A Single-Supplier, Multiple-Retailer Model with Single-Season, Multiple-Ordering Opportunities and Fixed Ordering Cost

Apurva Jain, Kamran Moinzadeh, Yong-Pin Zhou
Michael G. Foster School of Business
University of Washington, Seattle
{apurva, kamran, yongpin@uw.edu}

Abstract

In this paper, we discuss the replenishment decision of seasonal products in a two-echelon distribution system consisting of a supplier and multiple retailers. Due to long manufacturing lead time, the supplier orders its entire stock for the season well in advance. The retailers, on the other hand, can replenish their inventory from the supplier throughout the season as demand realizes. Demand at each retailer follows a Poisson process. Each retailer order incurs a fixed cost, and the usual understocking and overstocking costs occur. When retailer lead time is negligible, we show that it is optimal for the retailer to follow a time-based, order-up-to policy and order only when inventory is depleted. We also characterize the structure of the optimal policy and propose a number of heuristics for easier computation. For the supplier, we express the distribution of total demand. This allows the supplier to solve a classic newsvendor problem to determine the total stock for the season. We find that the optimal retailer policy can sometimes cause large demand variation for the supplier, resulting in lower supplier profit. In centralized settings, this may even result in lower system profit than some naïve retailer heuristics, creating inefficiency in the supply chain. We offer insights on potential causes and managerial implications.

1. Introduction and Literature Review

We develop a supplier-retailer inventory model for a single-season, random-demand product. At the end of the season, overstocking costs are charged on leftover inventory and understocking costs are charged on lost sales. Typically, the retailer's replenishment decision in such a setting is modeled as a newsvendor decision, with only one ordering opportunity at the beginning of the season. In our model, however, the retailer can place a replenishment order at any point during the season at a fixed ordering cost. With these continuous ordering opportunities, we determine the retailer's optimal ordering policy and study its implications for the supplier and the system.

The context of single-season, random-demand, and understocking and overstocking costs applies to a wide range of industries, such as fashion retailing, computer and consumer electronics, and pharmaceutical manufacturing. Typically, the product involved in these situations has a short life cycle (one season), and exhibits long manufacturing lead time. The combination of short season and long lead time usually leads to the assumption of a single beginning-of-season ordering opportunity and the application of the newsvendor model. We observe that while this may be true for the manufacturer of the item (called the "supplier" in our model), downstream retailers are free to replenish stocks from the supplier any time during the season. The retailer has to incur a fixed ordering cost for these replenishments. For other relevant costs, we simply follow the newsvendor model convention and charge the retailer overstocking cost on end-of-season leftover inventory and understocking cost on lost sales. Similar costs apply at the supplier echelon. Given that the supplier experiences a long lead time, we take the traditional approach of analyzing the supplier's ordering decision as a newsvendor model.

The fashion apparel industry provides a good example of the situation described here. Manufacturing lead times are long because most manufacturers are located in Asia (see Skinner 1992, Hammond 1991). Demand during the season is highly unpredictable. Orders are received (by the supplier in our model) at the beginning of the season and then used to satisfy retail orders placed (by the retailer in our model) during the season. The transportation lead time from supplier to retailer is short. Retailers incur ordering costs (e.g., shipping). Overstocking and understocking costs faced by retailers in such situations are severe; they are estimated to be around 35% of the sales (Frazier, 1986). Distribution of flu vaccines in the United States provides another example. Each year, the Centers for Disease Control and Prevention identifies the most likely flu strain for the upcoming flu season. Manufacturing the flu vaccine can take six months (Russell 2004), so the stock of vaccine doses is ordered well in advance of the season. Vaccines are received by suppliers at the beginning of the season and then are shipped during the season to health care providers (retailers) as they place orders.

The literature on the newsvendor model is vast. This paper differs from that literature in our modeling of multiple ordering opportunities in a single season. We provide a quick summary of papers that model at least one more ordering opportunity in the context of a seasonal product. Lau and Lau (1997) model a newsvendor problem with an opportunity to replenish in the middle of the period. In Gallego and Moon (1993), the second replenishment opportunity occurs at the end of the season. In our model, the retailer can order at any time. Thus, there are unlimited ordering opportunities.

Much of the literature that allows more than one ordering opportunity in a single season is motivated by Quick Response (QR) in the apparel industry (Hammond 1990). QR is a strategy under which manufacturing lead times can be shortened through better use of information. As a result, supply chains may benefit by more accurate forecasts. The modeling of QR and related contracting mechanisms naturally calls for more than one ordering opportunity. In their model of QR, Iyer and Bergen (1997) allow a second order in the season. In the time between the first and second orders, however, the model only captures the sales of "related" items. Fisher and Raman (1996) apply the notion of QR and examine the manufacturing, replenishment, and forecasting decisions in the fashion skiwear industry. Gurnani and Tang (1999) model a retailer who has two opportunities to order a product from the manufacturer, but the unit cost at the second opportunity is uncertain. A sub-problem in Donohue (2000) also models two-opportunity ordering with the cost at the second opportunity higher than the cost at the first. Tsay (1999) models a quantity flexibility contract in which the retailer is allowed to change the order size in the second period based on the demand signal received in the first period. More recent literature, such as Cachon and Swinney (2011), continues to model QR as another opportunity to place an order before the start of the selling season. None of these papers, however, explicitly model the sales of items between the first and second opportunities. In our model, sales may continuously occur between any two orders. Some other papers in this stream of research, such as Eppen and Iyer (1997) and Barnes-Schuster *et al.* (2002), do capture sales of items between two ordering opportunities. The motivation that drives these models – availability of updated demand signal in the second period – is not the same as ours. At the time retailers decide on the size of their first period order, models in these papers must take into account the possibility of a better demand signal later. In our model, there is no new demand signal, but earlier order size decisions must take into account the possibility that orders at any later time are allowed. We remove the restriction of only two ordering opportunities, allow orders at any time, and derive the optimal policy.

There are two other areas of inventory modeling research that may be interpreted as offering more than one ordering opportunity: emergency ordering and transshipment. When inventory obtained from the regular source runs out, the retailer can order replenishment from an alternative emergency source or request transshipment from an internal source. In a continuous-time, infinite-horizon setting, Moinzadeh and Schmidt (1991) and Axsater (2007) present models with emergency orders and Axsater (1990) considers transshipment. It is difficult to find optimal policies in this setting, and these papers propose and analyze heuristic policies. In a single-season, finite-horizon setting, Khouja (1996) develops an example of an emergency ordering model, but the size of the emergency order is limited to a fixed percentage of unsatisfied demand. For an example of a single-season transshipment model, see Rudi, Kapur and Pyke (2001) in which, after the initial order, there may be one opportunity of transshipment between the two firms. Recent research in this area, such as Shao *et al*. (2011) has focused on analyzing pricing policies for transshipment. The additional orders in these types of models are driven by the cost and lead time differences between different sources. In our single-season, continuous-time model, we focus on multiple orders from the same source. We put no constraints on the size of the order or the number of ordering opportunities.

The models that come closest in spirit to our retailer's model are the finite-horizon, periodic-review models. The classic dynamic formulation goes back to Arrow *et al.* (1951) and Scarf (1960) who proved the optimality of (s,S) policy for models with fixed ordering cost. For infinite-horizon models, these optimality results can be translated to continuous-review settings like ours (see Stidham 1986). Based on the similarity of cost functions in the infinite-horizon periodic and continuous settings, the efficient algorithms developed for the periodic settings can be applied to continuous models as well. See, for example, Federgruen and Zipkin (1985), Zheng and Federgruen (1991), and Zipkin (2000). We are not aware of similar correspondence results for computation of optimal policies in finite-horizon continuous models. Joint optimization of period-length and order-up-to level in such settings continues to be of interest; see, for example, Liu and Song (2011). The finite-horizon problem described in this paper can be viewed as a periodic inventory model with infinite periods, each with an infinitesimal period length but with end-of-horizon costs only. It is not clear how the existing periodic results can be translated to such a setting. In this paper, we work directly with the continuous-time, finite-horizon formulation. The literature on proving optimality of inventory policies directly within a continuous-time framework is scant. In our specific continuous-time setting, our paper offers a method to find the optimal policy.

Our model first addresses the stocking and replenishment decision of a retailer who experiences Poisson demand in a single season of finite length. (Our model also allows the possibility of a nonhomogeneous Poisson demand; that is, demand rate may change over time.) At the end of the season, an overstocking cost is charged on each unit of leftover inventory. An understocking cost is charged on each unit of sales that is lost due to lack of inventory. The differentiating feature of our model is that, unlike other single-season models cited above, we allow the retailer to place a replenishment order at any point during the season. The retailer incurs a fixed ordering cost for each order. We assume that the supplier has ample supply of the product or can instantly make up for any shortage with a unit penalty cost. Furthermore, lead time for the retailer’s orders is assumed to be negligible. Without a fixed shipping cost, the problem would be trivial since the retailer will employ an order-for-order policy. However, the fixed per-order shipping cost is present in most real-world applications. This complicates the problem because the retailer’s order quantity is not of unit size anymore and can vary from order to order, depending on when the order is placed. In addition, while making a decision on order size, the retailer must incorporate the possibility of placing future orders. The question of when to place a replenishment order is also open: Is it ever optimal to forego placing an order and incur lost sales? We contribute to the literature by characterizing the optimal ordering policy in such situations, which minimizes the retailer’s average total cost during the season. Our method is somewhat unusual because we convert the continuous-time problem to a discrete framework by identifying discrete points at which optimal decisions change. We devise an exact procedure for computing the optimal policy parameters and propose a number of heuristics for computing the policy parameters. We then evaluate the performance of such heuristics through a numerical experiment, which shows that some of our heuristics deliver near-optimal results.

Next, after solving the retailer's problem, we consider a distribution system consisting of one supplier and multiple retailers. We assume that the supplier replenishes stock only once at the beginning of the season and supplies the retailer's orders during the season. If the supplier runs out of stock during the season, the excess demand will be satisfied through a third party at some penalty to the supplier. We find the optimal order quantity for the supplier by expressing the exact distribution of the supplier's demand over the season. Our numerical results provide interesting managerial insights, such as the observation that as the retailer changes its policy to take advantage of multiple ordering opportunities, the supplier profit decreases in general.

The rest of the paper is organized as follows. In the next section, we consider the retailer's problem in isolation, derive the optimal policy, and present heuristics. In §3, we extend our model to a distribution system with one supplier and multiple retailers. Using the results developed in §2, we derive a method for computing the exact distribution of the supplier's demand and propose an approximation for it. Section 4 uses a numerical experiment to evaluate the retailer's policy and the effectiveness of the heuristics suggested in §2. We close this section by outlining some of the managerial implications resulting from operating such models in a distribution setting. Finally, in §5, we summarize our findings and suggest a number of possible extensions for future work.

2. The Retailer's Problem

Consider a single-season inventory system of a product. Let the length of the finite season be *T*. Since we will use backwards induction in our derivations, we follow the convention to let θ denote the time till the end of the season. So $\theta = T$ indicates the beginning of the season while $\theta = 0$ indicates the end. Using this terminology, an interval of $[\theta_1, \theta_2]$ refers to $\theta_1 \le \theta \le \theta_2$. In real time, θ_2 happens before θ_1. The retailer experiences unit-size demands for a product that occur randomly during the season, following a Poisson process with rate λ. The retailer can order at any time from a supplier. The fixed ordering and shipping cost is *K* for each order placed by the retailer. We assume that lead time is zero, and the supplier is always able to meet the order. As a result, an order placed at time θ can be used to satisfy a demand that occurred at time θ. At the end of the season, a per-unit overstocking cost *w* is charged on leftover inventory. All demands during the season that are not met from the retailer's inventory are lost and charged a per-unit understocking cost of π. As in a traditional newsvendor model, any end-of-season per-unit holding and salvage costs can be directly captured in the overstocking cost (see Nahmias 2009, pages 301-302). The retailer's objective function is to minimize the total expected cost, which is comprised of the fixed ordering cost and the overstocking and understocking costs.

It is a standard result (e.g., see Law and Kelton 2000, pages 486-489) that a nonhomogeneous Poisson process with rate $\lambda(\theta)$, $0 \leqslant \theta \leqslant T$, can be generated from a homogeneous Poisson process as follows: First, generate a homogeneous Poisson process with rate 1.Then, for any arrival in the homogeneous Poisson process at time $\tilde{\theta}$, invert the expectation function $T - \tilde{\theta} = \Lambda(T - \theta) = \int_{T-\theta}^{T} \lambda(s)ds$ to find the corresponding arrival time θ in the nonhomogeneous Poisson process. Essentially, by scaling time with the expectation function $\Lambda(\cdot)$, we have a one-to-one mapping between the two processes. Moreover, any ordering action taken in the homogeneous case at time $\tilde{\theta}$ can be taken in the nonhomogeneous case at time θ, and vice versa. The resulting system dynamics are identical. Hence, there is also a one-to-one mapping between the policies that can be employed between the two processes. Finally, because all the costs (ordering, overstocking, and understocking) are not time dependent, the corresponding policies will have identical costs. As a result of this one-to-one correspondence, in this paper we will focus only on the homogeneous Poisson process. All the major results of this paper continue to hold if the demand follows a nonhomogeneous Poisson process.

To gain insight into the structure of an optimal ordering policy, we first study the newsvendor optimal order quantity – which is derived under the condition that only one order is allowed – and analyze how it changes with the duration of time it covers. Denote by $D(\theta)$ the random Poisson demand during a period of length θ. Following the notation in Hadley and Whitin (1963), we denote the PDF and the complementary CDF of $D(\theta)$ as:

$$p(n,\theta) = \frac{e^{-\lambda\theta}(\lambda\theta)^n}{n!} \text{ and } P(n,\theta) = \sum_{m=n}^{\infty} p(m,\theta).$$

Let $TC(S,\theta)$ be the expected newsvendor cost; that is, the costs incurred at the end of the time horizon, if inventory at θ, after any demand and ordering, is S and no orders are placed during $[0,\theta)$. Since $np(n,\theta) = \lambda\theta\, p(n-1,\theta)$, we have:

$$\begin{aligned} TC(S,\theta) &= w(S-\lambda\theta) + (w+\pi)E\left(D(\theta)-S\right)^+ \\ &= w(S-\lambda\theta) + (w+\pi)[\lambda\theta P(S,\theta) - SP(S+1,\theta)]. \end{aligned} \tag{1}$$

Let $g(\theta) = \min_S TC(S,\theta)$ be the optimal newsvendor cost for the period $[0,\theta)$, and $\overline{S}(\theta)$ be the corresponding optimal newsvendor stocking level (the largest S that minimizes $TC(S,\theta)$). Next, we present some preliminaries. All the proofs in this paper are shown in the Appendix.

Lemma 1 There exist time points $0 = \overline{\theta}_{-1} < \overline{\theta}_0 < \overline{\theta}_1 < \ldots < \overline{\theta}_{l-1} < \overline{\theta}_l = T$ for some $l \geq 0$ such that $\overline{S}(\theta) = 0$ on $[\overline{\theta}_{-1}, \overline{\theta}_0)$, $\overline{S}(\theta) = n$ on $[\overline{\theta}_{n-1}, \overline{\theta}_n)$, $\forall 1 \leq n < l$, and $\overline{S}(\theta) = l$ on $[\overline{\theta}_{l-1}, \overline{\theta}_l]$.

Lemma 2 $g(\theta)$ is increasing in θ.

Lemmas 1 and 2 show that in a newsvendor setting, as the end of the season draws near (i.e., θ decreases), there is less future demand to cover and both the optimal newsvendor quantity and the optimal cost decrease accordingly. In particular, Lemma 1 shows that in the newsvendor setting, where only one order is placed, the optimal order quantity is an increasing (in θ) step function. We will show that this pattern continues to hold in the multiple-order setting, though the parameters will change.

Returning to the multi-order setting, we introduce some additional notation first (a summary is provided below). Let $I(\theta)$ represent on-hand inventory at θ before any demand occurrence or possible order, and $I^+(\theta)$ represent on-hand inventory after any demand and ordering at θ. Define an ordering policy φ to be a mapping from $[0,T] \times [0,\infty) \times \{0,1\}$ to $[0,\infty)$ such that, given the inventory at time $\theta \in [0,T]$ is $I(\theta) \in [0,\infty)$ and whether there is a demand arrival at that time ($A(\theta) = 1$ indicates an arrival at time θ and $A(\theta) = 0$ otherwise), policy φ determines whether and how much to order. This results in $I^+(\theta) \in [0,\infty)$. For every φ, there is a corresponding cost function $c^\varphi(\theta, I(\theta), A(\theta))$ that is a mapping from $[0,T] \times [0,\infty) \times \{0,1\}$ to $[0,\infty)$. It represents the sum of lost sales and fixed ordering costs under φ if $(\theta, I(\theta), A(\theta))$ occurs. We say policy φ is feasible if $c^\varphi(\theta, I(\theta), A(\theta))$ is a measureable function on $[0,T] \times [0,\infty) \times \{0,1\}$.

