

# On Linear Threshold Policies for Continuous-Time Dynamic Yield Management

Dipayan Banerjee      Alan Erera      Alejandro Toriello

H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology

**Abstract.** We study the finite-horizon continuous-time yield management problem with stationary arrival rates and two customer classes. We consider a class of linear threshold policies proposed by Hodge (2008), in which each *online* (i.e., less desirable) customer is accepted if and only if the remaining inventory at the customer’s arrival time exceeds a threshold that linearly decreases over the selling horizon. Using a discrete-time Markov chain representation of sample paths over inventory-time space, we show that a range of such linear threshold policies achieve uniformly bounded regret. We then generalize this result to analogous policies for the same problem with arbitrarily many customer classes. Numerical simulations demonstrate linear threshold policies’ competitiveness with existing heuristics and illustrate the effects of the linear threshold’s slope.

## 1. Background

A multichannel retailer sells a product through two different channels, denoted as *offline* and *online* channels for consistency, over a finite horizon of duration  $T$ . Customers arrive over the horizon and attempt to purchase one unit of the product. Offline customers arrive via a Poisson process with rate  $\lambda_1 > 0$  and pay  $p_1 > 0$  for each unit. Online customers arrive via a Poisson process with rate  $\lambda_2 > 0$  and pay the same price. The retailer incurs an additional fixed cost to fulfill each online order, and online customers receive free delivery; this reflects the situation faced by an industry partner. Equivalently, we instead assume without loss of generality that online customers pay  $p_2 \in (0, p_1)$  for each unit, and the retailer does not incur the fixed per-customer cost. Any inventory left over at the end of the selling horizon is discarded and provides no additional salvage value. For notational consistency with related work,  $t$  denotes the time left until the end of the horizon; thus,  $t = 0$  denotes the terminal time. The retailer has  $n$  units of starting inventory available at  $t = T$  with no replenishment opportunities during the selling horizon.

The objective is to maximize the revenue collected over the horizon. To this end, the retailer must immediately decide whether to accept an arriving customer’s order. It is clear that every offline customer should be accepted (thereby receiving  $p_1$  in revenue) while positive inventory remains. Backorders are not permitted, so all customers are rejected upon inventory depletion. Thus, the retailer faces the problem of determining whether to accept or reject each arriving online customer (thereby receiving  $p_2$  or 0 in revenue, respectively). Mathematically, this is the *continuous-time dynamic yield management problem* with two customer classes, stationary arrival rates, and a finite horizon.

As shown independently by Liang (1999), Zhao (1999), and Feng and Xiao (2001), a particular threshold policy maximizes expected revenue. Specifically, an online order is accepted at time  $t$  if (and only if) the inventory remaining exceeds  $\theta(t)$ , where the optimal threshold function  $\theta(\cdot)$  is non-decreasing in  $t$  and

independent of the starting inventory  $n$ . All offline orders are accepted while time remains. A continuous-time dynamic programming method is generally employed to compute the optimal threshold function  $\theta(\cdot)$ ; see Liang (1999) and Alishah et al. (2017) for details. Figure 1 illustrates numerically computed optimal threshold functions for two settings with  $\lambda_1 = 8$ ,  $\lambda_2 = 2$ ,  $p_1 = 25$ , and  $T = 5$ .

Exact computation of  $\theta(\cdot)$  requires recursion, repeated numerical integration, and floating-point calculations involving near-zero values. As a result, compounding numerical issues prevent reliable exact computation of thresholds as  $T$  grows. We are therefore interested in characterizing the behavior of  $\theta(t)$  as  $t \rightarrow \infty$ . Hodge (2008) conjectures that  $\theta(t)$  behaves linearly (that is, the function’s step widths converge to a finite positive constant) as  $t \rightarrow \infty$ . Despite empirical evidence in support of the conjecture, such as Figure 1 of this work and Figure 4 of Alishah et al. (2017), it remains unproven to our knowledge.

In this work, we instead show that a range of simple linear inventory threshold policies are asymptotically optimal. Formally, for each  $\beta \in (\lambda_1, \lambda_1 + \lambda_2)$ , the linear threshold acceptance policy with slope  $\beta$  (denoted  $\beta$ -LT) accepts an incoming online customer at time  $t$  if and only if the remaining inventory at that time is at least  $\beta t$ ; all offline orders are accepted while time remains. To show that  $\beta$ -LT policies are asymptotically optimal, we prove an even stronger result: for fixed  $\lambda_1$  and  $\lambda_2$ , any appropriately chosen  $\beta$ -LT policy achieves regret – the expected absolute gap between a policy’s revenue and that of the hindsight-optimal solution – that is bounded above by a uniform constant that depends on  $\beta$ ,  $\lambda_1$ , and  $\lambda_2$  but not on  $n$  and  $T$ . We then show that an extension of these policies also achieves uniformly bounded regret when more than two customer classes are present. Our policies and results are related to those of Arlotto and Gurvich (2019) and Vera and Banerjee (2021). In particular,  $\beta$ -LT policies are continuous-time generalizations of the policy proposed by Arlotto and Gurvich (2019) for the multi-secretary problem. Although the approach used by Arlotto and Gurvich (2019) could likely be modified to prove the results herein, our analysis (which directly uses Markov chain modeling and standard results from queueing theory) differs significantly from these prior works.

## 1.1 Related Work

Our problem is a special case of the broader class of continuous-time, quantity-based network revenue management problems, commonly referred to simply as network revenue management (NRM). Specifically, our main problem can be viewed as a NRM problem with one resource (the product to be sold), unit demands, and two customer classes having stationary Poisson arrival rates. Heuristics for NRM typically rely on the solution to a deterministic linear programming (DLP) approximation of the stochastic problem. The DLP associated with our specific two-class problem is

$$\max_{z_1, z_2 \in [0,1]} \lambda_1 T p_1 z_1 + \lambda_2 T p_2 z_2 \quad \text{s.t. } \lambda_1 T z_1 + \lambda_2 T z_2 \leq n, \quad (1)$$

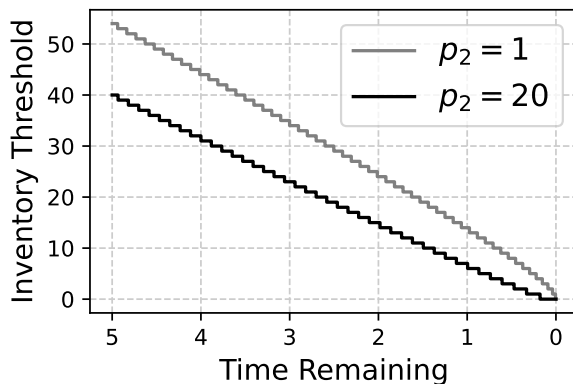


Figure 1: Example optimal threshold functions

where the variables  $z_1$  and  $z_2$  denote the proportion of offline and online orders accepted, respectively. One simple DLP-based heuristic accepts each incoming online customer with probability  $z_1$  and each incoming offline customer with probability  $z_2$  while inventory remains. This heuristic and many DLP-based heuristics for general NRM are known to achieve  $\Theta(\sqrt{T})$  regret when  $n$  and  $T$  are scaled up proportionally.

Jasin and Kumar (2012) develop an adaptive NRM heuristic policy, *Probabilistic Allocation Control* (PAC), that periodically re-solves the DLP at equally spaced points in time. Although the authors show the PAC policy to have  $\mathcal{O}(1)$  regret if the DLP is non-degenerate, Bumpensanti and Wang (2020) show that the PAC policy’s regret is  $\Theta(\sqrt{T})$  for certain degenerate DLPs. Bumpensanti and Wang (2020) develop an improved NRM heuristic policy that achieves  $\mathcal{O}(1)$  regret in the limit. Their approach, *Infrequent Re-Solving with Thresholding* (IRT), requires  $\mathcal{O}(\log \log T)$  re-optimizations of the underlying DLP during a horizon of length  $T$ . Unlike PAC, IRT’s re-optimizations do not occur at evenly spaced points in time. For the special case of a single resource (i.e., product) considered in this work, we show that linear threshold policies also achieve  $\mathcal{O}(1)$  regret without solving the DLP even once.

Arlotto and Gurvich (2019) construct a specific constant-regret *Budget Ratio* (BR) policy for a discrete-time multi-secretary problem. The continuous-time problem studied here can be viewed as a limiting case of their multi-secretary problem upon appropriate scaling. Vera and Banerjee (2021) construct a specific constant-regret *Fluid Bayes Selector* (FBS) policy, which uses dynamic thresholds, for a class of dynamic optimization problems of which NRM is a special case. In contrast to these works, our analysis explicitly assumes continuous-time arrivals, and our approach relies on characteristics of a particular discrete-time Markov chain (DTMC) model. Our approach shows that a range of  $\beta$ -LT policies each achieve  $\mathcal{O}(1)$  regret, suggesting that a range of policies similar to the authors’ specifically constructed BR and FBS policies might also achieve constant regret. Additionally, as we discuss later, the asymptotic optimality of our  $\beta$ -LT policies does not always require exact knowledge of the customer arrival rates. This is a continuous-time analog of the multi-secretary extension in which “the ability distribution is... unknown to the decision maker” noted by Arlotto and Gurvich (2019) as a potential direction for future work.

