Taming the Long Tail: The Gambler's Fallacy in Intermittent Demand Management

(Authors' names blinded for peer review)

Problem definition: "Long tail" products with intermittent demand often tie up valuable warehouse space and capital investment for many companies. Furthermore, the paucity of demand data poses additional challenges for model estimation and performance evaluation. Traditional inventory solutions are not designed for products with intermittent demand. In this paper, we propose a new framework to optimize the choice of "replenishment timing" and "replenishment quantity" for managing the inventory metrics of long tail products, when evaluated over a finite horizon.

Methodology/results: Our analysis is motivated by a recent interesting observation that the gambler's fallacy phenomenon actually holds in a finite number of coin tosses. We use this phenomenon to analyze the inventory problem for intermittent demand to demonstrate that classical inventory models using KPIs such as fill rate, average cost per cycle, or average cost per unit, etc., must necessarily bias the underlying demand distribution to account for the finite horizon effect. We provide the exact closed-form expressions of the biased distribution to account for this effect in performance evaluation. The results show that the choice of replenishment timing and replenishment quantity is essential to superior performance on several key inventory metrics.

Managerial implications: For long tail products, the belief that it is less likely for another demand to arrive shortly after a preceding one (the gambler's fallacy), turns out to be true when performance is tabulated over a finite horizon, even if demands across time are independent. So it pays to delay the replenishment of depleted stocks to save on holding cost and warehouse space. Managers can optimize the replenishment timing, besides choosing the replenishment quantity, to optimize the performance metrics of several classes of inventory problems. This is especially useful for companies managing a large number of long tail products.

Key words: intermittent demand, the gambler's fallacy, long tail products, staggered base stock policy, finite horizon

1. Introduction

Long tail, or slow-moving products, are inventory items with a low turnover that may stagnate business cash flow. However, as the rate of new product introduction in various industries has rapidly accelerated, this proliferation can also result in fewer sales per item, and higher demand variability. For example, Cornacchia and Shamir (2018) examine a large automotive aftermarket parts business wherein 98% of the products, contributing 62% of the sales revenue, are intermittent or long tail in nature. The authors also reveal that intermittent demand can account for 86% of the

SKUs and nearly half of the revenue, even for a branded, fast-moving consumer goods company. Similarly, a recent report by ToolsGroup[1] confirmed that long tail products are becoming more prominent and significant in many industries, with the long tail affecting a substantial portion of revenue (see Table 1 for the impact of long tail products in selected companies, where the columns "SKU's in Tail" and "Revenue in Tail" report the proportions of long tail products and their corresponding revenue contributions, respectively).

Table 1 **Impact of Long Tail in Selected Companies Featured in the Report by ToolsGroup**

Industry	Total SKU's	% SKU's in Tail	Revenue in Tail
Food and Beverage	3,245	44%	36%
Consumer Packaged Goods #1	6,700	74%	52%
Electronics	720	85%	44%
Consumer Packaged Goods #2	9,800	86%	46%
Automotive Aftermarkets Parts #1	5,627	92%	28%
Specialty retailers	7300	96%	71%
Automotive Aftermarkets Parts #2	18,200	98%	62%

The intermittent demand for slow-moving items is especially notorious in the spare parts industry. This is particularly evident in the aviation sector where spare parts are characterized by low demand but high value. In response to this challenge, leading suppliers like AAR Corp, AJW Aviation, and Boeing Distribution (formerly Aviall) have developed new business models to help airlines to grapple with their large capital investments in spare parts inventory through innovative spare parts delivery services. They offer repair management services to airlines, allowing them to access parts when needed, and reducing the cost of maintaining inventory. Consequently, ensuring the proper balance of inventory to meet customer demand while avoiding excessive investment in inventory is critical to operational excellence.

One of the performance metrics in inventory management is the *fill rate*, which is defined as the fraction of demand that can be met through immediate stock availability, without backorders or lost sales. Through performance contracts with customers or internal review, an inventory manager may be held accountable for meeting a target fill rate measured over a finite review horizon, e.g., monthly, quarterly, or yearly. The literature (cf. Thomas (2005)) often set inventory levels according to a constant base stock policy to achieve a target fill rate, assuming that the demand is stationary over an infinite horizon. However, this approach fails to acknowledge the gap between the fill rate measured in an infinite horizon and that measured in a finite horizon. Chen et al. (2003) note that the expected fill rate over a finite horizon is greater than or equal to the infinite horizon fill rate,

[1] `http://www.toolsgroup.com/images/wp-mastering-intermittent-demand.pdf`

which further implies that a short review horizon can benefit the inventory manager as pointed out in Thomas (2005).

Our paper provides an alternative perspective of the finite horizon effect on inventory management, motivated by the recent realization that the so-called *Gambler's Fallacy*, a belief that "if a particular event occurs more frequently than normal during the past, it is less likely to happen in the future (or vice versa)"[2], actually holds for events tabulated over a finite horizon. This observation has led to the re-examination of several empirical studies on the hot-hand versus the gambler's fallacy debate in contexts such as basketball, casino betting, and others. We refer the readers to Miller and Sanjuro (2018) for a thorough discussion on the implications of this finding in the economics literature.

We examine the implications of this phenomenon for long tail products, when demand epochs can be few and far between (e.g., demand for spares). Since the likelihood of having another demand occurring shortly after the preceding one is much lower due to the gambler's fallacy effect, should managers replenish the stock immediately after a demand epoch, or risk having an empty shelf for a short period of time following a demand epoch? This is tantamount to the claim that the safest time to fly is the day after a major plane crash, or the belief that "I Already Crashed Once, So Now I'm Safe".[3]

For several classes of inventory problems, particularly those with performance metrics that are evaluated over a finite horizon, we show that it is essential to optimize both replenishment timing and quantity to effectively leverage the gambler's fallacy effect. For instance, different from the constant base stock policy focused by Thomas (2005), we find that a staggered base stock (SBS) policy, which gradually increases the inventory level by delaying replenishment orders over multiple phases, can further improve the fill rate over a finite horizon even if the arising demands are independent and identically distributed (i.i.d) across time periods. At the same time, the SBS policy requires less warehouse storage space to hold the inventory, which may lead to a significant reduction in warehouse space if the policy is applied to the range of slow-moving items in a company (cf. Table 1).

The potential benefits of SBS are not limited to the metric on fill rate. Our analysis also shows that the staggered policy is beneficial for many other inventory metrics evaluated over a finite horizon. These metrics include *average cost per cycle* (the average cost during a cycle between two consecutive events) and *EBITDA margin* (the Earnings Before Interest, Tax, Depreciation, and Amortization percentage of total revenue).

[2] See `https://en.wikipedia.org/wiki/Gambler27s_fallacy`

[3] `https://www.psychologytoday.com/sg/blog/good-thinking/201305/i-already-crashed-once-so-now-i-m-safe`

1.1. The Finite Horizon Effect: The Gambler's Fallacy in Intermittent Demand

In its most basic form, intermittent demand prediction is tantamount to predicting the outcome in a sequence of independent coin tosses, with "H" (head) indicating a demand epoch, and "T" (tail) otherwise. While it is natural to believe that what happens in the next coin toss does not depend on the outcome of previous tosses, most people exhibit a cognitive bias known as the *Gambler's Fallacy*, believing that the next coin toss is unlikely to be "H" if it follows a streak of successive "H"s. In operations management, this fallacy assumes that another order (demand) for a long tail product arriving immediately after one has been received is less likely, even when the order arrival process is stationary and independent across time. While this is often couched as a mistaken belief or fallacy, Miller and Sanjuro (2018) prove that this phenomenon actually holds for independent Bernoulli data observed over a finite horizon!

We use a coin-flipping example to illustrate this phenomenon. Suppose we start with an initial state of "H" and proceed to flip a coin three more times. Our goal is to determine the proportion of the outcomes "HH" and "HT" in the experiment. Given the assumption of a fair coin, our intuition may suggest that we should see the same proportion of "HH" and "HT" on average, i.e., chances of seeing another "H" immediately after an "H" in the experiment should be 1/2. Surprisingly, for a fair coin, the expected proportion of "H" after another "H" turns out to be smaller than 1/2!

Table 2 The Gambler's Fallacy in 3-flip Experiments

Initial State	3-flip Sequence	Proportion of "H"s after an "H" in the Sequence
H	TTT	$0/1 = 0$
H	TTH	$0/1 = 0$
H	THT	$0/2 = 0$
H	HTT	$1/2$
H	THH	$1/2$
H	HTH	$1/2$
H	HHT	$2/3$
H	HHH	$3/3 = 1$
Average		$\frac{1}{8}(0{+}0{+}0{+}1/2 + 1/2 + 1/2 + 2/3 + 1) = 19/48$

Table 2 summarizes the possible outcomes in the 3-flip experiment[4]. Note that each outcome has an equal weight of 1/8, and the expected proportion of an "H" following another "H" is only 19/48! Thus, we are more likely to see a "T" following an "H", which coincides with the gambler's fallacy. The intuition behind this finite horizon effect is the fact that the expectation is weighted by the equal probability of each sequence rather than the proportion of "H" in that sequence.

[4] Note that Miller and Sanjuro (2018) did not assume an initial state in their analysis, but we have opted for this variant to facilitate the analysis in the inventory management context.

More specifically, in the case of n independent coin flips with a success probability $p \in (0,1)$, if $\hat{p}$ denotes the proportion of "H" that immediately follows an "H" observed in a n-flip experiment, then $\hat{p}$ is a random variable where its expectation $\mathbf{E}[\hat{p}]$ will be presented in Section 3.2. In Figure 1, we show the gap between p and $\mathbf{E}[\hat{p}]$, measured by $(p - \mathbf{E}[\hat{p}])/p$, for different finite horizon n. The gambler's fallacy effect is notably pronounced for small values of p (i.e., slow-moving product) and moderate values of n. For instance, if the intermittent demand has an arrival probability of $p = 0.1$ and we estimate the arrival probability based on $n = 52$ weekly observations in a year, the estimation error induced by the finite horizon effect can be as large as 14.27%.

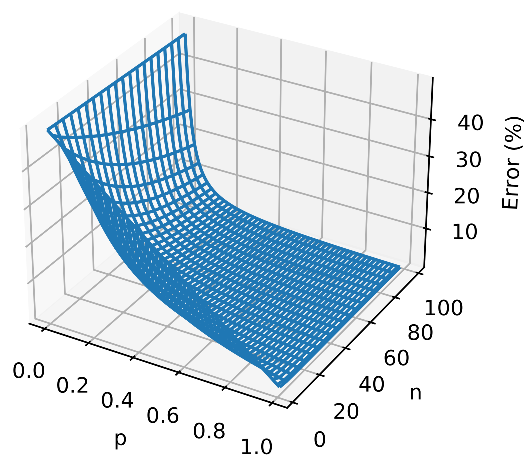

Figure 1 **Estimation error (in %) from the gambler's fallacy for different values of p and n**

Given the gap between p and the expected value of $\hat{p}$, Miller and Sanjuro (2018) assert that empirical work can be vulnerable to such bias, and the results in many empirical works may have to be reversed if the bias in the estimation has been corrected.

1.2. The Finite Horizon Effect in Fill Rate Performance

We use a simple inventory problem with stationary and independent demand to illustrate the subtlety of the fill rate performance measurement over a finite horizon. Suppose the intermittent demand is of unit size with an arrival probability of $p = 0.5$, with "1" indicating a unit demand arrival, "0" otherwise. The unfulfilled demand can be satisfied through emergency replenishment without incurring any additional shortage costs, however, the fill rate performance will be penalized. We compare the following inventory policies over a finite horizon with $n = 3$ periods with unit inventory holding cost of h per unit time:

- Policy A: A constant base stock policy with the order-up-to level of 3/7.

- Policy B: Instead of replenishing immediately after a demand epoch, this is a "delay ordering" policy where the replenishment is postponed for one period after a demand epoch, with an order-up-to level of 1.

Considering all possible demand arrival paths after an arrival (i.e., the initial state of "1"), both policies A and B achieve the same expected fill rate of 1/2, as shown in Table 3. More importantly, by delaying the replenishment order, policy B holds less inventory on average, resulting in a net saving of $1/7h$. This is a reduction of more than 20% of policy A's inventory cost without compromising on the fill rate performance over a finite horizon of 3 periods.

Table 3 Fill Rate and Cost Comparison of Policy A and Policy B

Initial State	Demand Path	Probability	Fill Rate (A)	Fill Rate (B)	Cost (A)	Cost (B)
1	000	1/8	1	1	$9/7h$	2h
1	001	1/8	3/7	1	$6/7h$	h
1	010	1/8	3/7	1	$6/7h$	0
1	100	1/8	3/7	0	$6/7h$	h
1	011	1/8	3/7	1/2	$3/7h$	0
1	101	1/8	3/7	1/2	$3/7h$	0
1	110	1/8	3/7	0	$3/7h$	0
1	111	1/8	3/7	0	0	0
Average			1/2	1/2	$9/14h$	$1/2h$

Note: when there is no demand (e.g., path "000"), we define the fill rate of any policy as 1.

Remark 1. Notice that demands are $\{0, 1\}$, but the base stock level for policy A is 3/7. If the inventory must be integral, we can interpret this as a randomized policy, with a probability of 3/7 choosing a base stock level of 1, and 0 otherwise.

1.3. Main Contributions

We summarize the main contributions of our paper.

- We extend the analysis of the gambler's fallacy phenomenon to inventory management in operations, and derive the closed form expressions for the biased estimates due to the finite horizon effect. We show theoretically that the effect of the gambler's fallacy is decreasing in both the demand occurrence probability and the time horizon, and discuss its implication for the management of long tail products. To the best of our knowledge, this comparative static result is new to the literature on the gambler's fallacy.
- We propose novel staggered base stock (SBS) policies that delay replenishment for products to optimize several performance metrics, including fill rate, cost per cycle, and EBITDA margin, measured over a finite horizon. The corresponding optimal inventory policies are obtained using the biased estimate for demand occurrence, instead of using the underlying (true) demand distribution.

- We extend the analysis to deal with the design of one-for-one inventory replenishment policies, for general inventory system with replenishment lead time. We show that the idea of delay ordering can be suitably incorporated into the one-for-one replenishment policies, and the gambler's fallacy phenomenon can be exploited to improve the performance metrics in this setting, both in terms of inventory cost and space requirement.

While the finite horizon effect is motivated by the observations of Miller and Sanjuro (2018), our results significantly extend the theoretical implications regarding intermittent demand management and provide improved inventory policies that are also of practical value in operations management.

The remainder of this paper starts with a review of related literature in Section 2. In Section 3, we present the gambler's fallacy phenomenon in relation to intermittent demand products and the manner in which it induces bias into the estimation of demand distribution in a finite horizon. We then introduce the demand model with the inter-arrival-demand joint distribution and propose the staggered base stock (SBS) policy for general demand arrivals in Section 4. The empirical validation of our SBS policy compared with other benchmarks is presented in Section 5. Finally, we provide concluding remarks in Section 6.

2. Related Literature

Our paper has connections with several streams of literature—the modeling of intermittent demand related to forecasting methods; inventory management with a focus on the long tail items; and the gambler's fallacy appears as cognitive bias or as the result of the small sample. We discuss these streams of literature below.

Intermittent demand forecasting

Almost no analysis can overcome the problem posed by poor input quality. The seminal paper Croston (1972) addresses the intermittent demand forecasting problem by modifying the traditional exponential smoothing to produce separate point estimates for nonzero demand size and inter-arrival time. This approach demonstrated that the "demand rate", defined as the ratio of nonzero demand and inter-arrival time, is a parameter that can be predicted more accurately for sporadic demand. Certain limitations are identified in Croston's method later, and many papers improved Croston's method, e.g., Syntetos and Boylan (2001, 2005). Nevertheless, as pointed out in Cornacchia and Shamir (2018), improving the forecast accuracy for the demand rate may not be helpful for any inventory models due to the intrinsic demand variability caused by the low order frequency. For example, one may produce a highly accurate estimation of the demand arrival probability, e.g., close to 0.1, for an item with an average sale of one unit every ten days. The issue is that knowing the well-intended demand rate does not elicit information regarding when the next demand will arrive. To address this issue, in this paper, we propose a joint distributional approach to model the intermittent demand, in terms of both size and inter-arrival time. Our approach is able to provide a reliable statistical description of the demand behavior leading to a reduction in perceived demand variability by utilizing information regarding the elapsed time since the last nonzero demand.

Inventory management of slow-moving items

In the related slow-moving inventory management literature, one line of research considers delayed replenishment for intermittent demand, as the demand does not occur frequently. Schultz (1989) revisits the (s, S) policy in a continuous review inventory system, introducing the concept of delayed replenishment orders. Closely related to our paper, Song and Zipkin (1993) describe the dynamics with two-state variables: the current inventory level and the state-of-the-world, to derive the optimal inventory policy. Akcay et al. (2015) quantify estimation inaccuracy for intermittent demand by improving inventory-target estimation, solving the modified newsvendor problem. Although this method mitigates the issue of estimation errors caused by limited data, an independence assumption on the demand size and arrival time and the demand distribution assumption is critically imposed. Syntetos et al. (2009) develop an replenishment policy that relies upon the separate estimates of the inter-demand intervals and demand sizes. In contrast, our paper does not assume any specific demand distribution, which is often unavailable due to limited historical data in the intermittent demand environment. Instead, we directly estimate the joint inter-arrival-demand distribution from data and show it is the correct input when optimizing for the considered performance metrics in a finite horizon.

The Gambler's Fallacy

The gambler's fallacy, also known as Monte Carlo fallacy or the fallacy of the maturity of chances, is the mistaken belief that if something happens more frequently than normal in a given time, it will happen less frequently in the future. The prevalence and magnitude of the gambler's fallacy can be extended into many activities of human life. For example, it has been found that gambler's playing roulette in casinos tend to guess a different color after a streak of the same color, although the presence of both colors is independent. Many papers have demonstrated the existence of the gambler's fallacy empirically, e.g., Clotfelter and Cook (1993) on lottery games, Terrell (1994) on horse racing, and Croson and Sundali (2005) on gambling. The above papers primarily focus on the formation of biased estimations in the presence of probabilities. Kong et al. (2020) examine how game design can shape the biased beliefs of the gambler's fallacy and hot-hand phenomenon.