Further, for any policy φ let $V^\varphi(I^+(\theta), \theta)$ be the expected total cost incurred during the interval $[0,\theta)$. To derive $V^\varphi\ (I^+(\theta), \theta)$ we note that $\int_0^{\theta^+} dc^\varphi\ (t, I(t), A(t))$ is the cost of applying policy φ under one sample path,$\{A(t)\}_{t=0}^{\theta}$, on $[0,\theta)$. Let $\Omega(0, \theta)$ be the set of all possible sample paths on $[0, \theta)$ and P its corresponding probability measure, then we have

$$V^\varphi\ (I^+(\theta), \theta) = \int_\Omega \int_0^{\theta^+} dc^\varphi\ (t, I(t), A(t))\, dP.$$

Define the optimal value function at time θ as $V^*(I^+(\theta),\theta)=\inf\{V^\varphi(I^+(\theta),\theta):\ \varphi\in\Phi\}$ where Φ is the set of all non-anticipatory policies such that $V^\varphi(I^+(\theta),\theta)$ exists for all θ, $I^+(\theta)$, and $\varphi\in\Phi$. To simplify exposition we use i, instead of $I^+(\theta)$, in the rest of the paper. A summary of additional notations is provided below:

φ	Ordering policy, $[0,T]\times[0,\infty)\times\{0,1\}\to[0,\infty)$
c^φ	Cost function corresponding to φ, $[0,T]\times[0,\infty)\times\{0,1\}\to[0,\infty)$
$V^\varphi(I^+(\theta),\theta)$	The expected total cost incurred under φ during $[0,\theta)$
Φ	$\{\varphi\mid\varphi \text{ is non}-\text{anticipatory and } V^\varphi(I^+(\theta),\theta) \text{ exists for all } \theta, I^+(\theta)\}$
$\Omega(0,\theta)$	$\{\text{all demand sample paths on } [0,\theta)\}$
P	probability measure corresponding to $\Omega(0,\theta)$

Let $\mathbb{N}$ denote the set of non-negative integers and $\mathbb{I}$ be any subset of $[0,T]$.

Definition A policy $\varphi\in\Phi$ is called optimal on $\mathbb{I}\subset[0,T]$ if it satisfies $V^\varphi(i,\theta)=V^*(i,\theta)$ for all $\theta\in\mathbb{I}$ and $i\in\mathbb{N}$. A policy φ is called optimal if it is optimal on $[0,T]$. Moreover, we call any ordering action that minimizes $V^*(i,\theta)$ the optimal action at time θ, and denote $S(\theta)=\arg\min_{i\in\mathbb{N}}V^*(i,\theta)$. If multiple minimizers exist, the largest value is chosen for $S(\theta)$.

The above definition focuses on the value function and the uniquely determined action. In terms of the policies that may correspond to the same value function, we note that multiple optimal polices may exist. For example, changing an ordering policy on an event with a measure of zero (e.g. that a demand occurs at a specific time) does not change the system cost, but it will result in a different policy. These policies all have the same value function for all time epochs and inventory levels. We say that any two policies φ and φ' are equivalent if $V^\varphi(I^+(\theta),\theta)=V^{\varphi'}(I^+(\theta),\theta)$ for all $(I^+(\theta),\theta)$. Henceforth, the optimal policy φ we discuss below represents all the policies that are equivalent to it. We begin with narrowing the set of policies that we must consider.

Proposition 1 Suppose that under policy φ there is a positive probability that an order is placed when inventory is positive. Then, policy φ is not optimal.

The next step is to further narrow the set of policies to those that only order when a demand occurs.

Proposition 2 Suppose that under policy φ there is a positive probability that an order is placed at a time when no demand occurs. Then, policy φ is not optimal.

We note that Proposition 2 is consistent with the result for infinite time horizon, where it is optimal to place orders only at demand epochs when the demand follows a Poisson process and lead time is zero (see Moinzadeh and Zhou 2008).

Propositions 1 and 2 imply that we need to only consider policies that order only when inventory is zero and a demand takes place. Next, we build on Propositions 1 and 2 to characterize the optimal action at any time. The expression for value function at time θ should include: (a) expected ordering cost for all the orders placed in $[0,\theta)$, (b) expected understocking cost for all the demands lost in $[0,\theta)$, and (c) expected overstocking cost for leftover inventory at 0. We accomplish this in two steps.

First, we show in Section 2.1 that there exists a point θ_0 after which it is optimal to not place an order. This determines the optimal actions for $\theta \le \theta_0$ and defines the value function at θ_0 as the newsvendor cost function. Next, in Section 2.2 we define the backwards value function induction for $\theta > \theta_0$ and use it to show the structure of the optimal action for the entire [0,*T*]. In Section 2.3, we show how to compute the optimal order quantities and develop heuristics.

2.1. Existence of θ_0

Lemma 1 shows that in the single-order, newsvendor setting, it is optimal not to order near the end of the time horizon. The next proposition will show that a similar result holds for our multiple-order setting.

We note that when $\pi(\lambda T + 1) \leq K$, it is optimal never to place any order. Thus, from now on, we will only consider the cases where $K < \pi(\lambda T + 1)$. We start the analysis at the end of the season ($\theta = 0$). If the ordering cost is reasonably high, then when there is relatively little time left in the season, it should be optimal not to place another order because the fixed ordering cost could outweigh the expected understocking cost. Therefore, there should exist a point in time, call it θ_0, after which the optimal action is to never place any order. On the other hand, when $K \leq \pi$, it's optimal to never lose any sale. Proposition 3 formalizes this idea.

Since we will always be focusing on the optimal policy, for the sake of convenience, we drop the * superscript from the value function. Consequently, the value function is expressed as $V(i,\theta)$ in the rest of the paper.

Proposition 3 Assume $\pi < K < \pi(\lambda T + 1)$. Let θ_0 be the unique non-negative solution to the equation $g(\theta_0) + K = \pi(\lambda\theta_0 + 1)$, and $S_0 = \overline{S}(\theta_0)$. Then $\theta_0 > 0$, $V^*(i,\theta) = TC(i,\theta)$ for $\theta \in [0, \min\{\theta_0, T\}]$, and any policy is optimal on $[0, \theta_0]$ if it never orders at time $\theta \in [0, \theta_0)$ and orders up to S_0 when inventory is zero and a demand occurs at θ_0. For $K \leq \pi$, we have $\theta_0 = 0$ and $S_0 = 0$.

Recall that the value function $V(i,\theta)$ is the expected total cost in $[0,\theta)$ with $I^+(\theta) = i \ (i \geq 0)$. It includes the ordering cost of all the orders placed in $[0,\theta)$, the understocking cost in $[0,\theta)$ and the overstocking cost incurred at 0. With the benefit of Proposition 3, we know that it is optimal to not place an order at

$\theta \le \theta_0$, so $V(i,\theta)$ for $\theta \le \theta_0$ is just the newsvendor cost function. That is: $V(i,\theta) = TC(i,\theta)$ for $\theta \le \theta_0$. Moreover, for $\pi < K < \pi(\lambda T + 1)$, we show, in the proof of Proposition 3, that $V(S_0,\theta_0) + K = V(0,\theta_0) + \pi$. This implies that when the on-hand inventory is zero and a demand occurs at θ_0, an optimal policy on $[0,\theta_0]$ is indifferent between not ordering and ordering up to S_0. We choose to order up to S_0 to be consistent with the assumption throughout the paper that in such cases the retailer will choose the largest order-up-to level.

2.2. Characterization of Optimal Policy

We are now ready to examine the optimal ordering quantities for $\theta > \theta_0$. From Propositions 1 and 2, we know that an optimal policy orders only when inventory is zero and a demand occurs; thus the next question is natural: when an order is placed at $\theta > \theta_0$, how much should be ordered? To answer this question, we let $I^+(\theta) = i$ and define, for any sample path ω, $L_\omega(i,\theta_0,\theta)$ and $N_\omega(i,\theta_0,\theta)$ to be the units of lost sales and the number of orders placed in $[\theta_0,\theta)$ respectively. Moreover, let the ending inventory at time θ_0 be j. Then,

$$V(i,\theta) = K\, E_\omega\left[N_\omega(i,\theta_0,\theta)\right] + \pi\, E_\omega\left[L_\omega(i,\theta_0,\theta)\right] + E_\omega\left[V(j,\theta_0)\right]. \tag{2}$$

(The ω subscript will be suppressed later in the paper when there is no confusion.) By definition, $S(\theta) = \arg\min_i V(i,\theta)$ denotes the optimal order-up-to level, if inventory is zero and a demand occurs at time θ. Again, if multiple optima exist, $S(\theta)$ is the largest one.

In the rest of this section, we will simplify the value function (2) and prove its properties. This will allow us to characterize $S(\theta)$ for all $\theta > \theta_0$.

Define $\theta_1 = \sup\{ \theta \mid \theta_0 < \theta \le T$, and if a demand occurs on $[\theta_0, \theta)$ and inventory is zero, the optimal action is to order up to $S_0 \}$. The following proposition first shows that the set is non-empty and $\theta_1 > \theta_0$. Moreover, at θ_1 both order-up-to levels S_0 and $S_0 + 1$ minimize the value function. By choosing the larger value as always, we then know the optimal order-up-to level at θ_1 is $S_0 + 1$. That is, $S_1 = S(\theta_1) = S_0 + 1$.

Proposition 4

(i) $\theta_1 > \theta_0$

(ii) For $\theta_0 \le \theta \le \theta_1$, the value function can be written recursively as:

$$V(i,\theta) = K \sum_{n=0}^{\infty} P\left(n(S_0+1)+i+1, \theta-\theta_0\right) + \sum_{j=0}^{S_0} \sum_{n=1}^{\infty} p\left(n(S_0+1)+i-j, \theta-\theta_0\right) V(j,\theta_0) + \sum_{j=0}^{i} p(j,\theta-\theta_0) V(i-j,\theta_0) \tag{3}$$

(iii) Any policy is optimal on $[0,\theta_1]$ if it satisfies the following: When a demand occurs at time θ,

- if $\theta_0 < \theta < \theta_1$ and $I(\theta) = 0$, order up to S_0,
- if $\theta = \theta_1$ and $I(\theta) = 0$, order up to $S_1 = S_0 + 1$,
- otherwise, do not order.

(iv) $V(S_1, \theta_1) = V(S_0, \theta_1)$.

The first term on the right hand side of (3) represents the total expected ordering cost between θ and θ_0, and the last two terms represent all the future expected cost starting from time θ_0 (starting inventory is *i* at time θ and ending inventory is *j* at time θ_0). The difference is that in the second term at least one order is placed between θ and θ_0, and in the last term no order is placed.

We can now characterize the optimal order-up-to levels on the entire [0,*T*].

Theorem 1 There exist an integer $n \geq 0$ and time points $0 \leq \theta_0 < \cdots < \theta_{n-1} < \theta_n = T$ such that the following policy is optimal on $[0,\theta_k]$, $0 \leq k \leq n$:

1) it orders only when inventory is zero and a demand occurs,
2) it never orders at time $\theta \in [0,\theta_0)$,
3) it orders up to $S_{i-1} = S_0 + (i-1)$ at time $\theta \in [\theta_{i-1}, \theta_i)$ for $1 \leq i \leq k < n$, and
4) it orders up to $S_n = S_0 + n$ at time $\theta = \theta_n = T$ for $k = n$.

Moreover, $V(i,\theta)$ is quasi-convex in *i* for all θ , and $V(S_0,\theta_0) = V(0,\theta_0) + \pi - K$ and $V(S_k,\theta_k) = V(S_{k-1},\theta_k)$ for all $1 \leq k \leq n-1$.

There may exist other policies that achieve the same value function as the policy in Theorem 1, but in the rest of the paper, we will refer to this policy as *the* optimal policy. As Theorem 1 shows, both S_k and S_{k-1} minimize the value function at time θ_k , $1 \leq k \leq n-1$. Again, we choose to order up to the larger value, S_k, to be consistent with the convention in the rest of the paper.

2.3 Computations of Retailer's Optimal Policy and Heuristics

Similar to (3), we have, for $1 \leq k \leq n-1$:

$$V(i,\theta_k) = KE[N(i,\theta_k,\theta_{k-1})] + \sum_{j=0}^{S_{k-1}} \sum_{n=1}^{\infty} V(j,\theta_{k-1})\, p\left(n(S_{k-1}+1) + i - j, \theta_k - \theta_{k-1}\right) + \sum_{j=0}^{i} V(i-j,\theta_{k-1})\, p(j,\theta_k - \theta_{k-1}),$$

where $E[N(i,\theta_k,\theta_{k-1})]$ is the expected number of orders (of size S_{k-1} +1) between θ_k and θ_{k-1}, and can be calculated as:

$$E[N(i,\theta_k,\theta_{k-1})] = \sum_{n=0}^{\infty} \Pr[N(i,\theta_k,\theta_{k-1}) > n] = \sum_{n=0}^{\infty} \Pr[D(\theta_k - \theta_{k-1}) > n(S_{k-1}+1) + i]$$
$$= \sum_{n=1}^{\infty} P\left((n-1)(S_{k-1}+1) + i + 1, \theta_k - \theta_{k-1}\right).$$

These expressions allow us to recursively calculate the value function by going back in time starting with θ_0. We now present a formal procedure for computing the optimal policy:

1. If $K \geq \pi(\lambda T + 1)$, optimal policy is to never order. Go to end.

2. If $K \leq \pi$, set $\theta_0 = 0$, $S_0 = 0$;

Else, set θ_0 as the unique solution to $g(\theta_0) + K = \pi(\lambda\theta_0 + 1)$.

Set $S_0 = \bar{S}(\theta_0)$ as the largest integer such that $P(S_0, \theta_0) \geq \frac{w}{w+\pi}$.

Compute $V(i,\theta_0) = TC(i,\theta_0) \quad \forall i = 0,..,S_0 + 1$; set $k = 1$.

3. Set θ_k as the unique root in $(\theta_{k-1}, T]$ of $V(S_{k-1}+1, \theta_k) = V(S_{k-1}, \theta_k)$ where we use recursion

$$V(S_{k-1}+1,\theta_k) = K\sum_{n=0}^{\infty} P\left(n(S_{k-1}+1) + S_{k-1}+1+1, \theta_k - \theta_{k-1}\right) + \sum_{j=0}^{S_{k-1}}\sum_{n=1}^{\infty} p\left(n(S_{k-1}+1) + S_{k-1}+1-j, \theta_k - \theta_{k-1}\right)V(j,\theta_{k-1})$$
$$+ \sum_{j=0}^{S_{k-1}+1} p(j, \theta_k - \theta_{k-1})V(S_{k-1}+1-j, \theta_{k-1}).$$

If no root is found, set $\theta_k = T, n = k$;

Else, set $S_k = S_{k-1} + 1, k = k + 1$, go to step 3.

Compute $V(i,\theta_k) \quad \forall i = 0,..,S_{k-1} + 1$.

End.

Note that uniqueness of θ_0 is due to Proposition 3, and uniqueness of θ_k is shown in step 3 of the proof of Proposition 4. We used the root finding algorithm code freely available under GNU general public license. In some cases, the algorithm may be computationally demanding as one needs to find θ_k by solving $V(S_{k-1}+1, \theta_k) = V(S_{k-1}, \theta_k)$. When many θ_k s are involved, the algorithm may run for a long time. For example, on a 2.33 GHz CPU Windows XP desktop with 4GB RAM, it took more than one hour to find the optimal policy in one of the cases presented in Section 4. Thus, it is desirable to develop effective heuristics that yield near-optimal results. We now outline four heuristics for solving the retailer's problem.

First, we start with the newsvendor model as the baseline case. Other heuristics can be compared with this baseline to assess the improvement one can achieve by allowing the retailer multiple order opportunities during the season. We denote by Heuristc H1 the policy that orders only once at the beginning of the season in the amount of $\bar{S}(T)$. (Recall that $\bar{S}(\theta)$ is the optimal newsvendor stocking level (the largest S that minimizes $TC(S,\theta)$) when time period $[0,\theta)$ is covered.)

Next, we devise heuristics which are mainly based on improvement to the newsvendor model. One plausible heuristic, referred to as H2, can be one which orders up to the newsvendor quantity $\bar{S}(\theta)$ when $I(\theta) = 0$ and a demand occurs ($\theta \geq \theta_0$). Note that H2 is still myopic since every time an order is placed, it is implicitly assumed that this will be the last one in the season.

