New Approximation Guarantees for The Inventory Staggering Problem

Noga Alon^{*} Danny Segev[†]

Abstract

Since its inception in the mid-60s, the inventory staggering problem has been explored and exploited in a wide range of application domains, such as production planning, stock control systems, warehousing, and aerospace/defense logistics. However, even with a rich history of academic focus, we are still very much in the dark when it comes to cornerstone computational questions around inventory staggering and to related structural characterizations, with our methodological toolbox being severely under-stocked.

The central contribution of this paper consists in devising a host of algorithmic techniques and analytical ideas — some being entirely novel and some leveraging well-studied concepts in combinatorics and number theory — for surpassing essentially all known approximation guarantees for the inventory staggering problem. In particular, our work demonstrates that numerous structural properties open the door for designing polynomial-time approximation schemes, including polynomially-bounded cycle lengths, constantly-many distinct time intervals, so-called nested instances, and pairwise coprime settings. These findings offer substantial improvements over currently available constant-factor approximations and resolve outstanding open questions in their respective contexts (Teo et al., 1998; Hum et al., 2005; Hochbaum and Rao, 2019, 2020). In parallel, we develop new theory around a number of yet-uncharted questions, related to the sampling complexity of peak inventory estimation as well as to the plausibility of groupwise synchronization. Interestingly, we establish the global nature of inventory staggering, proving that there are *n*-item instances where, for every subset of roughly \sqrt{n} items, no policy improves on the worst-possible one by a factor greater than $1 + \epsilon$, whereas for the entire instance, there exists a policy that outperforms the worst-possible one by a factor of nearly 2, which is optimal.

Keywords: Inventory theory, impossibility results, approximation schemes, LP-rounding

^{*}Department of Mathematics, Princeton University, Princeton, NJ, 08544. Email: nalon@math.princeton.edu. Research supported in part by NSF grant DMS-2154082.

[†]School of Mathematical Sciences and Coller School of Management, Tel Aviv University, Tel Aviv 69978, Israel. Email: segevdanny@tauex.tau.ac.il. Research supported by Israel Science Foundation grant 1407/20.

Contents

1	Introduction		1
	1.1	Model formulation	1
	1.2	Known results and main contributions	2
2	Impossibility Results		6
	2.1	Auxiliary claims	6
	2.2	The sample complexity of peak evaluation	8
	2.3	Infeasibility of groupwise synchronization	9
3	Approximation Scheme in Terms of Cycle Length		12
	3.1	$O(rac{1}{\epsilon})$ -sized discretization sets	13
	3.2	The LP-rounding algorithm	13
	3.3	Analysis	14
	3.4	Additional proofs	16
4	Арг	Approximation Scheme in Terms of Distinct Time Intervals	
	4.1	$O(n^{O(n)} \cdot \mathcal{I} ^{O(1)})$ -time peak evaluation oracle	17
	4.2	Mimicking partitions: Definitions and usefulness	19
	4.3	Mimicking partitions: Construction and testing	20
	4.4	Proof of Lemma 4.4	21
5	Approximation Scheme for Nested Instances		23
	5.1	Continuous-to-discrete reduction	23
	5.2	Well-separated partitions	23
	5.3	The Near-Additivity Theorem	24
	5.4	Algorithmic implications	26
6	Approximation Scheme for Continuous Pairwise Coprime Instances		26
	6.1	Instilling ϵ -shortness and near-additivity	27
	6.2	Algorithmic implications	28
7	Con	cluding Remarks	29
Bi	Bibliography		

1 Introduction

The principal goal of this paper is to deepen our understanding and to inject fresh algorithmic ideas into a foundational inventory management model, commonly known as the inventory staggering problem. Since its inception in the mid-60s (Homer, 1966; Page and Paul, 1976; Hartley and Thomas, 1982), this paradigm has been explored and exploited in a wide range of application domains, such as production planning, stock control systems, warehousing, and aerospace/defense logistics, just to name a few. Concurrently, on the theoretical front, we have been witnessing a steady stream of advances in this context, including an appreciable number of highly innovative papers; see, e.g., Zoller (1977), Gallego et al. (1992), Teo et al. (1998), Hum et al. (2005), and Hochbaum and Rao (2019). That said, even with a rich history of academic focus, we are still very much in the dark when it comes to cornerstone computational questions around inventory staggering and to related structural characterizations, with our methodological toolbox being severely under-stocked.

Broadly speaking, while inventory staggering problems present themselves in various configurations, they all fundamentally address the challenge of efficiently synchronizing multiple periodic replenishment policies via horizontal offsets, with the objective of minimizing peak storage requirement across a given planning horizon. In spite of its deceivingly-simple geometric interpretation, what renders this synchronization task particularly elusive is the interplay between numerous replenishment policies, whose individual cycle lengths and ordering quantities may be completely unrelated. Consequently, even very elementary questions, such as evaluating the peak storage requirement of a given solution in polynomial time, remain unresolved to date. Moreover, when we turn the spotlight to optimization-focused questions, nearly all known research directions continue to lack definitive answers. To delve into the finer details of these questions and to preface our main contributions, we proceed by providing a formal mathematical description of the inventory staggering problem in its most general form.

1.1 Model formulation

Input description and replenishment policies. Whether one considers discrete forms of the inventory staggering problem or their continuous counterparts, they are all defined with respect to a finite collection of n items. Across the continuous planning horizon $[0, \infty)$, the inventory level of each item $i \in [n]$ is replenished via an individual stationary order sizes and stationary intervals (SOSI) policy. Here, *i*-orders of identical quantities are evenly positioned along the timeline, each leaving us with zero inventory upon arriving at its subsequent order. Consequently, any such policy is fully characterized by the next three features:

- The uniform time interval T_i between successive *i*-orders.
- The number of units H_i comprising every *i*-order.
- The horizontal shift $\tau_i \in [0, T_i]$.

It is important to mention that both the time interval T_i between successive *i*-orders and the ordering quantity H_i are assumed to be integer-valued, specified as part of our input description.

By contrast, τ_i is a decision variable, which will be restricted to integer values in the discrete problem setting, and to real values in the continuous setting. Given these parameters, *i*-orders will be placed at the time points

...,
$$-2T_i + \tau_i$$
, $-T_i + \tau_i$, τ_i , $T_i + \tau_i$, $2T_i + \tau_i$, ..., (1)

each consisting of exactly H_i units, meaning that zero-inventory levels are reached at each of these points. For simplicity, this policy will be designated by $\mathcal{P}^i_{\tau_i}$, noting that we have extended the sequence (1) into the negative orthant just to avoid convoluted notation later on.

Inventory levels and objective value. Based on the preceding discussion, it is easy to verify that, when replenished according to $\mathcal{P}_{\tau_i}^i$, the inventory level of item *i* at any time $t \in \mathbb{R}$ is given by

$$I_i(\tau_i, t) = H_i \cdot \left(1 - \frac{t - \lfloor t \rfloor_{\mathcal{P}_{\tau_i}^i}}{T_i}\right) ,$$

where $\lfloor \cdot \rfloor_{\mathcal{P}^{i}_{\tau_{i}}}$ is an operator that rounds its argument down to the nearest ordering point in $\mathcal{P}^{i}_{\tau_{i}}$. As such, for any shift vector $\tau = (\tau_{1}, \ldots, \tau_{n})$, the overall inventory level of the joint policy $\mathcal{P}_{\tau} = (\mathcal{P}^{1}_{\tau_{1}}, \ldots, \mathcal{P}^{n}_{\tau_{n}})$ at any time $t \in \mathbb{R}$ is precisely $I_{\Sigma}(\tau, t) = \sum_{i \in [n]} I_{i}(\tau_{i}, t)$.

Now, let us observe that, since the time intervals T_1, \ldots, T_n are assumed to be integers, the policy \mathcal{P}_{τ} is necessarily cyclic, with a cycle length of $\Lambda = \text{LCM}(T_1, \ldots, T_n)$. In turn, the peak inventory level of this policy can be written as

$$I_{\max}(au) = \max_{t \in [0,\Lambda]} I_{\Sigma}(au,t)$$

As a side note, the above-mentioned maximum is indeed attained, since the policy \mathcal{P}_{τ} has a finite number of ordering points across $[0, \Lambda]$, and since the function $I_{\Sigma}(\tau, \cdot)$ is decreasing between any successive pair of these points.

In the discrete inventory staggering problem, our objective is to determine a shift vector $\tau \in \mathbb{Z}_+^n$ for which the peak inventory level $I_{\max}(\tau)$ of the policy \mathcal{P}_{τ} is minimized. Here, the optimum value is clearly attained at an integer time point, and for any given instance \mathcal{I} , this measure will be designated by

$$OPT^{disc}(\mathcal{I}) = \min_{\tau \in \mathbb{Z}^n_+} \max_{t \in [0,\Lambda] \cap \mathbb{Z}} I_{\Sigma}(\tau, t) .$$

In the continuous setting, we arrive at precisely the same formulation, except for expanding our solution space to $\tau \in \mathbb{R}^n_+$. The optimum peak inventory level of such instances, which is not necessarily attained at an integer time point, will be denoted by

$$OPT^{cont}(\mathcal{I}) = \min_{\tau \in \mathbb{R}^n_+} \max_{t \in [0,\Lambda]} I_{\Sigma}(\tau, t)$$

1.2 Known results and main contributions

The central contribution of this paper consists in devising a host of algorithmic techniques and analytical ideas — some being entirely novel and some leveraging well-studied concepts in combinatorics and number theory — for surpassing essentially all known approximation guarantees for the inventory staggering problem. Concurrently, we develop new theory around a number of yet-uncharted questions, related to the sampling complexity of peak inventory estimation as well as to the plausibility of groupwise synchronization.

For ease of exposition, we proceed by providing a formal account of our main findings, subsequent to a concise preamble on existing state-of-the-art results in each of these directions. As a side note, the upcoming presentation order is primarily designed to offer logical and user-friendly content; this order is uncorrelated with importance, uniqueness, or magnitude of improvement. Additionally, given the extensive body of work dedicated to studying inventory staggering problems, including rigorous methods, heuristics, and experimental evidence, we cannot do justice and exhaustively discuss these research domains. For an in-depth literature review, readers are referred to useful articles in this context (Gallego et al., 1996; Hall, 1998; Murthy et al., 2003; Boctor, 2010; Hochbaum and Rao, 2019) as well as to the references therein.

Main result 1: Sample complexity of peak evaluation. As previously mentioned, efficiently evaluating the peak inventory level $I_{\max}(\tau)$ of a given replenishment policy \mathcal{P}_{τ} has remained an unresolved issue since the very inception of inventory staggering, forcing most efforts around this problem to focus on stylized formulations. In Section 2, we provide the first rigorous evidence for the inherent difficulty of this question, by deriving an exponential lower bound on its sample complexity. Specifically, suppose we randomly draw M independent time points $X_1, \ldots, X_M \sim U\{0, \ldots, \Lambda - 1\}$, and consider the maximal inventory level $\tilde{I}_{\max}^M = \max_{m \in [M]} I_{\Sigma}(\tau^*, X_m)$ across these points as our estimator for $I_{\max}(\tau^*)$, where τ^* is a given optimal shift vector. Then, as an immediate conclusion of the next theorem, $2^{\Omega(\epsilon^2 n)}$ random points should generally be drawn in order to estimate $I_{\max}(\tau^*)$ within factor $\frac{1}{2} + \epsilon$ with constant probability.

THEOREM 1.1. There exists a discrete inventory staggering instance, together with an optimal shift vector $\tau^* \in \mathbb{Z}^n_+$, for which

$$\Pr\left[\tilde{I}_{\max}^{M} \ge \left(\frac{1}{2} + \epsilon\right) \cdot I_{\max}(\tau^{*})\right] \le M e^{-\epsilon^{2}n/6} .$$

Main result 2: Infeasibility of groupwise synchronization. Our next contribution is motivated by the yet-unknown usefulness of pairwise synchronization, recently investigated by Segev (2024) in the context of economic warehouse lot scheduling. Deferring the nutsand-bolts of this concept to Section 2, at a high level, suppose we initially compute a nearoptimal shift vector τ^{ij} for each pair of items $i \neq j$, hoping to meaningfully drop below the trivial peak inventory level of $H_i + H_j$ for a significant portion of these pairs. Then, the essential question is whether there exists a partition \mathcal{M} of the item set into pairs, such that $\sum_{(i,j)\in\mathcal{M}} I_{\max}(\tau^{ij}) \leq (1-\delta) \cdot \sum_{i\in[n]} H_i$, for some absolute constant $\delta > 0$. As explained later on, such a result would immediately lead to a $2(1-\delta)$ -approximation for the discrete inventory staggering problem in its utmost generality, beating the best known approximation guarantee for arbitrarily-structured instances. Surprisingly, we prove that this approach cannot succeed for pairs of items, for triplets, or for subsets of much larger size. Specifically, as a direct corollary of Theorem 1.2 below, even when optimal shift vectors are readily available for subsets of nearly \sqrt{n} items, gluing them together across all subsets would still be far from optimal by a factor of $2 - O(\epsilon)$.

THEOREM 1.2. There exists a discrete inventory staggering instance \mathcal{I} such that:

1. OPT^{disc}
$$(\mathcal{I}) \leq (\frac{1}{2} + \epsilon) \cdot \sum_{i \in [n]} H_i$$
.

2. OPT^{disc} $(\mathcal{I}_{\hat{\mathcal{F}}}) \ge (1-\epsilon) \cdot \sum_{i \in \hat{\mathcal{F}}} H_i$, for every subset $\hat{\mathcal{F}}$ of $O(\frac{\sqrt{n}}{\log n})$ items.

The boundary between trivial and non-trivial. Moving forward, to clearly mark an informative reference point for recognizing non-trivial approximation guarantees, it is worth keeping in mind that the peak inventory level of any shift vector is always within factor 2 of optimal. This property is a direct outcome of the simple classic average-space bound (see Lemma 2.1), stating that $OPT^{disc}(\mathcal{I}) \geq \frac{1}{2} \cdot \sum_{i \in [n]} H_i \cdot (1 + \frac{1}{T_i})$ and $OPT^{cont}(\mathcal{I}) \geq \frac{1}{2} \cdot \sum_{i \in [n]} H_i$.

Main result 3: Approximation scheme in terms of cycle length. As far as proximity to optimum is concerned, state-of-the-art performance guarantees for the discrete inventory staggering problem have recently culminated to polynomial-time approximation schemes (PTAS), when the cycle length $\Lambda = \text{LCM}(T_1, \ldots, T_n)$ we are facing is bounded by a constant, i.e., $\Lambda = O(1)$. In particular, when items are associated with identical time intervals, Hochbaum and Rao (2019) employed dynamic programming ideas to devise an approximation scheme in $O(n \cdot (\frac{1}{\epsilon})^{O(\Lambda+1/\epsilon)})$ time. Shortly thereafter, Hochbaum and Rao (2020) were successful in eliminating the identical-intervals assumption, albeit ending up with an $O(n^{O(\Lambda)} \cdot (\frac{1}{\epsilon})^{O(\Lambda)})$ -time approach. As part of their concluding remarks, the authors asked whether it is possible to obtain a fixed-parameter tractable PTAS, namely, one admitting an $f_1(\Lambda, \frac{1}{\epsilon}) \cdot |\mathcal{I}|^{f_2(1/\epsilon)}$ -time implementation, where $|\mathcal{I}|$ stands for our input length in its binary representation.

In Section 3, we resolve the above-mentioned question, in a much stronger sense: As long as Λ is polynomial in the input size (rather than a mere constant), we prove that the optimal peak inventory level can be efficiently approached within any degree of accuracy, as formally stated in Theorem 1.3 below. For convenience, we design our LP-based approximation scheme while shooting for a success probability of $\frac{1}{2}$; the latter can be arbitrarily amplified via independent repetitions. Yet another important remark is that this result can easily be adapted to the continuous problem formulation, as well as to exponentially-large identical time intervals, via relatively simple reductions to the discrete setting (see Sections 3.1 and 5.1).

THEOREM 1.3. For any $\epsilon > 0$, the discrete inventory staggering problem can be approximated within factor $1 + \epsilon$ of optimal. The running time of our randomized algorithm is $O(\Lambda^{\tilde{O}(1/\epsilon^3)} \cdot |\mathcal{I}|^{O(1)})$, and it is successful with probability at least $\frac{1}{2}$.