To analyze the effects of the biased probabilities on decision making, Chen et al. (2016) show that people underestimate the likelihood of sequential streaks, leading to negatively autocorrelated decisions. They provide real-world case evidence on refugee asylum court decisions, loan application reviews, and baseball umpire pitch calls. Rabin and Vayanos (2010) investigate the underlying bias of the gambler's fallacy, proposing prediction models under non-i.i.d. settings. It is widely believed that the decision-makers need to consider the cognitive bias caused by the gambler's fallacy to make wise choices. However, Miller and Sanjuro (2018) demonstrated that the gambler's fallacy exists in a common measure of the conditional dependence of present outcomes on past outcomes even for independent Bernoulli trials. This selection bias indicates that the gambler's fallacy is not a cognitive illusion. Our paper differs from the above literature on the usage of the gambler's fallacy. We show the existence of the gambler's fallacy in the demand of long tail items and adopt the presence of biased estimates as the input to solve for optimal inventory policy. Moreover, we significantly extend the theoretical results of this stream of literature by deriving the exact form of the general bias generated by the finite horizon, and establishing the comparative statics in the gambler's fallacy phenomenon.

3. The Gambler's Fallacy for Long Tail Items

We first demonstrate the existence of the gambler's fallacy in a general Markov chain model and the bias in estimating the inter-arrival time distribution. We then consider the special case of demand arrival being independent in each period, and present the exact form of the biased estimates. Some related propositions derived from the exact form are also discussed. We summarize the main notations in Appendix EC.1.

3.1. The Gambler's Fallacy from a (truncated) Markov Chain

We consider the order arrival process as a general Markov chain model. Let Z denote the random demand inter-arrival times of the system. For ease of exposition[5], suppose $\mathbf{P}(Z=j)=p_j>0$, for $j=1,2,\ldots,J$, and $\sum_{j=1}^{J} p_j=1$. If we sample the random variable Z repeatedly, stopping after the sum of the values sampled exceeds a threshold n, what is the probability that an event $\{Z=1\}$ has been sampled before the sum of the values exceeded n? We show in the rest of this section that this probability is actually smaller than $P(Z=1)$!

[5] More generally, if $p_k=0$ for $k\leq m$, and $p_m>0$, then our result applies to the case of $\mathbf{E}[\hat{p}_m(n)]$

We can simulate the random order arrival using the following Markov chain with J states in the space $\Omega = \left\{\{1\}, \{0_1\}, \{0_2\}, \dots, \{0_{J-1}\}\right\}$, where $\{1\}$ denotes an immediate arrival and $\{0_j\}$ denotes no arrival in the past j periods. The transition probabilities are specified as follows:

$$\begin{cases} p(\{1\},\{1\}) = p_1 & p(\{1\},\{0_1\}) = 1 - p_1 \\ p(\{0_1\},\{1\}) = \frac{p_2}{\sum_{j\geq 2} p_j} & p(\{0_1\},\{0_2\}) = 1 - p(\{0_1\},\{1\}) \\ p(\{0_2\},\{1\}) = \frac{p_3}{\sum_{j\geq 3} p_j} & p(\{0_2\},\{0_3\}) = 1 - p(\{0_2\},\{1\}) \\ \vdots & \vdots \\ p(\{0_{J-1}\},\{1\}) = 1 & 0 \text{ otherwise} \end{cases}$$

We assume the transition probability p_1 is strictly positive. Starting at an state "1", if we simulate n steps of the above Markov chain, tabulation of the empirical proportion of times the transition $\{1,1\}$ takes place immediately after observing $\{1\}$ in the simulation, which is denoted by $\hat{p}_1(n)$. The question that we are interested in is: *Do we have* $\mathbf{E}[\hat{p}_1(n)] < p_1$ *as a generalization to the gambler's fallacy phenomenon in Miller and Sanjuro (2018)?* The proposition below elicits an affirmative answer.

PROPOSITION 1. *For any general Markov chain starting at state "1", we have* $p_1(n) := \mathbf{E}[\hat{p}_1(n)] < p_1$ *for all* n.

Proof. Let $B_n(t)$ denote the indicator of the event that "1" is generated at the t-th position. We also define the indicator $G_n(t)$ for the event that there is a transition from the state "1" at position t to "1" at position $t+1$. Note that $B_n(0) = 1$ since we start at the initial state "1". Then we can write $\hat{p}_1(n)$ using the indicators as

$$\hat{p}_1(n) = \frac{\sum_{i=0}^{n-1} B_n(i) G_n(i)}{\sum_{i=0}^{n-1} B_n(i)}.$$

Note that for each $t \in [1, n-1]$, if $\mathcal{H}_t$ denote the possible histories up to position t, we have

$$\begin{aligned} &\mathbf{E}\left[\frac{B_n(t)G_n(t)}{\sum_{i=0}^{n-1} B_n(i)}\right] = \sum_{h\in\mathcal{H}_{t-1}} \mathbf{E}\left[\frac{B_n(t)G_n(t)}{\sum_{i=0}^{t} B_n(i) + \sum_{i=t+1}^{n-1} B_n(i)}\middle| h; B_n(t) = 1\right]\mathbf{P}[h; B_n(t) = 1] \\ &= \sum_{h\in\mathcal{H}_{t-1}} p_1 \times \mathbf{E}\left[\frac{B_n(t)}{\sum_{i=0}^{t} B_n(i) + 1 + \sum_{i=t+2}^{n-1} B_n(i)}\middle| h; B_n(t) = 1, B_n(t+1) = 1\right]\mathbf{P}[h; B_n(t) = 1] \\ &< \sum_{h\in\mathcal{H}_{t-1}} p_1 \times \mathbf{E}\left[\frac{B_n(t)}{\sum_{i=0}^{t} B_n(i) + \sum_{i=t+1}^{n-2} B_n(i) + B_n(n-1)}\middle| h; B_n(t) = 1\right]\mathbf{P}[h; B_n(t) = 1] \\ &= p_1 \times \mathbf{E}\left[\frac{B_n(t)}{\sum_{i=0}^{n-1} B_n(i)}\right]. \end{aligned}$$

The last inequality holds because conditional on $B_n(t) = 1$, the event $\sum_{i=t+1}^{n-1} B_n(i)$ is statistically equivalent to the event $\sum_{i=t+2}^{n} B_n(i)$ conditional on $B_n(t) = 1$ and $B_n(t+1) = 1$, whereas $B_n(n-1) < 1$ with non-zero probability.

The inequality $\mathbf{E}\left[\frac{B_n(t)G_n(t)}{\sum_{i=0}^{n-1}B_n(i)}\right] < p_1 \times \mathbf{E}\left[\frac{B_n(t)}{\sum_{i=0}^{n-1}B_n(i)}\right]$ above further implies that

$$p_1(n) := \mathbf{E}[\hat{p}_1(n)] = \mathbf{E}\left[\frac{\sum_{i=0}^{n-1}B_n(i)G_n(i)}{\sum_{i=0}^{n-1}B_n(i)}\right] < p_1 \times \mathbf{E}\left[\frac{\sum_{i=0}^{n-1}B_n(i)}{\sum_{i=0}^{n-1}B_n(i)}\right] = p_1.$$

Hence, the gambler's fallacy holds even when the states are generated by the general Markov chain. Q.E.D.

Proposition 1 establishes the bias in estimating the order arrival probability over a finite horizon. Suppose we observe t demand arrivals with inter-arrival times $Z_1, \ldots, Z_t$. We then have the unbiased estimate for the arrival probability with a given t as

$$\mathbf{E}\left[\frac{\chi(Z_1 = 1) + \cdots + \chi(Z_t = 1)}{t}\right] = p_1,$$

where $\chi(\cdot)$ is an indicator function. However, if we consider the order arrivals over a finite horizon of n periods, the sample path of order arrivals is often truncated after n-periods of observation and thus the number of arrivals is random. Starting with a demand arrival at time 0, let τ denote the number of arrivals by period $n-1$, i.e., $\tau = \max\{t : Z_1 + \cdots + Z_t \le n-1\}$. In this case, the estimate $p_1(n)$ using the n-period observation can be expressed in terms of Z_t as follows:

$$p_1(n) = \mathbf{E}\left[\frac{\chi(Z_1 = 1) + \cdots + \chi(Z_\tau = 1) + \chi(\sum_{s=1}^{\tau} Z_s = n-1)\chi(Z_{\tau+1} = 1)}{\tau + 1}\right] < p_1.$$

The last term $\chi(\sum_{s=1}^{\tau} Z_s = n-1)\chi(Z_{\tau+1} = 1)$ contributes to $p_1(n)$ if τ stops at period $n-1$, followed immediately by a demand arrival in period n. From this inequality, we can see the bias of the estimate is induced by the potential truncation of the $(\tau+1)$-th arrival beyond the n-period horizon. $p_1(n)$ is the proportion of times with successive arrivals, among $\tau + 1$ arrivals of demand.

More generally, the bias estimate of p_j can be represented similarly,

$$p_j(n) := \mathbf{E}\left[\frac{\chi(Z_1 = j) + \cdots + \chi(Z_\tau = j) + \chi(\sum_{s=1}^{\tau} Z_s = n-j)\chi(Z_{\tau+1} = j)}{\tau + 1}\right]. \quad (1)$$

Besides the bias in estimating order arrival probability, the gambler's fallacy phenomenon actually holds for other estimates regarding the demand distribution. We next examine the probability of observing at least k demand arrivals in the next ℓ periods immediately after a demand arrival. For any $k \le \ell \le n - t$, we first define the event

$$\mathcal{E}_{k\ell}(t) := \{\text{at least } k \text{ demand arrivals in the next } \ell \text{ periods, following a demand arrival at } t\}.$$

The set of events $\mathcal{E}_{k\ell}(t)$ includes the case of consecutive arrivals and a streak of k arrivals (if $\ell = k$). Then the true theoretical probability of observing at least k demand arrivals in the next ℓ periods immediately after a demand arrival can be denoted by $\mathbf{P}\left[\chi(\mathcal{E}_{k\ell}(0)) = 1\right]$. The estimation using the n-period sample is $\hat{p}_{\mathcal{E}_{k\ell}}(n)$. We then establish the estimation bias by showing $\mathbf{E}[\hat{p}_{\mathcal{E}_{k\ell}}(n)] < \mathbf{P}\left[\chi(\mathcal{E}_{k\ell}(0)) = 1\right]$ in the next theorem.

THEOREM 1. *For any general Markov chain starting at state "1", over a finite horizon n, we have*

$$\mathbf{E}[\hat{p}_{\mathcal{E}_{k\ell}}(n)] := \mathbf{E}\left[\frac{\sum_{i=0}^{n-\ell} B_n(i)\chi(\mathcal{E}_{k\ell}(i))}{\sum_{i=0}^{n-\ell} B_n(i)}\right] < \mathbf{P}\left[\chi(\mathcal{E}_{k\ell}(0)) = 1\right], \ \forall n.$$

Proof. We first show that for all $L > 0$, we have

$$\begin{aligned}
&\mathbf{E}\left[\frac{1}{L + B_n(t+1) + \cdots + B_n(t+\ell)}\middle| B_n(t) = 1\right] \\
=&\mathbf{E}\left[\frac{1}{L + B_n(t+1) + \cdots + B_n(t+\ell)}\middle| B_n(t) = 1, \chi(\mathcal{E}_{k\ell}(t)) = 1\right]\mathbf{P}\left[\chi(\mathcal{E}_{k\ell}(t)) = 1 \middle| B_n(t) = 1\right] \\
&+ \mathbf{E}\left[\frac{1}{L + B_n(t+1) + \cdots + B_n(t+\ell)}\middle| B_n(t) = 1, \chi(\mathcal{E}_{k\ell}(t)) = 0\right]\mathbf{P}\left[\chi(\mathcal{E}_{k\ell}(t)) = 0 \middle| B_n(t) = 1\right] \\
>&\mathbf{E}\left[\frac{1}{L + B_n(t+1) + \cdots + B_n(t+\ell)}\middle| B_n(t) = 1, \chi(\mathcal{E}_{k\ell}(t)) = 1\right].
\end{aligned} \tag{2}$$

The last inequality is due to $\mathbf{E}\left[\frac{1}{L + B_n(t+1) + \cdots + B_n(t+\ell)}\middle| B_n(t) = 1, \chi(\mathcal{E}_{k\ell}(t)) = 0\right] > \mathbf{E}\left[\frac{1}{L + B_n(t+1) + \cdots + B_n(t+\ell)}\middle| B_n(t) = 1, \chi(\mathcal{E}_{k\ell}(t)) = 1\right]$, because by definition $B_n(t+1) + \cdots + B_n(t+\ell) < k$ if $\chi(\mathcal{E}_{k\ell}(t)) = 0$ and $B_n(t+1) + \cdots + B_n(t+\ell) \geq k$ otherwise.

Then for all $t \in [0, n-\ell]$, we have

$$\begin{aligned}
\mathbf{E}\left[\frac{B_n(t)\chi(\mathcal{E}_{k\ell}(t))}{\sum_{i=0}^{n-\ell} B_n(i)}\right] &= \mathbf{E}\left[\frac{B_n(t)\chi(\mathcal{E}_{k\ell}(t))}{\sum_{i=0}^{n-\ell} B_n(i)}\middle| B_n(t) = 1\right]\mathbf{P}\left[B_n(t) = 1\right] \\
&= \mathbf{E}\left[\frac{1}{\sum_{i=0}^{n-\ell} B_n(i)}\middle| B_n(t) = 1, \chi(\mathcal{E}_{k\ell}(t)) = 1\right]\mathbf{P}\left[\chi(\mathcal{E}_{k\ell}(t)) = 1 \middle| B_n(t) = 1\right]\mathbf{P}\left[B_n(t) = 1\right] \\
&< \mathbf{E}\left[\frac{1}{\sum_{i=0}^{n-\ell} B_n(i)}\middle| B_n(t) = 1\right]\mathbf{P}\left[\chi(\mathcal{E}_{k\ell}(t)) = 1 \middle| B_n(t) = 1\right]\mathbf{P}\left[B_n(t) = 1\right] \\
&= \mathbf{E}\left[\frac{B_n(t)}{\sum_{i=0}^{n-\ell} B_n(i)}\right] \times \mathbf{P}\left[\chi(\mathcal{E}_{k\ell}(t)) = 1\right],
\end{aligned}$$

where the inequality holds due to the result in (2) that $\mathbf{E}\left[\frac{1}{\sum_{i=0}^{n-\ell} B_n(i)}\middle| B_n(t) = 1, \chi(\mathcal{E}_{k\ell}(t)) = 1\right] < \mathbf{E}\left[\frac{1}{\sum_{i=0}^{n-\ell} B_n(i)}\middle| B_n(t) = 1\right]$. Therefore, the following holds

$$\begin{aligned}
\mathbf{E}[\hat{p}_{\mathcal{E}_{k\ell}}(n)] = \mathbf{E}\left[\frac{\sum_{i=0}^{n-\ell} B_n(i)\chi(\mathcal{E}_{k\ell}(i))}{\sum_{i=0}^{n-\ell} B_n(i)}\right] &< \sum_{i=0}^{n-\ell} \mathbf{P}\left[\chi(\mathcal{E}_{k\ell}(i)) = 1\right] \times \mathbf{E}\left[\frac{B_n(i)}{\sum_{i=0}^{n-\ell} B_n(i)}\right] \\
&= \sum_{i=0}^{n-\ell} \mathbf{P}\left[\chi(\mathcal{E}_{k\ell}(0)) = 1\right] \times \mathbf{E}\left[\frac{B_n(i)}{\sum_{i=0}^{n-\ell} B_n(i)}\right] \\
&= \mathbf{P}\left[\chi(\mathcal{E}_{k\ell}(0)) = 1\right] \times \mathbf{E}\left[\frac{\sum_{i=0}^{n-\ell} B_n(i)}{\sum_{i=0}^{n-\ell} B_n(i)}\right] = \mathbf{P}\left[\chi(\mathcal{E}_{k\ell}(0)) = 1\right].
\end{aligned}$$

Q.E.D.

Theorem 1 indicates that the probability of having at least k demand arrivals in the next ℓ periods over a finite horizon n is strictly less than the theoretical probability due to the finite horizon effect. For intermittent demand, this theorem implies that the probability of seeing demand arrivals over a finite horizon is less than that of the theoretical probability. Therefore, it is wiser to replenish less inventory for a finite horizon, compared to the inventory solution derived based on the theoretical probability. This observation sheds light on the inventory planning for long tail items which will be discussed in Section 4. Before that, we conclude this section by providing closed-form characterization of the bias estimate for inter-arrival time distribution over a finite horizon when the arrivals are independent.

3.2. The Gambler's Fallacy for Independent and Identical Demand

The above exposition confirms that bias exists when estimating the inter-arrival time distribution in a finite horizon. While it is difficult if not impossible to quantify the bias in the general Markov chain model, we are able to derive the exact form for the bias estimate when demand arrivals are independent. Here, we analyze the exact magnitude of the biases for $p_j(n)$ defined in (1) under the finite horizon effect, when the probability of a demand arriving in each time period is p. We find the exact form of $p_j(n)$ using combinatorial arguments, using ideas from Rinott and Bar-Hillel (2015), who use a similar argument for the gambler's fallacy phenomenon. Note that $p_j(n)$ is simply the proportion of times, among $\tau+1$ arrivals of demand, that the inter-arrival time between successive arrivals is j. Different from Rinott and Bar-Hillel (2015), we assume an initial demand arrival at time 0, which triggers the inter-arrival process.

THEOREM 2. *The exact form of* $p_j(n)$ *is given by*

$$p_j(n) = \sum_{i=1}^{n-j} \binom{n-1-j}{i-1} p^i q^{n-1-i} \frac{p+i}{i+1}.$$

Theorem 2 provides the exact form of the biased estimates $p_j(n), \forall j$, when demand arrivals are independent. Specifically, when $j=1$, we can immediately establish that

$$\begin{aligned} p_1(n) &= \sum_{i=1}^{n-1} \binom{n-1}{i} p^i q^{n-1-i} \left[\frac{ip+i^2}{(i+1)(n-1)}\right] \\ &< \sum_{i=1}^{n-1} \binom{n-1}{i} p^i q^{n-1-i} \left[\frac{i+i^2}{(i+1)(n-1)}\right] \\ &= \sum_{i=1}^{n-1} \binom{n-1}{i} p^i q^{n-1-i} \left[\frac{i}{n-1}\right] = p, \end{aligned}$$

which is consistent with the result of the gambler's fallacy in Proposition 1.

Moreover, with the exact form of the biased estimate, we can characterize the relative bias defined as $\dfrac{p-p_1(n)}{p}$ in the following proposition.

Proposition 2. *The relative bias between p and $p_1(n)$, defined as $\frac{p - p_1(n)}{p}$, is (i) decreasing in p for all n, and (ii) decreasing in n for all p.*

Recall in Figure 1 we show the relative bias termed as the estimation error in percentage. Proposition 2 formalizes this observation that the relative bias is non-increasing in p for all n and decreasing in n when p is small. To the best of our knowledge, this comparative static result is new to the literature on the gambler's fallacy.