## 1.2 Definitions and Notation

Let the arrival of offline customers over the interval  $[T, 0]$  be denoted by the Poisson process  $\{N_{1,T}(t), t \in [T, 0]\}$ . Observe that, unlike a conventionally defined Poisson process,  $N_{1,T}(t)$  is non-increasing in  $t$  because  $t$  decreases as time elapses. Similarly, let the arrival of online customers during  $[T, 0]$  be denoted by the Poisson process  $\{N_{2,T}(t), t \in [T, 0]\}$ . For notational convenience, let  $\mathcal{N}_T = (\{N_{1,T}(t)\}, \{N_{2,T}(t)\})$ .

Fundamentally, the retailer seeks to sell all  $n$  units by the end of the horizon  $t = 0$  while selling to as many offline customers as possible. That is, selling a unit to an online customer is preferable to leaving it unsold, but selling the unit to an offline customer is ideal. Therefore, with hindsight knowledge of  $N_{1,T}(0)$  and  $N_{2,T}(0)$ , we can fully characterize the hindsight-optimal solution: sell  $\min\{n, N_{1,T}(0)\}$  units to offline customers, and  $\min\{n - \min\{n, N_{1,T}(0)\}, N_{2,T}(0)\}$  units to online customers. For given  $n$  and  $T$ , the *regret*  $\rho_P(n, T)$  associated with a policy  $P$  is the expected difference between the revenue gained by the policy and the revenue gained by the hindsight-optimal solution. Regret must always be non-negative because no policy can outperform hindsight optimality.

We use  $f(x, y)$  to denote the probability of a Poisson random variable with mean  $x$  taking the value  $y$ ,

and  $f^+(x, y)$  for the probability of a Poisson random variable with mean  $x$  taking a value of  $y$  or greater. The function  $h(w, x, y)$  denotes the probability density at  $y$  of an Erlang random variable with shape  $w$  and rate parameter  $x$ . We denote the positive integers by  $\mathbb{N}$  and the non-negative integers by  $\mathbb{N}_0$ .

## 2. Linear Threshold Policies

Hodge (2008) proposed the heuristic  $\beta$ -LT policy for two customer classes and empirically demonstrated its efficacy. We aim to show that, for any fixed  $\beta \in (\lambda_1, \lambda_1 + \lambda_2)$ , the  $\beta$ -LT policy's regret  $\rho_\beta(n, n/\alpha)$  is uniformly bounded above. This result is formalized in Theorem 1.

**Theorem 1.** *For each  $\beta \in (\lambda_1, \lambda_1 + \lambda_2)$ , there exists a finite positive constant  $\hat{\rho}_\beta$  such that  $\rho_\beta(n, n/\alpha) \leq \hat{\rho}_\beta$  for any  $n \in \mathbb{N}$  and  $\alpha > 0$ .*

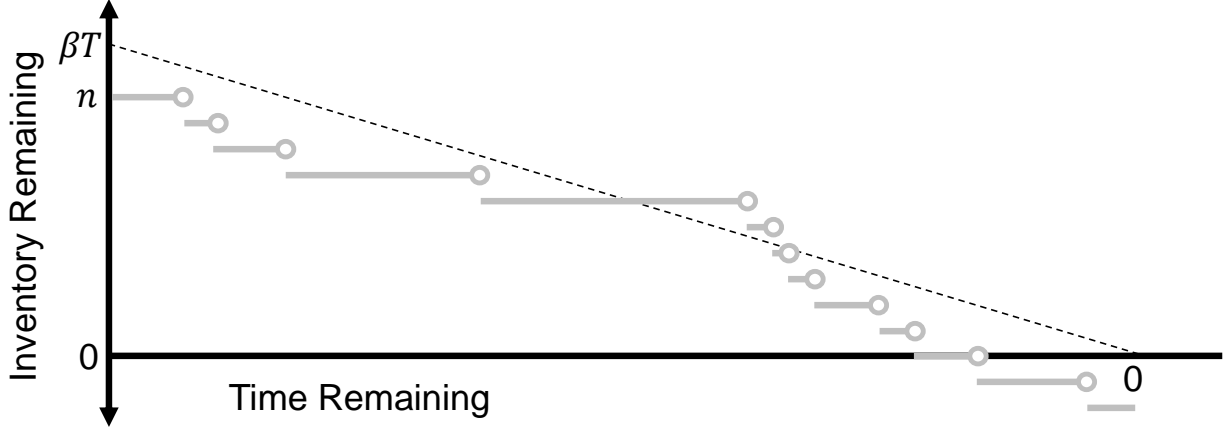
Theorem 1 is implied by Corollary 5, Proposition 8, and Proposition 10; these results correspond to the cases  $\alpha = \beta$ ,  $\alpha > \beta$ , and  $\alpha < \beta$ , respectively. It should be noted that  $\rho_\beta(\cdot, \cdot)$  and  $\hat{\rho}_\beta$ , as well as other functions and constants introduced later, depend on  $\lambda_1$  and  $\lambda_2$ . However, we assume that these rate parameters are fixed beforehand unless stated otherwise, so  $\lambda_1$  and  $\lambda_2$  are omitted from subscripts and function arguments for notational convenience.

### 2.1 Preliminaries

We consider an identical system in which the inventory position can be negative. Specifically, once the inventory depletes, we continue to track the number of offline orders rejected due to the stockout. The *final inventory position* is then the negative of the number of offline orders rejected after inventory depletion (if a stockout occurred) or the actual final positive inventory position otherwise. Let  $\Delta_\beta(s, t)$  denote the absolute expected final inventory position under the  $\beta$ -LT policy with  $s$  inventory remaining and  $t$  time remaining.

For a given policy and realization of  $\mathcal{N}_T$ , the absolute deviation from the hindsight-optimal objective value is determined by two factors: (i) the number of rejected customers of each type that should have been accepted while positive inventory remained, and (ii) the number of accepted customers of each type that should have been rejected. For a policy in which no offline customers are rejected until inventory depletion (e.g., a  $\beta$ -LT policy) a final inventory position of  $x > 0$  implies that achieving the hindsight-optimal objective value would have required accepting at most  $x$  additional online customers. Similarly, a final inventory position of  $x < 0$  indicates that achieving the hindsight-optimal objective value would have required rejecting at most  $y$  online customers and accepting at most  $y$  offline customers instead, where  $y \leq |x|$ . A final inventory position of zero indicates that hindsight optimality was achieved. Therefore, under any  $\beta$ -LT policy,  $\rho_\beta(n, T) \leq p_1 \Delta_\beta(n, T)$  for any  $n$  and  $T$ .

For a given  $\beta$ -LT policy, starting time  $T$ , initial inventory  $n$ , and realization of  $\mathcal{N}_T$ , we refer to the evolution of the inventory position over inventory-time space as the  $\beta$ -induced *sample path* or simply as the *sample path* as in Hodge (2008). Formally, the position of the sample path at time  $t$  can be represented by the function  $R(t, \beta, n, \mathcal{N}_T)$ . Figure 2 illustrates an example sample path and  $\beta$ -line in inventory-time space. By our bound, we know that hindsight optimality is achieved if the sample path ends at zero inventory. Lemmas 2 and 3 state that hindsight optimality is also achieved if the sample path never crosses the  $\beta$ -line during the horizon.



**Figure 2:** Example sample path (in grey) relative to  $\beta$ -line (dashed)

**Lemma 2.** *If  $n - (N_{1,T}(t) + N_{2,T}(t)) \geq \beta t$  for all  $t \in [0, T]$ , the  $\beta$ -LT policy achieves hindsight optimality.*

*Proof.* If  $n - (N_{1,T}(t) + N_{2,T}(t)) \geq \beta t$  for all  $t \in [0, T]$ , the  $\beta$ -LT policy and the hindsight-optimal policy both accept  $N_{1,T}(0)$  offline customers and  $N_{2,T}(0)$  online customers.  $\square$


**Lemma 3.** *If  $n - N_{1,T}(t) < \beta t$  for all  $t \in (0, T]$ , the  $\beta$ -LT policy achieves hindsight optimality.*

*Proof.* If  $n - N_{1,T}(t) < \beta t$  for all  $t \in (0, T]$ , the  $\beta$ -LT policy and the hindsight-optimal policy both accept  $n$  offline customers and no online customers.  $\square$

## 2.2 Starting on the $\beta$ -Line

Suppose first that the initial inventory lies on the  $\beta$ -line (i.e.,  $\alpha = \beta$ ); it will be evident later that this case is fundamental to analyzing the  $\alpha > \beta$  and  $\alpha < \beta$  cases. As in Hodge (2008), we define an *excursion* to mean a segment of the sample path (in inventory-time space) that begins and ends on the  $\beta$ -line and does not intersect the  $\beta$ -line at any point in between. An excursion is almost surely comprised of two consecutive phases. In the first phase, the sample path begins on the  $\beta$ -line and stays strictly above the  $\beta$ -line. The second phase begins when an arrival causes the sample path to jump strictly below the  $\beta$ -line; the sample path then stays strictly below the  $\beta$ -line until it intersects the  $\beta$ -line from below, completing the excursion. If  $\alpha = \beta$ , the sample path's behavior during the horizon can be viewed as a series of independent, complete excursions, likely followed by a partial excursion truncated by the end of the horizon (unless, by chance, an excursion is completed exactly at  $t = 0$ ). Figure 3 depicts a sample path starting on the  $\beta$ -line with two complete excursions during the selling horizon.

