

Journal Pre-proof

Dancing to the Rhythm of Demand: Inventory Control for Compound Renewal Demand

Sarah Van Der Auweraer, Thomas Van Pelt, Joachim Arts

PII: S0377-2217(26)00276-6
DOI: <https://doi.org/10.1016/j.ejor.2026.03.029>
Reference: EOR 20056

To appear in: *European Journal of Operational Research*

Received date: 27 May 2025
Accepted date: 20 March 2026

Please cite this article as: Sarah Van Der Auweraer, Thomas Van Pelt, Joachim Arts, Dancing to the Rhythm of Demand: Inventory Control for Compound Renewal Demand, *European Journal of Operational Research* (2026), doi: <https://doi.org/10.1016/j.ejor.2026.03.029>

This is a PDF of an article that has undergone enhancements after acceptance, such as the addition of a cover page and metadata, and formatting for readability. This version will undergo additional copyediting, typesetting and review before it is published in its final form. As such, this version is no longer the Accepted Manuscript, but it is not yet the definitive Version of Record; we are providing this early version to give early visibility of the article. Please note that Elsevier's sharing policy for the Published Journal Article applies to this version, see: <https://www.elsevier.com/about/policies-and-standards/sharing#4-published-journal-article>. Please also note that, during the production process, errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

© 2026 Published by Elsevier B.V.



Highlights

Dancing to the Rhythm of Demand: Inventory Control for Compound Renewal Demand

Sarah Van der Auweraer, Thomas van Pelt, Joachim Arts

- Data from practice show that the time between intermittent demand events is often not Markovian, but instead displays a rhythm.
- The time since the last demand occurrence can be a predictor for future demand.
- A state-dependent base-stock policy is optimal for rhythmic intermittent demand.

Journal Pre-proof

Dancing to the Rhythm of Demand: Inventory Control for Compound Renewal Demand

Sarah Van der Auweraer^a, Thomas van Pelt^b, Joachim Arts^b

^aIESEG School of Management, Univ. Lille, CNRS, UMR 9221 - LEM - Lille Economie Management, France

^bUniversity of Luxembourg, Luxembourg Centre for Logistics and Supply Chain Management, Luxembourg

Abstract

The demand for many items is intermittent. As a consequence, managing the inventories for these items efficiently and effectively can be challenging. Intermittent demand patterns are not well modeled by a single distribution of demand per period. It is therefore common to model the time between demand occurrences and the size of demand occurrences separately. Such models typically implicitly assume that times between demand occurrences are Markovian. Data from practice, however, indicate that the time between demand events is not necessarily Markovian but instead it can display a certain *rhythm*. Consequently, the time since the last demand occurrence is an important predictor of future demand for those items. We use the discrete compound renewal process to model such rhythmic intermittent demand. In a periodic review inventory system, we show that the optimal inventory policy is a state-dependent base-stock policy, where the order-up-to-levels are non-decreasing in the time since the last demand, regardless of the nature of the distribution of time between demand occurrences. We benchmark the performance of our approach against heuristic policies in terms of average inventory costs in both a numerical experiment, in which demand processes are known, and on real data, where the demand process is not known and needs to be estimated. We find that our approach adds the most value when demand occurs at regular intervals, with limited deviations, and when such intervals are long. The time since the last demand contains valuable information that our approach exploits.

Keywords: intermittent demand, inventory, optimal policy, rhythmic demand, renewal process

1. Introduction

Intermittent demand is characterized by many periods with no demand at all, that are seemingly randomly interspersed with demand occurrences. It is unknown when the next demand will occur, and there is uncertainty about the size of demand when it occurs. This compound nature of demand uncertainty complicates efficient inventory management.

Intermittent demand is commonly observed in practice, for example, in sectors such as manufacturing, aviation, IT (Willemain et al. 2004) and in the petrochemical industry (Molenaers et al. 2012). Although intermittent demand is typically associated with spare parts, it is also prevalent for other types of items, such as electrical products (Johnston et al. 2003), and the retail industry (see for example Zhang et al. 2021), as also indicated by the recent M5 competition (Makridakis et al. 2022). Some companies observe intermittent demand for more than half of the products in their inventory (e.g., Johnston et al. 2003).

Specific approaches have been developed to deal with such intermittent demand, initiated by the work of Croston (1972). In the existing literature, the analysis of intermittent demand time series is commonly decomposed to create separate estimates for the time between demand occurrences (the inter-demand interval) and the size of a demand occurrence. These analyses often implicitly assume that the time between demand occurrences is Markovian, i.e., the probability of observing a demand in a certain period is considered to be independent of the time since the last demand occurrence. However, data from practice indicate that the times between demand occurrences are not necessarily Markovian. Non-Markovian times between demand occurrences lead to *rhythmic* patterns in demand processes when the hazard rate is increasing, which has recently also been observed by Prak and Rogetzer (2022) and Turkmen et al. (2020). Reasons for demand to be evenly spaced in practice include (i) relatively few customers who order large quantities due to economies of scale, or (ii) spare part demand for a few machines that degrade over time.

In this paper we propose a demand model that accommodates such rhythmic intermittent demand. We model demand as a discrete compound renewal process. Unsatisfied demand is backordered at a penalty cost and inventory on-hand at the end of a period incurs a holding cost. We allow for positive leadtimes.

This research is strongly rooted in practice. It is the outcome of a joint research project with a company in the chemical industry, which produces specialty chemicals for its customers within a make-to-stock environment. The company experiences rhythmic intermittent demand for many of its items, as observed in their data. Nonetheless, the company currently does not consider this rhythm in its demand forecasting or inventory decision making. Instead, their current practice uses standard techniques to make forecasts and manage their inventories. In particular, they use simple exponential smoothing to create point forecasts of demand and make inventory decisions based on this point estimate and past forecast errors. A consequence of this practice is that forecasts are highest just after a demand occurs. Replenishments are triggered at this point in time, even though the forecast decreases after this moment until the next demand occurrence. Yet when demand occurrences are evenly spaced (i.e., they are rhythmic), the placement of a replenishment order can be delayed until demand is again likely to occur. This will reduce the inventory investment while maintaining the same service level to customers. The decision to delay the placement of a replenishment order should be a feature of the inventory control policy, not of the forecast updating procedure.

The model we developed in the context of this collaboration contributes to both the theory and practice of inventory management as we characterize the structure of optimal replenishment policies and design heuristics based on these structural results. The main contributions of our paper include descriptions of the form of optimal order policies. In particular, we show that the optimal policy is a state-dependent base-stock policy, where the state is the time since the last demand observation. We prove this by induction on the value function. We also show that there exist state-dependent base-stock policies for which the optimal base-stock level is non-decreasing in the time since the last demand. As such, any algorithm that searches for optimal (or good) base-stock levels can constrain the search space to non-decreasing base-stock levels only. Contrary to what may be expected, this result is not proven by showing submodularity of the value function. Instead we prove sensitivities of optimal base-stock levels with elementary arguments, which can provide new perspectives (see for example Perera et al. 2017). We characterize the sensitivities of the base-stock level to the current state and the proof exploits structure in compound renewal processes directly.

Moreover, we also contribute to inventory control practice. We demonstrate how the proposed approach can be implemented in practice through the development of numerically stable algorithms that can scale to large assortments as encountered in industry. Additionally, we introduce two heuristic policies that are more easy to implement, and we benchmark the performance of the different approaches.

We first benchmark the two fast heuristics against the optimal policy in a numerical experiment in which the demand process is assumed to be known. We identify settings where it is especially valuable to exploit the demand rhythm inherent in compound renewal processes to support companies in selecting suitable items and directing their funds and efforts. Moreover, we evaluate and benchmark the performance when the true demand process is not known but demand data is available. We apply the proposed approaches to the entire assortment of the chemical company. Additionally, to ensure that our findings can also be generalized to other industries, we validate the approach on items from other industries that are known in literature: (i) air force spare parts, (ii) car parts, and (iii) retail. Especially the retail industry is not typically known for having items with intermittent demand. However, shorter leadtimes and increased product differentiation/customization make intermittent demand also increasingly common in this industry.

We illustrate the research problem, proposed approach, and practical application using a running example throughout the paper. The example is drawn from the chemical company with which we collaborated for this research project. Subsequently, we apply the proposed method to the company's entire product assortment. While the example is not necessary for comprehension and may be omitted by readers, its inclusion enhances the reproducibility of the results. Nonetheless, incorporating the example has implications for the structure of the paper. Where a typical paper first presents the model and only then discusses estimation procedures, we have reversed the order in this work to simplify practical implementation.

Example 1 *The Specialty Chemicals Company (SCC) produces specialty chemicals in a make-to-stock environment. Each item has one or only a few customers. As a consequence, the company observes intermittent (lumpy) demand for many of its items. Managing inventories for such items is challenging, as there is uncertainty in both the timing of the next demand arrival and in the size of the demand when it occurs. Additionally, the SCC demand data suggest demand occurrences are rhythmic, i.e., the probability of observing a demand is not constant in every period, but instead it increases as the time since the last demand becomes longer. As an example, the demand pattern of SCC item ChemEx is presented in Figure 1. It seems evident that in such case it makes sense not to place a replenishment order immediately after a demand occurrence but instead to postpone order placements until demand becomes more likely to occur. However, it is not trivial to decide when to place this replenishment order and what needs to be the size of such order.*

The remainder of this paper is organized as follows. Section 2 discusses the relevant literature. In Section 3 we present the demand model. The demand estimation process is discussed in Section 4. Section 5 describes the inventory model and Section 6 explains the structure of the optimal ordering policy. Section 7 additionally discusses two heuristic policies and Section 8 presents results for a numerical experiment in which demand processes are known. Section 9 subsequently contains the evaluation based on the demand data of the chemical company that




Figure 1: Demand of ChemEx

inspired our research, as well as a replication study on four other real-life data sets from three different industries. Finally, Section 10 provides concluding remarks.

2. Related literature

It is quite common to observe intermittent demand in practice, especially in sectors such as manufacturing, aviation or IT (Willemain et al. 2004), or the petrochemical (Molenaers et al. 2012) and the automotive industry (Syntetos et al. 2005), with some companies observing this kind of demand for more than half of their products (Johnston et al. 2003, Molenaers et al. 2012). Intermittent demand is typically associated with spare parts. However, it is also common for other types of products, as illustrated by the recent M5 forecasting competition (Makridakis et al. 2022) which included intermittent demand from the retail industry. Syntetos et al. (2016a) state that as the granularity of demand analyses increases, we unavoidably end up with intermittence.

In a seminal paper, Croston (1972) tackles the issue of forecasting such intermittent demand by separately forecasting the two components of uncertainty, i.e., the size of the demand when it occurs and the time between demand occurrences. Since then, a number of extensions and improvements on Croston's approach have been proposed (see for example Johnston and Boylan 1996, Levén and Segerstedt 2004, Teunter et al. 2011, Babai et al. 2019), with the approximation by Syntetos and Boylan (2005) being the most widely known. Such approaches consider demand (during the leadtime) to be stationary and assume Markovian demand arrivals. As a result, they generate constant forecasts for the demand in each period of the forecast horizon. However, knowing the demand rate does not provide information on when the next demand will arrive (Bi et al. 2023). Moreover, in practice, we do not necessarily observe a constant demand rate (Smith and Dekker 1997, Syntetos et al. 2015) as recently also observed by Prak and Rogetzer (2022), Turkmen et al. (2020), and Syntetos et al. (2015).

Several approaches exist when the demand process is not an independently and identically distributed sequence of random variables, but the true process remains unspecified. For example, Willemain et al. (2004) and Porras and Dekker (2008) use bootstrapping and Gutierrez et al. (2008), Lolli et al. (2017), and Mukhopadhyay et al. (2012) use neural networks to forecast intermittent demand.

Pennings et al. (2017) develop methods to improve intermittent demand forecasting by estimating the correlation between demand size and interval. Babai et al. (2021) use Bayesian updating to estimate the parameters of a compound

Poisson demand process and test their procedure on intermittent data. Demand with geometric inter-arrival times can be modeled by an i.i.d. demand sequence (by putting probability mass in the origin). Nevertheless, it can be beneficial to isolate the geometric inter-arrival times explicitly, as is done in a compound Bernoulli process (Janssen et al. 1998, Teunter et al. 2010).

From an inventory perspective, Schultz (1987), Schultz (1989) and Syntetos et al. (2016b) motivate delaying replenishment orders when demand is intermittent, thereby assuming an increasing demand rate. The authors state that such order delay can reduce inventory holding costs with little additional risk of shortages. In a recent paper, Bi et al. (2023) similarly propose a gradual build-up of inventory by delaying replenishment orders for slow-moving items when evaluating inventory metrics over a finite horizon. The authors focus on optimizing the timing and quantity of replenishment orders when taking into account the finite horizon effect for long tail items.

According to Fildes et al. (2008), one of the key issues remaining for intermittent demand research is a better modeling of the demand characteristics underlying the intermittent demand in the context of inventory management. Similarly, Nasr (2022) points out that formulating models that capture the behavior of the demand process is one of the main challenges when modeling stochastic inventory systems. In our paper we therefore focus on modeling the demand process itself, rather than generating a forecast, and we derive an optimal order policy in the presence of intermittent demand.

Zipkin (2000) describes the concept of *world-driven demand*, where an exogenous system, the *world*, influences the demand rate and the state of the world can change over time. If the world is cyclic, it transitions with certainty from one state to another as time advances from one period to the next. Every period within the cycle induces a certain demand behavior, which is recurrent over time. Cyclic demand has been studied by, for example, Aviv and Federgruen (1997), Kapuściński and Tayur (1998), Karlin (1960), and Zipkin (1989). When a cycle comprises for example a year or a month, it is sometimes also called seasonal demand. The specific case of seasonal intermittent demand forecasting has been studied by Lindsey and Pavur (2013) and Kourentzes and Athanasopoulos (2021). Observe that seasonal or cyclic approaches consider deterministic cycle lengths and state transition probabilities equal to zero or one. In the case of spare parts, the state of the world may be the status of service contracts (Pinçe et al. 2015), or of maintenance plans (Arts et al. 2016).

More generally, we can consider an independent but *stochastic* modulating process that determines the demand distribution. This is referred to as Markov-modulated demand (see Beyer and Sethi 1997, Cheng and Sethi 1999, Chen and Song 2001, Sethi and Cheng 1997, Song and Zipkin 1993). Such an underlying process can for example represent economic fluctuations or the evolution of a product life-cycle (Song and Zipkin 1993). In such settings, the optimal policy is the state-dependent equivalent of the stationary policy, where the policy parameters must be adjusted to reflect the state-of-the-world (Chen and Song 2001). Note that Markov-modulated demand is dependent on the state-of-the-world, but the state itself is independent of the demand that occurred in the periods before.

In our paper, by contrast, we consider the modulating chain and the demand process to be intertwined: the (non-)occurrence of a demand influences the next state-of-the-world. As such, our work also relates to literature on Markovian Arrival Processes (MAP) (see for example Ibe 2013, Dizbin and Tan 2020) where a transition in the environment

process can coincide with an arrival. Although theoretically this concept is well known, few practical studies exist. In particular, almost no literature exists on modeling intermittent demand using a MAP. An exception is the work of Bi et al. (2023). Hence, our work contributes to this domain.

The literature discussed in this section can be categorized along several dimensions as is done in Table 1 (in the same order as papers appear in the review). The “inventory control” columns of this table show whether the paper proposes a stationary policy (a), state dependent policy (b) or characterizes the optimal policy (c). Note that some papers, notably the forecasting papers, do not propose any inventory policy.

The “Demand Model” columns show whether the time between demand arrival is assumed to be Markovian (d) or not (e) in a given paper. Columns (f) and (g) show whether the paper includes empirical validation in the form forecast accuracy or inventory performance respectively. Columns (h) and (i) indicate whether a paper includes simulation or numerical/analytical results to validate forecast accuracy or inventory performance respectively. Note that papers that consider non-Markovian demand typically focus exclusively on forecasting, or characterize the form of an optimal policy. This paper is the only paper that studies simple heuristics as well as optimal policies and performs both an empirical and numerical validation of the actual inventory performance of several approaches.

Paper	Inventory control			Demand Model		Empirical Validation		Numerical/Analytical Validation	
	(a)	(b)	(c)	(d)	(e)	(f)	(g)	(h)	(i)
Croston (1972)				✓				✓	
Johnston and Boylan (1996)				✓		✓			
Levén and Segerstedt (2004)	✓			✓					✓
Teunter et al. (2011)				✓				✓	
Babai et al. (2019)				✓				✓	
Syntetos and Boylan (2005)				✓		✓			
Bi et al. (2023)				✓					✓
Smith and Dekker (1997)	✓				✓				✓
Syntetos et al. (2015)					✓			✓	
Prak and Rogetzer (2022)					✓		✓		
Turkmen et al. (2020)					✓	✓		✓	
Willemain et al. (2004)					✓	✓			
Porras and Dekker (2008)	✓				✓		✓		
Gutierrez et al. (2008)					✓	✓			
Lolli et al. (2017)					✓	✓			
Mukhopadhyay et al. (2012)					✓	✓			
Penning et al. (2017)				✓		✓			
Dizbin and Tan (2020)			✓		✓				
Babai et al. (2021)					✓		✓		✓
Janssen et al. (1998)	✓			✓					✓
Teunter et al. (2010)				✓		✓		✓	
Schultz (1987)		✓			✓				
Schultz (1989)		✓			✓				
Syntetos et al. (2016b)		✓			✓				✓
Fildes et al. (2008)					✓				
Nasr (2022)	✓				✓				✓
Song and Zipkin (1993)			✓		✓				
Chen and Song (2001)			✓		✓				
Beyer and Sethi (1997)			✓		✓				
Cheng and Sethi (1999)			✓		✓				
Sethi and Cheng (1997)			✓		✓				
This paper	✓	✓	✓		✓	✓	✓		✓

Table 1: (a) Stationary Policy, (b) State Dependent Policy, (c) Optimal Policy, (d) Markovian, (e) non-Markovian, (f) Forecast Accuracy, (g) Inventory Performance, (h) Forecast Accuracy, (i) Inventory Performance

3. Compound renewal demand

We consider a single stock-point under periodic review with the review period equal to one. We consider the following order of events: At the beginning of each review epoch we observe the current inventory position and a replenishment order can be placed that arrives after a (non-negative) leadtime L that is an integer multiple of the review period. Next, we receive the orders placed one leadtime ago. Following, we satisfy outstanding backorders and we observe demand for the current period. Excess demand that cannot be filled by the end of a review period is

backordered and incurs a penalty cost of p per item per period. Similarly, inventory on-hand at the end of a review period incurs a holding cost of h per item per period.

