Inventory Control in Multi-Channel Retail

Frank Schneider
Supply Chain Management & Management Science
University of Cologne
Cologne, Germany
frank.schneider@uni-koeln.de

Diego Klabjan
Industrial Engineering and Management Sciences
Northwestern University
Evanston, IL, USA
d-klabjan@northwestern.edu

November 19, 2012

Abstract

In recent years multi-channel retail systems have received increasing interest. Partly due to growing online business that serves as a second sales channel for many firms, offering channel specific prices has become a common form of revenue management. We analyze conditions for known inventory control policies to be optimal in presence of two different sales channels. We propose a single item lost sales model with a lead time of zero, periodic review and nonlinear non-stationary cost components without rationing to realistically represent a typical web-based retail scenario. We analyze three variants of the model with different arrival processes: demand not following any particular distribution, Poisson distributed demand and a batch arrival process where demand follows a Pòlya frequency type distribution. We show that without further assumptions on the arrival process, relatively strict conditions must be imposed on the penalty cost in order to achieve optimality of the base stock policy. We also show that for a Poisson arrival process with fixed ordering costs the model with two sales channels can be transformed into the well known model with a single channel where mild conditions yield optimality of an (s, S) policy. Conditions for optimality of the base stock and (s, S) policy for the batch arrival process with and without fixed ordering costs, respectively, are presented together with a proof that the batch arrival process provides valid upper and lower bounds for the optimal value function.

1 Introduction

In recent years multi-channel retail systems have experienced increasing interest. Partly due to the growing online business that serves as a second sales channel for many firms, offering a product for different prices depending on the sales channel has become a common form of revenue management. Many retailers have opened an online channel in addition to brick-and-mortar. Another recent example are orders placed online and merchandise picked up in-store (e.g. Home Depot and Best Buy offer this service to customers). Additionally, many small businesses use platforms like eBay and Amazon Marketplace on the one hand, and a self-managed online store on the other hand, as sales channels.

In all these cases different prices can accompany the same product. There are various additional scenarios in which customer groups pay different prices for the same product. Examples include rebates for certain "valuable" customers, added-value products like cooled beverages that can be price-tagged using RFID technology and other price discrimination settings.

In general, we focus on the possibility of selling items managed under the same stock keeping unit (SKU) for different prices. Inventory control is typically performed at the SKU level and must take this aspect into account. In this work we concentrate on finding optimal ordering policies for the periodic review single item lost sales model with fixed setup costs, two different demand classes, and a lead time of zero in absence of inventory rationing.

This models a typical retail selling scenario with multiple (in our case two) sales channels. A firm orders a single SKU under periodic review from a supplier (single echelon). For each ordering decision a fixed setup and variable procurement cost is incurred. Items - managed under a single SKU - are then sold via two different sales channels for different selling prices to customers. Demand streams are stochastic, independent and follow non-stationary distributions. If demand cannot be satisfied, it is lost and a penalty cost is incurred.

We assume distinct penalty costs for both sales channels. This is consistent with the outlined setting of an online retailer that on the one hand sells through shopping platforms and, on the other hand, runs a self-managed online store. Stock-out penalties are typically higher for self-managed online stores, due to higher sales margins and loss of customer goodwill in case of low availability of products.

We assume prices for the different sales channels are exogeneous and we do not consider the pricing problem. Additionally, we do not allow rationing for any one of the sales channels. The reason for both aspects is the high (price) transparency customers typically have in online businesses. Availability of products in a single channel and stock out in the other is discovered frequently and causes complaints about this practice as well as bad ratings of stores. As a result many online retailers do not consider rationing to be a valid tool of revenue management. The assumption of zero lead time and only two sales channels is a limitation of our work owed to the complexity of lost sales models.

The contributions of this paper are the analyses of lost sales inventory models with two sales channels, nonlinear, non-stationary cost functions and non-stationary demand in a profit maximization scenario with fixed setup cost and no rationing. To the best of our knowledge, this is the first work addressing this setting by means of a rigorous mathematical analysis. We provide conditions for optimality of the base stock policy for a general arrival process and of the (s, S) policy for Poisson arrivals. We also model consecutive batch arrivals of orders in both sales channels and show that costs in this setting serve as valid bounds on the true cost.

The remainder of this paper is organized as follows. In Section 1.1 we provide a short summary of past literature in the field of lost sales models with various demand classes. This is followed by Section 1.2 which introduces background results and notation. Section 2 introduces the basic cost components and a simple arrival model, where customer demand does not follow a particular distribution and we present results of the model analysis in Section 3. In Section 4 we present and analyze the model including a Poisson arrival process of customers. This section also provides a numerical study based on real world data. Section 5 includes the model and analysis of the batch arrival process together with upper and lower bounds on the true value function. We conclude the article in Section 6.

1.1 Literature Review

Lost sales models have received considerably less attention than models with backlogging of demand, since their analysis is in general more involved. Karlin and Scarf [17] were the first to address the lost sales problem with non-zero lead time. They showed optimality of the (s, S) policy for the case of a lead time of one period. Their results are however limited to a single demand class, linear cost functions and stationary demand distributions. Morton [20] extended the work of Karlin and Scarf to models with arbitrary stochastic lead times, and derived various properties of an optimal policy without explicitly stating it. Again, only a single demand class was considered. Nahmias [21] and Johansen [11] introduced myopic heuristics for lost sales models with lead time, due to the absence of a provable optimal policy. Only one demand class with stationary demand is considered. Various papers (see Reiman [23], Hill [9]) consider continuous review models with demand following the Poisson distribution. A comparison between policies that place an order based on a fixed frequency and base stock policies is provided in Reiman [23]. Again, no fixed setup costs and only a single demand class are considered.

Another stream of literature considering lost sales models explicitely requires demand distribution to be Pòlya frequency functions of order 2, as we do in some of our results. Rosling [24] considers an inventory model with nonlinear (fixed plus linear) shortage cost. However, in this work backlogging of demand is assumed and all parameters are stationary. Huggins and Olsen [10] prove optimality of (s,S) policies in a lost sale model in an expediting context with zero lead times for general cost functions, but assume stationary cost and demand parameters for a single demand stream model. Li and Yu [19] show quasiconcavity for a number of lost sales models with a single demand stream, but do only consider lost revenue as penalty cost for lost sales. Additionally, holding costs are assumed to be linear. Also see Porteus [22] Sec. 9.1 for details on inventory control for Pòlya frequency distributions.

Early work addressing the problem of different demand classes was performed by Veinott [12],[13]. He presented a critical level policy for a periodic review model with several demand classes, a lead time of zero and partial backlogging of demand. Topkis [25] showed optimality of the policies by Veinott. Both papers concentrate mainly on finding critical levels for rationing sales to low class demand in order to meet defined service levels. Atkins and Katircioğlu [3] described a single item lost sales model with two demand classes in a periodic review setting, but with the focus on meeting service levels. Cohen, Kleindorfer and Lee [6] performed a comparable analysis in the case of inventory issued to two demand classes with different priority at the end of a period. The period's total demand is known when allocating inventory to the demand classes. Nahmias and Demmy [21] introduced a single period inventory model with rationing of sales for different demand classes. Demand only occurs at the end of a period and is known before the decision whether to satisfy demand or not is made. Ha [8] analyzes a stationary make-to-stock system with several demand classes and lost sales. Here, rationing levels are determined from which on demand from certain classes should be rejected. Downs, Metters and Semple [7] used a linear programming approach to calculate order quantities in a finite horizon non-zero lead time lost sales model for various products. Their approach is based on historic data. Fixed ordering cost (which adds nonlinearity) is not considered. Zipkin [28] presents bounds on the optimal policy for the periodic review single item lost sales model with lead time and also discusses an extension to multiple demand classes, although no optimal policies are mentioned or proven. Arslan, Graves and Roemer [2] study a single-product inventory model for multiple demand classes, but assume backlogging of demand. Zhou et al. [26] discuss a capacitated single item lost sales make-to-stock model with multiple demand classes, but again no fixed set-up costs are considered. Cattani and Souza [5] perform a numerical comparison of rationing policies with a pure first-come, first serve policy as modeled in our paper in the particular case of direct marketing in two sales channels. A general treatment of inventory theory is given in Zipkin [27] where fundamental results on lost sales models are summarized. An extensive review of work on inventory control and fulfillment in multi-channel distribution is given in Agatz et al. [1].

The presented past literature shares two aspects with our models: lost sales and multiple demand classes. Our models are unique due to the following characteristics. We are the first to incorporate general non-stationary non-linear penalty and holding cost functions in a lost sales setting with two demand streams. We concurrently consider fixed plus linear procurement cost as well as holding and penalty costs for both streams. We explicitely disallow inventory rationing for the two demand streams as this reflects many common multi-channel retail scenarios we investigated.

