of size $q_j^{\hat{\pi}} = U_H - x_j^\pi$ so that the net-inventory under $\hat{\pi}$ lands exactly at level U_H . Beyond A_j , there are two scenarios for π ; the inventory trajectory either (i.a) never falls back to level U_H again or (i.b) returns to level U_H at a future arrival epoch A_k , where $k > j$. For Scenario (i.a), we will let $\hat{\pi}$ follow the just-in-time policy—a policy that orders whenever a demand occurs—beyond A_j ; that is, let $q_n^{\hat{\pi}} = 1$ for all $n > j$. For Scenario (i.b), we will let $\hat{\pi}$ follow the just-in-time policy up to epoch A_k (i.e., let $q_n^{\hat{\pi}} = 1$ for all $j < n \leq k$) and then follow the same sequence of actions dictated by π until the next exit from set B occurs, if it exists. Scenarios (i.a) and (i.b) are illustrated in Figures A1 and A2, respectively.

Suppose on the other hand that Case (ii) applies at epoch A_j . Note that, by definition, π does not place a positive order when net-inventory decreases from U_L to $U_L - 1$ at A_j ; instead, we will let $\hat{\pi}$ place an order of size one at this epoch to bring the net-inventory back up to level U_L . Beyond A_j , we will again have two scenarios under π ; the inventory trajectory either (ii.a) never up-crosses level $U_L - 1$ or (ii.b) is elevated to level y_l^π at arrival epoch A_l , where $y_l^\pi \geq U_L$ and $l > j$. For Scenario (ii.a), we will let $\hat{\pi}$ follow the just-in-time policy. For Scenario (ii.b), let $\hat{\pi}$ also follow this policy

prior to A_l and then place an order at A_l to take the net-inventory up to level $\min\{y_l^\pi, U_H\}$. If $y_l^\pi \leq U_H$, let $\hat{\pi}$ mirror π beyond A_l until another exit of set B occurs, if it exists; otherwise, let $\hat{\pi}$ follow the prescription in

Figure A1 Construction of $\hat{\pi}$ for Scenario (i.a)

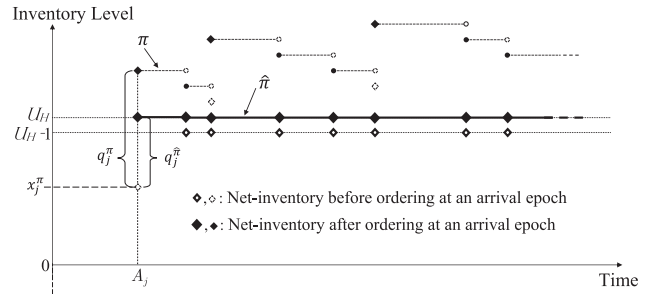


Figure A2 Construction of $\hat{\pi}$ for Scenario (i.b)

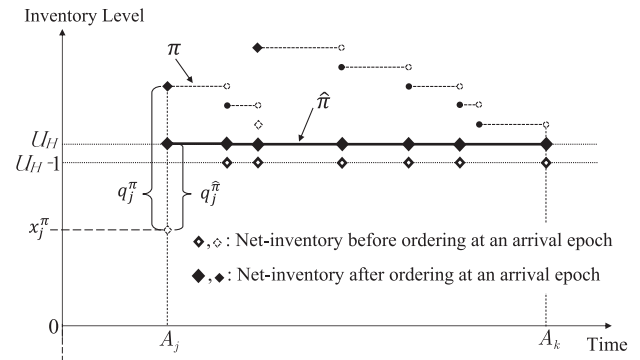


Figure A3 Construction of $\hat{\pi}$ for Scenario (ii.a)

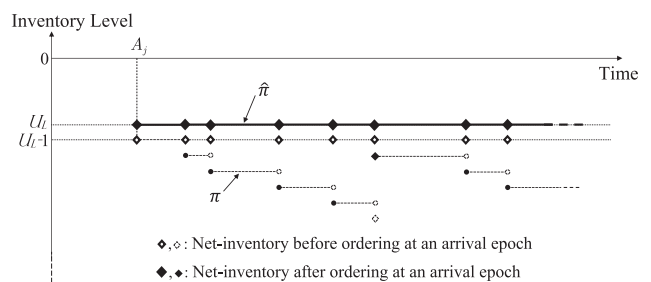


Figure A4 Construction of $\hat{\pi}$ for Scenario (ii.b) with $y_l^\pi \leq U_H$