So far the inventory standard conforms to the canonical backorder system studied first by Arrow et al. (1951); recent examples include Janakiraman et al. (2007) and van Jaarsveld and Arts (2024). Our model deviates from the canonical model as demand is not independent and identically distributed across periods. Rather demand will come from a discrete compound renewal process. Such a process can model demand that is intermittent, and for which the time between demand occurrences is not Markovian.

In particular, the demand process works as follows. We let T_i represent the time from (positive) demand occurrence $(i-1)$ till demand occurrence i and $\{T_i\}_{i=1}^{\infty}$ is a sequence of i.i.d random variables on $\mathbb{N} = \{1, 2, 3, \dots\}$ with finite mean. Note that the event $T_i = 1$ means that the i -th demand hit happens in the period immediately following the previous hit. Similarly, H_i is the i -th demand hit size and $\{H_i\}_{i=1}^{\infty}$ is a sequence of independent and identically distributed (i.i.d) random hit sizes on $\mathbb{N} = \{1, 2, 3, \dots\}$. We consider the distributions of H_i and T_i to be independent, to not confound the effect of the time since the last demand and the correlation with the size of the demand when it occurs. This assumption of independence is also supported by the company data used later in this paper (see Section 9.1.1). Let N_t denote the number of positive demand occurrences from the time origin until period t . N_t is a discrete renewal process and satisfies

$$N_t := \max \left\{ n \in \mathbb{N}_0 \mid \sum_{i=1}^n T_i \leq t \right\}. \quad (1)$$

Then the demand up to and including period t is denoted by D_t and is given by:

$$D_t := \sum_{i=1}^{N_t} H_i. \quad (2)$$

Later it will be convenient to have notation for the hazard (or mortality) rate of T_i . This hazard rate quantifies the probability of observing a demand in the current period given that we have already observed y periods without a demand, i.e., the probability of having *exactly* y periods between two consecutive demands given that we have already observed y periods since the last demand. This hazard rate $m(y)$ is defined as

$$m(y) := \mathbb{P}(T_i = y \mid T_i \geq y) = \frac{\mathbb{P}(T_i = y)}{\mathbb{P}(T_i \geq y)}.$$

If $m(y)$ is an increasing function, demand is considered as *rhythmic*, i.e., the probability of observing a demand increases as the time since the last demand becomes longer. Otherwise demand is said to be *arhythmic*, i.e., $m(y)$ is not increasing but instead it can be constant, decreasing, or fluctuating in y .

We define the random variable Y_t to be the time since the last demand occurrence at the beginning of period t (before demand realizes):

$$Y_t := t - \sum_{i=1}^{N_t-1} T_i, \quad (3)$$

and note that $Y_t \in \mathbb{N} := \{1, 2, 3, \dots\}$.

It is convenient to write the number of demand hits from a time period t_1 up to and including any other period t_2

as: $N_{t_1, t_2} := N_{t_2} - N_{t_1-1}$. Note that at time t_1 the distribution of N_{t_1, t_2} depends on the time from t_1 until the next hit arrives, which depends on how long ago the previous hit arrived, i.e. the distribution of N_{t_1, t_2} depends on Y_{t_1} . We let $N_{t_1, t_2}(y)$ denote the number of hits from period t_1 until period t_2 conditional on $Y_{t_1} = y$ so that

$$\mathbb{P}(N_{t_1, t_2}(y) = x) := \mathbb{P}(N_{t_1, t_2} = x \mid Y_{t_1} = y). \quad (4)$$

Analogously, let D_{t_1, t_2} denote demand from a period t_1 up to and including another period t_2 :

$$D_{t_1, t_2} := D_{t_2} - D_{t_1-1} = \sum_{i=1+N_{t_1-1}}^{N_{t_2}} H_i \quad (5)$$

Note that at time t_1 , the distribution of D_{t_1, t_2} depends on the time from t_1 until the next demand hit arrives, which depends on how long ago the previous hit arrived, i.e. the distribution of D_{t_1, t_2} depends on Y_{t_1} : We let $D_{t_1, t_2}(y)$ denote demand from period t_1 until period t_2 conditional on $Y_{t_1} = y$ so that

$$\mathbb{P}(D_{t_1, t_2}(y) = x) := \mathbb{P}(D_{t_1, t_2} = x \mid Y_{t_1} = y) \quad (6)$$

In inventory theory, orders are usually placed to cover demand over the leadtime (+ review period). The distribution of this demand over the leadtime (+ review period) will depend on how long ago the last demand hit occurred, i.e. $D_{t, t+L}$ will depend on Y_t for any $L \in \mathbb{N}_0$ as equation 6 shows. Good inventory policies will be based on the distribution of $D_{t, t+L}(y)$ and Proposition 1 shows how this distribution can be computed for different values of y . These expressions will be used later in the paper to compute various quantities of practical interest. The proof follows from the application of discrete renewal theory and is presented in Appendix A.

Proposition 1 *The distributions of demand and number of hits over a leadtime satisfy:*

- (i) $\mathbb{P}(N_{t, t+L}(y) \geq n) = \sum_{j=y}^{L+y} \mathbb{P}\left(\sum_{i=2}^n T_i \leq L + y - j\right) \frac{\mathbb{P}(T_1=j)}{\mathbb{P}(T_1 \geq y)}$
- (ii) $\lim_{t \rightarrow \infty} \mathbb{P}(N_{t, t+L} \geq n) = \sum_{y=1}^{\infty} \mathbb{P}(N_{t, t+L}(y) \geq n) \frac{\mathbb{P}(T_1 \geq y)}{\mathbb{E}[T_1]}$
- (iii) $\mathbb{P}(D_{t, t+L}(y) = x) = \sum_{n=0}^{L+1} \mathbb{P}(N_{t, t+L}(y) = n) \mathbb{P}\left(\sum_{i=1}^n H_i = x\right)$
- (iv) $\lim_{t \rightarrow \infty} \mathbb{P}(D_{t, t+L} = x) = \sum_{y=1}^{\infty} \mathbb{P}(D_{t, t+L}(y) = x) \frac{\mathbb{P}(T_1 \geq y)}{\mathbb{E}[T_1]}$

Example 2 *Figure 1 shows the demand realizations for ChemEx. We can observe the realizations of the time between demand hits and the demand hit sizes for ChemEx, if we assume that the start of our observations coincides with the introduction of this product: $H_1 = 3, H_2 = 5, H_3 = 5, H_4 = 5, H_5 = 5, H_6 = 6, H_7 = 6$, and $T_1 = 1, T_2 = 6, T_3 = 6, T_4 = 8, T_5 = 10, T_6 = 9, T_7 = 11$. Until time 15 we have observed 3 demands; namely at $t = 1, t = 7$, and $t = 13$. Thus, $N_{15} = 3$ and $D_{15} = H_1 + H_2 + H_3 = 13$. We can also observe that $Y_{15} = 2$, as we can count 2 periods since the last demand that occurred in period 13. Next, if we choose $t_1 = 16$ and $t_2 = 30$, then $N_{16, 30} = N_{30} - N_{15} = 4 - 3 = 1$ and $D_{16, 30} = D_{30} - D_{15} = 18 - 13 = 5$, which is also equal to H_4 (occurring at $t = 21$).*

To calculate the distribution of the demand, we need the distributions of the random variables T_i and H_i . In reality, these distributions are not known but instead they must be estimated from data. We will do so in Example 4 and calculate the distribution of the leadtime demand from Proposition 1 in Example 6

Example 3 We also empirically investigate whether T_i and H_i are indeed independent for ChemEx. For this, we calculate the Pearson correlation coefficient (ρ) for the inter-demand sequence $\{6, 6, 8, 10, 9, 11\}$ and the hit sizes $\{5, 5, 5, 5, 6, 6\}$. We evaluate whether ρ is significantly different from 0 ($H_0 : \rho = 0; H_a : \rho \neq 0$). We obtain a correlation coefficient of 0.6250 with a p -value of 0.1846. This implies that the correlation, albeit positive, is not significantly different from 0. Thus, we do not find evidence to reject the assumption of independence between the time between hits and the hit sizes.

4. Demand estimation

Where inventory theory often starts from a given demand process, inventory management in practice starts from demand data. The demand model introduced in Section 3 must be estimated from data in any practical setting. This section outlines how maximum likelihood estimation (MLE) can be used to estimate the distributions of T_i and H_i . The data will typically be given in the form of n sequential demands, i.e., in the form $d = (d_1, \dots, d_n)$ where many d_t will be zero due to the intermittent nature of demand. The hit sizes are observed directly and it is straightforward to find maximum likelihood estimates for the hit size distribution. Estimation of the inter-demand interval distribution is more involved.

The number of periods until the first demand observation in the data is called $\tilde{T}_1 := \min\{t \in \mathbb{N} \mid d_t > 0\}$. Note that in general the data set will not start with the introduction of a product and so $\tilde{T}_1 \leq T_1$, i.e. the first inter-demand interval is right-censored. Similarly, the final inter-demand interval is right-censored. Let $m := \sum_{t=1}^n \mathbf{1}_{\{d_t > 0\}} + 1$ denote the number of inter-demand intervals in the sample. Then the length of the final interval is given by $\tilde{T}_m := n + 1 - \max\{t \in \{1, \dots, n\} \mid d_t > 0\}$. The other (non-censored) inter-demand intervals are given by $\tilde{T}_i = \min\{t \in \mathbb{N} \mid d_{\tilde{T}_{i-1}+t} > 0\}$ for $i \in \{2, \dots, m-1\}$. Note that $\tilde{T}_i = T_i$ for $i \in \{2, \dots, m-1\}$ but $\tilde{T}_1 \leq T_1$ and $\tilde{T}_m \leq T_m$ due to censoring. We let $\tilde{T} := (\tilde{T}_1, \dots, \tilde{T}_m)$ for convenience.

Let $f_\theta(\cdot)$ and $F_\theta(\cdot)$ denote the likelihood and distribution functions of a parametric family of distributions where θ is the parameter (which may be a vector quantity). The log-likelihood function of the observed inter-demand intervals for a given parametric family of distributions is then given by

$$\mathcal{L}(\tilde{T}|\theta) = \ln(1 - F_\theta(\tilde{T}_1 - 1)) + \ln(1 - F_\theta(\tilde{T}_m - 1)) + \sum_{i=2}^{m-1} \ln(f_\theta(\tilde{T}_i)). \quad (7)$$

The maximum likelihood estimate is then the θ that maximizes the log-likelihood. The relative fit of different parametric families of distributions can be assessed by looking at their respective maximized log-likelihoods.

Example 4 In Example 2 we assume that the start of the observations coincides with the introduction of the product. In reality, however, this is not necessarily the case. For ChemEx, this implies $m = 8$ and $\tilde{T}_8 = 53 + 1 - 51 = 3$.

	Discrete Weibull	Mixture of binomial	Negative binomial	Poisson
T_i	12.25	12.27	12.97	12.97
H_i		8.77	12.16	12.16

Table 2: Negative log-likelihood for different discrete distributions for ChemEx.

Additionally, $T_1 \geq \tilde{T}_1 = 1$ and $T_8 \geq \tilde{T}_8 = 3$.

We compare the goodness-of-fit (GoF) of various discrete distributions on T_i and H_i : (i) a mixture of binomial distributions, suitable when data exhibit a low coefficient of variation (Adan et al. 1995), (ii) the negative binomial distribution, i.e., the discrete counterpart of the Erlang distribution and a very versatile distribution, and (iii) the Poisson distribution. We shift the distributions with one unit to exclude zero values. Additionally, we fit (iv) the discrete Weibull distribution (see Almalki and Nadarajah 2014) to the inter-demand intervals because of its closed form expression of the hazard rate, which is convenient for our study. We provide more details on these distributions in Appendix B.1, Appendix B.2 and Appendix B.3. The negative log-likelihood values, presented in Table 2, indicate that the mixture of binomial distributions provides the best fit for the hit sizes of ChemEx ($\hat{k} = 4$, $\hat{p} = 0.80$, and $\hat{q} = 10^{-5}8.12$ – see Appendix B.2), while the discrete Weibull distribution best fits the time between demand hits (scale $\hat{\alpha} \approx 8.57$ and shape $\hat{\beta} \approx 4.87 > 1$ – see Appendix B.1).

This analysis demonstrates GoF of these distributions to the data in the relative sense. We also assess the validity of these distributional assumptions through an adjusted Kolmogorov-Smirnov test, adapted to discrete distributions with censoring. This somewhat involved procedure uses Monte-Carlo simulation as the distribution of the test statistic is unknown, even asymptotically (see Appendix B.4 for further details). Using this procedure we do not find evidence to reject the distributional assumptions.

Example 5 Next, we test whether T_i exhibits an increasing hazard rate, as suggested by $\hat{\beta} > 1$. While a formal test for monotone increasing hazard rates exists for continuous distributions (Hall and Van Keilegom 2005), there is no equivalent for discrete distributions with censoring. Therefore, we examine whether $\hat{\beta}$ is significantly larger than 1 contingent on the suitability of the discrete Weibull model for the time between demands. We inspect the estimated standard error of $\hat{\beta}$, $SE_{\hat{\beta}} \approx 1.64$, and use that $\hat{\beta}$ is asymptotically normal by the theory of MLE. If $H_0 : \beta \leq 1$, then $z = \frac{\hat{\beta}-1}{SE_{\hat{\beta}}} \approx 2.36$ is a draw from a standard normal distribution under the null hypothesis. The corresponding p-value is 0.0090. Thus we confidently reject the hypothesis that $\beta \leq 1$ for the alternative hypothesis that demand is rhythmic.

Example 6 With estimated distributions for T_i and H_i in Example 4, it is possible to compute the distribution of the demand over the leadtime (+ review period) and other quantities of interest using Proposition 1. Figure 2 shows the probability mass function of $D_{t,t+L}(y)$ and $D_{t,t+L}$ of ChemEx for $L = 0$ for different values of y .

5. Inventory model

We let q_t denote the order placed in period $t - L$. This convention implies that q_t is the order that arrives in period t . Note that q_{t+L} is a decision made based on the state of the system at time t , which we describe next. Let I_t denote




Figure 2: Illustration of the estimated probabilities of demand during the leadtime (+ review period, $L = 0$): $D_{t,t}(y)$ and $D_{t,t}$.

the inventory level at the beginning of period t after receiving q_t and let J_t be the inventory level at the end of period t when demand has occurred. Furthermore, let X_t denote the inventory position at the beginning of period t before placing an order:

$$X_t := I_t + \sum_{\tau=t+1}^{t+L-1} q_\tau.$$

Note that $I_t, J_t, X_t \in \mathbb{Z}$. The inventory dynamics are

$$I_{t+1} = I_t - D_{t,t} + q_{t+1} \quad J_t = I_t - D_{t,t} \quad \text{and} \quad X_{t+1} = X_t - D_{t,t} + q_{t+L}.$$

Thus we assume that demand in excess of available inventory is backlogged. Note that the inventory position dynamics do not require knowledge of the full state $(I_t, q_{t+1}, \dots, q_{t+L-1})$; only the current inventory position and order are needed. Each item in inventory at the end of a period incurs a cost h and each backorder incurs a cost p such that the cost incurred in period t , C_t are given by

$$C_t = hJ_t^+ + p(-J_t)^+,$$

where we use the notation that $x^+ = \max(0, x)$. Next we observe that the ordering decision made in period t does not affect costs until period $t + L$. By the inventory balance equation we have that

$$X_t + q_{t+L} - D_{t,t+L} = J_{t+L}.$$

Therefore, the expected costs that will be incurred in period $t + L$ when the inventory position in period t is raised to $a = X_t + q_{t+L}$ and $Y_t = y$ are given by

$$c(a, y) = \mathbb{E}[C_{t+L} | Y_t = y, X_t + q_{t+L} = a] = p\mathbb{E}[(D_{t,t+L}(y) - a)^+] + h\mathbb{E}[(a - D_{t,t+L}(y))^+]. \quad (8)$$

Thus, in order to minimize expected costs over a horizon of any size, it suffices to know in each period the inventory position X_t and the time since the last demand hit Y_t , i.e. tuple (X_t, Y_t) is a sufficient state description to minimize costs. This observation holds in many systems (e.g. Song and Zipkin 1996, Chen and Song 2001, Arts et al. 2016) and more detailed arguments are available in textbooks and tutorials (e.g. Zipkin (2000) section 9.6 and Van Houtum

(2006) section 2.1-2.2).

Our objective will be to minimize the discounted costs over the planning horizon that can be affected by our ordering decisions. We let $\gamma \in (0, 1]$ denote the discount factor and $V_t(x, y)$ denote the optimal expected discounted cost over a planning horizon of length $t \in \mathbb{N}$ conditional on $(X_1, Y_1) = (x, y)$, that is

$$V_t(x, y) = \inf_{\pi \in \Pi} \mathbb{E}_{\pi} \left[\sum_{i=1}^t \gamma^{i-1} C_{i+L} \mid X_1 = x, Y_1 = y \right]$$

where Π is the class of non-anticipative ordering policies and the process C_t is understood to be induced by ordering policy $\pi \in \Pi$. V_t is called the value function and satisfies the Bellman optimality equations

$$V_t(x, y) = \min_{a \geq x} \left\{ c(a, y) + \gamma(1 - m(y))V_{t-1}(a, y + 1) + \gamma m(y)\mathbb{E} [V_{t-1}(a - H, 1)] \right\} \quad \forall (x, y) \in \mathbb{Z} \times \mathbb{N}, \quad (9)$$

where $V_0(x, y) := 0$ by convention. Note that we do not allow for negative orders in this setting, since $a \geq x$. Not allowing for negative orders stems from the make-to-stock production context, where items from a production order, once produced, cannot be returned (*un-produced*). They have to be kept in stock until there is a positive customer demand. The specific context in which product returns are allowed for (without costs associated) is studied in the work of Prak and Rogetzer (2022).