4. Inventory Control of Long Tail Items

This section illustrates how the gambler's fallacy affects the control policies for managing slow-moving items in a finite horizon under different performance metrics, e.g., cost per cycle and EBITDA margin. We consider a multi-period inventory control problem where the replenishment lead time is shorter than the review period (i.e., a replenishment order placed will arrive before the end of the current period). This setting (zero lead time) is practical for considering long tail items that are not frequently required. An extension with positive lead time will be discussed in Section 4.3. In each period, we need to determine whether to place an order before the demand realizes. Let the unit holding and shortage cost be h and b, respectively. To facilitate the exposition, we normalize the ordering cost to zero and similarly assume that any leftover inventory after a demand arrival can be disposed of at zero cost.

4.1. Application I: Average Cost per Cycle under Unit Demand Model

Cost per cycle is the inventory cost incurred during a cycle between two consecutive demand epochs. This is a common objective in inventory management, studied for instance in Schultz (1989), Tripathi and Mishra (2014), and Eftekhar et al. (2022). In particular, Schultz (1989) analyzed a unit demand system wherein inter-arrival time distribution satisfies the increasing-failure-rate (IFR) property, and found that delayed ordering is generally optimal. If the inter-arrival time distribution has a constant failure rate (e.g., exponential distribution), the optimal delay would be either zero or infinity, and delay ordering is therefore not needed. This is often used to justify the use of constant base stock policy for these problems. The results are obtained by minimizing the long run expected cost per cycle, without accounting for the finite horizon effect.

Similar results have been obtained in other settings. For instance, in humanitarian operations, the inventory positioning problem studied in Eftekhar et al. (2022) can also be viewed as a variant of intermittent demand management, with the arrival of disasters in a region corresponding to the arrival of demands of a slow-moving item. The authors used the long run average cost per cycle as the performance metric for their analysis, and concluded that as long as no additional funding will be received during the pre-positioning cycle of the relief items, and if the inter-arrival time to the next disaster is exponentially distributed, then the constant base stock (CBS) policy is optimal.

However, both results do not hold when the inventory metrics are evaluated over a finite horizon. In this case, the CBS policy may no longer be optimal, because of the gambler's fallacy. We establish this fundamental result in the rest of this section.

4.1.1. Unit Demand Model To investigate how the finite horizon effect influences the inventory control policy, we consider the case in which each arrival only has unit demand. A unit of inventory will be replenished upon a demand epoch, which consumes a unit of the product. For ease of exposition, we assume the replenishment lead time is zero. The challenge is to determine the timing of this replenishment order.

Consider the policy π^{adp}, in which the replenishment is delayed for W periods after a unit demand has arrived, say at time t. In this policy, the inventory level is kept at 0 from $t+1$ to $t+W$, and raised to a level of 1 at time $t+W+1$. In essence, this policy exploits the gambler's fallacy phenomenon that the next order will not arrive immediately after a demand epoch.

Let Z denote the (random) inter-arrival time of the next unit of demand. With a cycle defined by the periods between two demand arrivals, the cycle cost is then given by

$$C(\pi^{adp}) = b \times \chi(Z \leq W) + h \times \chi(Z = W+2) + 2h \times \chi(Z = W+3) + \cdots .$$

If the inventory metric is about the long-run average cycle cost, i.e., the expected cost per cycle in the long run, then

$$\mathbf{E}[\mathcal{C}(\pi^{adp})] = b \times \sum_{j=1}^{W} p_j + h \times p_{W+2} + 2h \times p_{W+3} + \cdots$$

by utilizing the inter-arrival time distribution with $E(\chi(Z=j)) = P(Z=j) := p_j$ (e.g., see Schultz, 1989). However, in a finite horizon problem with the cost evaluated over n time periods, the expected cycle cost defined above does not hold. To see this, let $\tau = \max\{t : Z_1 + \cdots + Z_t \leq n-1\}$, and assume the initial state is immediately after a demand epoch. The expected cost per cycle in the finite horizon case can be calculated as follows,

$$\mathbf{E}[\mathcal{C}(\pi^{adp})] = \mathbf{E}\left[\frac{C_1(\pi^{adp}) + C_2(\pi^{adp}) + \cdots + C_\tau(\pi^{adp}) + C_{tail}(\pi^{adp})}{\tau + 1}\right],$$

where $C_t(\pi^{adp})$ is the cycle cost of the t-th cycle with inter-arrival time Z_t, and $C_{tail}(\pi^{adp})$ is the cost from the last arrival period $\sum_{t=1}^{\tau} Z_t + 1$ to the end of horizon period n. The term $C_{tail}(\pi^{adp})$ is to include the possible inventory cost in the last cycle (i.e., the $\tau+1$-th cycle) from period $\sum_{t=1}^{\tau} Z_t + 1$ to period n.

In the following proposition, we demonstrate that the expected cost per cycle can be formulated with the biased estimate $p_j(n)$ in the unit demand model.

PROPOSITION 3. *In the unit demand model, the expected cost per cycle* $\mathbf{E}[\mathcal{C}(\pi^{adp})]$ *with the delay order policy specified by* W *periods can be expressed as*

$$\mathbf{E}[\mathcal{C}(\pi^{adp})] = b \times \sum_{j \le W} p_j(n) + h \sum_{j=W+2}^{n} (j - W - 1) p_j(n),$$

where $p_j(n) := \mathbf{E}\left[\frac{\chi(Z_1 = j) + \cdots + \chi(Z_\tau = j) + \chi(\sum_{s=1}^{\tau} Z_s = n - j)\chi(Z_{\tau+1} = j)}{\tau + 1}\right]$ *is the biased estimates in finite horizon previously defined in* (1).

Proof. For the t-th cycle, $t \in [\tau]$, we have

$$C_t(\pi^{adp}) = b \times \mathbf{E}\Big[\chi(Z_t \le W)\Big] + h \sum_{j=W+2}^{n} (j - W - 1)\mathbf{E}\Big[\chi(Z_t = j)\Big],$$

and

$$C_{tail}(\pi^{adp}) = b \times \mathbf{E}\Big[\sum_{j \le W} \chi(\sum_{s=1}^{\tau} Z_s = n - j)\chi(Z_{\tau+1} = j)\Big] + h \sum_{j=W+2}^{n} (j - W - 1)\mathbf{E}\Big[\chi(\sum_{s=1}^{\tau} Z_s = n - j)\chi(Z_{\tau+1} = j)\Big].$$

Following the above analysis, the expected cycle cost is

$$\begin{aligned}
\mathbf{E}[\mathcal{C}(\pi^{adp})] =& \mathbf{E}\Big[\frac{C_1(\pi^{adp}) + C_2(\pi^{adp}) + \cdots + C_\tau(\pi^{adp}) + C_{tail}(\pi^{adp})}{\tau + 1}\Big] \\
=& b \times \mathbf{E}\Big[\frac{\chi(Z_1 \le W) + \cdots + \chi(Z_\tau \le W) + \sum_{j \le W} \chi(\sum_{s=1}^{\tau} Z_s = n - j)\chi(Z_{\tau+1} = j)}{\tau + 1}\Big] \\
&+ h \times \mathbf{E}\Big[\frac{\chi(Z_1 = W + 2) + \cdots + \chi(Z_\tau = W + 2) + \chi(\sum_{s=1}^{\tau} Z_s = n - W - 2)\chi(Z_{\tau+1} = j)}{\tau + 1}\Big] \\
&+ \cdots + (n - W - 2) \times h \times \mathbf{E}\Big[\frac{\chi(Z_1 = n - 1) + \cdots + \chi(Z_\tau = n - 1) + \chi(\sum_{s=1}^{\tau} Z_s = 1)\chi(Z_{\tau+1} = j)}{\tau + 1}\Big] \\
=& b \times \sum_{j \le W} p_j(n) + h \sum_{j=W+2}^{n-1} (j - W - 1) p_j(n)
\end{aligned}$$

This is the expected cost per cycle model with $p_j(n)$ in place of p_j. Q.E.D.

To obtain the optimal adaptive policy in which the replenishment is delayed for W period, we need to minimize $\mathbf{E}[\mathcal{C}(\pi^{adp})]$ over W. The optimal solution can be obtained for discrete W by enumeration. Let $W^*(n)$ denote the optimal delay for an n-period problem, and W^* the optimal delay policy for a long run average cycle cost model. We next show that the optimal delay in a finite horizon cost per cycle model is actually longer than the delay in the long run average cycle cost model, showing that the gambler's fallacy effect actually leads to longer delay in the replenishment timing.

THEOREM 3. $W^*(n) \ge W^*$ *for all* n.

We note that after accounting for the finite horizon effect, the longer delay in $W^*(n)$ not only improves on the performance of the inventory metric, it also saves on warehouse space usage to hold the inventory. The aggregate savings for a company with a large number of slow-moving items can be significant using the optimal delay ordering policy.

4.1.2. Numerical Examples According to Schultz (1989), CBS policy is optimal when the failure rate of the inter-demand distribution is constant; however, due to the existence of the gambler's fallacy, the constant inventory policy fails to be optimal under the objective of cost per cycle (as well as the expected cycle cost) in a finite horizon, although the inter-arrival time distribution has a constant failure rate. Indeed, as highlighted by Theorem 3, replenishment should be delayed longer when the objective is measured in a finite horizon. In the following numerical example, we demonstrate the benefit of using the biased probability estimate $p_j(n)$ provided in Proposition 3, rather than the true probability p_j to identify the optimal delay ordering policy for inventory control in a finite horizon.

Example 1. We follow the example used for Figure 1 in Section 1 where the unit demand arrives independently with arrival probability p. Correspondingly, the biased probability $p(n)$ can be obtained through simulation. We test the performance (cost per cycle) of using the true probability p and biased probability $p(n)$ for the inventory problem with finite horizon n periods and $h = 1,\ b = 10$. Following the true distribution without considering the finite horizon would likely render the CBS policy optimal. It is not hard to show that when the ratio of h and b is less than $1/p$, then the constant order-up-to level should be 1; otherwise, it should be 0.

In this example, we compare the costs of two ordering policies, one that incorporates the finite horizon effects and one that does not, for a range of values of the parameters p and n. As shown in Figure (2a), using the biased estimates $p(n)$ can save up to 10% of the cost per cycle compared to using the true probability p. Nevertheless, even when the time horizon n is large, the cost savings remain significant as long as the occurrence probability p is small. We notice that the cost saving is significant when p is around 0.1, exactly the value of h/b. If the finite horizon is not considered, h/b is the threshold for a constant ordering policy with unit base stock. The cost savings are significant around this value due to the sudden switch of base stock level from 0 to 1 for the constant ordering policy. When p deviates from h/b, both policies exhibit similar behavior, leading to small cost savings. This non-monotonic relationship between cost savings and p for given time horizon is evident in Figure (2a). We provide more evidence of this phenomenon for different shortage costs with $b = 8$ and $b = 12$ in Appendix EC.4. Moreover, we demonstrate the corresponding optimal delays using biased estimates $p(n)$ in Figure (2b). The results suggest that the optimal delay generally decreases with p.

4.2. Application II: Cost per Unit Sold or EBITDA Margin under General Demand The "Earnings Before Interest, Tax, Depreciation, and Amortization" (EBITDA) margin is the EBITDA percentage of total revenue earned in a finite horizon. This is an important financial metric used by analysts and investors to measure and evaluate the ability of a company to generate income (profit) relative to revenue, balance sheet assets etc. It reveals how well a company

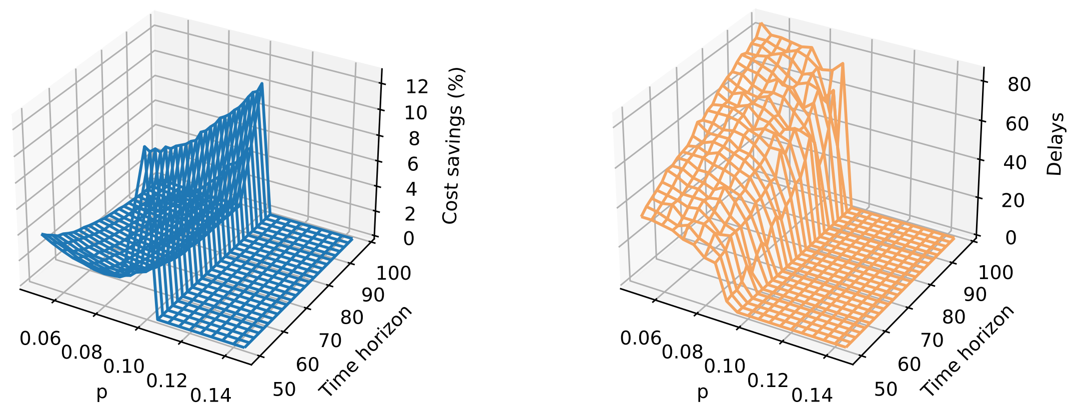

(a) The cost savings using biased estimates. (b) Delays using biased estimates.

Figure 2 The impact of biased estimates.

utilizes its assets to produce profit and value for shareholders. There are a couple of similar metrics capturing the balancing act faced by companies, like ROCE (return on capital employed) and operating margin, etc.[6] They are attached great importance by investors and managers.[7] Overall, the EBITDA margin, ROCE, and operating margin principle are two-dimensional performance metrics that reveal the trade-off between return and investment. They help to derive more realistic targets than one-dimensional performance metrics (e.g., total inventory cost) in practice. As noted in Rogers et al. (2010) and Hançerlioğullan et al. (2017), these metrics are important measures of organizational profitability.

The EBITDA margin varies in different industries.[8] Interestingly, this metric is intimately related to the "cost per unit of goods sold" (CPUGS) inventory metric. CPUGS can be derived from the total cost of inventory operations, divided by the total units of goods sold. This cost metric essentially measures the average cost per unit of goods sold over a finite horizon. The EBITDA margin and CPUGS are related in the following way:

$$\text{EBITDA margin}=1-\frac{\text{CPUGS}}{\text{price}}.$$

Given a product price, maximizing EBITDA margin is thus equivalent to minimizing CPUGS. In this section, we minimize the CPUGS alternatively for convenience. Note that under unit demand setting, the formula of cost per cycle is equivalent to CPUGS.

[6] ROCE is defined as the EBITDA generated over the capital employed, and the operating margin is calculated by dividing a companys operating income by its net sales.

[7] For example, the Japanese multinational electronics manufacturing corporation Casio shifted the positioning strategy based on the objective of increasing the operating margin. The strategic shift is well documented in the Casio annual report 2008. It states the management strategy is to ensure high profitability, achieving an overall operating income margin of 10% or more. In its 2016 annual report, Casio claims the company has made a shift with a 12% operating margin and achieved higher profitability.

[8] https://assetsamerica.com/ebitda-margin/

In the remainder of this section, we examine how a general demand model can be incorporated into our framework to exploit the gambler's fallacy effect. We use a joint distributional approach to model the intermittent demand, where the nonzero discrete demand is $D \in [K] = \{1, 2, \cdots, K\}$[9].

4.2.1. Intermittent Demand Modeling The size and timing of demand arrival are often related in the case of intermittent demand products. To better characterize the demand intermittency, the intermittent demand modeling in this paper is simple and practical—we represent the interaction between the demand size and inter-arrival time using their joint distribution, for which the empirical counterpart can be directly obtained from historical demand observations. Specifically, we represent the empirical distribution in a two-dimensional histogram (discrete or continuous). For illustration, we simulate synthetic time series with 200 periods[10] (including many zero-demand periods) in Figure (3a). The pattern between the demand size and inter-arrival time can be visualized by tabulating the empirical frequencies and representing it using a two-dimensional histogram as an estimate of the joint distribution in Figure (3b). This intermittent demand modeling approach explicitly captures the dependence between the size and timing of demand arrivals that often appears in practice.

Let p_{kj} denote the joint probability that demand of size k arrives with inter-arrival time of j periods, with $j \in [J]$ and $k \in [K]$. We use this joint distribution representation to estimate the probability of a demand arriving in each time period t, with an initial arrival at time $t = 0$. To this end, let z_t be the elapsed time since last demand epoch, measured at the beginning of period t. Let $z_{t+1} = 1$ if a nonzero demand arrives in period t, otherwise $z_{t+1} = z_t + 1$. By conditioning on the elapsed time z_t, the probability of a demand arriving in period t follows the Bayes' rule:

$$\mathbf{P}(z_{t+1} = 1 | z_t) = \frac{\sum_{k=1}^{K} p_{kz_t}}{\sum_{j=z_t}^{J} \sum_{k=1}^{K} p_{kj}}$$

The above statistics can be calibrated from the empirical data, e.g., using the historical demand estimates from the two-dimensional histogram.

4.2.2. Staggered Base Stock Policy The staggered base stock policy gradually builds up the inventory of slow-moving items, replenishing even if there is no demand arrival at the current time period. It can be described using the following state-dependent policy π^{sd}: At the beginning of each period, given the elapsed time $j \in [J]$, raise the inventory level up to S_j. The levels S_j are predetermined, with $S_{j-1} \leq S_j$ for all $j \in [J]$. Let $S_0 = 0$, and $S_j \geq 0, \forall j \in [J]$. We refer to this the **staggered base stock** policy, short as **SBS** policy.

[9] The notation $[N]$ for positive integer N stands for the set $\{1, 2, \cdots, N\}$ throughout the paper.

[10] In the simulation, the demand occurrence probability is $p = 0.1$ and the nonzero demand is Poisson distribution with mean $\lambda = 5$.