The original analysis of Hodge (2008) focused primarily on deriving various statistics about complete excursions, such as the expected length of a complete excursion and the expected proportion of time spent above the  $\beta$ -line during a complete excursion. Because we seek to leverage the  $\rho_\beta(n, T) \leq p_1 \Delta_\beta(n, T)$  bound, we are instead concerned with characterizing the expected inventory position upon truncation of the final partial excursion. To this end, we construct a DTMC model of the sample path by observing the system



**Figure 3:** Sample path with two excursions (red and blue) followed by a partial excursion (grey)

only at integer time points. Although our approach and intent are different from those of the original work, Hodge’s observation that sample paths exhibit queue-like behavior is useful in analyzing the DTMC model.

### 2.2.1 Inventory Position as a Discrete-Time Markov Chain

Assume without loss of generality that  $\alpha = \beta = 1$  and that all other parameters have been scaled correspondingly. For each  $T \in \mathbb{N}_0$ , we define the homogeneous DTMC  $(X_t^T) \triangleq (X_T^T, X_{T-1}^T, X_{T-2}^T, \dots)$  on  $\mathbb{Z}$  as the sample path’s position at time  $t$  relative to the  $\beta$ -line,  $R(t, \beta, T, \mathcal{N}_T) - \beta t = R(t, 1, T, \mathcal{N}_T) - t$ . As usual, inventory is allowed to be negative, and time counts backwards from  $T$ . Thus,  $X_T^T$  is the chain’s initial state, and  $X_0^T$  is the inventory position at the end of the horizon. For mathematical purposes, we allow the DTMCs to evolve past  $t = 0$  in the same manner so that  $X_{-1}^T, X_{-2}^T, X_{-3}^T, \dots$  are also well-defined.

For any chain  $(X_t^T)$ , the one-step transition probability from  $i \in \mathbb{Z}$  to  $j \in \mathbb{Z}$ , denoted  $q_{i,j}$ , is the probability that the sample path is  $j$  units above the  $\beta$ -line at time  $t - 1$  given that the sample path is  $i$  units above the  $\beta$ -line at time  $t$ . For  $i \geq 1$  and  $j \in \{1, 2, \dots, i + 1\}$ , we have  $q_{i,j} = f(\lambda_1 + \lambda_2, i - j + 1)$ . For  $i \leq -1$  and  $j \leq i + 1$ , we have  $q_{i,j} = f(\lambda_1, j - i - 1)$ . Finally, for  $i \geq 0$  and  $j \leq 0$ ,

$$q_{i,j} = \int_0^1 h(i + 1, \lambda_1 + \lambda_2, u) f((1 - u)\lambda_1, -j) du. \quad (2)$$

All other transition probabilities are zero. It is clear from the non-zero transition probabilities that each chain  $(X_t^T)$  is irreducible and aperiodic.

For all  $i \in \mathbb{Z}$ , let  $\mathbb{E}[S_i]$  denote the expected number of steps for the DTMC to return to state  $i$  after departing state  $i$ . Hodge (2008, p. 72) derives an expression for the expected duration of an excursion; in particular, the expression is always finite when  $\lambda_1 < \beta < \lambda_1 + \lambda_2$ . By definition, this expected duration is simply  $\mathbb{E}[S_0]$ , and it therefore follows that  $\mathbb{E}[S_0] < \infty$ . The chain  $(X_t^T)$  is therefore positive recurrent and ergodic, implying the existence of a unique stationary limiting distribution. Let  $\pi = (\pi_i)_{i \in \mathbb{Z}}$  denote the limiting distribution of  $(X_t^T)$ , where  $\pi_i = 1/S_i > 0$  for all  $i \in \mathbb{Z}$ . Observe that  $\pi$  does not depend on  $T$ , because all of the chains  $\{(X_t^0), (X_t^1), (X_t^2), \dots\}$  are time-homogeneous and have identical transition

probabilities.

For each  $T \in \mathbb{N}_0$ , let the DTMC  $(Y_\tau^T)$  denote the restriction of  $(X_t^T)$  to  $\mathbb{N}_0$ . That is, for each  $\tau = T, T-1, T-2, \dots$ , the value of  $Y_\tau^T$  is the value of  $(X_t^T)$  at the  $(T-\tau+1)$ -th occurrence of  $X_t^T \in \mathbb{N}_0$ . Let  $q_{i,j}^+$  denote the transition probabilities for this restricted chain, and let  $\pi^+$  denote the unique stationary limiting distribution of the restricted chain. For all integers  $i \geq 0$  and  $j \geq 1$ ,  $q_{i,j}^+ = q_{i,j}$ . For all  $i \geq 0$ , we have  $q_{i,0}^+ = \sum_{j=0}^{-\infty} q_{i,j} = f^+(\lambda_1 + \lambda_2, i+1)$ .

These are exactly the transition probabilities of the ergodic arrival-time Markov chain associated with a stable  $D/M/1$  queue having arrival rate  $\beta = 1$  and service rate  $\lambda_1 + \lambda_2$  (that is, the Markov chain that tracks the number of jobs in the system just before the arrival of each new job). This queue's limiting arrival-time length distribution is known to have finite expected value (see e.g., Gross and Harris (1998), pp. 248–252). Therefore,  $\sum_{i=1}^{\infty} |i| \pi_i^+ < \infty$ .

For each  $T \in \mathbb{N}_0$ , let the DTMC  $(Z_\tau^T)$  denote the restriction of  $(X_t^T)$  to  $\mathbb{Z} \setminus \mathbb{N}$ . Let  $q_{i,j}^-$  denote the transition probabilities for this restricted chain, and let  $\pi^-$  denote the unique stationary limiting distribution of the restricted chain. For all integers  $i \leq -1$  and  $j \leq 0$ ,  $q_{i,j}^- = q_{i,j}$ . Let the random variable  $A$  denote the duration of time that an arbitrary excursion spends above the  $\beta$ -line, and let  $g : [0, 1] \rightarrow \mathbb{R}^+$  denote the probability density function of  $\lceil A \rceil - A$ . Then, for all integers  $j \leq 0$ , we have  $q_{0,j}^- = \int_0^1 g(u) f(u\lambda_1, j) du$ .

For a stable  $M/D/1$  queue, the associated ergodic departure-time DTMC tracks the number of jobs in the system just after the departure of each job. Let  $\mathcal{Q}$  denote an  $M/D/1$  queue having job arrival rate  $\lambda_1$  and service rate  $\beta > \lambda_1$ . Let  $\psi$  denote the stationary distribution associated with the departure-time DTMC of  $\mathcal{Q}$ ; it is known that  $\sum_{i=1}^{\infty} |i| \psi_i < \infty$  (see e.g., Serfozo (2009), pp. 316–318). Let  $\mathcal{Q}'$  denote an  $M/D/1$  queue having job arrival rate  $\lambda_1$  and service rate  $\beta$  with the additional feature that the first arriving job after each idle period is a partial job whose size is distributed according to  $g(\cdot)$ . Let  $\omega$  denote the stationary distribution associated with the departure-time DTMC of  $\mathcal{Q}'$ . For any particular sequence of job arrival times and any  $k \in \mathbb{N}$ , the number of jobs in the system upon the  $k$ -th job departure in  $\mathcal{Q}'$  is no more than the number of jobs in the system upon the  $k$ -th job departure in  $\mathcal{Q}$ . Therefore,  $\omega_i \leq \psi_i$  for all  $i \in \mathbb{N}$ , and thus  $\sum_{i=1}^{\infty} i \omega_i \leq \sum_{i=1}^{\infty} i \psi_i < \infty$ .

Observe that the departure-time DTMC of  $\mathcal{Q}'$  has identical transition probabilities to the chains  $(|Z_\tau^T|)$ . It follows that  $\omega_i = \pi_{-i}^-$  for all  $i \in \mathbb{N}_0$ , implying that  $\sum_{i=-1}^{-\infty} |i| \pi_i^- < \infty$ . Because  $(Y_\tau^T)$  and  $(Z_\tau^T)$  are both restrictions of the corresponding  $(X_t^T)$ , the inequalities  $\sum_{i=1}^{\infty} |i| \pi_i^+ < \infty$  and  $\sum_{i=-1}^{-\infty} |i| \pi_i^- < \infty$  together imply  $\sum_{i \in \mathbb{Z}} |i| \pi_i < \infty$ . This allows us to prove that the expected absolute inventory position at  $t = 0$  is bounded, independent of  $T$ , given that the initial inventory is on the  $\beta$ -line (i.e., the initial state is  $X_T^T = 0$ ).