Main result 4: Approximation scheme in terms of distinct time intervals. The next thread of work we consider is that of approximating inventory staggering instances characterized by constantly-many distinct time intervals, noting that this setting will play an important role within the algorithmic advances presented in future sections. In spite of their misleadingly simple structure, such instances continue to elude truly near-optimal solution methods, with non-trivial guarantees currently available only for two distinct time intervals. In this scenario, Teo et al. (1998) showed that continuous inventory staggering can be efficiently approximated within factor $\frac{4}{3}$, when one time interval is an integer multiple of the other. Later on, Hum et al. (2005) eliminated the latter assumption, attaining an approximation ratio of roughly 1.34312 via a generalized Homer's policy (Homer, 1966).

In Section 4, we improve on the above-mentioned guarantees, further lifting these results beyond the two-interval scenario, along the lines of Theorem 1.4 below. Specifically, letting Kdesignate the underlying number of distinct time intervals, we devise a deterministic approach for discrete inventory staggering, arguing that the latter setting can be approximated within factor $1+\epsilon$ in $O(2^{\tilde{O}(K/\epsilon^2)} \cdot |\mathcal{I}|^{O(1)})$ time. Similarly to Theorem 1.3, this result seamlessly migrates to the continuous problem formulation. A particularly interesting tool we develop en route is an exact $O(n^{O(n)} \cdot |\mathcal{I}|^{O(1)})$ -time oracle for computing the peak inventory level of a given policy, thereby resolving the open question of peak evaluation for constantly-many items.

THEOREM 1.4. For any $\epsilon > 0$, the discrete inventory staggering problem can be approximated within factor $1 + \epsilon$ of optimal. The running time of our algorithm is $O(2^{\tilde{O}(K/\epsilon^2)} \cdot |\mathcal{I}|^{O(1)})$.

Main result 5: Approximation scheme for nested instances. Beyond the aforementioned scenarios, prior studies have uncovered non-trivial approximation guarantees under two additional structural features. The first such finding has been established for so-called nested instances; here, when the set of items is indexed such that $T_1 \leq \cdots \leq T_n$, each time interval divides its successor. This setting, which is highly compelling due to its emergence in power-of-2 policies (see, e.g., Roundy (1985, 1986); Jackson et al. (1985); Muckstadt and Roundy (1993)) is known to admit a polynomial-time $\frac{15}{8}$ -approximation (Teo et al., 1998). Our main result for nested instances, whose specifics are discussed in Section 5, resides in devising a polynomial-time approximation scheme.

THEOREM 1.5. For any $\epsilon > 0$, the nested inventory staggering problem can be approximated within factor $1 + \epsilon$ of optimal. The running time of our algorithm is $O(2^{\tilde{O}(1/\epsilon^3)} \cdot |\mathcal{I}|^{O(1)})$.

Main result 6: Approximation scheme for continuous pairwise coprime instances. The second setting for which a sub-2-approximation has been demonstrated to exist is centered around pairwise coprime instances of the continuous inventory staggering problem. Here, all pairs $T_i \neq T_j$ of time intervals are assumed to be relatively prime, meaning that $gcd(T_i, T_j) = 1$. It is worth noting that, by the Chinese Remainder Theorem, the peak inventory level of any replenishment policy for the discrete problem formulation is precisely $\sum_{i \in [n]} H_i$, rendering it completely trivial. However, when continuous inventory staggering is concerned, this property can easily be violated, making such instances particularly difficult to deal with. As outlined in Section 6, our final contribution comes in the form of a polynomial-time approximation scheme for pairwise coprime instances of the latter formulation, improving on its existing 1.4approximation due to Hum et al. (2005).

THEOREM 1.6. For any $\epsilon > 0$, the continuous coprime inventory staggering problem can be approximated within factor $1 + \epsilon$ of optimal. The running time of our algorithm is $O(\operatorname{tower}_4(O(\frac{1}{\epsilon}), O(\frac{1}{\epsilon})) \cdot |\mathcal{I}|^{O(1)}).$ As a side note, to avoid convoluted running time expressions, this theorem utilizes the function tower_{β}(h, x), where β is iteratively exponentiated h - 1 times, finally taken to the power of x, i.e.,

tower_{$$\beta$$}(h, x) = $\underbrace{\beta^{\beta}}_{h \text{ times}}^{h}$

As such, tower₄ $(O(\frac{1}{\epsilon}), O(\frac{1}{\epsilon}))$ is a function of $\frac{1}{\epsilon}$ and nothing more. To our knowledge, approximation schemes characterized by power-tower running times are few and far between, mostly arising as algorithmic applications of Szemerédi's Regularity Lemma; see, for instance, Alon et al. (1994) and Duke et al. (1995).

Hardness results. To better understand the nature of algorithmic outcomes one could be striving for, it is worth briefly mentioning known intractability results surrounding inventory staggering. Along these lines, Gallego et al. (1992) were the first to rigorously investigate how plausible it is to efficiently compute optimal shift vectors, proving that the continuous problem formulation is strongly NP-hard. Subsequently, Hall (1998) studied discrete inventory staggering, establishing weak NP-hardness even when all items share a common time interval of T = 2. Recently, Hochbaum and Rao (2019) showed that the latter setting becomes strongly NP-hard when T could be arbitrarily large. To our knowledge, whether the discrete formulation or its continuous counterpart are APX-hard in their full generality is still a major unresolved question.

2 Impossibility Results

In this section, we explore fundamental questions about estimating the peak inventory level of a given shift vector as well as about the plausibility of identifying non-trivial solutions via groupwise synchronization. Toward these objectives, Section 2.1 presents a number of known observations that will be helpful down the road. Then, Section 2.2 derives an exponential lower bound on the sample complexity of peak evaluation. Finally, Section 2.3 is devoted to arguing that even $\tilde{\Omega}(\sqrt{n})$ -synchronization is insufficient to obtain non-straightforward approximation guarantees.

2.1 Auxiliary claims

The average-space lower bound. Moving forward, it will be useful to keep in mind wellknown lower bounds on the peak inventory levels of optimal replenishment policies for the discrete and continuous formulations, $\text{OPT}^{\text{disc}}(\mathcal{I})$ and $\text{OPT}^{\text{cont}}(\mathcal{I})$. These bounds follow from a simple averaging argument, and we provide their proof for completeness. To avoid cumbersome expressions, $H_{\Sigma} = \sum_{i \in [n]} H_i$ will stand for the total ordering quantity of all items.

LEMMA 2.1. $\operatorname{OPT}^{\operatorname{disc}}(\mathcal{I}) \geq \frac{1}{2} \cdot \sum_{i \in [n]} H_i \cdot (1 + \frac{1}{T_i}) \text{ and } \operatorname{OPT}^{\operatorname{cont}}(\mathcal{I}) \geq \frac{H_{\Sigma}}{2}.$

PROOF. Recalling that $\Lambda = \text{LCM}(T_1, \ldots, T_n)$ is the cycle length of any replenishment policy, suppose we draw a random point $X \sim U\{0, 1, \ldots, \Lambda - 1\}$. Then, for any shift vector $\tau \in \mathbb{Z}_+^n$

and for any item $i \in [n]$, it is easy to verify that $I_i(\tau_i, X) \sim U\{\frac{H_i}{T_i}, \frac{2H_i}{T_i}, \dots, \frac{T_iH_i}{T_i}\}$. Therefore,

$$I_{\max}(\tau) \geq \mathbb{E}\left[I_{\Sigma}(\tau, X)\right] = \sum_{i \in [n]} \mathbb{E}\left[I_{i}(\tau_{i}, X)\right] = \frac{1}{2} \cdot \sum_{i \in [n]} H_{i} \cdot \left(1 + \frac{1}{T_{i}}\right) ,$$

implying that $\operatorname{OPT}^{\operatorname{disc}}(\mathcal{I}) \geq \frac{1}{2} \cdot \sum_{i \in [n]} H_i \cdot (1 + \frac{1}{T_i})$. In the continuous setting, we repeat precisely the same argument, except for sampling $X \sim U[0, \Lambda]$. As a result, $I_i(\tau_i, X) \sim U[0, H_i]$, ending up with $\operatorname{OPT}^{\operatorname{cont}}(\mathcal{I}) \geq \frac{1}{2} \cdot \sum_{i \in [n]} H_i = \frac{H_{\Sigma}}{2}$.

The near-optimality regime of random shifts. Let τ^R be a random shift vector where, for each item $i \in [n]$, we pick $\tau_i^R \sim U\{0, \ldots, T_i - 1\}$, independently of any other item. In what follows, we focus our attention on the next question: What conditions guarantee that τ^R is near-optimal with high probability, or at least with some positive probability? To derive such conditions, we first establish the next claim, providing an upper bound on the probability that the random inventory level of τ^R pointwise exceeds the average-space bound by a factor greater than $1 + \epsilon$.

LEMMA 2.2. For every time point $t \in [0, \Lambda] \cap \mathbb{Z}$,

$$\Pr\left[I_{\Sigma}(\tau^{R}, t) \geq \frac{1+\epsilon}{2} \cdot \sum_{i \in [n]} H_{i} \cdot \left(1 + \frac{1}{T_{i}}\right)\right] \leq \exp\left(-\frac{\epsilon^{2}}{6} \cdot \frac{H_{\Sigma}}{H_{\max}}\right)$$

PROOF. To derive the desired claim, for every item $i \in [n]$, let $Y_i^t = I_i(\tau_i^R, t)$ be its inventory level at time t in terms of the random policy \mathcal{P}_{τ^R} . It is easy to see that $Y_i^t \sim U\{\frac{H_i}{T_i}, \frac{2H_i}{T_i}, \dots, \frac{T_iH_i}{T_i}\}$, and consequently,

$$\Pr\left[I_{\Sigma}(\tau^{R}, t) \geq \frac{1+\epsilon}{2} \cdot \sum_{i \in [n]} H_{i} \cdot \left(1 + \frac{1}{T_{i}}\right)\right]$$
$$= \Pr\left[\sum_{i \in [n]} \frac{Y_{i}^{t}}{H_{\max}} \geq (1+\epsilon) \cdot \mathbb{E}\left[\sum_{i \in [n]} \frac{Y_{i}^{t}}{H_{\max}}\right]\right]$$
$$\leq \exp\left(-\frac{\epsilon^{2}}{6} \cdot \frac{H_{\Sigma}}{H_{\max}}\right).$$

The last inequality is obtained by employing the next Chernoff-Hoeffding bound (see Dubhashi and Panconesi (2009, Thm. 1.1 and Ex. 1.1)) with respect to $\{\frac{Y_i^t}{H_{\max}}\}_{i \in [n]}$: Let Z_1, \ldots, Z_n be independent [0, 1]-bounded random variables. Then, for every $\bar{\mu} \geq \mathbb{E}[\sum_{i \in [n]} Z_i]$ and $\xi \in (0, 1)$, we have

$$\Pr\left[\sum_{i\in[n]} Z_i > (1+\xi) \cdot \bar{\mu}\right] \leq \exp\left(-\frac{\xi^2 \bar{\mu}}{3}\right) . \tag{2}$$

The first implication of Lemma 2.2 is that, when $\Lambda < \exp(\frac{\epsilon^2}{6} \cdot \frac{H_{\Sigma}}{H_{\max}})$, by taking the union bound over all times point in $[0, \Lambda) \cap \mathbb{Z}$, we have $I_{\max}(\tau^R) \leq \frac{1+\epsilon}{2} \cdot \sum_{i \in [n]} H_i \cdot (1 + \frac{1}{T_i})$ with

positive probability. Consequently, $OPT^{disc}(\mathcal{I}) \leq \frac{1+\epsilon}{2} \cdot \sum_{i \in [n]} H_i \cdot (1 + \frac{1}{T_i})$ for any such instance \mathcal{I} , nearly matching the average-space lower bound of Lemma 2.1. The second implication of this result is that, when $\Lambda < \delta \cdot \exp(\frac{\epsilon^2}{6} \cdot \frac{H_{\Sigma}}{H_{\max}})$, the peak inventory level of τ^R is within factor $1 + \epsilon$ of optimal, with probability at least $1 - \delta$. It is worth mentioning that bounds of similar flavor have previously been obtained by Croot and Huang (2013); however, their precise form is somewhat cumbersome to exploit in our particular setting.

2.2 The sample complexity of peak evaluation

In what follows, we derive Theorem 1.1, showing that $2^{\Omega(\epsilon^2 n)}$ randomly drawn points are required in order to estimate the peak inventory level of a given shift vector within factor $\frac{1}{2} + \epsilon$ with constant probability.

Construction. Let us consider a discrete inventory staggering instance \mathcal{I} with $n \geq \frac{2}{\epsilon}$ items, whose ordering quantities are uniform, say $H_1 = \cdots = H_n = 1$. The time intervals T_1, \ldots, T_n of these items are distinct primes, each of value at least n; as a result, our cycle length is $\Lambda = \prod_{i \in [n]} T_i$. Focusing on the shift vector $\tau = \vec{0}$, we wish to estimate the peak inventory level of $\mathcal{P}_{\vec{0}}$, which is clearly $I_{\max}(\vec{0}) = n$. To this end, our randomized procedure operates as follows:

- We independently draw M time points, X_1, \ldots, X_M , such that $X_m \sim U\{0, \ldots, \Lambda 1\}$.
- Our estimate \tilde{I}_{\max}^M for $I_{\max}(\vec{0})$ is given by the maximum value of $I_{\Sigma}(\vec{0}, \cdot)$ across these points, i.e., $\tilde{I}_{\max}^M = \max_{m \in [M]} I_{\Sigma}(\vec{0}, X_m)$.

Analysis. The next claim shows that, with probability at least $1 - Me^{-\epsilon^2 n/6}$, the maximal inventory level inferred from our sample is at most $(\frac{1}{2} + \epsilon) \cdot n$. As previously noted, $I_{\max}(\vec{0}) = n$, meaning that the number of points required to attain a $(\frac{1}{2} + \epsilon)$ -estimate with constant probability is $2^{\Omega(\epsilon^2 n)}$.

Lemma 2.3. $\Pr[\tilde{I}_{\max}^M \ge (\frac{1}{2} + \epsilon) \cdot n] \le M e^{-\epsilon^2 n/6}$.

PROOF. By taking the union bound over X_1, \ldots, X_M , it suffices to show that $\Pr[I_{\Sigma}(\vec{0}, X_m) \ge (\frac{1}{2} + \epsilon) \cdot n] \le e^{-\epsilon^2 n/6}$ for every $m \in [M]$. To this end, letting $I_i^{X_m} = I_i(0, X_m)$ be the inventory level of item i at time X_m , it is easy to see that $I_i^{X_m} \sim U\{\frac{1}{T_i}, \frac{2}{T_i}, \ldots, \frac{T_i}{T_i}\}$, implying in particular that $\mathbb{E}[I_i^{X_m}] = \frac{1}{2}(1 + \frac{1}{T_i})$. The important observation is that $\{I_i^{X_m}\}_{i \in [n]}$ are mutually independent. To verify this claim, given a vector $s = (s_1, \ldots, s_n) \in [T_1] \times \cdots \times [T_n]$, let t_s be the unique solution modulo Λ to the system of congruences

$$t \equiv T_i - s_i \pmod{T_i} \quad \forall i \in [n]$$

Recalling that T_1, \ldots, T_n are distinct primes, by the Chinese Remainder Theorem (see, e.g., (Niven et al., 1991, Sec. 2.3)), we know that t_s indeed exists, and also, that it is unique (modulo Λ). Since we are considering the shift vector $\tau = \vec{0}$, it follows that $\bigwedge_{i \in [n]} \{I_i^{X_m} = \frac{s_i}{T_i}\} = \{X_m = I_i^{X_m} \in I_i^{X_m}\}$

 t_s , implying that $\{I_i^{X_m}\}_{i \in [n]}$ are mutually independent, as

$$\Pr\left[\bigwedge_{i\in[n]}\left\{I_i^{X_m} = \frac{s_i}{T_i}\right\}\right] = \Pr\left[X_m = t_s\right] = \frac{1}{\Lambda} = \frac{1}{\prod_{i\in[n]}T_i} = \prod_{i\in[n]}\Pr\left[I_i^{X_m} = \frac{s_i}{T_i}\right].$$

Based on these observations, we derive the desired upper bound on $\Pr[I_{\Sigma}(\vec{0}, X_m) \ge (\frac{1}{2} + \epsilon) \cdot n]$ by noting that

$$\Pr\left[I_{\Sigma}(\vec{0}, X_m) \ge \left(\frac{1}{2} + \epsilon\right) \cdot n\right] \le \Pr\left[\sum_{i \in [n]} I_i^{X_m} \ge (1 + \epsilon) \cdot \mathbb{E}\left[\sum_{i \in [n]} I_i^{X_m}\right]\right]$$
(3)

$$\leq \exp\left(-\frac{\epsilon^2}{3} \cdot \mathbb{E}\left[\sum_{i \in [n]} I_i^{X_m}\right]\right) \tag{4}$$

$$\leq e^{-\epsilon^2 n/6} . \tag{5}$$

Here, inequality (3) holds since $(1 + \epsilon) \cdot \mathbb{E}[\sum_{i \in [n]} I_i^{X_m}] = \frac{1+\epsilon}{2} \cdot \sum_{i \in [n]} (1 + \frac{1}{T_i}) \leq (\frac{1}{2} + \epsilon) \cdot n$, as $T_i \geq n \geq \frac{2}{\epsilon}$. Inequality (4) follows by specializing the Chernoff-Hoeffding bound (2) with respect to the random variables $\{I_i^{X_m}\}_{i \in [n]}$, which are independent and [0, 1]-bounded. Finally, inequality (5) follows by recalling that $\mathbb{E}[I_i^{X_m}] = \frac{1}{2}(1 + \frac{1}{T_i}) \geq \frac{1}{2}$.

