

Robustness of Online Inventory Balancing to Inventory Shocks

Yiding Feng

Industrial Engineering and Decision Analytics, Hong Kong University of Science & Technology (HKUST), ydfeng@ust.hk

Rad Niazadeh

The University of Chicago Booth School of Business, rad.niazadeh@chicagobooth.edu

Amin Saberi

Management Science and Engineering, Stanford University, saberi@stanford.edu

Abstract. In classic adversarial online resource allocation problems, such as matching, AdWords, or assortment planning, customers (demand) arrive online, while products (supply) are given offline with a fixed initial inventory. To ensure acceptable revenue guarantees given the uncertainty in future customer arrivals, the decision maker must balance consumption across different products. Motivated by this, the famous policy—*inventory balancing (IB)*—is introduced and studied in the literature and has proved to be optimal or near-optimal competitive in almost all classic settings. However, an important feature that these classic models do not capture are various forms of possible *inventory shocks* on the supply side, which plays an important role in several real-world applications of online assortment and could significantly impact the revenue performance of the IB algorithm.

Motivated by this paradigm, we introduce and study a new variant of online assortment planning with inventory shocks. Our model considers both adversarial exogenous shocks (where supply increases over time in unpredictable fashion) and allocation-coupled endogenous shocks (where an inventory reduction is triggered by the algorithms and is re-adjusted after a usage duration)—combination of which can cause non-monotonic inventory fluctuations. As our main result, we show the robustness of the inventory balancing strategy against inventory shocks and such fluctuations by designing a new family of optimal competitive IB-type algorithms, called *Batched Inventory Balancing (BIB)*. We develop a novel randomized primal-dual method to bound the competitive ratio of our BIB algorithm against any feasible policy. We show that with the proper choice of a certain parameter, this competitive ratio is asymptotically optimal and converges to $1 - 1/e$ as the initial inventories converge to infinity—in contrast to the original IB which no longer achieves the optimal competitive ratio in this new model. Moreover, we characterize BIB’s competitive ratio in its *general* form (parametric by the penalty function it uses) and show that it matches *exactly* the competitive ratio of IB in the special case without shocks. To this end, we use a refined analysis that reduces the dual construction to a combinatorial problem called the “interval assignment problem.” Our solution to this problem is algorithmic and might be of independent interest.

1. Introduction

Online assortment planning is a central problem in operations research, rooted in classical online resource allocation problems such as online matching (Karp et al., 1990) and AdWords (Mehta et al., 2007). Besides its theoretical appeal, this problem has also broad applications in the revenue management of retail platforms, as well as digital display and search advertising marketplaces. A common assumption in these applications is that each unit of a resource, once allocated and consumed, permanently leaves the inventory. Consequently, the planner begins with a fixed upfront inventory for different resources (e.g., advertisers with fixed daily budgets), which depletes over time due to allocations (e.g., once an impression is assigned to an advertiser, the budget decreases by the amount of the advertiser’s bid, and this money never returns to the advertiser). However, in several more modern online platforms—ranging from cloud computing and short- or long-term homestay marketplaces to online volunteer matching and freelance marketplaces—this assumption could be violated as inventory levels fluctuate over time due to various forms of *endogenous* and *exogenous* shocks.

An endogenous shock is typically triggered by an allocation decision, for example, when a user selects a reusable rental resource such as a cloud server on AWS or a listing on Airbnb (Gong et al., 2021; Goyal et al., 2025; Feng et al., 2024; Delong et al., 2024). Similarly, in volunteer matching platforms, when a volunteer is nudged to perform a newly arriving task, they become unavailable for a certain period of time—or the platform stops nudging them for some time, to avoid overwhelming the volunteers (Manshadi and Rodilitz, 2022). In all such cases, a unit of an offline resource temporarily leaves the inventory after an endogenous shock and returns after a period of usage. In contrast, an exogenous shock refers to an unpredictable and non-stationary replenishment of the inventory.¹ For instance, in omni-channel retail, unexpected product returns may occur when customers who purchased through other channels find the product does not meet their expectations, or when there is misalignment between inventory management and sale (de Brito and Dekker, 2003).² Another example arises in volunteer platforms such as Catchafire and Food Rescue U.S., where volunteers initially commit to a limited capacity but may later decide to spend additional working hours (thanks to engagement dynamics and incentive policies), allowing them to serve more than their initial commitment (Lo et al., 2024). In all these cases, one or several units of an offline resource are added to the inventory after an unpredictable exogenous shock and become available for future allocations.³

Motivated by the above paradigm, we introduce and study the problem of *online assortment with inventory shocks*. We consider a sequential allocation problem in which a sequence of consumers with arbitrary choice functions arrives over time. At each time step, the planner selects a subset of available products to display to the arriving consumer. The consumer then probabilistically chooses one of the displayed products

¹ One may also consider negative exogenous shocks. Although these are less common than positive shocks in motivating applications of our model, studying them remains an interesting direction; see Section 5 and Section EC.5 for further discussion.

² While returning a defective or undesired product occurs after the original sale, the presence of multiple return channels and the heterogeneity and non-stationarity of returns motivate an adversarial modeling; see Bruzzone et al. (2010) for discussion.

³ See Section EC.2 for a detailed discussion of various micro-foundations and applications of endogenous and exogenous shocks.

according to her choice model, consumes one unit of that product, and pays a product-dependent price. The planner faces both endogenous and exogenous inventory shocks. For endogenous shocks, once a unit of a product is purchased, it returns to the inventory after a fixed duration of usage, which may differ across products but remains constant over time. For exogenous shocks, we assume that at the beginning of each time step, a batch of products may be added to the inventory. Consumer choices and exogenous shocks are assumed to be arbitrary—chosen by an (adaptive) adversary—and are revealed to the platform only at the start of each time step, before the assortment is selected. In the absence of exogenous shocks, our model reduces to the online assortment with *reusable resources*, previously studied in the revenue management literature (Rusmevichientong et al., 2020; Gong et al., 2021; Goyal et al., 2025; Feng et al., 2024).

The goal of the assortment planner is to sequentially select inventory-feasible assortments while facing both endogenous and exogenous inventory shocks, in order to maximize the expected collected price (which we refer to as “revenue”). In principle, the one-shot assortment planning problem can be a computationally hard set-function optimization. To address this challenge, following the related literature (Golrezaei et al., 2014; Gong et al., 2021; Feng et al., 2024; Goyal et al., 2025), we assume an oracle access to an offline assortment solver given the consumer choice model. We also assume that consumers’ choice models satisfy the weak substitutability assumption—meaning that removing a product from the assortment can only weakly increase the probability of selecting another product—and the monotonicity assumption—meaning that adding a product to the assortment can only decrease the probability of selecting the outside option.

We aim to design online competitive algorithms for our problem and compare their performance to that of the optimal offline benchmark, which has complete knowledge of consumer choices and inventory shocks but not the exact choice realizations. This benchmark, commonly referred to as the *clairvoyant optimum* in the literature, provides an upper bound on the expected revenue achievable by any policy. We say that an algorithm is Γ -competitive for $\Gamma \in [0, 1]$ if, for every instance, the expected revenue of the algorithm is at least Γ times that of the clairvoyant optimum. We study the dependence of the competitive ratio on the minimum initial inventory across all products, denoted by c_0 , and focus primarily on the regime in which c_0 is large. Importantly, we make no assumptions on the exogenous shocks, other than that they are non-negative.

Summary of our contributions. Our work is inspired by the simple and elegant classic *BALANCE* algorithm—first introduced in Kalyanasundaram and Pruhs (2000) for the online bipartite b -matching problem, later generalized to the AdWords problem in Mehta et al. (2007), and subsequently to the non-reusable online assortment problem in Golrezaei et al. (2014). This algorithm, also referred to as *inventory balancing (IB)* in the assortment planning context, achieves the optimal $(1 - 1/e)$ competitive ratio in the fractional versions of these problems, or equivalently, in their integral versions under the “large inventory” regime (e.g., when c_0 is large in online assortment). Moreover, Golrezaei et al. (2014) consider a parametric version of IB (as a family of algorithms), where each algorithm is parameterized by a concave and increasing penalty function $\Psi : [0, 1] \rightarrow [0, 1]$ used to balance inventories. They provide an asymptotically tight characterization of the

competitive ratio Γ of an IB algorithm, given the penalty function Ψ and the minimum initial inventory c_0 , in the absence of endogenous or exogenous inventory shocks.

The above results are of paramount importance in the related application domains, as BALANCE/IB is simple, easy to implement, flexible to adapt to specific problem instances through the appropriate choice of penalty function, and transparent in its operations. Motivated by these properties, our goal is to address the following research question concerning the robustness of inventory balancing:

Is there a natural adaptation of the IB online algorithm that achieves the same asymptotic and non-asymptotic optimal competitive ratios established in [Golrezaei et al. \(2014\)](#), as if the online algorithm were not subject to endogenous and exogenous inventory shocks?