1.2 Background Results

Let $\delta : \mathbb{R} \to \{0,1\}$ be defined as $\delta(u) = 1$ for $u > 0$ and $\delta(u) = 0$ otherwise. Additionally, notation $x^+ = \max(x, 0)$ for $x \in \mathbb{R}$ is used. We use the term *discrete function* to denote a function $f : I \to \mathbb{R}$ where $I \subseteq \mathbb{Z}$. We refer to function $g : I \to \mathbb{R}$ defined by $g(u) = f(u+1) - f(u)$ as the *first difference* of function f. We refer to the first difference of function g as the *second difference* of function f.

Definition 1. Let v be an arbitrary finite sequence of real numbers $v_1, v_2, \dots, v_n$ with $v_i < v_{i+1}$ for all $i = 1, 2, \dots, n-1$. Then the number of sign changes of function $f : I \to \mathbb{R}$ is given by

$$\max_v \left\{ \sum_i |\operatorname{sgn}(f(v_i)) - \operatorname{sgn}(f(v_{i+1}))| \right\}$$

where sgn : $\mathbb{R} \to \{-1, 1\}$ is defined as $\operatorname{sgn}(x) = -1$ for $x < 0$, $\operatorname{sgn}(x) = 0$ for $x = 0$ and $\operatorname{sgn}(x) = 1$ otherwise.

In studying lost sales models with setup costs Pòlya frequency distributions play a vital role. We say that discrete function $f : I \to \mathbb{R}, I \subseteq \mathbb{Z}$ is PF_r if f is a Pòlya frequency sequence of order r in the sense of Karlin [15] Ch. 8.

Many known distributions (e.g. normal, binomial, Poisson, gamma) are Pòlya frequency distributions of infinite order (see e.g. Karlin [15]). For a thorough discussion of Pòlya frequency distributions and proofs of the presented propositions see Karlin [14] or Karlin and Proschan [16].

Definition 2. (adopted from Keilson and Gerber [18]) A discrete function is said to be strongly unimodal, if its convolution with any unimodal discrete function is unimodal.

Proposition 1. Let $f : \mathbb{Z} \to \mathbb{R}$ be a discrete function. If there exists $-\infty < n < \infty, n \in \mathbb{R}$ such that $\operatorname{sgn}(f(u)) = -1$ for all $u < n$ and $\operatorname{sgn}(f(u)) = 1$ for all $u > n$, then sequence g defined by $g(u) = \sum_{k=-\infty}^{u} f(k)$ is strongly unimodal. The first difference of any unimodal sequence has at most one sign change.

The proof of Proposition 1 is straight forward.

Proposition 2. (adopted from Barlow and Proschan [4] and Karlin [15]) Discrete functions that are at least PF_2 are strongly unimodal. The Poisson distribution is a Pòlya frequency function.

Proposition 3. Let $f : \mathbb{Z} \to \mathbb{R}$ be a discrete function with a single sign change as in Proposition 1. Then the convolution of f with a discrete distribution h defined by $g(u) = \sum_{k=-\infty}^{\infty} f(u-k)h(k)$, where $h \geq 0$ and h is at least PF_2, also has a single sign change. Additionally, the sign changes order in the same direction as f.

Proof. The first part directly follows from Definition 2 and Propositions 1 and 2. For the second part assume the sign of f changes order from $-$ to $+$ with increasing u. Then from $f(u) < 0$ for $u < n$, where n is defined as in Proposition 1, we get $\lim_{u \to -\infty} g(u) < 0$ and conversely we have $f(u) > 0$ for $u > n$ and therefore $\lim_{u \to \infty} g(u) > 0$. As a result the sign must change from $-$ to $+$ with increasing u. The same can be shown if the order of the sign change is reversed. $\square$

In some of the proofs we require convexity of the value function or one period cost functions. However, since we are analyzing a setting with discrete quantities, cost functions have domains $I \subseteq \mathbb{Z}$. When we deal with convexity of a function $f : I \to \mathbb{R}$, $I \subseteq \mathbb{Z}$, we refer to convexity of the linear interpolation function $\bar{f} : J \to \mathbb{R}$, with $J \subseteq \mathbb{R}$, defined by $\bar{f}(x) = f(u) + (x-u)(f(u+1) - f(u))$, where $u \in I$ is such that $u \leq x < u + 1$.

2 Modeling

In this section we present our modeling framework. In the general case customers purchase a product through two different sales channels. Customers place orders without a specific order and according to an arbitrary distribution of the order sequence. We call this scenario the general order process. All subsequent models are special cases of this setting. In all models we analyze a finite horizon setting with terminal costs assumed to be zero. Throughout the paper we denote the inventory on hand by $x_t \in \mathbb{N}_0$ and the order up to level by $y_t \in \mathbb{N}_0$. The actual order quantity, i.e., the decision taken in our model, is given by $y_t - x_t$ and always restricted to $y_t \geq x_t$.

The model considers two different sales channels of a single product. We refer to the channels as the high and low class channel (based on the corresponding high or low selling price). The inventory level on hand is reviewed periodically. The sequence of events is as follows. A certain amount of inventory x_t is observed at the beginning of time period t. Then a replenishment order is placed. After the replenishment order arrives with zero lead time nonnegative demand is realized and the cost is accounted for. Note that this is equivalent to placing a replenishment order at the end of a business day and receiving an overnight shipment on the next morning. This setting especially fits a conventional in-store retail business with fixed opening hours, but only approximates the setting of an online store or continuous 24/7 shopping. Typically, customers also place orders during the lead time of a shipment, which of course is non-zero in reality. However, depending on the distribution of the demand over time it can be the case that very low volumes are demanded between the time of order and shipment, which renders the model a good approximation.

2.1 Basic Cost Components

In all models demand exceeding inventory on hand is lost and incurs a penalty cost in addition to lost sales revenue. Inventory exceeding demand in a time period is carried over to the next time period and incurs a holding cost. The cost components considered in all models are

1. the procurement cost,
2. the penalty cost for lost sales,
3. the holding cost, and
4. sales revenue (modeled as negative cost).

As a result, the one period cost function has six components. Since multiple (in our case two) sales channels are considered, we have one penalty cost component for each channel, whereas only a single holding cost function is considered. For all time periods t of the planning horizon let $h_t^N : \mathbb{N}_0 \to \mathbb{R}^+$ be the inventory holding cost and let $h_t^H : \mathbb{N}_0 \to \mathbb{R}^+$, $h_t^L : \mathbb{N}_0 \to \mathbb{R}^+$ be the penalty cost functions for high and low class sales channels, respectively. The high (low) class channel is the sales channel associated with the higher (lower) sales price of the product. Herein we refer to the sum of the negative sales revenue and penalty cost as the *shortage cost*.

A few general standard assumptions with respect to the cost functions are made. Holding and penalty costs are assumed to be finite, nonnegative and are only defined on domain $\mathbb{N}_0$. Additionally, we assume $h_t^N(0) = h_t^L(0) = h_t^H(0) = 0$.

We introduce procurement cost in time period t as $c_t(y_t - x_t) + K_t\delta(y_t - x_t)$. The procurement cost component consists of two parts: a linear part, which includes the variable per item procurement cost c_t, and a fixed cost K_t, which is incurred once every time an order is placed irrespective of its quantity. We also consider sales revenue explicitely and assume it to be linear in the number of items sold through each sales channel. The expression for revenue (defined as a cost) is given by $-p_t^H A_t^H - p_t^L A_t^L$, where A_t^H and A_t^L denote the quantities sold through high and low class channel in period t, respectively. Similarly, p_t^H and p_t^L are the per unit selling prices of the two channels.