2.3 Infeasibility of groupwise synchronization

Our second impossibility result is motivated by the recent work of Segev (2024) on economic warehouse lot scheduling, where among other ideas, pairwise synchronization has been crucial for improving on the long-standing 2-approximation in this context (Anily, 1991; Gallego et al., 1996). Migrating this notion to inventory staggering, one would operate along the following lines:

- First, for each pair of items $i \neq j$, we compute an optimal shift vector τ^{ij} . This step is implementable in polynomial time, given the existence of closed-form solutions for two-item instances (Hartley and Thomas, 1982; Murthy et al., 2003). Our hope is to discover that, for a significant portion of these pairs, a peak inventory level of $I_{\max}(\tau^{ij}) \leq$ $(1 - \delta_{ij}) \cdot (H_i + H_j)$ can be attained, for some $\delta_{ij} > 0$.
- Then, the big unknown is whether there exists a partition \mathcal{M} of the overall set of items into pairs, such that $\sum_{(i,j)\in\mathcal{M}} I_{\max}(\tau^{ij}) \leq (1-\delta) \cdot H_{\Sigma}$, for some absolute constant $\delta > 0$. Combined with Lemma 2.1, such a result would immediately lead to a $2(1-\delta)$ approximation for the discrete inventory staggering problem in its full generality.

Unfortunately, we argue that this approach cannot succeed for pairs of items, for triplets, or for subsets of much larger size. As formally stated in Theorem 1.2, the remainder of this section is dedicated to presenting an inventory staggering instance \mathcal{I} where, on the one hand, $\text{OPT}^{\text{disc}}(\mathcal{I}) \leq (1+2\epsilon) \cdot \frac{H_{\Sigma}}{2}$, and on the other hand, $\text{OPT}^{\text{disc}}(\mathcal{I}_{\hat{\mathcal{F}}}) \geq (1-\epsilon) \cdot H(\hat{\mathcal{F}})$ for every subset $\hat{\mathcal{F}}$ of $O(\frac{\sqrt{n}}{\log n})$ items. Consequently, even when one is capable of computing optimal shift

vectors for subsets of nearly \sqrt{n} items, simply gluing them together across all subsets would still be far from the global optimum by a factor of $2 - O(\epsilon)$.

Construction. To describe our construction, it will be convenient to utilize a well-known result regarding sparse combinatorial designs. Specifically, suppose that \mathcal{F} is a family of qelement subsets of [n]. For $r \leq q$, we say that \mathcal{F} is r-sparse when any two subsets in \mathcal{F} intersect in at most r - 1 elements; $\mathcal{M}(n, q, r)$ will designate the maximum size of such a family. The next result is taken from the work of Alon et al. (2012, Sec. 2.2), noting that bounds in this spirit have also been suggested earlier, for example, by Hartman et al. (1982) and by Nisan and Wigderson (1994).

THEOREM 2.4. $\mathcal{M}(q^2, q, r) \ge q^r$, for every prime power q and $2 \le r \le q$.

Now, given an error parameter $\epsilon > 0$, the inventory staggering instance \mathcal{I} we consider is constructed as follows:

- For a prime $q \ge e^{1/(2\epsilon^2)}$, let $K = q^2$, and let p_1, \ldots, p_K be a collection of distinct primes within $[K, 2K \ln K]$. By the Prime Number Theorem (see, e.g., Davenport and Montgomery (2013, Sec. 7)), for sufficiently large K, we indeed have at least K such primes, with room to spare.
- *Items*: By Theorem 2.4, we know that $\mathcal{M}(K, \sqrt{K}, 3) \geq K^{3/2}$, implying the weaker claim that there exists a 3-sparse family \mathcal{F} of \sqrt{K} -element subsets of [K], with $|\mathcal{F}| = \lceil 12K \ln^2 K \rceil$. For every subset $S \in \mathcal{F}$, we introduce a corresponding item i_S , meaning that $n = |\mathcal{F}|$.
- Time intervals and ordering quantities: For every item i_S , the time interval between successive i_S -orders is $T_{i_S} = \prod_{k \in S} p_k$, with an ordering quantity of $H_{i_S} = 1$.

Upper bound on global optimum. With respect to this instance, we begin by arguing that its optimal peak inventory level is essentially $\frac{n}{2}$, up to ϵ -dependent terms. This finding will be deduced by showing that \mathcal{I} meets our sufficient condition for nearly matching the average-space lower bound of Lemma 2.1.

LEMMA 2.5. $\operatorname{OPT}^{\operatorname{disc}}(\mathcal{I}) \leq (1+2\epsilon) \cdot \frac{n}{2}.$

PROOF. Based on the concluding paragraph of Section 2.1, when $\Lambda < \exp(\frac{\epsilon^2}{6} \cdot \frac{H_{\Sigma}}{H_{\max}})$, we have

$$OPT^{disc}(\mathcal{I}) \leq \frac{1+\epsilon}{2} \cdot \sum_{S \in \mathcal{F}} H_{i_S} \cdot \left(1 + \frac{1}{T_{i_S}}\right) \leq \frac{(1+\epsilon)(1+e^{-2/\epsilon})}{2} \cdot \sum_{S \in \mathcal{F}} H_{i_S} \leq (1+2\epsilon) \cdot \frac{n}{2}.$$

To verify the above-mentioned condition for sufficiently large K, note that since p_1, \ldots, p_K are distinct primes within $[K, 2K \ln K]$, we have

$$\Lambda = \operatorname{LCM} \{ T_{i_S} : S \in \mathcal{F} \}$$
$$\leq \prod_{k \in [K]} p_k$$

$$< (2K \ln K)^{K}$$

$$\leq e^{2K \ln K}$$

$$\leq \exp\left(\frac{\epsilon^{2}n}{6}\right)$$
(6)

$$= \exp\left(\frac{\epsilon^2}{6} \cdot \frac{H_{\Sigma}}{H_{\max}}\right) . \tag{7}$$

Here, inequality (6) holds since $K \ge e^{1/\epsilon^2}$ and $n = \lceil 12K \ln^2 K \rceil$, whereas equality (7) is obtained by noting that $H_{\Sigma} = n$ and $H_{\max} = 1$, since $H_{i_S} = 1$ for every $S \in \mathcal{F}$.

Lower bound on $O(\frac{\sqrt{n}}{\log n})$ -wise optimum. Now, let us consider any subset $\hat{\mathcal{F}} \subseteq \mathcal{F}$ of size at most $\frac{\sqrt{K}}{2}$, noting that $\sqrt{K} = \Theta(\frac{\sqrt{n}}{\log n})$. We make use of $\mathcal{I}_{\hat{\mathcal{F}}}$ to designate the instance obtained by restricting \mathcal{I} to this subset. To conclude the proof of Theorem 1.2, we show that with respect to this instance, the optimal peak inventory level is nearly $|\hat{\mathcal{F}}|$, which identifies with the combined ordering quantity across all $\hat{\mathcal{F}}$ -items, and therefore, attainable by any shift vector.

LEMMA 2.6. $\operatorname{OPT}^{\operatorname{disc}}(\mathcal{I}_{\hat{\mathcal{F}}}) \geq (1-\epsilon) \cdot |\hat{\mathcal{F}}|.$

To establish this bound, we begin by deriving an auxiliary number-theoretic claim. Here, the primes p_1, \ldots, p_L are said to be unique divisors of the positive integers n_1, \ldots, n_L when, for every $\ell \in [L]$, the next two properties are satisfied: (1) p_ℓ is a divisor of n_ℓ ; and (2) p_ℓ is a not a divisor of n_k , for every $k \neq \ell$.

LEMMA 2.7. Suppose that p_1, \ldots, p_L are unique divisors of n_1, \ldots, n_L . Then, for any $\tau \in \mathbb{Z}^L$, there exists an integer t such that, for every $\ell \in [L]$,

$$(t - \tau_\ell) \mod n_\ell \in \left[0, \frac{n_\ell}{p_\ell} - 1\right]$$
.

PROOF. In what follows, we make use of $\alpha_{\ell} \geq 1$ to denote the largest integer for which $p_{\ell}^{\alpha_{\ell}}$ is a divisor of n_{ℓ} . With this notation, we will actually prove a stronger result, showing that there exists an integer t such that $(t - \tau_{\ell}) \mod n_{\ell} \in [0, \frac{n_{\ell}}{p_{\ell}^{\alpha_{\ell}}} - 1]$, for every $\ell \in [L]$.

To this end, for every $\ell \in [L]$ and $0 \leq r \leq p_{\ell}^{\alpha_{\ell}} - 1$, consider the system of congruences

$$(F_{\ell,r}) \begin{cases} t \equiv 0 \pmod{\frac{n_{\ell}}{p_{\ell}^{\alpha_{\ell}}}} \\ t \equiv r \pmod{p_{\ell}^{\alpha_{\ell}}} \end{cases}$$

Since $\frac{n_{\ell}}{p_{\ell}^{\alpha_{\ell}}}$ and $p_{\ell}^{\alpha_{\ell}}$ are relatively prime, by the Chinese Remainder Theorem, $(F_{\ell,r})$ has a unique integer solution $t_{\ell,r}$ modulo n_{ℓ} . It is easy to verify that $t_{\ell,0}, \ldots, t_{\ell,p_{\ell}^{\alpha_{\ell}}-1}$ are distinct multiples of $\frac{n_{\ell}}{p_{\ell}^{\alpha_{\ell}}}$ in $[0, n_{\ell} - 1]$, and clearly, there are exactly $p_{\ell}^{\alpha_{\ell}}$ such multiples. Therefore, $(t_{\ell,0} - \tau_{\ell}) \mod n_{\ell}, \ldots, (t_{\ell,p_{\ell}^{\alpha_{\ell}}-1} - \tau_{\ell}) \mod n_{\ell}$ necessarily stab each segment of $\frac{n_{\ell}}{p_{\ell}^{\alpha_{\ell}}}$ consecutive integers in $\mathbb{Z}_{n_{\ell}}$. In particular, there exists some index r_{ℓ} for which $(t_{\ell,r_{\ell}} - \tau_{\ell}) \mod n_{\ell} \in [0, \frac{n_{\ell}}{p_{\ell}^{\alpha_{\ell}}} - 1]$.

Now, to identify the desired integer t, letting $M = \frac{\text{LCM}(n_1,...,n_L)}{\prod_{\ell \in [L]} p_{\ell}^{\alpha_{\ell}}}$, consider the system of congruences

$$\begin{cases} t \equiv 0 \pmod{M} \\ t \equiv r_{\ell} \pmod{p_{\ell}^{\alpha_{\ell}}} & \forall \ell \in [L] \end{cases}$$

Once again, since $M, p_1^{\alpha_1}, \ldots, p_L^{\alpha_L}$ are pairwise coprime, this system has a unique integer solution $t \mod \operatorname{LCM}(n_1, \ldots, n_L)$. We conclude the proof by observing that $(t-\tau_\ell) \mod n_\ell \in [0, \frac{n_\ell}{p_\ell^{\alpha_\ell}} - 1]$ for every $\ell \in [L]$. Indeed, since $\frac{n_\ell}{p_\ell^{\alpha_\ell}}$ divides M, we know in particular that t is a solution to (F_{ℓ,r_ℓ}) . However, as explained above, the latter system has a unique solution $t_{\ell,r_\ell} \mod n_\ell$, implying that by our choice of r_ℓ ,

$$(t - \tau_{\ell}) \mod n_{\ell} \equiv (t_{\ell, r_{\ell}} - \tau_{\ell}) \mod n_{\ell} \in \left[0, \frac{n_{\ell}}{p_{\ell}^{\alpha_{\ell}}} - 1\right].$$

Given this claim, we proceed by arguing that $OPT^{disc}(\mathcal{I}_{\hat{\mathcal{F}}}) \geq (1-\epsilon) \cdot |\hat{\mathcal{F}}|$. For this purpose, let τ^* be an optimal shift vector with respect to $\mathcal{I}_{\hat{\mathcal{F}}}$, meaning that $I_{\max}(\tau^*) = OPT^{disc}(\mathcal{I}_{\hat{\mathcal{F}}})$. Focusing on a single subset $S \in \hat{\mathcal{F}}$, we observe that there are at least two elements in S that do not appear in any other subset in $\hat{\mathcal{F}}$, since

$$\left| S \setminus \left(\bigcup_{\hat{S} \in \hat{\mathcal{F}} \setminus \{S\}} \hat{S} \right) \right| \geq |S| - \sum_{\hat{S} \in \hat{\mathcal{F}} \setminus \{S\}} \left| S \cap \hat{S} \right| \geq \sqrt{K} - 2 \cdot \left(\frac{\sqrt{K}}{2} - 1 \right) = 2,$$

where the second inequality holds since \mathcal{F} is a 3-sparse family of \sqrt{K} -element subsets and since $|\hat{\mathcal{F}}| \leq \frac{\sqrt{K}}{2}$. We make use of $k_S \in S$ to denote one of these two elements.

As a result of this observation, it follows that the primes $(p_{k_S})_{S \in \hat{\mathcal{F}}}$ are unique divisors of $(T_{i_S})_{S \in \hat{\mathcal{F}}}$. Therefore, by Lemma 2.7, there exists an integer t such that $(t - \tau_{i_S}^*) \mod T_{i_S} \in [0, \frac{T_{i_S}}{p_{k_S}} - 1]$, for every $S \in \hat{\mathcal{F}}$. In other words, with respect the policy \mathcal{P}_{τ^*} , the time elapsing between the last i_S -order in [0, t] and the point t itself is at most $\frac{T_{i_S}}{p_{k_S}}$. In turn, the inventory level of this item at time t is

$$I_{i_S}(\tau_{i_S}^*, t) \geq H_{i_S} \cdot \left(1 - \frac{T_{i_S}/p_{k_S}}{T_{i_S}}\right) = 1 - \frac{1}{p_{k_S}} \geq 1 - \epsilon ,$$

where the last inequality holds since $p_{k_s} \ge K \ge e^{1/\epsilon^2}$. In conclusion,

$$OPT^{disc}(\mathcal{I}_{\hat{\mathcal{F}}}) = I_{\max}(\tau^*) \geq I_{\Sigma}(\tau^*, t) = \sum_{S \in \hat{\mathcal{F}}} I_{i_S}(\tau^*_{i_S}, t) \geq (1 - \epsilon) \cdot |\hat{\mathcal{F}}| .$$

3 Approximation Scheme in Terms of Cycle Length

This section is dedicated to presenting a randomized algorithm for discrete inventory staggering, showing that as long as our cycle length $\Lambda = \text{LCM}(T_1, \ldots, T_n)$ is not exponentially large, the optimal peak inventory level can be efficiently approached within any degree of accuracy. Formally, as stated in Theorem 1.3, we will argue that for any $\epsilon \in (0, 1)$, the discrete setting can be approximated within factor $1 + \epsilon$ of optimal via an $O(\Lambda^{\tilde{O}(1/\epsilon^3)} \cdot |\mathcal{I}|^{O(1)})$ -time algorithm, which is successful with probability at least $\frac{1}{2}$.

3.1 $O(\frac{1}{\epsilon})$ -sized discretization sets

For this purpose, rather than allowing the shift τ_i of each item $i \in [n]$ to take arbitrary integer values in $[0, T_i]$, we will restrict our choice to an $O(\frac{1}{\epsilon})$ -sized set of values \mathcal{D}_i as follows.