Our main result provides a compelling answer to the above question. We propose a new parametric family of online assortment algorithms, referred to as *batched inventory balancing (BIB)*, which extends the original “balance” idea to online assortment problems with both endogenous and exogenous inventory shocks. Using a novel randomized primal–dual analysis, we prove that this family achieves the asymptotically optimal competitive ratio of $1 - 1/e$ in the large-inventory regime. We further extend this result by characterizing an almost tight competitive ratio Γ for the general case, expressed as a function of the penalty function Ψ (a concave function analogous to that in the vanilla IB algorithm) and the minimum initial inventory c_0 in the non-asymptotic regime (see Theorem 1 for details). Notably, our bound exactly matches the one established in [Golrezaei et al. \(2014\)](#) for online assortment without inventory shocks.

The simple idea behind our algorithm is to treat different units of the same product separately. More precisely, for each product, we *assign* its units into what we call *batches* based on the time that a unit is added into the inventory due to an exogenous shock (or as part of the initial inventory). Each batch can be in one of two states: if it has accumulated a sufficiently large number of units, its available units can be used for the assortment—such a batch is called “*ready*”; otherwise, if it has not yet reached this size threshold, its available units cannot be used for the assortment until it grows larger—such a batch is called “*charging*”.

Our BIB algorithm operates by tracking (i) the assignment of units to “ready” batches and the potential “charging” batch for each product, and (ii) the inventory levels of the “ready” batches across all products. At each time step, when a new consumer arrives, the algorithm penalizes the price of each product based on the highest normalized inventory level among its “ready” batches. It then solves a one-shot assortment problem by making a customized query call to the oracle, where prices are replaced with penalized prices. When the consumer chooses a specific product in that assortment, the algorithm then allocates one unit of that product from the “ready” batch with the highest normalized inventory level.

New analysis framework: randomized primal–dual & interval assignment problem. We rely on a primal–dual analysis to establish the bounds on the competitive ratio of BIB. Rather than comparing directly to the clairvoyant benchmark, we follow prior work and consider an offline linear programming (LP) relaxation

of this policy. This LP benchmark has access to the same information as the clairvoyant policy about the input instance, but only requires that assortments respect inventory feasibility in expectation for each product unit. Unlike the classic LP benchmarks in the literature, it is important to note that our offline LP benchmark (named “*batch-specified LP*” and formulated in $\mathcal{P}_{\text{OPT}}(\gamma)$, where γ is the batch size threshold of the BIB algorithm) also depends on the execution of the BIB algorithm itself. In fact, similar to the BIB algorithm, this LP benchmark treats different units from different batches separately.

Our primal–dual method is randomized and, at least superficially, resembles the dual-fitting analysis of Devanur et al. (2013) for the RANKING algorithm, originally introduced in the seminal work of Karp et al. (1990) for the online bipartite matching problem. What makes the analysis of BIB particularly challenging—and distinguishes it from prior work—is twofold: (i) the times at which different units of the same product are added to the inventory due to exogenous shocks may vary, and since these units are treated differently by the primal BIB algorithm, they require distinct dual assignments; and (ii) the inventory level of an arbitrary batch is not necessarily *monotone decreasing* because of both endogenous and exogenous inventory shocks, in contrast to the classic online bipartite allocation problem without inventory shocks. This loss of monotonicity, a property central to previous analyses, introduces significant new challenges in our setting.

To address challenge (i), we develop the novel batch-specified LP benchmark $\mathcal{P}_{\text{OPT}}(\gamma)$, which, as alluded to earlier, is defined based on the execution path of the BIB algorithm. Furthermore, our primal–dual analysis constructs the duals at the batch level rather than at the product level (Section 4.2). To address challenge (ii), we employ a randomized dual construction that incorporates a *transformed version of the inventory levels* of “ready” batches through a particular combinatorial mapping. This mapping must satisfy certain structural properties to ensure that the resulting dual assignment is (approximately) feasible in expectation. We then restate our dual feasibility conditions in a purely combinatorial form and introduce a stand-alone combinatorial problem, which we call the *Interval Assignment Problem (IAP)*; see Lemma 2 and Section 4.3 for more details. Our final dual construction is obtained by developing an algorithmic solution to IAP (Algorithm 2).

The IAP and our analysis framework could be valuable tools for other related problems. In the literature on online resource allocation, various penalty functions beyond the exponential penalty function have been studied and shown to be optimal or maintain the state of the art in various models (e.g., Huang et al., 2019; Ma and Simchi-Levi, 2020; Jin and Williamson, 2021; Jin and Ma, 2022; Ekbatani et al., 2023). When incorporating these models with inventory shocks, our framework for general penalty functions, including the IAP, might be applicable. We leave any further developments to future research.

Finally, one important implication of our result is the *robustness of the performance guarantee* of the (batched) inventory balancing algorithm, even in complex scenarios where inventories change in non-monotone fashion due to both endogenous and exogenous inventory shocks. This is particularly relevant in applications, as simplicity and transparency of this algorithm, as well as the fact that it does not need to keep track of the full state of the restocking, makes it more practically appealing.

For a more in-depth discussion of related work, see Section EC.1 in the electronic companion.

2. Preliminaries

Problem formulation. We consider an online assortment planning problem subject to inventory constraints and inventory shocks. The platform manages inventories for n different products, indexed by $[n] = \{1, 2, \dots, n\}$. Each product i has an initial inventory of $c_i \in \mathbb{Z}_+$ units, a deterministic *usage duration* d_i (also referred to as the return duration), and a per-unit price r_i . We study a discrete, finite-horizon setting with T time periods, during which product inventories may experience both endogenous and exogenous shocks.⁴

For endogenous shocks, if a unit of product i is sold at time $t \in [T]$, it remains under usage for d_i periods and returns to the inventory at time $t + d_i$. For exogenous shocks, we assume that $\xi_{i,t} \in \mathbb{Z}_{\geq 0}$ units of each product i are added to the inventory at the beginning of period t . The sequence of exogenous shocks $\{\xi_{i,t}\}_{i \in [n], t \in [T]}$ is unknown to the platform in advance and is chosen by an adaptive adversary. We distinguish between the initial inventories $\{c_i\}_{i \in [n]}$ and the first-period shocks $\{\xi_{i,1}\}_{i \in [n]}$.⁵

Consumers who are interested in buying these products arrive sequentially in discrete times $t \in [T]$. Each consumer t has a choice model $\phi_t : \mathcal{S} \times [n] \rightarrow [0, 1]$, where $\mathcal{S} \subseteq 2^{[n]}$ is the collection of all feasible assortments that can be offered, and $\phi_t(S, i)$ is the probability that consumer t chooses product i when assortment set $S \in \mathcal{S}$ is offered. Upon the arrival of consumer t , her choice model ϕ_t and the adversarial exogenous shocks $\{\xi_{i,t}\}_{i \in [n]}$ are revealed to the platform. Given ϕ_t and the history up to time t , the platform offers an assortment $S_t \in \mathcal{S}$ of products from its available inventory. Based on the choice probabilities $\phi_t(S_t, i)$ for $i \in [n]$, consumer t randomly selects a product $i_t \in S_t$ or selects nothing. If i_t is selected, then she pays the price r_{i_t} to the platform. This unit sold returns to inventory at the beginning of time $t + d_i$.

We assume that consumers' choice models are chosen by an adversary (oblivious or adaptive). We further impose two assumptions on our choice models and feasible assortments, which have been commonly used in the past literature (Golrezaei et al., 2014; Rusmevichientong et al., 2020):

ASSUMPTION 1 (Weak substitutability & monotonicity). For all time periods $t \in [T]$, choice models ϕ_t and products $i \in [n]$, $\phi_t(\emptyset, i) = 0$. Moreover, for all $S \in \mathcal{S}$ and $j \in [n] \setminus \{i\}$, $\phi_t(S, i) \geq \phi_t(S \cup \{j\}, i)$.

ASSUMPTION 2 (Downward-closed feasibility). If $S \in \mathcal{S}$ and $S' \subseteq S$ then $S' \in \mathcal{S}$, i.e., a feasible assortment will remain feasible after removing any subset of its offered products.

Competitive ratio analysis. We evaluate the performance of the online algorithms compared to the *clairvoyant optimum offline benchmark*, which is the optimum offline solution maximizing expected revenue, knowing the sequence of consumer choices and exogenous shocks $\{\phi_t, \xi_{i,t}\}_{i \in [n], t \in [T]}$ but not knowing the exact realizations of the stochastic choices. More specifically, we study the worst case *competitive ratio* of our algorithm against this benchmark, defined formally as follows:

⁴ Our results also extend to *continuous* (i.e., real-valued) time. At each real-valued time $t \in [0, T]$, exogenous shocks $\{\xi_{i,t}\}_{i \in [n]}$ and a consumer may arrive. If the consumer selects product i , one unit is sold and becomes available again at real-valued time $t + d_i$.

⁵ We typically consider the regime in which the initial inventories $\{c_i\}$ are large, though we also analyze the small-inventory case. Nonetheless, we impose *no restrictions* on the exogenous shocks $\{\xi_{i,t}\}_{i \in [n], t \in [T]}$ other than non-negativity. The competitive ratio of our algorithm remains asymptotically optimal as $\min_i c_i$ tends to infinity.