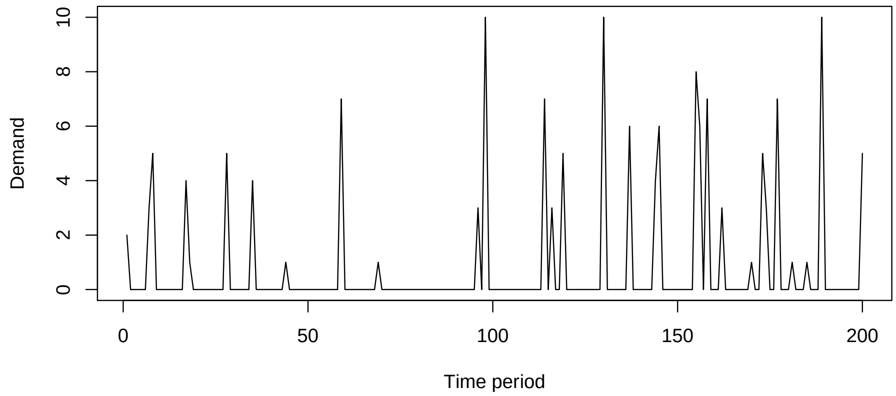

(a) Time series of demand in 200 periods

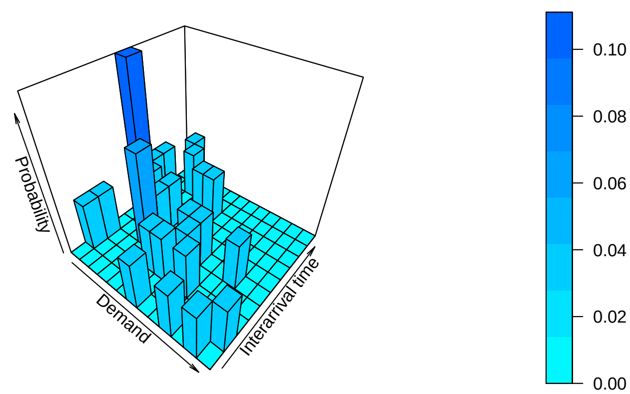

(b) Two-dimensional histogram

Figure 3 From time series to histogram: an example

Consider the n-period problem, where X_t is the demand in period t, for $t = 1, \ldots, n$. Let $\tau = \max\{t : Z_1 + \cdots + Z_t \leq n-1\}$. It is possible that there is no demand arrival from period $\sum_{t=1}^{\tau} Z_t + 1$ to period n when $X_n = 0$. In this case, it is possible that holding cost occurs without any goods sold in the last cycle. More generally, when there is no demand in all the n periods, the denominator of the expected CPUGS is zero. We define the CPUGS as follows:

$$\mathbf{E}[\mathcal{C}(\pi^{adp})] = \mathbf{E}\left[\frac{C_1(\pi^{adp}) + C_2(\pi^{adp}) + \cdots + C_\tau(\pi^{adp}) + C_{tail}(\pi^{adp})}{X_{Z_1} + \ldots + X_{Z_1 + \ldots + Z_\tau} + \max\{1, X_n(Z_1, \ldots, Z_\tau)\}}\right],$$

where $C_t(\pi^{adp})$ is the cycle cost of the inter-arrival time Z_t, and $C_{tail}(\pi^{adp})$ is the cost from period $\sum_{t=1}^{\tau} Z_t + 1$ to period n. Similar to the previous setting, $C_{tail}(\pi^{adp})$ is added to include the possible inventory cost in the last cycle. The demand size used to calculate the expected CPUGS in the last cycle from period $\sum_{t=1}^{\tau} Z_t + 1$ to period n is $\max\{1, X_n(Z_1, \ldots, Z_\tau)\}$ to avoid the denominator being zero in the case when there are no demand arrival in the n periods, and also to account for the initial state. Note that $X_n(Z_1, \ldots, Z_\tau)$ is independent of $Z_1, \ldots, Z_\tau$ if demand quantity is independent of the inter-arrival times (e.g., the i.i.d. case). In the rest of the paper, we write X_n

instead and omit the dependence on $Z's$, to simplify the exposition. The following theorem suggests that the expected CPUGS can also be expressed by the biased estimates $\tilde{p}_{kj}(n)$.

THEOREM 4. *For general demand setting, let $Y_1,\dots,Y_\tau \in [K]$ denote the random nonzero demand sizes at period $Z_1, Z_1+Z_2,\dots,Z_1+\cdots+Z_\tau$. Starting from an initial state after a nonzero demand, over a finite horizon of n periods, the expected CPUGS can be expressed by the cycle cost with the biased estimates $\tilde{p}_{kj}(n)$,*

$$\mathbf{E}[\mathcal{C}(\pi^{sd})] = \sum_{j\in[J]}\sum_{k\in[K]} \tilde{p}_{kj}(n)\frac{\sum_{\ell=1}^{j-1} hS_\ell + h(S_j-k)^+ + b(k-S_j)^+}{k}, \tag{3}$$

where $\tilde{p}_{kj}(n) := \mathbf{E}\left[\sum_{t\in[\tau]} \frac{\chi(Z_t=j,Y_t=k)}{(\sum_{s\in[\tau]^+, s\neq t} Y_s+k)/k} + \frac{\chi(\sum_{s=1}^{\tau} Z_s=n-j)\chi(Z_{\tau+1}=j,X_n=k)}{(\sum_{s\in[\tau]} Y_s+k)/k}\right]$, $[\tau]^+ = [\tau]\cup\{n\}$ *and* $Y_n = \max\{1, X_n\}$[11].

Following the discussions above, the joint inter-arrival-demand distribution $\tilde{\mathbb{Q}}$ with elements $\tilde{p}_{kj}(n)$ can be represented by a two-dimensional histogram. It can be estimated numerically via sampling. Based on the results in Theorem 4, our problem reduces to

$$\begin{aligned} \min_{S_j\geq 0} \quad & \sum_{j\in[J]}\sum_{k\in[K]} \tilde{p}_{kj}(n)\frac{\sum_{\ell=1}^{j-1} hS_\ell + h(S_j-k)^+ + b(k-S_j)^+}{k} \\ \text{s.t.} \quad & S_{j-1}\leq S_j, \forall j\in[J] \end{aligned} \tag{4}$$

If $Y_t = k$ is a constant for each demand arrival, then we can approximate $\tilde{p}_{kj}(n)$ as

$$p^c_{kj}(n) := \mathbf{E}\left[\frac{\sum_{t\in[\tau]}\chi(Z_t=j,Y_t=k) + \chi(\sum_{s=1}^{\tau} Z_s = n-j)\chi(Z_{\tau+1}=j, X_n=k)}{\tau+1}\right], \tag{5}$$

which is the empirical proportion of the event that the inter-arrival time is j and demand size is k. When the demand size does not vary by much, and when data is scarce as in the intermittent demand case, we can use $p^c_{kj}(n)$ as an approximation to $\tilde{p}_{kj}(n)$ in the optimization model.

Given the exact value of $\tilde{p}_{kj}(n)$, problem (4) is a linear program. The size of constraints and variables of the linear program is linear in J. Nevertheless, when J and K are too large to deal with, in practice, we can change the bin size by aggregating values to reduce the problem size. A similar analysis also applies to continuous demand.

4.2.3. Staggered Base stock Policy for Continuous Demand Our previous development of the staggered base stock policy is based on the discrete inter-arrival-demand distribution. In this subsection, we consider a more general case when demand is continuous.

Let D denote the random demand size with probability density function $f_j(\cdot)$ conditional on inter-arrival time j. We have the following theorem on CPUGS for the continuous demand case.

[11] We abuse the notation slightly for ease of exposition. Note that $Y_n \nsim Y_t$ for $t\in[\tau]$. Hence there is a need to split the calculation of $\tilde{p}_{kj}(n)$ into two ratios.

THEOREM 5. *For general continuous demand, let $Y_1, \ldots, Y_\tau$ denote the random nonzero demand size at time $Z_1, Z_1+Z_2, \ldots, Z_1+\cdots+Z_\tau$. Starting from an initial state after nonzero demand, over a finite horizon of n periods, the expected CPUGS can be expressed by the cycle cost with the biased estimates $\tilde{p}_j(n,\theta)$:*

$$\mathbf{E}[\mathcal{C}(\pi^{sd})] = \sum_{j\in[J]} \int_\theta \tilde{p}_j(n,\theta) \frac{\sum_{\ell=1}^{j-1} hS_\ell + h(S_j-\theta)^+ + b(\theta-S_j)^+}{\theta} d\theta,$$

where $\tilde{p}_j(n,\theta) := \mathbf{E}\left(\sum_{t\in[\tau]} \frac{\chi(Z_t=j)}{(\sum_{s\in[\tau]^+, s\neq t} Y_s+\theta)/\theta} + \frac{\chi(\sum_{s=1}^{\tau} Z_s=n-j)}{(\sum_{s\in[\tau]} Y_s+\theta)/\theta} \right) f_j(\theta)$

From Theorem 5, we can therefore transform the finite horizon cost model into a single cycle cost model, by replacing the density function $f_j(\theta)$ with $\tilde{p}_j(n,\theta)$ to account for the finite horizon effect. The optimal staggered base stock policy can be obtained by solving the following problem

$$\begin{aligned} \min_{S_j\geq 0} \quad & \sum_{j\in[J]} \int_\theta \tilde{p}_j(n,\theta) \frac{\sum_{\ell=1}^{j-1} hS_\ell + h(S_j-\theta)^+ + b(\theta-S_j)^+}{\theta} d\theta \\ \text{s.t.} \quad & S_{j-1} \leq S_j, \forall j \in [J] \end{aligned} \tag{6}$$

For both discrete and continuous demand cases, the formulations provided are data-driven, in the sense that historical samples are direct input in solving for the staggered base stock policy.

4.3. Application III: One-for-One Replenishment with Positive Lead Time

In many spares inventory system with the failures generated by Poisson processes and deterministic shipment lead times (from the repair depot to each site), a common approach used to replenish the repairable item at each site following a "one-for-one replenishment" policy. In the case when inventory metrics are evaluated over a finite horizon, we show that the analysis on the gambler's fallacy effect can be extended to this setting.

To this end, we consider the unit demand arrivals in a case with the replenishment lead time $L > 0$. We use a one-for-one replenishment policy with a buffer stock of B, i.e., once the inventory position drops to or below the reorder point $B-1$, the system will place an order (of one unit) to raise the inventory level to B. In other words, the one-for-one policy with a buffer stock of B can be seen as the $(B-1, B)$ replenishment policy, a variant of the popular (s, S) inventory policy.

In such a system, each unit of the replenished inventory will be used to satisfy the demand for the B-th unit demand arrival after the ordering. That is, from the replenishment unit's perspective, it faces a unit demand with an inter-arrival time of $Q^B \sim \sum_{s=1}^{B} Z_s$, which is similar to our discussion in Section 4.1. Upon the replenishment ordering, the following $B-1$ unit demands will be satisfied with the remaining buffer stock $B-1$ and the B-th unit demand is met with this replenished unit. In this case, we will show that it is still beneficial to delay the replenishment timing of each order, to exploit the gambler's fallacy effect in the demand process for Q^B. Therefore, the new delayed ordering policy determines not only the optimal buffer stock (also base stock level) B, but also the optimal delay interval W associated with the replenishment. To facilitate the search for the optimal parameters for W and B, we make the following assumption about demand distribution.

DEFINITION 1. (Log-Concavity) *Let the sample space* $\Omega = \{z_1, z_2, z_3, \ldots\}$ *be a countable subset of the Euclidean space* $\mathbb{R}$*, with* $z_1 < z_2 < z_3 < \cdots$. *Let* $\{p_i = \mathbf{P}(Z = z_i)\}$ *be the probability mass function for* Z*. A discrete random variable* Z *is log-concave if*

$$p_{i+1}^2 \geq p_i p_{i+2}, \quad \forall i. \tag{7}$$

Examples of discrete log-concave distributions include Poisson distributions, binomial distributions, negative binomial distributions, and geometric distributions, etc.

ASSUMPTION 1. *For each* t*, the distribution of the random variable* Z_t *is log-concave.*

ASSUMPTION 2. *For each* t*, the distribution of the random variable* Z_t *is concave.*

Under Assumption 1, since the sum of log concave functions is log-concave, for every B, the distribution of random variable Q^B is also log-concave. The concavity of Q^B holds similarly by Assumption 2. While Assumption 2 implies Assumption 1 as the domain is non-negative, we separate these two assumptions to facilitate our discussion afterwards.

Let Q_t^B be the t-th total inter-arrival time from Z_1 to Z_B. Let $\tau_Q = \max\{t : Q_1^B + \cdots + Q_t^B \leq n-1\}$. Similar to the previous argument in Section 3, we define the biased estimate of $\mathbf{P}(Q^B = j)$ on account of the finite horizon as

$$p_j(n,B) := \mathbf{E}\left[\frac{\chi(Q_1^B = j) + \ldots + \chi(Q_{\tau_Q}^B = j) + \chi(\sum_{s=1}^{\tau_Q} Q_s^B = n-j)\chi(Q_{\tau_Q+1}^B = j)}{\tau_Q + 1}\right]. \tag{8}$$

While assumption 1 ensures that $\mathbf{P}(Q^B = j)$ is log-concave, we need the stronger assumption 2 to ensure that the biased version $p_j(n,B)$ is log-concave.

Let $W^*(n,B)$ denote the optimal delay ordering policy with B, the buffer stock needed. The problem here is to jointly determine the optimal choice of B and $W^*(n,B)$ to minimize the expected cost per cycle given the finite horizon n and the lead time L. Note that, here, a cycle is the time interval between the ordering of a replenishment unit and the moment it is used to satisfy the demand. If W is the delay ordering period and B is the buffer stock, by Proposition 3, we can write the expected cost per cycle $\tilde{\mathcal{C}}(W,B)$ associated with a replenishment unit as

$$\tilde{\mathcal{C}}(W,B) := b \sum_{j \leq W+L} p_j(n,B) + h \sum_{j=W+L+1}^{n} (j - W - 1 - L) p_j(n,B). \tag{9}$$

For expositional clarity, we denote

$$f(W,B) \equiv \tilde{\mathcal{C}}(W+1,B) - \tilde{\mathcal{C}}(W,B) = b p_{W+L+1}(n,B) - h \sum_{j=W+L+2}^{n} p_j(n,B),$$

and

$$g(W,B) \equiv \frac{p_{W+L+1}(n,B)}{\sum_{j=W+L+2}^{n} p_j(n,B)} = \frac{p_{W+L+1}(n,B)}{\sum_{j=W+L+1}^{n} p_j(n,B) - p_{W+L+1}(n,B)}.$$

Notably, the objective function cost per cycle (9) is affected by the biased distribution of Q^B.

LEMMA 1. *Under Assumption 2, the biased distribution of Q^B defined in* (8) *is also log-concave.*

This allows us to streamline the search for the optimal delay with given buffer stock B in the following way:

PROPOSITION 4. *Under Assumption 2, for each given B, (a) if there exists a finite W such that $g(W-1,B) \le h/b,\ g(W,B) > h/b$, then $\tilde{C}(W,B)$ is minimized at $W^*(n,B) = \max\{W : g(W-1,B) \le h/b,\ g(W,B) > h/b\}$;*

(b) if $g(W,B) \le h/b$ for all W, then $\tilde{C}(W,B)$ is minimized at $W^*(n,B) = \infty$;

(c) if $g(W,B) \ge h/b$ for all W, the $\tilde{C}(W,B)$ is minimized at $W^*(n,B) = 0$.

Given this relationship between the optimal delay $W^*(n,B)$ and B, the optimal buffer stock B can be obtained with a simple line search procedure.

Proposition 4 implies that when $g(W,B)$ is less than a threshold h/b for all W, then the optimal policy is to delay forever; when the probability ratio is greater than a threshold h/b for all W, then the optimal policy is to avoid delaying; otherwise, the optimal delay ordering period is a positive integer. Note that $g(W,B)$ is simply the probability ratio of the B-th unit demand arriving at period $W+L+1$ and arriving after period $W+L+2$. We also notice that the optimal delay $W^*(n,B)$ is a function of B. The following proposition identifies a monotonicity property of $W^*(n,B)$.

PROPOSITION 5. *The optimal delay $W^*(n,B)$ is non-decreasing in B.*

We note that assumption 2 is more restrictive compared to assumption 1. If it is not satisfied, the approach described in Proposition 4 can be used as a heuristic algorithm to find the optimal delay and buffer stock. The following numerical examples shows that even if this assumption does not hold, delaying the replenishment according to the approach provided by Proposition 4 can still help to lower the cost compared to immediate replenishment.

4.3.1. Numerical Examples Consider an example where we let $n=50$, $h=1$, and $b=1.5$, with different lead time $L=0,4,8,12$. We assume the distribution of Z_t is Poisson with arrival rate $\lambda = 1/4$. The optimal buffer stock and optimal delay can be easily found through search along $\tilde{C}(W,B)$. Figures 4a and 4b show how they change as the lead time increases from 0 to 12. In general, the optimal[12] buffer stock increases with lead time, with $B^*=1$ up till lead time of 7,

[12] Note that under the current setting, Assumption 2 is not satisfied. The optimal delay and optimal buffer stock obtained through Proposition 4 are near optimal in this example. The term "optimal" in the remaining of this section refers to the optimal solution by Proposition 4.

and then increases to 2 when lead time increases further. Interestingly, the optimal delay interval decreases when the lead time increases from 0 to 3, ensuring that a replenishment order is always received exactly 3 time periods after the preceding demand epoch. The optimal solution reverts to the standard one-for-one replenishment system with $B = 1$, when lead time is between 4 to 7, after which the base stock increases to 2, and the optimal delay is such that the replenishment will be received exactly 10 time periods after the preceding demand epoch, until lead time increases to 10 or more.

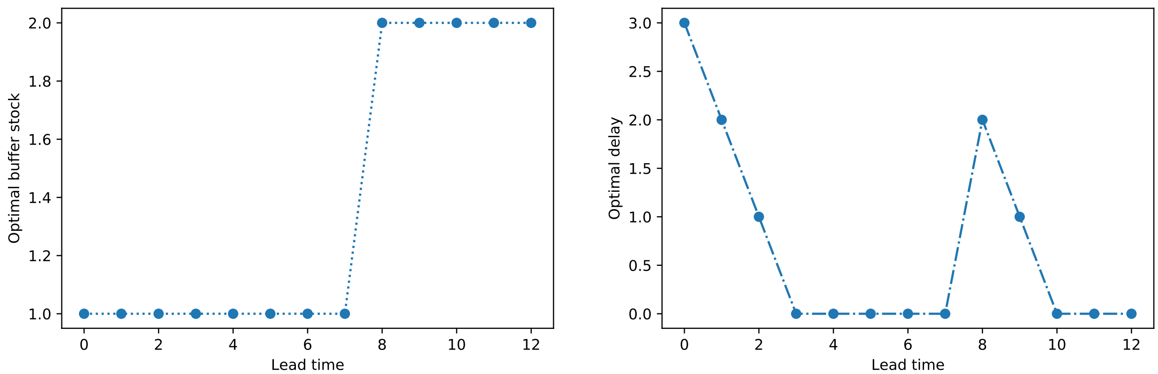

(a) Optimal buffer stock (b) Optimal delay

Figure 4 **Optimal buffer stock and optimal delay with different lead time**

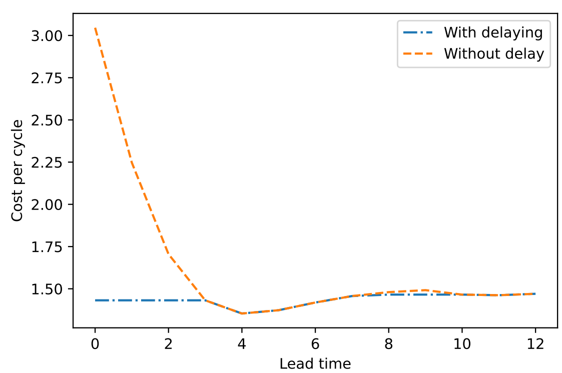

Figure 5 **Cost comparison between the replenishment policy with and without delaying**

Figure 5 shows that the cost per cycle for one-for-one replenishment without delay smallest when the lead time is 4, which is the mean inter-arrival time for a unit of demand. However, due to the finite horizon effects, delay ordering can play a significant role when the lead time is smaller than mean inter-arrival time, but with minimal impact when lead time is longer. This example demonstrates that the gambler's fallacy effect can also be leveraged, even if the replenishment lead time is positive, and suitably delaying the replenishment timing can help to lower the cost of a one-for-one replenishment policy for this class of inventory problems.