Construction of \mathcal{D}_i . Given an error parameter $\epsilon \in (0, 1)$, we say that item *i* is large when $T_i > \frac{1}{\epsilon}$; in the opposite case, this item is said to be small. From this point on, \mathcal{L} and \mathcal{S} will respectively denote the sets of large and small items. We proceed by associating each item $i \in [n]$ with a discrete set of values \mathcal{D}_i , defined according to the next case disjunction:

- When $i \in S$: Here, we keep our set of options as $\mathcal{D}_i = \{0, 1, \dots, T_i\}$.
- When $i \in S$: In this scenario, on top of the singular value T_i , only rounded-down multiples of ϵT_i will be allowed, specifically meaning that $\mathcal{D}_i = \{\lfloor \epsilon T_i \rfloor, \lfloor 2\epsilon T_i \rfloor, \ldots, \lfloor \lfloor \frac{1}{\epsilon} \rfloor \cdot \epsilon T_i \rfloor, T_i\}.$

In any case, we clearly have $|\mathcal{D}_i| \leq \lfloor \frac{1}{\epsilon} \rfloor + 1$. With these definitions, a shift vector $\tau = (\tau_1, \ldots, \tau_n)$ is called \mathcal{D} -restricted when $\tau_i \in \mathcal{D}_i$ for every item *i*.

Analysis. For the remainder of this section, τ^* will stand for some fixed optimal shift vector; without loss of generality, we assume that $\tau_i^* \in [0, T_i - 1] \cap \mathbb{Z}$ for every item $i \in [n]$. The next claim shows that, in terms of their peak inventory level, \mathcal{D} -restricted shift vectors are capable of competing against the unrestricted optimal vector τ^* .

LEMMA 3.1. There exists a \mathcal{D} -restricted shift vector τ for which $I_{\max}(\tau) \leq (1+4\epsilon) \cdot I_{\max}(\tau^*)$.

PROOF. Let us consider the shift vector τ , defined by setting $\tau_i = \lceil \tau_i^* \rceil^{\mathcal{D}_i}$ for every item $i \in [n]$. Here, $\lceil \cdot \rceil^{\mathcal{D}_i}$ is an operator that rounds its argument up to the nearest point in \mathcal{D}_i , implying that τ is \mathcal{D} -restricted. With respect to this definition, we clearly have $\tau_i = \tau_i^*$ for every item $i \in \mathcal{S}$, meaning in particular that $I_i(\tau_i, t) = I_i(\tau_i^*, t)$ at any time point $t \in [0, \Lambda]$. For every item $i \in \mathcal{L}$, while we generally have $\tau_i \geq \tau_i^*$, note that

$$I_i(\tau_i, t) \leq I_i(\tau_i^*, t) + \frac{H_i}{T_i} \cdot (\lceil \tau_i^* \rceil^{\mathcal{D}_i} - \tau_i^*)$$

$$\leq I_i(\tau_i^*, t) + \frac{H_i}{T_i} \cdot (\epsilon T_i + 1)$$

$$\leq I_i(\tau_i^*, t) + 2\epsilon H_i ,$$

where the last inequality holds since $T_i > \frac{1}{\epsilon}$. By summing these relations over all items, we get $I_{\Sigma}(\tau, t) \leq I_{\Sigma}(\tau^*, t) + 2\epsilon H_{\Sigma}$, and therefore, $I_{\max}(\tau) \leq I_{\max}(\tau^*) + 2\epsilon H_{\Sigma} \leq (1 + 4\epsilon) \cdot I_{\max}(\tau^*)$, noting that the last inequality follows from Lemma 2.1.

3.2 The LP-rounding algorithm

Step 1: Guessing the peak inventory level. Let $\tau^{\mathcal{D}}$ be a \mathcal{D} -restricted near-optimal shift vector, whose existence is guaranteed by Lemma 3.1, meaning that $I_{\max}(\tau^{\mathcal{D}}) \leq (1+4\epsilon) \cdot I_{\max}(\tau^*)$. As our first algorithm step, we guess an over-estimate \widetilde{OPT} for the peak inventory level $I_{\max}(\tau^{\mathcal{D}})$, such that $\widetilde{OPT} \in [1, 1+\epsilon] \cdot I_{\max}(\tau^{\mathcal{D}})$. Since we clearly have $I_{\max}(\tau^{\mathcal{D}}) \leq H_{\Sigma}$, and since $I_{\max}(\tau^{\mathcal{D}}) \geq I_{\max}(\tau^*) \geq I_{\max}(\tau^*) \geq \frac{H_{\Sigma}}{2}$ by Lemma 2.1, there are only $O(\frac{1}{\epsilon})$ candidate values to play the role of \widetilde{OPT} . Step 2: Guessing heavy contributions. Letting $\delta = \frac{\epsilon^3}{36 \ln(2\Lambda)}$, we say that item $i \in [n]$ is heavy when $H_i \geq \delta H_{\Sigma}$; otherwise, this item is said to be light. In what follows, \mathcal{H} and \mathcal{L} will respectively denote the sets of heavy and light items, noting that $|\mathcal{H}| \leq \frac{1}{\delta}$. Our next step consists of guessing the shift $\tau_i^{\mathcal{D}} \in \mathcal{D}_i$ of every item $i \in \mathcal{H}$ with respect to $\tau^{\mathcal{D}}$. Over all heavy items, the number of required guesses is only

$$\prod_{i \in \mathcal{H}} |\mathcal{D}_i| \leq \left(\left\lfloor \frac{1}{\epsilon} \right\rfloor + 1 \right)^{|\mathcal{H}|} = \left(\frac{1}{\epsilon} \right)^{O(1/\delta)} = \Lambda^{\tilde{O}(1/\epsilon^3)}$$

Step 3: The feasibility LP. With these parameters in place, we proceed by considering the following linear feasibility problem:

(1)
$$\sum_{i \in [n]} \sum_{\tau \in \mathcal{D}_{i}} I_{i\tau}^{t} x_{i\tau} \leq \widetilde{\text{OPT}} \qquad \forall t \in [0, \Lambda] \cap \mathbb{Z}$$

(2)
$$\sum_{\tau \in \mathcal{D}_{i}} x_{i\tau} = 1 \qquad \forall i \in [n]$$

(3)
$$x_{i\tau_{i}}^{\mathcal{D}} = 1 \qquad \forall i \in \mathcal{H}$$

(4)
$$x_{i\tau} \geq 0 \qquad \forall i \in [n], \tau \in \mathcal{D}_{i}$$

In an integer-valued solution to this formulation, each decision variable $x_{i\tau}$ is necessarily binary, indicating whether we choose a shift of $\tau \in \mathcal{D}_i$ for item *i*. As such, constraint (1) states that the combined inventory level at any point *t* is at most \overrightarrow{OPT} . Here, we make use of $I_{i\tau}^t$ as a shorthand notation for $I_i(\tau, t)$, which is the inventory level of item *i* at time *t* with respect to the policy \mathcal{P}_{τ}^i . Constraint (2) can be understood as choosing exactly one of the shifts in \mathcal{D}_i for each item *i*. Finally, constraint (3) forces any feasible solution to be aligned with our guesses for the shifts of heavy items. It is not difficult to verify that (LP) is feasible; for example, by setting $x_{i\tau} = 1$ if and only if $\tau_i^{\mathcal{D}} = \tau$, for every item $i \in [n]$ and shift $\tau \in \mathcal{D}_i$, we obtain a feasible solution.

Step 4: Randomized rounding. Let x^* be a feasible solution to (LP). For each item $i \in [n]$, we draw its random shift τ_i^R , such that each $\tau \in \mathcal{D}_i$ is selected with probability $x_{i\tau}^*$. This choice is made independently of any other item. As a side note, due to constraints (2) and (4), the quantities $\{x_{i\tau}^*\}_{\tau \in \mathcal{D}_i}$ are all non-negative and sum up to 1, meaning that they can indeed be utilized as probabilities.

Step 5: Peak evaluation. Due to the guessing procedure described in steps 1 and 2, we have constructed up until now a random collection of $\Lambda^{\tilde{O}(1/\epsilon^3)}$ shift vectors. Out of these vectors, we return one whose peak inventory level is minimal. Even though the question of efficient peak evaluation is still wide open (see Section 1.2), it can naively be resolved in $\Lambda n^{O(1)}$ time, simply by enumerating over all integer time points in $[0, \Lambda]$.

3.3 Analysis

As stated in Lemma 3.2 below, we proceed by proving that, when steps 3-5 are employed with the correct set of guesses, the inventory level of our random policy \mathcal{P}_{τ^R} at any given time point is at most $(1+2\epsilon) \cdot \widetilde{\text{OPT}}$, with probability at least $1 - \frac{1}{2\Lambda}$. By taking the union bound over all time points $t \in [0, \Lambda - 1] \cap \mathbb{Z}$, we infer that $\Pr[I_{\max}(\tau^R) \leq (1 + 2\epsilon) \cdot \widetilde{\text{OPT}}] \geq \frac{1}{2}$. As explained in step 1, we know that $\widetilde{\text{OPT}} \leq (1 + \epsilon) \cdot I_{\max}(\tau^{\mathcal{D}})$ and that $I_{\max}(\tau^{\mathcal{D}}) \leq (1 + 4\epsilon) \cdot I_{\max}(\tau^*)$. Therefore, with probability at least $\frac{1}{2}$, the peak inventory level of the random policy \mathcal{P}_{τ^R} is within factor $1 + 29\epsilon$ of optimal.

LEMMA 3.2. For every time point $t \in [0, \Lambda] \cap \mathbb{Z}$,

$$\Pr\left[I_{\Sigma}(\tau^{R}, t) \ge (1 + 2\epsilon) \cdot \widetilde{\text{OPT}}\right] \le \frac{1}{2\Lambda}$$

PROOF. For every item $i \in [n]$ and shift $\tau \in \mathcal{D}_i$, we introduce the random variable $X_{i\tau}$, as an indicator of the event $\{\tau_i^R = \tau\}$. In addition, the random variable Y_i^t will be defined as $Y_i^t = \sum_{\tau \in \mathcal{D}_i} I_{i\tau}^t X_{i\tau}$, which is the inventory level of item i at time t in terms of the random policy \mathcal{P}_{τ^R} . As such, the combined inventory level of \mathcal{P}_{τ^R} at time t can be written as $I_{\Sigma}(\tau^R, t) = \sum_{i \in [n]} Y_i^t$. With this notation, the bound we wish to prove is equivalent to

$$\Pr\left[\sum_{i \in [n]} Y_i^t \ge (1+2\epsilon) \cdot \widetilde{\text{OPT}}\right] \le \frac{1}{2\Lambda} .$$
(8)

To argue about $\sum_{i \in [n]} Y_i^t$, let us decompose this summation into the contributions of heavy and light items, i.e., $\sum_{i \in [n]} Y_i^t = \sum_{i \in \mathcal{H}} Y_i^t + \sum_{i \in \mathcal{L}} Y_i^t$. The next two claims, whose proofs are provided in Section 3.4, separately treat each of these terms.

CLAIM 3.3. $\sum_{i \in \mathcal{H}} Y_i^t = \sum_{i \in \mathcal{H}} \sum_{\tau \in \mathcal{D}_i} I_{i\tau}^t x_{i\tau}^*$ almost surely. CLAIM 3.4. $\Pr[\sum_{i \in \mathcal{L}} Y_i^t \ge \max\{\epsilon \widetilde{OPT}, (1+\epsilon) \cdot \sum_{i \in \mathcal{L}} \sum_{\tau \in \mathcal{D}_i} I_{i\tau}^t x_{i\tau}^*\}] \le \frac{1}{2\Lambda}$.

With these claims in place, we proceed to prove (8) by observing that

$$\Pr\left[\sum_{i\in[n]} Y_i^t \ge (1+2\epsilon) \cdot \widetilde{\operatorname{OPT}}\right]$$

$$= \Pr\left[\sum_{i\in\mathcal{H}} Y_i^t + \sum_{i\in\mathcal{L}} Y_i^t \ge (1+2\epsilon) \cdot \widetilde{\operatorname{OPT}}\right]$$

$$= \Pr\left[\sum_{i\in\mathcal{L}} Y_i^t \ge (1+2\epsilon) \cdot \widetilde{\operatorname{OPT}} - \sum_{i\in\mathcal{H}} \sum_{\tau\in\mathcal{D}_i} I_{i\tau}^t x_{i\tau}^*\right]$$
(9)

$$\leq \Pr\left[\sum_{i\in\mathcal{L}}Y_i^t \ge \widehat{\operatorname{OPT}} + (1+\epsilon) \cdot \sum_{i\in\mathcal{L}}\sum_{\tau\in\mathcal{D}_i}I_{i\tau}^t x_{i\tau}^*\right]$$
(10)

$$\leq \Pr\left[\sum_{i \in \mathcal{L}} Y_i^t \geq \max\left\{ \widehat{\epsilon \text{OPT}}, (1+\epsilon) \cdot \sum_{i \in \mathcal{L}} \sum_{\tau \in \mathcal{D}_i} I_{i\tau}^t x_{i\tau}^* \right\} \right]$$

$$\leq \frac{1}{2\Lambda} . \tag{11}$$

Here, equality (9) follows from Claim 3.3. Inequality (10) holds since $\sum_{i \in \mathcal{H}} \sum_{\tau \in \mathcal{D}_i} I_{i\tau}^t x_{i\tau}^* + \sum_{i \in \mathcal{L}} \sum_{\tau \in \mathcal{D}_i} I_{i\tau}^t x_{i\tau}^* \leq \widetilde{\text{OPT}}$, by constraint (1) of (LP). Finally, inequality (11) is precisely the

result stated in Claim 3.4.

3.4 Additional proofs

Proof of Claim 3.3. The important observation is that, for every item $i \in \mathcal{H}$, we have $Y_i^t = \sum_{\tau \in \mathcal{D}_i} I_{i\tau}^t X_{i\tau} = \sum_{\tau \in \mathcal{D}_i} I_{i\tau}^t x_{i\tau}^*$. Here, the second equality is obtained by noting that $X_{i\tau} = 1[\tau_i^R = \tau] = x_{i\tau}^*$ for every $\tau \in \mathcal{D}_i$, since constraint (3) is forcing us to choose $\tau_i^R = \tau_i^\mathcal{D}$ with probability $x_{i\tau}^* = 1$. It follows that $\sum_{i \in \mathcal{H}} Y_i^t = \sum_{\tau \in \mathcal{D}_i} I_{i\tau}^t x_{i\tau}^*$ almost surely.

Proof of Claim 3.4. Our proof considers two cases, depending on the relation between $\mathbb{E}[\sum_{i \in \mathcal{L}} Y_i^t]$ and $\widetilde{\text{OPT}}$.

Case 1: $\mathbb{E}[\sum_{i \in \mathcal{L}} Y_i^t] \leq \frac{2\epsilon}{3} \cdot \widetilde{\text{OPT}}$. In this scenario,

$$\Pr\left[\sum_{i\in\mathcal{L}} Y_i^t \ge \widehat{\epsilon \text{OPT}}\right] = \Pr\left[\sum_{i\in\mathcal{L}} \frac{Y_i^t}{\delta H_{\Sigma}} \ge \left(1 + \frac{1}{2}\right) \cdot \frac{\frac{2\epsilon}{3} \cdot \widetilde{\text{OPT}}}{\delta H_{\Sigma}}\right] \le \exp\left(-\frac{\epsilon}{18\delta} \cdot \frac{\widetilde{\text{OPT}}}{H_{\Sigma}}\right)$$
(12)

$$\leq \exp\left(-\frac{\epsilon}{36\delta}\right) \tag{13}$$

$$\leq \frac{1}{2\Lambda}$$
 (14)

Here, inequality (12) follows by specializing the Chernoff-Hoeffding bound (2) with respect to the random variables $\{\frac{Y_i^t}{\delta H_{\Sigma}}\}_{i\in\mathcal{L}}$, which are independent and [0, 1]-bounded. Indeed, the former property is satisfied since the random shift of each item is independently selected, whereas the latter property is obtained by observing that $Y_i^t \leq H_i \leq \delta H_{\Sigma}$, since item *i* is light. In addition, by the case hypothesis, $\mathbb{E}[\sum_{i\in\mathcal{L}} Y_i^t] \leq \frac{2\epsilon}{3} \cdot \widetilde{OPT}$. Inequality (13) is obtained by noting that $\widetilde{OPT} \geq I_{\max}(\tau^{\mathcal{D}}) \geq \frac{H_{\Sigma}}{2}$, where the last inequality follows from Lemma 2.1, stating that $I_{\max}(\tau^*) \geq \frac{H_{\Sigma}}{2}$. Finally, inequality (14) holds since $\delta = \frac{\epsilon^3}{36 \ln(2\Lambda)}$.