As opposed to H1, where only one order is placed at the beginning of the season, under H2 whenever inventory is zero and a demand occurs with at least θ_0 left, the optimal newsvendor quantity at that point is ordered. The procedure is easy to employ, but is expected only to perform reasonably in situations where the fixed order cost is significant. In Proposition 5, we prove that the order up to level resulted from H2 serves as an upper bound to that of the optimal policy.

Proposition 5 The optimal order-up-to level for H2 is greater than or equal to that for the optimal policy**.**

Suppose we are to place an order at time θ. To determine the order quantity, H2 uses π as the marginal understocking cost in the newsvendor calculations. Under the optimal policy, however, if inventory runs out in the future the retailer still has the option to make another order. Therefore, the marginal understocking cost should be less than π. The next two heuristics aim to capture this in the calculations.

That is, suppose an order is placed at time θ. At a future time $x > \theta_0$, if $I(x) = 0$ and a demand occurs, an order will be placed to bring inventory back to $S(x)$. The expected total cost is $K + V(S(x), x)$. We will approximate $V(S(x), x)$ in heuristics H3 and H4.

For instance, one may estimate $V(S(x), x)$ by $g(\theta_0)$, the optimal newsvendor cost at time θ_0. That is, the heuristic assumes that expected total future cost at x is the same as the expected total future cost at θ_0, $K + g(\theta_0)$. We call this Heuristic H3. Given this approximation, Heuristic H3 can be specified as follows:

(H3) If at time θ $(\theta > \theta_0)$, $I(\theta) = 0$ and a demand occurs, order up to $S(\theta)$, which is the largest integer *S* satisfying the condition $\sum_{j=0}^{S} [w - (w + \pi) P(S - j, \theta_0)] p(j, \theta - \theta_0) \le 0$.

The derivation of (H3) is available in the appendix.

Heuristic H3 can be further improved by using a higher order approximation of $V(S(\theta), \theta)$. Define:

$\beta(\theta) = \dfrac{g(\theta) - g(\theta_0)}{\theta - \theta_0}$, then $V(S(x), x) \approx g(\theta_0) + \beta(\theta)(x - \theta_0) = \pi(\lambda\theta_0 + 1) - K + \beta(\theta)(x - \theta_0)$.

We call this Heuristic H4:

(H4) If at time θ $(\theta > \theta_0)$, $I(\theta) = 0$ and a demand occurs, order up to $S(\theta)$, which is the largest integer *S* satisfying the condition

$$\frac{1}{\lambda}\beta(\theta) P(S + 1, \theta - \theta_0) + \sum_{j=0}^{S} [w - (w + \pi) P(S - j, \theta_0)] p(j, \theta - \theta_0) \le 0 .$$

The derivation of (H4) is also available in the appendix.

Before proceeding further, some discussion is in order. First, as stated in Proposition 5, the order up to level resulted by employing H2 serves as an upper bound to that of the optimal policy. Second, our numerical tests indicate that the order up to levels resulted by employing H3 and H4 are upper and lower bounds to those of the optimal policy, respectively. In closing, we would like to stress that the latter observations were only based on our limited numerical experimentation outlined in Section 4.

3. The Supplier's Problem

In this section, we consider a distribution system with one supplier and many independent retailers. The supplier replenishes and receives its stock at the beginning of the season. The supplier's stock is then used to satisfy retailers' orders placed during the season. If the supplier runs out, excess demand will be satisfied instantly at a penalty through a third party. We emphasize that while this assumption may be restrictive, it represents some situations in practice. Clearly, situations in which the supplier's back ordering will result in delayed delivery to the retailer are interesting and challenging, but we leave them as extensions for future work in this area. Let w_0 and π_0 denote the supplier's unit overstocking and understocking costs, respectively. Since the supplier's and the retailers' problems are decoupled, one can use the results developed in the previous section for the retailers. The supplier, however, faces a newsvendor problem. The key here is to find the demand distribution at the supplier, which is comprised of all retailers' orders placed during the season.

We will focus on one retailer for now. Consider any of the intervals $[\theta_{i-1}, \theta_i)$ in the retailer's policy. Given a starting inventory of j at time θ_i and l orders during that interval, the probability that ending inventory at time θ_{i-1} is k is:

$$q^i_{jkl} = \Pr(\text{ending inventory at } \theta_{i-1} = k \,|\, \text{starting inventory at } \theta_i = j,\ \text{number of orders during } [\theta_{i-1}, \theta_i) = l)$$
$$= \begin{cases} p(j-k, \theta_i - \theta_{i-1}) & \text{if } j \geq k > S_{i-1} \text{ and } l = 0, \\ p(j + l(S_{i-1}+1) - k, \theta_i - \theta_{i-1}) & \text{if } k \leq S_{i-1} \text{ and } j + l(S_{i-1}+1) - k \geq 0, \\ 0 & \text{otherwise.} \end{cases}$$

The distribution of retailer's total order quantity over the period, Q, can be expressed as:

$$\Pr(Q = x) = \sum_{\{l_n, l_{n-1}, \ldots, l_1 \mid l_n, l_{n-1}, \ldots, l_1 \geq 0,\ l_n(S_{n-1}+1) + l_{n-1}(S_{n-2}+1) + \ldots + l_1(S_0+1) = x\}} \left(\sum_{\{j_n, j_{n-1}, \ldots, j_0 \mid j_n = S_{n-1},\ 0 \leq j_{n-1}, j_{n-2}, \ldots, j_1, j_0 \leq S_{n-1}\}} q^{n}_{j_n j_{n-1} l_n}\, q^{n-1}_{j_{n-1} j_{n-2} l_{n-1}} \cdots q^{1}_{j_1 j_0 l_1} \right).$$

The evaluation of the above expression relies on considering all possible cases that would add up to retailer ordering x units over the season. First, we enumerate all possible combinations of numbers of orders l_i over different time-segments $[\theta_{i-1}, \theta_i)$ such that total order quantity for all these orders adds up to x units over the season. We can do this because, given retailer's policy structure, the size of order in each time-segment is known. Second, for each of these combinations, we consider all possible values of ending inventories in each time-segment. Finally, the probability for each case can be specified using multiplication of q^{i}_{jkl} values derived above.

As an example, consider n=2. In such a case, time horizon [0,T] is divided into $[0, \theta_0)$, $[\theta_0, \theta_1)$, and $[\theta_1, T = \theta_2]$. The probability that retailer will places l_2 orders of size $S_1 + 1$ and l_1 orders of size $S_0 + 1$ can be expressed as $\sum_{\{j_2, j_1, j_0 \mid j_2 = S_1,\ 0 \leq j_1, j_0 \leq S_1\}} q^{2}_{j_2 j_1 l_2}\, q^{1}_{j_1 j_0 l_1}$ where the summation is taken over all possible j_1 and j_0 (since the starting inventory at time T will be S_1, j_2 is fixed). Therefore, for the total order quantity to be x, we simply sum over all possible l_2 and l_1 that satisfy $l_2(S_1 + 1) + l_1(S_0 + 1) = x$:

$$\Pr(Q = x) = \sum_{\{l_2, l_1 \mid l_2, l_1 \geq 0,\ l_2(S_1+1) + l_1(S_0+1) = x\}} \left(\sum_{\{j_2, j_1, j_0 \mid j_2 = S_1,\ 0 \leq j_1, j_0 \leq S_1\}} q^{2}_{j_2 j_1 l_2}\, q^{1}_{j_1 j_0 l_1} \right).$$

If there are multiple, but not necessarily identical, independent retailers, the total demand from all retailers is simply the convolution of all the distributions. Note that we have intentionally let the order-up-to levels, S_i, $0 \leq i \leq n-1$ be general and not require $S_i = S_0 + i$. This way, the above $\Pr(Q = x)$ expression can be used to evaluate any policy of this form.

Computing the exact distribution of the supplier demand can be cumbersome as it requires convolutions of distributions across retailers. When there are M independent retailers, one easy approximation is to assume that the total order from all the retailers to the supplier is normally distributed. Such approximations have been employed in previous works (See Duermeyer and Schwarz 1981, Graves 1985 and Zipkin and Svoronos 1986). The mean and variance of this normal distribution is simply the sum of all the individual retailers' order mean and variance, which can be calculated from the mass function $\Pr(Q = x)$ given above.

4. Numerical Results

Our first objective in this section is to compare the performance of heuristics H1-H4 with the optimal solution of the retailer's problem and to understand the impact of parameter values on the solution. We will then focus on developing insights into the supplier's problem.

4.1. The Retailer's Problem

As we stated earlier, the major analytical results in Section 2 hold for nonhomogeneous Poisson retailer demand as well as homogeneous Poisson retailer demand. Therefore in our discussion below we will focus on the homogeneous Poisson demand case.

To evaluate the performance of the heuristics, we assign three parameters, π , K and λ , values at different levels ranging from low to high. The other two parameters are held at constant values without loss of generality: $w = 1$, $T = 1$. Understocking penalty cost π takes values of 0.5, 1, 3, and 9, covering a range of newsvendor fractile values: 0.333, 0.5, 0.75, and 0.9. The ordering cost K takes values 1, 5, and 25. Finally, the base arrival rate λ is set at 50, 100, and 200. This results in a total of 36 cases. For one of the 36 cases, where the demand and under-stocking penalty are the lowest but ordering cost is the highest

(π=0.5, K=25, λ=50), none of the heuristics place any orders. We will exclude this case in our analysis below.

Since demand is exogenous, to solve the retailer's problem it is well known that maximizing expected profit and minimizing expected total cost are equivalent. Here we use the total expected cost as comparison measure for two reasons:

1. It follows the tradition in inventory literature.
2. The comparison is cleaner when retailer sales price is not involved.

The performance of heuristics is measured by the % deviation = 100 x (expected cost under heuristic policy – expected optimal cost) / (expected optimal cost). Results of our numerical experiment are reported in Table 1.

	max dev.	min dev.	mean dev.
H1	461.87%	0.00%	81.60%
H2	62.77%	0.00%	7.28%
H3	11.74%	0.00%	1.93%
H4	2.35%	0.00%	0.31%

Table 1: Cost Deviation from the Optimal Policy.

From Table 1, we can make several observations. First, the minimum deviation of all the heuristics is at 0%. This means that each heuristic can perform as well as the optimal policy in some scenarios. For example, this occurs when ordering cost K is very high. As we can see from Proposition 3, when K goes to infinity, θ_0 approaches T. Therefore, the optimal policy and the other heuristics behave like the newsvendor policy H1, ordering only once at $\theta = T$. Their costs are thus the same.

Under H1, the retailer orders the newsvendor quantity only once, at the beginning of the season. Its comparison with the other policies, which may order multiple times, can tell us how valuable it may be to have multiple order opportunities within the season. Our second observation from Table 1 is that this benefit is high, as H1 significantly underperforms the other heuristics. On a practical level, it may be costly to reconfigure supply chain and implement necessary information technologies to enable retailers to order at any time with a negligible lead time. Nevertheless, by providing important information about the benefit of such systems, our numerical results can help in making such managerial decisions.

Once the retailer can reorder anytime during the period, it is possible that he or she approaches this problem using a heuristic similar to H2. That is, the retailer will order the newsvendor quantity anytime the inventory runs out with more than θ_0 left, when it becomes more economical not to order. Note that H2 is still a myopic policy because the retailer doesn't take future order opportunities into consideration when making the current order. Our third observation from Table 1 is that, just by utilizing the reorder opportunities, the retailer can significantly improve its performance with H2, essentially the same myopic ordering policy as that of H1.

H3 and H4 are non-myopic heuristics. They improve H2 by refining the understocking cost to take into consideration any future order opportunities. Our fourth observation from Table 1 is that on average they both perform very well against the optimal policy. H4 has the better performance – it is nearly optimal in many cases, and has a small maximum deviation. This is consistent with the fact that H4 employs a higher-order approximation of $V(S(\theta),\theta)$ than H3, and is thus more accurate. The three worst-performance cases for H4 occur when understocking penalty cost is high ($\pi = 9$) and ordering cost is in the middle of our range ($K = 5$), irrespective of the value of arrival rate λ. The overall performance of H4 is so good that, beyond these three cases, the % deviation for H4 is less than 1% for all other instances. The three worst-performance cases for H3 occur when $K = 1, \lambda = 200$, irrespective of the value of π. The next three worst performances for H3 occur at $K = 1, \lambda = 100$, irrespective of the value of π. In Tables 2 and 3 below, we further discuss the effect of parameter values on the performance of different heuristics.

Figure 1 illustrates the order-up-to levels and the corresponding interval break points for all the heuristics and provides further insights into the performance of the various heuristics. For H1, there is only one order-up-to level at time $\theta=T$. Another way to analyze policies is to focus on the average total order quantity across the horizon T. A table, available in the appendix, reports average order quantities across horizon under different policies for all parameter combinations.

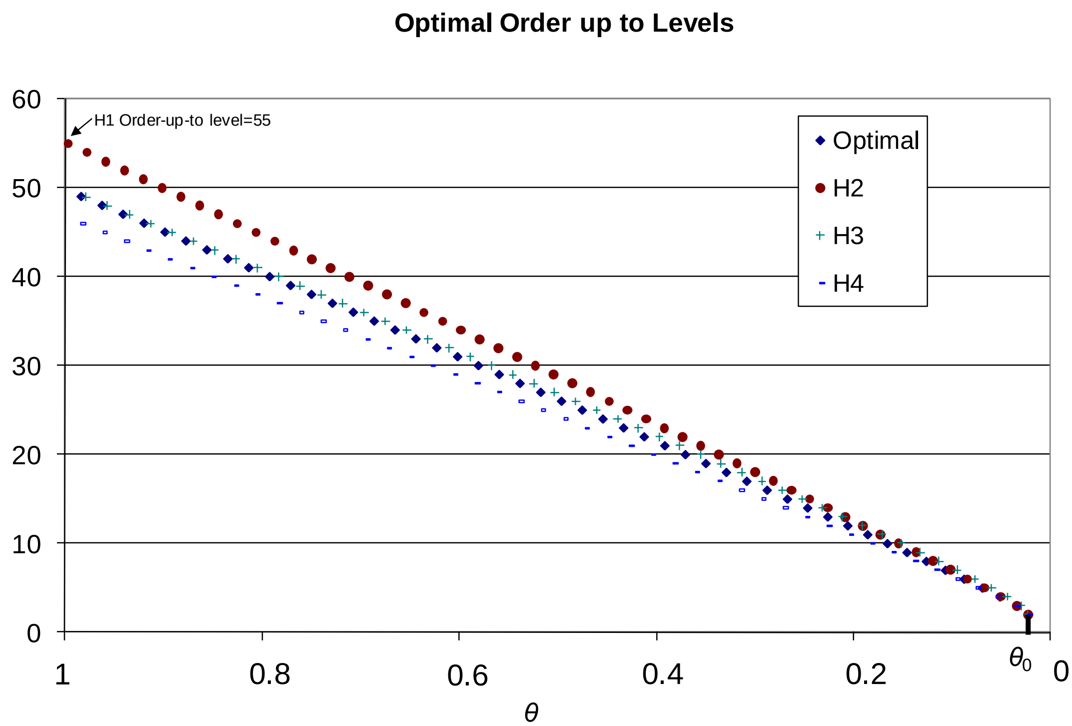

Figure 1: Comparison of heuristics with optimal policy: w=1, π=3, T=1, K=5, λ=50.

As we discussed earlier, H1 and H2 make myopic decisions by ignoring the possibility that the system will be able to reorder if it runs out. As a result, H1 and H2 always order more than the optimal policy. In our experiment, across all cases, on average the order up to level for H1 and H2 at time T is 10.69% more than the optimal order size.

By revising marginal penalty cost from π to something that acknowledges the possibility of reordering, H3 and H4 not only improve the performance of H1 and H2, they also bring the order up to levels close to near-optimal sizes. For the case represented in the graph, H3 order up to levels appear to be closer to the optimal than H4 order up to levels. However, taking an average across all cases, the order up to level at time *T* is 10.54% less than the optimal for H3 and 0.77% less than the optimal for H4.

We now discuss the effect of parameters on the solution. The ordering cost, *K*, clearly has a significant impact on the solution. Intuitively, when *K* is huge, H3, H4, and the optimal policy will make ordering decisions to avoid future orders. In such cases, they behave like H1 and H2, and the differences will diminish. This is confirmed by results in Table 2.

K	1	5	25
H1	196.48%	38.46%	3.33%
H2	20.13%	1.10%	0.00%
H3	5.40%	0.22%	0.01%
H4	0.35%	0.55%	0.00%

Table 2: Impact of *K* on the Average Cost Deviation from the Optimal Policy.