In Equation (9), we differentiate two scenarios; one in which a demand hit occurs in period t , with probability $m(y)$, and one in which no such hit occurs. This is why $m(y)$ appears in the expression.

Additionally, we define the discounted cost over an infinite horizon for $\gamma \in (0, 1)$ as

$$V(x, y) := \lim_{t \rightarrow \infty} V_t(x, y).$$

The limit above exists because the policy that only clears backorders already has a finite discounted cost over an infinite horizon, and so any optimal policy will also have a finite cost for $\gamma \in (0, 1)$.

Example 7 Figure 3 illustrates the value function $V_t(x, y)$ of ChemEx for different values of t for $y = 1$ and $y = 7$, with $V_0(x, y) := 0$ by convention. We let $p = 9$, $h = 1$, $L = 0$, and $\gamma = 1$. We use the parameter estimates from Example 4 to characterize the distributions.

6. Structure of optimal ordering policies

The structure of optimal ordering policies can be investigated via properties of the value functions $V_t(x, y)$. To express such properties it will be convenient to define the finite differences

$$\Delta V_t(x, y) = V_t(x + 1, y) - V_t(x, y) \quad \text{and} \quad \Delta V(x, y) = V(x + 1, y) - V(x, y).$$

Lemma 1 The value functions $V_t(x, y)$ and $V(x, y)$ have the following properties:




Figure 3: Illustration of the value function $V_t(x, y)$ for $y = 1$ and $y = 7$.

(i) $\Delta V_t(x + 1, y) \geq \Delta V_t(x, y)$ for all $y \in \mathbb{N}$ and $t \in \mathbb{N}$, i.e., $V_t(x, y)$ is convex in x .

(ii) $\Delta V(x + 1, y) \geq \Delta V(x, y)$ for all $y \in \mathbb{N}$, i.e., $V(x, y)$ is convex in x .

The proof is in Appendix A.2. From Lemma 1 it follows that the optimal ordering policy is a state-dependent base-stock policy where the state is the time since the last demand occurrence. We will let $a_t^*(x, y)$ denote the optimal action at time t in state (x, y) .

Theorem 1 *There is an optimal state-dependent base-stock policy to place replenishment orders. For the finite horizon problem, this policy places orders as*

$$a_t^*(x, y) = \max(S_t(y), x),$$

and for the infinite horizon discounted cost criterion this policy places orders as

$$a^*(x, y) = \max(S(y), x)$$

where $S_t(y)$ and $S(y)$ are optimal base-stock levels when the last demand occurrence was y periods ago (in period t).

Theorem 1 is not a surprising result, and the proof is standard (see Appendix A.1), but it is a new result for the case of compound renewal demand.

A natural question is how optimal base-stock levels $S_t(y)$ behave as a function of y . One might expect that if the time between demand occurrences has an increasing hazard rate, i.e., if $m(y)$ is an increasing function, then $S_t(y)$ will increase in y . Thus, as demand becomes imminent, one may expect that optimal policies stock more. One way to prove this would be to show that $V_t(x, y)$ is submodular, i.e. $\Delta V_t(x, y + 1) \leq \Delta V_t(x, y)$. As it turns out, the value function has




Figure 4: Counterexample of the submodularity of the value function.

both submodular and supermodular terms (the inequality that defines submodularity is reversed in supermodularity). With the following counterexample we provide evidence that the value function is not in general submodular even for rhythmic demand.

Counterexample 1 Consider the following instance: The random variables T_i follow a discrete Weibull distribution, with the average time between demand hits $\bar{T} = 4$ and the coefficient of variation of the time between demand hits $CV_T = 0.2$. The random variables H_i follow a shifted negative binomial distribution with average size $\bar{H} = 3$ and standard deviation $\sigma_H = 1.5$ (thus $CV_H = 0.5$), i.e., $H \sim A + 1$ where A follows a negative binomial distribution with $\bar{A} = 2$ and $\sigma_A = 1.5$ ($CV_A = 0.75$). The leadtime L equals 1 period, and the penalty and holding costs are 4 and 1 respectively. Figure 4 illustrates the value function for $t = 2$, for different values of y . It is clear from the figure that $V_2(x, y)$ is not submodular for $x = 4$ and $y = 4$.

Moreover, one may expect the reverse result to hold when the hazard rate of time between occurrences is decreasing. Indeed, the direct cost function $c(x, y)$ is supermodular when $m(y)$ decreases in y . Nevertheless we can characterize how the optimal base-stock levels behave as a function of the time since the last demand occurrence by exploiting the structure of the discrete compound renewal demand process, regardless of whether $m(y)$ is increasing, decreasing, or fluctuates.

Theorem 2 There exist state-dependent optimal base-stock levels $S_t(y)$ and $S(y)$ that satisfy

$$(i) S_t(y) \leq S_{t+1}(y + 1)$$

$$(ii) S(y) \leq S(y + 1) \text{ provided } Y_1 = 1$$

for all $y \in \mathbb{N}$ and $t \in \mathbb{N}$.

Note that Theorem 2 asserts the existence of optimal non-decreasing base-stock levels, but does not exclude the existence of optimal non-increasing base-stock levels. For $m(y)$ decreasing, optimal non-increasing base-stock levels

usually do exist. However, any base-stock level $S_{t+1}(y) \leq S_t(1)$ can never be attained; rather $S_t(1)$ will be attained. As such the base-stock levels $S_t(y) = S_t(1)$ give the exact same decisions and are thus also optimal. For the infinite horizon discounted cost criterion we also require that $Y_1 = 1$ so that optimal base-stock levels that decrease in y cannot be attained. Any algorithm that searches for good/optimal base-stock levels can thus constrain the search space to increasing base-stock levels only.

When demand intervals are evenly spaced, it is intuitive to delay replenishment orders until demand becomes more likely to occur. Theorem 2 formalizes this intuition. This intuition also forms the basis of the proof in Appendix A.4.

In practice, Theorem 2 limits the search efforts for appropriate base-stock levels to non-decreasing ones. This is particularly interesting when the state space is large. Additionally, it implies that an order policy which replenishes the inventory immediately after a demand hit is not a smart policy in this case and it can thus be improved. The myopic policy, which is defined in Section 7.1, is informed by this result.

To compute the optimal policy in practice, we can –for example– use the well-known Value Iteration algorithm (see for example Tijms 2003). In this case, when the discount factor $\gamma \in (0, 1)$, we calculate $V_t(x, y)$ for every state (x, y) until $|V_{t+1}(x, y) - V_t(x, y)| < \varepsilon$ for all (x, y) . When $\gamma = 1$, we calculate $M_t = \max_{(x,y)}\{V_t(x, y) - V_{t-1}(x, y)\}$ and $m_t = \min_{(x,y)}\{V_t(x, y) - V_{t-1}(x, y)\}$. We increase t until $(M_t - m_t) < \varepsilon m_t$, because then we know that $|(M_t + m_t)/2 - g^*| < \varepsilon m_t$ so that g^* is given by $(M_t + m_t)/2$ within a relative precision of at least ε . We use $\varepsilon = 10^{-3}$.

Example 8 *Setting $p = 9$, $h = 1$, $L = 0$, and $\gamma = 1$, we can calculate $V_t(x, y)$ of ChemEx for increasing values of t (using the parameter estimates from Example 4). When $t = 78$, we calculate $M_{78} \approx 2.3105$ and $m_{78} \approx 2.3082$. $M_{78} - m_{78} \approx 0.0023 \geq 10^{-3}m_{78} \approx 0.0023$. Next, we increase t to 79: $M_{79} \approx 2.3104$ and $m_{79} \approx 2.3082$. $M_{79} - m_{79} \approx 0.0022 < 10^{-3}m_{79} \approx 0.0023$ and we conclude that the Value Iteration algorithm has converged. The resulting optimal decisions are illustrated in Figure 5.*

7. Fast Heuristics

The computation of optimal base-stock levels requires solving the Bellman equations or another involved algorithm such as the one of Zipkin (1989). Such algorithms require iterative optimization. In this section, we present two myopic heuristics that only require the solution of one or several newsvendor problems. The first heuristics applies to rhythmic demand, where the hazard rate is increasing in the time since the last demand, while the second applies to the case where the hazard rate is decreasing in the time since the last demand. These heuristics apply also when one is interested in the average cost criterion.

7.1. Myopic policy

The myopic policy sets base-stock levels to minimize the immediate cost function $c(x, y)$ (see Eq. 8). This does not lead to optimal base-stock levels in general. However, when demand is rhythmic, the myopic policy inherits the property of optimal policies that base-stock levels are non-decreasing in the time since the last demand occurrence. The proof is in Appendix A.5.

Proposition 2 *If demand is rhythmic, i.e. when $m(y)$ is increasing in y , then*

(i) $c(x, y)$ is submodular.

(ii) the myopic policy that sets base-stock levels as

$$S^M(y) := \operatorname{argmin}_x c(x, y) = \inf \left\{ x \mid \mathbb{P}(D_{t,t+L}(y) \leq x) \geq \frac{p}{p+h} \right\}$$

satisfies $S^M(y) \leq S^M(y+1)$ for all $y \in \mathbb{N}$.

The myopic policy is a new policy: it is simpler than the optimal policy from Section 6, but it also uses information on the time since the last demand y myopically. In Section 8 we show that the myopic policy performs quite well when demand is rhythmic. For the special case that demand is rhythmic and hit sizes are deterministic with size 1 the myopic policy is optimal and the model reduces to that of Schultz (1989).

7.2. Stationary base-stock policy

The converse of the first part of Proposition 2 also holds: When $m(y)$ is decreasing in y , then $c(x, y)$ is supermodular.

In this case, there is an optimal stationary base-stock policy:

Proposition 3 *If $m(y)$ is decreasing in y , then there exists an average optimal stationary base-stock policy with base-stock level:*

$$S^S := \operatorname{argmin}_x \left(\sum_{y=1}^{\infty} c(x, y) \lim_{t \rightarrow \infty} \mathbb{P}(Y_t = y) \right) = \inf \left\{ x \mid \lim_{t \rightarrow \infty} \mathbb{P}(D_{t,t+L} \leq x) \geq \frac{p}{p+h} \right\}.$$

The proof is in Appendix A.6. It relies on combining Theorem 2 and the observation that an unconstrained minimizer $x^*(y)$ of $\sum_{y=1}^{\infty} c(x(y), y) \lim_{t \rightarrow \infty} \mathbb{P}(Y_t = y)$ will be decreasing due to the supermodularity of $c(x, y)$ when $m(y)$ is decreasing. The assumption that there exists a strictly increasing state dependent base-stock policy then leads to a contradiction.

Note that $\lim_{t \rightarrow \infty} \mathbb{P}(D_{t,t+L} \leq x)$ can be determined using Proposition 1 part (iv). Stationary base-stock policies are not a new concept, but the computation of the optimal stationary base-stock level is involved for compound renewal demand and requires the use of Proposition 1 part (iv).

Example 9 *Figure 5 illustrates the three order policies: the optimal policy (see Section 6), the myopic policy (see Section 7.1), and the stationary policy (see Section 7.2), considering the parameter estimates from Example 4 and parameters $L = 0$, $p = 9$, $h = 1$, and $\gamma = 1$. The stationary policy proposes a constant base-stock level (equal to 4). The optimal and myopic policy propose non-decreasing base-stock levels: The base-stock level is zero up until $y = 6$, and increases from $y = 7$ onward. The myopic policy consistently requires a base-stock level that is at least the level of the optimal policy. The stationary policy requires higher base-stock levels than the optimal policy in early stages (until and including $y = 6$) and lower ones when y is further increased.*




Figure 5: Illustration of the different order policies for ChemEx

Parameter	Values
$\mathbb{E}[T]$	4, 6, 8, 10
CV_T	0.2, 0.4, 0.6, 0.8
$\mathbb{E}[H]$	3, 5, 10
CV_H	0.75, 1.25
Leadtime	0, 1, 2
Penalty cost	4, 9, 19, 49

Table 3: Parameter values for the controlled experiment

8. Numerical study

In this numerical study, we compare the performance of the heuristic policies from Section 7 with that of the optimal policy in a controlled numerical experiment. The inventory systems are characterized by the distributions of the time between demand occurrences and the hit sizes, the inventory cost parameters, and the leadtime. Table 3 contains the parameter values that are used in a full factorial, with $4^3 \times 3^2 \times 2 = 1152$ scenarios.

The time between demand hits follow a discrete Weibull distribution (see Appendix B.1), parametrized by the expected value $\mathbb{E}[T]$ and the coefficient of variation CV_T . Although a discrete Weibull distribution is not commonly used in literature, it has three main advantages in the context of our work: First, a discrete Weibull distribution has a closed form expression for the hazard rate (see Appendix B.1). Second, the discrete Weibull distribution is versatile, similar to its continuous counterpart, as it can model an increasing ($\beta > 1$), decreasing ($\beta < 1$), and a constant hazard rate ($\beta = 1$, simplifying to the geometric distribution). Third, the choice of this distribution to model the time between demand occurrences is also supported by the company data used later in this paper. We refer the reader to Section 9.1.3 for more details.

The hit sizes follow a negative binomial distribution, where the support is shifted to strictly positive integers. That is, $H \sim X + 1$, with $X \sim NB(r, p)$. Consequently, $\mathbb{E}[H] = \mathbb{E}[X] + 1$ and $\mathbf{Var}[H] = \mathbf{Var}[X]$ (see also Appendix B.2). Table 6 in Section 9.1.3 shows that the use of this distribution is supported by real company demand data.

We analyze three leadtime settings, where leadtimes equal 0, 1, and 2 periods respectively. The penalty costs can take values 4, 9, 19, and 49, and the holding costs are kept constant at $h = 1$. Doing so, we obtain interpretable and realistic values for the cycle service level target of 80%, 90%, 95%, and 98% respectively. We consider a discount factor $\gamma = 1$ and we calculate the average per period cost for each method.

We assess the performance of the heuristics, and measure the solution quality by comparing their inventory cost

Parameter name	Parameter value	Myopic policy (%)		Stationary policy (%)	
		Mean	Max	Mean	Max
Expected time between demands $\mathbb{E}[T]$	4	1.65	20.16	5.1	39.32
	6	3.25	30.33	7.31	39.89
	8	4.58	36.81	8.51	40.71
	10	5.77	41.64	9.42	40.21
Coefficient of variation CV_T	0.2	12.60	41.65	18.42	40.71
	0.4	2.43	8.77	8.58	20.96
	0.6	0.23	1.28	2.30	9.52
	0.8	0.00	0.04	0.41	3.02
Expected hit size $\mathbb{E}[H]$	3	3.71	41.64	7.90	40.71
	5	3.83	41.29	7.54	40.41
	10	3.90	39.81	7.36	40.32
Coefficient of variation CV_H	0.75	2.59	25.03	9.51	40.71
	1.25	5.04	41.64	5.70	28.12
Penalty costs p	4	4.18	30.58	8.45	39.33
	9	4.43	41.29	8.95	40.71
	19	3.77	41.64	7.44	39.11
	49	2.87	30.10	5.57	30.31
Leadtime L	0	6.09	41.64	10.00	40.71
	1	3.51	29.15	7.69	39.33
	2	1.84	20.49	5.12	34.61
Total		3.81		7.60	

Table 4: Optimality gap of the heuristic policies (1152 scenarios).

performance with the optimal policy. We compute the optimality gap to the optimal solution as:

$$Gap_{policy} = 100 \times (g_{policy} - g^*) / g^*,$$

where g^* is the optimal cost per period, as defined in Section 6. The cost-rate of a given policy, g_{policy} , is determined exactly by Value Iteration; see section 8.2 of Puterman (2005).

The results for the 1152 instances are summarized in Table 4.

Table 4 shows that the optimality gap of both the myopic and the stationary policy increases when the expected time between demands, $\mathbb{E}[T]$, increases, that is, when demand becomes more intermittent. The myopic policy does not consider the decrease in demand probability after a demand occurrence, as we consider an increasing hazard rate and stochastic demand sizes in our experiment. It consequently ignores that on average it will take a longer time after a potential demand before a new demand is likely to happen. The optimal base-stock level is always less than or equal to the base-stock level of the myopic policy: after a demand occurrence, the leftover stock could be higher than desired by the respective policy. The optimal policy considers this possibility when determining the required order-up-to level, but the myopic policy does not.

The stationary policy, by contrast, does not consider the changing hazard rate and keeps a constant base-stock level. As a result, it overstocks immediately after a demand has occurred, and stocks too little when a demand becomes more likely. This effect becomes more pronounced when the average time between demand occurrences increases.

When the coefficient of variation of T is small, the optimality gaps are at their largest values. This is intuitive, as a low coefficient of variation indicates that demand occurs at regular intervals, with only limited deviations, i.e., demand is more rhythmic. As such, demand occurrences are more predictable in terms of the time since the last demand and including information on the renewal process and time since the last demand becomes more valuable.

By contrast, when the coefficient of variation of T increases, i.e., when demand arrivals become less predictable, both the myopic and the stationary policy perform close to optimal. In fact, we observe the three policies behave almost identically, with optimality gaps of on average less than 1%.

We can also see from Table 4 that when the leadtime increases, the performance gap between both the myopic and stationary policy with the optimal policy decreases. The rhythm in demand plays a larger role if we consider the renewal demand over a relatively short period of time. Over a longer horizon, the rhythm averages out and the exact timings of the renewal demand become less important.

When the coefficient of variation of H increases, we observe two effects. First, we observe that the optimality gap of the myopic policy increases when CV_H increases. The myopic policy consistently attains higher base-stock levels than the optimal policy, as it does not consider that the probability of observing a demand decreases right after a positive demand (under an increasing hazard rate). As this policy only protects against uncertainty in a myopic way, it ignores that leftover inventory might be kept for multiple periods following the demand. When the uncertainty in the demand size increases, the myopic policy will stock more, potentially resulting in larger leftover inventory. Second, uncertainty in the demand process originates from two sources: variability in the renewal process and variability in the hit sizes H . The stationary policy can consider the variability originating from the variability in H , but not that of the renewal process. Consequently, when H becomes more variable (as indicated by an increasing value of CV_H), a larger part of the total variability will be caused by the variability in hit sizes relative to the renewal process.