Case 2: $\mathbb{E}[\sum_{i \in \mathcal{L}} Y_i^t] > \frac{2\epsilon}{3} \cdot \widetilde{\text{OPT}}$. To arrive at the desired bound, note that since $Y_i^t = \sum_{\tau \in \mathcal{D}_i} I_{i\tau}^t X_{i\tau}$ and since $\mathbb{E}[X_{i\tau}] = \Pr[\tau_i^R = \tau] = x_{i\tau}^*$, we have

$$\Pr\left[\sum_{i\in\mathcal{L}} Y_i^t \ge (1+\epsilon) \cdot \sum_{i\in\mathcal{L}} \sum_{\tau\in\mathcal{D}_i} I_{i\tau}^t x_{i\tau}^*\right]$$

=
$$\Pr\left[\sum_{i\in\mathcal{L}} \frac{Y_i^t}{\delta H_{\Sigma}} \ge (1+\epsilon) \cdot \mathbb{E}\left[\sum_{i\in\mathcal{L}} \frac{Y_i^t}{\delta H_{\Sigma}}\right]\right]$$

$$\le \exp\left(-\frac{\epsilon^2}{3\delta} \cdot \frac{\mathbb{E}[\sum_{i\in\mathcal{L}} Y_i^t]}{H_{\Sigma}}\right)$$
(15)

$$\leq \exp\left(-\frac{2\epsilon^3}{9\delta} \cdot \frac{\widetilde{\text{OPT}}}{H_{\Sigma}}\right) \tag{16}$$

$$\leq \exp\left(-\frac{\epsilon^3}{9\delta}\right) \tag{17}$$

$$\leq \quad \frac{1}{2\Lambda} \ . \tag{18}$$

As in case 1, inequality (15) follows by specializing (2) with respect to $\{\frac{Y_i^t}{\delta H_{\Sigma}}\}_{i \in \mathcal{L}}$. Inequality (16) holds due to our case hypothesis. Inequality (17) is obtained by noting once again that $\widetilde{OPT} \geq \frac{H_{\Sigma}}{2}$. Finally, we arrive at inequality (18) by plugging in $\delta = \frac{\epsilon^3}{36 \ln(2\Lambda)}$.

4 Approximation Scheme in Terms of Distinct Time Intervals

In this section, we devise a deterministic approach for discrete inventory staggering, culminating to a polynomial-time approximation scheme in the scenario of constantly-many distinct time intervals. Technically speaking, as stated in Theorem 1.4, we will show that for any $\epsilon \in (0, 1)$, the discrete setting can be approximated within factor $1 + \epsilon$ of optimal in $O(2^{\tilde{O}(K/\epsilon^2)} \cdot |\mathcal{I}|^{O(1)})$ time, where K stands for the number of distinct time intervals across all items.

4.1 $O(n^{O(n)} \cdot |\mathcal{I}|^{O(1)})$ -time peak evaluation oracle

As previously mentioned, given an arbitrary shift vector τ , the question of evaluating the peak inventory level $I_{\max}(\tau)$ of the replenishment policy \mathcal{P}_{τ} in polynomial time is still open. That said, since we will be dealing with $O(\frac{K}{\epsilon})$ super-items in Sections 4.2 and 4.3, designing an exponential-in-*n* peak evaluation oracle will suffice for our particular purposes.

THEOREM 4.1. For any integer-valued shift vector τ , we can compute the peak inventory level of \mathcal{P}_{τ} in $O(n^{O(n)} \cdot |\mathcal{I}|^{O(1)})$ time.

ILP-formulation. To derive this result, given a shift vector $\tau \in \mathbb{Z}^n$, we assume without loss of generality that $\tau_i \in [0, T_i - 1]$ for every item $i \in [n]$. Let us consider the following integer linear program, where our decision variables are $p \in \mathbb{Z}$ and $x \in \mathbb{Z}^n$:

$$\max \sum_{i \in [n]} H_i \cdot \left(1 - \frac{p - (\tau_i + x_i T_i)}{T_i} \right)$$
s.t. (1) $\tau_i + x_i T_i \leq p$ $\forall i \in [n]$ (IP)
(2) $0 \leq p \leq \Lambda$
(3) $p \in \mathbb{Z}, x \in \mathbb{Z}^n$

To develop some basic intuition around this formulation, one can think of p as the time point at which the peak inventory level $I_{\max}(\tau)$ is attained. From this perspective, constraints (2) and (3) jointly ensure that p is an integer-valued point within $[0, \Lambda]$. Regarding the x-variables, for every item $i \in [n]$, the value of $\tau_i + x_i T_i$ represents the last *i*-ordering point in $(-\infty, p]$ with respect to $\mathcal{P}^i_{\tau_i}$. As such, the inventory level of this item at time p is given by $H_i \cdot (1 - \frac{p - (\tau_i + x_i T_i)}{T_i})$, which is precisely the term we are seeing in the objective function. Along these lines, constraint (1) guarantees that $\tau_i + x_i T_i$ falls within $(-\infty, p]$, and moreover, since x_i is coupled with a coefficient of $H_i > 0$ in the objective function, an optimal solution will indeed set $\tau_i + x_i T_i$ as the last *i*-ordering point in $(-\infty, p]$. **Basic properties of (IP).** The next claim explicitly constructs a feasible solution to (IP) and prescribes lower and upper bounds on the optimum value of this program.

LEMMA 4.2. (IP) is feasible, and moreover, $OPT(IP) \in [I_{max}(\tau), H_{\Sigma}]$.

PROOF. To argue about the feasibility of (IP), let us construct a candidate solution (\hat{p}, \hat{x}) as follows. First, \hat{p} is chosen as one of the points in $[0, \Lambda] \cap \mathbb{Z}$ for which $I_{\Sigma}(\tau, \hat{p}) = I_{\max}(\tau)$. Next, for every item $i \in [n]$, let $\hat{x}_i = \max\{x \in \mathbb{Z} : \tau_i + xT_i \leq \hat{p}\}$; this choice ensures that $\tau_i + \hat{x}_iT_i$ is precisely the time point where the last *i*-order in $(-\infty, \hat{p}]$ occurs, with respect to the policy $\mathcal{P}^i_{\tau_i}$. Clearly, (\hat{p}, \hat{x}) is a feasible solution to (IP), with an objective value of

$$\sum_{i \in [n]} H_i \cdot \left(1 - \frac{\hat{p} - (\tau_i + \hat{x}_i T_i)}{T_i} \right) = \sum_{i \in [n]} I_i(\tau_i, \hat{p}) = I_{\Sigma}(\tau, \hat{p}) = I_{\max}(\tau) ,$$

implying that we indeed have $OPT(IP) \ge I_{max}(\tau)$.

In the opposite direction, to prove that $OPT(IP) \leq H_{\Sigma}$, suppose that (p, x) is a feasible solution to (IP). Then, its objective value can be upper-bounded by noting that

$$\sum_{i \in [n]} H_i \cdot \left(1 - \frac{p - (\tau_i + x_i T_i)}{T_i} \right) \leq \sum_{i \in [n]} H_i = H_{\Sigma}$$

where the inequality above holds since $p \ge \tau_i + x_i T_i$, by constraint (1).

Inventory peaks via (IP). By Lemma 4.2, we know that (IP) is feasible and bounded, meaning that this formulation indeed has one or more optimal solutions. In what follows, we prove that with respect to any optimal solution (p^*, x^*) , the time point p^* must be one where the inventory level of \mathcal{P}_{τ} is maximized.

LEMMA 4.3. Let (p^*, x^*) be an optimal solution to (IP). Then, $I_{\Sigma}(\tau, p^*) = I_{\max}(\tau)$.

PROOF. Suppose on the contrary that $I_{\Sigma}(\tau, p^*) < I_{\max}(\tau)$, and let us examine the implications of (p^*, x^*) being an optimal solution to (IP). In particular, for every item $i \in [n]$, we must have $x_i^* = \max\{x \in \mathbb{Z} : \tau_i + xT_i \leq p^*\}$. Otherwise, x_i^* can be incremented by 1, to obtain a feasible solution whose objective value is strictly greater than that of (p^*, x^*) . As a result, $\tau_i + x_i^*T_i$ is precisely the time point where the last *i*-order in $(-\infty, p^*]$ occurs, with respect to the policy $\mathcal{P}_{\tau_i}^i$, and in turn, $H_i \cdot (1 - \frac{p^* - (\tau_i + x_i^*T_i)}{T_i}) = I_i(\tau_i, p^*)$. We are now ready to reveal the resulting contradiction, arguing that (p^*, x^*) cannot be an optimal solution, by proving that its objective value is strictly smaller than OPT(IP). To this end, note that

$$\sum_{i \in [n]} H_i \cdot \left(1 - \frac{p^* - (\tau_i + x_i^* T_i)}{T_i} \right) = \sum_{i \in [n]} I_i(\tau_i, p^*)$$
$$= I_{\Sigma}(\tau, p^*)$$
$$< I_{\max}(\tau)$$
(19)

$$\leq \text{OPT}(\text{IP})$$
. (20)

Here, inequality (19) corresponds to the scenario we are considering, where $I_{\Sigma}(\tau, p^*) < I_{\max}(\tau)$. Inequality (20) is precisely our lower bound on OPT(IP), as stated in Lemma 4.2. **Solving (IP).** In summary, to compute the peak inventory level of \mathcal{P}_{τ} , it remains to identify p^* , for some optimal solution (p^*, x^*) . The next claim, whose proof is provided in Section 4.4, argues that this goal can be accomplished in $O(n^{O(n)} \cdot |\mathcal{I}|^{O(1)})$ time.

LEMMA 4.4. For some optimal solution (p^*, x^*) to (IP), we can determine the value of p^* in $O(n^{O(n)} \cdot |\mathcal{I}|^{O(1)})$ time.

4.2 Mimicking partitions: Definitions and usefulness

As previously alluded to, our algorithmic approach will rely on aggregating the overall set of items into a small-sized collection of $O(\frac{K}{\epsilon})$ super-items, which will be coupled with uniform shifts. For this purpose, let $T_{(1)}, \ldots, T_{(K)}$ be the set of distinct time intervals across T_1, \ldots, T_n . For every $k \in [K]$, we use $S_k = \{i \in [n] : T_i = T_{(k)}\}$ to denote the collection of items associated with a time interval of $T_{(k)}$. Circling back to Section 3.1, we remind the reader that there exists a \mathcal{D} -restricted shift vector $\tau^{\mathcal{D}}$ for which $I_{\max}(\tau^{\mathcal{D}}) \leq (1 + 4\epsilon) \cdot \text{OPT}^{\text{disc}}(\mathcal{I})$. Moreover, for every $k \in [K]$, all items in S_k share the same discretization set, which will be designated by $\mathcal{D}_{(k)}$, noting that $|\mathcal{D}_{(k)}| \leq \lfloor \frac{1}{\epsilon} \rfloor + 1$.

Mimicking partitions. For every $k \in [K]$ and $\tau \in \mathcal{D}_{(k)}$, let $S_{k,\tau}^{\mathcal{D}} = \{i \in S_k : \tau_i^{\mathcal{D}} = \tau\}$ be the set of items in S_k whose shift with respect to $\tau^{\mathcal{D}}$ is precisely τ . Clearly, the sets $\{S_{k,\tau}^{\mathcal{D}}\}_{k\in[K],\tau\in\mathcal{D}_{(k)}}$ form a partition of the underlying items, which is unknown from an algorithmic perspective. To go around this obstacle, we say that a partition $\{S_{k,\tau}^{\text{APX}}\}_{k\in[K],\tau\in\mathcal{D}_{(k)}}$ is mimicking $\{S_{k,\tau}^{\mathcal{D}}\}_{k\in[K],\tau\in\mathcal{D}_{(k)}}$ when it satisfies the next two properties:

- 1. Marginal partition: $\{S_{k,\tau}^{\text{APX}}\}_{\tau \in \mathcal{D}_{(k)}}$ is a partition of S_k , for every $k \in [K]$.
- 2. Cumulative ordering quantity: For every $k \in [K]$ and $\tau \in \mathcal{D}_{(k)}$,

$$H(S_{k,\tau}^{\text{APX}}) \leq (1+\epsilon) \cdot H(S_{k,\tau}^{\mathcal{D}}) + \frac{2\epsilon^2}{K} \cdot H_{\Sigma}.$$

Usefulness. To better understand the motivation behind introducing this notion, suppose that $\{S_{k,\tau}^{APX}\}_{k\in[K],\tau\in\mathcal{D}_{(k)}}$ is mimicking $\{S_{k,\tau}^{\mathcal{D}}\}_{k\in[K],\tau\in\mathcal{D}_{(k)}}$. In this case, let use define a shift vector τ^{APX} such that, for every index $k \in [K]$ and time interval $\tau \in \mathcal{D}_{(k)}$, all items $i \in S_{k,\tau}^{APX}$ are given a uniform shift of $\tau_i^{APX} = \tau$. The next claim shows that τ^{APX} is near-optimal, meaning that we have just reduced our original inventory staggering problem to the computational question of identifying a mimicking partition.

LEMMA 4.5. $I_{\max}(\tau^{APX}) \leq (1+17\epsilon) \cdot \text{OPT}^{\text{disc}}(\mathcal{I}).$

PROOF. To establish the desired claim, consider some index $k \in [K]$, time interval $\tau \in \mathcal{D}_{(k)}$, and integer-valued point $t \in [0, \Lambda]$. Let t^- be the last $S_{k,\tau}^{\text{APX}}$ -order in $(-\infty, t]$ with respect to τ^{APX} , namely, $t^- = \max\{\tau + mT_{(k)} : \tau + mT_{(k)} \leq t, m \in \mathbb{Z}\}$. According to this definition, $I_i(\tau_i^{\text{APX}}, t) = H_i \cdot (1 - \frac{t-t^-}{T_{(k)}})$ for every item $i \in S_{k,\tau}^{\text{APX}}$. Therefore,

$$\sum_{i \in S_{k,\tau}^{\text{APX}}} I_i(\tau_i^{\text{APX}}, t) = \left(1 - \frac{t - t^-}{T_{(k)}}\right) \cdot H(S_{k,\tau}^{\text{APX}})$$

$$\leq \left(1 - \frac{t - t^{-}}{T_{(k)}}\right) \cdot \left((1 + \epsilon) \cdot H(S_{k,\tau}^{\mathcal{D}}) + \frac{2\epsilon^{2}}{K} \cdot H_{\Sigma}\right)$$
(21)

$$\leq (1+\epsilon) \cdot \sum_{i \in S_{k,\tau}^{\mathcal{D}}} I_i(\tau_i^{\mathcal{D}}, t) + \frac{2\epsilon^2}{K} \cdot H_{\Sigma} , \qquad (22)$$

where inequality (21) follows from property 2, and inequality (22) holds since $I_i(\tau_i^{\mathcal{D}}, t) = H_i \cdot (1 - \frac{t-t^-}{T_{(k)}})$ for every item $i \in S_{k,\tau}^{\mathcal{D}}$. By summing this relation across all $k \in [K]$ and $\tau \in \mathcal{D}_{(k)}$, we have

$$I_{\Sigma}(\tau^{\text{APX}}, t) = \sum_{k \in [K]} \sum_{\tau \in \mathcal{D}_{(k)}} \sum_{i \in S_{k,\tau}^{\text{APX}}} I_i(\tau_i^{\text{APX}}, t)$$
(23)

$$\leq \sum_{k \in [K]} \sum_{\tau \in \mathcal{D}_{(k)}} \left((1+\epsilon) \cdot \sum_{i \in S_{k,\tau}^{\mathcal{D}}} I_i(\tau_i^{\mathcal{D}}, t) + \frac{2\epsilon^2}{K} \cdot H_{\Sigma} \right)$$
(24)

$$\leq (1+\epsilon) \cdot I_{\max}(\tau^{\mathcal{D}}) + 4\epsilon H_{\Sigma}$$
⁽²⁵⁾

$$\leq (1+17\epsilon) \cdot \operatorname{OPT}^{\operatorname{disc}}(\mathcal{I}) . \tag{26}$$

Here, equality (23) holds since $\{S_{k,\tau}^{\text{APX}}\}_{k\in[K],\tau\in\mathcal{D}_{(k)}}$ is a partition of the entire item set, by property 1. Inequality (24) follows by plugging in (22). Inequality (25) is obtained by recalling that $|\mathcal{D}_{(k)}| \leq \lfloor \frac{1}{\epsilon} \rfloor + 1 \leq \frac{2}{\epsilon}$. Finally, to arrive at inequality (26), note that $I_{\max}(\tau^{\mathcal{D}}) \leq (1+4\epsilon) \cdot \text{OPT}^{\text{disc}}(\mathcal{I})$ and that $H_{\Sigma} \leq 2 \cdot \text{OPT}^{\text{disc}}(\mathcal{I})$, by Lemma 2.1.