The impact of understocking and overstocking costs is similar. The optimal policy, as well as H3 and H4, recognize future reordering opportunities in making ordering decisions. H1 and H2, however, tend to over-order because they don't account for future ordering opportunities. The over-ordering becomes more severe as the newsvendor ratio $\pi/(\pi+w)$ becomes larger. This is also confirmed by the results in Table 3.

$\pi / (\pi+w)$	1/3	0.5	0.75	0.9
H1	31.01%	48.90%	94.49%	146.36%
H2	0.72%	2.35%	8.33%	16.98%
H3	0.16%	2.12%	2.52%	2.71%
H4	0.12%	0.19%	0.19%	0.72%

Table 3: Impact of $\pi/(\pi+w)$ on the Average Cost Deviation from the Optimal Policy.

The impact of the demand rate is given in Table 4. As λ increases over the fixed time horizon (T=1), under the optimal policy the number of orders are expected to go up. Again, because H1 and H2 do not account for future ordering opportunities, this drives up their cost deviation from the optimal policy. H3 and H4 improve on H2 by incorporating future orders, but they only recognize the first future order and ignore the rest (recall that we approximate the first future reorder point by θ_0). Therefore, as the number of orders goes up, their deviations from the optimal policy also go up. However, it should be noted that even for λ=200, H3 and H4 perform quite well.

λ	50	100	200
H1	56.95%	75.71%	111.32%
H2	5.39%	6.81%	9.87%
H3	1.13%	1.78%	2.88%
H4	0.22%	0.30%	0.41%

Table 4: Impact of λ on the Average Cost Deviation from the Optimal Policy.

4.2. The Supplier's Problem

Next, we study the impact of the retailer's various policies on the performance of the supplier. Note that all the heuristics in Section 2, H1-H4, are developed with only one objective: to minimize the retailer's cost (or, equivalently, to maximize retailer's profit). Therefore, it is especially interesting to see how they impact the supplier. To underscore the fact that the optimal policy in Section 4.1 is optimal only for the retailer, we will call it the retailer-optimal policy in this section.

Recall that the supplier solves a newsvendor problem. We assume that the supplier has enough data to know the distribution of each retailer's total order in the period. For example, suppliers can achieve this by knowing the end demand pattern and which heuristic the retailers are using. The supplier places only one order at the beginning of the period, and its total demand is the aggregate order quantity across all of the retailers. However, different policies used by the retailers will result in different demand distributions at the supplier, so in order to compare the supplier and system performance across various policies, we will use profit as the measure.

If we denote the supplier's demand by D_S, then the supplier maximizes its profit $= p\mathbf{E}[D_S] - \mathrm{TC}_{D_S}(0, T)$, where p is the unit wholesale price and TC is the newsvendor function as defined in equation (1) when the demand is D_S.

The 36 problem instances we use are the same as those used in Section 4.1. In addition, we assume there are eight identical retailers. We assume that the supplier simplifies the calculations by assuming that the total demand follows a normal distribution (with eight identical retailers, this is a reasonable assumption). In each problem instance, we use the order-up-to quantities developed in Section 4.1 for the heuristics. The supplier then finds the first two moments of its total order and uses the newsvendor model to calculate its own order quantity.

For supplier-specific parameters, we assume the following values for unit overstocking (w_0) and understocking (π_0) costs: $w_0 = w = 1$, and $\pi_0/\pi = 0.5$, 1, or 2 representing cases where the supplier's understocking cost is lower than, equal to, or higher than the retailer's understocking penalty cost per unit. Finally, to calculate profits, we let the wholesale price p be 125%, 150%, and 200% of w_0.

The results across all the possible cases are presented in Table 5. For each heuristic, we compute the profit deviation from the retailer-optimal policy, both for the supplier and for the system.

	Supplier Profit			Total System Profit		
	max dev.	min dev.	mean dev.	max dev.	min dev.	mean dev.
H1	81.71%	-6.55%	7.31%	2.41%	-4.68%	-0.35%
H2	65.10%	0.00%	5.60%	3.49%	0.00%	0.92%
H3	6.70%	-34.24%	-1.85%	1.64%	-0.73%	0.12%
H4	5.82%	-28.89%	-2.89%	0.16%	-1.38%	-0.15%

Table 5: Profit Deviation from the Retailer-Optimal Policy.

It is interesting to note the following:

- While the retailer-optimal policy naturally dominates H1-H4 for the retailers, H1-H2 mostly outperform the retailer-optimal policy for the supplier.
- Going from H1 to H4, while the retailers' performance improves, the supplier's profit decreases in general – in this situation, what's good for the retailers is bad for the supplier.
- The result for total system profit is mixed.
    - Using H1, retailers do not take advantage of multiple order opportunities, thus resulting in much higher retailer cost than the retailer-optimal policy. This also means, however, that the retailer's order to the supplier is a fixed number without variance (i.e., it is determined at the beginning of the period, not subject to the realization of random demand). This reduces the supplier's overstocking and understocking costs to zero. Overall, however, these two costs offset each other, and the total system performance of H1 is very close to that of the retailer-optimal policy. For H3 and H4 the opposite is true: while they perform well for the retailers they result in lower profit for the supplier. Their overall total system profit is close to that of the retailer-optimal policy.
    - H2 is quite different. Because H2 increases the supplier's profit much more significantly than H3-H4, this more than compensates for the slightly bigger retailer cost increase. Consequently, H2 results in the highest system profit overall.

It is instructive to contrast the retailer-optimal policy and the heuristics based on it (H3 and H4), with the heuristics based on the newsvendor model (H1 and H2). With the former group, the goal is to optimize retailer performance. This is achieved through dynamically adjusting retailer order size over time. While this dynamic adjustment makes the retailer better off, it creates higher demand uncertainty at the supplier, who makes a newsvendor order decision. It is this increased demand variability that can sometimes significantly reduce the supplier's profit. H1, on the other hand, performs poorly for the retailer because it foregoes the multiple order opportunities, but it works best for the supplier. We can see this interesting tension between what's good for the retailer and what's good for the supplier. It seems that H2 is a good policy that strikes the balance: It is based on H1 but it also modifies H1 to account for multiple order opportunities. Among the heuristics we study, it is indeed the best in terms of total system profit. The overall optimal policy for the centralized system (i.e., one that maximizes total system profit) is unknown, and would indeed be a good topic for future research.

5. Conclusion and Future Research

In this paper we analyze a finite-horizon distribution system with one supplier and multiple retailers. Our work applies to a wide variety of situations in which the supplier must get all supply before the season starts, but the retailers can place orders any time during the season. The retailers must carefully decide when to reorder and how much to order each time by balancing the fixed ordering cost against overstocking and understocking penalty costs. We show that the optimal choice is a time-based, order-up-to policy in which the order-up-to level decreases over time. Moreover, we are able to calculate the optimal order-up-to level at any time and find the time break points at which order-up-to levels change. This completely characterizes the optimal retailer policy. It is worth noting that even though we present our results in terms of homogeneous Poisson demand, they apply to non-homogenous Poisson demand as well. This is especially important for products with a short life cycle, where demand often changes over time (e.g., demand diminishes towards the end of the season). For practical purposes, we also propose heuristics H3 and H4, which are easier to calculate.

The supplier faces a classic newsvendor problem. We express the probability distribution of supplier's demand. However, evaluation of the demand distribution, which involves conditioning over time intervals and convolutions across retailers, can be complex.

Our numerical experiment results in interesting findings. First, we show that by having multiple order opportunities within the same season, retailers can significantly reduce their costs (H2 versus H1). If retailers use more sophisticated ordering policies (H3, H4, and optimal), they can achieve further cost reductions. When retailers place orders using a myopic model, such as the newsvendor model H1, the model overestimates understocking cost because it assumes that once the current stock runs out, shortages will occur. The heuristics we propose, H3 and H4, revise the understocking cost to reflect the fact that future orders can be placed to mitigate shortages. They include only the effect of the first possible future order, but they already achieve most of the improvements. Therefore, we conclude that they are effective heuristics.

Our numerical experiment also reveals that good policies for retailers may not be good policies for the supplier, and possibly the whole system. A good policy, such as the retailer-optimal policy, H3, or H4, will adjust order quantities as demand unfolds over time. This results in uncertain order quantities by the retailers (i.e., uncertain demand at the supplier). On the other hand, the newsvendor policy H1 or its improved version, H2 reduces supplier cost by lowering the variation in the order quantity. This tension between retailer and supplier is reflected in the total system profit. Even though H1 results in no supplier overstocking and understocking costs, its retailer costs are so high that overall it is still a poor policy for the system. H2, on the other hand, can increase supplier profit enough that it sometimes results in a higher system profit than the retailer-optimal policy. This comparison of system profit is made in a decentralized setting. Nevertheless, it points out the supply chain inefficiencies that can be present. In the centralized setting, the optimal policy may lie between the order-for-order and the newsvendor policies. That is, the optimal policy in the centralized setting must balance retailer and supplier profits. This is an interesting topic for future research.

So far, we have not included the holding cost and have considered only the understocking and overstocking costs and fixed ordering cost. This is reasonable in the context of the newsvendor setting because the horizon under which the problem is considered is usually short, making the holding cost negligible. Whether the structural results obtained in this paper extend to holding cost is another interesting topic for future research.

Another direction for future research is to consider joint replenishment in a centralized setting. When one retailer places an order, it may be beneficial to replenish other retailers as well, even though they may not need the additional inventory at that time due to savings in ordering and shipping costs. In such a case, it will be interesting to find out what triggers a joint replenishment and which retailers to replenish.

Acknowledgements

The authors acknowledge the referees for their many suggestions that helped to improve the paper substantially. In particular, one of the referees' contribution to Propositions 1 and 2 is gratefully acknowledged. We thank another referee for suggesting a clever alternative approach.

References:

Arrow, K.J., T. Harris, J. Merchak. 1951. Optimal Inventory Policy. *Econometrica*, 19(3) 250-272.

Axsater, S. 1990. Modeling Emergency Lateral Transshipments in Inventory Systems. *Management Science*, 36(11) 1329-1338.

Axsater, S. 2007. A Heuristic for Triggering Emergency Orders in an Inventory System. *European Journal of Operational Research*, 176 880-891.

Barnes-Schuster, D., Y. Bassok, R. Anupindi. 2002. Coordination and Flexibility in Supply Chain Contracts with Options. *Manufacturing and Service Operations Management*, 4(3) 171-207.

Cachon, G.P., R. Swinney. 2011. The Value of Fast Fashion: Quick Response, Enhanced Design, and Strategic Consumer Behavior. *Management Science*, 57(4) 778-795.

Deuermeyer, B.L., L.B. Schwarz. 1981. A Model for the Analysis of System Service Level in Warehouse-Retailer Distribution Systems: The Identical Retail Case. In Multi-Level Production/Inventory Control Systems: Theory and Practice, edited by L.B. Schwarz, 163-194.

Eppen, G.D., A.V. Iyer. 1997. Backup Agreements in Fashion Buying-The Value of Upstream Flexibility. *Operations Research*, 43(11) 1469-1484.

Federgruen, A., P. Zipkin. 1985. Computing Optimal (s,S) Policies in Inventory Models with Continuous Demand. *Advances in Applied Probability*, 17(2) 424-442.

Fisher, M.A., A. Raman. 1996. Reducing the Cost of Demand Uncertainty through Accurate Response to Early Sales. *Operations Research*, 44(1) 87-99.

Fisher, M.A., J.H. Hamond, W.R. Obermeyer, A. Raman. 1994. Making Supply Meet Demand in an Uncertain World. *Harvard Business Review*, May-June 83-93.

Frazier, R.M. 1986. Quick Response in Soft Lines Discount. Merchand. January, 42-56.

Gallego, G., I. Moon. 1993. The Distribution Free Newsboy Problem: Review and Extensions. *The Journal of the Operational Research Society*, 44(8) 825-834.

Ghemawat, P., J.L. Nueno. 2003. Zara: Fast Fashion. *Harvard Business School* 9-703-497.

Graves, S.C. 1985. Multi-Echelon Inventory Model for a Repairable Item with One-for-One Replenishment. *Management Science*, 31, 1247-1256.

Gurnani, H., C.S. Tang. 1999. Optimal Ordering Decisions with Uncertain Costs and Demand Forecast Updating. *Management Science*, 45(10) 1456-1462.

Hadley, G., T.M. Whitin. 1963. Analysis of Inventory Systems. Prentice Hall.

Hammond, J.H. 1990. Quick Response in Apparel Industry. *Harvard Business School Note* N9-690-038, Cambridge, Mass.

Iyer, A.V., M.E. Bergen. 1997. Quick Response in Manufacturer-Retailer Channels. *Management Science*, 43(4) 559-570.

Khouja, M. 1996. A Note on the Newsboy Problem with Emergency Supply Option. *Journal of Operational Research Society*, 47 1530-1534.

Lau, H.S., H.L. Lau. 1997. A Semi-Analytical Solution for a Newsboy Problem with Mid-Period Replenishment. *The Journal of the Operational Research Society*, 48(12) 1245-1253.

Law, A.M., W.D. Kelton. 2000. Simulation Modeling and Analysis. McGraw Hill. New York.

Lee, H.L., V. Padmanabhan, S. Whang. 2004. Information Distortion in a Supply Chain: The Bullwhip Effect. *Management Science*, 50(12_supplement) 1875-1886.

Liu, F., J.-S. Song. 2011. Good and Bad News about the (S,T) Policy. *Manufacturing and Service Operations Management*, Article in Advance, published online October 28, 2011.

Moinzadeh, K., C. Schmidt. 1991. An (S-1,S) Inventory System with Emergency Orders. *Operations Research*, 2 308-321.

Moinzadeh, K., Y. Zhou. 2008. Incorporating Delay Mechanism in Ordering Policies in Multi-Echelon Distribution Systems. *IIE Transactions*, 40(4) 445-458.

Nahmias, S. 2009. Production and Operations Analysis. McGraw-Hill Irwin, New York. Sixth Edition.

Rudi, N., S. Kapur, D. F. Pyke. 2001. A Two-Location Inventory Model with Transshipment and Local Decision Making. *Management Science*, 47(12) 1668-1680.

Russell, S. 2004. Problems Have Beset Flu Program for Years as Vaccinations Zoom, Production Bottleneck Worsens, *San Francisco Chronicle.*

Scarf, H. 1960. The Optimality of (s,S) Policies in the Dynamic Inventory Problem. In *The Mathematical Methods in the Social Sciences* edited by K.A. Arrow, S. Karlin and P. Suppes, Stanford University Press, Stanford, CA.

Shao, J., H. Krishnan, S.T. McCormick. 2011. Incentives for Transshipment in a Supply Chain with Decentralized Retailers. *Manufacturing and Service Operations Management*, 13(3) 361-372.

Skinner, R.C. 1992. Fashion Forecasting at Oxford Shirting. *Proc. Quick Response Conference*, Chicago, IL. 90-107.

Stidham, S. 1986. Clearing Systems and (s,S) Inventory Systems with Non-linear Costs and Positive Lead Times. *Operations Research*, 34 276-280.

Tsay, A.A. 1999. The Quantity Flexibility Contract and Supplier-Customer Incentives. *Management Science*, 45(10) 1339-1358.

Whitin, T.M. 1955. Inventory Control and Price Theory. *Management Science*, 2 61-68.

Zheng, Y.S., A. Federgruen. 1991. Finding Optimal (s,S) Policies Is about as Simple as Evaluating a Single Policy. *Operations Research*, 39 654-665.

Svoronos, A., P. Zipkin. 1988. Estimating the Performance of Multi-Level Inventory Systems. *Operations Research*, 36(1), 57-72.

Zipkin, P. 2000. Foundations of Inventory Management. McGraw-Hill, New York.

APPENDIX

Proof of Lemma 1 Recall that as the optimal newsvendor quantity, $\bar{S}(\theta)$ is the largest integer S such that $P(S,\theta) \geq \frac{w}{w+\pi}$. Let's examine the properties of $P(n,\theta)$ with respect to n and θ for $0 \leq \theta \leq T$. Clearly $P(0,\theta) = 1$ for all θ. Now suppose n>0. If

1. $P(1,T) < \frac{w}{w+\pi}$, then $P(n,\theta) < \frac{w}{w+\pi}$ for all n and θ as $P(n,\theta)$ is decreasing in n and increasing in θ. Thus $\bar{S}(\theta) = 0$ for all $0 \leq \theta \leq T$.