DEFINITION 1 (COMPETITIVE RATIO). Fix n products with parameters $\{r_i, d_i, c_i\}_{i \in [n]}$. An online algorithm $\mathcal{A}$ is α -competitive against the clairvoyant optimum offline benchmark if

$$\inf_{T \geq 1} \inf_{\{\phi_t, \xi_{i,t}\}_{i \in [n], t \in [T]}} \frac{\mathbf{Rev}_{\mathcal{A}}[\{\phi_t, \xi_{i,t}\}_{i \in [n], t \in [T]}}]{\mathbf{OPT}[\{\phi_t, \xi_{i,t}\}_{i \in [n], t \in [T]}}} \geq \alpha,$$

where $\mathbf{Rev}_{\mathcal{A}}[\cdot]$ and $\mathbf{OPT}[\cdot]$ are the expected revenues of algorithm $\mathcal{A}$ and the clairvoyant optimum offline benchmark, respectively (with input being the sequence of consumer choices and exogenous shocks).

In simple words, the competitive ratio above is defined as the worst-case ratio between the expected total revenue of the online algorithm and the clairvoyant optimum offline benchmark, where the worst case is over all sequences of consumer choice models and exogenous inventory shocks over time.⁶

Oracle access to offline assortment solver. Similar to prior work, we assume having access to an oracle algorithm that solves the one-shot offline assortment problem. For simplicity, we assume the solver is exact throughout the paper, but all of our results still hold with a multiplicative loss of β in the competitive ratios if the solver is a β -approximation algorithm for some $0 < \beta < 1$.

ASSUMPTION 3 (Offline assortment oracle). For all $t \in [T]$ and non-negative reward $\hat{\mathbf{R}} \in \mathbb{R}_+^n$, we have oracle access to an algorithm that finds an assortment $\hat{S} \in \mathcal{S}$ such that: $\hat{S} \in \arg \max_{S \in \mathcal{S}} \sum_{i=1}^n \hat{R}_i \phi_t(S, i)$ and has the smallest cardinality among all such reward-maximizing assortments (for tie-breaking).

3. Asymptotically Optimal Algorithm: Batched Inventory Balancing

In this section, we propose our new online assortment algorithm, termed *batched-inventory-balancing* (BIB). We also summarize our main result, which shows that with appropriately chosen parameters, the BIB algorithm is not only asymptotically optimal in terms of its competitive ratio but also achieves a competitive ratio that *exactly* matches that of the vanilla Inventory Balancing (IB) algorithm in online assortment optimization, as established in the setting without endogenous or exogenous shocks (Golrezaei et al., 2014). We prove this result in the next section (Section 4) by developing a novel and refined primal–dual analysis for our BIB algorithm, which may be of independent interest.

3.1. An Overview of the BIB Algorithm

The BIB algorithm begins with a simple idea. Instead of monitoring the total inventory level—that is, the total number of available units for each product i —BIB treats the units of product i individually and dynamically assigns them to *batches* as new units are added to the inventory, either through exogenous shocks or as part of the initial inventory. Let $B_{i,j}$ denote the j^{th} batch for product i . The construction of batches $\{B_{i,j}\}_{j=1,2,\dots}$ for a specific product $i \in [n]$ proceeds as follows. At time period 0, all initial units of the product are grouped

⁶ Notably, due to exogenous inventory shocks, we introduce and analyze a batch-specified linear-programming relaxation of the clairvoyant optimum offline (formally defined in Section 4.1), as a proxy to obtain competitive ratios against this benchmark.

into the first batch, $B_{i,1}$, which is marked as “*ready*”. Simultaneously, an empty batch, $B_{i,2}$, is created and marked as “*charging*”. For each subsequent time period $t \in [T]$, any new unit of product i added through an exogenous inventory shock is permanently assigned to the current “charging” batch.

The BIB algorithm uses a predetermined *batch-size threshold* $\gamma \in \mathbb{N}$. When the number of units in the current “charging” batch, denoted by $|B_{i,j}|$, reaches this threshold, the batch is switched to “ready,” and a new empty batch marked as “charging” is created. Importantly, when a sold unit of a product returns to the inventory (an endogenous shock), it is added back to the same batch to which it was originally assigned—either when it first entered the inventory through an exogenous shock or as part of the initial inventory.

The BIB algorithm is also parameterized by a *penalty function* $\Psi : [0, 1] \rightarrow [0, 1]$ satisfying $\Psi(0) = 0$ and $\Psi(1) = 1$. During each time period t , while performing the batch construction process described above for each product upon the arrival of exogenous shocks, the BIB algorithm maintains a record of the *normalized inventory levels* $\{f_{i,j,t}\}$, where $f_{i,j,t}$ denotes the fraction of units of product i currently available at time t among all units assigned to the “ready” batch $B_{i,j}$ before and at time t .⁷ The algorithm then uses these normalized inventory levels to determine the final assortment. Specifically, it selects the assortment S_t by invoking the offline assortment oracle with a reduced price $\hat{R}_i \leftarrow r_i \Psi(\max_j f_{i,j,t})$ for each product i , where the maximization is taken over all “ready” batches of that product. If product i^* is chosen by consumer t , then an available unit from the “ready” batch B_{i^*,j^*} with $j^* = \arg \max_j f_{i^*,j,t}$ is allocated.⁸ See Algorithm 1 for a formal description of the BIB algorithm.

3.2. Main Result: Competitive Ratio of BIB

Our main theorem establishes competitive ratio bounds for the general BIB algorithm

THEOREM 1. *For any monotone non-decreasing, concave and differentiable penalty function Ψ , and any batch-size scalar $\gamma \in [1, c_0]$, the competitive ratio of the Batched Inventory Balancing algorithm is at least $\Gamma(\Psi, \gamma) = \min\{\Gamma_1(\Psi, \gamma), \Gamma_2(\Psi, \gamma)\}$, where*

$$\Gamma_1(\Psi, \gamma) = \min_{x \in [0, 1 - \frac{1}{c_0}]} \left\{ \frac{1 - x}{\frac{1}{c_0} + (1 + \frac{\gamma}{c_0})(1 - \Psi(x)) + \int_{x + \frac{1}{c_0}}^1 \Psi(y) dy} \right\},$$

$$\Gamma_2(\Psi, \gamma) = \min_{x \in [0, 1 - \frac{1}{\gamma}]} \left\{ \frac{1 - x}{\frac{1}{\gamma} + 1 - \Psi(x) + \int_{x + \frac{1}{\gamma}}^1 \Psi(y) dy} \right\},$$

and $c_0 = \min_{i \in [n]} c_i$ is the smallest initial inventory.

We provide a comprehensive analysis and prove Theorem 1 in Section 4. Before that, we next elaborate on some theoretical and practical corollaries and implications of our main result.

⁷ Let $\mathcal{E}(i, j, t')$ denote the event that an available unit from product i 's “ready” batch $B_{i,j}$ is allocated (or sold) at time t' . Then the normalized inventory level is defined as $f_{i,j,t} \triangleq 1 - \sum_{t' \in [t - d_i + 1 : t - 1]} \mathbb{1}\{\mathcal{E}(i, j, t')\} / |B_{i,j}|$.

⁸ There always exists an available unit in batch B_{i^*,j^*} . To see this, suppose otherwise. Then $f_{i^*,j^*,t} = 0$, and the reduced price $\hat{R}_{i^*} = 0$ since $\Psi(0) = 0$. Hence, the assortment $S_t \setminus \{i^*\}$, which remains feasible by Assumption 2, would also maximize the reduced total revenue, contradicting the tie-breaking rule in Assumption 3.

Algorithm 1: BATCHED INVENTORY BALANCING (BIB)**Input:** penalty function Ψ , batch-size threshold γ .

```

1 for each product  $i \in [n]$  do
2   initialize batch  $B_{i,1} \leftarrow [c_i]$  as the “ready” batch containing all  $c_i$  initial units.
3   initialize batch  $B_{i,2} \leftarrow \emptyset$  as the “charging” batch.
4   initialize “ready” batch counter  $h_i \leftarrow 1$ .
5 for each time period  $t \in [T]$  do
6   for each product  $i \in [n]$  do
7     for each new unit of product  $i$  added from exogenous shock do
8       assign this new unit into batch  $B_{i,h_i+1}$ 
9       if  $|B_{i,h_i+1}| = \gamma$  then
10        update  $h_i \leftarrow h_i + 1$ 
11        initialize  $B_{i,h_i+1} \leftarrow \emptyset$  as the new “charging” batch.
12   /* consumer  $t$  with choice model  $\phi_t$  arrives. */
13   /* some sold units of products are returned to inventory due to
14      endogenous shocks (i.e., restocking). */
15   update normalized inventory levels  $\{f_{i,j,t}\}$  of “ready” batches.
16   solve  $S_t \leftarrow \arg \max_{S \in \mathcal{S}} \sum_{i \in S} r_i \Psi(\max_{j \in [h_i]} f_{i,j,t}) \cdot \phi_t(S, i)$  with offline assortment oracle.
17   offer assortment  $S_t$  to consumer  $t$ .
18   /* suppose consumer  $t$  chooses product  $i^*$ . */
19   allocate an available unit of product  $i^*$  from batch  $B_{i^*,j^*}$  with  $j^* = \arg \max_{j \in [h_{i^*}]} f_{i^*,j,t}$ .
```

3.3. Corollaries and Implications of Theorem 1

Based on Theorem 1, it can be verified that the competitive ratio of the BIB algorithm is asymptotically optimized by selecting the exponential penalty function $\Psi(x) = \frac{e^{1-x}-e}{1-e}$ and $\gamma = \Theta(\sqrt{c_0})$, which approaches $(1 - 1/e) \approx 63\%$ as c_0 goes to infinity. This bound, which is now obtained in the presence of endogenous and exogenous shocks, is the asymptotic optimal bound even for the special cases of the online assortment (Golrezaei et al., 2014) and the online vertex-weighted b-matching (Karp et al., 1990; Birnbaum and Mathieu, 2008; Aggarwal et al., 2011) without inventory shocks.