**Proposition 4.**  $\sup_{T \in \mathbb{N}} \{ \mathbb{E} [|X_0^T| | X_T^T = 0] \} < \infty$ .

*Proof.* For any  $T \in \mathbb{N}$ , we have  $\mathbb{E} [|X_0^T| | X_T^T = 0] = \mathbb{E} [|X_{-T}^0| | X_0^0 = 0]$  by definition. Hence, it suffices to show  $\sup_{t \in \mathbb{Z} \setminus \mathbb{N}} \{ \mathbb{E} [|X_t^0| | X_0^0 = 0] \} < \infty$ .

Let  $\pi$  denote the stationary distribution, which we know to satisfy  $\sum_{i \in \mathbb{Z}} |i| \pi_i < \infty$ . Then, for any  $t \in \mathbb{Z} \setminus \mathbb{N}$ ,

$$\mathbb{E} [|X_t^0| | X_0^0 \sim \pi] = \sum_{i \in \mathbb{Z}} |i| \mathbb{P}(X_t^0 = i | X_0^0 \sim \pi) = \sum_{i \in \mathbb{Z}} |i| \pi_i \quad (3)$$

and

$$\mathbb{E} [|X_t^0| | X_0^0 \sim \pi] = \sum_{i \in \mathbb{Z}} \mathbb{E} [|X_t^0| | X_0^0 = i] \mathbb{P}(X_0^0 = i) \quad (4a)$$

$$= \sum_{i \in \mathbb{Z}} \mathbb{E} [|X_t^0| | X_0^0 = i] \pi_i \geq \mathbb{E} [|X_t^0| | X_0^0 = 0] \pi_0, \quad (4b)$$

implying  $\mathbb{E} [|X_t^0| | X_0^0 = 0] \leq \frac{1}{\pi_0} \sum_{i \in \mathbb{Z}} |i| \pi_i < \infty$ .  $\square$

The Markov chain model was constructed to ensure  $|\mathbb{E} [|X_0^n| | X_n^n = 0]| = \Delta_\beta(n, n)$  for all  $n \in \mathbb{N}$ . As a result, Proposition 4 directly implies Corollary 5.

**Corollary 5.** *For each  $\beta \in (\lambda_1, \lambda_1 + \lambda_2)$ , there exists  $\eta_{\beta,0} < \infty$  such that  $\Delta_\beta(n, n) \leq \eta_{\beta,0}$  for all  $n \in \mathbb{N}$ . Additionally,  $\rho_\beta(n, n) \leq \rho_\beta^-$  for any  $n \in \mathbb{N}$ , where  $\rho_\beta^- = p_1 \eta_{\beta,0}$ .*

### 2.3 Starting Above the $\beta$ -Line

We next study the case in which the sample path starts above the  $\beta$ -line. As a useful preliminary, we begin by showing that a similar uniform bound exists when the sample path begins no more than one unit below the  $\beta$ -line.

**Lemma 6.** *For each  $\beta \in (\lambda_1, \lambda_1 + \lambda_2)$ , there exists  $\eta_{\beta,1} < \infty$  such that  $\{\Delta_\beta(s-1, s/\beta)\} \leq \eta_{\beta,1}$  for all  $s \in \mathbb{N}$ .*

*Proof.* The argument follows Section 2.2.1 except that the DTMCs begin in state -1 at time  $t = T$ .  $\square$

**Lemma 7.** *For each  $\beta \in (\lambda_1, \lambda_1 + \lambda_2)$ , there exists  $\hat{\eta}_\beta < \infty$  such that  $\max_{s \in \mathbb{R}^+} \{\Delta_\beta(\lfloor s \rfloor, s/\beta)\} \leq \hat{\eta}_\beta$ .*

*Proof.* Let  $s > 0$ , and assume that  $s \notin \mathbb{N}$  to avoid the case from Corollary 5. Without loss of generality, assume that  $\beta = 1$  and all that other parameters have been scaled appropriately. Then,

$$\begin{aligned} \Delta_\beta(\lfloor s \rfloor, s/\beta) &= \sum_{k=0}^{\infty} f((s - \lfloor s \rfloor)\lambda_1, k) \Delta_\beta(\lfloor s \rfloor - k, \lfloor s \rfloor/\beta) \\ &\leq f((s - \lfloor s \rfloor)\lambda_1, 0) \Delta_\beta(\lfloor s \rfloor, \lfloor s \rfloor/\beta) + \sum_{k=1}^{\infty} f((\lceil s \rceil - \lfloor s \rfloor)\lambda_1, k) \Delta_\beta(\lfloor s \rfloor - k, \lfloor s \rfloor/\beta) \\ &\leq \Delta_\beta(\lfloor s \rfloor, \lfloor s \rfloor/\beta) + \Delta_\beta(\lfloor s \rfloor, \lceil s \rceil/\beta). \end{aligned}$$

By Corollary 5 and Lemma 6,  $\hat{\eta}_\beta = \eta_{\beta,0} + \eta_{\beta,1} < \infty$  gives the desired result.  $\square$

If the initial inventory  $n$  is above the  $\beta$ -line (i.e.,  $\alpha > \beta$ ), a sample path can evolve in one of two ways, both of which entail bounded regret. The first possibility is that the sample path ‘jumps’ below the  $\beta$ -line prior to the end of the horizon at least once. If it does so at some time  $t > 0$ , the inventory position is  $\lfloor \beta t \rfloor$ ; thus, Lemma 7 implies that the expected final absolute inventory position is at most  $\hat{\eta}_\beta$  in this case. The other possibility is that the sample path stays at or above the  $\beta$ -line for the entire horizon; in this case, the  $\beta$ -LT policy gives the hindsight-optimal solution (Lemma 2). Proposition 8 formalizes this argument.



**Proposition 8.** Let  $\beta \in (\lambda_1, \lambda_1 + \lambda_2)$ . There exists  $\rho_\beta^+ < \infty$  such that  $\rho_\beta(n, n/\alpha) \leq \rho_\beta^+$  for all  $n \in \mathbb{N}$  and  $\alpha > \beta$ .

*Proof.* Let  $\beta \in (\lambda_1, \lambda_1 + \lambda_2)$ ,  $n \in \mathbb{N}$ , and  $\alpha > \beta$  with  $T = n/\alpha$ . Let  $X(t)$  denote the sample path. Then,

$$\rho_\beta(n, n/\alpha) \leq 0 \cdot \mathbb{P}(X(t) \geq \beta t \ \forall t \in [0, T]) + p_1 \hat{\eta}_\beta \cdot \mathbb{P}(\exists t \in [0, T] : X(t) < \beta t) \leq p_1 \hat{\eta}_\beta. \quad (5)$$

Thus,  $\rho_\beta^+ = p_1 \hat{\eta}_\beta = p_1(\eta_{\beta,0} + \eta_{\beta,1}) < \infty$  is the desired uniform bound.  $\square$

The proof also implies that, if the probability of crossing the  $\beta$ -line were to approach zero as  $n \rightarrow \infty$ , the regret would also tend to zero. This is indeed the case when  $\alpha > \lambda_1 + \lambda_2$ .

**Corollary 9.** For any  $\alpha > \lambda_1 + \lambda_2$  and  $\beta \in (\lambda_1, \lambda_1 + \lambda_2)$ ,  $\lim_{n \rightarrow \infty} \rho_\beta(n, n/\alpha) = 0$ .

*Proof.* It remains to be shown that the probability of the sample path crossing the  $\beta$ -line at least once during the horizon approaches zero as  $n \rightarrow \infty$  and  $T = n/\alpha \rightarrow \infty$ ; denote this probability  $\xi_\beta(n, n/\alpha)$ . Because  $\lambda_1 + \lambda_2 > \beta$ ,  $\xi_\beta(n, n/\alpha)$  is bounded above by the probability of the sample path crossing the  $(\lambda_1 + \lambda_2)$ -line (defined analogously to the  $\beta$ -line) at least once during the horizon.