2. $P(1,T) \geq \frac{w}{w+\pi}$, then since $P(n,T) \to 0$ as $n \to \infty$, there must exist an $l \geq 1$ such that $P(n,T) \geq \frac{w}{w+\pi}$ for $1 \leqslant n < l$, and $P(n,T) < \frac{w}{w+\pi}$ for $n \geq l$. Because $P(n,\theta)$ is increasing in θ, we conclude that $P(n,\theta) < \frac{w}{w+\pi}$ for all $n \geq l$ and $0 \leq \theta \leq T$. Now we will focus on $1 \leqslant n < l$. Since $P(n,\theta)$ function is continuous in θ and approaches 0 as $\theta \to 0$, let $\bar{\theta}_{n-1} \in (0,T]$ be the unique solution to

$$P(n,\theta) = \frac{w}{w+\pi}. \tag{A.1}$$

In this case, $\bar{S}(\bar{\theta}_{n-1}) = n$. Note that the left side of (A.1) is decreasing in n and increasing in θ.

Therefore, we conclude $\bar{\theta}_n > \bar{\theta}_{n-1}$ for $1 \leqslant n < l$. For any θ in $[\theta_{n-1}, \theta_n)$, $\bar{S}(\theta) = n$.

Hence, we have proved that there exist time points $0 = \bar{\theta}_{-1} < \bar{\theta}_0 < \bar{\theta}_1 < \ldots < \bar{\theta}_{l-1} < \bar{\theta}_l = T$ for some $l \geq 0$ such that $\bar{S}(\theta) = 0$ on $[\bar{\theta}_{-1}, \bar{\theta}_0)$, $\bar{S}(\theta) = n$ on $[\bar{\theta}_{n-1}, \bar{\theta}_n)$, $\forall 1 \leq n < l$, and $\bar{S}(\theta) = l$ on $[\bar{\theta}_{l-1}, \bar{\theta}_l]$. □

Proof of Lemma 2 On any interval $[\bar{\theta}_n, \bar{\theta}_{n+1})$, the optimal newsvendor stocking level $\bar{S}$ is constant and $g(\theta) = TC(\bar{S}, \theta)$. From (1), we have $\frac{\partial g(\theta)}{\partial \theta} = -w\lambda + (w+\pi)\lambda P(\bar{S}, \theta) \geq 0$. Since $g(\theta)$ is continuous and non-decreasing in θ on each interval $[\bar{\theta}_n, \bar{\theta}_{n+1})$, then it is non-decreasing in θ on the whole interval $[0,T]$. □

Proof of Proposition 1 Let $p > 0$ be the probability that an order is placed when inventory is positive under policy φ. Construct a new policy φ' as follows. Whenever policy φ places an order while inventory is positive, policy φ' delays that order until the next demand takes place. Moreover, if the end of the season is reached before the next demand occurs, policy φ' does not place an order; if policy φ orders again when the next demand occurs, the two orders are combined. Then on any sample path, the two policies will have the exact same number of lost sales, and the ending inventory under policy φ' will be no more than that under policy φ. Moreover, each time policy φ places an order when inventory is positive, there is a positive probability that the next demand will not occur by the end of the season, thus saving one order for policy φ'. Due to the Poisson arrival, this probability is at least $e^{-\lambda T}$. Therefore, E[number of orders placed under φ'] $\leq$ E[number of orders placed under φ] $- pe^{-\lambda T}$. It follows that the total expected cost incurred by policy φ is strictly larger than that under policy φ', hence policy φ is not optimal. □

Proof of Proposition 2 In view of Proposition 1, we can assume that policy φ only orders when inventory is zero. Construct a new policy φ' as follows. Policy φ' places the same orders as policy φ, except that when policy φ places an order when no demand takes place, policy φ' delays that order until the next demand takes place (or does not place an order at all if no further demand takes place during the season). Ending inventory is no larger under policy φ' than under policy φ, lost sales are identical for both policies, and there is a probability of at least $pe^{-\lambda T}$ that policy φ' places one fewer order than policy φ. It again follows that the total expected cost for policy φ is strictly larger than that for policy φ', so policy φ is not optimal. □

Proof of Proposition 3 The case of $K \le \pi$ is easy to see, so we will focus on the case of $\pi < K < \pi(\lambda T + 1)$. As $\pi < K$ and $\pi(\lambda T + 1) > K$, there must exist a $\upsilon_0 \in (0, T)$ such that:

$$\pi(\lambda \upsilon_0 + 1) = K \ . \tag{A.2}$$

Thus, for $\upsilon \le \upsilon_0$ the ordering cost alone at time υ exceeds the expected cost of not ordering at all (i.e., all the future demands are lost). So the optimal action at υ is to not order. Next, consider:

$$g(\upsilon_0) + K = \pi(\lambda \upsilon + 1) . \tag{A.3}$$

When $\upsilon = 0$, the right hand side of (A.3) is less than the left hand side, but as $\upsilon \to \infty$, the right hand side exceeds the left hand side. Furthermore, the right hand side is continuous and strictly increasing in υ. Thus, there exists a unique solution, $\upsilon_1 > 0$, to (A.3). Combining (A.2) and (A.3), we get $\upsilon_1 > \upsilon_0$. Next, consider any time $\upsilon \in [\upsilon_0, \upsilon_1]$. Suppose inventory is zero and a demand occurs at time υ. There are two possible sets of actions:

1. We place an order, and an ordering cost of K is incurred. Then, for any inventory level at time υ_0, the future expected total cost from that point on is at least, $g(\upsilon_0)$, because by definition $g(\upsilon_0)$ is the minimum cost at time υ_0. So the total expected future cost is at least $K + g(\upsilon_0)$.
2. We do not place an order. We lose the immediate demands and incur a cost of π right away. Then the future total cost is at most $\pi\lambda\upsilon$, when no future orders are placed and all future demands are lost.

For $\upsilon \in [\upsilon_0, \upsilon_1]$, because $g(\upsilon_0) + K = \pi(\lambda\upsilon_1 + 1) > \pi(\lambda\upsilon + 1)$, it is clear that the optimal action is to not place any order.

We note that if $\upsilon_1 > T$, then we have indeed characterized an optimal policy on [0,*T*]. Therefore, from now on, we consider $\upsilon_1 \leq T$.

We now repeat this process to define υ_2 as the solution to $g(\upsilon_1) + K = \pi(\lambda \upsilon_2 + 1)$. Since $g(\upsilon_0) + K = \pi(\lambda \upsilon_1 + 1)$, by Lemma 2, we must have $\upsilon_2 > \upsilon_1$. Let $\upsilon \in [\upsilon_1, \upsilon_2]$, and suppose inventory is zero and a demand occurs at time υ . Again, there are two possible actions:

1. We place an order, and an ordering cost of *K* is immediately incurred. Because the expected total future cost is at least $g(\upsilon_1)$, the total cost associated with this action is at least $g(\upsilon_1) + K$.
2. We don't place an order. Then similar to the previous argument, we can show that the expected total future cost is at most $\pi(\lambda \upsilon + 1)$.

Again, for $\upsilon \in [\upsilon_1, \upsilon_2]$, $g(\upsilon_1) + K = \pi(\lambda \upsilon_2 + 1) \geq \pi(\lambda \upsilon + 1)$ and we conclude that the optimal action at time υ is to not place any order.

We can continue this process to define $\upsilon_3 < \upsilon_4 < \ldots$ which is an increasing series. If at any point the series goes above (or equal to) *T*, then it is optimal to never order at any time. In such a case, we simply set $\theta_0 = T$. On the other hand, if the series is always below *T*, then let its limit be θ_0 . It must satisfy $0 < \theta_0 \leqslant T$. In this case, we have just shown that the optimal action is to never place an order when the time left is less than θ_0 . That is, any policy that never orders on $[0, \theta_0)$ is an optimal policy on $[0, \theta_0)$.

By taking the limit of $g(\upsilon_n) + K = \pi(\lambda \upsilon_{n+1} + 1)$ (note that *g* is continuous), we get $g(\theta_0) + K = \pi(\lambda \theta_0 + 1)$. Therefore, θ_0 is a solution to $g(\theta) + K = \pi(\lambda \theta + 1)$. Moreover, from the proof of Lemma 2 we know that on any interval $[\overline{\theta}_i, \overline{\theta}_{i+1})$,

$$\frac{\partial\left\{g(\theta)+K-\pi\left(\lambda\theta+1\right)\right\}}{\partial\theta}=(w+\pi)\lambda\left(P(\bar{S},\theta)-1\right)<0\ . \qquad \text{(A.4)}$$

Thus, $g(\theta)+K-\pi(\lambda\theta+1)$ is strictly monotone and θ_0 is the unique solution to $g(\theta)+K=\pi(\lambda\theta+1)$.

As described in Section 2, the optimal value function at time θ is represented by $V(i,\theta)$ (recall that for convenience, we drop the * superscript). Also, define $\Delta V(i,\theta)=V(i+1,\theta)-V(i,\theta)$. Since it is optimal to never order at $\theta<\theta_0$, $V(i,\theta_0)$ equals the newsvendor cost function $TC(i,\theta_0)$, and thus is convex in *i*. Specifically, because by definition $S_0=\bar{S}(\theta_0)$, the optimal newsvendor order quantity with θ_0 time left,

- $\Delta V(i,\theta_0)\le 0$ for $i<S_0$.
- $\Delta V(i,\theta_0)>0$ for $i\geqslant S_0$,
- $\Delta V(i,\theta_0)$ is increasing for $i\geqslant S_0$.

Furthermore, since no order will be placed at $\theta<\theta_0$, we know $V(0,\theta_0)=\pi\lambda\theta_0$. Moreover,

$$V(S_0,\theta_0)=g(\theta_0)=\pi\left(\lambda\theta_0+1\right)-K=V(0,\theta_0)+\pi-K\ , \qquad \text{(A.5)}$$

where the first equality comes from the definition of S_0, and we have just proved the second equality above. So, $V(0,\theta_0)-V(S_0,\theta_0)-K=-\pi<0$. Finally, we have

$$\begin{aligned}
&\Delta V(i,\theta_0)=V(i+1,\theta_0)-V(i,\theta_0)=TC(i+1,\theta_0)-TC(i,\theta_0)\\
&=wE\left[i+1-D(\theta_0)\right]^{+}+\pi E\left[D(\theta_0)-(i+1)\right]^{+}-wE\left[i-D(\theta_0)\right]^{+}-\pi E\left[D(\theta_0)-i\right]^{+}\\
&=w\left(\sum_{j=0}^{i}(i+1-j)\,p(j,\theta_0)-\sum_{j=0}^{i-1}(i-j)\,p(j,\theta_0)\right)+\pi\left(\sum_{j=i+2}^{\infty}(j-i-1)\,p(j,\theta_0)-\sum_{j=i+1}^{\infty}(j-i)\,p(j,\theta_0)\right)\\
&=w\sum_{j=0}^{i}p(j,\theta_0)-\pi\sum_{j=i+1}^{\infty}p(j,\theta_0)=w\left(1-P(i+1,\theta_0)\right)-\pi P(i+1,\theta_0).
\end{aligned}$$

Because $P(i+1,\theta_0)<1$, it follows that $\Delta V(i,\theta_0)>-\pi$ for $i<S_0$.

In summary, we have proved:

$$-\pi < \Delta V(i,\theta_0) \le 0 \text{ for } i < S_0 \tag{A.6}$$

$$\Delta V(i,\theta_0) > 0 \text{ for } i \geqslant S_0 \tag{A.7}$$

$$\Delta V(i,\theta_0) \text{ is increasing for } i \geqslant S_0 \tag{A.8}$$

$$V(0,\theta_0) < V(S_0,\theta_0) + K \tag{A.9}$$

$$V(0,\theta_0) + \pi = V(S_0,\theta_0) + K \tag{A.10}$$

At time θ_0, if a demand occurs, there are three possible scenarios depending on the level of on-hand inventory, $I(\theta_0)$:

a) $I(\theta_0) \ge S_0 + 1$. The immediate demand will be satisfied from on-hand inventory, and based on (A.7) the optimal action is to not order.

b) $1 \le I(\theta_0) \le S_0$. The immediate demand will be satisfied from on-hand inventory. Moreover, if an order is placed, the order up to level is S_0, and the total cost till the end of the season will be $K + V(S_0,\theta_0)$. If an order is not placed, the total cost till the end of the season will be $V(I(\theta_0) - 1, \theta_0)$. Based on (A.6) and (A.9), the latter cost is lower. Thus, the optimal action is to not place an order.

c) $I(\theta_0) = 0$. The optimal cost until time 0 will be $K + V(S_0,\theta_0)$ if an order is placed, and $\pi + V(0,\theta_0)$ if an order is not placed. Based on (A.10), the costs of the two actions are equal. We have consistently picked the larger value in such cases, so we will say that the optimal order-up-to level is S_0.

This concludes the proof of Proposition 3. □

Proof of Proposition 4 The proof is rather long and we start with a brief roadmap for the proof.

Step 1: We start with rigorously defining the existence of a point θ_1 such that $\theta_1 = \sup\{\theta_0 < \theta \le T \mid$ if a demand occurs on $[\theta_0, \theta)$ and inventory is zero, the optimal order-up-to level is $S_0\}$. This allows us to show the reduction of value function recursion (2) to equation (3) for $\theta_0 < \theta \le \theta_1$.

Step 2: We prove some properties (see (A.14)-(A.17) below) of $V(i,\theta)$ for $\theta_0 < \theta \le \theta_1$. This step is further sub-divided into the proof of (A.14) through (A.17) separately.

Step 3: We prove an additional property (see (A.23) below) of $V(i,\theta)$ for $\theta = \theta_1$ which will then allow us to prove the optimal order-up-to level at $\theta = \theta_1$.

Step 1: Definition of θ_1

We begin by showing the existence of a neighborhood of θ_0 such that the optimal order-up-to level at θ_0 continues to remain optimal in this neighborhood. This will allow us to rigorously define θ_1 and prove that $\theta_1 > \theta_0$. Formally, we prove the following:

> There exists an $\varepsilon > 0$ such that if a demand occurs at any $\theta \in [\theta_0, \theta_0 + \varepsilon]$ and the inventory is zero, the optimal order-up-to level is S_0.

To demonstrate this, consider the value function definition in (2). Under any optimal policy, both lost sales and the number of orders are bounded by the demand during the same time period almost surely. The former is obvious and the latter holds because it's optimal to order only at demand epochs almost surely. Therefore, we must have $\lim_{\theta \to \theta_0^+} E_\omega[N_\omega(i, \theta_0, \theta)] = 0$ and $\lim_{\theta \to \theta_0^+} E_\omega[L_\omega(i, \theta_0, \theta)] = 0$. Furthermore, since the probability of any demand occurring between θ and θ_0 goes to zero as $\theta \to \theta_0$, we must also have: $\lim_{\theta \to \theta_0^+} E_\omega [V(i,\theta_0)] = V(i,\theta_0)$. Hence, $\lim_{\theta \to \theta_0^+} V(i,\theta) = V(i,\theta_0)$. That is, the value function $V(i,\theta)$ is right-continuous at θ_0.

By definition of S_0, $V(S_0,\theta_0) < V(S,\theta_0), \forall S > S_0$. Now, there exists a neighborhood $[\theta_0, \theta_0 + \varepsilon]$ for some $\varepsilon > 0$, where $V(S_0,\theta) < V(S,\theta), \forall S > S_0$ for all $\theta \in [\theta_0, \theta_0 + \varepsilon]$. Therefore, at any point in $[\theta_0, \theta_0 + \varepsilon]$, if the inventory is zero and a demand occurs, either ordering up to some $S \leq S_0$ or never ordering will minimize the value function. Below, we will show that the optimal order-up-to level is always S_0 during $[\theta_0, \theta_0 + \varepsilon]$. To that end, let Policy O be the policy that order up to S_0 during $[\theta_0, \theta_0 + \varepsilon]$, if inventory is zero and a demand occurs. Moreover, let Policy P be any other policy that violates this. That is, there exists time epochs on $[\theta_0, \theta_0 + \varepsilon]$ where, if inventory is zero and a demand occurs, Policy P chooses to either not order and incur the lost sale, or order up to some $S < S_0$. It suffices to show that the system cost is higher under Policy P than under Policy O. For the rest of the proof, we will use superscript ^{O} to represent Policy O and superscript ^{P} to represent Policy P.

We now need to define some notations. Let ω be any sample path on $[\theta_0, \theta_0 + \varepsilon]$. Define:

- $m^O(\omega)$: the number of orders of size $S_0 + 1$ placed on $[\theta_0, \theta_0 + \varepsilon]$ under Policy O (i.e., when inventory is zero and a demand occurs, order up to $S_0 + 1$).
- $m_j^P(\omega)$: the number of orders of size $j + 1$ placed on $[\theta_0, \theta_0 + \varepsilon]$ under Policy P (i.e., when inventory is zero and a demand occurs, order up to j), $0 \leq j \leq S_0$.
- $i^z(\omega)$: on hand inventory at time θ_0 under Policy $z \in \{O, P\}$. Note that $0 \leq i^z(\omega) \leq S_0$.