In addition to establishing the asymptotic optimality of the BIB algorithm under the exponential penalty function, the competitive ratio bound for a general concave penalty function Ψ in Theorem 1 has three key implications. First, when compared with the competitive ratio bound of the vanilla IB algorithm with an arbitrary concave penalty function Ψ —as proved in Golrezaei et al. (2014) for the special case of online assortment optimization *without* inventory shocks—we show that the BIB algorithm achieves the *exact same* asymptotic competitive ratio as IB for *any* choice of concave penalty function Ψ , even in the presence of both

exogenous and endogenous shocks.⁹ This result highlights the robustness and power of inventory-balancing-style algorithms in online resource allocation: when properly adapted, they preserve their performance even under dynamic inventory environments with endogenous and exogenous shocks. Second, the precise competitive ratio bound in Theorem 1 also characterizes the dependence on the initial (finite) inventory levels, offering a sharper perspective beyond the asymptotic regime potentially useful for practitioners. Lastly, Theorem 1 provides a best-of-both-worlds guarantee: it ensures robust performance under arbitrary concave penalty functions in general settings, while allowing practitioners to leverage application-specific structures by tailoring the penalty function for potentially stronger guarantees (e.g., [Feng and Niazadeh, 2025](#)).

4. Competitive Ratio Analysis: Proof of Theorem 1

As mentioned earlier in the introduction, our goal is to use the primal–dual analysis to prove Theorem 1. In contrast to prior work that employs similar approaches, our analysis faces two main technical challenges: (i) different units of the same product require distinct dual assignments, since the times at which they are added to the inventory due to exogenous shocks may differ; and (ii) the inventory level of an arbitrary batch is no longer monotonically decreasing because of both endogenous and exogenous shocks—a property that was crucial for the success of prior primal–dual analyses in both the online matching problem and its generalizations to online assortment.

Overview of the analysis. To address the challenge (i), we come up with a batch-specified LP revenue benchmark that relies on the execution path of the BIB algorithm (Section 4.1), and our primal-dual analysis designs a dual construction at the batch level rather than at the product level (Section 4.2). To address the challenge (ii), roughly speaking, the dual construction uses a transformed version of inventory levels of “ready” batches using a particular mapping. The mapping is combinatorial and constructing it is rather involved; it uses a chain of techniques to reduce the problem of finding an appropriate inventory mapping for our purpose to solving a sub-problem called *Interval Assignment Problem (IAP)* (Lemma 2 and Section 4.3). Given the solution to IAP, our final dual construction uses transformed inventory levels and incorporates randomness in a specific fashion to provide a certificate for the approximate optimality of the primal algorithm.

4.1. Batch-Specified Assortment and LP Revenue Benchmark

We first define the revenue benchmark used in analyzing the competitive ratio of the BIB algorithm (Algorithm 1). Recall that the BIB algorithm views different units of the same product as separate entities rather than identical ones. This distinction is important, as different units may become available at different times due to replenishments. Therefore, to make the comparison meaningful, the benchmark should also distinguish between units of a fixed product. With this motivation, we extend the standard concept of an assortment

⁹ In Section EC.3, we examine several natural generalizations of the vanilla Inventory Balancing (IB) algorithm that are considerably simpler than BIB. We construct specific instances demonstrating that these generalizations fail to achieve the asymptotic optimal competitive ratio of $1 - 1/e$, even in the special case with only exogenous inventory shocks. In particular, we establish competitive ratio upper bounds of 0.552, 0.53, and 0.5 for these simple IB variants in Section EC.3.1.

(which is defined over products) to the notion of *batch-specified assortments*. Building on this, we then introduce the *batch-specified LP benchmark*, where both concepts are defined with respect to the product batches generated by the BIB algorithm.

Batch-specified assortments. Though all units of a fixed product i have the same price r_i and are identical from the consumers' perspective, it is useful to consider them separately.¹⁰ More precisely, our analysis partitions all units of each product i based on their assigned batches in the BIB algorithm, and units in each batch are considered separately.

We introduce an important auxiliary concept, which we term as *batch-specified assortment*, defined as follows. A batch-specified assortment $L = \{(i, j_i) : i \in S\}$ encodes a subset of products S (same as a typical assortment) together with a specific batch B_{i,j_i} for each product $i \in S$. With slight abuse of notation, we define $\phi_t(L, i) \triangleq \phi_t(S, i)$ for every batch-specified assortment $L = \{(i, j_i) : i \in S\}$. By offering batch-specified assortment L to consumer t , a unit from batch B_{i,j_i} of each product $i \in S$ is selected with probability $\phi_t(L, i)$.

Recall the (dynamic) construction of batches in the BIB algorithm, described formally in Algorithm 1. When the algorithm terminates at the end of the horizon T , there are h_i “ready” batches for each product $i \in [n]$. We denote by $\tau_{i,j}$ as the time period when batch $B_{i,j}$ is switched from “charging” to “ready”. As a sanity check, $\tau_{i,1} = 1$ and $\tau_{i,j}$ is weakly increasing in j . We say a batch-specified assortment $L = \{(i, j_i) : i \in S\}$ is *feasible* in time period t if its encoded assortment S is feasible (i.e., $S \in \mathcal{S}$) and $\tau_{i,j_i} \leq t$ for every product $i \in S$. We use $\mathcal{L}_t$ to denote the set of all feasible batch-specified assortments in time period t . Moreover, we use $\mathcal{L}_{i,j,t}$ to denote all feasible batch-specified assortments in time period t that contain batch $B_{i,j}$ of product i . As a sanity check, the BIB algorithm is essentially offering batch-specified assortment $L_t = \{(i, j_i) : i \in S_t, j_i = \arg \max_j f_{i,j,t}\}$ for each consumer $t \in [T]$ and each batch-specified assortment L_t is feasible in time period t , i.e., $L_t \in \mathcal{L}_t$.

Batch-specified LP benchmark. Similar to the batch-specified assortment introduced above, the revenue benchmark for our competitive ratio analysis also relies on the execution of the BIB algorithm, and in particular, the batch formulation. Consider the following linear program that knows the exact sequence of $\{\phi_t, \xi_{i,t}\}_{i \in [n], t \in [T]}$, but not the realization of consumers' choices:

$$\begin{aligned}
 \max_{\mathcal{X} \geq 0} \quad & \sum_{t=1}^T \sum_{L \in \mathcal{L}(t)} \sum_{i=1}^n r_i \phi_t(L, i) \mathcal{X}_{L,t} & \text{s.t.} \\
 & \sum_{t=t^\dagger}^{t^\dagger+d_i-1} \sum_{L \in \mathcal{L}_{i,1,t}} \phi_t(L, i) \mathcal{X}_{L,t} \leq c_i + \gamma & i \in [n], t^\dagger \in [\tau_{i,1} : T] \\
 & \sum_{t=t^\dagger}^{t^\dagger+d_i-1} \sum_{L \in \mathcal{L}_{i,j,t}} \phi_t(L, i) \mathcal{X}_{L,t} \leq \gamma & i \in [n], j \in [2 : h_i], t^\dagger \in [\tau_{i,j} : T] \\
 & \sum_{L \in \mathcal{L}_t} \mathcal{X}_{L,t} \leq 1 & t \in [T].
 \end{aligned} \tag{P_{\text{OPT}}(\gamma)}$$

¹⁰ In fact, ignoring endogenous effects of return, one can think of the inventory at each time being formed by the initial inventory and added inventory due to exogenous shocks at different times. With this view, considering units separately (as will be clear later) helps with keeping track of the time they are added for the first time to the inventory.