For all integers  $n \geq 0$ , let  $n\delta = n(1 - (\lambda_1 + \lambda_2)/\alpha) = n - (\lambda_1 + \lambda_2)n/\alpha > 0$  denote the expected final inventory position if all customers were to be accepted (allowing for negative inventory positions as usual). For all  $t \in [0, T]$ , observe that  $(\lambda_1 + \lambda_2)t = n - \mathbb{E}[N_{1,n/\alpha}(t) + N_{2,n/\alpha}(t)] - n\delta$  for all  $t \in [0, T]$ . Without loss of generality, assume for convenience that  $\alpha = 1$  (so that  $n = T$ ) and all other parameters have been scaled appropriately. Then, for any integer  $n \geq 1$ , we have

$$\begin{aligned} \xi_\beta(n, n/\alpha) &\leq \mathbb{P}(\exists t \in [0, n] : n - (N_{1,n/\alpha}(t) + N_{2,n/\alpha}(t)) < (\lambda_1 + \lambda_2)t) \\ &= \mathbb{P}(\exists t \in \{0, 1, 2, \dots, n-1\} : n - (N_{1,n/\alpha}(t) + N_{2,n/\alpha}(t)) \leq (\lambda_1 + \lambda_2)t) \\ &= \mathbb{P}(\exists t \in \{0, 1, 2, \dots, n-1\} : N_{1,n/\alpha}(t) + N_{2,n/\alpha}(t) \geq n - (\lambda_1 + \lambda_2)t) \\ &= \mathbb{P}(\exists t \in \{0, 1, 2, \dots, n-1\} : N_{1,n/\alpha}(t) + N_{2,n/\alpha}(t) \geq \mathbb{E}[N_{1,n/\alpha}(t) + N_{2,n/\alpha}(t)] + n\delta) \\ &\leq \sum_{t=0}^{n-1} \mathbb{P}(N_{1,n/\alpha}(t) + N_{2,n/\alpha}(t) \geq \mathbb{E}[N_{1,n/\alpha}(t) + N_{2,n/\alpha}(t)] + n\delta) \\ &\leq \sum_{t=0}^{n-1} \exp\left(\frac{-\delta^2 n^2}{2(\mathbb{E}[N_{1,n/\alpha}(t) + N_{2,n/\alpha}(t)] + n\delta)}\right) \end{aligned} \quad (6a)$$

$$\leq n \exp\left(\frac{-\delta^2 n^2}{2((\lambda_1 + \lambda_2)n + n\delta)}\right) \quad (6b)$$

$$\leq n / \exp(\delta^2 n / (2\lambda_1 + 2\lambda_2 + \delta)), \quad (6c)$$

where (6a) uses the Poisson upper tail bound in Theorem A.8 of Canonne (2022), and (6c) uses the fact that  $\mathbb{E}[N_{1,n/\alpha}(t) + N_{2,n/\alpha}(t)] \leq (\lambda_1 + \lambda_2)n$  for all  $t \in [0, n]$ . The final expression goes to zero as  $n \rightarrow \infty$ , giving the desired result and an upper bound on the rate of convergence.  $\square$

## 2.4 Starting Below the $\beta$ -Line

If the initial inventory  $n$  is below the  $\beta$ -line (equivalently,  $\alpha < \beta$ ), a sample path can evolve in one of two ways, both of which entail bounded regret. The first possibility is that the sample path intersects the  $\beta$ -line prior to the end of the horizon at least once; in this case, Corollary 5 implies that the expected final absolute inventory position is at most  $\eta_{\beta,0}$ . The other possibility is that the sample path never intersects the  $\beta$ -line prior to the end of the horizon; in this case, the  $\beta$ -LT policy gives the hindsight-optimal solution (Lemma 3). The proof of Proposition 10 formalizes this argument.

**Proposition 10.** *Let  $\beta \in (\lambda_1, \lambda_1 + \lambda_2)$ . There exists  $\rho_{\beta}^- < \infty$  such that  $\rho_{\beta}(n, n/\alpha) \leq \rho_{\beta}^-$  for all  $n \in \mathbb{N}$  and  $\alpha < \beta$ .*

*Proof.* Let  $\beta \in \lambda_1, \lambda_1 + \lambda_2$ ,  $n \in \mathbb{N}$ , and  $\alpha < \beta$  with  $T = n/\alpha$ . Let  $X(t)$  denote the  $\beta$ -induced sample path. Then,

$$\rho_{\beta}(n, n/\alpha) \leq 0 \cdot \mathbb{P}(X(t) < \beta t \ \forall t \in (0, T]) + p_1 \eta_{\beta,0} \cdot \mathbb{P}(\exists t \in (0, T] : X(t) = \beta t) \leq p_1 \eta_{\beta,0}. \quad (7)$$

Thus,  $\rho_{\beta}^- = p_1 \eta_{\beta,0} = \rho_{\beta}^- < \infty$  is the desired uniform bound.  $\square$

Theorem 1 is thus proved with  $\hat{\rho}_{\beta} = \max\{\rho_{\beta}^-, \rho_{\beta}^+, \rho_{\beta}^-\}$ .

As before, the proof of Proposition 10 implies that regret would approach zero if the probability of intersecting the  $\beta$ -line were to approach zero. This is the case when  $\alpha < \lambda_1$ . The proof of Corollary 11 mirrors that of Corollary 9 (albeit using lower tail bounds) and is omitted for brevity.

**Corollary 11.** *For any  $\alpha \in (0, \lambda_1)$  and  $\beta \in (\lambda_1, \lambda_1 + \lambda_2)$ ,  $\lim_{n \rightarrow \infty} \rho_{\beta}(n, n/\alpha) = 0$ .*

## 2.5 Discussion

A benefit of  $\beta$ -LT policies is that they may achieve asymptotic optimality even when  $\lambda_1$  and  $\lambda_2$  are not known with certainty. For example, if  $\lambda_1 \in [18, 22]$  and  $\lambda_2 \in [15, 20]$  but their exact values are unknown, then any  $\beta$ -LT policy with  $\beta \in (22, 33)$  is asymptotically optimal. Observe that the distributions of  $\lambda_1$  and  $\lambda_2$  were not required to deduce the viable range of  $\beta$  — only their supports. In contrast, other constant-regret policies require exact knowledge of  $\lambda_1$  and  $\lambda_2$ . Although the IRT (and other asymptotically optimal policies) could potentially be modified to include online estimation of  $\lambda_1$  and  $\lambda_2$ , a correctly chosen  $\beta$ -LT policy does not require any online modification to guarantee asymptotic optimality.

We conclude our discussion by highlighting a connection between the  $\beta$ -LT policy and the DLP (1). Hodge (2008) derives the expected duration spent above and below the  $\beta$ -line during an excursion as follows.

**Theorem 12** (Hodge 2008, p. 69). *Let  $\nu < 1$  denote the solution to  $\nu e^{-\nu} = (\lambda_1 + \lambda_2)\beta^{-1} e^{-(\lambda_1 + \lambda_2)/\beta}$ , and let  $\kappa = \lambda_1 + \lambda_2 - \beta\nu$ . The expected lengths of time above and below the  $\beta$ -line during an excursion are  $\mathbb{E}[A] = 1/\kappa$  and  $\mathbb{E}[B] = (\lambda_1 + \lambda_2 - \beta)/\kappa(\beta - \lambda_1)$ , respectively.*

The ratio between the expected time above the line (i.e., the expected time during which online orders are accepted) during an excursion and the expected total duration of an excursion is then

$$\frac{\mathbb{E}[A]}{\mathbb{E}[A] + \mathbb{E}[B]} = \frac{1}{\kappa} \left( \frac{1}{\kappa} + \frac{\lambda_1 + \lambda_2 - \beta}{\kappa(\beta - \lambda_1)} \right)^{-1} = \frac{1}{\kappa} \left( \frac{\lambda_2}{\kappa(\beta - \lambda_1)} \right)^{-1} = \frac{\beta - \lambda_1}{\lambda_2}. \quad (8)$$

This is exactly the optimal value of  $z_2$  in (1) when the initial inventory lies on the  $\beta$ -line (i.e.,  $T = n/\beta$ ). Thus, assuming  $T = n/\beta$  and a sufficiently long horizon, the optimal  $z_2$  roughly coincides with the expected proportion of time during which online orders are accepted by the  $\beta$ -LT policy. Although this result is not surprising, it is interesting to observe how we arrive at the optimal DLP proportion via an atypical route.

The arguments underlying Corollaries 9 and 11 also imply that the  $\beta$ -LT policy's online acceptance proportions approach one and zero when  $\alpha > \lambda_1 + \lambda_2$  and  $\alpha < \lambda_1$ , respectively; these are again the respective optimal  $z_2$  proportions in each case. If we had a better understanding of the first crossing times of increasing linear boundaries by Poisson processes, particularly in regard to distributional behavior upon scaling, we would expect analogous results for  $\alpha \in [\lambda_1, \beta) \cup (\beta, \lambda_1 + \lambda_2]$ . Unfortunately, we presently lack the necessary analytical results and/or bounds to formalize the argument for such  $\alpha$ . Because it is not a focus of this work, we defer to Chapter 3 of Zacks (2017) for a summary of existing results on the topic of crossing time distributions for Poisson processes.