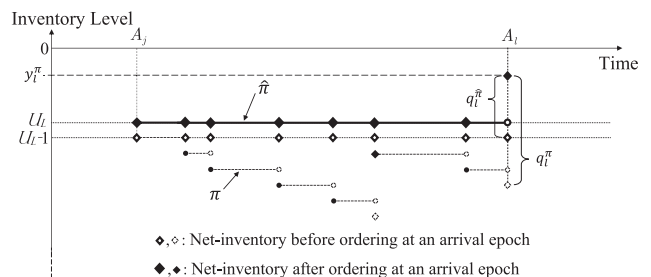
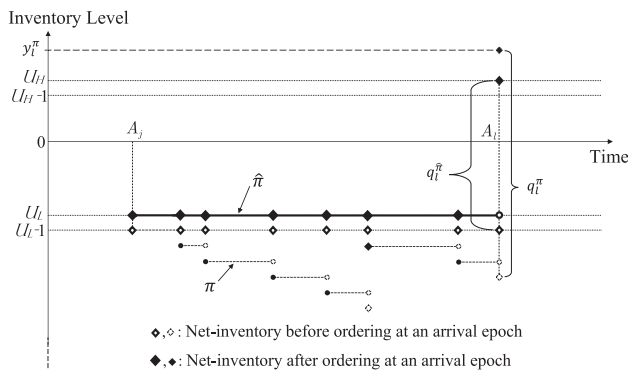


Figure A5 Construction of $\hat{\pi}$ for Scenario (ii.b) with $y_i^\pi > U_H$ 

Case (i), as we have an exit of set B at epoch A_j . Scenarios (ii.a) and (ii.b) are illustrated in Figures A3–A5.

Clearly, the above construction yields a policy $\hat{\pi}$ whose inventory trajectory always stays within set B ; that is, we have $\hat{\pi} \in \Pi^U$. We will next compare $C_{U^L}^{\hat{\pi}}[0, T)$ against $C^\pi[0, T)$ for any $T > 0$. From Figures A1–A5, we see that the on-hand inventory and the backorder levels under policy $\hat{\pi}$ are never greater than that under π ; hence, Assumption 2 implies that the cumulative holding and backordering costs under $\hat{\pi}$ are no greater than that under π within any time interval. Furthermore, whenever $\hat{\pi}$ places an order at an arrival epoch where the inventory trajectories under $\hat{\pi}$ and π are not in agreement, the cost of that order is zero. It follows that $C^\pi[0, T) \geq C_{U^L}^{\hat{\pi}}[0, T)$ holds for all $T > 0$. Dividing both sides of this inequality by T and letting $T \rightarrow \infty$ then yields $f^\pi \geq f_{U^L}^{\hat{\pi}}$; this proves Lemma 1.

Notes

¹This was brought to our attention while our work was in progress.

²A real-valued function f on the set of integers, \mathbb{Z} , is quasi-convex if there exists $z^* \in \mathbb{Z}$ such that $g(\cdot)$ is non-increasing on $\{\dots, z^*-2, z^*-1, z^*\}$ and non-decreasing on $\{z^*, z^*+1, z^*+2, \dots\}$.

³While $c(0) = 0$ is assumed, a finite $c(0)$ is actually sufficient in our analysis. Moreover, $g(0) < \infty$ is not explicitly needed for (s, S) -optimality; this is because if $g(0)$ is not finite, then any (s, S) policy will be optimal with an infinite average cost.