  Moreover, let $l^z(\omega)$ be the units of demand that are lost on $[\theta_0, \theta_0 + \varepsilon]$ under policy z. Because both policies face the same demand, we must have

$$\sum_{j=0}^{S_0} \left[m_j^P(\omega)(j+1) \right] + l^P(\omega) - i^P(\omega) = m^O(\omega)(S_0 + 1) - i^O(\omega).$$

$$\Rightarrow l^P(\omega) = [i^P(\omega) - i^O(\omega)] + \left[m^O(\omega)(S_0 + 1) - \sum_{j=0}^{S_0} \left[m_j^P(\omega)(j+1) \right] \right]. \qquad \text{(A.11)}$$

We calculate the cost of either policy at time $\theta_0 + \varepsilon$ as the sum of the sample-path cost on $[\theta_0, \theta_0 + \varepsilon]$ and the expected cost on $[0, \theta_0)$. That is, the total cost at time $\theta_0 + \varepsilon$ for Policies O and P, based on the sample path ω, are $K \sum_{j=0}^{S_0} m_j^P(\omega) + \pi l^{\mathrm{P}}(\omega) + V(i^{\mathrm{P}}(\omega), \theta_0)$ and $m^O(\omega)K + V(i^O(\omega), \theta_0)$, respectively. Furthermore, the cost difference between these policies will be:

$$\Delta \text{ Cost} = \left[K \sum_{j=0}^{S_0} m_j^P(\omega) + \pi l^{\mathrm{P}}(\omega) + V\big(i^{\mathrm{P}}(\omega), \theta_0\big)\right] - [m^O(\omega)K + V(i^O(\omega), \theta_0)\,]$$

$$= K\left[\sum_{j=0}^{S_0} m_j^P(\omega) - m^O(\omega)\right] + \pi l^{\mathrm{P}}(\omega) + \left[V\big(i^{\mathrm{P}}(\omega), \theta_0\big) - V(i^O(\omega), \theta_0)\,\right]. \qquad \text{(A.12)}$$

Depending on the values of $m(\omega)$ and $l(\omega)$ for the policies, we have three possible cases:

1. When $\sum_{j=0}^{S_0} m_j^P(\omega) = m^{\mathrm{O}}(\omega)$: (A.12) becomes

$$\Delta \text{ Cost} = \pi l^{\mathrm{P}}(\omega) + \left[V\big(i^{\mathrm{P}}(\omega), \theta_0\big) - V(i^O(\omega), \theta_0)\,\right]. \qquad \text{(A.13)}$$

If $i^{\mathrm{P}}(\omega) \le i^O(\omega)$, then $V\big(i^{\mathrm{P}}(\omega), \theta_0\big) \ge V(i^O(\omega), \theta_0)$, since both $i^{\mathrm{P}}(\omega)$ and $i^O(\omega)$ are no more than S_0, and Δ Cost ≥ 0, with the equality holding only if $l^{\mathrm{P}}(\omega) = 0$ and $i^{\mathrm{P}}(\omega) = i^O(\omega)$. This happens with probability less than one.

If $i^{\mathrm{P}}(\omega) > i^O(\omega)$, then from (A.6), we have: $V\big(i^{\mathrm{P}}(\omega), \theta_0\big) - V(i^O(\omega), \theta_0) > -[i^{\mathrm{P}}(\omega) - i^O(\omega)]\pi$, and by employing (A.11), (A.13) can be written as:

$$\Delta \text{ Cost} = \left[\, m^O(\omega)(S_0 + 1) - \sum_{j=0}^{S_0} \left[m_j^P(\omega)(j + 1)\right]\right] \pi$$

$$+ [i^{\mathrm{P}}(\omega) - i^O(\omega)]\pi + \left[V\big(i^{\mathrm{P}}(\omega), \theta_0\big) - V(i^O(\omega), \theta_0)\,\right] > 0.$$

2. When $\sum_{j=0}^{S_0} m_j^P(\omega) > m^{\mathrm{O}}(\omega)$: Since, by Propositions 1 and 2, $V\big(i^{\mathrm{O}}(\omega), \theta_0\big) < K + V(S_0, \theta_0)$, (A.12) becomes:

$$\Delta \text{ Cost} = K\left[\sum_{j=0}^{S_0} m_j^P(\omega) - m^O(\omega)\right] + \pi l^{\mathrm{P}}(\omega) + \left[V\big(i^{\mathrm{P}}(\omega), \theta_0\big) - V(i^O(\omega), \theta_0)\right]$$

$$> K + \pi l^{\mathrm{P}}(\omega) + \left[V\big(i^{\mathrm{P}}(\omega), \theta_0\big) - K - V(S_0, \theta_0)\right] \ge 0..$$

The last inequality holds because $i^{\mathrm{P}}(\omega) \le S_0$ implies $V(i^{\mathrm{P}}(\omega), \theta_0) \ge V(S_0, \theta_0)$.

3. When $\sum_{j=0}^{S_0} m_j^P(\omega) < m^{\mathrm{O}}(\omega)$: using (A.11), (A.12) will be:

$$\Delta \text{ Cost} = K\left[\sum_{j=0}^{S_0} m_j^P(\omega) - m^O(\omega)\right] + \pi l^{\mathrm{P}}(\omega) + \left[V\left(i^{\mathrm{P}}(\omega), \theta_0\right) - V\left(i^O(\omega), \theta_0\right)\right]$$

$$= K\left[\sum_{j=0}^{S_0} m_j^P(\omega) - m^O(\omega)\right] + \pi\left[m^O(\omega)(S_0+1) - \sum_{j=0}^{S_0}\left[m_j^P(\omega)(j+1)\right]\right]$$

$$+\pi\left[i^{\mathrm{P}}(\omega) - i^O(\omega)\right] + \left[V\left(i^{\mathrm{P}}(\omega), \theta_0\right) - V\left(i^O(\omega), \theta_0\right)\right]$$

$$\geq K\left[\sum_{j=0}^{S_0} m_j^P(\omega) - m^O(\omega)\right] + \pi\,(S_0+1)\left[m^O(\omega) - \sum_{j=0}^{S_0} m_j^P(\omega)\right]$$

$$+\pi\left[i^{\mathrm{P}}(\omega) - i^O(\omega)\right] + \left[V\left(i^{\mathrm{P}}(\omega), \theta_0\right) - V(S_0, \theta_0) - (S_0 - i^O(\omega))\pi\right]$$

(as $V(i^O(\omega), \theta_0) \leq V(S_0, \theta_0) + (S_0 - i^O(\omega))\pi$ from (A.6))

$$> [\pi\,(S_0+1) - K]\left[m^O(\omega) - \sum_{j=0}^{S_0} m_j^P(\omega)\right] + \pi\left[i^{\mathrm{P}}(\omega) - i^O(\omega)\right]$$

$$+\left[V\left(i^{\mathrm{P}}(\omega), \theta_0\right) - V(0, \theta_0) - \pi + K - (S_0 - i^O(\omega))\pi\right]$$

(as $V(S_0, \theta_0) + K = V(0, \theta_0) + \pi$)

$$\geq [\pi\,(S_0+1) - K] + \pi\left[i^{\mathrm{P}}(\omega) - i^O(\omega)\right]$$

$$+\left[V\left(i^{\mathrm{P}}(\omega), \theta_0\right) - V(0, \theta_0) - \pi + K - (S_0 - i^O(\omega))\pi\right]$$

(as $(S_0 + 1) > K$)

$$= V\left(i^{\mathrm{P}}(\omega), \theta_0\right) - V(0, \theta_0) + \pi\left[i^{\mathrm{P}}(\omega)\right] \geq 0,$$

where the last inequality is due to (A.6).

By integrating over all the sample paths ω on $[\theta_0, \theta_0 + \varepsilon]$, it can be verified that the overall cost under Policy O is strictly lower than that under Policy P. Therefore, the optimal order-up-to level is always S_0 when inventory is zero and a demand occurs for all $\theta \in [\theta_0, \theta_0 + \varepsilon]$.

The proof above allows us to define

$\theta_1 = \sup\{ \theta \mid \theta_0 < \theta \le T$, and if a demand occurs on $[\theta_0, \theta)$ and inventory is zero, the optimal action is to order up to $S_0\}$ and know that since the set is non-empty, θ_1 must be strictly larger than θ_0. Now, consider any $\theta \in (\theta_0, \theta_1]$. By the definition of θ_1, there will not be any lost sales between θ and θ_0, so $E[L(i, \theta_0, \theta)] = 0$. Moreover, any order between θ and θ_0 will be of size $S_0 + 1$, so $E[N(i, \theta_0, \theta)] = \sum_{n=0}^{\infty} P(n(S_0+1)+i+1, \theta - \theta_0)$. Therefore, for $\theta_0 < \theta \le \theta_1$,

$$V(i,\theta) = K\sum_{n=0}^{\infty} P\left(n(S_0+1)+i+1, \theta-\theta_0\right) + \sum_{j=0}^{S_0}\sum_{n=1}^{\infty} p\left(n(S_0+1)+i-j, \theta-\theta_0\right)V(j,\theta_0) + \sum_{j=0}^{i} p(j,\theta-\theta_0)V(i-j,\theta_0).$$

This is (3).

Step 2: Properties of $V(i,\theta)$ for $\theta_0 < \theta \le \theta_1$

In what follows, we prove the following properties of $V(i,\theta)$ for $\theta_0 < \theta \le \theta_1$:

$$-\pi < \Delta V(i,\theta) < 0 \quad \text{for } i < S_0 \tag{A.14}$$

$$\Delta V(i,\theta) \text{ is increasing for } i \ge S_0 \tag{A.15}$$

$$V(0,\theta) < V(S_0,\theta) + K \tag{A.16}$$

$$V(0,\theta) + \pi > V(S_0,\theta) + K. \tag{A.17}$$

Proof of (A.14)

For $i \le S_0$ and $\theta_0 < \theta \le \theta_1$, by using $E[N(i,\theta_0,\theta)] = \sum_{n=0}^{\infty} P(n(S_0+1)+i+1, \theta-\theta_0)$ and rearranging terms, (3) can be written as:

$$V(i,\theta) = K\sum_{n=0}^{\infty}\sum_{j=i+1}^{\infty} p\left(n(S_0+1)+j, \theta-\theta_0\right)$$
$$+\sum_{n=0}^{\infty}\left\{\sum_{j=0}^{i} p\left(n(S_0+1)+j,\theta-\theta_0\right)V(i-j,\theta_0) + \sum_{j=i+1}^{S_0} p\left(n(S_0+1)+j,\theta-\theta_0\right)V(S_0+1+i-j,\theta_0)\right\}.$$

Recall that $\Delta V(i-1,\theta) = V(i,\theta) - V(i-1,\theta)$ is the first-order difference function. Then, for $i \le S_0$:

$$
\begin{aligned}
&\Delta V(i-1,\theta) \\
&= \sum_{n=0}^{\infty}\left(\begin{array}{l}
K\sum_{j=i+1}^{\infty} p\left(n(S_0+1)+j,\theta-\theta_0\right) - K\sum_{j=i}^{\infty} p\left(n(S_0+1)+j,\theta-\theta_0\right) \\
+\sum_{j=0}^{i} p\left(n(S_0+1)+j,\theta-\theta_0\right)V(i-j,\theta_0) - \sum_{j=0}^{i-1} p\left(n(S_0+1)+j,\theta-\theta_0\right)V(i-1-j,\theta_0) \\
+\sum_{j=i+1}^{S_0} p\left(n(S_0+1)+j,\theta-\theta_0\right)V(S_0+1+i-j,\theta_0) - \sum_{j=i}^{S_0} p\left(n(S_0+1)+j,\theta-\theta_0\right)V(S_0+i-j,\theta_0)
\end{array}\right).
\end{aligned}
$$

Grouping $V(.,.)$ terms with the same probability weights, we get:

$$
\begin{aligned}
&\Delta V(i-1,\theta) \\
&= \sum_{n=0}^{\infty}\left(\begin{array}{l}
-Kp\left(n(S_0+1)+i,\theta-\theta_0\right) \\
+\sum_{j=0}^{i-1} p\left(n(S_0+1)+j,\theta-\theta_0\right)\left[V(i-j,\theta_0)-V(i-1-j,\theta_0)\right] + p(n(S_0+1)+i,\theta-\theta_0)V(0,\theta_0) \\
+\sum_{j=i+1}^{S_0} p(n(S_0+1)+j,\theta-\theta_0)\left[V(S_0+1+i-j,\theta_0)-V(S_0+i-j,\theta_0)\right] - p(n(S_0+1)+i,\theta-\theta_0)V(S_0,\theta_0)
\end{array}\right) \\
&= \sum_{n=0}^{\infty}\left(\begin{array}{l}
p\left(n(S_0+1)+i,\right)\left[V(0,\theta_0)-V(S_0,\theta_0)-K\right] \\
+\sum_{j=0}^{i-1} p(n(S_0+1)+j,\theta-\theta_0)\left[V(i-j,\theta_0)-V(i-1-j,\theta_0)\right] \\
+\sum_{j=i+1}^{S_0} p(n(S_0+1)+j,\theta-\theta_0)\left[V(S_0+1+i-j,\theta_0)-V(S_0+i-j,\theta_0)\right]
\end{array}\right).
\end{aligned}
$$

The first term in the expression above is negative due to (A.9). Furthermore, for $i \leq S_0$, due to (A.6), we have:

$$V(i-j,\theta_0) - V(i-1-j,\theta_0) \leq 0 \quad \text{for } 0 \leq j \leq i,$$

and for $i+1 \leq S_0$

$$V(S_0+1+i-j,\theta_0) - V(S_0+i-j,\theta_0) \leq 0 \quad \text{for } i+1 \leq j \leq S_0.$$

Together, these imply $\Delta V(i-1,\theta) < 0$ for $i \leq S_0$ and $\theta_0 < \theta \leq \theta_1$.

Now, from (A.10) we have $[V(0,\theta_0) - V(S_0,\theta_0) - K] = -\pi$ and from (A.6) we have $\Delta V(i-1,\theta_0) > -\pi$ for $1 \le i \le S_0$. Thus, the expression for $\Delta V(i-1,\theta)$ above is a weighted average of three terms, one of which is equal to $-\pi$ and other two are strictly greater than $-\pi$. All weights are positive. Therefore, we have $\Delta V(i-1,\theta) > -\pi$ for $i \le S_0$ and $\theta_0 < \theta < \theta_1$. This completes the proof of (A.14).

Proof of (A.16)-(A.17)

Note that for $\theta_0 < \theta \le \theta_1$,

$$V(0,\theta) = K\sum_{n=0}^{\infty}\sum_{j=1}^{\infty} p\left(n(S_0+1)+j,\theta-\theta_0\right)$$
$$+\sum_{n=0}^{\infty}\left(p\left(n(S_0+1),\theta-\theta_0\right)V(0,\theta_0)+\sum_{j=1}^{S_0} p\left(n(S_0+1)+j,\theta-\theta_0\right)V(S_0+1-j,\theta_0)\right),$$

and,

$$V(S_0,\theta) = K\sum_{n=0}^{\infty}\sum_{j=S_0+1}^{\infty} p\left(n(S_0+1)+j,\theta-\theta_0\right)+\sum_{n=0}^{\infty}\left(\sum_{j=0}^{S_0} p\left(n(S_0+1)+j,\theta-\theta_0\right)V(S_0-j,\theta_0)\right).$$

Thus, we have:

$$\begin{aligned}
&V(0,\theta) - V(S_0,\theta) - K \\
&= K\left(\sum_{n=0}^{\infty}\sum_{j=1}^{S_0} p\left(n(S_0+1)+j,\theta-\theta_0\right)-1\right) \\
&+\sum_{n=0}^{\infty}\left(\begin{array}{l} p\left(n(S_0+1),\theta-\theta_0\right)V(0,\theta_0)+\sum_{j=1}^{S_0} p\left(n(S_0+1)+j,\theta-\theta_0\right)V(S_0+1-j,\theta_0) \\ \\ -\sum_{j=0}^{S_0} p\left(n(S_0+1)+j,\theta-\theta_0\right)V(S_0-j,\theta_0) \end{array}\right) \\
&= \sum_{n=0}^{\infty} p\left(n(S_0+1),\theta-\theta_0\right)\left[V(0,\theta_0)-V(S_0,\theta_0)-K\right]+\sum_{n=0}^{\infty}\sum_{j=1}^{S_0} p\left(n(S_0+1)+j,\theta-\theta_0\right)\Delta V(S_0-j,\theta_0).
\end{aligned}$$

Both terms in the expression above are negative due to (A.9) and (A.6), respectively. This proves (A.16) for $\theta_0 < \theta \le \theta_1$.