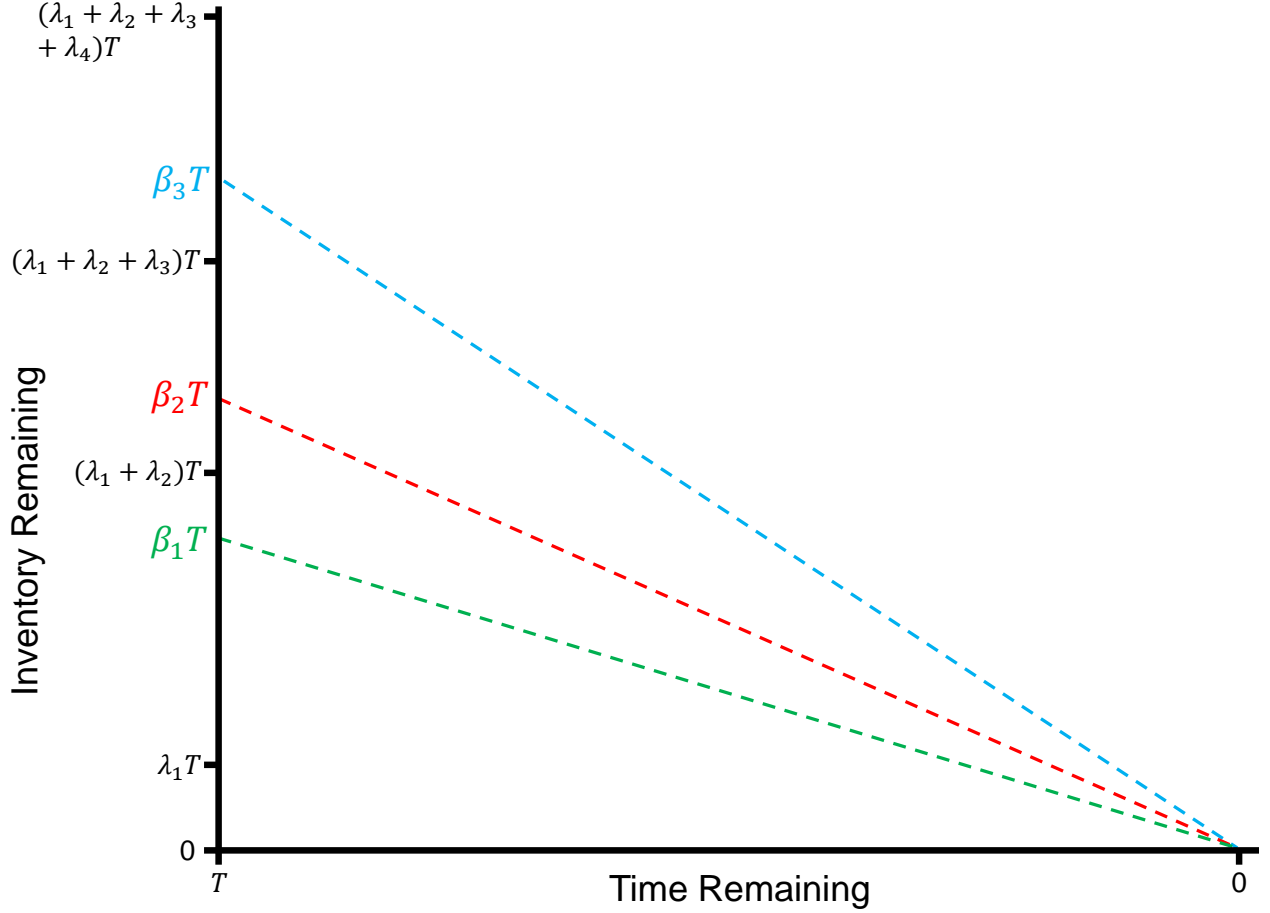
### 3. Generalization to Multiple Customer Classes

Suppose now that there are  $K > 2$  distinct customer classes (denoted type-1, type-2, ..., type- $K$ ) with corresponding known arrival rates  $\lambda_1, \lambda_2, \dots, \lambda_K$  and selling prices  $p_1 > p_2 > \dots > p_K$ . For each  $j \in [K]$ , the arrival of class- $j$  customers during  $[T, 0]$  is given by the Poisson process  $\{N_{j,T}(t), t \in [T, 0]\}$ . In this  $K$ -class setting, we propose a generalization of the  $\beta$ -LT policy. We define a  $\beta$ -LT policy with a vector  $\beta = (\beta_1, \beta_2, \dots, \beta_{K-1})$  where each  $\beta_j$  satisfies  $\sum_{i=1}^{j-1} \lambda_i < \beta_j < \sum_{i=1}^j \lambda_i$ . Unlike in the two-class setting, the operation of the  $\beta$ -LT policy depends on the initial inventory-time ratio  $\alpha = n/T$ .

Suppose that  $\alpha \in [\beta_{k-1}, \beta_k]$  for some  $k \in \{1, 2, \dots, K-1\}$ . The  $\beta$ -LT policy in this scenario proceeds as follows. We begin by accepting all customers of classes  $\{1, 2, \dots, k\}$  and no other customers until a stopping time  $S$ . We define  $S$  as the earliest time at which the resulting sample path is not in the open cone between the  $\beta_{k-1}$ -line and  $\beta_k$ -line. That is,

$$S = \max_{t \leq T} \left\{ t : n - \sum_{j=1}^k N_{j,T}(t) \notin (\beta_{k-1}t, \beta_k t) \right\}. \quad (9)$$

Observe that  $S > 0$ . After time  $t = S$ , the policy behaves in one of two ways. If  $n - \sum_{j=1}^k N_{j,T}(S) \leq \beta_{k-1}S$ , we operate analogously to the two-class policy with the  $\beta_{k-1}$ -line guiding the acceptance decisions. That is, for all  $t < S$ , customers of classes  $\{1, 2, \dots, k-1\}$  are always accepted, and customers of class  $k$  are accepted if and only if the sample path is at or above the  $\beta_{k-1}$ -line. Customers of classes  $k+1$  and above are never accepted. Because  $\sum_{i=1}^{k-1} \lambda_i < \beta_{k-1} < \sum_{i=1}^k \lambda_i$ , the analysis in Section 2 implies that the expected final inventory position in this case is bounded above by some constant  $\zeta_{k-1}^+$  that does not depend on the exact value of  $n$  or  $T$ . In this case, a final inventory position of zero implies that hindsight optimality was achieved; the structure of the policy ensured that the optimal number of class- $k$  customers were accepted and none of the more profitable customers were rejected. A final inventory position of  $x > 0$  implies that at most  $x$  additional customers of classes  $\{k, k+1, \dots, K\}$  should have been accepted to achieve hindsight optimality. A final inventory position of  $y < 0$  implies that at most  $y$  of the accepted customers should have been rejected instead to allow for at most  $y$  additional sales to more profitable customers. Thus, the regret



**Figure 4:** Graphical representation of example four-class  $\beta$ -LT policy

in this case is bounded above by  $p_1 \zeta_{k-1}^+$ .

The other case occurs if  $n - \sum_{j=1}^k N_{j,T}(S) \geq \beta_k S$ . In this case, the  $\beta_k$ -line guides the acceptance decisions. That is, for all  $t < S$ , customers of classes  $\{1, 2, \dots, k\}$  are always accepted, and customers of class  $k+1$  are accepted if and only if the sample path is at or above the  $\beta_k$ -line. Customers of classes  $k+2$  and above are never accepted. Because  $\sum_{i=1}^k \lambda_i < \beta_k < \sum_{i=1}^{k+1} \lambda_i$ , the analysis in Section 2 implies that the expected final inventory position in this case is bounded above by some constant  $\zeta_k^-$  that does not depend on  $n$  or  $T$ . Repeating the preceding analysis implies that the regret in this case is bounded above by  $p_1 \zeta_k^-$ .

It remains to consider the initial states for which  $\alpha \notin [\beta_{k-1}, \beta_k]$ . When  $\alpha < \beta_1$ , the  $\beta$ -LT policy accepts only class-1 customers until the sample path intersects the  $\beta_1$ -line, at which point we simply follow the  $\beta_1$ -LT policy from Section 2 for the remainder of the horizon. As before, the  $\beta$ -LT policy achieves hindsight optimality if the sample path stays below the  $\beta_1$ -line for the entire horizon. Thus, when  $\alpha < \beta_1$ , the regret is bounded above by some finite  $p_1 \zeta_1^-$ . When  $\alpha > \beta_K$ , the  $\beta$ -LT policy accepts all customers until the sample path crosses or intersects the  $\beta_K$ -line for the first time. After this time, we reject class- $K$  customers if and only if we are below the  $\beta_K$ -line; all other customers are always accepted. The  $\beta$ -LT policy achieves hindsight optimality if the sample path stays above the  $\beta_K$ -line for the entire horizon. Thus, when  $\alpha > \beta_K$ ,

the regret is bounded above by some finite  $p_1\zeta_K^+$ . Generalizing over all  $\alpha > 0$ , the  $\beta$ -LT policy’s regret is therefore bounded above by  $\max\{p_1\zeta_1^-, p_1\zeta_1^+, p_1\zeta_2^-, p_1\zeta_2^+, \dots, p_1\zeta_K^-, p_1\zeta_K^+\}$  for any choice of  $n$  and  $T$ .