2.2 General Order Process Model

In general orders are assumed to be placed in an arbitrary unordered sequence in both channels. Random vector W_t denotes the arrival sequence of orders during time period t. Each coordinate of W_t encodes a binary random variable and the i^{th} coordinate encodes the channel in which order number i was placed. The number of coordinates of W_t is itself random. An example of a realization is given by

$$w_t = \{LHLLHHLHLLLHLHL...L\},$$

which encodes that the first order is placed in the low class sales channel, the second one is placed in the high class sales channel, and so forth (see Figure 1). We denote the i^{th} order of period t by w_{ti} (or W_{ti} when we refer to random variables). Quantity $|w_t|$ denotes the number of coordinates in w_t, which is the number of orders in time period t. The total number of orders for each channel in period t is given by

$$D_t^H(W_t) = |\{i|W_{ti} = H\}|$$

and

$$D_t^L(W_t) = |\{i|W_{ti} = L\}|.$$

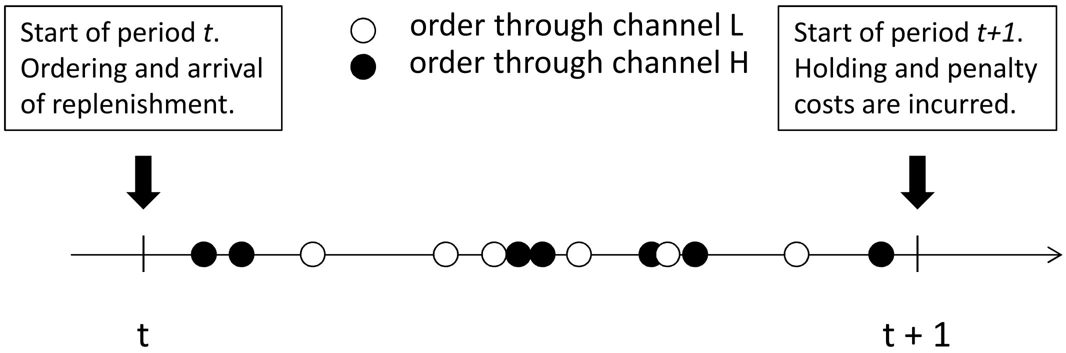

Figure 1: Timeline for one period of the general order process

We also define

$$A_t^H(y_t, W_t) = |\{i|W_{ti} = H, i \leq y_t\}|$$

and

$$A_t^L(y_t, W_t) = |\{i|W_{ti} = L, i \leq y_t\}|.$$

as the amount of the goods sold to high and low class channel, respectively, given inventory level y_t. Note that $A_t^L(y_t, w_t) + A_t^H(y_t, w_t) \leq y_t$ and $D_t^L \geq A_t^L$, $D_t^H \geq A_t^H$ for $y_t \geq 0$ and any realization w_t of W_t.

Using the above definitions we can write the one period cost function $Q_t : \mathbb{N}_0 \to \mathbb{R}$, which is defined by

$$\begin{aligned} Q_t(y_t) = c_t y_t + \sum_{w_t} \Big(& - p_t^L A_t^L(y_t, w_t) - p_t^H A_t^H(y_t, w_t) \\ & + h_t^N (y_t - A_t^L(y_t, w_t) - A_t^H(y_t, w_t)) \\ & + h_t^L (D_t^L(w_t) - A_t^L(y_t, w_t)) \\ & + h_t^H (D_t^H(w_t) - A_t^H(y_t, w_t)) \Big) P(W_t = w_t) \end{aligned}$$

and denotes the controllable part of expected variable cost in period t for given order up to level y_t. Value function $G_t : \mathbb{N}_0 \to \mathbb{R}$ represents expected cost-to-go until the end of the time horizon for current inventory level x_t and is defined by the optimality equation

$$G_t(x_t) = \min_{y_t \geq x_t} \Big\{ -c_t x_t + K_t \delta(y_t - x_t) + Q_t(y_t) + \gamma \sum_{w_t} G_{t+1}(y_t - A_t^H(y_t, w_t) - A_t^L(y_t, w_t)) P(W_t = w_t) \Big\}, \quad (1)$$

where $0 \leq \gamma \leq 1$ is a discount factor.

3 Inventory Control for the General Order Process Model

The presented model is very general and thus - without further assumptions - does not exhibit structured optimal policies (even in the abscence of fixed setup costs). This is mainly due to the fact that selling quantities cannot be written in a closed form expression, which makes analysis of the value function tedious.

We note that even the standard single item lost sales model with a single demand stream requires additional assumptions (e.g. monotone cost functions and the sum of penalty and holding cost must exceed the discounted procurement cost, see Zipkin [27] Sec. 9.4.6) in order to prove optimality of base stock policies in the abscence of fixed costs.

In the next theorem we give conditions for optimality of base stock policies under the abscence of fixed costs, but for the general arrival sequence.

Theorem 1. If in each period t

1. fixed ordering cost $K_t = 0$,
2. penalty cost functions are linear and their slopes b_t^H and b_t^L obey

$$\gamma c_{t+1} - h_t^N(1) \leq b_t^L + p_t^L = b_t^H + p_t^H$$

3. and holding costs are convex,

then the base stock policy is optimal for the general arrival process.

The theorem's conditions can be expressed as follows. Three conditions must be met: (a) penalty costs must be linear in the amount of lost sales, (b) shortage cost per item must be equal for both sales channels, and (c) this quantity must not be below next time period's discounted per unit procurement cost minus holding cost for the first unit. If this is the case, then - in the abscence of fixed ordering costs - the base stock policy is optimal.

Condition 2 of the theorem might not always hold in practice. It requires the sum of the per unit penalty cost and selling price to be equal in both channels, whereas it is expected that lost sales penalties are higher for customers willing to pay a higher price per unit. However, in special cases (e.g. a sales channel for key accounts with rebates) the condition is likely to hold.

The detailed proof is provided in Appendix A. Here, we outline the main ideas at a high level. Optimality of base stock policies is shown by induction. A sufficient condition is convexity of the value function in each time period. We first rewrite (1) in terms of $\bar{G}_t(x_t) = G_t(x_t) + c_t x_t$. It then suffices to show that for any given w_t expression

$$Q_t(y_t) - \gamma c_{t+1}(y_t - |W_t|)^+ \tag{2}$$

is convex. To show this we distinguish case $y_t < |w_t|$, where the penalty and procurement costs are incurred (i.e., demand is higher than inventory), case $y_t > |w_t|$, where the holding and procurement costs are incurred, and case $y_t = |w_t|$, where only the procurement costs are incurred. The procurement costs are linear and can be omitted in the analysis. In the first case all possible arrival sequences of two subsequent orders (HH, HL, LH, LL) are analyzed. Here, the equality in condition 2 of the theorem ensures nondecreasing first differences of (2). The inequality of condition 2 provides this in the remaining two cases.

Note that the conditions of Theorem 1 do not restrict the nature of the underlying arrival sequence, i.e., stochastic process W_t. It follows from the proof that under convexity of G_t our assumptions are sufficient and necessary. However, base stock policies do not require convexity of G_t (but convexity of G_t implies optimality of base stock policies) and therefore in the abscence of this convexity our conditions might be too strict or not sufficient.

Example: To illustrate the conditions of our theorems we present numerical examples based on the following setting. As in our original motivation, we consider an online retailer with two sales channels. According to our model a single product is sold through these channels (e.g. websites) and can be purchased through either channel as long as at least one item of this SKU is in stock.

For illustration of Theorem 1 let us assume that no fixed ordering cost is incurred as required by condition 1, and thus $K_t = 0$ for all t. Daily or at least weekly shipments are common in many online retail segments and we therefore assume a discount factor $\gamma = 1$. Furthermore, let us assume stationary cost parameters for ease of exposition. We assume a unit retail price of $p_t^L = 0.75$ for the first channel and $p_t^H = 1.00$ for the second channel. The equality in condition 2 then requires the unit penalty cost for the first channel to be 0.25 below the unit penalty cost for items sold through the second channel, e.g. $b_t^H = 0.10$ and $b_t^L = 0.35$. Assuming a procurement cost of $c_t = 0.20$, any (nonnegative) convex holding cost function satisfies the inequality in condition 2, rendering base stock policy optimal for the outlined arrival process.

4 Poisson Arrivals

An interesting special case of the general arrival process is the case of Poisson arrivals, which is discussed in this section.

4.1 Model

For the Poisson arrival process order placement in both sales channels is assumed to be Poisson distributed. Let τ be the length of a time period and λ_t be the arrival rate for orders in period t. The total number of orders per period then follows

$$P(|W_t| = k) = e^{-\lambda_t \tau} \frac{(\lambda_t \tau)^k}{k!},$$

where $|W_t|$ is the total number of orders in period t and $\lambda_t \tau$ is mean and variance of the underlying Poisson distribution. For every order there exists certain probability $0 \leq v \leq 1$ that it was placed through the high class channel, and probability $1 - v$ that it comes from the low class channel.

This process is equivalent to the following process. Let high class orders be placed based on rate λ_t^H and low class orders be placed based on rate λ_t^L. If the two processes are independent, then this process is identical to the above process with $\lambda_t = \lambda_t^H + \lambda_t^L$ and $v = \frac{\lambda_t^H}{\lambda_t^H + \lambda_t^L}$.

Expected orders through high and low class channel for a particular realization w_t of W_t are given by $v|w_t|$ and $(1-v)|w_t|$, respectively.