It can be observed that program $\mathcal{P}_{\text{OPT}}(\gamma)$ depends on the definition of BIB algorithm. Specifically, γ is the batch-size threshold of the BIB algorithm, $\{h_i\}_{i \in [n]}$ is the number of “ready” batches in the algorithm, and $\{\tau_{i,j}, \mathcal{L}_t, \mathcal{L}_{i,j,t}\}$ are all induced by the BIB algorithm. To highlight this dependence on the BIB algorithm, we parameterized program $\mathcal{P}_{\text{OPT}}(\gamma)$ by γ .

In program $\mathcal{P}_{\text{OPT}}(\gamma)$, each variable $\mathcal{X}_{L,t}$ can be interpreted as the probability of offering a batch-specified assortment L for consumer t . With this interpretation, the objective captures the expected revenue. Furthermore, the first (second) set of constraints captures a *relaxed* inventory feasibility constraint for all units in the first batch (second and later batches) of each product.¹¹ The third set of constraints ensures that the sum of allocation probabilities is at most one for each consumer.

As the main technical lemma in this subsection, we formally show that for any execution of the BIB algorithm, its induced program $\mathcal{P}_{\text{OPT}}(\gamma)$ is an upper bound on the revenue of the clairvoyant optimum policy that knows the adversarial sequence of consumer choice models and amounts of exogenous inventory shocks (and hence a relaxation to any feasible policy):

LEMMA 1. *For any sequence $\{\phi_t, \xi_{i,t}\}_{i \in [n], t \in [T]}$, the expected revenue of the clairvoyant optimum policy (and hence any online algorithm) is at most the optimal objective value in program $\mathcal{P}_{\text{OPT}}(\gamma)$ with any choice of batch-size threshold $\gamma \in \mathbb{N}$.*

Proof. *Proof.* Given any $\gamma \in \mathbb{N}$, we partition all units of each product into batches based on the batch construction/assignment of the BIB algorithm (with batch-size threshold γ) and equivalently consider the batch-specified assortment as we defined above. In the following, we construct a feasible solution for program $\mathcal{P}_{\text{OPT}}(\gamma)$ using the execution path of the clairvoyant optimum policy.

For each time period t , let $\tilde{L}_t$ denote the batch-specified assortment offered by the clairvoyant policy, where $\tilde{L}_t$ is a sample-path dependent random subset. For each unit allocation of product i in batch $B_{i,j}$ by the clairvoyant policy, we reassign it to the earlier batch $B_{i,j-1}$ (with $B_{i,0}$ interpreted as $B_{i,1}$). Let $\tilde{L}_t^\dagger$ denote the batch-specified assortment obtained by applying this batch-shifting operation to $\tilde{L}_t$, i.e., for each $(i,j) \in \tilde{L}_t$, include $(i, j-1 \vee 1)$ in $\tilde{L}_t^\dagger$. We then set $\mathcal{X}_{\tilde{L}_t^\dagger, t} = 1$ for all t , and $\mathcal{X}_{L,t} = 0$ for all other $L \in \mathcal{L}_t$. We next verify the constructed solution is feasible and finally analyze its objective value.

First set of constraints. Consider any product i and any interval of length d_i . The clairvoyant policy can allocate at most c_i units from the initial inventory plus those from the first exogenous batch. By construction, such allocations are shifted to $B_{i,1}$ in our constructed solution. Since $B_{i,1}$ has total size $c_i + \gamma$, the number of shifted allocations never exceeds this amount. Hence, the first family of constraints in $\mathcal{P}_{\text{OPT}}(\gamma)$ is satisfied.

¹¹ Fixing a product, the first constraint is an aggregated capacity constraint for all the units in the initial ready batch and the batch right after that (imagine those units are merged with the first batch). Therefore, the right-hand side is $c_i + \gamma$. Then the constraint corresponding to each index $j \in [2 : h_i]$ is a capacity constraint for units in batch $j+1$ (imagine that those units are shifted to batch j , as if they have arrived earlier). As shifting units to arrive earlier makes the problem more relaxed, this program is a relaxation. See the proof of Lemma 1 for a formal argument.

Second set of constraints. For each $j \geq 2$, each allocation of a unit from $B_{i,j+1}$ in the clairvoyant policy is shifted to $B_{i,j}$ in our constructed solution. As each batch has size at most γ , in any interval of length d_i there can be at most γ such effective allocations. Thus, the second family of constraints in $\mathcal{P}_{\text{OPT}}(\gamma)$ holds.

Third set of constraints. For each time t , we set $\mathcal{X}_{\tilde{L}_t^\dagger, t} = 1$ and all other allocations at time t to zero. Therefore, $\sum_{L \in \mathcal{L}_t} \mathcal{X}_{L, t} = 1$ for each t , which satisfies the third family of constraints.

Objective value. Finally, note that for every period t , the shifted assortment $\tilde{L}_t^\dagger$ yields the same choice probabilities as the clairvoyant assortment $\tilde{L}_t$. Therefore, the expected objective value of our constructed solution matches that of the clairvoyant policy. $\square$

4.2. Competitive Ratio Analysis of Batched Inventory Balancing

In this section, we present the competitive ratio analysis of the BIB algorithm using a primal-dual framework. Before delving into analysis, we introduce a key combinatorial construct underpinning our dual construction.

Interval assignment problem (IAP). Consider a sequence of s intervals indexed by $[s] \triangleq \{1, 2, \dots, s\}$, denoted $\{[a_i, b_i]\}_{i \in [s]}$, whose left endpoints satisfy $a_1 < a_2 < \dots < a_s$. For each interval index $i \in [s]$, define the *coverage set* $\mathcal{N}_i = \{j \in [s] : a_i \in [a_j, b_j]\}$, i.e., the set of interval indices whose intervals cover the left endpoint a_i , and let $|\mathcal{N}_i|$ denote its cardinality. The *Interval Assignment Problem (IAP)* asks to assign a positive integer $\hat{I}_i \in \mathbb{N}$ to each interval index i such that the resulting assignment $(\hat{I}_1, \dots, \hat{I}_s)$ satisfies the following three structural properties:

- (i) (LOCAL-DOMINANCE): For each interval index $i \in [s]$, there exists a bijection $\rho_i : \mathcal{N}_i \rightarrow \{1, 2, \dots, |\mathcal{N}_i|\}$ such that $\hat{I}_j \geq \rho_i(j)$ for all $j \in \mathcal{N}_i$.
- (ii) (GLOBAL-DOMINANCE): For every decreasing and concave function $\bar{\Psi}$, we have $\sum_{i \in [s]} \bar{\Psi}(\hat{I}_i) \geq \sum_{i \in [s]} \bar{\Psi}(|\mathcal{N}_i|)$.
- (iii) (PARTITION-MONOTONICITY): The interval index set can be partitioned into disjoint subgroups $A_1, \dots, A_\kappa$ such that for each $\ell \in [\kappa]$, $\{\hat{I}_i : i \in A_\ell\} = \{1, 2, \dots, |A_\ell|\}$.

We now state our main result about the IAP: a polynomial-time algorithm exists that always produces a valid assignment satisfying all three properties. Illustrative examples for the IAP can be found in Section EC.4.

LEMMA 2. *For any input instance of the IAP as defined above, there exists a polynomial-time algorithm (Algorithm 2; see Section 4.3 for details) that computes an assignment $(\hat{I}_1, \dots, \hat{I}_s) \in \mathbb{N}^s$ satisfying LOCAL-DOMINANCE, GLOBAL-DOMINANCE, and PARTITION-MONOTONICITY.*

Overview of the solution to IAP and its applications. In the primal-dual analysis of the BIB algorithm, a central difficulty arises because the actual number of units that have been allocated from a batch—but not yet returned—may increase and decrease over time due to overlapping rental periods. This non-monotonicity prevents a direct application of standard dual constructions, which rely on inventory trajectories that only

decrease.¹² To resolve this, for each “ready” batch, we construct a *transformed* inventory trajectory by solving an instance of IAP. In this instance, each allocation from the batch corresponds to an interval spanning from its allocation time to its return time. The IAP maps these intervals to integer labels that serve as surrogate inventory levels. These labels are carefully designed to satisfy three key properties, each of which enables a critical step in the primal-dual analysis: **LOCAL-DOMINANCE** ensures that, at the moment any unit is allocated, the transformed inventory levels of all units currently in use are sufficiently large to support the dual feasibility condition. Intuitively, it guarantees that the dual variables can be set high enough to “cover” the primal revenue of the allocation. **GLOBAL-DOMINANCE** guarantees that the total penalty (i.e., the sum of Ψ evaluated at the transformed inventory levels) is no worse than if we had used the real, non-monotone consumption levels. This allows us to relate the algorithm’s actual revenue to the dual objective without loss. **PARTITION-MONOTONICITY** reorganizes the allocations into groups where the transformed inventory levels behave like those in the classic setting—strictly decreasing by one after each allocation. This structure is essential for recovering the tight competitive ratio for general penalty function Ψ in Golrezaei et al. (2014). The details of how to solve IAP and the proof of Lemma 2 are deferred to Section 4.3.

We now present our primal-dual analysis for the BIB algorithm using the key Lemma 2.