## 4. Numerical Results

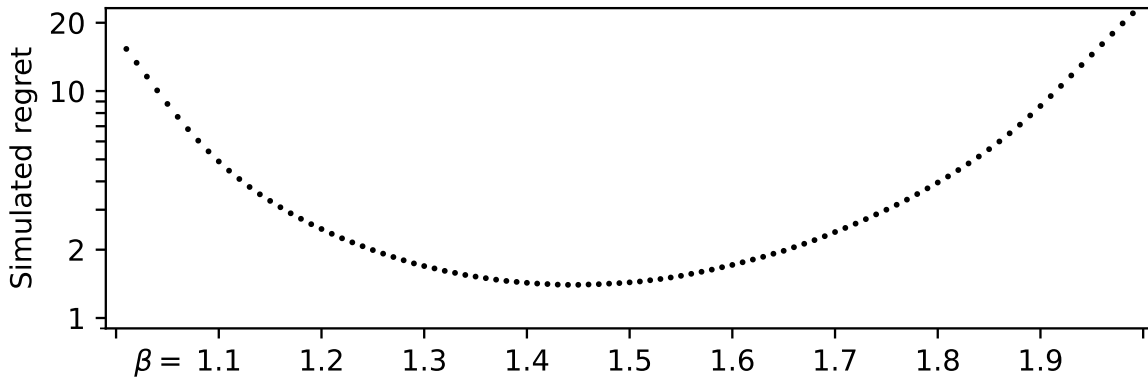
### 4.1 Policy Comparison

We consider a two-class problem with offline and online arrival rates of  $\lambda_1 = \lambda_2 = 1$ . Offline and online selling prices are set at  $p_1 = 2$  and  $p_1 = 1$ , respectively, and  $\alpha = n/T = 1.5$ . We begin by simulating 10000 realizations for each  $T \in \{50, 100, 500, 1000, 5000, 10000, 25000\}$ , and we compare seven choices of  $\beta$  between  $\lambda_1$  and  $\lambda_1 + \lambda_2$ . Table 1 summarizes the results; the ‘Average HO’ column gives the simulations’ average hindsight-optimal objective.

	Average HO	Simulated $\beta$ -LT Regret						
		$\beta = 1.05$	$\beta = 1.1$	$\beta = 1.25$	$\beta = 1.5$	$\beta = 1.75$	$\beta = 1.9$	$\beta = 1.95$
<b><math>T = 50</math></b>	124.9353	3.1432	2.7147	1.7761	1.4060	2.4515	3.6997	4.1805
<b>100</b>	250.1043	4.3772	3.5370	1.9691	1.4428	2.7614	4.9310	5.8886
<b>500</b>	1250.2120	7.5088	4.6858	1.9436	1.4416	2.9732	7.8447	11.5840
<b>1000</b>	2500.0233	8.7768	4.8966	1.9924	1.4356	3.0006	8.5941	14.4756
<b>5000</b>	12500.1551	9.7739	4.8452	1.9338	1.4080	2.9284	8.7940	18.3506
<b>10000</b>	25000.2015	9.9066	4.8818	1.9672	1.4514	2.9583	8.6791	18.6464
<b>25000</b>	62501.7251	9.8401	4.8190	1.9618	1.4529	2.9747	8.7149	18.5961

**Table 1:** Simulated performance of  $\beta$ -LT policies

As expected, regret tends to increase as  $\beta \searrow \lambda_1 = 1$  and  $\beta \nearrow \lambda_1 + \lambda_2 = 2$ . Additionally, we need larger values of  $T$  for the regret to stabilize when  $\beta \searrow 1$  and  $\beta \nearrow 2$ . These phenomena are not symmetric over the interval  $\beta \in (\lambda_1, \lambda_1 + \lambda_2)$ , however. Figure 5 plots the simulated regret for  $T = 1000$  and  $\beta \in \{1.01, 1.02, \dots, 1.99\}$  across 10000 simulations. Of these choices,  $\beta = 1.44$  performs the best with a simulated regret of 1.4001. All  $\beta$  between 1.36 and 1.53 achieve regret less than 1.5 for  $T = 1000$ .



**Figure 5:** Simulated  $\beta$ -LT regrets,  $T = 1000$  (logarithmic vertical axis)

We compare these results to Figure EC.1(e) of Bumpensanti and Wang (2020) that summarizes policies’ empirical performance for a problem with the same parameters. First, we observe that the Fluid Bayes

Selector (FBS) policy (Vera and Banerjee 2021) results in that figure are competitive with our best  $\beta$ -LT policies. Second, the Budget Ratio (BR) policy (Arlotto and Gurvich 2019) results in that figure are also competitive with our best  $\beta$ -LT policies. This is as expected, because our  $\beta$ -LT policy with  $\beta = 1.5$  is essentially the BR policy appropriately scaled to this continuous-time problem. Third, we observe that  $\beta$ -LT policies with  $\beta \approx 1.1$  and  $\beta \approx 1.8$  entail regrets that are similar to those produced by the Infrequent Re-Solving with Thresholding (IRT) policy (Bumpensanti and Wang 2020) for these parameters. All choices of  $\beta$  between these values give even lower regrets, with a wide range of  $\beta \in [1.25, 1.65]$  entailing regrets that are no more than half of those produced by IRT.

Next, we compare  $\beta$ -LT policies to the known optimal thresholding policy. Figure 6 illustrates the optimal policy for  $t \leq 50$ . To avoid computational issues for large  $t$ , we define an extrapolated optimal (EO) heuristic: we follow the exactly computed optimal policy for  $t \leq 100$ , and we linearly extrapolate the optimal threshold function for  $t > 100$  based on the function’s behavior at  $t \approx 100$ . The slope of the linear extrapolation is approximately 1.4420. Although this slope is close to the best-performing  $\beta = 1.44$ , this is merely a coincidence resulting from our choices of  $p_1$  and  $p_2$ . For the values of  $T$  in Table 1, EO achieves regrets of 1.3549, 1.4079, 1.4001, 1.3950, 1.3652, 1.4069, and 1.4091. These are all at most 0.01 lower than the corresponding regrets achieved by the  $\beta$ -LT policy with  $\beta = 1.44$ .

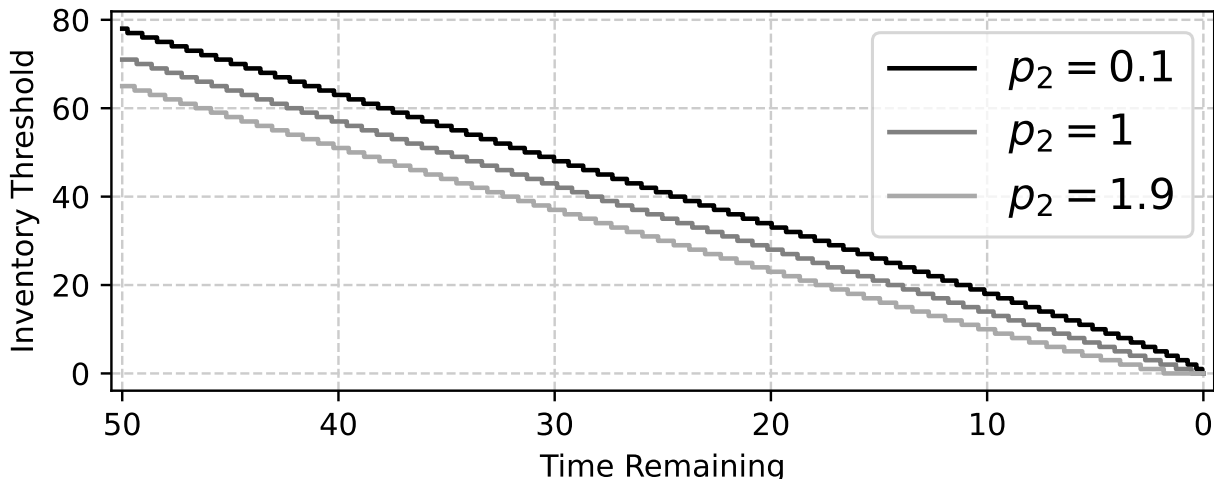


Figure 6: Optimal threshold functions,  $p_2 \in \{0.1, 1, 1.9\}$

## 4.2 Effects of Initial Inventory

Our analysis indicates that sample paths are memoryless with respect to their initial conditions. Thus, one would expect that two values of  $\alpha$  that are both associated with a high probability of crossing a particular  $\beta$ -line would entail similar regrets. We next conduct a brief numerical experiment to illustrate this idea. For each  $\alpha \in \{1, 1.25, 1.75, 2\}$  and  $\beta \in \{1.25, 1.75\}$ , we simulate 10000 realizations for  $T \in \{100, 1000, 10000\}$ . Table 2 displays the simulated regrets.