⁴The adoption of an open right endpoint in $C^\pi[0, T)$ here is only for convenience. This can be argued as follows. Observe that the inequality $C^\pi[0, T) \geq C^\pi[0, T)$ holds, and that these costs could differ only if there is an order at epoch T . Note that, in general, we do not have a uniform upper bound on this potential cost difference for an arbitrary policy π . However, it follows from the above inequality that if an (s, S) policy is optimal with respect to the choice $C^\pi[0, T)$, which is established in this paper, then, the same policy is also optimal with respect to $C^\pi[0, T]$.

⁵For a simple example, let: $I_0 = -1$; $c(q) < \infty$ holds only for $q = 0$ and for $q \geq 4$; and $g(z) = \infty$ for all $|z| \geq 2$.

⁶See Assumption 3 below, as well as related discussions surrounding that assumption.

⁷It is important to note that we do not assume that the ordering costs are *finite* for all order sizes; see Assumption 1. Hence, starting from an arbitrarily given I_0 , it is not guaranteed that a single admissible order could be placed to take the net-inventory up to level S , for every $S > I_0$.

⁸If $I_0 \geq 0$ with $g(I_0) < \infty$, then the $(-1, 0)$ policy is optimal; and if $I_0 < 0$ with $g(I_0 + 1) < \infty$, then the $(I_0, I_0 + 1)$ policy is optimal.

⁹If $z_H - z_L + 1 = q_L$, then a single order is sufficient. Otherwise, wait until the net-inventory drops down to level z_L and then place another order, resulting in a net increment of size one beyond the level right after the placement of the first order. Repeat this process if necessary.

¹⁰For the scenario with $z_H - q_L < I_0 < z_H < \infty$, let net-inventory drop down to level $z_H - q_L$ and then place an order of size q_L . Otherwise, place an order of size q_L immediately and then follow a prescription similar to Footnote 9.

¹¹Note that the analysis for the deterministic EOQ setting in Perera et al. (2017) is based on the idea of decomposing the total cost over an interval into those in successive “ordering cycles.” It turns out that a parallel argument does not apply in our setting here. To see this, recall that the ordering costs in Model U and hence Model D are state dependent. This implies that Lemma 1 in Perera et al. (2017) does not have an obvious counterpart in Model D.

¹²For given U_L and U_H , the boundaries of the set of possible initial states differ slightly for the four scenarios stated at the beginning of section 3.1. However, this has no bearing on our argument.

¹³An (s, S) policy is optimal if and only if the average-cost function $\alpha(x, y)$ over all (x, y) policies has a minimizer.

¹⁴Noted by a reviewer of this paper.

¹⁵Brought to our attention by a reviewer of this paper.

¹⁶This statement is applicable only if $j > 0$.

References

- Alp, O., W. T. Huh, T. Tan. 2014. Inventory control with multiple setup costs. *Manuf. Serv. Oper. Manag.* **16**(1): 89–103.
- Altintas, N., F. Erhun, S. Tayur. 2008. Quantity discounts under demand uncertainty. *Management Sci.* **54**(4): 777–792.
- Bather, J. A. 1966. A continuous time inventory model. *J. Appl. Prob.* **3**(2): 538–549.
- Beckmann, M. 1961. An inventory model for arbitrary interval and quantity distributions of demand. *Management Sci.* **8**, 35–57.
- Benkherouf, L., A. Bensoussan. 2009. Optimality of an (s, S) policy with compound Poisson and diffusion demands: A quasi-variational inequalities approach. *SIAM J. Cont. Optim.* **48**(2): 756–762.
- Bensoussan, A., R. H. Liu, S. P. Sethi. 2005. Optimality of an (s, S) policy with compound Poisson and diffusion demands: A quasi-variational inequalities approach. *SIAM J. Cont. Optim.* **44**(5): 1650–1676.
- Benton, W. C., S. Park. 1996. A classification of literature on determining the lot size under quantity discounts. *Eur. J. Oper. Res.* **92**(2): 219–238.
- Bertsekas, D. P. 2001. *Dynamic Programming and Optimal Control*, Vol. 2, 2nd edn. Athena Scientific, Belmont, MA.
- Beyer, D., S. P. Sethi. 1998. A proof of the EOQ formula using quasi-variational inequalities. *Int. J. Syst. Sci.* **29**(11): 1295–1299.