The expected number of goods sold is composed of two terms. The first term captures all cases where total demand exceeds y_t, i.e., $|w_t| \geq y_t$, while the second treats the case where $|w_t| < y_t$. The total number of goods sold through high class channel is given by

$$y_t v \sum_{k=y_t}^{\infty} P(|W_t| = k) + v \sum_{k=0}^{y_t - 1} k P(|W_t| = k). \tag{3}$$

The corresponding expression for low class channel can be obtained by replacing v by $1 - v$ in (3).

Incurred penalty costs in a time period depend on the difference between realized demand and the number of goods sold. Since demand is a random variable specified by the arrival sequence W_t, the amount of unsatisfied demand is also stochastic. Let us denote these quantities for high and low class channel by

$$N_t^H(y_t, W_t) = D_t^H(W_t) - A_t^H(y_t, W_t)$$

and

$$N_t^L(y_t, W_t) = D_t^L(W_t) - A_t^L(y_t, W_t).$$

Note that $N_t^L \geq 0$ and $N_t^H \geq 0$.

The expected penalty costs for the high class channel in period t is given by

$$\sum_{k=0}^{\infty} h_t^H(k) P(N_t^H(y_t, W_t) = k).$$

Conditioning $P(N_t^H(y_t, W_t) = k)$ on the total number of orders yields

$$P(N_t^H(y_t, W_t) = k) = \sum_{i=0}^{\infty} P(N_t^H(y_t, W_t) = k \| w_t| = i) P(|w_t| = i).$$

In the next step term $P(N_t^H(W_t) = k \| w_t| = i)$ is developed. Two cases are distinguished.

1. Let first $y_t \geq i$.
   In this case inventory exceeds i. Therefore, all demand is satisfied and thus

$$P(N_t^H(y_t, W_t) = k \| w_t| = i) = \begin{cases} 1 & \text{if } k = 0, \\ 0 & \text{if } k > 0. \end{cases} \tag{4}$$

2. Let now $y_t < i$.
   In this case inventory is not sufficient to satisfy all demand. The probability of not fulfilling exactly k high class orders is equal to the probability of k orders being placed in the high class channel after depletion of inventory. If $|w_t|$ exceeds inventory y_t by less than k, the probability of not fulfilling k orders is 0. If $|w_t|$ exceeds the inventory by k or more units, the probability follows the Binomial distribution as indicated earlier. To summarize,

$$P(N_t^H(y_t, W_t) = k \| w_t| = i) = \begin{cases} 0 & \text{if } i < k + y_t, \\ \binom{i-y_t}{k} v^k (1-v)^{i-y_t-k} & \text{if } i \geq k + y_t. \end{cases}$$

Probability $P(N_t^H(y_t, W_t) = k)$ can be written as the combination of the two cases discussed above and it reads

$$\begin{aligned} P(N_t^H(y_t, W_t) = k) = & \sum_{i=0}^{y_t-1} P(N_t^H(y_t, W_t) = k \| w_t| = i) P(|w_t| = i), \\ & + \sum_{i=y_t}^{\infty} P(N_t^H(y_t, W_t) = k \| w_t| = i) P(|w_t| = i). \end{aligned}$$

The expected value for the penalty costs is then given by

$$\begin{aligned} E(h_t^H(y_t, W_t)) = \sum_{k=0}^{\infty} h_t^H(k) \Bigg[& \sum_{i=0}^{y_t-1} P(N_t^H(y_t, W_t) = k \| w_t| = i) P(|w_t| = i) \\ & + \sum_{i=k+y_t}^{\infty} P(N_t^H(y_t, W_t) = k \| w_t| = i) P(|w_t| = i) \Bigg]. \end{aligned}$$

The expression for expected penalty costs in the low class channel can be derived accordingly.

Because of (4) and $h_t^H(0, W_t) = 0$, the first summation term in the brackets is always equal to zero. The final expression for the expected value of the penalty costs is then given by

$$E(h_t^H(y_t, W_t)) = \sum_{k=0}^{\infty} h_t^H(k) \sum_{i=k+y_t}^{\infty} \binom{i-y_t}{k} v^k (1-v)^{i-y_t-k} P(|w_t| = i).$$

Equivalently, the expected penalty costs for the low class channel can be written by exchanging v and $1-v$ in the above formulae.

Holding costs are determined by the amount of inventory held at the end of the period after satisfying realized demand. If demand exceeds inventory on hand, clearly, no holding costs are incurred. Expected holding costs are given by

$$E(h_t^N(y_t, W_t)) = \sum_{k=0}^{y_t} P(|w_t| = k) h_t^N(y_t - k).$$

Sales revenue is determined by the total number of items sold through high and low class channels. Rewriting (3) for low class demand and multiplying by the selling price yields expected revenue (defined as a cost)

$$\begin{aligned} E(r_t(y_t, W_t)) = & - p_t^H \left(y_t v \sum_{k=y_t}^{\infty} P(|W_t| = k) + v \sum_{k=0}^{y_t-1} k P(|W_t| = k) \right) \\ & - p_t^L \left(y_t (1-v) \sum_{k=y_t}^{\infty} P(|W_t| = k) + \sum_{k=0}^{y_t-1} k (1-v) P(|W_t| = k) \right) \\ = & - (v p_t^H + (1-v) p_t^L) \left(y_t \sum_{k=y_t}^{\infty} P(|W_t| = k) + \sum_{k=0}^{y_t-1} k P(|W_t| = k) \right). \end{aligned}$$

Having all cost components specified, the optimality equation $G_t : \mathbb{N}_0 \to \mathbb{R}$ for the model can now be written as

$$G_t(x_t) = \min_{y_t \geq x_t} \left\{-c_t x_t + K_t \delta(y_t - x_t) + Q_t(y_t) + \gamma E\left[G_{t+1}\left((y_t - |W_t|)^+\right)\right]\right\}$$

with

$$Q_t(y_t) = E\left(h_t^H(y_t, W_t)\right) + E\left(h_t^L(y_t, W_t)\right) + E\left(h_t^N(y_t, W_t)\right) + E\left(r_t(y_t, W_t)\right) + c_t y_t. \tag{5}$$

4.2 Analysis

For the Poisson model (as well as the model with batch arrivals, which is discussed later in Section 4), two sets of sufficient conditions for optimality of (s_t, S_t)-type policies are established. We present both sets, since their conditions are not subsets of each other. The first set of conditions provides convexity of the one period cost function and preserves K_t-convexity of the value function. It is given in Theorem 2.

Theorem 2. If for all t

1. penalty and holding cost functions are convex and nondecreasing,
2. $K_t \geq \gamma K_{t+1}$, and
3. $p_t^H v + p_t^L(1 - v) \geq -(h_t^H(1)v + h_t^L(1)(1 - v))$,

then the (s_t, S_t) policy is optimal for the Poisson model.

Condition 1 requires cost functions for penalty and holding costs to be convex and nondecreasing, which is a common assumption in inventory control. Condition 2 is fulfilled if discounted setup costs for placing an order do not increase over time. Assuming that the process of placing and processing an order is not subject to change, we expect this condition to be fulfilled in practice. The last condition is always satisfied in case of nonnegative sales revenue. We can conclude that Theorem 2 holds under very mild conditions.

We return to our example from Section 3 to illustrate that the theorem's conditions can commonly be satisfied. We choose e.g. $K_t = 25$ for all t and a discount factor of $0 \leq \gamma \leq 1$, which satisfies condition 2. All other parameters can take arbitrary values, e.g., those given in our example of Section 3, as long as cost functions are convex and take a common value for zero items.

The proof of Theorem 2 is provided in Appendix B. It uses the fact that expected sales revenue as well as penalty costs in both sales channels can be expressed as simple functions, which are then shown to be convex if h_t^L and h_t^H are convex functions. The model then reduces to a single item lost sales model for which optimality of the (s_t, S_t) policy is known for a lead time of zero and convex cost functions (see e.g. Porteus [22], Sec. 7.1).

We provide a second set of optimality conditions for (s_t, S_t) policy relying on unimodality of the value function for each time period, which is another sufficient condition for optimality of (s_t, S_t) type policies. It is presented in Theorem 3.

Theorem 3. Let M be the binomially distributed random variable with parameters $k \in \mathbb{N}$ and $0 \leq v \leq 1$, i.e., $M \sim B(k, v)$. Also define $a_t^L(u) = h_t^L(u+1) - h_t^L(u)$ and $a_t^H(u) = h_t^H(u+1) - h_t^H(u)$. If for all t

1. $p_t^H v + p_t^L(1 - v) > c_t$,
2. $c_{t-1} + \min_{u \geq 0}(h_{t-1}^N(u+1) - h_{t-1}^N(u)) \geq \gamma\left(p_t^H v + p_t^L(1 - v) + \max_k E^M\left[v a_t^H(k - M) + (1 - v)a_t^L(M)\right]\right)$ and
3. penalty costs are nondecreasing,

then the (s_t, S_t) policy is optimal for the Poisson model.