4.3 Mimicking partitions: Construction and testing

In what follows, we construct a family of $2^{O(K/\epsilon^2)}$ partitions, out of which at least one is guaranteed to be mimicking $\{S_{k,\tau}^{\mathcal{D}}\}_{k\in[K],\tau\in\mathcal{D}_{(k)}}$; this construction can be implemented in $O(2^{O(K/\epsilon^2)} \cdot n^{O(1)})$ time. Subsequently, we will explain how to utilize this family in order to compute a near-optimal shift vector.

Step 1: Guessing. As our first algorithmic step, for every $k \in [K]$ and $\tau \in \mathcal{D}_{(k)}$, we guess an over-estimate $H_{k,\tau}^{\text{APX}}$ for $H(S_{k,\tau}^{\mathcal{D}})$, such that

$$H(S_{k,\tau}^{\mathcal{D}}) < H_{k,\tau}^{\text{APX}} \leq H(S_{k,\tau}^{\mathcal{D}}) + \frac{\epsilon^2}{K} \cdot H_{\Sigma}$$
 (27)

To this end, it suffices to enumerate over integer multiples of $\frac{\epsilon^2}{K} \cdot H_{\Sigma}$ within the interval $[0, H_{\Sigma}]$. Based on elementary balls-into-bins arguments, across all combinations of $k \in [K]$ and $\tau \in \mathcal{D}_{(k)}$, the required number of guesses is $2^{O(K/\epsilon^2)}$.

Step 2: Load balancing formulation. For every $k \in [K]$, we proceed by separately partitioning the set of items S_k into $\{S_{k,\tau}^{\text{APX}}\}_{\tau \in \mathcal{D}_{(k)}}$. For this purpose, we create the following instance of makespan minimization on related (uniform) machines:

• Jobs: Each item $i \in S_k$ corresponds to a job, requiring H_i amount of work.

- Machines: For every $\tau \in \mathcal{D}_{(k)}$, we represent the subset $S_{k,\tau}^{\text{APX}}$ via a corresponding machine $\mathcal{M}_{k,\tau}$, whose speed is $H_{k,\tau}^{\text{APX}}$. As such, when job *i* is assigned to this machine, it incurs a processing time of $\frac{H_i}{H_{\lambda}^{\text{APX}}}$.
- *Goal*: Our objective is to compute a job-to-machine assignment whose makespan is minimized.

It is easy to see that the optimal makespan for this particular instance is at most 1. Indeed, when for every $\tau \in \mathcal{D}_{(k)}$, the unknown set of jobs $S_{k,\tau}^{\mathcal{D}}$ is assigned to machine $\mathcal{M}_{k,\tau}$, the latter incurs a total processing time of $\frac{H(S_{k,\tau}^{\mathcal{D}})}{H_{h}^{\text{APX}}} < 1$, due to inequality (27).

Step 3: Approximating the load balancing problem. Luckily, makespan minimization on related machines admits an efficient polynomial-time approximation scheme (EPTAS). Specifically, Jansen et al. (2020) showed how to compute an assignment whose makespan is within factor $1 + \epsilon$ of optimal; their approach can be implemented in $O(2^{\tilde{O}(1/\epsilon)} + n^{O(1)})$ time. This result improves on a slightly less efficient approach by Jansen (2010). Migrated to our context, it follows that within this running time, we can partition S_k into $\{S_{k,\tau}^{\text{APX}}\}_{\tau \in \mathcal{D}_{(k)}}$, such that for every $\tau \in \mathcal{D}_{(k)}$,

$$H(S_{k,\tau}^{\text{APX}}) \leq (1+\epsilon) \cdot H_{k,\tau}^{\text{APX}} \leq (1+\epsilon) \cdot H(S_{k,\tau}^{\mathcal{D}}) + \frac{2\epsilon^2}{K} \cdot H_{\Sigma} ,$$

where the last transition follows from inequality (27). In other words, $\{S_{k,\tau}^{\text{APX}}\}_{k\in[K],\tau\in\mathcal{D}_{(k)}}$ is a mimicking partition for $\{S_{k,\tau}^{\mathcal{D}}\}_{k\in[K],\tau\in\mathcal{D}_{(k)}}$.

Testing. Up until now, we have constructed a family of $2^{O(K/\epsilon^2)}$ partitions, containing at least one which is guaranteed to be mimicking $\{S_{k,\tau}^{\mathcal{D}}\}_{k\in[K],\tau\in\mathcal{D}_{(k)}}$. As shown in Section 4.2, the shift vector τ^{APX} corresponding to such a partition, $\{S_{k,\tau}^{\text{APX}}\}_{k\in[K],\tau\in\mathcal{D}_{(k)}}$, forms a $(1+17\epsilon)$ -approximation. Therefore, the remaining question is: Given τ^{APX} , how do we compute its peak inventory level, $I_{\text{max}}(\tau^{\text{APX}})$?

To resolve this question, we remind the reader that τ^{APX} is defined such that, for every index $k \in [K]$ and time interval $\tau \in \mathcal{D}_{(k)}$, all items in $S_{k,\tau}^{APX}$ are given a uniform shift of τ . Thus, to evaluate $I_{\max}(\tau^{APX})$, we can view $S_{k,\tau}^{APX}$ as a super-item, (k, τ) , associated with a time interval of $T_{(k)}$ and with an ordering quantity of $H(S_{k,\tau}^{APX})$. Consequently, it remains to compute the peak inventory level of a given shift vector, when the underlying number of items is $O(\frac{K}{\epsilon})$. By Theorem 4.1, $I_{\max}(\tau^{APX})$ can be evaluated in $O((\frac{K}{\epsilon})^{O(K/\epsilon)} \cdot |\mathcal{I}|^{O(1)}) = O(2^{\tilde{O}(K/\epsilon)} \cdot |\mathcal{I}|^{O(1)})$ time.

4.4 Proof of Lemma 4.4

Integer linear feasibility. For the purpose of computing an optimal solution to (IP), we will rely on the work of Lenstra (1983), Kannan (1987), and Frank and Tardos (1987) for solving integer linear feasibility problems in fixed dimension. Specifically, given an integer-valued matrix $A \in \mathbb{Z}^{m \times n}$ along with an integer-valued vector $b \in \mathbb{Z}^m$, suppose we wish to decide whether there exists a vector $x \in \mathbb{Z}^n$ satisfying $Ax \leq b$. Then, the above-mentioned papers aggregately proved that this problem is decidable in $O(n^{O(n)} \cdot L_{A,b})$ time, where $L_{A,b}$ represents the combined number of bits required to specify A and b.

The feasibility version of (IP). Prior to proceeding any further, we mention in passing that the upcoming procedure for optimizing (IP) is rather standard, and we refer interested readers to two excellent sources — Lokshtanov (2009, Sec. 2.8) and Cygan et al. (2015, Sec. 6.2) — for additional explanations about obtaining an optimal solution to a given integer linear program via its feasibility counterpart. In order to convert (IP) into a feasibility problem of the form described above, rather than having $\sum_{i \in [n]} H_i \cdot (1 - \frac{p - (\tau_i + x_i T_i)}{T_i})$ as our objective function, when this term is scaled by a factor of Λ and lower-bounded by ψ , we arrive at the next integer linear feasibility problem:

(1)
$$\Lambda \cdot \sum_{i \in [n]} H_i \cdot \left(1 - \frac{p - (\tau_i + x_i T_i)}{T_i} \right) \ge \psi$$
(2)
$$\tau_i + x_i T_i \le p \qquad \forall i \in [n] \qquad (IP_{\text{feasible}})$$
(3)
$$0 \le p \le \Lambda$$
(4)
$$p \in \mathbb{Z}, x \in \mathbb{Z}^n$$

Optimizing (IP) via (IP_{feasible}). Let us observe that, as long as ψ is an integer, the constraint matrix A of this formulation and its right-side vector b are integer-valued. Therefore, subsequently to the opening paragraph of this section, we know that (IP_{feasible}) can be decided in $O(n^{O(n)} \cdot L_{A,b,\psi})$ time. Here, the absolute value of every entry in A and b is at most $\Lambda nH_{\max}\psi$, with room to spare, implying that $L_{A,b,\psi} = O(n^2 \log(\Lambda nH_{\max}\psi))$.

Now, by Lemma 4.2, we know that $OPT(IP) \in [I_{max}(\tau), H_{\Sigma}] \subseteq [\frac{H_{\Sigma}}{2}, H_{\Sigma}]$. During the construction of (IP_{feasible}) , we have scaled our objective function by Λ , meaning that the exact optimum of (IP) can be identified via binary search over the interval $[\frac{\Lambda H_{\Sigma}}{2}, \Lambda H_{\Sigma}]$. This search will consist of $O(\log(\Lambda H_{\Sigma}))$ iterations, each employing a single call to (IP_{feasible}) with a threshold of $\psi = \Theta(\Lambda H_{\Sigma})$. As a result, OPT(IP) can be computed in time

$$O\left(n^{O(n)} \cdot \log^2(\Lambda H_{\max})\right) = O\left(n^{O(n)} \cdot \log^2(T_{\max}H_{\max})\right) = O\left(n^{O(n)} \cdot |\mathcal{I}|^{O(1)}\right) .$$

Next, our objective is to determine the value of p^* , for some optimal solution (p^*, x^*) to (IP). For this purpose, we augment (IP_{feasible}) with two constraints:

- First, in place of $\Lambda \cdot \sum_{i \in [n]} H_i \cdot (1 \frac{p (\tau_i + x_i T_i)}{T_i}) \ge \psi$, since OPT(IP) is already known, we focus on optimal solutions by writing $\Lambda \cdot \sum_{i \in [n]} H_i \cdot (1 \frac{p (\tau_i + x_i T_i)}{T_i}) = \Lambda \cdot \text{OPT}(\text{IP})$.
- Our second constraint is of the form $p \ge \psi$, intended to conduct a binary search for the value of p^* .

Similarly to the line of reasoning described above, p^* can be computed in $O(n^{O(n)} \cdot |\mathcal{I}|^{O(1)})$ time as well.

5 Approximation Scheme for Nested Instances

In this section, we develop a polynomial-time approximation scheme for nested instances of the inventory staggering problem, thereby establishing Theorem 1.5. Specifically, for any $\epsilon > 0$, our decomposition-based approach will approximate such instances within factor $1 + \epsilon$ of optimal in $O(2^{\tilde{O}(1/\epsilon^3)} \cdot |\mathcal{I}|^{O(1)})$ time. For ease of exposition, these ideas will be described in terms of the discrete problem formulation; we will therefore open by explaining how continuous instances can be reduced to discrete ones, with negligible loss in optimality.

5.1 Continuous-to-discrete reduction

Interestingly, the discretization ideas presented in Section 3.1 allow us to reduce any instance $\mathcal{I}_{\mathbb{R}^n}$ of the continuous model to a discrete instance, $\mathcal{I}_{\mathbb{Z}^n}$, losing a factor of $1+O(\epsilon)$ in its objective value. For this purpose, all time intervals T_1, \ldots, T_n are uniformly scaled up by a factor of $\lceil \frac{1}{\epsilon} \rceil$ and nothing more. This modification clearly preserves nestedness.

To analyze the relationship between $\mathcal{I}_{\mathbb{R}^n}$ and $\mathcal{I}_{\mathbb{Z}^n}$, let us first observe that due to uniformly up-scaling all time intervals, when scaled down by a factor of $\lceil \frac{1}{\epsilon} \rceil$, the peak inventory level of any integer-valued shift vector τ for $\mathcal{I}_{\mathbb{Z}^n}$ translates to precisely the same peak for $\mathcal{I}_{\mathbb{R}^n}$. In other words, $I_{\max}^{\mathcal{I}_{\mathbb{R}^n}}(\frac{\tau}{\lceil 1/\epsilon \rceil}) = I_{\max}^{\mathcal{I}_{\mathbb{Z}^n}}(\tau)$, implying that $OPT^{cont}(\mathcal{I}_{\mathbb{R}^n}) \leq OPT^{disc}(\mathcal{I}_{\mathbb{Z}^n})$. To relate these measures in the opposite direction, let $\tilde{T}_1, \ldots, \tilde{T}_n$ be the time intervals of $\mathcal{I}_{\mathbb{Z}^n}$, and suppose we associate each item $i \in [n]$ with a discrete set of integer-valued shifts, $\mathcal{D}_i = \{\lfloor \epsilon \tilde{T}_i \rfloor, \lfloor 2\epsilon \tilde{T}_i \rfloor, \ldots, \lfloor \lfloor \frac{1}{\epsilon} \rfloor \cdot \epsilon \tilde{T}_i \rfloor, \tilde{T}_i\}$. Then, minor modifications to the proof of Lemma 3.1 show that there exists a \mathcal{D} -restricted shift vector for $\mathcal{I}_{\mathbb{Z}^n}$ whose peak inventory level is $(1 + O(\epsilon)) \cdot OPT^{cont}(\mathcal{I}_{\mathbb{R}^n})$.

5.2 Well-separated partitions

Our approach for devising a polynomial-time approximation scheme is driven by the notion of well-separated partitions. To gradually introduce these objects, given a nested instance \mathcal{I} of the discrete inventory staggering problem, suppose that $\mathcal{S} = (S_1, \ldots, S_M, S_\infty)$ is an ordered partition of its item set. Then, \mathcal{S} is called well-separated when it satisfies the next three properties:

- 1. Width: $\frac{\max_{i \in S_m} T_i}{\min_{i \in S_m} T_i} \leq (\frac{1}{\epsilon})^{1/\epsilon}$, for every $m \in [M]$.
- 2. Separation: $T_{i_1} \leq \epsilon T_{i_2}$, for every $(i_1, i_2) \in S_{m_1} \times S_{m_2}$ with $1 \leq m_1 < m_2 \leq M$.
- 3. S_{∞} is negligible: $H(S_{\infty}) \leq \epsilon H_{\Sigma}$.

In what follows, we prove the existence of well-separated partitions and propose an efficient construction through a simple application of the probabilistic method (Alon and Spencer, 2016).

LEMMA 5.1. A well-separated partition exists, and can be deterministically computed in polynomial time.

PROOF. We begin by defining the sequence of time interval subsets

$$\mathcal{T}_1 = \left\{ i \in [n] : T_i \in \left[1, \frac{1}{\epsilon}\right] \right\}, \qquad \mathcal{T}_2 = \left\{ i \in [n] : T_i \in \left(\frac{1}{\epsilon}, \frac{1}{\epsilon^2}\right] \right\}, \qquad \dots$$

so on and so forth, where in general $\mathcal{T}_q = \left\{i \in [n] : T_i \in \left(\frac{1}{\epsilon^{q-1}}, \frac{1}{\epsilon^q}\right]\right\}$ for every $q \geq 2$. This sequence proceeds up to \mathcal{T}_Q , where Q is the smallest integer for which $\frac{1}{\epsilon^Q} \geq T_{\max}$, meaning that $Q = \left\lceil \log_{1/\epsilon} T_{\max} \right\rceil = O(\log T_{\max}).$

Assuming without loss of generality that $\frac{1}{\epsilon}$ is an integer, we randomly pick $\xi \sim U\{0, \ldots, \frac{1}{\epsilon} - 1\}$. With respect to this choice, any subset \mathcal{T}_q with $q \equiv \xi \pmod{\frac{1}{\epsilon}}$ is said to be marked, while all other subsets are unmarked. We use S_1^{ξ} to denote the collection of time intervals belonging to the (unmarked) subsets appearing prior to the first marked subset (i.e., $\mathcal{T}_1, \ldots, \mathcal{T}_{\xi-1}$). Then, S_2^{ξ} will denote the collection of time intervals in the (unmarked) subsets appearing between the first and second marked subsets (i.e., $\mathcal{T}_{\xi+1}, \ldots, \mathcal{T}_{\xi+\frac{1}{\epsilon}-1}$), so on and so forth, up until the last such set, S_M^{ξ} . Finally, time intervals within marked subsets are collected into S_{∞}^{ξ} .