5. Staggered Base Stock Policy: Computational Results

To empirically validate our methods, we test the performance using a heavy machinery parts industry dataset, where the demand is intermittent. The company, which operates as a distributor of heavy machinery parts and diesel engine components, provided monthly sales data for $M = 31$ products over a 38-month time period. The selected products those that had at least one recorded demand during the horizon. The time series of demand arrivals for these products are shown in Figure 6, where the x-axis represents the time periods, the y-axis represents the individual products, and the bubble size represents the demand size, if any, in a given month. There is no demand arrival in nearly 77.8% of the months (i.e., time periods) with an average inter-arrival time of 4.08 months. The demand occurrence probabilities in each month for all the products are less than 0.6. The mean and variance of the positive (monthly) demand size are 5.4 and 33.1, respectively. Moreover, Figure 6 indicates that the demand arrival patterns vary across products.

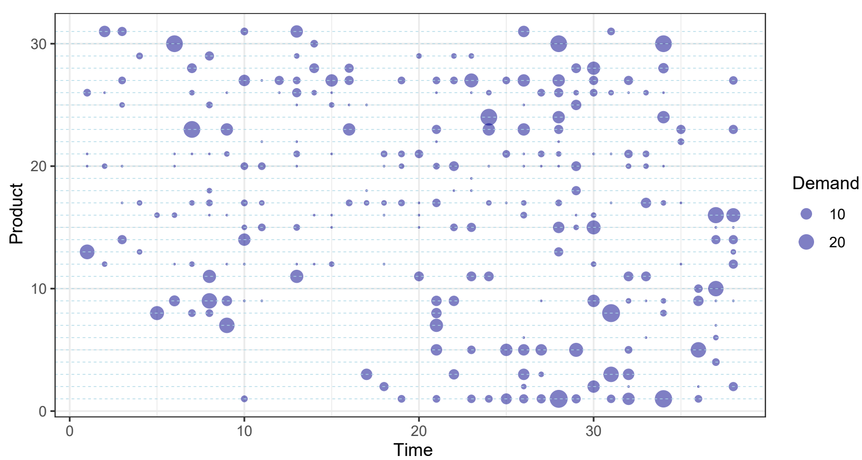

Figure 6 **Demand arrivals and demand size**

5.1. Demand Modeling and Data Generation

This rather small data set for each product prohibits a reliable performance evaluation using standard validation techniques such as training and test split. Moreover, using real data (i.e., one sample path) to test the performance of any algorithm will not be convincing, and the outcomes may not be reproducible. In order to overcome this challenge, we first construct demand distributions based on the real data and then regenerate sample paths from these demand distributions for computation and performance evaluation. As such, while the actual data did not serve as direct input to the inventory control and numerical experiments, the findings presented in the subsequent results remain largely grounded in real demand patterns.

Specifically, we set the Gaussian kernel smoothing of the empirical distribution (from the real data) as the ground truth distribution where both the training and test data sets are sampled. The inter-arrival-demand distributions for two products are presented in Figure 7. For each product, we generate a single sample path with a length of 500 periods from the product's inter-arrival-demand distribution as the training data. We then generate from the same distribution $N_{\text{test}} = 100$ sample paths with $n = 38$ periods of observations as the test data. It is worth noting that the chosen length of $n = 38$ is the length of the finite horizon in the considered setting.

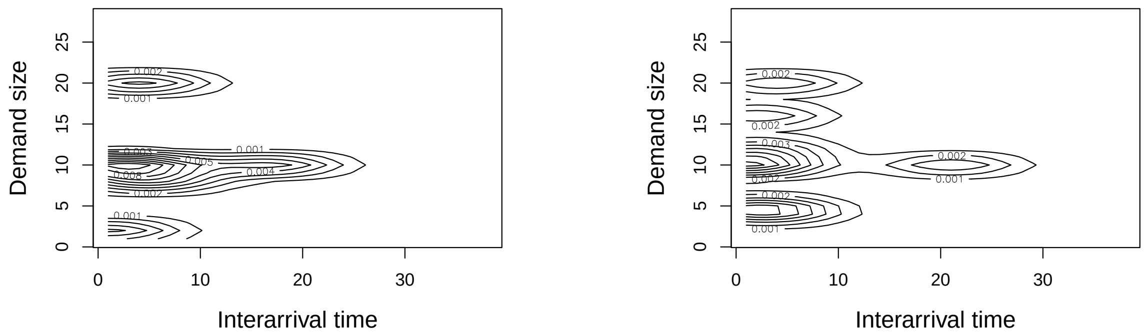

Figure 7 **Two examples of inter-arrival-demand distributions.**

In order to derive inventory control policies, we need to estimate $\tilde{p}_{kj}(n)$, the probability of having a nonzero demand of size k and inter-arrival time j in a finite horizon with n periods. To this end, we transform the sample path with 500 periods of training data, into a two-dimensional histogram by partitioning it into non-overlapping intervals with $n = 38$ periods to account for the finite horizon effect. This partitioning process is illustrated in Figure 8 where a sample path of 20 periods can be partitioned into 3 non-overlapping intervals with a horizon of $n = 5$ periods. Note that each interval begins after a demand arrival such that we can regard each interval as a sample run of our inventory control setting with the n-period horizon. Then $\hat{p}_{kj}(n)$, the estimation of $\tilde{p}_{kj}(n)$, are obtained from the realizations in these n-period sample intervals and then used as inputs for the optimization for the inventory control policies. For simplicity, we use the estimate $p^c_{kj}(n)$ defined in Equation (5) for the calculation in part because the demand size does not vary much in our dataset.

5.2. Implementation and Benchmarks

To obtain the SBS policy, we solve the discrete demand formulation provided in equation (4), because of the relatively small demand range observed in the data. There are also two related benchmarks used for comparison, i.e., the ITE policy (Akcay et al., 2015) and the approach by Croston(Croston, 1972), developed specifically for inventory problems with intermittent demand. The ITE policy considers a single-period inventory problem with an adjusted newsvendor quantity

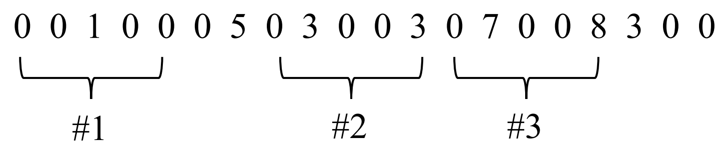

Figure 8 **This sample path with** 20 **periods of training data is partitioned into** 3 **non-overlapping intervals with** $n = 5$ **periods.**

hedging against both demand uncertainty and parameter uncertainty. To account for the estimation error of the parameters, Akcay et al. (2015) quantify the expected cost of incorrectly estimating the parameters of an intermittent demand process as a function of critical fractile. On the other hand, the Croston policy uses exponential smoothing to produce separate point estimates for nonzero demand size and inter-arrival time. The rules for replenishment of Croston policy are based upon a linear combination of the estimates of average demand, and the mean absolute deviation of one period ahead forecasting. The details for the construction of these policies are provided in Appendix EC.3.

We evaluate the CPUGS under different inventory control policies using the test data. In our calculation of CPUGS, we also account for the carryover inventory between successive cycles to reflect the actual inventory cost over the planning horizon in practice. This setting is not exactly identical to our theoretical model in Section 4.2 where we assume the leftover inventory in each cycle can be disposed of at zero cost. Nevertheless, we will demonstrate that the SBS policy still outperforms both benchmarks that do not consider the finite horizon effect, despite the slightly unrealistic assumption. Note that both the ITE and Croston policy remain static during each demand cycle. Their order-up-to levels will only be updated when a new demand arrives. On the other hand, the SBS policy will change the order-up-to level according to the elapsed time, even within a demand cycle. In a sense, the SBS policy is more dynamic and can adapt to new information in the operating environment, even when there is no demand arrival in the ensuing period. Finally, throughout the experiments, we fix the holding cost at $h = 1$, and vary the shortage cost $b \in \{6, 14, 22, 30\}$ across four scenarios. Additional numerical results for more values of shortage cost are available in Appendix EC.5.

5.3. Average performance

For each product, we evaluate the average CPUGS based on $N_{\text{test}} = 100$ sample paths for each inventory control policy respectively and present the results as a boxplot in Figure 9. The plot reveals that the SBS policy has the smallest mean compared to other benchmarks. We also notice that as the unit shortage cost b becomes larger, the gap between SBS and Croston reduces, but the gap between SBS and ITE widens. That is because, when b is small, both SBS and ITE do not hold excessive inventory to prevent high inventory costs, resulting in similar performances. As b increases, both SBS and ITE policy will increase the base stock levels. Unlike the constant base stock level of ITE, the SBS policy gradually increases the inventory level, saving inventory holding cost at the beginning and still achieving a lower CPUGS. Regarding the Croston policy, because the cost parameters do not directly affect its base stock levels as well as due to the lags and stability of exponential smoothing, the average CPUGS changes the least among all the policies. Among all the cases with different shortage costs, the SBS policy achieves smaller CPUGS than the Croston policy in 85.5% of products, and smaller than the ITE policy in 94.4% of products.

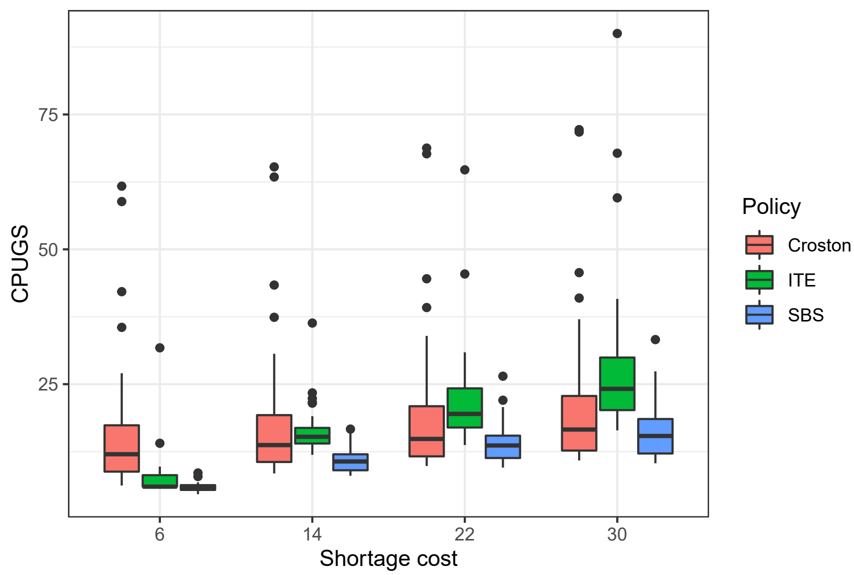

Figure 9 Average CPUGS of the 31 products.

Furthermore, we investigate the total inventory levels of all 31 products combined over 38 periods in Figure 10. Overall, the ITE and Croston policy tend to hold more inventory than the SBS policy. Notably, the total inventory levels of Croston remain the same when the shortage cost changes. This phenomenon coincides with the Croston policy's inability to adapt to cost parameters. While more inventory will be held under higher shortage costs for both SBS and ITE, it is evident that ITE is more sensitive to the increase in shortage costs. In the SBS policy, the inventory levels of the first few periods are significantly smaller than those of the later periods. It ensures the avoidance of unnecessary inventory in the beginning but gradually raises the inventory level. Again, this observation explains the cost saving advantage of the SBS policy indicated in Figure 9 above. Moreover, keeping a low inventory level and eliminating any items deemed unnecessary will not only save warehouse space but will also improve capital utilization.

In addition to the average CPUGS, we also extend the performance metric to the total cost, which is prevalent in the inventory management literature. Figure 11 demonstrates that the SBS policy

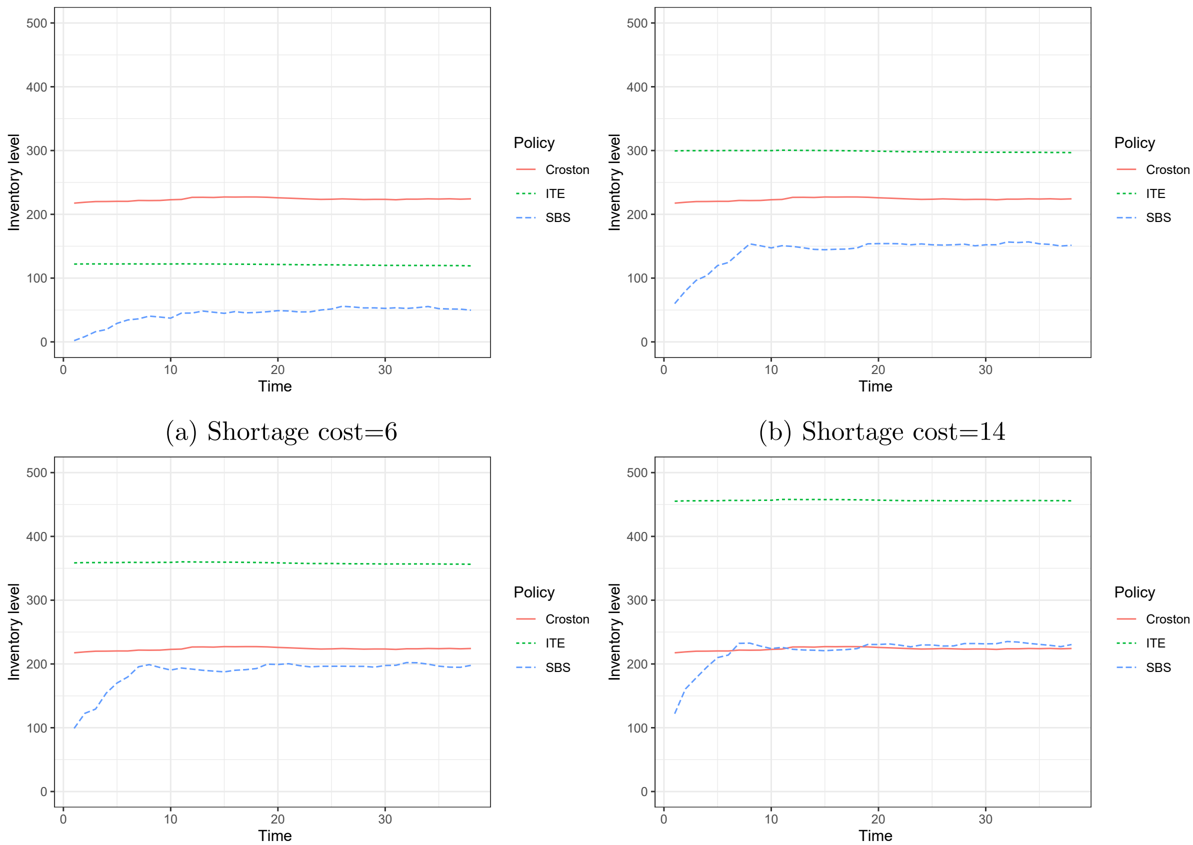

(a) Shortage cost=6

(b) Shortage cost=14

(c) Shortage cost=22

(d) Shortage cost=30

Figure 10 Total inventory levels of 31 products over time under different shortage costs

achieves lower total costs than both the Croston and ITE policy. The reason for such performance is similar to the above discussion that the SBS policy prevents unnecessary inventory and shortages by gradually raising the inventory levels.

Note that, in the implementation, we use the empirical estimate $\hat{p}_{kj}$ from the single training sample path in place of the bias estimate $\tilde{p}_{kj}$. Therefore, we are also interested in how this estimation error (i.e., the difference between $\hat{p}_{kj}$ and $\tilde{p}_{kj}$) would result in deviation from the optimal inventory policy. To quantify the optimality gap, we compare the performance of the SBS policy using $\hat{p}_{kj}$ and $\tilde{p}_{kj}$, respectively. The result in Figure 12 shows that using the estimation will not deviate from the optimal inventory policy too much, suggesting that the gambler's fallacy is significant even when the demand distribution is estimated based on data. Furthermore, as the number of data points increases, the optimality gap caused by the empirical estimation error would vanish.

5.4. The performance on a randomly selected sample path

The results above for each product are calculated based on the average performance of the $N_{\text{test}} = 100$ sample paths. In this experiment, we extract two randomly selected sample paths as examples, showing the CPUGS for all 31 products in the boxplots in Figure 13. The performance comparisons

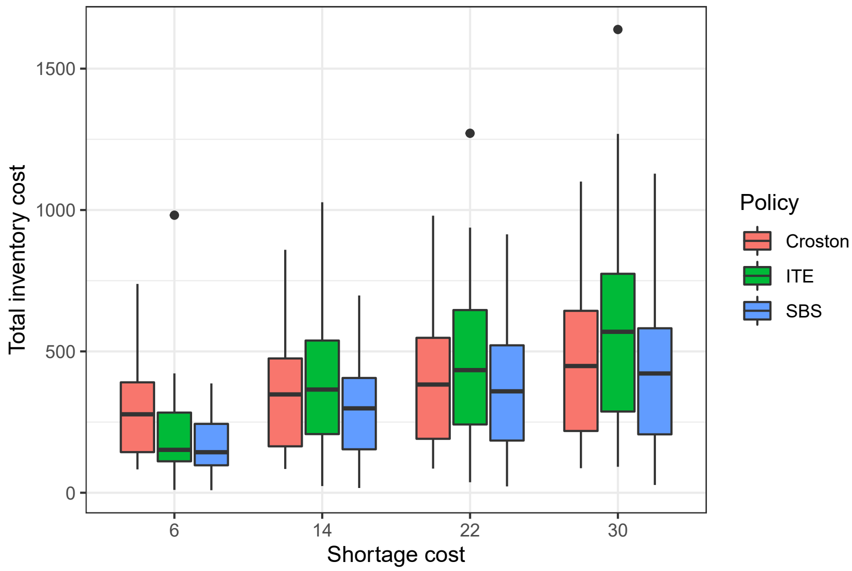

Figure 11 **Total costs of the 31 products**

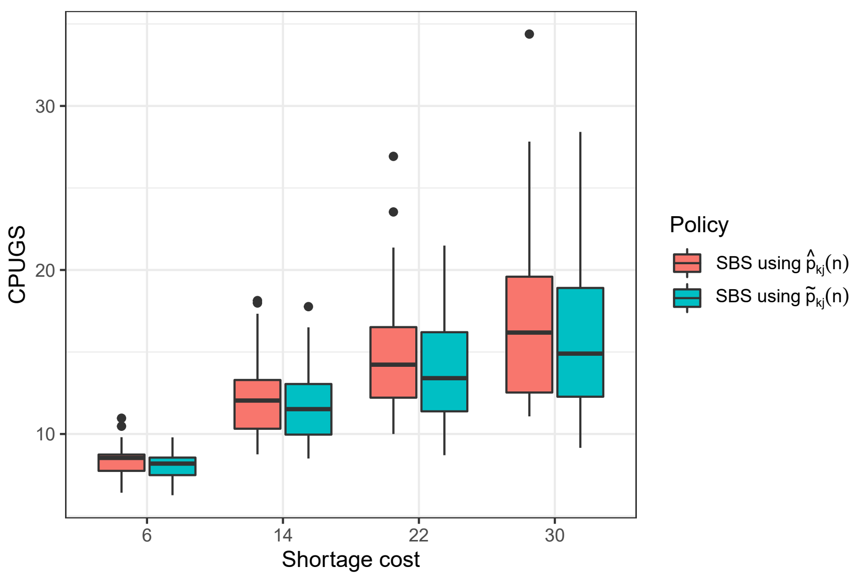

Figure 12 **Optimality gap between the true $\tilde{p}_{kj}$ and the empirical one $\hat{p}_{kj}$ for $n = 38$ when calibrated from the single training sample path with $n_{\text{training}} = 500$ data points.**

in the two sample paths are consistent with the observation of the average performance. Therefore, the SBS policy not only performs well on average but can also outperform the benchmarks in individual sample paths. If we compare the performance on each sample path in terms of the total CPUGS of all 31 products, then the SBS policy outperforms the Croston policy and ITE policy in 66.1% and 83.1% of the test sample paths, respectively.