Overall, we can see that both heuristic policies are outperformed in all scenarios. In addition, we see that the myopic state-dependent policy consistently outperforms the stationary policy. In some scenarios, it even performs close to optimal.

In conclusion, from this experiment, we observe that incorporating information on the demand rhythm and on the time since the last demand occurrence in the order decisions can greatly improve the inventory cost performance. Even when this information is only included in a myopic way, we can already observe large gains compared to when a stationary policy is applied. Nonetheless, the myopic policy is sub-optimal in a number of scenario settings.

9. Real-world applications

The results in Section 8 show that our approach holds promise to improve inventory decisions for rhythmic intermittent demand. In real life applications however, the demand process is unknown and parameters of a hypothesized process must be estimated from data.

In this section we benchmark our approach to other approaches with real data from different industries. This allows us to answer the following questions throughout the following subsections: (i) Among the approaches to make decisions in real life setting, which one(s) work best and under what circumstance(s), (ii) how well do heuristics perform relative to our approach, both in mean and in distribution, and (iii) to what extent do the results generalize across different industries?

Sections 9.1 – 9.3 discuss the application of the proposed approach to the full data set of the Specialty Chemicals Company, while Section 9.4 briefly discusses replication studies in other industries, and Section 9.5 discusses practical

	Demand intervals		Demand sizes	
	Mean	CV	Mean	CV
Min	1.01	0.10	1.00	0.00
25%ile	3.53	0.68	1.76	0.34
Mean	6.16	1.12	6.89	0.49
75%ile	8.15	1.41	8.70	0.64
Max	15.00	3.23	144.80	1.48

Table 5: Descriptive statistics of SCC data.

implications.

9.1. Data description

The demand process is unknown and must be estimated from data. We use data from the Specialty Chemicals Company (SCC), whose operations motivated our research. The illustrative example used throughout this paper is also drawn from this data set. The data set contains demand records for 865 items and covers 105 periods. The descriptive statistics are presented in Table 5. We include both leading and trailing zeros of the series. We split the data into a training set and a test set, where the first 50% of the data (rounded up to the nearest integer, i.e., 53 periods) is designated as training data to initialize different methods. We choose a 50-50 split of the data to ensure a sufficiently large amount of positive demands available for both training *and* testing. Given the intermittent nature of demand, using a too small test set can make it difficult to draw any general conclusions. Nonetheless, to leverage the information that becomes available as time progresses, we use an expanding window approach in Section 9.3 and we re-estimate the distribution parameters whenever a positive demand is observed.

9.1.1. Correlation of H and T

We empirically evaluate the correlation between the time between hits and the hit sizes for the SCC items, using the observations in the training data and removing the censored first and last interval. As in Example 3, we calculate the Pearson correlation coefficient ρ and evaluate whether it differs significantly from zero ($H_0 : \rho = 0$). We reject H_0 for 50 items at the 0.05 level, indicating a correlation that is significantly different from 0 for 5.8% of the items. Out of those, 36 have a positive and 14 a negative correlation coefficient. Thus at the aggregate level, there is no evidence of positive correlation between demand hit sizes and the time between hits in this data. Therefore the assumption of independence between both variables is justified.

9.1.2. Prevalence of rhythmic demand

We investigate the prevalence of an increasing hazard rate in our data, as an increasing hazard rate indicates rhythmic intermittent demand. We therefore fit a discrete Weibull distribution (see Appendix B.1) to the observations of the demand intervals in the first half of the data for each series, using MLE (as done in Example 4). Note that the first and last inter-demand intervals are right-censored. We identify 505 out of 865 items for which $\hat{\beta} > 1$. Note that this is conditional on the Weibull distribution being an appropriate distribution to model the items' demand inter-arrival pattern. We also report the results of significance tests where $H_0 : \beta \leq 1$ (analogous to Example 4). For 287 series we reject H_0 at the 0.1 significance level (see also Table C.14 in Appendix C.4).

Data set	Series	Discrete Weibull	Mix. of binomials	Negative binomial	Poisson
Time between hits	296	243	47	6	0
Hit sizes	296		175	113	8

Table 6: Number of SCC series for which a certain discrete distribution is selected in the training data based on the negative log-likelihood.

Data set	Distribution	Series	\hat{p}			
			< 0.01	(0.01, 0.05]	(0.05, 0.1]	> 0.1
Time between hits		296	3	5	4	284
	Discrete Weibull	249	0	4	4	241
	Mixture of binomials	47	3	1	0	43
Hit sizes		296	34	25	12	225
	Mixture of binomials	175	16	6	7	146
	Negative binomial	121	18	19	5	79

Table 7: GoF results using Monte Carlo simulation. The number of series for different ranges of p-values are presented for the considered distributions.

In addition, we observe that it is not always possible to characterize the distribution of the time between demands using a discrete Weibull distribution. For example, for a series with demand intervals $\{10, 10, 11\}$ or $\{1, 1, 1, 1, 1\}$ the demand pattern is clearly rhythmic (almost deterministic), but the shape parameter of the discrete Weibull distribution estimated by MLE can become infinitely large. In such cases, we fit a mixture of binomial distributions to the data, which has a finite support and can model such demand patterns. In the remainder of this numerical study, we only consider those time series where the empirical hazard rate is increasing (at the 0.1 significance level – 287 series) or where a mixture of binomial distributions is fit (9 items), as our approach should only improve performance for these series.

9.1.3. Validation of distributional assumptions

We compare for each of the series a number of discrete distributions to model the demand inter-arrival times and the hit sizes, similar to Example 4. For each demand series we select the best distribution fit based on the negative log-likelihood. Table 6 presents an overview of the results. The table suggests that the discrete Weibull distribution is a reasonable distribution to model the time between demand arrivals. When the discrete Weibull distribution is not appropriate, a (shifted) mixture of binomial distributions is often the preferred distribution. We also observe that the (shifted) negative binomial distribution is an appropriate distribution to model the hit sizes, and that the (shifted) mixture of binomial distributions is a good alternative for the majority of the series. We also validate these distributional assumptions (see Appendix B.4). Table 7 presents the p-values of this analysis. We observe that for the majority of items there is no sufficient evidence to reject the assumption that the data is generated by the respective fitted processes. Consequently, we continue using those distributions in our further analyses.

9.2. Methods

We compare the performance of three different methods with that of the optimal policy described in Section 6, which we will refer to as the *sophisticated approach* from this point on as we can no longer guarantee its optimality when parameters are estimated from data. The demand process parameters are re-estimated with MLE after each positive demand occurrence in the test data. Subsequently, we use Value Iteration to determine the state-dependent

order-up-to levels for this policy (see Section 6). We truncate the state-space at sufficiently large values and add self-transitions at the boundary. Whenever the 99.7 percentile of the fitted distribution of the hit sizes (defined by the expected value of the hit sizes plus three times the standard deviation) exceeds 100, and the hit sizes follow a shifted negative binomial distribution, we scale the distribution such that the 99.7 percentile equals 100. We approximate the required order-up-to levels using this scaled distribution. More details on this approach are in Appendix B.5. The discount factor γ is set to one.

We include the two heuristic policies from Section 7, i.e., the myopic state-dependent policy and the stationary policy, as benchmarks. We call these policies the *myopic* and *stationary 1* heuristic respectively. We also include a third heuristic that sets a stationary base-stock level using the estimated leadtime demand distribution. This leadtime demand distribution is estimated directly by using MLE to fit a negative binomial distribution (see also Turrini and Meissner 2017) and a mixture of binomial distributions (in both cases allowing for zero demand) to the per-period demand in the training set. The parameters of the leadtime demand can be obtained through the convolutions of the respective distributions. The desired base-stock level is set at the critical fractile $\left(\frac{p}{p+h}\right)$ of the obtained distribution. This heuristic is called the *stationary 2* heuristic. The *stationary 2* heuristic is most closely related to the inventory policy currently used by the Specialty Chemicals Company that motivated our research. In practice, a number of legacy systems are in place at the company that we cannot disclose in detail in this paper.

Example 10 *Figure 6 illustrates the evolution of inventory on-hand over time for the four approaches, considering $L = 0, p = 9, h = 1$ and $\gamma = 1$. The first panel shows the inventory on-hand at the beginning of the period after receiving incoming orders and satisfying outstanding backorders, but before any demand is observed in that period. The second panel shows the on-hand inventory right after demand has occurred, if a positive demand occurs in the period. We can clearly observe that both the sophisticated and myopic approach gradually increase their inventories in the anticipation of demand, whereas the stationary policies carry too much inventory when it is not needed and not enough when a demand occurs.*

9.3. Results

For every item, we investigate the performance of the different methods under different parameter settings. Since the true values for the penalty and holding costs as well as the leadtimes are not known for all items, we investigate different scenarios to cover different realistic settings and evaluate sensitivities of the results. We set the holding cost equal to one, and let the penalty cost take different values: $p \in \{4, 9, 19, 49\}$. This coincides with enforcing different non-stockout probability constraints that are often used in practice (80%, 90%, 95%, and 98% respectively). Additionally, we vary the leadtime: $L \in \{0, 1, 2\}$.

We determine the order-up-to levels required by each method in each state and use those to make order decisions. We initialize the inventory model as follows. At the beginning of the first period of the test set, the inventory on-hand is set equal to the base-stock level required by the respective methods in this state. Next, we run through every period of the test data. We observe the state (i.e., the inventory position and the time since the last demand) and make an order decision. Next, we receive the orders placed one leadtime ago. We satisfy outstanding backorders and

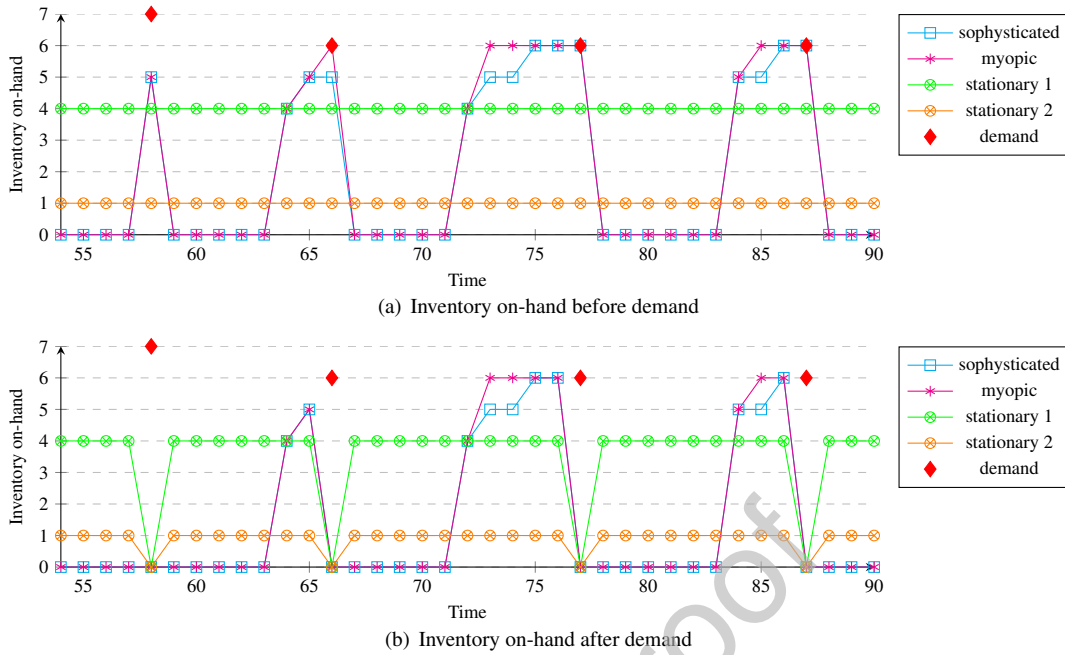


Figure 6: Illustration of the evolution of on-hand inventory under different order policies for ChemEx. The red diamond indicates a positive demand occurrence.

	Penalty costs				Leadtime		
	4	9	19	49	0	1	2
Myopic	2.41	1.55	2.35	3.19	2.81	2.19	2.13
Stationary 1	-0.08	3.76	1.88	-0.51	1.51	1.87	0.41
Stationary 2	0.60	9.38	18.29	38.68	9.30	17.62	23.30

Table 8: Average percentage difference with the cost performance of the sophisticated method for different values of the leadtime and penalty costs for the SCC.

observe demand for the current period. Finally, we register leftover inventory and backorders and we account for the corresponding costs. We only start to account for backorder and holding costs after the first $L + 1$ periods in the test data, to initialise the inventory systems.

Tables 8, 9, and 10 report the performance gap of the different heuristics relative to the approach in which the sophisticated method is used to set state-dependent base-stock levels. A positive difference implies lower costs for the sophisticated approach. Note that we can no longer speak about optimality gaps in this case; since the parameters of the demand process are not known, we cannot guarantee optimality of the sophisticated approach. Figure 7 presents boxplots of the performance differences. Table 8 stratifies the results by the backorder penalty (p) and the leadtime (L). Table 9 stratifies the results by average time between demand hits (using the quartiles). Finally Table 10 stratifies the results by empirical coefficient of variation of the times between demand occurrences.

9.3.1. Best approach under different circumstances

On average, the sophisticated method outperforms the other approaches. Nonetheless, we can identify differences depending on different circumstances. The performance gap between the sophisticated approach and the myopic policy increases when the leadtime is shorter (see Table 8). This is in line with the results from Section 8. For

	Myopic		Stationary 1		Stationary 2	
	avg	max	avg	max	avg	max
	2.38	504.73	1.26	270.00	16.74	448.22
$\bar{T} \leq Q_1(4.08)$ (75 items)	1.46	125.31	-1.11	66.43	20.45	285.71
$Q_1 < \bar{T} \leq Q_2(5.89)$ (75 items)	2.18	116.86	0.38	115.33	16.19	302.14
$Q_2 < \bar{T} \leq Q_3(8.32)$ (72 items)	3.33	504.73	-0.80	95.53	10.79	385.22
$\bar{T} > Q_3$ (74 items)	2.58	224.62	6.58	270.00	19.32	448.22

Table 9: Average and maximum percentage difference with the cost performance of the sophisticated method for the SCC data for increasing average time between positive demands \bar{T} . The values of the respective quartiles are presented between brackets.

the two stationary policies, however, we see an unexpected effect of the leadtime relative to the results from the controlled experiment in Section 8 (see Tables 4 and 8); There is no clear directional impact for the performance of the stationary 1 policy. The performance gap with the stationary 2 policy increases when the leadtime is longer. This suggests that the sophisticated approach may overestimate the degree of rhythmicity in the demand process, and this overestimation is tempered when the leadtime increases.

When the average time between demand occurrences (\bar{T}) increases, the performance gap between the sophisticated method and the myopic policy increases (see Table 9). This is in line with the results from Section 8. For the stationary 1 policy, the results are more mixed. Nonetheless, the largest performance gap in terms of average time between arrivals is observed for $\bar{T} > Q_3$. For the stationary 2 policy, we observe the largest performance gaps for both the lowest and highest values of \bar{T} . We believe this happens because the SCC data is rather intermittent, with $Q_1 = 4.08$. Thus, relatively few data point are available to estimate the distribution of T when \bar{T} is large. The stationary policy 2 has fewer parameters to estimate from data and in the face of data scarcity this seems to be an advantage. However, when the average time between demand hits further increases, both the sophisticated and stationary 2 approach suffer from the lack of data and capturing the rhythmic behavior, even imperfectly, becomes increasingly important.

When the coefficient of variation of the time between demand occurrences (CV_T) is lower, the performance gap between the sophisticated method and the three heuristics tends to increase (see Table 10). This is in line with the results from Section 8.

The stationary 1 policy performs well on average. And even the stationary 2 policy performs well, especially in low service level regimes (see Table 8). In practice, we observe that it is thus not consistently beneficial to use advanced approaches. For example, sufficient historical demand data is required for good performance of approaches that need estimations of many parameters. There seems to be a trade-off between the value of accurately specifying the demand process and data availability.

9.3.2. Performance in mean and distribution

The averages reported in Tables 8, 9, and 10 do not convey how robust the performance of the different approaches is across SKUs. This robustness is demonstrated in the boxplots in Figure 7. From this plot it is clear that the heuristic policies do not have consistent performance relative to the sophisticated approach: in some scenarios they outperform the sophisticated approach, but equally often they perform much worse. Especially the stationary 2 heuristic, and to a lesser extent the stationary 1 heuristic, have quite large interquartile ranges (long boxes) and the long whiskers in Figure 7. This means that those heuristics overperform on some SKUs and underperform on others to achieve a net

	Myopic		Stationary 1		Stationary 2	
	avg	max	avg	max	avg	max
	2.19	504.73	1.26	270.00	16.74	448.22
$CV_T \leq 0.2$ (2 items)	0	0	-0.02	0.14	-4.76	2.63
$0.2 < CV_T \leq 0.4$ (16 items)	6.70	504.73	34.83	196.00	64.10	448.22
$0.4 < CV_T \leq 0.6$ (97 items)	3.83	357.56	7.49	270.00	26.93	385.22
$0.6 < CV_T \leq 0.8$ (78 items)	0.96	104.73	-1.76	82.61	10.41	217.54
$0.8 < CV_T \leq 1$ (47 items)	1.51	178.35	-8.22	63.27	3.70	268.13
$CV_T \leq 1$ (240 items)	2.60	504.73	3.16	270.00	19.22	448.250
$CV_T > 1$ (56 items)	1.41	65.00	-6.88	54.02	6.10	217.74

Table 10: Average and maximum percentage difference with the cost performance of the sophisticated method for the SCC data for increasing CV_T .




Figure 7: Boxplots of the performance gaps per policy (blue – myopic, orange – stationary 1, green – stationary 2) for the SCC data.

negative effect. The implicit assumption here is that all SKUs are equally important. The myopic policy, by contrast, has a more robust performance relative to the sophisticated approach across the different SKUs.