The theorem's conditions can be illustrated as follows. First, the average per item sales price in both channels must exceed the per item procurement cost. In practice this is commonly the case; otherwise the retailer is better off by not selling the product in one or both of the two channels. The second condition of Theorem 3 is more involved. Its interpretation is that the maximum discounted expected per unit shortage cost in the next period must be bounded from above by the minimum per item holding cost plus the procurement cost in the current time period. This can easily be checked in certain cases (e.g., for convex holding costs and concave penalty costs by inspecting incremental costs for the first item). In practice condition 2 is satisfied in a number of scenarios. Examples include high holding costs (e.g. for perishable goods), products with low profit margins and low penalty costs, or in the case of rapidly dropping sales prices. The last condition simply requires penalty costs to be nondecreasing, which is a trivial requirement. Note that shortage costs need not be equal in both channels as was the case for the general order process.

To provide evidence that the conditions are satisfied in certain cases, we state a numerical example. We assume stationary parameters and omit time indices. Consider the example from Section 3 and assume retail prices of $p^L = 0.75$ and $p^H = 1.00$ for the two sales channels, respectively. For $v = 0.3$ and $c = 0.70$ the first condition is clearly satisfied. For a perishable good relatively high holding costs are a reasonable assumption. We assume minimum per item holding costs to be 0.25 and linear penalty costs. The second term on the right hand side of condition 2 then reduces to $\gamma(a^L(1-v) + a^H v)$, where a^L and a^H are the slopes of the penalty cost functions. Penalty costs typically represent loss of goodwill induced by stock outs. We assume values of $a^L = 0.1$ and $a^H = 0.2$ and constant retail prices for our good. condition 2 then requires a discount factor of $\gamma \leq 0.99$ to provide optimality of the (s_t, S_t) policy.

The proof of Theorem 3 is provided in Appendix C. Here we outline the proof at a high level. We prove optimality of the (s_t, S_t) policy by showing unimodality of $L_t = Q_t(y_t) + \gamma E[G_{t+1}((y_t - W_t)^+)]$. This is achieved by proving that precisely one sign change of the first difference of L_t exists and it occurs from - to +. We rewrite the first difference of $L_t(y_t)$ as a convolution of a piecewise defined function $g(k)$ with the Poisson distribution. After showing that $g(k)$ must have precisely one sign change that occurs from - to +, we use Definition 2 and Proposition 2 to show that $L_t(y_t + 1) - L_t(y_t)$ must also have exactly one sign change from - to +. This provides unimodality of L_t and standard arguments show optimality of the (s_t, S_t) policy.

Remark 1. Note that both theorems of this section do not require the underlying distribution of total demand in a period, i.e., random variable $|W_t|$, to be Poisson as their proofs do not rely on unique characteristics of this distribution. For Theorem 2 it is sufficient that for a given total demand the distribution between the two sales channels is binomial, whereas $|W_t|$ can follow an arbitrary distribution. Theorem 3 additionally requires the distribution of $|W_t|$ to be at least PF_2. However, we formulated the theorems in a Poisson based framework, since it is commonly assumed that demands follow a Poisson distribution in our outlined setting.

4.3 Discussion

The two sets of conditions from the previous section cover different cost configurations. While the first set requires convex cost functions and nonincreasing discounted setup costs, the second set of conditions does not impose any restrictions on the setup costs. Instead, the minimum cost incurred by ordering a unit exceeding demand must be higher than the maximum discounted shortage cost in the next time period. Condition 2 of Theorem 3 ensures unimodality of the value function. If it does not hold, there can be cases in which buying an extra unit results in higher expected costs, as the holding costs exceed potential savings in penalty cost, but buying two extra units reduces total expected cost. This results in a policy struture that is more complex than (s_t, S_t). Table 1 provides an overview of cost configurations that provide optimality of the (s_t, S_t) policy.

Cost type	Theorem 2	Theorem 3
procurement (linear)	free	below average revenue
procurement (fixed)	nondecreasing from t to $t+1$	free
penalty	convex, nondecreasing and average for first item $\geq$ average revenue	free[1]
holding	convex, nondecreasing	nondecreasing[1]

Table 1: Optimality conditions imposed by Theorem 2 and 3 for (s_t, S_t) policy

4.4 Numerical Study

To assess the value of following an (s_t, S_t) type policy in a practical setting, we performed a number of computational experiments with cost parameters not satisfying the conditions of Theorems 1 to 3. In order to compare an optimal solution against an optimal (s_t, S_t) type policy, we first solved the dynamic program optimally by backwards recursion. We also obtained values for s_t and S_t in each period by following these steps.

1. Set $S_t = \arg\min_{y_t}\{Q_t(y_t) + \gamma E\left[G_{t+1}\left((y_t - |W_t|)^+\right)\right]\}$.
2. Set $s_t = \max \bar{y}_t$ with

$$\bar{y}_t \in \{y_t | y_t \leq S_t, K_t + Q_t(S_t) + \gamma E\left[G_{t+1}\left((S_t - |W_t|)^+\right)\right] \leq Q_t(y_t) + \gamma E\left[G_{t+1}\left((y_t - |W_t|)^+\right)\right]\}$$

3. If such an s_t does not exist, set $s_t = 0$.

In these steps G_{t+1} is the function in time period $t+1$, when following the (s_t, S_t) policy as described above in each period.

There are two main reasons rendering (s_t, S_t) type policies non-optimal: the demand distribution is not log-concave (as required by Theorem 3) or cost functions are not convex (as required by Theorem 2).

While there is a number of well-known non-log-concave distributions (e.g. Student's t, Pareto, Cauchy, F), we focus on the more practical setting of non-convex cost functions. Commonly observed are non-convex holding or procurement costs, e.g. for renting storage in a warehouse or bulk purchasing, where holding and procurement cost is a staircase function.

We performed a number of experiments with selling prices, variable procurement cost, holding and penalty cost held constant. We allowed expected demand to vary over time and used non-stationary fixed setup cost in some of the instances. Cost parameters from a local retailer for pet supplies who sells through eBay as well as a self-managed online store were used. Sales through the self-managed online store yield higher sales margins. The self-managed online store is therefore considered as the high class channel. Sales through eBay generate higher absolute revenue, but margins are lower due to fees incurred and better comparability of competing offers.

For the experiments we chose a single SKU, in our case a transport box for pets. We set sales prices to $p^H = 6.05$, $p^L = 5.25$ and a high class to low class ratio of $v = 0.25$ as observed from historic sales data for the boxes. Procurement cost for the product is $c = 3$ and penalty costs in both channels were estimated by the sales margin of an average shopping cart including our product and set to $h^L(k) = 3.7k$ and $h^H(k) = 4.5k$. We set the discount factor $\gamma = 0.99995$, representing an annual inflation rate of 2 %, as we assumed period length of one day.

We performed experiments with fixed ordering cost $K = [0, 5, 10, 15]$ and holding cost as a fraction of procurement cost of $h = [0.1, 0.2, 0.3, 0.4]$, as these were parameters for which the retailer was not able to provide specific estimates. We investigated a problem with $T = 30$ periods. For each time period we randomly generated a mean value $\lambda \in [10, 300]$ for the Poisson distribution with equal probability. We generated and solved 40 instances and report the average and maximum cost increase of following an (s_t, S_t) policy over an optimal policy as a percentage of total cost for each parameter combination.

[1]Additionally, condition 2 of Theorem 3 must hold.

h		K 0	5	10	15	20
0.1	avg	1.25	1.04	0.87	0.73	0.62
	max	1.73	1.42	1.14	0.94	0.81
0.2	avg	0.74	0.62	0.53	0.46	0.41
	max	1.15	0.98	0.81	0.68	0.66
0.3	avg	0.56	0.48	0.43	0.39	0.35
	max	0.99	0.82	0.69	0.60	0.59
0.4	avg	0.52	0.44	0.38	0.34	0.30
	max	1.32	1.09	0.88	0.70	0.56

(a) linear holding cost, staircase procurement cost

h		K 0	5	10	15	20
0.1	avg	1.31	1.09	0.92	0.77	0.66
	max	1.73	1.44	1.20	1.01	0.93
0.2	avg	0.89	0.74	0.64	0.57	0.53
	max	1.19	0.97	0.86	0.86	0.85
0.3	avg	0.81	0.70	0.63	0.58	0.54
	max	1.40	1.19	1.03	0.89	0.80
0.4	avg	0.92	0.80	0.70	0.63	0.57
	max	2.41	2.26	2.06	1.83	1.59

(b) staircase holding and procurement cost

Table 2: Cost increase (s_t, S_t) vs. optimal policy for 30 periods in percent

Table 2 shows results of our experiments. We observed suboptimal performance in all instances with staircase procurement cost. Higher fixed setup cost reduces the cost difference in all instances, since the fixed setup cost essentially smoothes out ripples in the value function introduced by the staircase shaped cost functions. Higher holding cost also decreases the cost difference in most instances, mainly due to the higher total cost. For linear procurement cost and stationary fixed ordering cost, (s_t, S_t) was optimal for all instances, including staircase holding cost.