So far, we have evaluated the inventory policies using abundant test data generated from the calibrated inter-arrival-demand distributions. Note that the original data contain one sample path

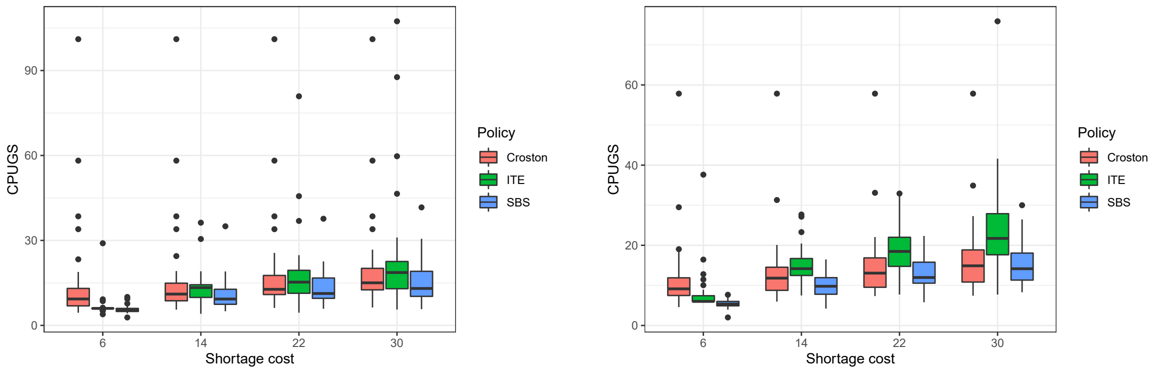

(a) Sample path #1. (b) Sample path #2.

Figure 13 CPUGS of the 31 products evaluated in a single sample path.

of $n = 38$ periods for each product. To further explore the performance, we evaluate CPUGS by testing each policy under the original data and show the comparison in Figure 14. Similar to the average performance, the SBS policy outperforms the Croston and ITE policy in all the cases with different shortage costs. This establishes additional evidence to support the efficacy of the SBS policy.

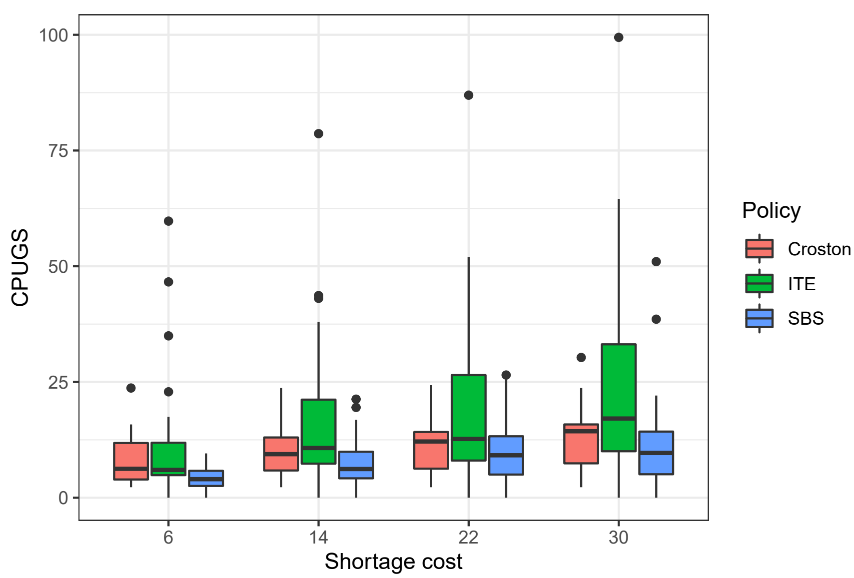

Figure 14 CPUGS tested on the original data for the 31 products.

6. Concluding Remarks

It is common to formulate models based on some cost performance metrics to obtain a recommended solution to a problem. However, the fact that such performance metrics are calculated based on data collected over a finite horizon is often disregarded. This induces a truncation effect on our observations, introducing biases into calculations. The recommended solution to the problem needs to anticipate such issues, and correct for the biases. This is the basis of the emergence of the gambler's fallacy in a finite sequence of coin tosses.

This paper investigates the effects of the finite horizon and the gambler's fallacy on inventory models to obtain the optimal SBS policy. Our performance metric is based on fill rate, cost per cycle, and the average cost per unit sold or cost per cycle measurements. These challenges are exceedingly difficult to address for products in the long tail, because of infrequent demand arrivals, and the difficulty of obtaining the true demand distribution. Instead, we use a biased version of the joint distribution between demand and inter-arrival time, as the input of our optimization model, demonstrating how this addresses the finite horizon effect in the traditional inventory control approach for intermittent demand products. The SBS policy we propose is supported by the gambler's fallacy phenomenon, which analyzes the impact of statistical biases and provides justification for both the use of replenishment-delay policies and biased estimation based on empirical distribution, obtained from finite-horizon data. We implement the SBS policy and other benchmarks on real data, finding that the SBS policy outperforms the benchmarks in all of our experiments.

There are several managerial insights that we can draw from this work:

- *Do not ignore the finite horizon effect*, especially if the performance metric is measured from data collected over a finite interval. This paper analyzed the case in which the performance metric is measured by the fill rate, cost per unit sold, and cost per cycle. These metrics are routinely reported by companies based on data collected over a fixed interval. Other inventory metrics that need to account for the finite horizon effect include: inventory turnover ratio, gross margin percent, return on investment, etc.

- *Biased estimation of the demand probability distribution can perform better than true distribution*, when used as inputs in the performance metric optimization model. This brings a new twist to the relentless efforts of companies to estimate the true demand distribution—based on the performance metric used, it may be more advantageous for companies to use the biased demand distribution instead!

- In the management of intermittent demand products, *delayed ordering and staggered base stock policy generally perform better than constant base stock policy with immediate replenishment*, when the performance metric is measured over a finite horizon. This holds even in cases in which demands are identical and independent across time.

Last but not least, there are several future research directions that are interesting yet challenging. A natural extension is to manage inventory with a continuous time dimension, compared to the discrete-time setting in this paper. Calculating the effect of the finite horizon in this setting appears to be difficult. Another possible research direction is to consider cases in which demand is not stationary, so that the joint distribution may change accordingly over time. In such instance, we would need to use a feature-based demand model, and the challenge lies in uncovering the underlying base demand model while accounting for the effects due to the features. It is unclear how the finite horizon effect would skew the base demand model and the associated effects due to the features. We can also investigate the inventory management problem for multiple slow-moving items. The presence of multiple items brings opportunities to consider risk-pooling strategy and data aggregation methods to mitigate stock-outs and excess supply, and explore the demand correlation between products.

References

Akcay, Alp, Bahar Biller, Sridhar R Tayur. 2015. Managing inventory with limited history of intermittent demand. *Available at SSRN 2710282* .

An, Mark Yuying. 1997. Log-concave probability distributions: Theory and statistical testing. *Duke University Dept of Economics Working Paper* (95-03).

Chen, Daniel L, Tobias J Moskowitz, Kelly Shue. 2016. Decision making under the gambler's fallacy: Evidence from asylum judges, loan officers, and baseball umpires. *The Quarterly Journal of Economics* **131**(3) 1181–1242.

Chen, Jiahua, Dennis KJ Lin, Douglas J Thomas. 2003. On the single item fill rate for a finite horizon. *Operations Research Letters* **31**(2) 119–123.

Clotfelter, Charles T, Philip J Cook. 1993. The gambler's fallacy in lottery play. *Management Science* **39**(12) 1521–1525.

Cornacchia, Eugenio, Joseph Shamir. 2018. Forecasting the long tail and intermittent demand. URL `https://www.toolsgroup.com/resources/white-papers/forecasting-long-tail/`.

Croson, Rachel, James Sundali. 2005. The gambler's fallacy and the hot hand: Empirical data from casinos. *Journal of risk and uncertainty* **30**(3) 195–209.

Croston, John D. 1972. Forecasting and stock control for intermittent demands. *Journal of the Operational Research Society* **23**(3) 289–303.

Eftekhar, Mahyar, Jing-Sheng Jeannette Song, Scott Webster. 2022. Prepositioning and local purchasing for emergency operations under budget, demand, and supply uncertainty. *Manufacturing & Service Operations Management* **24**(1) 315–332.

Hançerlioğullan, Gülşah, Emrah Köksalmiş, Fethi Çalişir. 2017. How does inventory performance affect operational performance? an example from us space vehicles and aircraft parts industries. *2017 8th International Conference on Recent Advances in Space Technologies (RAST)*. IEEE, 291–294.

Kong, Qingxia, Georg D. Granic, Nicolas S. Lambert, Chung Piaw Teo. 2020. Judgment error in lottery play: When the hot hand meets the gambler's fallacy. *Management Science* **66**(2) 844–862.

Miller, Joshua B., Adam Sanjuro. 2018. Surprised by the hot hand fallacy? a truth in the law of small numbers. *Econometrica* **86**(6) 2019–2047.

Rabin, Matthew, Dimitri Vayanos. 2010. The gambler's and hot-hand fallacies: Theory and applications. *The Review of Economic Studies* **77**(2) 730–778.

Rinott, Yosef, Maya Bar-Hillel. 2015. Comments on a 'hot hand' paper by miller and sanjurjo (2015). *Available at SSRN 2642450* .

Rogers, Mary Margaret, John Schatzberg, Laurie Schatzberg. 2010. The relationship between inventory management and firm profitability: Sector consequences of catastrophic supply chain disruptions. *California Journal of Operations Management, 8 (2): 39* **54**.

Saumard, Adrien, Jon A Wellner. 2014. Log-concavity and strong log-concavity: a review. *Statistics surveys* **8** 45.

Schultz, Carl R. 1989. Replenishment delays for expensive slow-moving items. *Management Science* **35**(12) 1454–1462.

Song, Jing-Sheng, Paul Zipkin. 1993. Inventory control in a fluctuating demand environment. *Operations Research* **41**(2) 351–370.

Syntetos, Aris A, M Zied Babai, Yves Dallery, R Teunter. 2009. Periodic control of intermittent demand items: theory and empirical analysis. *Journal of the Operational Research Society* **60**(5) 611–618.

Syntetos, Aris A, John E Boylan. 2001. On the bias of intermittent demand estimates. *International journal of production economics* **71**(1-3) 457–466.

Syntetos, Aris A, John E Boylan. 2005. The accuracy of intermittent demand estimates. *International Journal of forecasting* **21**(2) 303–314.

Terrell, Dek. 1994. A test of the gambler's fallacy: Evidence from pari-mutuel games. *Journal of risk and uncertainty* **8**(3) 309–317.

Thomas, Douglas J. 2005. Measuring item fill-rate performance in a finite horizon. *Manufacturing & Service Operations Management* **7**(1) 74–80.

Tripathi, R, S Mishra. 2014. Inventory model with inventory-dependent demand for deteriorating items in a single warehouse system. *Uncertain Supply Chain Management* **2**(4) 209–218.

Appendix – Taming the Long Tail: The Gambler's Fallacy in Intermittent Demand Management

EC.1. List of Notations

Table EC.1 List of notations

Notation	Definition
h	Unit holding cost.
b	Unit shortage cost.
$\hat{p}_j(n)$	The empirical proportion of transition $\{1, 0_1, \ldots, 0_{j-1}, 1\}$ in n-sample 0-1 sequence.
$p_1(n) := \mathbf{E}[\hat{p}_1(n)]$	The expected empirical proportion of transition $\{1, 1\}$ in n-sample 0-1 sequence.
$p_j(n) := \mathbf{E}[\hat{p}_j(n)]$	The expected empirical proportion of transition $\{1, 0_1, \ldots, 0_{j-1}, 1\}$ in n-sample 0-1 sequence.
$B_n(t)$	The indicator of the event that "1" is generated at the t-th position.
$G_n(t)$	The event that there is a transition from "1" at position t to "1" at position $t+1$.
$\mathcal{E}_{k\ell}(t)$	The event that at least k demand arrivals in the next ℓ periods, following a demand arrival at t.
W	The replenishment is delayed for W periods after a unit demand has arrived in the delay ordering policy.
Z_j	The random inter-arrival time of j-th nonzero demand.
Y_j	The random demand size of j-th nonzero demand.

EC.2. Proofs of Statements

EC.2.1. Proof of Theorem 2

Let $m = n-1$. With probability $\binom{m}{i} p^i q^{m-i}$, there are i nonzero demands arriving in m periods.

Choose one of the i unit demands randomly. We consider the following three events:

- Event $\mathcal{A}_{ij}$ is a unit demand falls in the $(n-j)$th position and it will be followed by $j-1$ successive 0's and a 1 in the nth position, illustrated in the following figure.

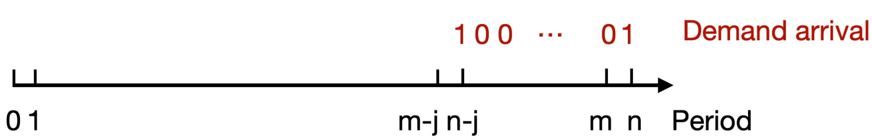

Event $\mathcal{A}_{ij}$ happens with probability $\dfrac{1}{m} \times p \times \binom{m-j}{i-1} / \binom{m-1}{i-1}$, i.e.,

$$\mathbf{P}(\mathcal{A}_{ij}) = \begin{cases} \dfrac{p}{m} & j = 1 \\ \dfrac{p}{m} \times \dfrac{(m-i)(m-i-1)\cdots(m-i-j+2)}{(m-1)\cdots(m-j+1)} & j \geq 2 \end{cases}$$

- Event $\mathcal{B}_{ij}$ denotes the case when this unit falls in one of the first $m-j$ positions and it will be followed by $j-1$ successive 0's and a 1.

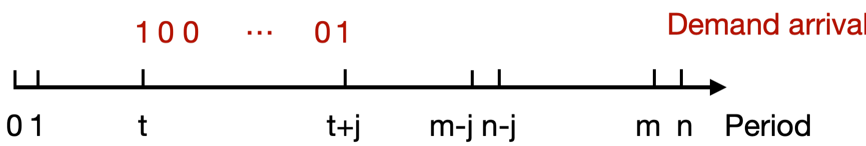

Event $\mathcal{B}_{ij}$ happens with probability $\frac{m-j}{m} \times \binom{m-j-1}{i-2} / \binom{m-1}{i-1}$, i.e.,

$$\mathbf{P}(\mathcal{B}_{ij}) = \begin{cases} \frac{m-1}{m} \times \frac{i-1}{m-1} & j=1 \\ \frac{m-j}{m} \times \frac{(m-i)\cdots(m-i-j+2)(i-1)}{(m-1)(m-2)\cdots(m-j)} & j \geq 2 \end{cases}$$

- Event $\mathcal{C}_{ij}$ denotes the case when the unit arrives at period 0 and $Z_1 = j$.

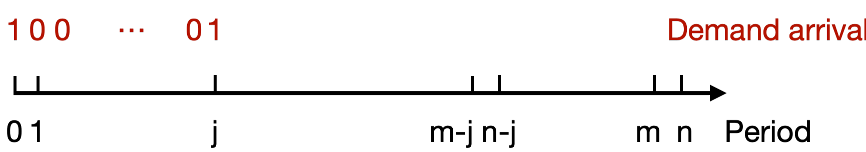

Event $\mathcal{C}_{ij}$ happens with probability $\binom{m-j}{i-1} / \binom{m}{i}$, i.e.,

$$\mathbf{P}(\mathcal{C}_{ij}) = \begin{cases} \frac{i}{m} & j=1 \\ \frac{(m-i)\cdots(m-i-j+2)i}{m(m-1)\cdots(m-j+1)} & j \geq 2 \end{cases}$$

Note that when calculating $p_j(n)$, according to the pigeonhole principle, we have $i \leq m-j+1$; otherwise, there is no probability of a 1 following $j-1$ successive 0's, i.e., the events $\mathcal{A}_{ij}$, $\mathcal{B}_{ij}$ and $\mathcal{C}_{ij}$ would not happen.