9.4. Generalization across industries

For the SCC, the results illustrate that on average cost savings can be generated by applying the sophisticated method instead of the heuristic policies. Especially when the coefficient of variation of the time between hits (CV_T) is lower than 0.6, we observe improved performance for the sophisticated method relative to the benchmarks in Table 10. To ensure that our findings also generalize to other industries and to investigate potential differences between performance across industries, we validate the approach on real data from other industries that are known in the literature: the UK Royal Air Force spare parts demand data (see for example Syntetos et al. 2009, Teunter and Duncan 2009); the Car Parts data (Snyder et al. 2012), the Auto data (e.g., Pennings et al. 2017), and the M5 data (Makridakis et al. 2022) that contains Walmart retail demand data. Detailed information on the data sets can be found in Appendix C.

Similar to the SCC data, we split the data into a training set (50%) and a test set (50%) and we select those series with at least 4 demand occurrences in the training data so that all methods are initialized with sufficient data. Additionally, we only select those series that have at least 3 demand occurrences in the test set to reduce the risk of selecting obsolete items. As for the SCC data, the demand process parameters are re-estimated with MLE after each positive demand occurrence, except for some long time series in the M5 data set where they are updated only every 10 demand occurrences. The evaluation set-up and process are analogous to before.

The detailed results of these replication studies can be found in Appendix C. They confirm the validity of the

most important conclusions from the SCC case study: When the coefficient of variation of the time between hits (CV_T) is small, the sophisticated method performs well. In contrast, when the variability of the inter-demand intervals increases, the stationary 1 policy outperforms alternative approaches. Moreover, when the average time between demand hits T decreases, it is better to use a simpler approach, such as the myopic or stationary 1 policy.

9.5. Practical Implications

From the analyses in the previous sections we can derive a number of practical guidelines for companies. First, we observe there is value in correctly specifying the demand process by characterizing the distributions of the time between demand hits and of the hit sizes separately. Especially when the penalty costs are high, directly estimating the demand distribution and using this for inventory decisions using a stationary policy, results in higher costs.

Second, when inter-demand arrival times demonstrate low variability around the average ($CV_T \leq 0.6$), it becomes increasingly important to collect and use information on the demand rhythm and the time since the last demand occurrence in the decision making, and to specify an advanced demand model to determine order quantities. When, by contrast, demand arrivals are highly variable, a simple stationary policy is more appropriate.

Third, using a sophisticated approach that makes use of the demand process can result in lower costs when the time between consecutive demands is long, on the condition that this process can be properly specified (i.e., if sufficient data is available or if it is known). When the demand is not very intermittent (or *fast intermittent*), it is beneficial to implement a less complicated policy such as the myopic and stationary 1 policy. Their performance is in this case consistently close to the sophisticated approach, but their implementation is much simpler.

Fourth, although there is value in the proposed sophisticated approach, it should not be applied without caution. For example, the approach could be subject to overfitting, especially when the size of the training data is small. Besides, the approach implicitly assumes that the demand rhythm from the past continues in the future. This is not necessarily the case as the demand process can change over time. It is therefore important to not only continuously update or re-estimate the distribution parameters when new observations become available, but also to closely monitor the development of the demand process. In particular, when additional external information on the demand (generating) process becomes available (such as for example product phase-outs or client contract terminations), the estimates from the past could no longer hold in the future.

10. Concluding remarks

Intermittent demand complicates efficient inventory management, as there are two separate sources of uncertainty present. On the one hand, there is uncertainty as to when the next demand will occur, and on the other hand there is also uncertainty about the size of the demand when it occurs. Existing approaches for demand forecasting and inventory management often implicitly assume that the time between demand occurrences is Markovian, i.e., the probability of observing a demand in a certain period is assumed independent of the time since the last demand occurrence. However, data from practice, indicate that the times between demand events are not necessarily memoryless. Non-Markovian times between demand occurrences lead to *rhythm* in the demand process.

This paper is the outcome of a joint research project with a company in the specialty chemicals industry. This company faces issues in the efficient and effective inventory management of its items. It observes intermittent demand for many SKUs. Moreover, real company data demonstrate a certain *rhythm* in the demand occurrences. In such cases, it intuitively makes sense to postpone order placements until demand becomes more likely to occur. Nonetheless, the company currently does not consider the intermittence of the demand nor the rhythmic pattern in the demand occurrences in its inventory decision making. Instead, their current practice uses standard techniques that are built-in in commonly used planning software to make forecasts and manage their inventories.

We propose a demand model that accommodates such rhythmic intermittent demand. We model demand as a discrete compound renewal process. This paper contributes to both the theory and practice of inventory management. The main theoretical contributions of this paper include characterization of optimal order policies, i.e., we show that the optimal policy is a state-dependent base-stock policy, where the state is the time since the last demand observation. Furthermore, we show that there exist optimal base-stock levels that are non-decreasing in the time since the last demand arrival. We demonstrate the implementation of such optimal policy in a numerically stable manner.

We propose a myopic state-dependent policy and two stationary policies that can be easily computed and compare performance relative to optimal policies on real data. There are opportunities for inventory cost performance improvements when incorporating information on the demand rhythm and the time since the last demand in the order decisions. In particular, and not surprisingly, when inter-demand arrival times have low variability relative to the mean, it is important to collect and use information on the demand rhythm and the time since the last demand occurrence in the decision making. By contrast, when demand arrivals are highly variable, it is sufficient to implement (simpler) stationary policies. In addition, as the mean time between demand hits increases, the value of using the time since the last demand hit for inventory decision making increases.

The current research can be extended in a number of directions. First, the benchmarks used in this study include solely parametric approaches. Nonetheless, non-parametric approaches or machine learning could also be used to identify (additional) structure in the demand process. Second, our approach assumes stationary parameters of the compound renewal demand process. Thus, non-stationary phenomena such as SKUs that phase in and out of the assortment cannot easily be accommodated. Thirdly, since Maximum Likelihood estimators can perform poorly in small samples, and they could be prone to overfitting, it may be worthwhile to investigate whether a different choice of parameter estimator could improve the results. Lastly, our study only considers the time since the last observed demand hit as an indicator for future demand. When there is a single customer, it is clear why the time since the last demand is useful information. When there are several customers, it may be beneficial to model demand as a superposition of several compound renewal processes in which the time since the last demand of each individual customer is tracked.

References

Adan, I., van Eenige, M., Resing, J., 1995. Fitting discrete distributions on the first two moments. *Probability in the engineering and informational sciences* 9, 623–632.

- Almalki, S.J., Nadarajah, S., 2014. Modifications of the weibull distribution: A review. *Reliability Engineering & System Safety* 124, 32–55.
- Arrow, K.J., Harris, T., Marschak, J., 1951. Optimal inventory policy. *Econometrica: Journal of the Econometric Society* , 250–272.
- Arts, J., Basten, R., Van Houtum, G.J., 2016. Repairable stocking and expediting in a fluctuating demand environment: Optimal policy and heuristics. *Operations Research* 64, 1285–1301.
- Aviv, Y., Federgruen, A., 1997. Stochastic inventory models with limited production capacity and periodically varying parameters. *Probability in the Engineering and Informational Sciences* 11, 107–135.
- Babai, M.Z., Chen, H., Syntetos, A.A., Lengu, D., 2021. A compound-poisson bayesian approach for spare parts inventory forecasting. *International Journal of Production Economics* 232.
- Babai, M.Z., Dallery, Y., Boubaker, S., Kalai, R., 2019. A new method to forecast intermittent demand in the presence of inventory obsolescence. *International Journal of Production Economics* 209, 30–41.
- Beyer, D., Sethi, S.P., 1997. Average cost optimality in inventory models with markovian demands. *Journal of Optimization Theory and Applications* 92, 497–526.
- Bi, S., He, L., Teo, C.P., 2023. Taming the long tail: The gambler's fallacy in intermittent demand management. *Manufacturing & Service Operations Management* , 1692–1710.
- Chen, F., Song, J.S., 2001. Optimal policies for multiechelon inventory problems with markov-modulated demand. *Operations Research* 49, 226–234.
- Cheng, F., Sethi, S.P., 1999. Optimality of state-dependent (s, s) policies in inventory models with markov-modulated demand and lost sales. *Production and Operations Management* 8.
- Croston, J.D., 1972. Forecasting and stock control for intermittent demands. *Operational Research Quarterly* 23, 289–303.
- Dizbin, N.M., Tan, B., 2020. Optimal control of production-inventory systems with correlated demand inter-arrival and processing times. *International Journal of Production Economics* 228, 107692.
- Fildes, R., Nikolopoulos, K., Crone, S.F., Syntetos, A.A., 2008. Forecasting and operational research: a review. *Journal of the Operational Research Society* 59, 1150–1172.
- Gutierrez, R.S., Solis, A.O., Mukhopadhyay, S., 2008. Lumpy demand forecasting using neural networks. *International Journal of Production Economics* 111, 409–420.
- Hall, P., Van Keilegom, I., 2005. Testing for monotone increasing hazard rate. *Annals of Statistics* 33, 1109–1137.
- Hyndman, R., Koehler, A.B., Ord, J.K., Snyder, R.D., 2008. *Forecasting with exponential smoothing: the state space approach*. Springer Science & Business Media.
- Ibe, O.C., 2013. 12 - markovian arrival processes, in: Ibe, O.C. (Ed.), *Markov Processes for Stochastic Modeling (Second Edition)*. second edition ed.. Elsevier, Oxford, pp. 349–376.
- van Jaarsveld, W., Arts, J., 2024. Projected inventory-level policies for lost sales inventory systems: Asymptotic optimality in two regimes. *Operations Research* .
- Janakiraman, G., Seshadri, S., Shanthikumar, J.G., 2007. A comparison of the optimal costs of two canonical inventory systems. *Operations Research* 55, 866–875.
- Janssen, F., Heuts, R., De Kok, T., 1998. On the (r, s, q) inventory model when demand is modelled as a compound bernoulli process. *European Journal of Operational Research* 104, 423–436.
- Johnston, F., Boylan, J.E., Shale, E.A., 2003. An examination of the size of orders from customers, their characterisation and the implications for inventory control of slow moving items. *Journal of the Operational Research Society* 54, 833–837.
- Johnston, F.R., Boylan, J.E., 1996. Forecasting for items with intermittent demand. *Journal of the Operational Research Society* 47, 113–121.
- Kaplan, E.L., Meier, P., 1958. Nonparametric estimation from incomplete observations. *Journal of the American statistical*

- association 53, 457–481.
- Kapuściński, R., Tayur, S., 1998. A capacitated production-inventory model with periodic demand. *Operations Research* 46, 899–911.
- Karlin, S., 1960. Dynamic inventory policy with varying stochastic demands. *Management Science* 6, 231–258.
- Koole, G., 2006. Monotonicity in Markov reward and decision chains: Theory and applications. volume 1. Now Publishers.
- Kourntzes, N., Athanasopoulos, G., 2021. Elucidate structure in intermittent demand series. *European Journal of Operational Research* 288, 141–152.
- Levén, E., Segerstedt, A., 2004. Inventory control with a modified Croston procedure and Erlang distribution. *International Journal of Production Economics* 90, 361–367.
- Lindsey, M., Pavur, R., 2013. Assessing a modification to croston’s method to incorporate a seasonal component, in: *Advances in Business and Management Forecasting*. Emerald Group Publishing Limited, pp. 185–195.
- Lolli, F., Gamberini, R., Regattieri, A., Balugani, E., Gatos, T., Gucci, S., 2017. Single-hidden layer neural networks for forecasting intermittent demand. *International Journal of Production Economics* 183, 116–128.
- Makridakis, S., Spiliotis, E., Assimakopoulos, V., 2022. M5 accuracy competition: Results, findings and conclusions. *International Journal of Forecasting* 38, 1346–1364. Special Issue: M5 competition.
- Molenaers, A., Baets, H., Pintelon, L., Waeyenbergh, G., 2012. Criticality classification of spare parts: A case study. *International Journal of Production Economics* 140, 570–578.
- Mukhopadhyay, S., Solis, A.O., Gutierrez, R.S., 2012. The accuracy of non-traditional versus traditional methods of forecasting lumpy demand. *Journal of Forecasting* 31, 721–735.
- Nakagawa, T., Osaki, S., 1975. The discrete weibull distribution. *IEEE transactions on reliability* 24, 300–301.
- Nasr, W.W., 2022. Inventory systems with stochastic and batch demand: computational approaches. *Annals of Operations Research* 309, 163–187.
- Pennings, C.L.P., van Dalen, J., van der Laan, E.A., 2017. Exploiting elapsed time for managing intermittent demand for spare parts. *European Journal of Operational Research* 258, 958–969.
- Perera, S., Janakiraman, G., Niu, S.C., 2017. Optimality of (s, s) inventory policies under renewal demand and general cost structures. *Production and Operations Management* 27, 368–383.
- Pinçe, C., Frenk, J.B., Dekker, R., 2015. The role of contract expirations in service parts management. *Production and Operations Management* 24, 1580–1597.
- Porras, E., Dekker, R., 2008. An inventory control system for spare parts at a refinery: An empirical comparison of different re-order point methods. *European Journal of Operational Research* 184, 101–132.
- Prak, D., Rogetzer, P., 2022. Timing intermittent demand with time-varying order-up-to levels. *European Journal of Operational Research* 303, 1126–1136.
- Puterman, M.L., 2005. Markov decision processes: discrete stochastic dynamic programming. John Wiley & Sons.
- Ross, S.M., 1996. Stochastic processes. volume 2. Wiley, New York.
- Schultz, C.R., 1987. Forecasting and inventory control for sporadic demand under periodic review. *Journal of the Operational Research Society* 38, 453–458.
- Schultz, C.R., 1989. Replenishment delays for expensive slow-moving items. *Management Science* 35, 1454–1462.
- Sethi, S.P., Cheng, F., 1997. Optimality of (s, s) policies in inventory models with markovian demand. *Operations Research* 45, 931–939.
- Shaked, M., Shanthikumar, J.G., 2007. Stochastic orders. Springer Science & Business Media.
- Smith, M.A., Dekker, R., 1997. On the $(s-1, s)$ stock model for renewal demand processes: Poisson’s poison. *Probability in the*

- Engineering and informational Sciences 11, 375–386.
- Snyder, R.D., Ord, J.K., Beaumont, A., 2012. Forecasting the intermittent demand for slow-moving inventories: A modelling approach. *International Journal of Forecasting* 28, 485–496.
- Song, J.S., Zipkin, P., 1993. Inventory control in a fluctuating demand environment. *Operations Research* 41, 351–370.
- Song, J.S., Zipkin, P.H., 1996. Inventory control with information about supply conditions. *Management science* 42, 1409–1419.
- Syntetos, A.A., Babai, M.Z., Dallery, Y., Teunter, R., 2009. Periodic control of intermittent demand items: theory and empirical analysis. *Journal of the Operational Research Society* 60, 611–618.
- Syntetos, A.A., Babai, M.Z., Luo, S., 2015. Forecasting of compound erlang demand. *Journal of the Operational Research Society* 66, 2061–2074.
- Syntetos, A.A., Babai, Z., Boylan, J.E., Kolassa, S., Nikolopoulos, K., 2016a. Supply chain forecasting: Theory, practice, their gap and the future. *European Journal of Operational Research* 252, 1–26.
- Syntetos, A.A., Boylan, J.E., 2005. The accuracy of intermittent demand estimates. *International Journal of Forecasting* 21, 303–314.
- Syntetos, A.A., Boylan, J.E., Croston, J., 2005. On the categorization of demand patterns. *Journal of the operational research society* 56, 495–503.
- Syntetos, A.A., Teunter, R.H., Babai, M.Z., Transchel, S., 2016b. On the benefits of delayed ordering. *European Journal of Operational Research* 248, 963–970.
- Teunter, R., Duncan, L., 2009. Forecasting intermittent demand: A comparative study. *The Journal of the Operational Research Society* 60, 321–329.
- Teunter, R.H., Syntetos, A.A., Babai, M.Z., 2010. Determining order-up-to levels under periodic review for compound binomial (intermittent) demand. *European Journal of Operational Research* 203, 619–624.
- Teunter, R.H., Syntetos, A.A., Babai, M.Z., 2011. Intermittent demand: Linking forecasting to inventory obsolescence. *European Journal of Operational Research* 214, 606–615.
- Tijms, H.C., 2003. *A first course in stochastic models*. John Wiley and sons.
- Turkmen, A.C., Januschowski, T., Wang, Y., Cengil, A.T., 2020. Intermittent demand forecasting with renewal processes doi:<https://doi.org/10.48550/arXiv.2010.01550>. working paper.
- Turrini, L., Meissner, J., 2017. Spare parts inventory management: new evidence from distribution fitting. *European Journal of Operational Research* .
- Van Houtum, G.J., 2006. Multiechelon production/inventory systems: optimal policies, heuristics, and algorithms, in: *Models, methods, and applications for innovative decision making*. Informs, pp. 163–199.
- Willemain, T.R., Smart, C.N., Schwarz, H.F., 2004. A new approach to forecasting intermittent demand for service parts inventories. *International Journal of Forecasting* 20, 375–387.
- Zhang, N., Kannan, K., Shanthikumar, G., 2021. Nudging a slow-moving high-margin product in a supply chain with constrained capacity. *Production and Operations Management* 30, 11–27.
- Zipkin, P., 1989. Critical number policies for inventory models with periodic data. *Management Science* 35, 71–80.
- Zipkin, P.H., 2000. *Foundations of inventory management*. McGraw-Hill, New York.