We also performed a number of experiments with non-stationary fixed cost and randomly set K_t to values between 0 and 150. In none of the investigated instances (s_t, S_t) was optimal and we report the maximum and average deviation for different shapes of the cost functions in Table 3 for the holding cost fraction of 0.1. In the case of non-stationary fixed cost, even with all costs linear, (s_t, S_t) is not optimal in all instances. This is caused by the violation of condition 2 of Theorem 3 by our parameter choice in all experiments, which enables varying K_t to cause sub-optimality of (s_t, S_t) as indicated by condition 2 of Theorem 2; otherwise all conditions of Theorem 3 would be met, and K_t had no influence on optimality. In all instances the deviation with non-stationary fixed cost was higher than in the case of stationary fixed cost.

In all our real-world instances we observed that the cost incurred when following an (s_t, S_t) policy is within 3% of the minimum cost. The additional effort to compute the optimal policy is most likely not justifiable in many practical cases. However, note that it is possible to choose parameters that result in an arbitrarily bad performance of the (s_t, S_t) policy.

Procurement Cost	Holding Cost	average	maximum
linear	linear	0.25	0.81
linear	staircase	0.56	1.23
staircase	linear	1.63	1.89
staircase	staircase	1.79	1.97

Table 3: Cost increase of (s_t, S_t) vs. optimal policy for 30 periods and non-stationary fixed cost in percent

5 Demand Channel Batch Arrivals

The main difficulty of the previously described models is the fact that they rely on the actual arrival sequence (for the general order arrival model) or on a particular distribution function (as is the case in the Poisson arrival model). To circumvent this, we study in this section the full batch cases where all orders through one channel are placed at once (a valid interpretation is that demand from one channel has priority and must be satisfied first). This is clearly not completely realistic in our outlined setting and therefore only

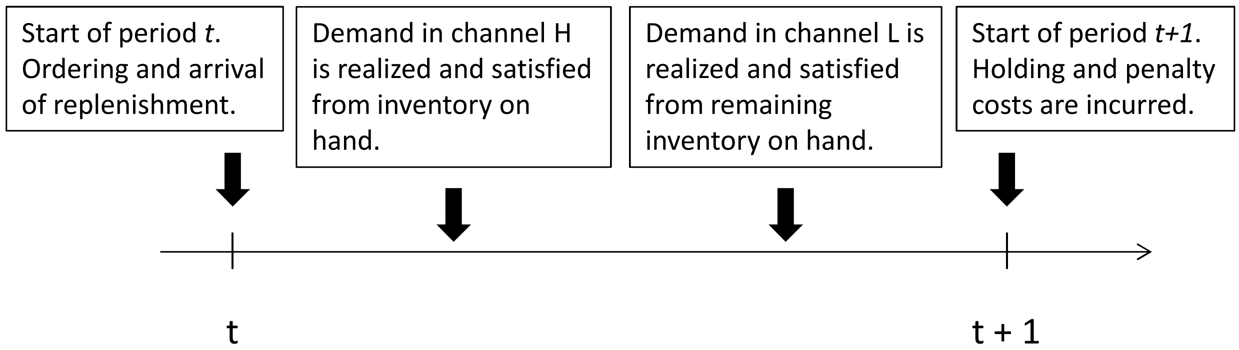

Figure 2: Timeline for one period of the HL variant of the Batch Arrival Process

approximates the previous models. In this section we study policies of this approximate model and provide a relationship to the general order process model. Figure 2 illustrates the batch arrival process.

We consider two cases and assume that either

1. all orders through high class channel arrive and are fulfilled first; left over inventory is then used to fulfill orders through low class channel (this model and the underlying quantities are denoted by superscript HL), or
2. this order is reversed and all orders through the low class channel must be satisfied first, before left over inventory is assigned to satisfy orders from the high class channel (this model and the underlying quantities are denoted by superscript LH).

Total demand in each channel is distributed according to an arbitrary discrete distribution. Distributions of high and low class demand are assumed to be independent of each other. Let $P(D_t^H = h)$ and $P(D_t^L = l)$ be the probability of having exactly h and l orders from the high and low class channel, respectively.

We first state the two models. The one period cost function $Q_t : \mathbb{N}_0^3 \to \mathbb{R}$ for the HL case reads

$$\begin{aligned} Q_t^{HL}(y_t, D_t^H, D_t^L) = &- p_t^H \left(\min(y_t, D_t^H)\right) - p_t^L \left(\min\left(y_t - \min(y_t, D_t^H), D_t^L\right)\right) \\ &+ h_t^N \left((y_t - D_t^H - D_t^L)^+\right) + h_t^L \left(D_t^L - \min\left(y_t - \min(y_t, D_t^H), D_t^L\right)\right) \\ &+ h_t^H \left(D_t^H - \min(y_t, D_t^H)\right) + c_t y_t \end{aligned} \tag{6}$$

and the optimality equation $G_t^{HL} : \mathbb{N}_0 \to \mathbb{R}$ is stated as

$$\begin{aligned} G_t^{HL}(x_t) = \min_{y_t \geq x_t} \Big\{ \sum_{h=0}^{\infty} \sum_{l=0}^{\infty} \Big[&-c_t x_t + K_t \delta(y_t - x_t) + Q_t^{HL}(y_t, h, l) \\ &+ \gamma G_{t+1}^{HL}\left((y_t - h - l)^+\right) \Big] P(D_t^L = l) P(D_t^H = h) \Big\}. \end{aligned} \tag{7}$$

On the other hand, the LH setting is

$$\begin{aligned} Q_t^{LH}(y_t, D_t^L, D_t^H) = &- p_t^L \left(\min(y_t, D_t^L)\right) - p_t^H \left(\min\left(y_t - \min(y_t, D_t^L), D_t^H\right)\right) \\ &+ h_t^N \left((y_t - D_t^H - D_t^L)^+\right) + h_t^H \left(D_t^H - \min\left(y_t - \min(y_t, D_t^L), D_t^H\right)\right) \\ &+ h_t^L \left(D_t^L - \min(y_t, D_t^L)\right) + c_t y_t. \end{aligned}$$

with the optimality equation

$$G_t^{LH}(x_t) = \min_{y_t \geq x_t} \Big\{ \sum_{l=0}^{\infty} \sum_{h=0}^{\infty} \Big[-c_t x_t + K_t \delta(y_t - x_t) + Q_t^{LH}(y_t, l, h) + \gamma \sum_{l=0}^{\infty} \sum_{h=0}^{\infty} G_{t+1}^{LH}\left((y_t - h - l)^+\right) \Big] P(D_t^H = h) P(D_t^L = l) \Big\}.$$

In this analysis we consider model HL, i.e., demand from the high class channel demand is satisfied first. The results can easily be adapted for model LH. We first study optimal policies without the fixed ordering cost.

Theorem 4 (Base stock policy)**.** If for all t

1. $p_t^H + \big(h_t^H(1) - h_t^H(0)\big) \geq p_t^L + \lim_{u\to\infty}\big(h_t^L(u+1) - h_t^L(u)\big)$,
2. $K_t = 0$ and
3. holding and penalty costs are convex and nondecreasing,

then the base stock policy is optimal for model HL.

Theorem 4 has a simple interpretation. The expression on the left-hand side in condition 1 represents the shortage cost for the first lost sale in the high class channel. The expression on the right depicts the maximum cost of a lost sale in the low class channel. Condition 1 simply requires the per item shortage cost in the high class channel to be always equal to or higher than in the low class channel.

The result from Theorem 4 is not surprising. The analyzed scenario can be transformed into a single item, single demand stream model by serially joining the penalty cost and sales revenue functions at quantity D_t^H. Conditions 1 and 3 of the theorem then assures convexity of the shortage cost at this breakpoint.