With the initial nonzero demand, there are i+1 nonzero demand periods in total. Among these periods, event $\mathcal{A}_{ij}$ and $\mathcal{B}_{ij}$ can be chosen from any from 2nd to $(i+1)$-th nonzero demand periods, and event $\mathcal{C}_{ij}$ can only choose the initial one. Therefore, we have

$$\begin{aligned}
p_j(n) &= \sum_{i=1}^{m-j+1} \binom{m}{i} p^i q^{m-i} \left(\frac{i}{i+1} \mathbf{P}(\mathcal{A}_{ij}) + \frac{i}{i+1} \mathbf{P}(\mathcal{B}_{ij}) + \frac{1}{i+1} \mathbf{P}(\mathcal{C}_{ij}) \right) \\
&= \sum_{i=1}^{m-j+1} \binom{m}{i} p^i q^{m-i} \left(\frac{i}{i+1} \frac{1}{m} \times p \times \frac{\binom{m-j}{i-1}}{\binom{m-1}{i-1}} + \frac{i}{i+1} \frac{m-j}{m} \times \frac{\binom{m-j-1}{i-2}}{\binom{m-1}{i-1}} + \frac{1}{i+1} \frac{\binom{m-j}{i-1}}{\binom{m}{i}} \right) \\
&= \sum_{i=1}^{m-j+1} \binom{m}{i} p^i q^{m-i} \left[\frac{i(m-j)!\,(m-i)!\,(p+i)}{(i+1)(m-j-i+1)!\,m!} \right] \\
&= \sum_{i=1}^{m-j+1} p^i q^{m-i} \left[\frac{(m-j)!}{(i-1)!\,(m-j-i+1)!} \frac{p+i}{i+1} \right] \\
&= \sum_{i=1}^{m-j+1} \binom{m-j}{i-1} p^i q^{m-i} \frac{p+i}{i+1}
\end{aligned}$$

Q.E.D.

EC.2.2. Proof of Proposition 2

(i) According to the form of $p_1(n)$ given in Theorem 2, we have

$$\frac{p_1(n)}{p} = \sum_{i=1}^{m} \binom{m-1}{i-1} p^{i-1} q^{m-i} \frac{p+i}{i+1}.$$

Let $f(p) = \frac{p_1(n)}{p}$, the first order derivatives is

$$\begin{aligned}
f'(p) &= \sum_{i=1}^{m} \binom{m-1}{i-1} p^{i-2} (1-p)^{m-i-1} \frac{-mp^2 + (2i-mi)p + i^2 - i}{i+1} \\
&= \sum_{i=1}^{m} \binom{m-1}{i-1} p^{i-2} (1-p)^{m-i-1} \frac{p(i+1) + (i-1)(i+1) - mp(p+i) - q(i-1)}{i+1} \\
&= p^{-1} q^{-1} \left(p + (m-1)p + \sum_{i=1}^{m} \binom{m-1}{i-1} p^{i-1} (1-p)^{m-i} \frac{-mp(p+i) - q(i-1)}{i+1} \right) \\
&= p^{-1} q^{-1} \left(mp \sum_{i=1}^{m} \binom{m-1}{i-1} p^{i-1} (1-p)^{m-i} + \sum_{i=1}^{m} \binom{m-1}{i-1} p^{i-1} (1-p)^{m-i} \frac{-mp(p+i) - q(i-1)}{i+1} \right) \\
&= p^{-1} q^{-1} \left(\sum_{i=1}^{m} \binom{m-1}{i-1} p^{i-1} (1-p)^{m-i} \frac{mp(i+1) - mp(p+i) - q(i-1)}{i+1} \right) \\
&= p^{-1} \left(\sum_{i=1}^{m} \binom{m-1}{i-1} p^{i-1} (1-p)^{m-i} \frac{mp - (i-1)}{i+1} \right) \\
&= p^{-1} \left(\sum_{i=1}^{m} \binom{m-1}{i-1} p^{i-1} (1-p)^{m-i} \frac{mp+2}{i+1} - 1 \right)
\end{aligned}$$

Here, the third, fourth, and last equality are due to $\sum_{i=1}^{m} \binom{m-1}{i-1} p^{i-1} (1-p)^{m-i} = 1$ and $\sum_{i=1}^{m} \binom{m-1}{i-1} p^{i-1} (1-p)^{m-i} (i-1) = (m-1)p$ from Binomial distribution $B(m-1,p)$. With slight abuse of notation, we denote $\rho_i = \binom{m-1}{i-1}p^{i-1}(1-p)^{m-i}$ as the probability in the Binomial distribution $B(m-1,p)$ with $\sum_{i=1}^m \rho_i = 1$ and $\sum_{i=1}^m (i-1)\rho_i = (m-1)p$. Then we show $f'(p) > 0$ by $\sum_{i=1}^m \binom{m-1}{i-1}p^{i-1}(1-p)^{m-i}\frac{mp+2}{i+1} = (mp+2)\sum_{i=1}^m \rho_i \frac{1}{i+1} > 1$ as follows:

$$\begin{aligned}
(mp+2)\sum_{i=1}^m \rho_i \frac{1}{i+1} &> [(m-1)p+2]\sum_{i=1}^m \rho_i \frac{1}{i+1} \\
&= \left(\sum_{i=1}^m \rho_i (i+1)\right)\left(\sum_{i=1}^m \rho_i \frac{1}{i+1}\right) \\
&= \sum_{i=1}^m \sum_{j=1}^m \rho_i \rho_j \frac{i+1}{j+1} \\
&= \frac{1}{2}\sum_{i=1}^m \sum_{j=1}^m \rho_i \rho_j \left(\frac{i+1}{j+1} + \frac{j+1}{i+1}\right) \\
&\geq \sum_{i=1}^m \sum_{j=1}^m \rho_i \rho_j = \left(\sum_{i=1}^m \rho_i\right)^2 = 1.
\end{aligned} \tag{EC.1}$$

Hence, for a given n, $f'(p) > 0$ and $\frac{p_1(n)}{p}$ is increasing in p, and $\frac{p - p_1(n)}{p}$ is decreasing in p.

(ii) According to the form of $p_1(n)$, we have

$$\begin{aligned}
p_1(n+1) - p_1(n) &= \sum_{i=1}^{m+1} \binom{m}{i-1} p^i q^{m-i+1} \frac{p+i}{i+1} - \sum_{i=1}^{m} \binom{m-1}{i-1} p^i q^{m-i} \frac{p+i}{i+1} \\
&= \sum_{i=0}^{m} \binom{m}{i} p^{i+1} q^{m-i} \frac{p+i+1}{i+2} - \sum_{i=0}^{m-1} \binom{m-1}{i} p^{i+1} q^{m-i-1} \frac{p+i+1}{i+2} \\
&> \sum_{i=0}^{m} \binom{m}{i} p^{i+1} q^{m-i} \frac{p+i+1}{i+2} \left(1 - \frac{m-i}{mq}\right) \\
&= \frac{p}{mq} \sum_{i=0}^{m} \binom{m}{i} p^{i} q^{m-i} \left(1 - \frac{q}{i+2}\right)(i - mp) \\
&= \frac{p}{mq} \sum_{i=0}^{m} \binom{m}{i} p^{i} q^{m-i} \frac{q(mp-i)}{i+2} \qquad \text{(EC.2)} \\
&= \frac{pq}{mq} \sum_{i=0}^{m} \binom{m}{i} p^{i} q^{m-i} \frac{mp-i-2+2}{i+2} \\
&= \frac{pq}{mq} \sum_{i=0}^{m} \binom{m}{i} p^{i} q^{m-i} \left(\frac{mp+2}{i+2} - 1\right) \\
&= \frac{pq}{mq} \left[\sum_{i=0}^{m} \binom{m}{i} p^{i} q^{m-i} \frac{mp+2}{i+2} - 1\right] \qquad \text{(EC.3)} \\
&\geq 0.
\end{aligned}$$

The equality (EC.2) is due to $\sum_{i=0}^m \binom{m}{i}p^i q^{m-i} i = mp$ and $\sum_{i=0}^m \binom{m}{i}p^i q^{m-i} = 1$ from the Binomial distribution $B(m,p)$. Again, the equality (EC.3) is due to $\sum_{i=0}^m \binom{m}{i}p^i q^{m-i} = 1$. For the last inequality, we use the same result from (EC.1) above such that

$$\begin{aligned}\sum_{i=0}^{m}\binom{m}{i}p^i q^{m-i}\frac{mp+2}{i+2} &= (mp+2)\sum_{i=0}^{m}\binom{m}{i}p^i q^{m-i}\frac{1}{i+2}\\ &= \left[\sum_{i=0}^{m}\binom{m}{i}p^i q^{m-i}(i+2)\right]\left[\sum_{i=0}^{m}\binom{m}{i}p^i q^{m-i}\frac{1}{i+2}\right]\\ &\geq 1.\end{aligned}$$

Thus, we have $\frac{p-p_1(n+1)}{p} < \frac{p-p_1(n)}{p}$, which indicates $\frac{p-p_1(n)}{p}$ is decreasing in n for all p. Q.E.D.

EC.2.3. Proof of Theorem 3

Let $k=1$ and $\ell=W$ in the definition of the event $\mathcal{E}_{k\ell}(t)$ in the proof of Theorem 1, i.e. having at least one demand in the next W periods, conditional on a demand arrival in the current period t. By definition, we have

$$\mathbf{E}[\hat{p}_{\mathcal{E}_{1W}}(n)] = \sum_{j=1}^{W} p_j(n), \text{ and } \mathbf{P}\left[\chi(\mathcal{E}_{1W}(0)) = 1\right] = \sum_{j=1}^{W} p_j.$$

Note that the long run average model uses the true probability p_j while the finite horizon model uses the biased estimate $p_j(n)$. In the following, we compare the cost components in both models given the same delay W. From Theorem 1, we also have the comparison between the backorder costs

$$b\sum_{j=1}^{W} p_j(n) < b\sum_{j=1}^{W} p_j, \text{ for all } W \geq 1.$$

Furthermore, the comparison between the inventory holding costs

$$h\sum_{j=W+2}^{n-1}(j-W-1)p_j(n) = h\sum_{j=W+2}^{n-1}\sum_{k\geq j} p_k(n) > h\sum_{j=W+2}^{n-1}\sum_{k\geq j} p_k = h\sum_{j=W+2}^{n-1}(j-W-1)p_j.$$

From the comparison, we see that the back order cost has more weight in the long run average model compared to the finite horizon model when the delay W is fixed. Therefore, in the trade-off between the back order cost and the inventory holding cost, the long run average model tends to choose a smaller delay W^* compared to the delay $W^*(n)$ chosen by the finite horizon model.

Q.E.D.

EC.2.4. Proof of Theorem 4

For the t-th ($t \in [\tau]$) cycle, when $Z_t = j$ and $Y_t = k$, i.e., $\chi(Z_t = j, Y_t = k) = 1$, we have

$$C_t(\pi^{sd}) = C_{kj} := h\left[(S_j - k)^+ + \sum_{\ell=1}^{j-1} S_\ell\right] + b(k - S_j)^+.$$

Similarly, when $\chi(\sum_{s=1}^{\tau} Z_s = n-j)\chi(Z_{\tau+1}=j, X_n=k)=1$, we also have $C_{tail}(\pi^{sd})=C_{kj}$.

$$\begin{aligned}
\mathbf{E}[\mathcal{C}(\pi^{sd})] =& \mathbf{E}\left[\frac{C_1(\pi^{sd})+\cdots+C_\tau(\pi^{sd})+C_{tail}(\pi^{sd})}{Y_1+\cdots+Y_\tau+\max\{1,X_n\}}\right] \\
=& \mathbf{E}\sum_{j\in[J]}\sum_{k\in[K]}\left(\sum_{t\in[\tau]}\frac{\chi(Z_t=j,Y_t=k)}{\sum_{s\in[\tau]^+,s\neq t}Y_s+k}+\frac{\chi(\sum_{s=1}^{\tau}Z_s=n-j)\chi(Z_{\tau+1}=j,X_n=k)}{\sum_{s\in[\tau]}Y_s+k}\right)C_{kj} \\
=& \sum_{j\in[J]}\sum_{k\in[K]}\frac{C_{kj}}{k}\times\mathbf{E}\left[\sum_{t\in[\tau]}\frac{\chi(Z_t=j,Y_t=k)}{(\sum_{s\in[\tau]^+,s\neq t}Y_s+k)/k}+\frac{\chi(\sum_{s=1}^{\tau}Z_s=n-j)\chi(Z_{\tau+1}=j,X_n=k)}{(\sum_{s\in[\tau]}Y_s+k)/k}\right] \\
=& \sum_{j\in[J]}\sum_{k\in[K]}\tilde{p}_{kj}(n)\frac{\sum_{\ell=1}^{j-1}hS_\ell+h(S_j-k)^++b(k-S_j)^+}{k}
\end{aligned}$$

where $\tilde{p}_{kj}(n):=\mathbf{E}\left[\sum_{t\in[\tau]}\frac{\chi(Z_t=j,Y_t=k)}{(\sum_{s\in[\tau]^+,s\neq t}Y_s+k)/k}+\frac{\chi(\sum_{s=1}^{\tau}Z_s=n-j)\chi(Z_{\tau+1}=j,X_n=k)}{(\sum_{s\in[\tau]}Y_s+k)/k}\right]$. Q.E.D.

EC.2.5. Proof of Theorem 5

For the t-th cycle, when $Z_t=j$ and $Y_t=\theta$ with conditional density function f_j, we have the realized cycle cost

$$C_t(\pi^{sd})=C_{\theta j}:=\sum_{\ell=1}^{j-1}hS_\ell+h(S_j-\theta)^++b(\theta-S_j)^+.$$

Similarly, when $\sum_{s=1}^{\tau}Z_s=n-j$ and $X_n=\theta$ with conditional density function f_j, we also have $C_{tail}(\pi^{sd})=C_{\theta j}$.

$$\begin{aligned}
\mathbf{E}[\mathcal{C}(\pi^{sd})] =& \mathbf{E}\left[\frac{C_1(\pi^{sd})+\cdots+C_\tau(\pi^{sd})+C_{tail}(\pi^{sd})}{Y_1+\cdots+Y_\tau+Y_n}\right] \\
=& \mathbf{E}\left[\sum_{j\in[J]}\int_\theta\left(\sum_{t\in[\tau]}\frac{\chi(Z_t=j)}{\sum_{s\in[\tau]^+,s\neq t}Y_s+\theta}C_{\theta j}+\frac{\chi(\sum_{s=1}^{\tau}Z_s=n-j)}{\sum_{s\in[\tau]}Y_s+\theta}C_{\theta j}\right)f_j(\theta)d\theta\right] \\
=& \mathbf{E}\left[\sum_{j\in[J]}\int_\theta\left(\sum_{t\in[\tau]}\frac{\chi(Z_t=j)}{(\sum_{s\in[\tau]^+,s\neq t}Y_s+\theta)/\theta}+\frac{\chi(\sum_{s=1}^{\tau}Z_s=n-j)}{(\sum_{s\in[\tau]}Y_s+\theta)/\theta}\right)f_j(\theta)\frac{C_{\theta j}}{\theta}d\theta\right] \\
=& \sum_{j\in[J]}\int_\theta\mathbf{E}\left(\sum_{t\in[\tau]}\frac{\chi(Z_t=j)}{(\sum_{s\in[\tau]^+,s\neq t}Y_s+\theta)/\theta}+\frac{\chi(\sum_{s=1}^{\tau}Z_s=n-j)}{(\sum_{s\in[\tau]}Y_s+\theta)/\theta}\right)f_j(\theta)\frac{C_{\theta j}}{\theta}d\theta \\
=& \sum_{j\in[J]}\int_\theta\tilde{p}_j(n,\theta)\frac{\sum_{\ell=1}^{j-1}hS_\ell+h(S_j-\theta)^++b(\theta-S_j)^+}{\theta}d\theta,
\end{aligned}$$

where $\tilde{p}_j(n,\theta):=\mathbf{E}\left(\sum_{t\in[\tau]}\frac{\chi(Z_t=j)}{(\sum_{s\in[\tau]^+,s\neq t}Y_s+\theta)/\theta}+\frac{\chi(\sum_{s=1}^{\tau}Z_s=n-j)}{(\sum_{s\in[\tau]}Y_s+\theta)/\theta}\right)f_j(\theta)$, $[\tau]^+=[\tau]\cup\{n\}$ and $Y_n=\max\{1,X_n\}$. Q.E.D.

EC.2.6. Proof of Lemma 1

The following proofs rely on the established results below.

Lemma EC.1. *(i) (Saumard and Wellner, 2014) If two random variables X and Y are log-concave, then the product of the two random variables XY is also log-concave.*

(ii) (Saumard and Wellner, 2014) If two independent random variables X and Y are log-concave, then the summation of the two random variables $X+Y$ is also log-concave.

(iii) (An, 1997) If the distribution of a random variable X is log-concave, then the hazard rate function $h_\alpha \equiv \dfrac{\mathbf{P}(X=\alpha)}{\mathbf{P}(X \geq \alpha)}$ is non-decreasing in α.

The biased estimate in Equation (8) can be rewrite as

$$\begin{aligned} p_j(n,B) &= \mathbf{E}\Big[\frac{\chi(Q_1^B=j)+\ldots+\chi(Q_{\tau_Q}^B=j)+\chi(\sum_{s=1}^{\tau_Q} Q_s^B=n-j)\chi(Q_{\tau_Q+1}^B=j)}{\tau_Q+1}\Big] \\ &= \mathbf{E}\Big[\frac{\mathbf{P}(Q_1^B=j)+\ldots+\mathbf{P}(Q_{\tau_Q}^B=j)+\mathbf{P}(\sum_{s=1}^{\tau_Q} Q_s^B=n-j)\mathbf{P}(Q_{\tau_Q+1}^B=j)}{\tau_Q+1}\Big] \end{aligned} \tag{EC.4}$$

By the log-concavity of inter-arrival time $Z_1,\ldots,Z_B$, their sums $Q_1^B,\ldots,Q_{\tau_Q}^B,Q_{\tau_Q+1}^B$ are also log-concave. Note that $Q_1^B,\ldots,Q_{\tau_Q}^B,Q_{\tau_Q+1}^B$ are independent and identically distributed following a log-concave distribution. By Lemma EC.1 (i) and (ii), the product $\mathbf{P}(\sum_{s=1}^{\tau_Q} Q_s^B=n-j)\mathbf{P}(Q_{\tau_Q+1}^B=j)$ is also log-concave. Based on these results, we can show that the following four inequalities hold.