Recall that $[V(0,\theta_0)-V(S_0,\theta_0)-K]=-\pi$. Furthermore, from (A.6), we have $\Delta V(i-1,\theta_0)>-\pi$ for $1\le i\le S_0$. Now $V(0,\theta)-V(S_0,\theta)-K$ above is the weighted average of two terms greater than or equal to $-\pi$. All weights are positive. Therefore, $V(0,\theta)-V(S_0,\theta)-K>-\pi$. This proves (A.17) for $\theta_0<\theta\le\theta_1$.

Proof of (A.15)

In the demand process, the number of demand arrivals during a fixed time interval follows a Poisson distribution. In the proof below we need the probability mass function (pmf) of the Poisson distribution to be decreasing. This holds only if the interval length is sufficiently small. Hence, in our proof we need to divide $(\theta_0,\theta_1]$ into several smaller intervals, and apply our proof to each interval. Because the intervals are of fixed length, this will be repeated only a fixed number of times.

To that end, let δ be a small time interval such that $\lambda\delta\le 1$. Poisson-distributed number of arrivals in any time interval of length δ or less will have a decreasing probability mass function (pmf). That is,

$$p(j,\theta_y-\theta_x)>p(j+1,\theta_y-\theta_x)\quad \forall j,\forall\theta_x<\theta_y<\theta_x+\delta\ . \tag{A.18}$$

First we show that (A.15) holds on $(\theta_0,\min(\theta_0+\delta,\theta_1)]$. Then, we show that we can repeat the argument and prove that (A.15) holds on $(\min(\theta_0+\delta,\theta_1),\min(\theta_0+2\delta,\theta_1)]$, $(\min(\theta_0+2\delta,\theta_1),\min(\theta_0+3\delta,\theta_1)]$, etc. In this way, we show that (A.15) holds for all $\theta_0<\theta\le\theta_1$.

Let $\theta_0^{'}=\min(\theta_0+\delta,\theta_1)$. For $i\ge S_0$ and $\theta\in(\theta_0,\theta_0^{'}]$, we can rewrite $V(i,\theta)$ function as:

$$V(i,\theta)$$
$$= K\sum_{n=0}^{\infty} P\left(n(S_0+1)+i+1,\theta-\theta_0\right)+\sum_{j=0}^{S_0}\sum_{n=1}^{\infty} V(j,\theta_0)\,p\left(n(S_0+1)+i-j,\theta-\theta_0\right)+\sum_{j=0}^{i} V(i-j,\theta_0)\,p(j,\theta-\theta_0)$$
$$= K\sum_{n=0}^{\infty}\sum_{j=i+1}^{\infty} p\left(n(S_0+1)+j,\theta-\theta_0\right)+\sum_{n=1}^{\infty}\sum_{j=i-S_0}^{i} p(n(S_0+1)+j,\theta-\theta_0)V(i-j,\theta_0)+\sum_{j=0}^{i} p(j,\theta-\theta_0)V(i-j,\theta_0)$$
$$= K\sum_{n=0}^{\infty}\sum_{j=i+1}^{\infty} p\left(n(S_0+1)+j,\theta-\theta_0\right)+\sum_{n=0}^{\infty}\sum_{j=i-S_0}^{i} p(n(S_0+1)+j,\theta-\theta_0)V(i-j,\theta_0)+\sum_{j=0}^{i-S_0-1} p(j,\theta-\theta_0)V(i-j,\theta_0).$$

The first-order difference is as follows:

$$\begin{aligned}\Delta V(i,\theta) &= V(i+1,\theta)-V(i,\theta)\\ &= \sum_{j=0}^{i-S_0} p(j,\theta-\theta_0)V(i+1-j,\theta_0)-\sum_{j=0}^{i-S_0-1} p(j,\theta-\theta_0)V(i-j,\theta_0)\\ &\quad + K\sum_{n=0}^{\infty}\left(\sum_{j=i+2}^{\infty} p\left(n(S_0+1)+j,\theta-\theta_0\right)-\sum_{j=i+1}^{\infty} p\left(n(S_0+1)+j,\theta-\theta_0\right)\right)\\ &\quad + \sum_{n=0}^{\infty}\left(\sum_{j=i+1-S_0}^{i+1} p(n(S_0+1)+j,\theta-\theta_0)V(i+1-j,\theta_0)-\sum_{j=i-S_0}^{i} p(n(S_0+1)+j,\theta-\theta_0)V(i-j,\theta_0)\right).\end{aligned}$$

Grouping $V(.,.)$ terms with same probability weights, we get:

$$\begin{aligned}\Delta V(i,\theta) = &\sum_{j=0}^{i-S_0-1} p(j,\theta-\theta_0)\left[V(i+1-j,\theta_0)-V(i-j,\theta_0)\right]+p(i-S_0,\theta-\theta_0)V(S_0+1,\theta_0)\\ &-K\sum_{n=0}^{\infty} p\left(n(S_0+1)+i+1,\theta-\theta_0\right)+\sum_{n=0}^{\infty}\left(\sum_{j=i+1-S_0}^{i} p(n(S_0+1)+j,\theta-\theta_0)\left[V(i+1-j,\theta_0)-V(i-j,\theta_0)\right]\right)\\ &+\sum_{n=0}^{\infty} p(n(S_0+1)+i+1,\theta-\theta_0)V(0,\theta_0)-\sum_{n=1}^{\infty} p((n-1)S_0+n+i,\theta-\theta_0)V(S_0,\theta_0)-p(i-S_0,\theta-\theta_0)V(S_0,\theta_0)\end{aligned}$$

Using $\sum_{n=1}^{\infty} p((n-1)S_0+n+i,\theta-\theta_0)=\sum_{n=0}^{\infty} p(n(S_0+1)+i+1,\theta-\theta_0)$ and simplifying, we get:

$$\begin{aligned}\Delta V(i,\theta) = &\sum_{j=0}^{i-S_0-1} p(j,\theta_0,\theta)\Delta V(i-j,\theta_0)+\sum_{n=0}^{\infty}\left(\sum_{j=i+1-S_0}^{i} p(n(S_0+1)+j,\theta-\theta_0)\Delta V(i-j,\theta_0)\right)\\ &+\sum_{n=0}^{\infty} p(n(S_0+1)+i+1,\theta-\theta_0)\left[V(0,\theta_0)-V(S_0,\theta_0)-K\right]+p(i-S_0,\theta-\theta_0)\Delta V(S_0,\theta_0).\end{aligned}$$

A change of index in the summation yields:

$$\Delta V(i,\theta) = \sum_{j=0}^{i-S_0-1} p(j,\theta-\theta_0)\Delta V(i-j,\theta_0) + \sum_{n=0}^{\infty}\left(\sum_{j=1}^{S_0} p((n-1)(S_0+1)+i+j+1,\theta-\theta_0)\Delta V(S_0-j,\theta_0)\right)$$
$$+\sum_{n=0}^{\infty}\left(p(n(S_0+1)+i+1,\theta-\theta_0)\left[V(0,\theta_0)-V(S_0,\theta_0)-K\right]\right)+p(i-S_0,\theta-\theta_0)\Delta V(S_0,\theta_0).$$

This allows us to write the second order difference as:

$$\Delta V(i+1,\theta)-\Delta V(i,\theta)$$
$$=\sum_{j=0}^{i-S_0} p(j,\theta-\theta_0)\left(\Delta V(i+1-j,\theta_0)-\Delta V(i-j,\theta_0)\right) \qquad (A.19)$$
$$+\,p(i+1-S_0,\theta-\theta_0)\Delta V(S_0,\theta_0) \qquad (A.20)$$
$$+\sum_{n=0}^{\infty}\left(\sum_{j=1}^{S_0}\left[p((n-1)(S_0+1)+i+j+2,\theta-\theta_0)-p((n-1)(S_0+1)+i+j+1,\theta-\theta_0)\right]\Delta V(S_0-j,\theta_0)\right) \qquad (A.21)$$
$$+\sum_{n=0}^{\infty}\left(\left[p(n(S_0+1)+i+2,\theta-\theta_0)-p(n(S_0+1)+i+1,\theta-\theta_0)\right]\left(V(0,\theta_0)-V(S_0,\theta_0)-K\right)\right) \qquad (A.22)$$
$$>0,$$

where (A.19) is strictly positive due to (A.8), (A.20) is strictly positive due to (A.7), (A.21) is non-negative due to (A.6) and (A.18), and (A.22) is strictly positive due to (A.9) and (A.18). This proves (A.15) for $\theta_0 < \theta \le \theta_0^{'}$.

Now, recall that the optimal order-up-to level at $\theta_0^{'}$ is S_0. This implies $\Delta V(S_0,\theta_0^{'}) > 0$. This observation, combined with the above proof that (A.15) holds for $\theta_0 < \theta \le \theta_0^{'}$, allows us to write:

$$\Delta V(i,\theta_0^{'}) > 0 \text{ for } i \ge S_0$$

In addition, since $\theta_0^{'} \le \theta_1$:

Based on (A.14), we get $\quad -\pi < \Delta V(i,\theta_0^{'}) < 0 \quad$ for $i < S_0$

Based on (A.15), we get $\quad \Delta V(i,\theta)$ is increasing for $i \ge S_0$

Based on (A.16), we get $\quad V(0,\theta_0^{'}) < V(S_0,\theta_0^{'}) + K$

These four conditions show that (A.6), (A.7),(A.8), and (A.9) also hold true at $\theta_0^{'}$. Now, consider a point $\theta_0^{''} = \min(\theta_0^{'}+\delta,\theta_1)$ such that $\lambda\delta \le 1$. We can repeat the entire proof of (A.15) and with the help of above four conditions show that (A.15) holds on $\theta_0^{'} < \theta \leq \theta_0^{''}$. We can repeat the above argument for $\theta_0^{'''} = \min(\theta_0^{''} + \delta, \theta_1)$ and continue on. Eventually, this proves (A.15) for $\theta_0 < \theta \leq \theta_1$.

Step 3: Optimal Action at θ_1

So far we have shown that (A.14), (A.15), (A.16), and (A.17) hold for all $\theta_0 < \theta \leq \theta_1$. We now prove one more property of $\Delta V(S_0, \theta_1)$ which is $\Delta V(S_0, \theta_1) = 0$.

From the definitions of S_0 and θ_1, S_0 is the largest optimal order quantity on $\theta \in [\theta_0, \theta_1)$. This fact and (A.15) imply $\Delta V(S_0, \theta) > 0$ for all $\theta \in [\theta_0, \theta_1)$.

Since (A.17) holds on θ_1, we know that if inventory is zero and a demand occurs, it is better to order up to S_0 than not ordering. However, by the definition of θ_1, the optimal order-up-to level cannot be S_0 at θ_1. Now, since $\Delta V(S_0, \theta) > 0$ for all $\theta \in [\theta_0, \theta_1)$ and $V(i, \theta)$ (hence, $\Delta V(i, \theta)$) is a continuous function of θ for any fixed *i*, we must also have $\Delta V(S_0, \theta_1) \geq 0$.

Now suppose $\Delta V(S_0, \theta_1) > 0$. Together with (A.15), this implies $\Delta V(i, \theta_1) > 0$ for $i \geq S_0$. Then we can find a small $\epsilon_1 > 0$ such that $\Delta V(i, \theta) > 0$ for all $\theta \in [\theta_1, \theta_1 + \epsilon_1]$ and $i \geq S_0$.

Since we have shown $\Delta V(i, \theta_1) < 0$ for all $i < S_0$, we can also find a small $\epsilon_2 > 0$ such that $\Delta V(i, \theta) < 0$ for $i < S_0$ and $\theta \in [\theta_1, \theta_1 + \epsilon_2]$. Thus, letting $\epsilon = \min\{\epsilon_1, \epsilon_2\}$, we have proved that there is a small neighborhood to the left of θ_1, $[\theta_1, \theta_1 + \epsilon]$, such that for any θ in this interval, $\Delta V(i, \theta) < 0$ for $i < S_0$ and $\Delta V(i, \theta) > 0$ for all $i \geq S_0$. Hence, the optimal order-up-to level is S_0 during the interval $[\theta_1, \theta_1 + \epsilon]$, violating the definition of θ_1. This leaves us with:

$$\Delta V(S_0, \theta_1) = 0. \tag{A.23}$$

(A.14) shows $V(i, \theta_1)$ is decreasing for $i < S_0$, (A.15) and (A.23) show $V(S_0 + 1, \theta_1) = V(S_0, \theta_1)$ and $V(i, \theta_1)$ is increasing for $i > S_0$. Therefore, $S_1 = S_0 + 1$. This concludes the proof of Proposition 4. □

Proof of Theorem 1 Proof of Proposition 4 shows that properties (A.6)-(A.10) of $V(i, \theta_0)$ are needed to prove the form of optimal actions on $\theta_0 < \theta \leq \theta_1$ and properties (A.14)-(A.17) and (A.23). These properties can in turn be translated into:

$-\pi < \Delta V(i, \theta_1) \leq 0$ for $i < S_1$;

$\Delta V(i, \theta_1) > 0$ for $i \geq S_1$;

$\Delta V(i, \theta_1)$ is increasing for $i \geq S_1$;

$V(0, \theta_1) < V(S_1, \theta_1) + K$;

$V(0, \theta_1) + \pi > V(S_1, \theta_1) + K$.

These properties at $\theta = \theta_1$ are the counterparts to properties (A.6)-(A.10) at $\theta = \theta_0$. Therefore, if we repeat the steps in the proof of Proposition 4, we can prove the optimal order-up-to level at $\theta = \theta_2$ and a similar set of properties for $V(i, \theta_2)$. Repeated over and over, these allow us to prove Theorem 1. The fact that $V(i, \theta)$, for any fixed θ, is decreasing in *i* for $i < S(\theta)$ and convex increasing in $i > S(\theta)$ shows it is quasi-convex. □

Proof of Proposition 5 Recall that $TC(i, \theta)$ is the expected total (newsvendor) cost function under which the optimal order-up-to level for H2 is determined, and it can be written as:

$$TC(i, \theta) = \sum_{j=0}^{i} p(j, \theta - \theta_0) V(i - j, \theta_0) + \pi \sum_{j=i+1}^{\infty} p(j, \theta - \theta_0)(j - i + \lambda \theta_0).$$

Therefore,

$$\begin{aligned}\Delta TC(i,\theta) &= TC(i+1,\theta) - TC(i,\theta)\\
&= \sum_{j=0}^{i+1} p(j,\theta-\theta_0)V(i+1-j,\theta_0) + \pi\sum_{j=i+2}^{\infty} p(j,\theta-\theta_0)(j-i-1+\lambda\theta_0)\\
&\quad - \sum_{j=0}^{i} p(j,\theta-\theta_0)V(i-j,\theta_0) - \pi\sum_{j=i+1}^{\infty} p(j,\theta-\theta_0)(j-i+\lambda\theta_0)\\
&= \sum_{j=0}^{i} p(j,\theta-\theta_0)\left[V(i+1-j,\theta_0) - V(i-j,\theta_0)\right] + p(i+1,\theta-\theta_0)V(0,\theta_0)\\
&\quad - \pi P(i+1,\theta-\theta_0) - \pi\lambda\theta_0 p(i+1,\theta-\theta_0)\\
&= \sum_{j=0}^{i} p(j,\theta-\theta_0)\left[V(i+1-j,\theta_0) - V(i-j,\theta_0)\right] - \pi P(i+1,\theta-\theta_0),\end{aligned}$$

as $V(0,\theta_0) = \pi\lambda\theta_0$.