The detailed proof is provided in Appendix D. It is based on showing convexity of the value function in (7) by induction. First, (6) is inspected for an arbitrary fixed realization of the random demand variables. As mentioned, the one period cost function then reduces to a simple piecewise defined function with two break points. Convexity between breakpoints follows from the general assumption that the holding and penalty costs are convex. Condition 1 of the theorem provides increasing first differences at the breakpoints. As a result, (6) - which is a weighted sum over fixed realizations of the random variables - is convex. Together with the induction assumption that G_{t+1} is convex and monotone, it can be established that (7) is convex and the base stock policy is optimal.

We again give a numerical example of these conditions. Condition 1 requires sales revenue plus the penalty cost for the first unit for the high class channel to be greater than or equal to the sales revenue plus the maximum per unit penalty cost for the low class channel. For ease of exposition we assume the penalty cost functions to be linear and choose $p_t^H = 1.00$, $h_t^H(k) = 0.2k$, $p_t^L = 0.75$ and $h_t^L(k) = 0.1k$ as in the example of Section 4.2. Additionally, we assume no setup costs are incurred. This satisfies all of the conditions of the theorem and renders the base stock policy optimal.

In presence of fixed ordering costs, we can state two sets of conditions, under which (s_t, S_t)-type policies are optimal. The sets - as in Section 4.2 - rely on different properties of the value function and are not just subsets of each other. The first is a generalization of Theorem 4 that allows for $K_t > 0$. It is based on convexity of the one period cost function and K_t-convexity of the value function.

Proposition 4. If for all t

1. holding and penalty cost functions are convex and nondecreasing,
2. $p_t^H + h_t^H(1) - h_t^H(0) \geq p_t^L + \lim_{u\to\infty}(h_t^L(u+1) - h_t^L(u))$ and

3. $K_t \geq \gamma K_{t+1}$,

then the optimal policy for model HL is of (s_t, S_t) type.

Proof. We know from Theorem 4 and its corresponding proof that given condition 1 of Theorem 4, G_t^{HL} and Q_t^{HL} are convex in all periods in the abscence of fixed costs, i.e., $K_t = 0$ for all t. For $K_t > 0$, let

$$f_t(y_t) = E\big[Q_t^{HL}(y_t, h, l) + \gamma G_{t+1}^{HL}\left((y_t - D_t^H - D_t^L)^+\right)\big].$$

The optimality equation now reads

$$G_t^{HL}(x_t) = -c_t x_t + \min\left\{f_t(x_t), \min_{y_t \geq x_t}\left[K_t + f_t\left(y_t\right)\right]\right\}.$$

Following the proof of Theorem 2, we note that assuming G_{t+1}^{HL} is K_{t+1}-convex, f_t must be γK_{t+1}-convex. Assuming $G_{T+1}^{HL} = 0$, it is then easy to show by backwards induction that an (s_t, S_t)-type policy is optimal for the batch arrival model with fixed ordering costs, if $K_t \geq \gamma K_{t+1}$. □

The second set of conditions provides unimodality of the value function and is in the same spirit as Theorem 3.

Theorem 5. If for all t

1. $c_{t-1} + \min_u(h_{t-1}^N(u+1) - h_{t-1}^N(u)) > \gamma \max\left(p_t^H + \max_u(h_t^H(u+1) - h_t^H(u)), p_t^L + \max_u(h_t^L(u+1) - h_t^L(u))\right)$,
2. $c_t < p_t^L + h_t^L(u+1) - h_t^L(u)$ and $c_t < p_t^H + h_t^H(u+1) - h_t^H(u)$ for all u and
3. D_t^H and D_t^H are random variables with a distribution function that is at least PF_2,

then the optimal policy for model HL is of (s_t, S_t) type.

The crucial requirement of Theorem 5 lies in condition 1. It requires the procurement cost plus the minimum per item holding cost in the current period to be higher than the next period's discounted maximum per item shortage cost (sum of lost revenue and penalty cost). This can be the case for rapidly dropping sales prices or high holding costs (e.g. for perishable goods). Low discount factors have a similar interpretation and can also lead to this condition to hold.

Additionally, we need the per item shortage cost to be greater than the per item procurement cost. Setting retail prices above the procurement cost and assuming nonnegative penalty costs suffices to fulfill this condition. The last condition is a technical condition requiring the distribution functions of the demand streams to be at least PF_2, which is the case for many common distributions (e.g. Poisson, binomial, gamma, normal, etc.). Note that Theorem 5 does not formulate any conditions including K_t, which is typical for theorems based on unimodality of the value function.

The proof for this theorem is provided in Appendix E. Again, we outline it at a high level. The proof is carried out by induction. We inspect the expression

$$R_t(y_t) = \sum_{h=0}^{\infty}\sum_{l=0}^{\infty}\left[Q_t^{HL}(y_t, h, l) + \gamma G_{t+1}^{HL}\left((y_t - D_t^H - D_t^L)^+\right)\right] P(D_t^l = l) P(D_t^H = h).$$

It can be shown that conditions 2 and 3 of Theorem 5 assert that the first difference of $R_t(y_t)$ is the convolution of a function with a single sign change from - to + with a strongly unimodal discrete function. According to Proposition 3 the first difference of $R_t(y_t)$ must also have a single sign change from - to +. It then follows that $R_t(y_t)$ is unimodal and standard arguments show that the (s_t, S_t) policy is optimal.

For a numerical example we use the parameters from the corresponding example of Theorem 3. With $p^L = 0.75$, $p^H = 1.00$, $a^L = 0.1$ and $a^H = 0.2$, condition 3 of the theorem is satisfied. Condition 2 requires $0.7 + \min_{u \geq 0}(h_N(u+1) - h^N(u)) \geq 1.2\gamma$. Setting $\gamma = 0.99$ as before requires minimum per item holding costs of 0.488, whereas in the Poisson arrival case a value of 0.25 was sufficient.

As for the Poisson arrival process the presented sets of optimality conditions rely on different fundamental properties of the cost functions. The first set requires nonincreasing setup costs and convex holding and shortage costs. It uses the fact that shortage costs are simple to define in a closed form expression in the case of batch arrival, as there is a single breakpoint at which the function changes its form.

The second set does not constrain values for setup costs, but needs shortage costs to be nondecreasing and limited by the minimum holding plus procurement cost in the previous time period. The reason for the last part is that otherwise there could be inventory levels at which it would be cost optimal to order up to an inventory level greater than S_t.

5.1 Upper and Lower Bounds

Intuitively it is clear that the case in which high class orders are placed and satisfied first represents a lower bound on the real cost in a time period, if we have higher penalty costs and sales margins in the high class channel. In the same manner, case LH where all low class orders are placed and satisfied first intuitively represents an upper bound. Theorem 6 provides a formal proof of necessary conditions on selling prices and penalty cost functions for the batch model to be valid lower and upper bounds on the true cost function.

Theorem 6. If for all t

1. $h_t^L(u_1+1) - h_t^L(u_1) \leq h_t^H(u_2+1) - h_t^H(u_2)$ for all $u_1, u_2 \in \mathbb{N}_0$ and
2. $p_t^L \leq p_t^H$

hold, then G_t^{HL} and G_t^{LH} are lower and upper bounds for G_t, i.e.,

$$G_t^{HL}(x_t) \leq G_t(x_t) \leq G_t^{LH}(x_t) \tag{8}$$

holds.

The theorem's conditions are in line with the intuitive idea that a firm could maximize its profit, if it was able to always satisfy demand associated with the higher shortage cost first. However, it is surprising that we cannot simply require shortage cost to be higher in one channel than in the other. We explicitely need the condition that sales prices in the high channel must be equal to or exceed sales prices in the low channel. Otherwise, in the case of sufficient inventory on hand, we could achieve lower total cost by selling to the low class channel first.

The proof is provided in Appendix F. The proof is based on the fact that given $Q_t^{HL} \leq Q_t \leq Q_t^{LH}$ for all t, $G_t^{HL} \leq G_t \leq G_t^{LH}$ does hold. This can easily be seen by induction. Given an arbitrary demand realization in a time period, we can omit procurement cost in our analysis, since it is simply a linear term in $Q_t(y_t)$ and identical for G_t^{HL}, G_t^{LH} and G_t. We then fix y_t and inspect possible realizations of demand for a period. Two cases are inspected. First we assume all demand can be satisfied. Then no penalty costs are incurred and condition 2 of the theorem asserts that $Q_t^{HL} \leq Q_t \leq Q_t^{LH}$. In the case when demand exceeds supply, condition 1 of the theorem asserts $Q_t^{HL} \leq Q_t^{LH}$ for any arbitrary realization of demand. It can then be argued that Q_t must always lie within the given bounds, since total demand is always equal for G_t^{HL}, G_t^{LH} and G_t and the minimum and maximum penalty costs that can be incurred given a fixed total demand are given by the assumed batch arrivals. This completes the proof.