It is easy to verify that, for any choice of ξ , the ordered partition $S^{\xi} = (S_1^{\xi}, \ldots, S_M^{\xi}, S_{\infty}^{\xi})$ satisfies properties 1 and 2. In addition,

$$\mathbb{E}\left[H(S_{\infty}^{\xi})\right] = \mathbb{E}\left[\sum_{i\in[n]} H_i \cdot 1[i\in S_{\infty}^{\xi}]\right] = \epsilon H_{\Sigma},$$

since each item $i \in [n]$ ends up in S_{∞}^{ξ} if and only if the unique subset \mathcal{T}_q containing this item is marked, which happens with probability ϵ . As an immediate consequence, it follows that there exists some choice of $\xi \in \{0, \ldots, \frac{1}{\epsilon} - 1\}$ for which $H(S_{\infty}^{\xi}) \leq \epsilon H_{\Sigma}$, meaning that property 3 holds as well.

5.3 The Near-Additivity Theorem

Given a well-separated partition $S = (S_1, \ldots, S_M, S_\infty)$, our next step consists of arguing that, by temporarily leaving the subset S_∞ aside, we are left with a residual instance whose optimal peak inventory level is nearly additive with respect to S_1, \ldots, S_M . To formalize this claim, for every $m \in [M]$, we make use of \mathcal{I}_m to denote the instance obtained by restricting \mathcal{I} to the set of items S_m . Similarly, $\mathcal{I}_{[m]}$ will designate its restriction to $S_{[m]} = \bigcup_{\mu \in [m]} S_{\mu}$. With this notation, it is easy to verify that $OPT^{disc}(\mathcal{I}_{[M]}) \leq \sum_{m \in [M]} OPT^{disc}(\mathcal{I}_m)$. However, we will show that well-separated partitions, in conjunction with nested time intervals, lead to a nearly-tight inequality in the opposite direction.

THEOREM 5.2. OPT^{disc} $(\mathcal{I}_{[M]}) \ge (1 - 2\epsilon) \cdot \sum_{m \in [M]} \text{OPT}^{\text{disc}}(\mathcal{I}_m).$

The basic claim. In what follows, we argue that for every $m \in [M]$,

$$OPT^{disc}(\mathcal{I}_{[m]}) \geq OPT^{disc}(\mathcal{I}_{[m-1]}) + (1 - 2\epsilon) \cdot OPT^{disc}(\mathcal{I}_{m}) , \qquad (28)$$

with the convention that $OPT^{disc}(\mathcal{I}_{[0]}) = 0$. By expanding this inequality one term after the other, we indeed get $OPT^{disc}(\mathcal{I}_{[M]}) \ge (1 - 2\epsilon) \cdot \sum_{m \in [M]} OPT^{disc}(\mathcal{I}_m)$, as desired. Noting that inequality (28) is trivial for m = 1, we proceed by considering the general case of $m \ge 2$.

The ϵ -shortness relation. To this end, suppose that $\mathcal{I}_{\mathcal{A}}$ and $\mathcal{I}_{\mathcal{B}}$ are two inventory staggering instances, whose corresponding sets of items \mathcal{A} and \mathcal{B} are disjoint. For $\epsilon \in (0, 1)$, we say that

 $\mathcal{I}_{\mathcal{A}}$ is ϵ -shorter than $\mathcal{I}_{\mathcal{B}}$ when LCM($\{T_i\}_{i\in\mathcal{A}}$) $\leq \epsilon \cdot \min_{i\in\mathcal{B}} T_i$. Put differently, the cycle length of any policy in regard to $\mathcal{I}_{\mathcal{A}}$ is at most an ϵ -fraction of the minimal time interval in $\mathcal{I}_{\mathcal{B}}$. The next claim shows that, in this scenario, the best-possible peak inventory level for the merged instance $\mathcal{I}_{\mathcal{A}\cup\mathcal{B}}$ cannot be much smaller than the sum of optimal peaks for $\mathcal{I}_{\mathcal{A}}$ and $\mathcal{I}_{\mathcal{B}}$. It is important to point out that this result applies to both versions of the inventory staggering problem, which is why we make use of $\text{OPT}(\cdot)$, rather than $\text{OPT}^{\text{disc}}(\cdot)$ or $\text{OPT}^{\text{cont}}(\cdot)$.

LEMMA 5.3. When $\mathcal{I}_{\mathcal{A}}$ is ϵ -shorter than $\mathcal{I}_{\mathcal{B}}$,

$$OPT(\mathcal{I}_{\mathcal{A}\cup\mathcal{B}}) \geq OPT(\mathcal{I}_{\mathcal{A}}) + (1-2\epsilon) \cdot OPT(\mathcal{I}_{\mathcal{B}}).$$

PROOF. Let $\tau^{*\mathcal{A}\cup\mathcal{B}}$ be an optimal shift vector with respect to $\mathcal{I}_{\mathcal{A}\cup\mathcal{B}}$. In terms of this vector, the total inventory level of \mathcal{B} -items at any time point $t \in \mathbb{R}$ will be denoted by $I_{\Sigma}^{\mathcal{B}}(\tau^{*\mathcal{A}\cup\mathcal{B}},t) = \sum_{i\in\mathcal{B}} I_i(\tau_i^{*\mathcal{A}\cup\mathcal{B}},t)$, with its peak level being $I_{\max}^{\mathcal{B}}(\tau^{*\mathcal{A}\cup\mathcal{B}}) = \max_{t\in\mathbb{R}} I_{\Sigma}^{\mathcal{B}}(\tau^{*\mathcal{A}\cup\mathcal{B}},t)$. Now, suppose that the latter max is attained at \hat{t} , and let $\Delta = \epsilon T_{\min}^{\mathcal{B}}$, where $T_{\min}^{\mathcal{B}} = \min_{i\in\mathcal{B}} T_i$. Our first claim is that, for every time point $t \in [\hat{t}, \hat{t} + \Delta]$, we have

$$I_{\Sigma}^{\mathcal{B}}(\tau^{*\mathcal{A}\cup\mathcal{B}},t) \geq (1-2\epsilon) \cdot I_{\max}^{\mathcal{B}}(\tau^{*\mathcal{A}\cup\mathcal{B}}) .$$
⁽²⁹⁾

For this purpose, note that when there are no \mathcal{B} -orders in $[\hat{t}, t]$,

$$I_{\Sigma}^{\mathcal{B}}(\tau^{*\mathcal{A}\cup\mathcal{B}},t) = I_{\Sigma}^{\mathcal{B}}(\tau^{*\mathcal{A}\cup\mathcal{B}},\hat{t}) - (t-\hat{t}) \cdot \sum_{i\in\mathcal{B}} \frac{H_{i}}{T_{i}}$$

$$\geq I_{\max}^{\mathcal{B}}(\tau^{*\mathcal{A}\cup\mathcal{B}}) - \epsilon T_{\min}^{\mathcal{B}} \cdot \sum_{i\in\mathcal{B}} \frac{H_{i}}{T_{i}}$$
(30)

$$\geq I_{\max}^{\mathcal{B}}(\tau^{*\mathcal{A}\cup\mathcal{B}}) - \epsilon \cdot H(\mathcal{B})$$

$$\geq (1-2\epsilon) \cdot I_{\max}^{\mathcal{B}}(\tau^{*\mathcal{A}\cup\mathcal{B}}) .$$
(31)

Here, inequality (30) holds since $I_{\Sigma}^{\mathcal{B}}(\tau^{*\mathcal{A}\cup\mathcal{B}},\hat{t}) = I_{\max}^{\mathcal{B}}(\tau^{*\mathcal{A}\cup\mathcal{B}})$, by definition of \hat{t} , and since $t - \hat{t} \leq \Delta = \epsilon T_{\min}^{\mathcal{B}}$. Inequality (31) is obtained by specializing Lemma 2.1 to the set of items \mathcal{B} , implying that $H(\mathcal{B}) \leq 2 \cdot \operatorname{OPT}(\mathcal{I}_{\mathcal{B}}) \leq 2 \cdot I_{\max}^{\mathcal{B}}(\tau^{*\mathcal{A}\cup\mathcal{B}})$. In the complementary case, where there are one or more \mathcal{B} -orders in $[\hat{t}, t]$, the inventory level $I_{\Sigma}^{\mathcal{B}}(\tau^{*\mathcal{A}\cup\mathcal{B}}, t)$ can only be larger.

Now, let us recall that $\mathcal{I}_{\mathcal{A}}$ is ϵ -shorter than $\mathcal{I}_{\mathcal{B}}$, meaning that the cycle length of any policy with respect to $\mathcal{I}_{\mathcal{A}}$ is at most $\epsilon T^{\mathcal{B}}_{\min} = \Delta$. In particular, focusing on the restriction of $\tau^{*\mathcal{A}\cup\mathcal{B}}$ to these items, within the interval $[\hat{t}, \hat{t} + \Delta]$, there exists at least one point \bar{t} with

$$I_{\Sigma}^{\mathcal{A}}(\tau^{*\mathcal{A}\cup\mathcal{B}},\bar{t}) = I_{\max}^{\mathcal{A}}(\tau^{*\mathcal{A}\cup\mathcal{B}}) .$$
(32)

Consequently, to derive the desired lower bound on $OPT(\mathcal{I}_{\mathcal{A}\cup\mathcal{B}})$, we observe that

$$OPT(\mathcal{I}_{\mathcal{A}\cup\mathcal{B}}) = I_{\max}^{\mathcal{A}\cup\mathcal{B}}(\tau^{*\mathcal{A}\cup\mathcal{B}})$$

$$\geq I_{\Sigma}^{\mathcal{A}\cup\mathcal{B}}(\tau^{*\mathcal{A}\cup\mathcal{B}},\bar{t})$$

$$= I_{\Sigma}^{\mathcal{A}}(\tau^{*\mathcal{A}\cup\mathcal{B}},\bar{t}) + I_{\Sigma}^{\mathcal{B}}(\tau^{*\mathcal{A}\cup\mathcal{B}},\bar{t})$$

$$\geq I_{\max}^{\mathcal{A}}(\tau^{*\mathcal{A}\cup\mathcal{B}}) + (1-2\epsilon) \cdot I_{\max}^{\mathcal{B}}(\tau^{*\mathcal{A}\cup\mathcal{B}})$$
(33)

$$\geq \operatorname{OPT}(\mathcal{I}_{\mathcal{A}}) + (1 - 2\epsilon) \cdot \operatorname{OPT}(\mathcal{I}_{\mathcal{B}})$$

where inequality (33) is obtained by plugging in (29) and (32).

Concluding the basic claim (28). As our last step, we explain how to derive inequality (28) for the general case of $m \geq 2$. To this end, let us recall that S is a well-separated partition, meaning that for every pair of items $i_1 \in S_{[m-1]}$ and $i_2 \in S_m$, we necessarily have $T_{i_1} \leq \epsilon T_{i_2}$, by property 2. Consequently, since we are considering a nested instance, the cycle length of any policy with respect to $\mathcal{I}_{[m-1]}$ is given by $\max_{i_1 \in S_{[m-1]}} T_{i_1} \leq \epsilon \cdot \min_{i_2 \in S_m} T_{i_2}$, implying that $\mathcal{I}_{[m-1]}$ is ϵ -shorter than \mathcal{I}_m . As a result, by Lemma 5.3,

$$\operatorname{OPT}^{\operatorname{disc}}(\mathcal{I}_{[m]}) \geq \operatorname{OPT}^{\operatorname{disc}}(\mathcal{I}_{[m-1]}) + (1 - 2\epsilon) \cdot \operatorname{OPT}^{\operatorname{disc}}(\mathcal{I}_m) .$$

5.4 Algorithmic implications

From an algorithmic perspective, in light of the preceding discussion, we first compute a $(1 + \epsilon)$ approximate shift vector τ^m with respect to \mathcal{I}_m , for every $m \in [M]$. To this end, by property 1,
we know that $\frac{\max_{i \in S_m} T_i}{\min_{i \in S_m} T_i} \leq (\frac{1}{\epsilon})^{1/\epsilon}$. Additionally, since \mathcal{I}_m is a nested instance, the multiplicative
gap between any pair of distinct time intervals in $\{T_i\}_{i \in S_m}$ is at least 2, implying that there
are only $O(\log_2((\frac{1}{\epsilon})^{1/\epsilon})) = O(\frac{1}{\epsilon}\log\frac{1}{\epsilon})$ such values. Therefore, as stated in Theorem 1.4, we can
compute τ^m in $O(2^{\tilde{O}(1/\epsilon^3)} \cdot |\mathcal{I}|^{O(1)})$ time. For the singular instance \mathcal{I}_∞ , we pick an arbitrary
shift vector τ^∞ . Our final solution τ is formed by gluing $\tau^1, \ldots, \tau^M, \tau^\infty$ together. The next
claim proves that this vector is indeed near-optimal in terms of the original instance \mathcal{I} .

LEMMA 5.4. $I_{\max}(\tau) \leq (1+8\epsilon) \cdot \text{OPT}^{\text{disc}}(\mathcal{I}).$

1

PROOF. To upper-bound the peak inventory level of the policy \mathcal{P}_{τ} , we observe that

$$I_{\max}(\tau) \leq \sum_{m \in [M]} I_{\max}(\tau^m) + I_{\max}(\tau^\infty)
 \leq (1+\epsilon) \cdot \sum_{m \in [M]} OPT^{\operatorname{disc}}(\mathcal{I}_m) + H(S_\infty)$$
(34)

$$\leq \frac{1+\epsilon}{1-2\epsilon} \cdot \operatorname{OPT}^{\operatorname{disc}}(\mathcal{I}_{[M]}) + H(S_{\infty})$$
(35)

$$\leq \left(\frac{1+\epsilon}{1-2\epsilon}+2\epsilon\right) \cdot \operatorname{OPT}^{\operatorname{disc}}(\mathcal{I}) \tag{36}$$

 $\leq (1+8\epsilon) \cdot \operatorname{OPT}^{\operatorname{disc}}(\mathcal{I}) .$

Here, inequality (34) holds since τ^m is a $(1 + \epsilon)$ -approximate vector with respect to \mathcal{I}_m , for every $m \in [M]$. Inequality (35) follows from Theorem 5.2. Finally, inequality (36) is obtained by observing that $H(S_{\infty}) \leq \epsilon H_{\Sigma} \leq 2\epsilon \cdot \text{OPT}^{\text{disc}}(\mathcal{I})$, due to property 3 and Lemma 2.1.

6 Approximation Scheme for Continuous Pairwise Coprime Instances

Our final contribution comes in the form of a polynomial-time approximation scheme for pairwise coprime instances of the continuous inventory staggering problem. As stated in Theorem 1.6,

by repurposing some of the main ideas behind Section 5, we will approximate such instances within factor $1 + \epsilon$ of optimal in $O(\operatorname{tower}_4(O(\frac{1}{\epsilon}), O(\frac{1}{\epsilon})) \cdot |\mathcal{I}|^{O(1)})$ time.

6.1 Instilling ϵ -shortness and near-additivity

Decomposition. Given an instance \mathcal{I} of the continuous inventory staggering problem, comprised of pairwise coprime time intervals $T_1 < \cdots < T_n$, our first objective is to decompose \mathcal{I} into two essentially independent instances, with one being ϵ -shorter than the other. For this purpose, assuming without loss of generality that $\frac{1}{\epsilon}$ takes an integer value, we define the infinite sequence Ψ_1, Ψ_2, \ldots , such that

$$\Psi_0 = 1, \qquad \Psi_1 = \frac{1}{\epsilon}, \qquad \Psi_2 = 4^{1/\epsilon}, \qquad \Psi_3 = 4^{4^{1/\epsilon}}, \qquad \dots$$

where in general, $\Psi_m = \text{tower}_4(m, \frac{1}{\epsilon})$ for all $m \ge 1$. In correspondence to this sequence, we additionally define the items sets S_1, S_2, \ldots , where $S_m = \{i \in [n] : T_i \in [\Psi_{m-1}, \Psi_m)\}$. Clearly, there exists some $2 \le \hat{m} \le \frac{1}{\epsilon} + 1$ for which $H(S_{\hat{m}}) \le \epsilon H_{\Sigma}$, and we make use of this index to introduce three related instances:

- Small time intervals: \mathcal{I}_{-} is the restriction of \mathcal{I} to the set of items $S_{-} = \bigcup_{m < \hat{m}} S_m$.
- Large time intervals: \mathcal{I}_+ is its restriction to $S_+ = \bigcup_{m > \hat{m}} S_m$.
- Merged instance: \mathcal{I}_{\pm} is the instance obtained by unifying the item sets S_{-} and S_{+} .