- Beyer, D., F. Cheng, S. P. Sethi, M. I. Taksar. 2010. *Markovian Demand Inventory Models*. Springer, New York.
- Caliskan-Demirag, O., Y. Chen, Y. Yang. 2012. Ordering policies for periodic-review inventory systems with quantity-dependent fixed costs. *Oper. Res.* **60**(4): 785–796.
- Chan, L. M. A., A. Muriel, Z. Shen, D. Simchi-Levi. 2002. On the effectiveness of zero-inventory-ordering policies for the economic lot-sizing model with a class of piecewise linear cost structures. *Oper. Res.* **50**(6): 1058–1067.
- Chao, X., P. Zipkin. 2008. Optimal policy for a periodic-review inventory system under a supply capacity contract. *Oper. Res.* **56**(1), 59–68.
- Chen, F. 2000. Optimal policies for multi-echelon inventory problems with batch ordering. *Oper. Res.* **48**(3): 376–389.
- Chen, X. 2009. Inventory centralization games with price-dependent demand and quantity discount. *Oper. Res.* **57**(6): 1394–1406.
- Constantinides, G. M., S. F. Richard. 1978. Existence of optimal simple policies for discounted-cost inventory and cash management in continuous time. *Oper. Res.* **4**: 620–636.
- Federgruen, A., C. Lee. 1990. The dynamic lot size model with quantity discount. *Nav. Res. Log.* **37**(5): 707–713.
- Federgruen, A., Z. Schechner. 1983. Cost formulas for continuous review inventory models with fixed delivery lags. *Oper. Res.* **31**(5): 957–965.
- Federgruen, A., Y. S. Zheng. 1992. An efficient algorithm for computing an optimal (r, Q) policy in continuous review stochastic inventory systems. *Oper. Res.* **40**(4): 808–813.
- Finch, P. 1961. Some probability theorems in inventory control. *Publ. Math. Debrecen* **8**, 241–261.
- Gallego, G., G. J. van Ryzin. 1994. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Sci.* **40**(8): 999–1020.
- Harris, F. W. 1913. How many parts to make at once. *Fact. Mag. Manag.* **10**(2): 135–136. 152.
- He, S., D. Yao, H. Zhang. 2017. Optimal ordering policy for inventory systems with quantity-dependent setup costs. *Math. Oper. Res.* **42**(4): 979–1006. <https://doi.org/10.1287/moor.2016.0833>
- Huh, W. T., G. Janakiraman, A. Muharremoglu, A. Sheopuri. 2011. Inventory systems with a generalized cost model. *Oper. Res.* **59**(4): 1040–1047.
- Iglehart, D. 1963a. Optimality of (s, S) policies in the infinite-horizon dynamic inventory problem. *Management Sci.* **9**(2): 259–267.
- Iglehart, D. 1963b. Dynamic programming and stationary analysis of inventory problems. H. Scarf, D. Gilford, M. Shelly, eds. *Chapter 1 in Multistage Inventory Models and Techniques*. Stanford University Press, Stanford, CA, 1–31.
- Keilson, J., H. Gerber. 1971. Some results for discrete unimodality. *J. Am. Stat. Assoc.* **66**(334): 386–389.
- Koca, E., H. Yaman, M. S. Aktürk. 2014. Lot sizing with piecewise concave production costs. *INFORMS J. Comput.* **26**(4): 767–779.
- Li, C., V. N. Hsu, W. Xiao. 2004. Dynamic lot sizing with batch ordering and truckload discounts. *Oper. Res.* **52**(4): 639–654.
- Li, C., J. Ou, V. N. Hsu. 2012. Dynamic lot sizing with all-units discount and resales. *Nav. Res. Log.* **59**(3–4): 230–243.
- Lippman, S. A. 1969. Optimal inventory policy with multiple setup costs. *Management Sci.* **16**(1): 118–138.
- Lippman, S. A. 1971. Economic order quantities and multiple setup costs. *Management Sci.* **18**(1): 39–47.
- Lu, Y., M. Song. 2014. Inventory control with a fixed cost and a piecewise linear convex cost. *Prod. Oper. Manag.* **23**: 1966–1984.