Appendix A. Proofs

Appendix A.1. Proof of Proposition 1

Proposition 1 *The distributions of demand and number of hits over a leadtime satisfy:*

$$(i) \mathbb{P}(N_{t,t+L}(y) \geq n) = \sum_{j=y}^{L+y} \mathbb{P}\left(\sum_{i=2}^n T_i \leq L + y - j\right) \frac{\mathbb{P}(T_1=j)}{\mathbb{P}(T_1 \geq y)}$$

$$(ii) \lim_{t \rightarrow \infty} \mathbb{P}(N_{t,t+L} \geq n) = \sum_{y=1}^{\infty} \mathbb{P}(N_{t,t+L}(y) \geq n) \frac{\mathbb{P}(T \geq y)}{\mathbb{E}[T]}$$

$$(iii) \mathbb{P}(D_{t,t+L}(y) = x) = \sum_{n=0}^{L+1} \mathbb{P}(N_{t,t+L}(y) = n) \mathbb{P}\left(\sum_{i=1}^n H_i = x\right)$$

$$(iv) \lim_{t \rightarrow \infty} \mathbb{P}(D_{t,t+L} = x) = \sum_{y=1}^{\infty} \mathbb{P}(D_{t,t+L}(y) = x) \frac{\mathbb{P}(T \geq y)}{\mathbb{E}[T]}$$

Proof of Proposition 1 The proof is a direct proof. For part (i) we have:

$$\begin{aligned} \mathbb{P}(N_{t,t+L}(y) \geq n) &= \mathbb{P}(N_{t,t+L} \geq n \mid Y_t = y) && \text{(By definition)} \\ &= \mathbb{P}(N_{t+L} - N_{t-1} \geq n \mid Y_t = y) && \text{(By definition)} \\ &= \mathbb{P}(N_{L+1} \geq n \mid Y_1 = y) && \text{(Renewals are i.i.d.)} \\ &= \mathbb{P}\left(\sum_{i=1}^n T_i \leq L + y \mid T_1 \geq y\right) && \text{(By definition)} \\ &= \mathbb{P}\left(T_1 + \sum_{i=2}^n T_i \leq L + y \mid T_1 \geq y\right) \\ &= \sum_{j=y}^{L+y} \mathbb{P}(T_1 = j \mid T_1 \geq y) \mathbb{P}\left(\sum_{i=2}^n T_i \leq L + y - j\right) \\ &= \sum_{j=y}^{L+y} \frac{\mathbb{P}(T_1 = j)}{\mathbb{P}(T_1 \geq y)} \mathbb{P}\left(\sum_{i=2}^n T_i \leq L + y - j\right) && \text{(By definition)} \end{aligned}$$

For part (ii) we have:

$$\begin{aligned} \lim_{t \rightarrow \infty} \mathbb{P}(N_{t,t+L} \geq n) &= \lim_{t \rightarrow \infty} \sum_y \mathbb{P}(N_{t,t+L}(y) \geq n) \mathbb{P}(Y_t = y) \\ &= \sum_y \mathbb{P}(N_{t,t+L}(y) \geq n) \lim_{t \rightarrow \infty} \mathbb{P}(Y_t = y) && \text{(By (i))} \\ &= \sum_y \mathbb{P}(N_{t,t+L}(y) \geq n) \frac{\mathbb{P}(T \geq y)}{\mathbb{E}[T]} && \text{(By Prop. 3.4.5 in Ross 1996)} \end{aligned}$$

The final step follows from the fact that Y_t follows a residual service time distribution, as we consider a renewal process with inter-arrival times T_i . Results for the residual service time in such case exist in literature, for example in Ross (1996).

For part (iii) we have:

$$\begin{aligned}
 \mathbb{P}(D_{t,t+L}(y) = x) &= \mathbb{P}(D_{t,t+L} = x \mid Y_t = y) && \text{(By definition)} \\
 &= \mathbb{P}\left(\sum_{i=1+N_{t-1}}^{N_{t+L}} H_i = x \mid Y_t = y\right) && \text{(By definition)} \\
 &= \mathbb{P}\left(\sum_{i=1}^{N_{t+L}} H_i = x \mid Y_t = y\right) \\
 &= \sum_{n=0}^{L+1} \mathbb{P}\left(\sum_{i=1}^{N_{t+L}} H_i = x \mid Y_t = y, N_{t,t+L} = n\right) \mathbb{P}(N_{t,t+L} = n \mid Y_t = y) \\
 &= \sum_{n=0}^{L+1} \mathbb{P}\left(\sum_{i=1}^n H_i = x\right) \mathbb{P}(N_{t,t+L} = n \mid Y_t = y) \\
 &= \sum_{n=0}^{L+1} \mathbb{P}\left(\sum_{i=1}^n H_i = x\right) \mathbb{P}(N_{t,t+L}(y) = n) && \text{(By definition)}
 \end{aligned}$$

By convention, we have $\sum_a^b = 0$ for $a > b$.

The proof of (iv) is analogous to (ii). ■

Appendix A.2. Proof of Lemma 1

Lemma 1 *The value functions $V_t(x, y)$ and $V(x, y)$ have the following properties:*

- (i) $\Delta V_t(x+1, y) \geq \Delta V_t(x, y)$ for all $y \in \mathbb{N}$ and $t \in \mathbb{N}$, i.e., $V_t(x, y)$ is convex in x .
- (ii) $\Delta V(x+1, y) \geq \Delta V(x, y)$ for all $y \in \mathbb{N}$, i.e., $V(x, y)$ is convex in x .

Proof of Lemma 1 We first prove convexity of $c(x, y)$. Note that $(x - D(y))^+ = (D(y) - x)^+ + x - D(y)$, $c(y, x) = (p + h) \cdot \mathbb{E}[(D(y) - x)^+] + h \cdot (x - \mathbb{E}[D(y)])$. Since $(D(y) - x)^+$ is convex in x , and expectation preserves convexity, $\mathbb{E}[(D(y) - x)^+]$ is convex in x , and thus $c(x, y)$ is convex in x .

We now demonstrate by induction that $V_t(x, y)$ is convex in x . Observe that $V_0(x, y) := 0$ is convex. Suppose V_{n-1} is convex. $V_t(x, y)$ is convex in x because convex combination and translation preserve convexity, and

$$\min_{a \geq x} f(a) = \begin{cases} f(a^*) & \text{for } x \leq a^* \\ f(x) & \text{otherwise,} \end{cases}$$

where $a^* = \operatorname{argmin} f(a)$.

Part (ii) of the lemma follows directly from part (i) and the fact that limits of convex functions are convex. ■

Appendix A.3. Proof of Theorem 1

Theorem 1 *There is an optimal state-dependent base-stock policy to place replenishment orders. For the finite horizon problem, this policy places orders as*

$$a_t^*(x, y) = \max(S_t(y), x),$$

and for the infinite horizon discounted cost criterion this policy places orders as

$$a^*(x, y) = \max(S(y), x)$$

where $S_t(y)$ and $S(y)$ are optimal base-stock levels when the last demand occurrence was y periods ago (in period t).

Proof of Theorem 1

Let $S_t(y) = \operatorname{argmin}_{a \in \mathbb{Z}} \{c(a, y) + \gamma(1 - m(y))V_{t-1}(a, y + 1) + \gamma m(y)\mathbb{E}[V_{t-1}(a - H, 1)]\}$. If $x \leq S_t(y)$, then the minimizing $a \geq x$ is at $a = S_t(y)$. Alternatively, if $x \geq S_t(y)$, then the minimizing a is at $a = x$. Thus, a state-dependent base-stock policy with base-stock level $S_t(y)$ is optimal for period t for every y .

Let $S(y) = \operatorname{argmin}_{a \in \mathbb{Z}} \{c(a, y) + \gamma(1 - m(y))V(a, y + 1) + \gamma m(y)\mathbb{E}[V(a - H, 1)]\}$. If $x \leq S(y)$, then the minimizing $a \geq x$ is at $a = S(y)$. Alternatively, if $x \geq S(y)$, then the minimizing a is at $a = x$. Thus, for every y , there is an optimal base-stock policy with base-stock level $S(y)$ for period t . ■

Appendix A.4. Proof of Theorem 2

Theorem 2 *There exist state-dependent optimal base-stock levels $S_t(y)$ and $S(y)$ that satisfy*

$$(i) S_t(y) \leq S_{t+1}(y + 1)$$

$$(ii) S(y) \leq S(y + 1) \text{ provided } Y_1 = 1$$

for all $y \in \mathbb{N}$ and $t \in \mathbb{N}$.

Proof of Theorem 2

We use the structural properties of our problem to show that there exist optimal base-stock levels $S_t(y) \leq S_{t+1}(y + 1)$. When the process is in state (x, y) at time t , it can only move to a state $y + 1$ at $t + 1$ if no demand occurs in period t . Otherwise, if demand does occur, the process ends up in state $y = 1$. When no demand occurs (and since we do not allow for items in stock to be discarded), the inventory position x cannot be reduced to any number $s < x$ going from period t to the next. As such, in period $t + 1$, any inventory position $S'_{t+1}(y + 1) < S_t(y)$ can never be attained. That is, the inventory position after ordering is greater than $S_t(y)$ for any base-stock level $S'_{t+1}(y + 1) \leq S_t(y)$.

For the infinite horizon discounted cost criterion, there may exist optimal base-stock levels $S(y) > S(y + 1)$, but the attainment argument in (i) still applies when $Y_1 = 1$. Thus, when $Y_1 = 1$, there also exist optimal base-stock levels that satisfy $S(y) \leq S(y + 1)$. ■

Appendix A.5. Proof of Proposition 2

Proposition 2 *If $m(y)$ is increasing in y , then*

(i) $c(x, y)$ is submodular.

(ii) the myopic policy that sets base-stock levels as

$$S^M(y) := \operatorname{argmin}_x c(x, y) = \inf \left\{ x \mid \mathbb{P}(D_{t,t+L}(y) \leq x) \geq \frac{p}{p+h} \right\}$$

satisfies $S^M(y) \leq S^M(y+1)$ for all $y \in \mathbb{N}$.

Proof of Proposition 2

In this proof we use stochastic orders as in the following definition.

Definition 1 (1.A.1 Definition, Shaked and Shanthikumar (2007)) *If X and Y are discrete random variables taking on values in \mathbb{N} , then X is stochastically smaller than Y in the usual stochastic order, denoted by $X \leq_{st} Y$, if, and only if $\mathbb{P}(X > i) \leq \mathbb{P}(Y > i)$.*

Define the random variables T_{iy} such that $\mathbb{P}(T_{iy} = x) := \mathbb{P}(T_i - y + 1 = x \mid T_i \geq y)$. Note that T_{iy} denotes the residual time until the i -th demand given that the $(i-1)$ -th demand occurred y periods ago.

Lemma 2 *If the hazard rate $m(y)$ is increasing in y , then the random variables T_{iy} are stochastically decreasing in y , i.e., $\mathbb{P}(T_{iy} \leq x) \leq \mathbb{P}(T_{i(y+1)} \leq x)$ for all $x, y \in \mathbb{N}$.*

Proof of Lemma 2

$$\begin{aligned} \mathbb{P}(T_{iy} \leq x) - \mathbb{P}(T_{i(y+1)} \leq x) &= \mathbb{P}(T_i \leq x + y - 1 \mid T_i \geq y) - \mathbb{P}(T_i \leq x + y \mid T_i \geq y + 1) \\ &= \mathbb{P}(T_i > x + y \mid T_i \geq y + 1) - \mathbb{P}(T_i > x + y - 1 \mid T_i \geq y) \\ &= \frac{\mathbb{P}(T_i > x + y)}{\mathbb{P}(T_i \geq y + 1)} - \frac{\mathbb{P}(T_i > x + y - 1)}{\mathbb{P}(T_i \geq y)} \\ &= \prod_{i=1}^x (1 - m(y + i)) - \prod_{i=0}^{x-1} (1 - m(y + i)) \\ &= (1 - m(y + x) - 1 + m(y)) \prod_{i=1}^{x-1} (1 - m(y + i)) \\ &= (m(y) - m(x + y)) \prod_{i=1}^{x-1} (1 - m(y + i)) \end{aligned}$$

Observe that $m(y) - m(x + y) \leq 0$ as the hazard rate is increasing and that $\prod_{i=1}^{x-1} (1 - m(y + i)) \geq 0$. Consequently, $\mathbb{P}(T_{iy} \leq x) - \mathbb{P}(T_{i(y+1)} \leq x) \leq 0$ and thus $\mathbb{P}(T_{iy} \leq x) \leq \mathbb{P}(T_{i(y+1)} \leq x)$ for all $x \in \mathbb{N}$. ■

Lemma 3 *If the hazard rate $m(y)$ is increasing in y , then the random variables $N_{t,t+L}(y)$ are stochastically increasing in y , i.e., $\mathbb{P}(N_{t,t+L}(y) \leq n) \geq \mathbb{P}(N_{t,t+L}(y+1) \leq n), \forall n, y \in \mathbb{N}$.*

Proof of Lemma 3

$$\begin{aligned}
 \mathbb{P}(N_{t,t+L}(y) \leq n) &= \mathbb{P}(N_{t,t+L} \leq n \mid Y_t = y) = \mathbb{P}(N_{t+L} - N_t \leq n \mid Y_t = y) && \text{(By definition)} \\
 &= \mathbb{P}(N_{L+1} \leq n \mid Y_1 = y) \\
 &= \mathbb{P}\left(\sum_{i=1}^n T_i \geq L + y \mid T_1 \geq y\right) && \text{(By definition)} \\
 &= \mathbb{P}\left(T_1 + \sum_{i=2}^n T_i \geq L + y \mid T_1 \geq y\right) \\
 &= \mathbb{P}\left(T_{1|y} + \sum_{i=2}^n T_i \geq L + 1\right) \\
 &\geq \mathbb{P}\left(T_{1|y+1} + \sum_{i=2}^n T_i \geq L + 1\right) && \text{(Lemma 2)} \\
 &= \mathbb{P}(N_{t,t+L}(y+1) \leq n) && \text{(By definition)}
 \end{aligned}$$

■

Lemma 4 *If the hazard rate $m(y)$ is increasing in y , then the random variables $D_{t,t+L}(y)$ are stochastically increasing in y , i.e., $\mathbb{P}(D_{t,t+L}(y) \leq x) \geq \mathbb{P}(D_{t,t+L}(y+1) \leq x), \forall x, y \in \mathbb{N}$.*

Proof of Lemma 4 By definition:

$$\mathbb{P}(D_{t,t+L}(y) \leq x) - \mathbb{P}(D_{t,t+L}(y+1) \leq x) = \mathbb{P}\left(\sum_{i=1}^{N_{t,t+L}(y)} H_i \leq x\right) - \mathbb{P}\left(\sum_{i=1}^{N_{t,t+L}(y+1)} H_i \leq x\right).$$

As we know from Lemma 3, $N_{t,t+L}(y) \leq_{st} N_{t,t+L}(y+1)$ when $m(y)$ is increasing in y . Consequently,

$$\mathbb{P}\left(\sum_{i=1}^{N_{t,t+L}(y)} H_i \leq x\right) \geq \mathbb{P}\left(\sum_{i=1}^{N_{t,t+L}(y+1)} H_i \leq x\right).$$

This implies that $\mathbb{P}(D_{t,t+L}(y) \leq x) - \mathbb{P}(D_{t,t+L}(y+1) \leq x) \geq 0$, which proves our lemma. ■

Definition 2 (Definition 6.3, Koole (2006)) *A function $f : \mathbb{N}_0^m \rightarrow \mathbb{R}$ is submodular, i.e., $f \in \text{Sub}(i, j)$ if*

$$f(x) + f(x + e_i + e_j) \leq f(x + e_i) + f(x + e_j)$$

for all x and $1 \leq i \leq j \leq m$.

We now proceed with the proof of Proposition 2(i):

$$\begin{aligned}
 c(x+1, y) - c(x, y) &= p \cdot \mathbb{E}[(D_{t,t+L}(y) - x - 1)^+] + h \cdot \mathbb{E}[(x + 1 - D_{t,t+L}(y))^+] \\
 &\quad - (p \cdot \mathbb{E}[(D_{t,t+L}(y) - x)^+] + h \cdot \mathbb{E}[(x - D_{t,t+L}(y))^+]) \\
 &= -p \cdot \mathbb{P}(D_{t,t+L}(y) \geq x + 1) + h \cdot \mathbb{P}(D_{t,t+L}(y) \leq x) \\
 &\geq -p \cdot \mathbb{P}(D_{t,t+L}(y+1) \geq x + 1) + h \cdot \mathbb{P}(D_{t,t+L}(y+1) \leq x) \quad (\text{Lemma 4}) \\
 &= c(x+1, y+1) - c(x, y+1)
 \end{aligned}$$

Following Definition 2, $c(x, y)$ is submodular when $m(y)$ is increasing.

For part (ii) observe that the minimization of $c(x, y)$ for given y is a newsvendor problem. The minimizers of this newsvendor problem satisfy $S^M(y+1) \geq S^M(y)$ when $m(y)$ is increasing due to submodularity (see, e.g., Lemma 4.7.1 in Puterman 2005). ■

Appendix A.6. Proof of Proposition 3

Proposition 3 *If $m(y)$ is decreasing in y , then there exists an average optimal stationary base-stock policy with base-stock level:*

$$S^S := \operatorname{argmin}_x \left(\sum_{y=1}^{\infty} c(x, y) \lim_{t \rightarrow \infty} \mathbb{P}(Y_t = y) \right) = \inf \left\{ x \mid \lim_{t \rightarrow \infty} \mathbb{P}(D_{t,t+L} \leq x) \geq \frac{p}{p+h} \right\}.$$

Proof of Proposition 3

Observe first that $c(x, y)$ is supermodular when $m(y)$ is decreasing in y . The proof of this is analogous to the proof of Proposition 2(i) as all inequalities reverse. By Theorem 2, there exists an optimal non-decreasing base-stock policy $S(y)$. (The same argument easily extends to the average cost criterion.) The average cost-rate of such a policy can be expressed as

$$H(S) := \sum_{y \in \mathbb{N}} c(S(y), y) f(y) \quad (\text{A.1})$$

where $f(y) \equiv \lim_{t \rightarrow \infty} \mathbb{P}(Y_t = y)$. Thus we seek a policy $S : \mathbb{N} \rightarrow \mathbb{Z}$ that minimizes (A.1) subject to the constraints that $S(y) \leq S(y+1)$ for all $y \in \mathbb{N}$. Assume that there exists an optimal solution S and an integer \tilde{y} such that $S(y^-) < S(y^+)$ for all $y^- \leq \tilde{y}$ and $y^+ > \tilde{y}$, i.e., assume that the base-stock level increases at least once from \tilde{y} to $\tilde{y} + 1$. Consider now the following base-stock policies in which we move the changes in the base-stock level left or right respectively:

$$S^l(y) = \begin{cases} S(y) & \text{if } y \neq \tilde{y} \\ S(\tilde{y} + 1) & \text{if } y = \tilde{y}, \end{cases}$$

and

$$S^r(y) = \begin{cases} S(y) & \text{if } y \neq \tilde{y} + 1 \\ S(\tilde{y}) & \text{if } y = \tilde{y} + 1, \end{cases}$$

We also define $x^*(y) \in \operatorname{argmin}_x c(x, y)$. Consider now the differences

$$H(S^l) - H(S) = f(\tilde{y})(c(S(\tilde{y} + 1), \tilde{y}) - c(S(\tilde{y}), \tilde{y})) \quad (\text{A.2})$$

$$H(S^r) - H(S) = f(\tilde{y} + 1)(c(S(\tilde{y}), \tilde{y} + 1) - c(S(\tilde{y} + 1), \tilde{y} + 1)) \quad (\text{A.3})$$

If $S(\tilde{y}) < x^*(\tilde{y})$, then $H(S^l) - H(S)$ will be non-negative by the convexity of $c(x, y)$ in x and the changepoint \tilde{y} can be moved to the left without loss of optimality. However, if $S(\tilde{y}) \geq x^*(\tilde{y})$ then we must have that $H(S^r) - H(S)$ is non-negative because $x^*(\tilde{y} + 1) \leq x^*(\tilde{y})$ by the supermodularity of $c(x, y)$. Therefore, any changepoint in base-stock level can be moved either left or right. This argument can be repeated to show that any changepoint must occur either at $\tilde{y} = 1$ or $\tilde{y} \rightarrow \infty$ which implies that there is an optimal base-stock policy that does not change.