As expected, any particular  $\beta$ -LT policy performs similarly for each  $\alpha \in \{1.25, 1.5, 1.75\}$ . This effect is even more pronounced for larger values of  $T$ . When  $T = 10000$ , each  $\alpha \in \{1.25, 1.5, 1.75\}$  results in exactly the same simulated regret of 1.9672 for  $\beta = 1.25$ . Also, when  $T = 10000$ , each  $\alpha \in \{1.25, 1.5, 1.75\}$

**Table 2:** Simulated comparison of initial relative inventory levels

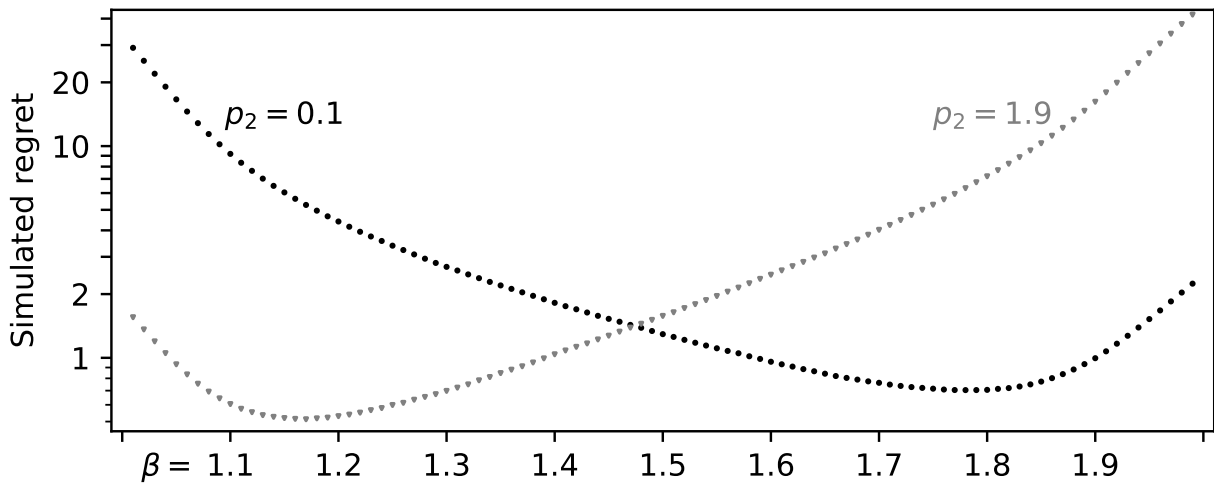
	Simulated $\beta$ -LT Regret							
	$\beta = 1.25$				$\beta = 1.75$			
	$\alpha = 1$	$\alpha = 1.25$	$\alpha = 1.75$	$\alpha = 2$	$\alpha = 1$	$\alpha = 1.25$	$\alpha = 1.75$	$\alpha = 2$
$T = 100$	0.8196	1.9506	1.7675	0.7479	0.8213	2.4067	2.6624	1.0741
<b>1000</b>	0.9424	1.9924	1.9921	0.8804	1.1384	2.9997	3.0006	1.3637
<b>10000</b>	0.9678	1.9672	1.9672	0.9509	1.3522	2.9583	2.9583	1.4354

results in exactly the same simulated regret of 1.9672 for  $\beta = 1.25$ . Each experiment uses the same set of 10000 random streams of customer arrivals. Thus, for each individual trial, these results indicate that the three sample paths beginning at each  $\alpha \in \{1.25, 1.5, 1.75\}$  typically coincide at some point and move together for the remainder of the selling horizon (i.e., the sample paths couple) for large values of  $T$ .

For any  $\alpha \in (\lambda_1, \lambda_1 + \lambda_2)$  and any  $\beta \in (\lambda_1, \lambda_1 + \lambda_2)$ , the probability of crossing the  $\beta$ -line during the horizon approaches one as  $T$  grows. This is not the case when  $\alpha = \lambda_1$  or  $\alpha = \lambda_1 + \lambda_2$ . Recall that a sample path achieves hindsight optimality if it never crosses the  $\beta$ -line during the horizon. Thus, the simulated regrets for  $\alpha \in \{1, 2\}$  are much lower than those for  $\alpha \in \{1.25, 1.5, 1.75\}$ , even for large values of  $T$ .


### 4.3 Effects of Selling Prices

The risk associated with accepting an online customer decreases as  $p_2/p_1$  grows. Hence, we expect lower ideal values of  $\beta$  for higher  $p_2/p_1$  and vice versa. Empirical evidence supports this idea. Returning to the original setting of  $\alpha = 1.5$ , Figure 7 shows the regret across 10000 simulations for  $T = 1000$ ,  $p_1 = 2$ , and  $p_2 \in \{0.1, 1.9\}$ . For  $p_2 = 0.1$ , the lowest regret of 0.7041 is achieved at  $\beta = 1.78$ . For  $p_2 = 1.9$ , the lowest regret of 0.5127 is achieved at  $\beta = 1.17$ . These two  $\beta$ -LT policies are again competitive with their corresponding EO policies, with each achieving regret no higher than 0.3 above EO for all of the values of  $T$  in Table 1. These numerical results for  $p_2 \in \{0.1, 1, 1.9\}$  suggest that the expected performance of an appropriately chosen  $\beta$ -LT policy may be nearly indistinguishable from that of the optimal threshold policy.



**Figure 7:** Simulated  $\beta$ -LT regrets with varying offline selling price,  $T = 1000$  (logarithmic vertical axis)

Figure 6 displays the exactly computed threshold functions for  $t \leq 50$  and  $p_2 \in \{0.1, 1, 1.9\}$ . Although the optimal threshold functions' behavior near the end of the horizon is clearly influenced by  $p_2$ , the three functions' slopes seem to coincide as  $t \rightarrow \infty$ ! For these values of  $p_2$ , the optimal threshold functions' step sizes all appear to converge to approximately 0.69 (corresponding to a slope of approximately 1.44) as  $t$  grows. We observe the same phenomenon for various other choices of  $\lambda_1$  and  $\lambda_2$ , albeit with convergence to slopes other than 1.44. For example, Figure 8 displays optimal threshold functions for  $\lambda_1 = 0.5$  and  $\lambda_2 = 1.5$ . These optimal threshold functions' step sizes all appear to converge to approximately 0.925 (corresponding to a slope of approximately 0.075) as  $t$  grows.



**Figure 8:** Optimal threshold functions,  $\lambda_1 = 0.5$  and  $\lambda_2 = 1.5$

Recall the original conjecture of Hodge (2008) in Section 1. We conclude this work by conjecturing an even stronger statement: as  $t \rightarrow \infty$ , the step widths of the optimal threshold function  $\theta(t)$  converge to a finite positive constant that depends on  $\lambda_1$  and  $\lambda_2$  but not on  $p_1$  or  $p_2$  (assuming  $p_1 > p_2 > 0$ ).

## Acknowledgements

Dipayan Banerjee's work was supported by the U.S. National Science Foundation (DGE-2039655). The authors thank He Wang for helpful comments.

## References

- Alishah EJ, Moinzadeh K, and Zhou YP (2017) Store fulfillment strategy for an omni-channel retailer. Working paper, URL <https://faculty.washington.edu/yongpin/omnichannel%20inventory%20paper.pdf>.
- Arlotto A and Gurvich I (2019) Uniformly bounded regret in the multisecretary problem. *Stochastic Systems*, 9(3):231–260.
- Bumpensanti P and Wang H (2020) A re-solving heuristic with uniformly bounded loss for network revenue management. *Management Science*, 66(7):2993–3009.
- Canonne CL (2022) *Topics and Techniques in Distribution Testing*. now Publishers, Hanover, MA.



- Feng Y and Xiao B (2001) A dynamic airline seat inventory control model and its optimal policy. *Operations Research*, 49(6):938–949.
- Gross D and Harris CM (1998) *Fundamentals of Queueing Theory*. Wiley, New York, NY, 3rd edition.
- Hodge DJ (2008) *Problems of Stochastic Optimal Control and Yield Management*. Ph.D. thesis, University of Cambridge.
- Jasin S and Kumar S (2012) A re-solving heuristic with bounded revenue loss for network revenue management with customer choice. *Mathematics of Operations Research*, 37(2):313–345.
- Liang Y (1999) Solution to the continuous time dynamic yield management model. *Transportation Science*, 33(1):117–123.
- Serfozo R (2009) *Basics of Applied Stochastic Processes*. Springer-Verlag, Berlin, Germany.
- Vera A and Banerjee S (2021) The Bayesian Prophet: a low-regret framework for online decision making. *Management Science*, 67(3):1368–1391.
- Zacks S (2017) *Sample Path Analysis and Distributions of Boundary Crossing Times*. Springer International Publishing AG, Cham, Switzerland.
- Zhao W (1999) *Dynamic and Static Yield Management Models*. Ph.D. thesis, University of Pennsylvania.