Proof. Proof of Theorem 1. We start our analysis by writing down the dual program of our LP benchmark $\mathcal{P}_{\text{OPT}}(\gamma)$:

$$\begin{aligned} \min_{\theta, \lambda \geq 0} \quad & \sum_{i=1}^n \left(\sum_{t=\tau_{i,1}}^T (c_i + \gamma) \theta_{i,1,t} + \sum_{j=2}^{h_i} \sum_{t=\tau_{i,j}}^T \gamma \theta_{i,j,t} \right) + \sum_{t=1}^T \lambda_t \quad \text{s.t.} \\ & \lambda_t + \sum_{(i,j) \in L} \sum_{t'=t-d_i+1}^t \phi_t(L, i) \theta_{i,j,t'} \geq \sum_{(i,j) \in L} r_i \phi_t(L, i) \quad t \in [T], L \in \mathcal{L}_t. \end{aligned}$$

We construct a randomized dual assignment based on the execution path of the BIB algorithm together with the random realization of consumers’ choices as follows. First, set $\lambda_t \leftarrow 0$ and $\theta_{i,j,t} \leftarrow 0$ for all $t \in [T], i \in [n], j \in [h_i]$. For each time period $t \in [T]$, suppose consumer t chooses a unit of product i from batch $B_{i,j}$, update dual variables as follows:

$$\lambda_t \leftarrow r_i \Psi(f_{i,j,t}) \quad \text{and} \quad \theta_{i,j,t} \leftarrow r_i (\Psi(g_{i,j,t}) - \Psi(g_{i,j,t} - (1/|B_{i,j}|))) .$$

Here $\{g_{i,j,t}\}_{i \in [n], j \in [h_i], t \in [T]}$ is a sequence of values constructed based on the interval assignment problem (IAP) (Lemma 2). Specifically, for each product $i \in [n]$ and its “ready” batch $B_{i,j}$ with $j \in [h_i]$, let $\mathcal{T}_{i,j} = \{t_1, t_2, \dots, t_s\}$ be the time periods when a unit of product i from $B_{i,j}$ is chosen by consumers. Consider an instance of IAP with s intervals $[t_1, t_1 + d_i], \dots, [t_s, t_s + d_i]$. Let $\{\hat{I}_k\}_{k \in [s]}$ be the assignment solved in IAP for this batch. With slight abuse of notation, we also denote the assignment solved in this IAP instance as $\{\hat{I}_{i,j,t}\}_{t \in \mathcal{T}_{i,j}}$ to emphasize its connection to the BIB algorithm’s execution. We set $g_{i,j,t} \triangleq 1 - (\hat{I}_{i,j,t} - 1)/|B_{i,j}|$ for every time $t \in \mathcal{T}_{i,j}$ and $g_{i,j,t} \triangleq 0$ for all the remaining time $t \in [T] \setminus \mathcal{T}_{i,j}$.

¹² cf. Example EC.1, where the naive assignment based on real consumption fails.

The analysis proceeds in two steps. First, we show that the randomized dual assignments are feasible in expectation. Second, we show that the expected dual objective provides a $\Gamma(\Psi, \gamma)$ approximation to the expected revenue of the BIB algorithm. Combining these two results with the strong duality establishes the competitive ratio of the algorithm, as stated in Theorem 1.

[Step i] *Checking the feasibility of dual in expectation.* Since penalty function $\Psi(\cdot)$ of the BIB algorithm is non-negative and monotone non-decreasing, the constructed dual assignment is non-negative. We next show that the dual constraints corresponding to each $\mathcal{X}_{L,t}$ is satisfied in expectation. Note that this is enough, because we can later use the expectation of the dual variables as the actual dual certificate. First, we establish the following inequality for every product $i \in [n]$, batch $j \in [h_i]$, and time $t \in [\tau_{i,j} : T]$:

$$r_i \Psi(f_{i,j,t}) + \sum_{t'=t-d_i+1}^{t-1} \theta_{i,j,t'} \geq r_i. \quad (1)$$

To prove inequality (1), consider on-hand units together with units of product i that have not yet returned in batch $B_{i,j}$ at time t . Since the normalized inventory level of batch $B_{i,j}$ at time t is $f_{i,j,t}$, there are $m \triangleq (1 - f_{i,j,t})|B_{i,j}|$ units of product i in this batch that have not yet returned at time t . Suppose that these units were sold at times $t'_1 \leq t'_2 \leq \dots \leq t'_m$. By construction, all these time periods are between $t - d_i + 1$ and $t - 1$ and belong to $\mathcal{T}_{i,j}$. We now show inequality (1) as follows:

$$\begin{aligned} & r_i \Psi(f_{i,j,t}) + \sum_{t'=t-d_i+1}^{t-1} \theta_{i,j,t'} \stackrel{(a)}{=} r_i \Psi(f_{i,j,t}) + \sum_{k=1}^m \theta_{i,j,t'_k} \\ & \stackrel{(b)}{=} r_i \Psi(f_{i,j,t}) + \sum_{k=1}^m r_i \left(\Psi(g_{i,j,t'_k}) - \Psi(g_{i,j,t'_k} - (1/|B_{i,j}|)) \right) \\ & \stackrel{(c)}{=} r_i \Psi(f_{i,j,t}) + \sum_{k=1}^m r_i \left(\Psi\left(1 - ((\hat{I}_{i,j,t'_k} - 1)/|B_{i,j}|)\right) - \Psi\left(1 - (\hat{I}_{i,j,t'_k}/|B_{i,j}|)\right) \right) \\ & \stackrel{(d)}{\geq} r_i \Psi(1 - (m/|B_{i,j}|)) + \sum_{k=1}^m r_i \left(\Psi(1 - ((k-1)/|B_{i,j}|)) - \Psi(1 - (k/|B_{i,j}|)) \right) \stackrel{(e)}{=} r_i \Psi(1) \stackrel{(f)}{=} r_i. \end{aligned}$$

Here, equality (a) holds due to the definition of m , $\{t'_k\}_{k \in [m]}$ and the dual assignment construction; Equality (b) holds due to the dual construction; Equality (c) holds due to the construction of $\{g_{i,j,t'_k}\}$. For inequality (d), we use the definition of m , and the rest of inequality holds due to **LOCAL-DOMINANCE** of IAP (Lemma 2) at time t'_m , and the concavity of the penalty function Ψ .¹³ Finally, equality (e) holds due to the telescopic sum and equality (f) holds since $\Psi(1) = 1$.

Fixing an arbitrary time period t and a feasible batch-specified assortment $L \in \mathcal{L}_t$, we establish the following inequality for every time period t , which takes the expectation over the randomized choice of the

¹³ Let $m = (1 - f_{i,j,t})|B_{i,j}|$ be the number of units from batch $B_{i,j}$ in use at time t , allocated at times $t'_1 < \dots < t'_m$. Consider the last such time t'_m ; the coverage set $\mathcal{N}_{t'_m}$ includes all m intervals, so $|\mathcal{N}_{t'_m}| = m$. By **LOCAL-DOMINANCE**, the IAP assignment satisfy $\hat{I}_{(k)} \geq k$ for all $k \in [m]$, where $\hat{I}_{(k)}$ is the k -th smallest label. Since Ψ is concave and non-decreasing, its marginal difference is non-increasing. Thus, replacing the IAP assignment $\hat{I}_{i,j,t'_k}$ with the smaller values k can only increase each marginal term. Consequently, the sum on the LHS—using the IAP assignment—is at least as large as the RHS—using the idealized sequence $k = 1, \dots, m$. This justifies inequality (d).

consumer t' , while conditioning on the realized sample path $\mathcal{F}_{t-1}$ up to the end of time period t and the realization of the batch-specified assortment L_t displayed to consumer t :

$$\mathbb{E}[\lambda_t \mid \mathcal{F}_{t-1}, L_t] \stackrel{(a)}{=} \sum_{(i,j) \in L_t} r_i \phi_t(L_t, i) \Psi(f_{i,j,t}) \stackrel{(b)}{\geq} \sum_{(i,j) \in L} r_i \phi_t(L, i) \Psi(f_{i,j,t}). \quad (2)$$

Here, equality (a) holds due to the dual assignment construction and the expectation over consumer t 's choice; and inequality (b) holds due to the greedy decision for selecting L_t in BIB.

Combining inequalities (1) and (2), we are able to verify the feasibility of the dual constraint corresponding to primal variable $\mathcal{X}_{L,t}$ as follows:

$$\begin{aligned} & \mathbb{E} \left[\lambda_t + \sum_{(i,j) \in L} \sum_{t'=t-d_i+1}^t \phi_t(L, i) \theta_{i,j,t'} \right] \stackrel{(a)}{\geq} \mathbb{E} \left[\sum_{(i,j) \in L} r_i \phi_t(L, i) \Psi(f_{i,j,t}) + \sum_{(i,j) \in L} \sum_{t'=t-d_i+1}^t \phi_t(L, i) \theta_{i,j,t'} \right] \\ & \stackrel{(b)}{\geq} \sum_{(i,j) \in L} \phi_t(L, i) \mathbb{E} \left[r_i \Psi(f_{i,j,t}) + \sum_{t'=t-d_i+1}^t \theta_{i,j,t'} \right] \stackrel{(c)}{\geq} \sum_{(i,j) \in L} r_i \phi_t(L, i), \end{aligned}$$

where inequality (a) holds due to inequality (2); inequality (b) holds due to the linearity of expectation; and inequality (c) holds due to inequality (1).