We now prove the result by induction. First, consider $\theta_0 < \theta \le \theta_1$.

$$E[N(i,\theta,\theta_0)] = \sum_{n=0}^{\infty}\Pr[N(i,\theta,\theta_0) > n] = \sum_{n=0}^{\infty}\Pr[D(\theta-\theta_0) > n(S_0+1)+i] = \sum_{n=0}^{\infty}\sum_{j=i+1}^{\infty} p\left(n(S_0+1)+j,\theta-\theta_0\right),$$

so from (3) we get:

$$\begin{aligned}V(i,\theta) &= K\sum_{n=0}^{\infty}\sum_{j=i+1}^{\infty} p\left(n(S_0+1)+j,\theta-\theta_0\right)\\
&\quad+\sum_{n=0}^{\infty}\left\{\sum_{j=0}^{i} p\left(n(S_0+1)+j,\theta-\theta_0\right)V(i-j,\theta_0) + \sum_{j=i+1}^{S_0} p\left(n(S_0+1)+j,\theta-\theta_0\right)V(S_0+1+i-j,\theta_0)\right\},\end{aligned}\tag{A.24}$$

and,

$$\begin{aligned}&\Delta V(i,\theta) = V(i+1,\theta) - V(i,\theta)\\
&= \sum_{n=0}^{\infty} p\left(n(S_0+1)+i+1,\theta-\theta_0\right)\left[V(0,\theta_0) - V(S_0,\theta_0) - K\right] + \sum_{j=0}^{i} p(j,\theta-\theta_0)\left[V(i-j+1,\theta_0) - V(i-j,\theta_0)\right]\\
&\qquad + \sum_{n=1}^{\infty}\sum_{j=0}^{i} p(n(S_0+1)+j,\theta-\theta_0)\left[V(i+1-j,\theta_0) - V(i-j,\theta_0)\right]\\
&\qquad + \sum_{n=0}^{\infty}\sum_{j=i+2}^{S_0} p(n(S_0+1)+j,\theta-\theta_0)\left[V(S_0+2+i-j,\theta_0) - V(S_0+i+1-j,\theta_0)\right]\\
&\ge \sum_{j=0}^{i} p(j,\theta-\theta_0)\left[V(i+1-j,\theta_0) - V(i-j,\theta_0)\right] - \pi P(i+1,\theta-\theta_0)\\
&= \Delta TC(i,\theta),\end{aligned}$$

where we have used the following facts:

- $\left[V(0,\theta_0) - V(S_0,\theta_0) - K\right] = -\pi$ due to (A.10)
- $\Delta V(i,\theta_0) > -\pi$ for $1 \le i < S_0$ due to (A.6), and
- $$\begin{aligned}
&\sum_{n=0}^{\infty} p(n(S_0+1)+i+1,\theta-\theta_0) + \sum_{n=1}^{\infty}\sum_{j=0}^{i} p(n(S_0+1)+j,\theta-\theta_0) + \sum_{n=0}^{\infty}\sum_{j=i+2}^{S_0} p(n(S_0+1)+j,\theta-\theta_0) \\
&= \sum_{n=0}^{\infty} p(n(S_0+1)+i+1,\theta-\theta_0) + \sum_{n=0}^{\infty}\sum_{j=0}^{i} p(n(S_0+1)+j,\theta-\theta_0) \\
&\quad + \sum_{n=0}^{\infty}\sum_{j=i+2}^{S_0} p(n(S_0+1)+j,\theta-\theta_0) - \sum_{j=0}^{i} p(j,\theta-\theta_0) \\
&= \sum_{n=0}^{\infty}\sum_{j=0}^{S_0} p(n(S_0+1)+j,\theta-\theta_0) - \sum_{j=0}^{i} p(j,\theta-\theta_0) \\
&= P(i+1,\theta-\theta_0).
\end{aligned} \tag{A.25}$$

This means that the optimal order-up-to level for $TC(i,\theta)$ is greater than or equal to that for $V(i,\theta)$, proving the result for $\theta_0 < \theta \le \theta_1$.

Next, suppose $\Delta V(i,\theta_m) \ge \Delta TC(i,\theta_m), \forall i$. We then prove that $\Delta V(i,\theta) \ge \Delta TC(i,\theta)$, for all i and for all $\theta_m \le \theta < \theta_{m+1}$. Similar to (A.24), we have:

$$\begin{aligned}
V(i,\theta) &= K\sum_{n=0}^{\infty}\sum_{j=i+1}^{\infty} p\left(n(S_m+1)+j,\theta-\theta_m\right) \\
&+ \sum_{n=0}^{\infty}\left\{\sum_{j=0}^{i} p\left(n(S_m+1)+j,\theta-\theta_m\right)V(i-j,\theta_m) + \sum_{j=i+1}^{S_m} p\left(n(S_m+1)+j,\theta-\theta_m\right)V(S_m+1+i-j,\theta_m)\right\}.
\end{aligned}$$

Now,

$$\Delta V(i,\theta)$$
$$= \sum_{n=0}^{\infty}\left(\begin{aligned}
&K\sum_{j=i+2}^{\infty} p\left(n(S_m+1)+j,\theta-\theta_m\right) - K\sum_{j=i+1}^{\infty} p\left(n(S_m+1)+j,\theta-\theta_m\right) \\
&+ \sum_{j=0}^{i+1} p\left(n(S_m+1)+j,\theta-\theta_m\right)V(i+1-j,\theta_m) - \sum_{j=0}^{i} p\left(n(S_m+1)+j,\theta-\theta_m\right)V(i-j,\theta_m) \\
&+ \sum_{j=i+2}^{S_m} p\left(n(S_m+1)+j,\theta-\theta_m\right)V(S_m+2+i-j,\theta_m) - \sum_{j=i+1}^{S_m} p\left(n(S_m+1)+j,\theta-\theta_m\right)V(S_m+1+i-j,\theta_m)
\end{aligned}\right).$$

Grouping $V(.,.)$ terms with the same probability weights, we get:

$$
\begin{aligned}
&\Delta V(i,\theta) \\
&= \sum_{n=0}^{\infty} \begin{pmatrix} -Kp\left(n(S_m+1)+i+1,\theta-\theta_m\right) \\ +\sum_{j=0}^{i} p\left(n(S_m+1)+j,\theta-\theta_m\right)\left[V(i+1-j,\theta_m)-V(i-j,\theta_m)\right]+p(n(S_m+1)+i+1,\theta-\theta_m)V(0,\theta_m) \\ +\sum_{j=i+2}^{S_0} p(n(S_m+1)+j,\theta-\theta_m)\left[V(S_m+2+i-j,\theta_m)-V(S_m+1+i-j,\theta_m)\right]-p(n(S_m+1)+i+1,\theta-\theta_m)V(S_m,\theta_m) \end{pmatrix} \\
&= \sum_{n=0}^{\infty} \begin{pmatrix} p\left(n(S_m+1)+i+1,\theta-\theta_m\right)\left[V(0,\theta_m)-V(S_m,\theta_m)-K\right] \\ +\sum_{j=0}^{i} p(n(S_m+1)+j,\theta-\theta_m)\left[V(i+1-j,\theta_m)-V(i-j,\theta_m)\right] \\ +\sum_{j=i+2}^{S_m} p(n(S_m+1)+j,\theta-\theta_m)\left[V(S_m+2+i-j,\theta_m)-V(S_m+1+i-j,\theta_m)\right] \end{pmatrix} \\
&= \sum_{n=0}^{\infty} p\left(n(S_m+1)+i+1,\theta-\theta_m\right)\left[V(0,\theta_m)-V(S_m,\theta_m)-K\right] \\
&\quad +\sum_{j=0}^{i} p(j,\theta-\theta_m)\left[V(i+1-j,\theta_m)-V(i-j,\theta_m)\right] \\
&\quad +\sum_{n=1}^{\infty}\sum_{j=0}^{i} p(n(S_m+1)+j,\theta-\theta_m)\left[V(i+1-j,\theta_m)-V(i-j,\theta_m)\right] \\
&\quad +\sum_{n=0}^{\infty}\sum_{j=i+2}^{S_m} p(n(S_m+1)+j,\theta-\theta_m)\left[V(S_m+2+i-j,\theta_m)-V(S_m+1+i-j,\theta_m)\right] \\
&> -\sum_{n=0}^{\infty} p\left(n(S_m+1)+i+1,\theta-\theta_m\right)\pi+\sum_{j=0}^{i} p(j,\theta-\theta_m)\Delta TC(i-j,\theta_m) \\
&\quad -\sum_{n=1}^{\infty}\sum_{j=0}^{i} p(n(S_m+1)+j,\theta-\theta_m)\pi-\sum_{n=0}^{\infty}\sum_{j=i+2}^{S_m} p(n(S_m+1)+j,\theta-\theta_m)\pi \\
&= \sum_{j=0}^{i} p(j,\theta-\theta_m)\Delta TC(i-j,\theta_m)-\pi P(i+1,\theta-\theta_m) \\
&= \Delta TC(i,\theta),
\end{aligned}
$$

where we have used the following facts:

- $V(0,\theta)-V(S,\theta)-K\geq-\pi$ is satisfied for $S=S_m,\theta=\theta_m$, from the proof of Theorem 1,
- $\Delta V(i,\theta)>-\pi$ for $1\leq i<S_0$ is satisfied for $\theta=\theta_m$, from the proof of Theorem 1,
- $\Delta V(i-j,\theta_m)\geq\Delta TC(i-j,\theta_m)$ from induction assumption, and similar to (A.26):

- just like (A.26),

$$\sum_{n=0}^{\infty} p(n(S_m+1)+i+1,\theta-\theta_m)+\sum_{n=1}^{\infty}\sum_{j=0}^{i} p(n(S_m+1)+j,\theta-\theta_m)+\sum_{n=0}^{\infty}\sum_{j=i+2}^{S_m} p(n(S_m+1)+j,\theta-\theta_m)$$
$$= P(i+1,\theta-\theta_m).$$

This means that the optimal order-up-to level for $TC(i,\theta)$ is greater than or equal to that for $V(i,\theta)$ for $\theta_m \le \theta < \theta_{m+1}$, thus completing the induction proof. □

Derivation of the base stock condition for Heuristic H3 Suppose we are at θ. Then we base our ordering decision on the expectation that if we run out at x, $\theta_0 < x < \theta$, an average cost of $K + V(S(x),x)$ will be charged. We approximate $V(S(x),x)$ by the optimal newsvendor cost at time θ_0,

$g(\theta_0) = \pi(\lambda\theta_0 + 1) - K$.

Let $G(S,\theta)$ be the expected total cost for the heuristic that needs to be minimized. Then:

$$G(S,\theta) = [K + g(\theta_0)]P(S+1,\theta-\theta_0) + \sum_{j=0}^{S}[wE_Y(S-j-Y)^+ + \pi E_Y(j+Y-S)^+]p(j,\theta-\theta_0),$$

where, to simplify exposition, we have used Y to represent the random demand in $(0,\theta_0)$. Now:

$$G(S,\theta)-G(S-1,\theta)$$
$$= -(\pi\lambda\theta_0+\pi)p(S,\theta-\theta_0)+\pi\lambda\theta_0 p(S,\theta-\theta_0)+\sum_{j=0}^{S-1}[w-(w+\pi)P(S-j,\theta_0)]p(j,\theta-\theta_0)$$
$$= \sum_{j=0}^{S}[w-(w+\pi)P(S-j,\theta_0)]p(j,\theta-\theta_0).$$

As a heuristic, we use first order condition to choose the order up to level at time θ, $S(\theta)$, to be the largest S such that $\sum_{j=0}^{S}[w-(w+\pi)P(S-j,\theta_0)]p(j,\theta-\theta_0) \le 0$. □

Derivation of the base stock condition for Heuristic H4 Suppose we are at θ. Then, we base our ordering decision on the expectation that if we run out at x, $\theta_0 < x < \theta$, an average cost of $K + V(S(x),x)$ will be charged. We approximate $V(S(x),x)$ by applying the mean value theorem as follows:

$$V(S(x),x) \approx V(S_0,\theta_0) + \frac{V(S(\theta),\theta) - V(S_0,\theta_0)}{\theta - \theta_0}(x - \theta_0)$$
$$= \pi(\lambda\theta_0 + 1) - K + \frac{V(S(\theta),\theta) - V(S_0,\theta_0)}{\theta - \theta_0}(x - \theta_0).$$

Furthermore, we approximate $V(S(\theta),\theta)$ by $g(\theta)$. Define $\beta(\theta) = \dfrac{g(\theta) - g(\theta_0)}{\theta - \theta_0}$, then

$$V(S(x),x) \approx \pi(\lambda\theta_0 + 1) - K + \beta(\theta)(x - \theta_0).$$

Note that by Lemma 2 $\beta(\theta) \le 0$. Again, let $G(S,\theta)$ be the expected total cost for the heuristic that needs to be minimized. Then:

$$G(S,\theta) = \int_{\theta_0}^{\theta} \left(\pi\lambda\theta_0 + \pi - \beta(\theta)(\theta_0 - x)\right)\lambda p(S,\theta - x)dx + \sum_{j=0}^{S}[wE_Y(S - j - Y)^+ + \pi E_Y(j + Y - S)^+]p(j,\theta - \theta_0)$$
$$= \int_0^{\theta-\theta_0} \left(\pi\lambda\theta_0 + \pi - \beta(\theta)(\theta_0 - \theta + x)\right)\lambda p(S,x)dx + \sum_{j=0}^{S}[wE_Y(S - j - Y)^+ + \pi E_Y(j + Y - S)^+]p(j,\theta - \theta_0)$$
$$= \left(\pi\lambda\theta_0 + \pi - \beta(\theta)(\theta - \theta_0)\right)P(S+1,\theta - \theta_0) + \frac{S+1}{\lambda}\beta(\theta)P(S+2,\theta - \theta_0)$$
$$+ \sum_{j=0}^{S}[hE_Y(S - j - Y)^+ + \pi E_Y(j + Y - S)^+]p(j,\theta - \theta_0),$$

where, again, Y represents the random demand during $(0,\theta_0)$. Then:

$$\Delta G(S-1,\theta) = G(S,\theta) - G(S-1,\theta)$$
$$= -\left[\pi\lambda\theta_0 + \pi - \beta(\theta)(\theta - \theta_0)\right]p(S,\theta - \theta_0) - \frac{S+1}{\lambda}\beta(\theta)p(S+1,\theta - \theta_0)$$
$$+ \frac{1}{\lambda}\beta(\theta)P(S+1,\theta - \theta_0) + \pi\lambda\theta_0 p(S-1,\theta - \theta_0) + \sum_{j=0}^{S-1}[w - (w+\pi)P(S - j,\theta_0)]p(j,\theta - \theta_0).$$

Since $(S+1)p(S+1,\theta - \theta_0) = \lambda(\theta - \theta_0)p(S,\theta - \theta_0)$, we get

$$\Delta G(S-1,\theta) = \frac{1}{\lambda}\beta(\theta)P(S+1,\theta - \theta_0) + \sum_{j=0}^{S}[w - (w+\pi)P(S - j,\theta_0)]p(j,\theta_0 - \theta).$$

As a heuristic, we use first order condition to choose the order up to level at time θ, $S(\theta)$, to be the largest S such that $\dfrac{1}{\lambda}\beta(\theta)P(S+1,\theta - \theta_0) + \sum_{j=0}^{S}[w - (w+\pi)P(S - j,\theta_0)]p(j,\theta - \theta_0) \le 0$. □

Table for retailer average order quantities mentioned in Section 4.1

π		λ=50			λ=100			λ=200		
		K=1	K=5	K=25	K=1	K=5	K=25	K=1	K=5	K=25
0.50	Optimal	49.64	47.73	0.00	99.77	97.78	96.00	199.56	197.62	194.00
	H1	47.00	47.00	0.00	96.00	96.00	96.00	194.00	194.00	194.00
	H2	50.91	47.73	0.00	101.84	98.80	96.00	202.64	199.62	194.00
	H3	49.60	48.53	0.00	99.67	98.80	96.00	199.46	199.21	194.00
	H4	49.42	47.01	0.00	99.56	97.36	96.00	199.38	197.42	194.00
1.00	Optimal	50.75	50.57	50.00	100.87	100.95	100.01	200.80	201.08	200.56
	H1	50.00	50.00	50.00	100.00	100.00	100.00	200.00	200.00	200.00
	H2	53.05	51.61	50.00	104.39	103.11	100.01	206.17	205.02	200.56
	H3	50.32	51.63	50.00	100.39	102.21	100.01	200.28	202.31	201.48
	H4	50.59	50.14	50.00	100.70	100.36	100.01	200.70	200.55	200.56
3.00	Optimal	50.97	52.89	54.30	101.25	103.26	106.12	201.19	203.37	207.88
	H1	55.00	55.00	55.00	107.00	107.00	107.00	209.00	209.00	209.00
	H2	56.28	56.24	55.21	108.90	108.73	107.75	211.90	211.74	210.74
	H3	50.48	53.12	55.21	100.71	103.35	108.62	200.63	203.08	211.53
	H4	51.04	51.90	54.30	101.45	102.13	106.12	201.43	202.15	207.21
9.00	Optimal	51.17	53.59	57.75	101.39	103.54	109.54	201.34	204.00	211.30
	H1	59.00	59.00	59.00	113.00	113.00	113.00	218.00	218.00	218.00
	H2	59.54	59.56	59.42	113.77	113.71	113.58	219.06	219.00	218.95
	H3	50.69	52.64	59.44	100.88	102.60	111.89	200.73	202.63	214.15
	H4	51.62	52.02	56.25	101.83	102.05	107.50	201.94	202.37	207.96

This table presents the average order quantity ordered by the retailer across the whole horizon for different combination of parameter values.