Appendix B. Distributions

Appendix B.1. Discrete Weibull distribution

We use a discrete version of the Weibull distribution. A discrete Weibull random variable (of which the support does not include zero) is defined with the probability mass function (see also Nakagawa and Osaki 1975)

$$f(x) = q^{(x-1)^\beta} - q^{x^\beta},$$

where $q = e^{-(\alpha^{-\beta})}$, β is the shape parameter, and α the scale parameter that are conventionally used to characterize a continuous Weibull random variable. Alternatively, when the support can include zero, it is characterized as $f(x) = q^{(x)^\beta} - q^{(x+1)^\beta}$. The cumulative distribution function is defined as:

$$F(x) = \begin{cases} 1 - q^{x^\beta} & \text{for } x = 1, 2, 3, \dots \\ 0 & \text{for } x < 1. \end{cases}$$

Using a discrete Weibull distribution for the time between demand occurrences allows to formulate a parametric expression for the hazard rate:

$$m(x) = 1 - (q^{x^\beta} - q^{(x-1)^\beta}).$$

The discrete Weibull distribution, like its continuous counterpart, is versatile as it can model different types of behavior: a parameter $\beta > 1$ indicates an increasing hazard rate, whereas $\beta < 1$ represents a decreasing hazard rate. When β equals 1, the hazard rate is constant and the distribution simplifies to the geometric distribution. For more details on the different variants of the Weibull distribution, we refer an interested reader to Almalki and Nadarajah (2014).

Appendix B.2. Negative Binomial distribution

We consider a shifted version of the negative binomial distribution to model the demand hit sizes. Recall that the conventional negative binomial distribution is defined by the shape parameter $r \in (0, \infty)$ and the probability parameter $p \in (0, 1)$. We have

$$\mathbb{P}(X = x) = \binom{r+x-1}{x} p^r (1-p)^x, \quad x \in \mathbb{N}_0.$$

The parameters p and r can be obtained from the expected value $\mathbb{E}[X]$ and standard deviation σ_X of the random variable:

$$r = \frac{(\mathbb{E}[X])^2}{\sigma_X^2 - \mathbb{E}[X]} \text{ and } p = \frac{\mathbb{E}[X]}{\sigma_X^2}.$$

We shift the support of the distribution to allow only for strictly positive integers. The random variable $Y \sim X + 1$ follows a shifted negative binomial distribution where $\mathbb{E}[Y] = \mathbb{E}[X] + 1$ and $\sigma_Y = \sigma_X$.

The negative binomial distribution is closed under convolution, i.e., the sum of n independent negative binomial random variables X_i with shape r and probability p also has a negative binomial distribution with shape nr and

probability p . Define $Z_X^n = \sum_{i=1}^n X_i$ the n -th convolution of X . The probabilities $\mathbb{P}(Z_X^n = x)$ can then be computed recursively as:

$$\mathbb{P}(Z_X^n = x) = \frac{nr + x - 1}{x} (1 - p) \mathbb{P}(Z_X^n = x - 1), \quad \mathbb{P}(Z_X^n = 0) = p^{nr}, \quad x \in \mathbb{N}.$$

For the convolutions of the shifted negative binomial random variable Y , define $Z_Y^n = \sum_{i=1}^n Y_i$. Then

$$\mathbb{P}(Z_Y^n = x) = \mathbb{P}(Z_X^n = x - n), \text{ for } x \geq n.$$

Appendix B.3. Mixture of Binomial distributions

The random variable W follows a mixture of binomial distributions with parameters k , p , and q (see for example Adan et al. 1995):

$$W = \begin{cases} \text{BIN}(k, p) & \text{w.p. } q \\ \text{BIN}(k + 1, p) & \text{w.p. } (1 - q). \end{cases}$$

Let X follow a binomial distribution for which the density function is denoted by $\mathbb{P}(X = x; k, p)$, with the number of trials equal to k and success probability p . Then the probability density function of W is defined as:

$$\mathbb{P}(W = w) = q\mathbb{P}(X = w; k, p) + (1 - q)\mathbb{P}(X = w; k + 1, p).$$

This probability can be calculated recursively as:

$$\mathbb{P}(W = w) = \frac{p}{(1 - p)^w} ((k - w + 1)\mathbb{P}(W = w - 1) + (1 - q)\mathbb{P}(X = w - 1; k + 1, p)).$$

The hazard rate of W , $m_W(w)$, can be obtained recursively through the general result for any discrete random variable:

$$m_W(w) = \frac{\mathbb{P}(W = w)}{\mathbb{P}(W \geq w)} = \frac{\mathbb{P}(W = w)}{\prod_{i=1}^{w-1} (1 - m_W(i))}.$$

The convolutions of W can be obtained as follows. Define $Z_W^n = \sum_{i=1}^n W_i$ the n -th convolution of W . Then

$$\mathbb{P}(Z_W^n = w) = \sum_{i=0}^n \binom{n}{i} \mathbb{P}(X = w; nk + i, p) q^{n-i} (1 - q)^i.$$

In our work, we also consider a shifted version of the mixture of binomial distributions to allow only for strictly positive integers. The random variable $Y \sim W + 1$ follows a shifted mixture of binomial distributions where $\mathbb{E}[Y] = \mathbb{E}[W] + 1$ and $\sigma_Y = \sigma_W$. The hazard rate and convolutions of Y can be calculated from the results of W :

$$m_Y(w) = m_W(w - 1)$$

Define $Z_Y^n = \sum_{i=1}^n Y_i$. Then

$$\mathbb{P}(Z_Y^n = w) = \mathbb{P}(Z_W^n = w - n), \text{ for } w \geq n.$$

Appendix B.4. Goodness-of-fit evaluation

We assess the validity of the distributional assumptions. However, our current setting presents multiple complexities for using standard Goodness-of-Fit (GoF) tests: (1) the time between hits can be censored, (2) we evaluate GoF of discrete distributions, whereas standard tests evaluate continuous ones, (3) these distributions (such as the discrete Weibull distribution) are not commonly included in existing goodness-of-fit packages, (4) the available sample sizes are small, preventing the use of tests such as the chi-square test, and (5) we are estimating distribution parameters of the hypothesized distribution from the data which causes problems for the validity of tests such as the Kolmogorov-Smirnov test. Existing GoF tests that deal with some of these aspects typically are not able to handle all at once. Therefore, we developed a customized testing procedure that simultaneously deals with all the issues raised above.

Our null hypothesis is that the data is generated by a member of some parametric family, F_0 . We compare this *hypothesized* distribution with the empirical distribution. The empirical distribution is determined using the *Product-Limit Estimator* (PLE), also known as the *Kaplan-Meier Estimator* (Kaplan and Meier 1958), which is a non-parametric estimator of the distribution function: $\hat{F}_n(x)$. The test statistic is calculated as $D = \sup_{x \in \mathbb{N}} [F_0(x) - \hat{F}_n(x)]$. When F_0 is a continuous function and there is no censoring, this test statistic corresponds to that of the well-known *Kolmogorov-Smirnov* test. However, in the case of a discrete distribution with censoring, the distribution of this test statistic is not known, not even asymptotically.

Appendix B.4.1. Procedure

We rely on Monte Carlo simulation to estimate the distribution of the test statistic under the null hypothesis. We used the following approach for the inter-arrival times:

1. We start from the distribution indicated to be the *best fit* using the negative log-likelihood evaluation and we use the corresponding MLE parameters.
2. We obtain the product-limit estimator (PLE) from the observations of the time between hits, using the *survfit()* function from the survival package in the R statistical software.
3. We then compute the maximum difference D between the hypothesized distribution using the MLE (see Step 1), F_0 , and the empirical distribution using the PLE, \hat{F}_n : $D = \sup |F_0 - \hat{F}_n|$.
4. We generate 1000 sequences of random values for the time between hits from the selected distribution (see Step 1), until the total number of periods for each sequence is equal to or exceeds 70 periods.
5. To obtain a sequence of length 53, equal to the observed number of periods for ChemEx and also equaling length of the SCC training data, we generate a random number from a discrete uniform distribution $U(1,17)$. This number indicates the starting period of the sequence. The sequence runs for exactly 53 periods, thus potentially censoring intervals at the beginning and the end of the series.
6. For each sample path, we (1) estimate the distribution parameters using MLE, and (2) calculate the product-limit estimator (PLE).

7. We compute the test statistic for each sample path i : $D_i = \sup |F_{i,0} - \hat{F}_{i,n}|$, which is the maximum difference between the theoretical distribution $F_{i,0}$ (using the MLE parameters, obtained in the previous step) and the empirical distribution $\hat{F}_{i,n}$ (using the PLE). We thus obtain 1000 values of the test statistic D_i (one for each sample path i).
8. We evaluate how many of the 1000 sample path test statistics D_i are larger than the test statistic D , and we divide this number by 1000. The resulting value \hat{p} is an unbiased estimator for the p-value of our hypothesis test, and has standard error of at most $1/(2\sqrt{1000})$. Large \hat{p} -values indicate that we find no support to reject H_0 .

Unlike the time between hits, the observed hit sizes are not censored. Nonetheless, their analysis does suffer from most of the same limitations to the use of traditional goodness-of-fit tests. We thus perform a similar test to validate the distributions of the hit sizes. The main differences are in Step 4, where we sample a number of hit sizes equal to the number of hit sizes in the original series (training data). Step 5 is omitted, since there is no censoring. The other steps are identical to those for the analysis of the time between hits.

Appendix B.4.2. Analysis for ChemEx

The time between hits sequence of ChemEx is $\{1, 6, 6, 8, 10, 9, 11, 3\}$, where the first and last inter-arrival interval are censored. The hit size sequence is $\{3, 5, 5, 5, 5, 6, 6\}$, with no censoring. The hypothesized distribution F_0 is obtained using the MLE from Example 4. We obtain the PLE for both sequences as in Tables B.11 and B.12. Note that the PLE provides the empirical survival function \hat{S}_n , and the empirical distribution function is $\hat{F}_n = 1 - \hat{S}_n$.

To compute the test statistic $D = \sup |F_0 - \hat{F}_n|$, we calculate $D^+ = \max_x [\hat{F}_n(x) - F_0(x)]$ and $D^- = \max_x [F_0(x) - \hat{F}_n(x-1)]$, with $x \in \{1 \dots 11\}$ and assuming $\hat{F}_n(0) = 0$. For the time between hits, $D^{T,+} = 0.1717$ and $D^{T,-} = 0.2188$. Thus, $D^T = 0.2188$. For the hit sizes, the test statistic $D^H = 0.5294$.

Time	Time between hits				
	At risk	Events	\hat{S}_n^T	\hat{F}_n^T	F_0^T
1	8	0	1.00	0.00	0.00
3	7	0	1.00	0.00	0.01
6	6	2	0.67	0.33	0.16
8	4	1	0.50	0.50	0.51
9	3	1	0.33	0.67	0.72
10	2	1	0.17	0.83	0.88
11	1	1	0.00	1.00	0.97

Table B.11: F_n and F^0 for the time between hits of ChemEx.

Time	Hit sizes				
	At risk	Events	\hat{S}_n^H	\hat{F}_n^H	F_0^H
3	7	1	0.86	0.14	0.06
5	6	4	0.29	0.71	0.67
6	2	2	0.00	1.00	1.00

Table B.12: F_n and F^0 for the hit sizes of ChemEx.




Figure B.8: Plot of the 1000 sample test statistics D_i for the time between hits of ChemEx.

We sample 1000 sample paths using the discrete Weibull distribution with scale parameter $\alpha \approx 8.57$ and $\beta \approx 4.87$ for the time between hits and the shifted mixture of binomial distributions for the hit sizes ($k = 4, p = 0.80, q = 10^{-5}8.12$), as determined in Example 4. For example, we sample the sequence $\{5, 9, 10, 6, 10, 6, 9, 9, 10\}$, with a total length 74 periods, for the time between hits. Next, we generate a random value from a discrete uniform distribution $U(1, 17)$. As an example, the random value equals 3. This implies the series will start in the 4th period (thus removing the first 3 periods). Since we require exactly 53 periods, our final sequence will be $\{2, 9, 10, 6, 10, 6, 9, 1\}$ and both the first and last interval are censored. For this sample path, we estimate the distribution parameters using MLE ($\hat{\alpha}_i = 8.50, \hat{\beta}_i = 5.97$) and we calculate the PLE.

We calculate the test statistic D_i^T for each of the 1000 sample paths i . A histogram of these test statistics can be found in Figure B.8. We count that 945 of the 1000 sample path test statistics D_i are larger than $D^T = 0.2188$. Thus, the corresponding estimated p-value is $\hat{p}_T = 0.945$. Consequently, we do not find evidence to reject the null hypothesis that the time between hits data is generated by a discrete Weibull distribution. Similarly, for the hit sizes, we find that $\hat{p}_H = 0.462$ and we thus do not find evidence to reject the hypothesis that the hit sizes are generated by a mixture of binomial distributions.

Appendix B.5. Approximation using distribution scaling

The state space can become too large for the Value Iteration algorithm to converge within reasonable time. Whenever the 99.7 percentile of the hit size distribution exceeds 100, and the hit sizes follow a shifted negative binomial distribution, we approximate the required order-up-to levels using a scaled distribution for the random hit sizes. This reduces the run time. In particular, we scale the distribution of H such that this percentile equals 101.

Consider the distribution of a random variable X , defined by an expected value μ and a standard deviation σ . When 99.7 percentile of the distribution, approximated by $\mu + 3\sigma$, is larger than 100, we scale the distribution of X to the distribution of \tilde{X} , defined by $\tilde{\mu}$ and $\tilde{\sigma}$, for this percentile to equal 100.

We can define

$$\tilde{\sigma} = \frac{\sigma}{a}, \tilde{\mu} = \frac{\mu}{a}$$

with a such that

$$\frac{\mu + 3\sigma}{a} = 100, \text{ i.e., } a = \frac{\mu + 3\sigma}{100}$$

When the hit sizes H are distributed as $H \sim X + 1$ and $\hat{\mu}_x + 3\hat{\sigma}_x > 100$ we first scale the distribution of H , such that the hit sizes are distributed as $\tilde{H} \sim (\tilde{X} + 1)$, where $\tilde{X} \sim (\tilde{\mu}, \tilde{\sigma})$ as defined before. We then find the required order-up-to levels using the distribution of the hit sizes \tilde{H} . Next, we multiply the resulting order-up-to levels with a , and round them to the nearest integer value, to obtain an approximation of the required order-up-to levels for H .

Appendix C. Results on different data sets

Appendix C.1. Data sets

- The **RAF** data set consists of 5000 SKUs from the UK Royal Air Force. It includes 84 months of data, and has been used before by, among others, Syntetos et al. (2009) and Teunter and Duncan (2009).
- The **Car Parts** data consists of 2674 monthly time series, previously used in, for example, Snyder et al. (2012). As suggested by Hyndman et al. (2008) we consider the 2509 series which have complete data for the 51 observation periods, removing those series with missing values.
- The **Auto** data, previously used in, for example, Pennings et al. (2017), contains monthly spare part demand data for 3000 items in the automotive industry, for 24 observation periods.
- The **M5** data set (Makridakis et al. 2022) consists of 30490 time series of daily unit sales at Walmart. We include both the training and evaluation data, as they are publicly available. This results in over 5 years of daily demand observations. We remove for all items the leading zeros of the demand series such that every series starts with a demand (see also Makridakis et al. 2022). Following the approach of Prak and Rogetzer (2022), we delete periods of 30 or more consecutive days without demand to avoid these periods being (incorrectly) considered as long inter-demand intervals. Note that different time series consequently have different lengths.

We split the data into a training set (50%) and a test set (50%) and we select those series with at least 4 demand occurrences in the training data so that all methods are initialized with sufficient data. Additionally, we only select those that have at least 3 demand occurrences in the test set to reduce the risk of selecting obsolete items. This results in respectively 2641 series for the RAF, 1142 for the Car Parts, 3000 for the Auto, and 30490 for the M5 data. The descriptive statistics for each series are presented in Table C.13. We include both leading and trailing zeros of the series (except for the M5 data, where we only include trailing zeros).

We also investigate the prevalence of increasing hazard rates in the training data and only keep those series where we reject $H_0 : \beta \leq 1$ at a 0.1 significance level or where a mixture of binomial distributions is fit (see Table C.14). Additionally, we compare for each of the series in the different data sets the discrete distributions from Section 9.1.3 to model the demand inter-arrival times and the hit sizes. An overview of the results can be found in Table C.15 and C.16. Table C.15 confirms the choice of the discrete Weibull distribution (and mixture of binomial distributions) to model the time between demand arrivals and Table C.16 confirms that the (shifted) negative binomial distribution and mixture of binomial distributions are appropriate distributions to model the hit sizes for the majority of the series.

The demand process parameters are re-estimated with MLE after each positive demand occurrence, except for some long time series in M5 data set where they are updated only every 10 demand occurrences. The evaluation set-up and process are analogous to before.