(i) For any $t\in[1,\ldots,\tau_Q]$, we have $\mathbf{P}(Q_t^B=j+1)^2 \geq \mathbf{P}(Q_t^B=j)\mathbf{P}(Q_t^B=j+1)$ because the distribution of Q_t is log-concave;

(ii) For any $s,\ t\in[1,\ldots,\tau_Q]$, we have $\mathbf{P}(Q_s^B=j+1)\mathbf{P}(Q_t^B=j+1)\geq \mathbf{P}(Q_s^B=j)\mathbf{P}(Q_t^B=j+2)$ because the distributions of Q_s and Q_t are identical and log-concave;

(iii) The inequality $\mathbf{P}(\sum_{s=1}^{\tau_Q} Q_s^B=n-j-1)^2\mathbf{P}(Q_{\tau_Q+1}^B=j+1)^2 \geq \mathbf{P}(\sum_{s=1}^{\tau_Q} Q_s^B=n-j)\mathbf{P}(Q_{\tau_Q+1}^B=j)\mathbf{P}(\sum_{s=1}^{\tau_Q} Q_s^B=n-j-2)\mathbf{P}(Q_{\tau_Q+1}^B=j+2)$ holds because $\mathbf{P}(\sum_{s=1}^{\tau_Q} Q_s^B=n-j)\mathbf{P}(Q_{\tau_Q+1}^B=j)$ is log-concave;

(iv) Note that the distribution of $\sum_{s=1}^{\tau_Q} Q_s^B$ is concave by Assumption 2, i.e., $2\mathbf{P}(\sum_{s=1}^{\tau_Q} Q_s^B=n-j-1)\geq \mathbf{P}(\sum_{s=1}^{\tau_Q} Q_s^B=n-j)+\mathbf{P}(\sum_{s=1}^{\tau_Q} Q_s^B=n-j-2)$. For any $t\in[1,\ldots,\tau_Q]$, we have

$2\mathbf{P}(\sum_{s=1}^{\tau_Q} Q_s^B=n-j-1)\mathbf{P}(Q_{\tau_Q+1}^B=j+1)\mathbf{P}(Q_t^B=j+1)\geq \mathbf{P}(\sum_{s=1}^{\tau_Q} Q_s^B=n-j)\mathbf{P}(Q_{\tau_Q+1}^B=j)\mathbf{P}(Q_t^B=j+2)+\mathbf{P}(\sum_{s=1}^{\tau_Q} Q_s^B=n-j-2)\mathbf{P}(Q_{\tau_Q+1}^B=j+2)\mathbf{P}(Q_t^B=j)$ because the distributions of $Q_{\tau_Q+1}^B$ and Q_t^B are identical and log-concave.

We can show that $p_j(n,B)$ satisfies the log-concave condition $p_{j+1}^2(n,B)\geq p_j(n,B)p_{j+2}(n,B)$ using the Definition 1 and the summation of the above inequalities in the four cases. Thus, the biased distribution defined in (8) is also log-concave. Q.E.D.

EC.2.7. Proof of Proposition 4

According to Assumption 1 and Lemma EC.1 (iii), the hazard rate function $h_\alpha(B)=p_\alpha(n,B)/\sum_{j\geq\alpha} p_j(n,B)$ is non-decreasing in α. Then, the function $g(W,B)=\dfrac{1}{1/h_{W+L+1}(B)-1}$ is also non-decreasing in W. Note that the cost difference $f(W,B)>0$ if and only if $g(W,B)>\frac{h}{b}$. The optimal $W^*(n,B)$ is identified when $f(W-1,B)\leq 0$ and $f(W,B)>0$, i.e., $g(W-1,B)\leq h/b$ and $g(W,B)>h/b$. The results in Proposition 4 hold. Q.E.D.

EC.2.8. Proof of Proposition 5

We first show that $g(W,B)$ is non-increasing in B. According the definition of $g(W,B)$, we have

$$g(W,B) = \frac{p_{W+L+1}(n,B)}{\sum_{j=W+L+1}^{n} p_j(n,B) - p_{W+L+1}(n,B)} = \frac{1}{1/h_{W+L+1}(B) - 1},$$

and

$$g(W,B+1) = \frac{1}{1/h_{W+L+1}(B+1) - 1}.$$

Recall that $h_\alpha(B+1) = p_\alpha(n,B+1)/\sum_{j\geq\alpha} p_j(n,B+1)$, where $p_j(n,B+1) = \mathbf{E}\left[\frac{\chi(Q_1^B = j - Z_{B+1}) + \ldots + \chi(Q_{\tau_Q}^B = j - Z_{B+1}) + \chi(\sum_{s=1}^{\tau_Q} Q_s^B = n - j - Z_{B+1})\chi(Q_{\tau_Q+1}^B = j - Z_{B+1})}{\tau_Q + 1}\right]$. Comparing with the form of $p_j(n,B)$, we note that $p_j(n,B+1) = p_{j'}(n,B)$ for some $j' \leq j$. Therefore, $h_\alpha(B+1) = h_{\alpha'}(B)$ with $\alpha' \leq \alpha$ which further implies that $h_\alpha(B)$ is non-increasing in B by the non-decreasing hazard rate property in Lemma EC.1 (iii). As a result, we have $h_{W+L+1}(B+1) \leq h_{W+L+1}(B)$. Thus $g(W,B+1) \leq g(W,B)$, which indicates $g(W,B)$ is non-increasing in B.

Now we prove $W^*(n,B)$ is non-decreasing in B by contradiction. Let $W^*(n,B)$ and $W^*(n,B+1)$ be the optimal delay given the buffer stock B and $B+1$, respectively. Suppose there exists a B such that $W^*(n,B+1) < W^*(n,B)$.

Due to the optimality of $W^*(n,B+1)$ and $W^*(n,B)$, the following results hold,

$$g(W^*(n,B)-1,B) \leq h/b, \;\; g(W^*(n,B),B) > h/b;$$

$$g(W^*(n,B+1)-1,B+1) \leq h/b, \;\; g(W^*(n,B+1),B+1) > h/b.$$

Because $g(W,B)$ is non-increasing in B and non-decreasing in W, we have

$$g(W^*(n,B)-1,B+1) \leq g(W^*(n,B)-1,B) \leq h/b$$

$$g(W^*(n,B),B+1) \geq g(W^*(n,B+1),B+1) > h/b,$$

which indicates that $W^*(n,B)$ also satisfy the optimality condition. However, we know $W^*(n,B+1) < W^*(n,B)$, then $W^*(n,B+1)$ can not be the optimal delay given the buffer stock as $B+1$, as the definition of optimal delay is $W^*(n,B) = \max\{W : g(W-1,B) \leq h/b, \; g(W,B) > h/b\}$. This result contradicts the fact that $W^*(n,B+1)$ is the optimal delay. Therefore, for every B, we have $W^*(n,B+1) \geq W^*(n,B)$. Q.E.D.

EC.3. Benchmarks

EC.3.1. The Croston Policy

Croston's method counts nonzero demand arrivals η and uses two separate forecasts—$\hat{y}_t$ for the inter-arrival time, i.e., from period t to next nonzero demand arrival, and $\hat{z}_\eta$ for the size of η-th nonzero demand. Both are estimated by the exponential weighted moving average method.

Let d_t be the observed demand in period t, $\hat{y}'_\eta$ be the forecast inter-arrival time between η-th and $(\eta+1)$-th nonzero demand arrivals, and τ_t be the elapsed time since the last nonzero demand. With the smoothing coefficient α, the updating procedure is as follows:

- When $d_t \neq 0$, $\eta \leftarrow \eta + 1$,

$$\begin{cases} \hat{z}_\eta = (1-\alpha)\hat{z}_{\eta-1} + \alpha d_t \\ \hat{y}'_\eta = (1-\alpha)\hat{y}'_{\eta-1} + \alpha \tau_t \\ \hat{y}_t = \hat{y}'_\eta \\ \tau_t = 1 \end{cases}$$

- When $d_t = 0$, η and $\hat{z}_\eta$ remain unchanged,

$$\begin{cases} \tau_{t+1} \leftarrow \tau_t + 1 \\ \hat{y}_t = \hat{y}_{t-1} \end{cases}$$

In period t, the demand rate, i.e., the expected demand per period, is given by $\hat{z}_\eta / \hat{y}_t$.

The stock replenishment is a linear combination of the estimates of expected demand, and the mean absolute deviation of the one period ahead forecasting.

EC.3.2. The ITE Policy

They consider a discrete-time model with inventory review periods that are often shorter than the times between successive demand arrivals. Therefore, there are periods in which no demand is received. They model the randomness in demand arrivals by a Bernoulli process with parameter p; i.e., the probability of a positive demand is equal to p in any period. Let the distribution of positive demand size, denoted by X, be a member of the location-scale family of distributions with location parameter τ and scale parameter θ with cdf $F(x;\tau,\theta)$.

If the distribution is known, the problem is a classic newsvendor problem. Let γ_0 denote $h/(h+b)$. For $p > \gamma_0$, the optimal inventory target q^* is given by $\tau + \eta(p,\gamma_0)\theta$, where $\eta(p,\gamma_0)$ is equal to $F^{-1}(1-\gamma_0/p;0,1)$ For $p \leq \gamma_0, q^*$ is zero and it is optimal not to carry any inventory.

However, if we need to estimate the parameters $\hat{p}, \hat{\tau}, \hat{\theta}$ through historical data, then the resulting inventory-target estimator is

$$\hat{Q}(\hat{p},\gamma_o) = \begin{cases} \hat{\tau} + \eta(\hat{p},\gamma_0)\hat{\theta} & \text{if } \hat{p} > \gamma_0, \\ 0 & \text{if } \hat{p} \leq \gamma_0. \end{cases}$$

Next, they quantify the expected cost of incorrectly estimating the parameters in inventory-target estimation, i.e., $C\left(\hat{Q}(\hat{p},\gamma_o);p,\tau,\theta\right) - C(q^*;p,\tau,\theta)$. The expected cost of incorrectly estimating the unknown parameters p, τ, and θ, which is given by $E\left(\Delta_{\hat{\tau},\hat{\theta}}(\gamma;p,\tau,\theta)\right)$, can be written as $\theta E(\Delta_{U,V}(\gamma;p,0,1))$, where $\Delta_{U,V}(\cdot;p,0,1)$ is a function of the random variables $U := (\hat{\tau}-\tau)/\theta$ and $V := \hat{\theta}/\theta$, and the demand-occurrence probability p.

For example, for threshold γ and exponentially distributed positive demand with mean θ, $E(\Delta_{\hat{\theta}}(\gamma;p,\tau,\theta))$ corresponding to the expected cost of incorrectly estimating the parameters p and θ is given by

$$\theta \sum_{w>\gamma} \left\{ h \log\left(\frac{w}{\gamma}\right) + p(h+b)(\ell(n,\gamma,w)-1) \right\} P(\hat{p}=w).$$

The decision maker solves the optimization problem

$$\min_{\gamma \in [0,1]} \max_{p \in \mathcal{P}} E\left(\Delta_{\hat{\tau},\hat{\theta}}(\gamma; p, \tau, \theta)\right),$$

where $\mathcal{P} := \left[\max\left(0, \hat{p} - v_{1-\frac{\alpha}{2}}\sqrt{\frac{\hat{p}(1-\hat{p})}{n}}\right), \min\left(1, \hat{p} + v_{1-\frac{\alpha}{2}}\sqrt{\frac{\hat{p}(1-\hat{p})}{n}}\right)\right]$.

EC.4. The impact of biased estimates under different shortage costs

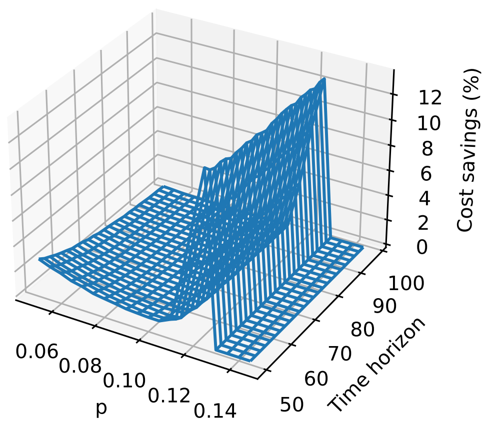

(a) The cost savings using biased estimates ($b = 8$).

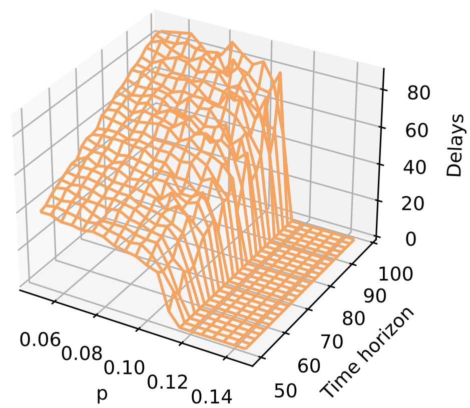

(b) Delays using biased estimates ($b = 8$).

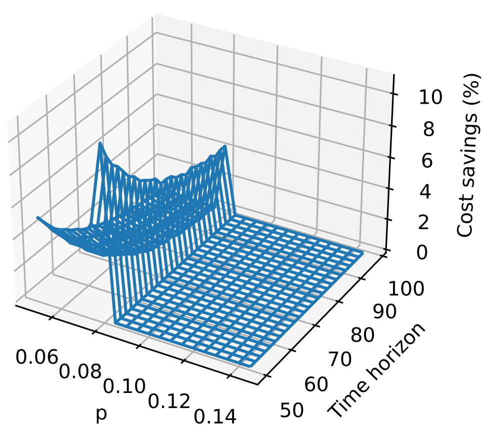

(c) The cost savings using biased estimates ($b = 12$).

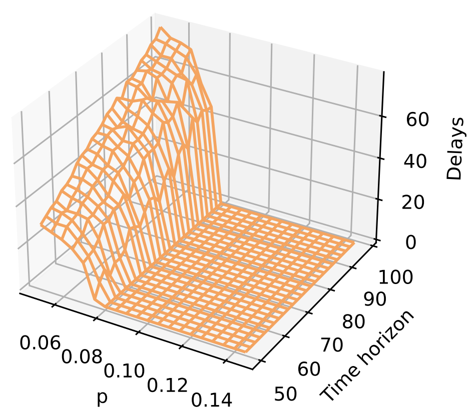

(d) Delays using biased estimates ($b = 12$).

Figure EC.1 The impact of biased estimates.

EC.5. Additional numerical results under different shortage costs

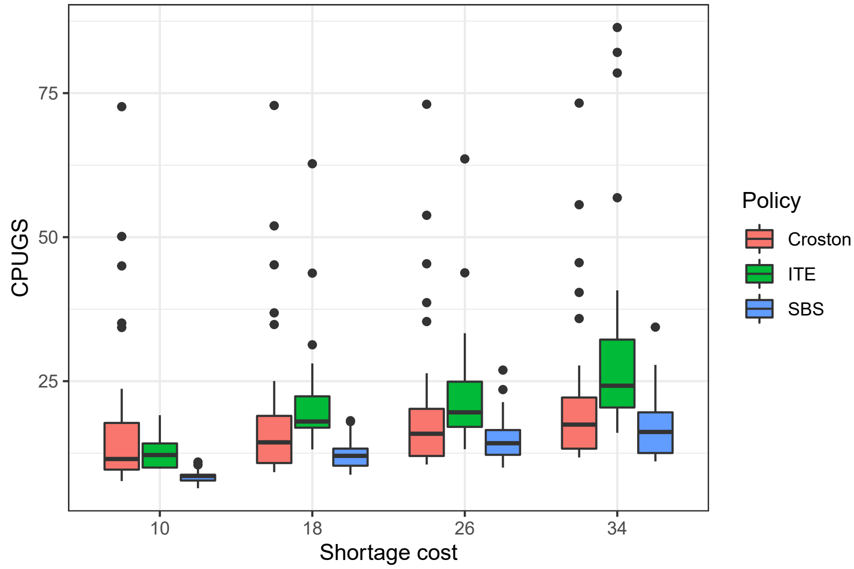

Figure EC.2 **Average CPUGS of the 31 products.**

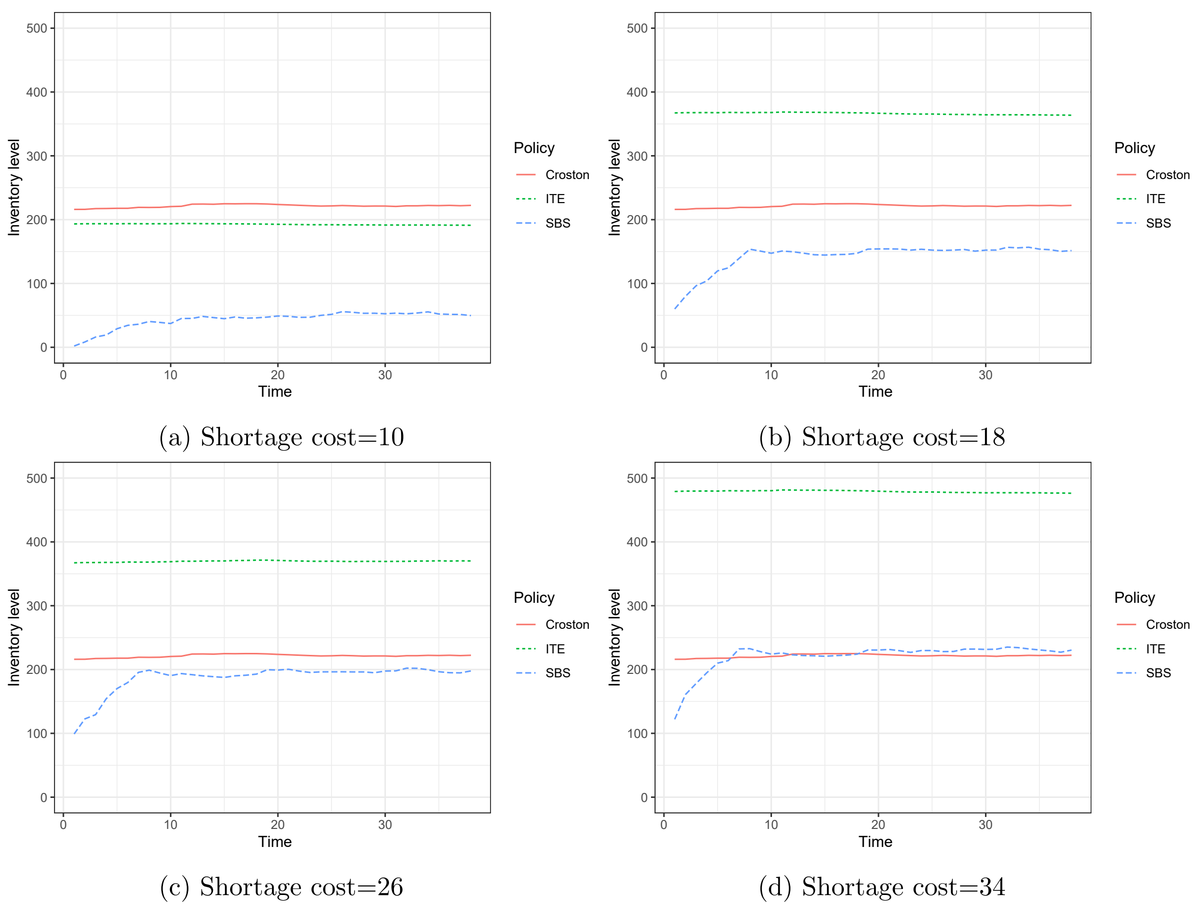

(a) Shortage cost=10

(b) Shortage cost=18

(c) Shortage cost=26

(d) Shortage cost=34

Figure EC.3 **Total inventory levels of 31 products over time under different shortage costs**