[Step ii] *Comparing objective values of primal and dual.* We compare the contribution of batch $B_{i,j}$ in the BIB algorithm and the contribution in the dual objective for each “ready” batch $B_{i,j}$ separately. Recall $\mathcal{T}_{i,j}$ is the set of time periods when a unit of product i from batch $B_{i,j}$ has been selected. The revenue contribution in the BIB algorithm (aka., primal) corresponding to batch $B_{i,j}$ can be expressed as $\text{Primal}_{i,j} = r_i |\mathcal{T}_{i,j}|$. Meanwhile, the contribution of the dual objective based on the constructed dual assignment can be upper-bounded as

$$\begin{aligned} \text{Dual}_{i,j} & \stackrel{(a)}{=} r_i \sum_{t \in \mathcal{T}_{i,j}} (\Psi(f_{i,j,t}) + (c_i \cdot \mathbb{1}\{j=1\} + \gamma) (\Psi(g_{i,j,t}) - \Psi(g_{i,j,t} - (1/|B_{i,j}|)))) \\ & \stackrel{(b)}{\leq} r_i \sum_{t \in \mathcal{T}_{i,j}} (\Psi(g_{i,j,t}) + (c_i \cdot \mathbb{1}\{j=1\} + \gamma) (\Psi(g_{i,j,t}) - \Psi(g_{i,j,t} - (1/|B_{i,j}|)))) , \end{aligned} \quad (3)$$

where equality (a) holds due to the constructed dual assignment; and inequality (b) holds due to the construction of $\{g_{i,j,t}\}$ and **GLOBAL-DOMINANCE** of IAP (Lemma 2).¹⁴

Now, we partition time periods in $\mathcal{T}_{i,j}$ into subgroups $A_{i,j,1}, A_{i,j,2}, \dots, A_{i,j,\kappa}$ induced by **PARTITION-MONOTONICITY** of IAP (Lemma 2). Our remaining analysis compares the contribution of batch $B_{i,j}$ from time periods $A_{i,j,\ell}$ in the BIB algorithm (denoted by $\text{Primal}_{i,j,\ell}$) and the corresponding dual assignments' contribution in the dual objective (denoted by $\text{Dual}_{i,j,\ell}$) for each $\ell \in [\kappa]$ separately. Define auxiliary notation $\tilde{c}_{i,j} \triangleq c_i \cdot \mathbb{1}\{j=1\} + \gamma$. Note that $\text{Primal}_{i,j,\ell} = r_i |A_{i,j,\ell}|$ and

$$\text{Dual}_{i,j,\ell} \stackrel{(a)}{\leq} r_i \sum_{t \in A_{i,j,\ell}} (\Psi(g_{i,j,t}) + \tilde{c}_{i,j} (\Psi(g_{i,j,t}) - \Psi(g_{i,j,t} - (1/|B_{i,j}|))))$$

¹⁴ Recall that $f_{i,j,t} = 1 - m_t/|B_{i,j}|$, where m_t is the number of units from batch $B_{i,j}$ in use at time t , and $g_{i,j,t} = 1 - (\hat{I}_{i,j,t} - 1)/|B_{i,j}|$, where $\hat{I}_{i,j,t}$ is the IAP assignment. Define $\bar{\Psi}(x) \triangleq \Psi(1 - (x-1)/|B_{i,j}|)$, which is decreasing and concave because Ψ is concave and non-decreasing. Since $m_t + 1 = |\mathcal{N}_t|$ (the coverage size at the allocation time t), we have $\Psi(f_{i,j,t}) = \bar{\Psi}(|\mathcal{N}_t|)$ and $\Psi(g_{i,j,t}) = \bar{\Psi}(\hat{I}_{i,j,t})$. **GLOBAL-DOMINANCE** then implies $\sum_{t \in \mathcal{T}_{i,j}} \bar{\Psi}(\hat{I}_{i,j,t}) \geq \sum_{t \in \mathcal{T}_{i,j}} \bar{\Psi}(|\mathcal{N}_t|)$, i.e., $\sum_{t \in \mathcal{T}_{i,j}} \Psi(g_{i,j,t}) \geq \sum_{t \in \mathcal{T}_{i,j}} \Psi(f_{i,j,t})$, which justifies inequality (b).

$$\begin{aligned}
&\stackrel{(b)}{\leq} r_i \sum_{t \in A_{i,j,\ell}} \left(\Psi \left(1 - ((\hat{I}_{i,j,t} - 1)/|B_{i,j}|) \right) + \tilde{c}_{i,j} \left(\Psi \left(1 - ((\hat{I}_{i,j,t} - 1)/|B_{i,j}|) \right) - \Psi \left(1 - (\hat{I}_{i,j,t}/|B_{i,j}|) \right) \right) \right) \\
&\stackrel{(c)}{=} r_i \sum_{k=1}^{|A_{i,j,\ell}|} \left(\Psi \left(1 - ((k-1)/|B_{i,j}|) \right) + \tilde{c}_{i,j} \left(\Psi \left(1 - ((k-1)/|B_{i,j}|) \right) - \Psi \left(1 - (k/|B_{i,j}|) \right) \right) \right) \\
&= r_i \sum_{k=|B_{i,j}| - |A_{i,j,\ell}| + 1}^{|B_{i,j}|} \left(\Psi \left(k/|B_{i,j}| \right) + \tilde{c}_{i,j} \left(\Psi \left(k/|B_{i,j}| \right) - \Psi \left((k-1)/|B_{i,j}| \right) \right) \right) \\
&\stackrel{(d)}{=} r_i \sum_{k=|B_{i,j}| - |A_{i,j,\ell}| + 1}^{|B_{i,j}|} \Psi \left(k/|B_{i,j}| \right) + r_i \tilde{c}_{i,j} \left(\Psi \left(1 \right) - \Psi \left(1 - (|A_{i,j,\ell}|/|B_{i,j}|) \right) \right) \\
&\stackrel{(e)}{\leq} r_i + r_i \int_{1 - (|A_{i,j,\ell}| - 1)/|B_{i,j}|}^1 \Psi(y) dy + r_i \tilde{c}_{i,j} \left(1 - \Psi \left(1 - (|A_{i,j,\ell}|/|B_{i,j}|) \right) \right)
\end{aligned}$$

where inequality (a) holds due to inequality (3); inequality (b) holds due to the construction of $g_{i,j,t}$; equality (c) holds due to **PARTITION-MONOTONICITY** of IAP (Lemma 2); equality (d) holds due to the telescopic sum; and inequality (e) holds since Ψ is weakly increasing and $\Psi(1) = 1$.

Note that by setting $x \triangleq 1 - (|A_{i,j,\ell}|/|B_{i,j}|)$, the ratio between $\text{Primal}_{i,j,\ell}$ and $\text{Dual}_{i,j,\ell}$ has the same closed form as $\Gamma_1(\Psi, \gamma)$ for “ready” batch $B_{i,1}$, and $\Gamma_2(\Psi, \gamma)$ for “ready” batch $B_{i,j}$ with $j \in [2 : h_i]$. Thus, we conclude that for every sample path, the revenue in the BIB algorithm is a $\Gamma(\Psi, \gamma)$ approximation of the dual objective value of the constructed dual assignment. $\square$

4.3. Interval Assignment Problem

We now prove Lemma 2 by explicitly constructing an assignment that satisfies all three required properties. We first describe its core idea and then verify the properties using its structural invariants. A reference implementation is provided in Algorithm 2, and illustrative executions (Figures EC.2 and EC.3) of the algorithm on representative IAP instances (Examples EC.1 and EC.2) can be found in Section EC.4.

Algorithmic idea. The assignment is constructed by processing intervals in a specific order dictated by their *current coverage size*: at each step, among all unassigned intervals (denoted by $\mathcal{W} \subseteq [s]$), we identify the one (denoted by i) that is covered by the *largest* number of still-unassigned intervals (i.e., $\mathcal{N}_i \cap \mathcal{W}$), and assign the *leftmost* interval j in that coverage set a label equal to the current coverage size $|\mathcal{N}_i \cap \mathcal{W}|$. This greedy rule ensures that intervals appearing earlier (with smaller indices) are never “starved” of large labels, which is essential for **LOCAL-DOMINANCE**. Moreover, because the algorithm always assigns labels based on maximal coverage sets, the resulting assignment incurs no more penalty (under any decreasing concave function $\bar{\Psi}(\cdot)$) than the real consumption levels, i.e., satisfies **GLOBAL-DOMINANCE**.