5.2 Numerical Study

To assess the quality of the upper and lower bounds, we performed a number of numerical experiments based on the real-world setting from Section 4.4, which satisfies the conditions of Theorem 6. We performed experiments with fixed ordering cost $K = [0, 5, 10, 15]$ and holding cost as a fraction of procurement cost of $h = [0.1, 0.2, 0.3, 0.4]$. We investigated a problem with $T = 30$ periods and $\lambda = 30$. In our first study we used a high class to low class ratio of $v = 0.25$ and all costs were assumed to be linear. The difference in shortage cost $\Delta = p^h + h^H - p^L - h^L = 1.8$ was chosen to match our real-world example. In the second study we performed a number of experiments for $h = 0.1$ and $K = 10$ and varied v as well as Δ. Table 4 reports the maximum relative gap in expected cost between upper and lower bound, i.e., $\max_{x_0 \geq 0} \frac{G_0^{LH}(x_0) - G_0^{HL}(x_0)}{G_0^{LH}(x_0)}$.

	K				
h	0	5	10	15	20
0.1	0.18	0.19	0.21	0.24	0.26
0.2	0.22	0.22	0.23	0.23	0.23
0.3	0.04	0.03	0.03	0.02	0.01
0.4	0.01	0.00	0.00	0.00	0.00

(a) varying K and h for $v = 0.25$, $\Delta = 1.8$

	v				
Δ	0.1	0.3	0.5	0.7	0.9
1	0.24	0.29	0.27	0.25	0.17
2	0.41	0.43	0.38	0.34	0.21
3	0.54	0.55	0.46	0.39	0.24
4	0.66	0.60	0.49	0.41	0.24

(b) varying v and Δ for $K = 10$, $h = 0.1$

Table 4: Maximum gap in expected cost between upper and lower bound for $T = 30$ in percent

Table 4 shows that the gap is below 1% for all computed instances. Table 4 (a) indicates that for increasing holding cost the relative gap decreases, as a result of higher total cost incurred. For holding cost rates of 0.1 and 0.2 the gap increases in K, but for holding cost rates of 0.3 and 0.4 the gap decreases in K. This is caused by the fact that for high holding cost the product is ordered in every period in both models, whereas for low holding cost the ordering decisions are different in both models. If high class demand is satisfies first, the policy is to always order. If low class demand is satisfied first, the product is not ordered in every period for high values of K.

Table 4 (b) shows that increasing the difference in sales prices and penalty cost increases the relative gap for given v. This is intuitively clear, because the difference in revenue and cost in both models must be higher if the difference in cost of a lost sale increases. Increasing the fraction of high class demand of total demand induces two competing effects. On the one hand, intuitively the gap takes its maximum at $v = 0.5$ if the same ordering decisions are taken in both models. On the other hand, total profit increases when the fraction of high class demand is increased such that for high Δ the relative gap decreases monotonely for increasing v.

6 Concluding Remarks

Although the model with only two sales channels analyzed in this work is very basic, the conditions listed for optimality of basic inventory control policies are relatively restrictive for general order processes. This is a property inherent to lost sales models, which are complex due to their special nonlinear structure in system dynamics. However, we succeeded to analyze models that closely parallel true circumstances in a real world retail scenario with two sales channels. As expected, there is a trade-off between restrictiveness of assumptions regarding the arrival process on the one hand and restrictiveness of optimality conditions on the other hand. While the absence of assumptions on the arrival process leads to requesting the penalty costs to be linear and to have the difference in slopes equal to the difference in selling prices, modeling order arrivals by a Poisson process mitigates this requirements and allows general convex penalty cost functions. We also have been able to provide two independent sets of optimality conditions for the (s_t, S_t)-policy under Poisson arrival and batch arrival processes; one based on providing convexity of the one period cost function and the other one based on unimodality of the value function.

In short, we have provided useful theory and insights into a practical retail setting with multiple sales channels or, equivalently, price discrimination. We hope that our results can be of value to practitioners in inventory control managing multiple sales channels. Our results shed light on structure of optimal inventory control policies and the corresponding restrictions on model parameters. To the best of our knowledge, there has been no work including nonlinear non-stationary cost components, two sales channels and lost sales in a revenue maximization setting without inventory rationing.

References

[1] N.A. Agatz, M. Fleischmann, and J.A. van Nunen. E-fulfillment and multi-channel distribution a review. *European Journal of Operational Research*, 187(2):339 – 356, 2008.

[2] H. Arslan, S.C. Graves, and T.A. Roemer. A single-product inventory model for multiple demand classes. *Management Science*, 53(9):1486–1500, 2007.

[3] D. Atkins and K.K. Katircioğlu. Managing inventory for multiple customers requiring different levels of service. Technical Report Working Paper 94-MSC-015, University of British Columbia, Vancouver, B.C., 1995.

[4] R.E. Barlow and F. Proschan. *Mathematical Theory of Reliability.* John Wiley & Sons, New York, 1965.

[5] K.D. Cattani and G.C. Souza. Inventory rationing and shipment flexibility alternatives for direct marketing firms. *Production and Operations Management*, 11(4):441–457, 2002.

[6] M.A. Cohen, P.R. Kleindorfer, and H.L. Lee. Service constrained (s, S) inventory systems with priority demand classes and lost sales. *Management Science*, 34:482–499, 1983.

[7] B. Downs, R. Metters, and J. Semple. Managing inventory with multiple products, lags in delivery, resource constraints, and lost sales: A mathematical programming approach. *Management Science*, 47(3):464–479, 2001.

[8] A.Y. Ha. Inventory rationing in a make-to-stock production system with several demand classes and lost sales. *Management Science*, 43(8):1093–1104, 1997.

[9] R.M. Hill. Continuous-review, lost-sales inventory models with Poisson demand, a fixed lead time and no fixed order cost. *European Journal of Operational Research*, 176:956–963, 2007.

[10] E.L. Huggins and T.L. Olsen. Inventory control with generalized expediting. *Operations Research*, 58(5):1414 – 1426, 2010.

[11] S. G. Johansen. Pure and modified base-stock policies for the lost sales inventory system with negligible set-up costs and constant lead times. *International Journal of Production Economics*, 71:391–399, 2001.

[12] A.F. Veinott Jr. Optimal policy in a dynamic, single product, non-stationary inventory model with several demand classes. *Operations Research*, 13:761–778, 1965.

[13] A.F. Veinott Jr. The status of mathematical inventory theory. *Management Science*, 12:745–777, 1966.

[14] S. Karlin. Polya type distributions II. *Annals of Mathematical Statistics*, 28(2):281–308, 1957.

[15] S. Karlin. *Total Positivity. Volume 1.* Stanford University Press, 1968.

[16] S. Karlin and F. Proschan. Polya type distributions of convolutions. *Annals of Mathematical Statistics*, 31(3):721–736, 1960.

[17] S. Karlin and H. Scarf. Inventory models of the Arrow-Harris-Marschak type with time lag. *Studies in the mathematical theory of inventory and production*, pages 155–178, 1958.

[18] J. Keilson and H. Gerber. Some results for discrete unimodality. *Journal of the American Statistical Association*, 66(334):386–389, 1971.

[19] Q. Li and P. Yu. On the quasiconcavity of lost-sales inventory models with fixed costs. *Operations Research*, 60(2):286 – 291, 2012.

[20] T.E. Morton. Bounds on the solution of the lagged optimal inventory equation with no demand backlogging and proportional costs. *SIAM Review*, 11:527–596, 1969.

[21] S. Nahmias and S. Demmy. Operating characteristics of an inventory system with rationing. *Management Science*, 27:1236–1245, 1981.

[22] E.L. Porteus. *Foundations of Stochastic Inventory Theory*. Stanford University Press, 2002.

[23] M.I. Reiman. A new simple policy for a continuous review lost-sales inventory model. Unpublished, 2004.

[24] K. Rosling. Inventory cost rate functions with nonlinear shortage costs. *Operations Research*, 50(6):1007 – 1017, 2002.

[25] D.M. Topkis. Optimal ordering and rationing policies in a nonstationary dynamic inventory model with n demand classes. *Management Science*, 15:160–176, 1968.

[26] W. Zhou, C.-Y. Lee, and D. Wu. Optimal Control of a Capacitated Inventory System with Multiple Demand Classes. *Naval Research Logistics*, 58:43–58, 2011.

[27] P.H. Zipkin. *Foundations of Inventory Management*. McGraw-Hill, 2000.

[28] P.H. Zipkin. On the structure of lost-sales inventory models. *Operations Research*, 56(4):937–944, 2008.