The Near-Additivity Theorem. Similarly to Section 5.3, we proceed by arguing that the best-possible peak inventory level for \mathcal{I}_{\pm} cannot be much smaller than the sum of these measures for \mathcal{I}_{-} and \mathcal{I}_{+} .

THEOREM 6.1. $\operatorname{OPT}^{\operatorname{cont}}(\mathcal{I}_{\pm}) \ge \operatorname{OPT}^{\operatorname{cont}}(\mathcal{I}_{-}) + (1 - 2\epsilon) \cdot \operatorname{OPT}^{\operatorname{cont}}(\mathcal{I}_{+}).$

PROOF. The inequality in question is an immediate consequence of Lemma 5.3, as our construction guarantees that \mathcal{I}_{-} is ϵ -shorter than \mathcal{I}_{+} . To verify this claim, note that

$$LCM(\{T_i\}_{i\in S_-}) = \prod_{i\in S_-} T_i$$
(37)

$$\leq \Psi_{\hat{m}-1}^{\pi(\Psi_{\hat{m}-1})}$$
 (38)

$$\leq \Psi_{\hat{m}-1}^{1.2551 \cdot \frac{\Psi \hat{m}-1}{\ln \Psi_{\hat{m}-1}}} \tag{39}$$

$$= 4^{1.2551 \cdot \frac{\Psi \hat{m} - 1}{\ln 4}}$$

$$\leq \frac{4^{\Psi \hat{m} - 1}}{\Psi \hat{m} - 1}$$

$$\leq \epsilon \cdot \Psi \hat{m}$$

$$\leq \epsilon \cdot \min_{i \in S_{+}} T_{i} . \qquad (40)$$

Here, equality (37) holds since $\{T_i\}_{i \in S_-}$ are pairwise coprime. Inequality (38) is obtained by observing that all time intervals in S_- are upper-bounded by $\Psi_{\hat{m}-1}$. In addition, there could

be at most $\pi(\Psi_{\hat{m}-1})$ such values greater than 1, where $\pi(\cdot)$ is the prime-counting function. Inequality (39) follows by employing the upper bound $\pi(x) \leq 1.2551 \cdot \frac{x}{\ln x}$ for x > 1, due to Dusart (2018, Cor. 5.2). Finally, inequality (40) holds since $\{T_i\}_{i \in S_+}$ are lower-bounded by $\Psi_{\hat{m}}$.

6.2 Algorithmic implications

Given this result, we begin by computing a $(1 + \epsilon)$ -approximate shift vector τ^- with respect to \mathcal{I}_- . For this purpose, since $T_i < \Psi_{\hat{m}} \leq \Psi_{\frac{1}{\epsilon}+1}$ for every item $i \in S_-$, the number of distinct time intervals across this set is only $O(\Psi_{\frac{1}{\epsilon}+1})$. Therefore, as argued in Theorem 1.4, we can compute τ^- in time

$$O\left(2^{\tilde{O}(\Psi_{\frac{1}{\epsilon}+1}/\epsilon^2)} \cdot |\mathcal{I}|^{O(1)}\right) = O\left(\operatorname{tower}_4\left(O\left(\frac{1}{\epsilon}\right), O\left(\frac{1}{\epsilon}\right)\right) \cdot |\mathcal{I}|^{O(1)}\right) .$$

For the instance \mathcal{I}_+ , as well as for the set of items $S_{\hat{m}}$, we pick arbitrary shift vectors, τ^+ and $\tau^{\hat{m}}$. Our final solution τ is formed by gluing τ^- , τ^+ , and $\tau^{\hat{m}}$ together.

Large pairwise coprime time intervals. While τ^- is a $(1+\epsilon)$ -approximate shift vector with respect to \mathcal{I}_- , the performance guarantee of τ^+ in terms of \mathcal{I}_+ is still unclear. Recalling that all time intervals of the latter instance are lower-bounded by $\Psi_{\hat{m}} \geq \frac{1}{\epsilon}$, the next claim shows that any shift vector is $(1+\epsilon)$ -approximate in this context, since we obtain $\text{OPT}^{\text{cont}}(\mathcal{I}_+) \geq (1-\epsilon) \cdot H(S_+)$ as a direct implication.

LEMMA 6.2. Let \mathcal{I} be a pairwise coprime inventory staggering instance. Then,

$$OPT^{cont}(\mathcal{I}) \geq \left(1 - \frac{1}{T_{\min}}\right) \cdot H_{\Sigma}$$

PROOF. Our proof is based on arguing that $I_{\max}(\tau) \geq (1 - \frac{1}{T_{\min}}) \cdot H_{\Sigma}$, for every real-valued shift vector $\tau \in \mathbb{R}^n$. To this end, let $\hat{\tau} \in \mathbb{Z}^n$ be the vector specified by $\hat{\tau}_i = \lceil \tau_i \rceil$, for every item $i \in [n]$. Since the time intervals T_1, \ldots, T_n are pairwise coprime, by the Chinese Remainder Theorem, the system of congruences

$$t \equiv \hat{\tau}_i \pmod{T_i} \quad \forall i \in [n]$$

has a unique solution t modulo $\prod_{i \in [n]} T_i$. As such, for every $i \in [n]$, the policy \mathcal{P}_{τ} has an *i*-order at time $t - (\hat{\tau}_i - \tau_i)$, implying that the inventory level of this item at time t is

$$I_i(\tau_i, t) \geq H_i \cdot \left(1 - \frac{\hat{\tau}_i - \tau_i}{T_i}\right) = H_i \cdot \left(1 - \frac{\lceil \tau_i \rceil - \tau_i}{T_i}\right) \geq H_i \cdot \left(1 - \frac{1}{T_{\min}}\right).$$

It follows that the maximum inventory level of the policy \mathcal{P}_{τ} is $I_{\max}(\tau) \geq I_{\Sigma}(\tau, t) \geq (1 - \frac{1}{T_{\min}}) \cdot H_{\Sigma}$.

Final performance guarantee. We conclude the proof of Theorem 1.6 by observing that

$$I_{\max}(\tau) \leq I_{\max}^{S_{-}}(\tau^{-}) + I_{\max}^{S_{+}}(\tau^{+}) + I_{\max}^{S_{\hat{m}}}(\tau^{\hat{m}})$$

$$\leq (1+\epsilon) \cdot \operatorname{OPT}^{\operatorname{cont}}(\mathcal{I}_{-}) + H(S_{+}) + H(S_{\hat{m}})$$

$$\tag{41}$$

$$\leq (1+2\epsilon) \cdot \left(\operatorname{OPT}^{\operatorname{cont}}(\mathcal{I}_{-}) + \operatorname{OPT}^{\operatorname{cont}}(\mathcal{I}_{+}) \right) + \epsilon H_{\Sigma}$$

$$\tag{42}$$

$$\leq (1+2\epsilon)(1+3\epsilon) \cdot \operatorname{OPT}^{\operatorname{cont}}(\mathcal{I}_{\pm}) + 2\epsilon \cdot \operatorname{OPT}^{\operatorname{cont}}(\mathcal{I})$$
(43)

 $\leq (1+13\epsilon) \cdot \operatorname{OPT}^{\operatorname{cont}}(\mathcal{I})$.

Here, inequality (41) holds since τ^- is a $(1 + \epsilon)$ -approximate shift vector with respect to \mathcal{I}_- . Inequality (42) is obtained by recalling that $\text{OPT}^{\text{cont}}(\mathcal{I}_+) \ge (1 - \epsilon) \cdot H(S_+)$ and $H(S_{\hat{m}}) \le \epsilon H_{\Sigma}$. Finally, inequality (43) follows from Theorem 6.1 and Lemma 2.1.

7 Concluding Remarks

We conclude our work by briefly discussing two pivotal questions that lie at the heart of inventory staggering. The first question is concerned with the long-standing approximability gap of this problem in its full generality, whereas the second touches upon peak inventory evaluation, which has been circumvented in earlier papers by focusing on stylized formulations.

Sub-2-approximation/APX-hardness for general instances? As explained in Section 1.2, the classic average-space bound (Lemma 2.1) informs us that the peak inventory level of any shift vector is within factor 2 of optimal, regardless of how a given inventory staggering instance is structured. In the opposite direction, this problem has been shown to be strongly NP-hard; see Gallego et al. (1992) and Hochbaum and Rao (2019) for intractability proofs applying to the continuous and discrete settings, respectively. Without any simplifying assumptions, we believe that meaningfully narrowing the approximability region [PTAS, 2] would be a very challenging avenue for future research, probably requiring technical developments that are unavailable within this domain at present time.

Algorithmic/hardness results for peak evaluation? As mentioned in Section 1.2, efficiently evaluating the peak inventory level attained by a given shift vector remains a central question to be addressed. Toward this objective, Theorem 1.1 excludes standard sampling ideas, showing that $2^{\Omega(\epsilon^2 n)}$ random points should be drawn in order to estimate peak inventory levels within factor $\frac{1}{2} + \epsilon$ with constant probability. Along these lines, it would be interesting to examine whether exact/approximate peak evaluation is actually #P-hard. On a brighter note, Theorem 4.1 presents an exact evaluation oracle in $O(n^{O(n)} \cdot |\mathcal{I}|^{O(1)})$ time, by means of integer linear programming in fixed dimension (Lenstra, 1983; Kannan, 1987; Frank and Tardos, 1987). Substantial improvements on the latter running time, as well as faster approximate procedures in this context, would be challenging directions to pursue.

References

- Noga Alon and Joel H. Spencer. *The Probabilistic Method.* John Wiley & Sons, fourth edition, 2016. 23
- Noga Alon, Richard A. Duke, Hanno Lefmann, Vojtěch Rödl, and Raphael Yuster. The algorithmic aspects of the regularity lemma. *Journal of Algorithms*, 16(1):80–109, 1994. 6

- Noga Alon, Keith E. Mellinger, Dhruv Mubayi, and Jacques Verstraëte. The de Bruijn–Erdős theorem for hypergraphs. *Designs, Codes and Cryptography*, 65(3):233–245, 2012. 10
- Shoshana Anily. Multi-item replenishment and storage problem (MIRSP): Heuristics and bounds. *Operations Research*, 39(2):233–243, 1991. 9
- Fayez F. Boctor. Offsetting inventory replenishment cycles to minimize storage space. European Journal of Operational Research, 203(2):321–325, 2010. 3
- Ernest Croot and Kai Huang. A class of random algorithms for inventory cycle offsetting. International Journal of Operational Research, 18(2):201–217, 2013. 8
- Marek Cygan, Fedor V. Fomin, Łukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michał Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer, 2015. 22
- Harold Davenport and Hugh L. Montgomery. *Multiplicative Number Theory*. Graduate Texts in Mathematics. Springer New York, second edition, 2013. 10
- Devdatt P. Dubhashi and Alessandro Panconesi. Concentration of Measure for the Analysis of Randomized Algorithms. Cambridge University Press, 2009. 7
- Richard A. Duke, Hanno Lefmann, and Vojtěch Rödl. A fast approximation algorithm for computing the frequencies of subgraphs in a given graph. SIAM Journal on Computing, 24 (3):598–620, 1995. 6
- Pierre Dusart. Explicit estimates of some functions over primes. *The Ramanujan Journal*, 45: 227–251, 2018. 28
- András Frank and Éva Tardos. An application of simultaneous diophantine approximation in combinatorial optimization. *Combinatorica*, 7:49–65, 1987. 21, 29
- Guillermo Gallego, Dong Shaw, and David Simchi-Levi. The complexity of the staggering problem, and other classical inventory problems. Operations Research Letters, 12(1):47–52, 1992. 1, 6, 29
- Guillermo Gallego, Maurice Queyranne, and David Simchi-Levi. Single resource multi-item inventory systems. *Operations Research*, 44(4):580–595, 1996. 3, 9
- Nicholas G. Hall. A comparison of inventory replenishment heuristics for minimizing maximum storage. American Journal of Mathematical and Management Sciences, 18(3-4):245–258, 1998. 3, 6
- Roger Hartley and Lyn C. Thomas. The deterministic, two-product, inventory system with capacity constraint. Journal of the Operational Research Society, 33(11):1013–1020, 1982. 1, 9
- Allan Hartman, Ronald C. Mullin, and Douglas R. Stinson. Exact covering configurations and Steiner systems. Journal of the London Mathematical Society, 2(2):193–200, 1982. 10

- Dorit S. Hochbaum and Xu Rao. The replenishment schedule to minimize peak storage problem: The gap between the continuous and discrete versions of the problem. *Operations Research*, 67(5):1345–1361, 2019. I, 1, 3, 4, 6, 29
- Dorit S. Hochbaum and Xu Rao. A fully polynomial time approximation scheme for the replenishment storage problem. *Operations Research Letters*, 48(6):835–839, 2020. I, 4
- Eugene D. Homer. Space-limited aggregate inventories with phased deliveries. Journal of Industrial Engineering, 17(6):327, 1966. 1, 5
- Sin-Hoon Hum, Moosa Sharafali, and Chung-Piaw Teo. Staggering periodic replenishment in multivendor JIT environments. Operations Research, 53(4):698–710, 2005. I, 1, 5
- Peter Jackson, William Maxwell, and John Muckstadt. The joint replenishment problem with a powers-of-two restriction. *IIE Transactions*, 17(1):25–32, 1985. 5
- Klaus Jansen. An EPTAS for scheduling jobs on uniform processors: Using an MILP relaxation with a constant number of integral variables. SIAM Journal on Discrete Mathematics, 24(2): 457–485, 2010. 21
- Klaus Jansen, Kim-Manuel Klein, and José Verschae. Closing the gap for makespan scheduling via sparsification techniques. *Mathematics of Operations Research*, 45(4):1371–1392, 2020. 21
- Ravi Kannan. Minkowski's convex body theorem and integer programming. Mathematics of Operations Research, 12(3):415–440, 1987. 21, 29
- Hendrik W. Lenstra. Integer programming with a fixed number of variables. *Mathematics of Operations Research*, 8(4):538–548, 1983. 21, 29
- Daniel Lokshtanov. New Methods in Parameterized Algorithms and Complexity. PhD thesis, University of Bergen, Norway, 2009. 22
- John A. Muckstadt and Robin O. Roundy. Analysis of multistage production systems. In Stephen C. Graves, Alexander H. G. Rinnooy Kan, and Paul H. Zipkin, editors, *Handbooks* in Operations Research and Management Science, volume 4, chapter 2, pages 59–131. Elsevier, 1993. 5
- Nagesh N. Murthy, W. C. Benton, and Paul A. Rubin. Offsetting inventory cycles of items sharing storage. *European Journal of Operational Research*, 150(2):304–319, 2003. 3, 9
- Noam Nisan and Avi Wigderson. Hardness vs randomness. Journal of Computer and System Sciences, 49(2):149–167, 1994. 10
- Ivan Niven, Herbert S. Zuckerman, and Hugh L. Montgomery. An Introduction to the Theory of Numbers. John Wiley & Sons, fifth edition, 1991.
- E. Page and Ray J. Paul. Multi-product inventory situations with one restriction. Journal of the Operational Research Society, 27(4):815–834, 1976. 1

- Robin Roundy. 98%-effective integer-ratio lot-sizing for one-warehouse multi-retailer systems. Management Science, 31(11):1416–1430, 1985. 5
- Robin Roundy. A 98%-effective lot-sizing rule for a multi-product, multi-stage production/inventory system. *Mathematics of Operations Research*, 11(4):699–727, 1986. 5
- Danny Segev. New approximation guarantees for the economic warehouse lot scheduling problem, 2024. Working paper. Available online as arXiv preprint 2412.11184. 3, 9
- Chung-Piaw Teo, Jihong Ou, and Kok-Choon Tan. Multi-item inventory staggering problems: Heuristic and bounds. In Proceedings of the 9th Annual ACM-SIAM Symposium on Discrete Algorithms, pages 584–593, 1998. I, 1, 5
- Klaus Zoller. Deterministic multi-item inventory systems with limited capacity. *Management Science*, 24(4):451–455, 1977. 1