- Perera, S., G. Janakiraman, S.-C. Niu. 2017. Optimality of (s, S) policies in EOQ models with general cost structures. *Int. J. Prod. Econ.* **187**: 216–228.
- Porteus, E. 1971. On the optimality of generalized (s, S) policies. *Management Sci.* **17**, 411–426.
- Porteus, E. 1990. Stochastic inventory theory. D. Heyman, M. Sobel, eds. *Handbooks in Operations Research and Management Science*, Vol. 2. Elsevier B.V., North-Holland, 605–652.
- Porteus, E. L. 2002. *Foundations of Stochastic Inventory Theory*. Stanford University Press, Stanford, CA.
- Presman, E., S. P. Sethi. 2006. Inventory models with continuous and Poisson demands and discounted and average costs. *Prod. Oper. Manag.* **15**(2): 279–293.
- Rosling, K. 2002. Inventory cost rate functions with nonlinear shortage costs. *Oper. Res.* **50**(6): 1007–1017.
- Ross, S. 1996. *Stochastic Processes*, 2nd edn. John Wiley & Sons, New York.
- Rubalskiy, G. 1972a. On the level of supplies in a warehouse with a log in procurement. *Eng. Cyber.* **10**: 52–57.
- Rubalskiy, G. 1972b. Calculation of optimum parameters in an inventory control problem. *Eng. Cyber.* **10**: 182–187.
- Sahin, I. 1979. On the stationary analysis of continuous review (s, S) inventory systems with constant lead times. *Oper. Res.* **27**(4): 717–729.
- Sahin, I. 1983. On the continuous review (s, S) inventory model under compound renewal demand and random lead times. *J. Appl. Prob.* **20**: 213–219.
- Scarf, H. 1960. The optimality of (S, s) policies in dynamic inventory problems. K. J. Arrow, S. Karlin, P. Suppes, eds. *Mathematical Methods in the Social Sciences*. Stanford University Press, Stanford, CA, 196–202.
- Scarf, H. 1963. A survey of analytic techniques in inventory theory. H. Scarf, D. Gilford, M. Shelly, eds. *Chapter 7 in Multistage Inventory Models and Techniques*. Stanford University Press, Stanford, CA, 185–225.
- Sivazlian, B. 1974. A continuous review (s, S) inventory system with arbitrary interarrival distribution between unit demand. *Oper. Res.* **22**: 65–71.
- Sulem, A. 1986. A solvable one-dimensional model of a diffusion inventory system. *Math. Oper. Res.* **11**(1): 125–133.
- Tijms, H. C. 1972. *Analysis of (s, S) Inventory Models*. *Mathematical Centre Tracts* **40**. Mathematisch Centrum, Amsterdam.
- Tunc, H., O. A. Kilic, S. A. Tarim, B. Eksioğlu. 2016. The stochastic lot sizing problem with piecewise linear concave ordering costs. *Comput. Oper. Res.* **65**: 104–110.
- Veinott, A. F. 1966. On the optimality of (s, S) inventory policies: new conditions and a new proof. *SIAM J. Appl. Math.* **14**(5): 1067–1083.
- Yao, D., X. Chao, J. Wu. 2015. Optimal control policy for a Brownian inventory system with concave ordering cost. *J. Appl. Prob.* **52**(4): 909–925.
- Yao, D., X. Chao, J. Wu. 2017. Optimal policies for Brownian inventory systems with a piecewise linear ordering cost. *IEEE Trans. Autom. Control* **62**(7): 3235–3248.
- Zhao, Y., M. N. Katehakis. 2006. On the structure of optimal ordering policies for stochastic inventory systems with minimum order quantity. *Prob. Eng. Infor. Sci.* **20**(2): 257–270.
- Zheng, Y.-S. 1991. A simple proof of the optimality of (s, S) policies in infinite-horizon inventory systems. *J. Appl. Prob.* **28**(4): 802–810.
- Zipkin, P. 2000. *Foundations of Inventory Management*. McGraw-Hill, Boston, MA.