Appendix C.2. Analysis of results

We investigate the performance of the four methods. The results can be found in Table C.17, C.18, and C.19, as well as Figure C.9, all in Appendix C.

The RAF data set contains spare part demand data from the UK military. It consists of very intermittent time series that oftentimes contain only few demand observations. The performance of the myopic policy is in line with the expectations from Section 8 and 9.3: it is outperformed on average, and the performance gap increases for shorter leadtimes, larger values of \bar{T} , and lower values of CV_T (see Table C.18 and C.19). The stationary policies, however, perform well on this data. On average, the sophisticated method does not even outperform the stationary 1 policy. Nonetheless, when the $CV_T < 0.6$, we observe from Table C.19 that the stationary 2 policy is on average outperformed by the sophisticated method for the RAF data. When the CV_T increases, however, both stationary policies outperform the sophisticated approach. We note that the RAF data is known for containing items with a high inventory obsolescence risk. Hence, the demand rhythm from the training data may not necessarily continue in the test set. It is thus not consistently beneficial to use an advanced approach for this data set.

The Car Parts data consists of demand observations of car parts sales. It contains time series with rather variable demand arrivals, indicated by the absence of series with a CV_T lower than 0.4 (see Table C.19). Because the time between demand arrivals is rather variable for this data set, the value of using the sophisticated approach decreases.

The Auto data contains so called *fast intermittent* series of spare part demand (Syntetos and Boylan 2005). Both the myopic and stationary 1 policy perform very similar to the sophisticated approach for these series. Note that the average time between demands is never more than 2 periods in this data set (see Table C.13), implying that sufficient data is available for parameter estimations despite the short series lengths. The stationary 2 policy is in this case almost always outperformed by all other policies. Nonetheless, the value of using complex approaches is limited. Because the performance of the myopic and stationary 1 policy are consistently close to the sophisticated approach, but their implementation is much simpler, using either of those policies would be advisable for this data set.

Finally, the M5 data contains demand data from the retail sector. On average, none of the heuristic policies are outperformed by the sophisticated approach for this data. We do note that the maximum performance gap is attained for $\bar{T} > Q_3$, which seems to indicate that, when sufficient data is available to estimate the distribution parameters accurately, incorporating demand rhythm can become increasingly important when demand is more intermittent. The sophisticated, myopic, and stationary 1 policy result in almost identical order processes (see also Figure C.9). This could be due to the relatively high CV_T : only 56 out of 645 items have a $CV_T \leq 0.6$.

Appendix C.3. Descriptive statistics

	Series length	Demand intervals		Demand sizes	
		Mean	CV	Mean	CV
SCC (865 series)					
Min		1.01	0.10	1.00	0.00
25%ile		3.53	0.68	1.76	0.34
Mean	105	6.16	1.12	6.89	0.49
75%ile		8.15	1.41	8.70	0.64
Max		15.00	3.23	144.80	1.48
RAF (2641 series)					
Min		4.05	0.28	1.00	0.00
25%ile		6.54	0.68	1.62	0.49
Mean	84	7.78	0.79	14.92	0.77
75%ile		8.50	0.90	12.38	0.98
Max		12.00	1.36	483.56	3.45
Car Parts (1142 series)					
Min		1.33	0.42	1.00	0.00
25%ile		1.96	0.68	1.50	0.47
Mean	51	2.86	0.85	1.94	0.58
75%ile		3.47	0.96	2.12	0.67
Max		6.50	2.06	8.89	1.90
Auto (3000 series)					
Min		1.04	0.20	1.00	0.00
25%ile		1.14	0.31	2.05	0.51
Mean	24	1.31	0.42	5.42	0.63
75%ile		1.47	0.52	5.00	0.70
Max		2.00	1.02	193.75	3.75
M5 (30490 series)					
Min	49	1.00	0.04	1.00	0.00
25%ile	922	1.53	0.80	1.31	0.48
Mean	1299.37	3.02	0.89	2.47	0.59
75%ile	1728	3.78	1.00	2.40	0.68
Max	1969	13.51	2.03	159.82	10.34

Table C.13: Descriptive statistics of the data sets.

Appendix C.4. Prevalence of increasing hazard rates

Data set	Series	$\hat{\beta}_T > 1$	$\hat{\beta}_T > 1(*)$	$\hat{\beta}_T > 1(**)$	$\hat{\beta}_T > 1(***)$	Mix. of bin. dist.	Nb. retained items
SCC	865	505	287	201	57	9	296
RAF	2641	1688	394	149	2	18	412
Car Parts	1142	519	88	35	12	0	88
Auto	3000	1912	776	731	731	390	1166
M5	30 490	3718	644	323	104	1	645

Table C.14: Prevalence of increasing hazard rates in the training data, where (*) indicates significance at 0.1, (**) at 0.05, and (***) at 0.01 level.

Appendix C.5. Validation of distributional assumptions

Data set	Series	Discrete Weibull	Mix. of Binomial	Negative Binomial	Poisson
SCC	296	241	49	6	0
RAF	412	339	40	29	4
Car Parts	88	45	42	1	0
Auto	1166	7	1158	0	1
M5	645	584	51	10	0

Table C.15: Number of series for which a certain discrete distribution is selected for the time between demand hits in the training data based on the negative log-likelihood.

Data set	Series	Mix. of Binomial	Negative Binomial	Poisson
SCC	296	175	113	8
RAF	412	134	271	7
Car Parts	88	38	47	3
Auto	1166	188	941	37
M5	645	134	507	4

Table C.16: Number of series for which a discrete distribution is selected for the demand hit sizes in the training data based on the negative log-likelihood.

Appendix C.6. Results for different leadtimes and penalty costs

		Penalty costs				Leadtime		
		4	9	19	49	0	1	2
SCC	Myopic	2.41	1.55	2.35	3.19	2.81	2.19	2.13
	Stationary 1	-0.08	3.76	1.88	-0.51	1.51	1.87	0.41
	Stationary 2	0.60	9.38	18.29	38.68	9.30	17.62	23.30
RAF	Myopic	6.40	5.33	4.12	4.93	7.66	4.59	3.33
	Stationary 1	-6.36	-2.68	-3.48	-2.57	-6.51	-2.60	-2.20
	Stationary 2	-7.21	-2.51	0.48	15.00	-5.62	2.48	7.46
Car Parts	Myopic	1.05	1.12	0.32	0.55	1.88	0.45	-0.05
	Stationary 1	-3.16	-3.21	-1.67	0.53	-3.36	-1.70	-0.57
	Stationary 2	-1.20	1.98	5.07	11.34	-1.45	5.25	9.10
Auto	Myopic	-0.03	0.03	0.01	0.02	0.05	-0.02	-0.02
	Stationary 1	-0.05	-0.16	-0.06	-0.06	-0.17	-0.09	0.03
	Stationary 2	0.19	0.05	0.90	1.51	0.41	0.82	0.75
M5	Myopic	-0.01	0.00	-0.02	0.00	-0.02	0.02	-0.02
	Stationary 1	-0.41	-0.48	-0.37	-0.31	-0.67	-0.34	-0.16
	Stationary 2	-0.38	-0.41	-0.52	-0.56	-0.56	-0.39	-0.45

Table C.17: Average percentage difference with the cost performance of the sophisticated method for different values of the leadtime and penalty costs.

Appendix C.7. Results as a function of \bar{T}

		Myopic		Stationary 1		Stationary 2	
		avg	max	avg	max	avg	max
SCC		2.38	504.73	1.26	270.00	16.74	448.22
(75 items)	$\bar{T} \leq Q_1(4.08)$	1.46	125.31	-1.11	66.43	20.45	285.71
(75 items)	$Q_1 < \bar{T} \leq Q_2(5.89)$	2.18	116.86	0.38	115.00	16.19	302.14
(72 items)	$Q_2 < \bar{T} \leq Q_3(8.32)$	3.33	504.73	-0.80	95.53	10.79	385.22
(74 items)	$\bar{T} > Q_3$	2.58	224.62	6.58	270.00	19.32	448.22
RAF		5.20	922.09	-3.77	177.78	1.44	268.52
(120 items)	$\bar{T} \leq Q_1(6.54)$	2.43	175.00	-2.95	132.59	2.76	268.52
(130 items)	$Q_1 < \bar{T} \leq Q_2(7.73)$	3.84	320.79	-5.24	177.78	0.79	254.46
(59 items)	$Q_2 < \bar{T} \leq Q_3(8.71)$	6.93	218.85	-4.33	93.85	0.11	158.27
(103 items)	$\bar{T} > Q_3$	9.14	922.09	-2.55	172.73	1.49	234.39
Car Parts		0.76	84.90	-1.88	250.00	4.30	167.71
(22 items)	$\bar{T} \leq Q_1(2.03)$	0.17	9.86	-1.33	17.19	0.37	42.31
(22 items)	$Q_1 < \bar{T} \leq Q_2(3.22)$	0.04	28.52	-3.78	26.09	1.95	48.67
(25 items)	$Q_2 < \bar{T} \leq Q_3(4)$	0.84	57.89	-1.63	208.70	5.32	94.93
(19 items)	$\bar{T} > Q_3$	2.16	84.90	-0.63	250.00	10.21	167.71
Auto		0.01	27.97	-0.08	119.35	0.66	315.25
(379 items)	$\bar{T} \leq Q_1(1.04)$	0.00	0.00	-0.02	5.56	0.61	240.00
(261 items)	$Q_1 < \bar{T} \leq Q_2(1.09)$	0.01	16.67	0.03	52.13	1.30	315.25
(312 items)	$Q_2 < \bar{T} \leq Q_3(1.20)$	0.00	23.68	-0.04	80.00	-0.04	200.00
(214 items)	$\bar{T} > Q_3$	0.03	27.97	-0.36	119.35	0.98	220.00
M5		-0.01	16.81	-0.39	36.08	-0.47	54.85
(162 items)	$\bar{T} \leq Q_1(2.78)$	0.00	7.27	-0.11	14.99	-0.10	18.08
(161 items)	$Q_1 < \bar{T} \leq Q_2(4.45)$	0.03	9.33	-0.30	34.14	-0.18	34.12
(161 items)	$Q_2 < \bar{T} \leq Q_3(6.49)$	-0.04	8.90	-0.28	24.17	-0.45	32.24
(161 items)	$\bar{T} > Q_3$	-0.02	16.81	-0.88	36.08	-1.13	54.85

Table C.18: Average and maximum percentage difference with the cost performance of the sophisticated method per data set for increasing average time between positive demands \bar{T} . The values of the respective quartiles are presented between brackets.

Appendix C.8. Results as a function of CV_T

		Myopic		Stationary 1		Stationary 2	
		avg	max	avg	max	avg	max
SCC		2.38	504.73	1.26	270.00	16.74	448.22
(2 items)	$CV_T \leq 0.2$	0	0	-0.02	0.14	-4.76	2.63
(16 items)	$0.2 < CV_T \leq 0.4$	6.70	504.73	34.83	196.00	64.10	448.22
(97 items)	$0.4 < CV_T \leq 0.6$	3.83	357.56	7.49	270.00	26.93	385.22
(78 items)	$0.6 < CV_T \leq 0.8$	0.96	104.73	-1.76	82.61	10.41	217.54
(47 items)	$0.8 < CV_T \leq 1$	1.51	178.35	-8.22	63.27	3.70	268.13
(240 items)	$CV_T \leq 1$	2.60	504.73	3.16	270.00	19.22	448.250
(56 items)	$CV_T > 1$	1.41	65.00	-6.88	54.02	6.10	217.74
RAF		5.20	922.09	-3.77	177.78	1.44	268.52
	$CV_T \leq 0.2$	/	/	/	/	/	/
(8 items)	$0.2 < CV_T \leq 0.4$	24.27	218.85	19.13	172.73	24.99	172.73
(130 items)	$0.4 < CV_T \leq 0.6$	7.05	922.09	-0.09	177.78	5.82	268.52
(201 items)	$0.6 < CV_T \leq 0.8$	4.25	320.79	-5.05	132.59	0.11	234.39
(62 items)	$0.8 < CV_T \leq 1$	2.79	209.86	-8.75	61.57	-4.89	140.20
(401 items)	$CV_T \leq 1$	5.33	922.09	-5.53	177.78	1.68	268.52
(11 items)	$CV_T > 1$	0.28	50.00	-12.43	46.41	-7.46	148.54
Car Parts		0.76	84.90	-1.88	250.00	4.30	167.71
	$CV_T \leq 0.4$	/	/	/	/	/	/
(25 items)	$0.4 < CV_T \leq 0.6$	0.78	84.90	1.36	250.00	6.24	167.71
(47 items)	$0.6 < CV_T \leq 0.8$	0.97	57.89	-2.70	167.86	4.24	150.00
(8 items)	$0.8 < CV_T \leq 1$	-0.14	18.95	-5.71	8.01	3.56	119.28
(80 items)	$CV_T \leq 1$	0.8	84.90	1.73	250.00	4.79	167.71
(8 items)	$CV_T > 1$	0.33	14.13	-3.35	26.90	-0.68	26.09
Auto		0.01	27.97	-0.08	119.35	0.66	315.25
(379 items)	$CV_T \leq 0.2$	0.00	0.00	-0.02	5.56	0.61	240.00
(496 items)	$0.2 < CV_T \leq 0.4$	-0.01	23.68	0.01	81.97	0.81	315.25
(254 items)	$0.4 < CV_T \leq 0.6$	0.04	27.97	-0.27	119.35	0.62	220.00
(25 items)	$0.6 < CV_T \leq 0.8$	0.11	16.67	-0.42	28.57	-0.48	84.31
(10 items)	$0.8 < CV_T \leq 1$	0.04	9.09	-0.70	2.33	-0.70	10.81
(1164 items)	$CV_T \leq 1$	0.01	27.97	-0.08	119.35	0.66	315.25
(2 items)	$CV_T > 1$	0.00	0.00	-0.12	0.00	-1.80	2.13
M5		-0.01	16.81	-0.39	36.08	-0.47	54.85
(27 items)	$CV_T \leq 0.2$	0.00	0.49	0.00	0.43	0.28	12.00
(10 items)	$0.2 < CV_T \leq 0.4$	-0.08	0.27	0.00	0.53	-0.47	4.88
(19 items)	$0.4 < CV_T \leq 0.6$	-0.01	2.91	-0.11	4.01	0.33	17.50
(228 items)	$0.6 < CV_T \leq 0.8$	-0.01	3.79	-0.21	36.08	-0.32	54.85
(354 items)	$0.8 < CV_T \leq 1$	0.00	16.81	-0.56	24.91	-0.66	35.13
(638 items)	$CV_T \leq 1$	0.00	16.81	-0.39	36.08	-0.47	54.85
(7 items)	$CV_T > 1$	-0.49	8.90	-0.58	1.34	-0.32	7.04

Table C.19: Average and maximum percentage difference with the cost performance of the sophisticated method per data set for increasing CV_T .

Appendix C.9. Results in mean and in distribution

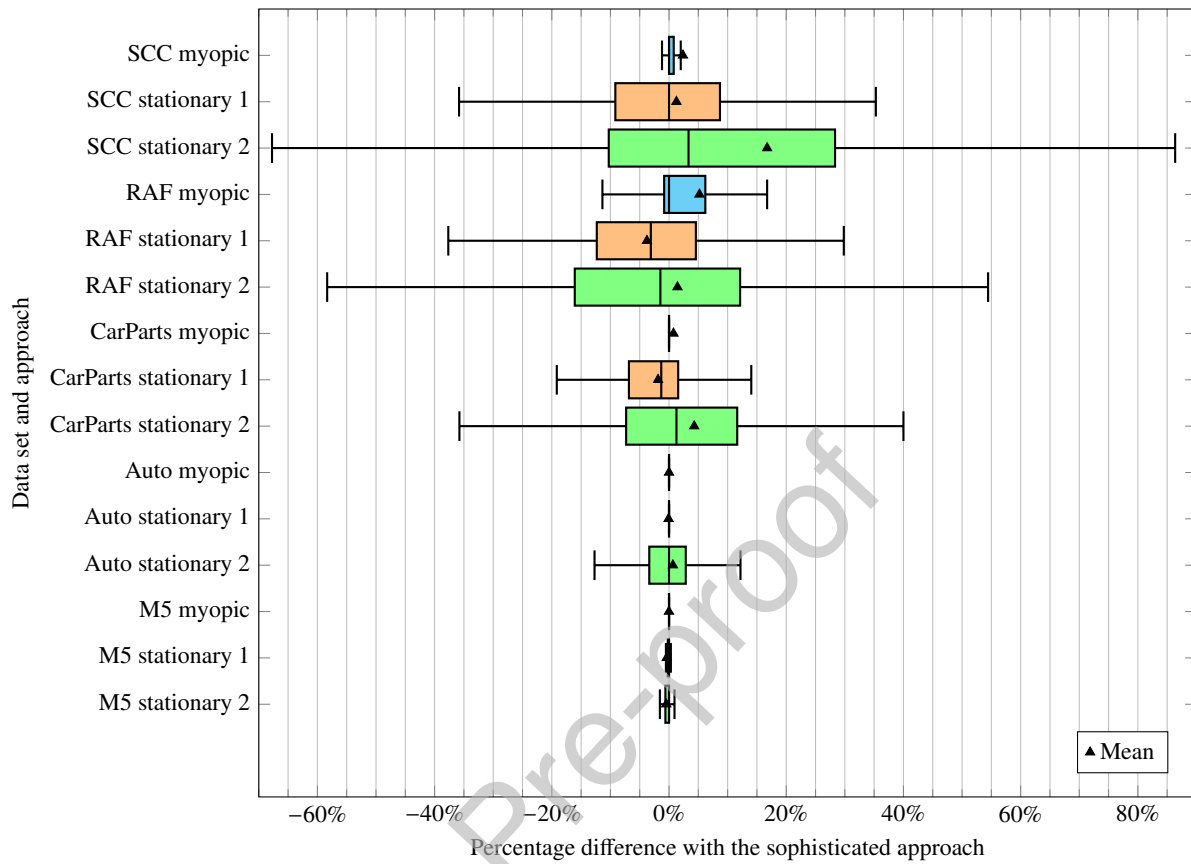


Figure C.9: Boxplots of the performance gaps per data set and policy (blue – myopic, orange – stationary 1, green – stationary 2).