To enable **PARTITION-MONOTONICITY**, the algorithm also builds a predecessor mapping during execution. Specifically: (i) mapping q_j records which interval’s coverage set triggered the assignment of $\hat{I}_j$. This is used to determine which intervals overlap in a way that justifies linking them into a partition group; (ii) mapping p_k records a predecessor link: when interval j is assigned label m , it becomes the predecessor of any previously assigned interval k with $\hat{I}_k = m + 1$ that overlaps appropriately with j ’s coverage context (i.e., $j \in \mathcal{N}_{q_k}$). These links induce a forest of chains, each forming a valid partition group. These mappings are not used during assignment decisions—they exist solely to support the analysis.

Algorithm 2: INTERVAL ASSIGNMENT SOLVER**Input:** sequence of s intervals $\{[a_1, b_1], \dots, [a_s, b_s]\}$ **Output:** assignment $(\hat{I}_1, \dots, \hat{I}_s)$

```

1 initialize assignment  $\hat{I}_i \leftarrow 0$ , mapping  $p_i \leftarrow -1$ , and mapping  $q_i \leftarrow -1$  for all  $i \in [s]$ .
2 initialize remaining interval set  $\mathcal{W} \leftarrow \{1, 2, \dots, s\}$ . /* define  $\mathcal{N}_i^{\mathcal{W}} = \mathcal{N}_i \cap \mathcal{W}$  */
3 while  $\mathcal{W}$  is not empty do
4   find interval  $i$  with the largest index in remaining interval set  $\mathcal{W}$  where  $|\mathcal{N}_i^{\mathcal{W}}|$  is maximized, i.e.,
      $i \leftarrow \arg \max_{i \in \mathcal{W}} |\mathcal{N}_i^{\mathcal{W}}|$ .
5   find interval  $j$  with the smallest index in remaining interval set  $\mathcal{W}$  which covers position  $a_i$ , i.e.,
      $j \leftarrow \min \mathcal{N}_i^{\mathcal{W}}$ .
6   set assignment  $\hat{I}_j \leftarrow |\mathcal{N}_i^{\mathcal{W}}|$  and mapping  $q_j \leftarrow i$ .
7   for each interval  $k \in [s] \setminus \mathcal{W}$  such that  $j \in \mathcal{N}_{q_k}$  and  $\hat{I}_k = \hat{I}_j + 1$  do
8     set mapping  $p_k \leftarrow j$ .
9   remove interval  $j$  from remaining interval set  $\mathcal{W}$  i.e., update  $\mathcal{W} \leftarrow \mathcal{W} \setminus \{j\}$ .
10 return  $\{\hat{I}_1, \dots, \hat{I}_s\}$ 

```

We now prove Lemma 2 by formally verifying the three properties using our constructed procedure.

Proof. *Proof of Lemma 2.* The constructed procedure (Algorithm 2) terminates after s steps, producing a complete assignment $(\hat{I}_1, \dots, \hat{I}_s)$. We next verify that this assignment satisfies all three desired properties in the statement of the lemma.

(i) **LOCAL-DOMINANCE:** For each interval $i \in [s]$, consider $j_1, j_2, \dots \in \mathcal{N}_i$ in the order in which they are removed from the algorithm: for the k -th removed interval $j_k \in \mathcal{N}_i$, at that time the cardinality $|\mathcal{N}_i^{\mathcal{W}}|$ is $|\mathcal{N}_i| - k + 1$, which guarantees that $\hat{I}_{j_k}$ is at least $|\mathcal{N}_i| - k + 1$ by construction. Thus, **LOCAL-DOMINANCE** is satisfied.

(ii) **GLOBAL-DOMINANCE:** We verify this property by the following induction argument. Consider the reverse order of intervals removed from the remaining interval set $\mathcal{W}$ in the algorithm, i.e., we set $\mathcal{W} = \emptyset$ as the base case and then add intervals back to $\mathcal{W}$ in the reverse order in which they have been removed in the algorithm. Our induction hypothesis is that restricting to intervals in $\mathcal{W}$, **GLOBAL-DOMINANCE** holds, i.e., for any decreasing and concave function $\bar{\Psi}$,

$$\sum_{i \in \mathcal{W}} \bar{\Psi}(\hat{I}_i) \geq \sum_{i \in \mathcal{W}} \bar{\Psi}(|\mathcal{N}_i^{\mathcal{W}}|). \quad (4)$$

• *Base Case* ($\mathcal{W} = \emptyset$): In this case, the induction hypothesis holds trivially.

• *Inductive Step:* Suppose the induction hypothesis holds before we add interval i with assignment $\hat{I}_i$ back to $\mathcal{W}$. Let $m \triangleq \hat{I}_i - 1$. To make it clear, let $\mathcal{W}$ to be the remaining interval set before we add interval i , and $\mathcal{W}'$ to be the remaining interval set after we add interval i , i.e., $\mathcal{W}' = \mathcal{W} \cup \{i\}$. Next we analyze the increment of both sides in inequality (4) after we add interval i into $\mathcal{W}$. The left-hand side of inequality (4)

increases by $\bar{\Psi}(m+1)$. For the right-hand side, since $\hat{I}_i = m+1$, there exist m intervals $j_1, j_2, \dots, j_m \in \mathcal{W}$ which cover position a_{q_i} , and $i < j_k$ (otherwise, j_k should be removed earlier than i in the algorithm) for all $k \in [m]$. This implies that interval i covers position a_{j_k} for all $k \in [m]$. Without loss of generality, we assume $j_1 < j_2 < \dots < j_m$. We have $|\mathcal{N}_{j_k}^{\mathcal{W}}| \geq k$ for all $k \in [m]$. Hence, after adding interval i into $\mathcal{W}$, the right-hand side of inequality (4) increases by

$$\begin{aligned} & \bar{\Psi}(|\mathcal{N}_i^{\mathcal{W}'}|) + \sum_{j \in \mathcal{W}} \left(\bar{\Psi}(|\mathcal{N}_j^{\mathcal{W}'}|) - \bar{\Psi}(|\mathcal{N}_j^{\mathcal{W}}|) \right) \stackrel{(a)}{\leq} \bar{\Psi}(|\mathcal{N}_i^{\mathcal{W}'}|) + \sum_{k \in [m]} \left(\bar{\Psi}(|\mathcal{N}_{j_k}^{\mathcal{W}'}|) - \bar{\Psi}(|\mathcal{N}_{j_k}^{\mathcal{W}}|) \right) \\ & \stackrel{(b)}{=} \bar{\Psi}(|\mathcal{N}_i^{\mathcal{W}'}|) + \sum_{k \in [m]} \left(\bar{\Psi}(|\mathcal{N}_{j_k}^{\mathcal{W}}| + 1) - \bar{\Psi}(|\mathcal{N}_{j_k}^{\mathcal{W}}|) \right) \stackrel{(c)}{\leq} \bar{\Psi}(1) + \sum_{k \in [m]} (\bar{\Psi}(k+1) - \bar{\Psi}(k)) = \bar{\Psi}(m+1) \end{aligned}$$

where inequality (a) holds since $|\mathcal{N}_{j_k}^{\mathcal{W}'}| \geq |\mathcal{N}_{j_k}^{\mathcal{W}}|$ and $\bar{\Psi}$ is decreasing; equality (b) holds since $|\mathcal{N}_{j_k}^{\mathcal{W}'}| = |\mathcal{N}_{j_k}^{\mathcal{W}}| + 1$ by construction; and inequality (c) holds since $\bar{\Psi}$ is decreasing, concave and $|\mathcal{N}_{j_k}^{\mathcal{W}}| \geq k$ for all $k \in [m]$. Applying the induction hypothesis (before we add interval i), the inductive statement (after we add interval i) holds, which finishes the induction as desired.

(iii) **PARTITION-MONOTONICITY**: We first show that the predecessor mapping $\{p_i\}_{i \in [s]}$ satisfies two key properties: (a) each interval has at most one successor j (i.e., $p_j = i$), and (b) every interval i with $\hat{I}_i > 1$ has a predecessor (i.e., $p_i \in [s]$). Consequently, the p -links decompose the set of intervals into disjoint chains along which the assigned labels decrease by exactly one. This yields a valid partition $\{A_1, \dots, A_\kappa\}$ such that $\{\hat{I}_i : i \in A_\ell\} = \{1, 2, \dots, |A_\ell|\}$ for each ℓ , thereby guaranteeing **PARTITION-MONOTONICITY**. Next, we establish claims (a) and (b) separately.

Proof of claim (a). We show claim (a) by contradiction. Suppose there exist intervals i, j_1 and j_2 such that $p_{j_1} = p_{j_2} = i$. By the construction of Algorithm 2, $\hat{I}_{j_1} = \hat{I}_{j_2} = \hat{I}_i + 1$. Now, consider case by case.

- *Case 1 ($i < j_1$ or $i < j_2$):* Without loss of generality, we assume $i < j_1$, and the case $i < j_2$ is symmetric. By the construction of Algorithm 2, $p_{j_1} = i$ implies $i \in \mathcal{N}_{q_{j_1}}$. This leads to a contradiction because if $i < j_1$, the interval i should be removed earlier than the interval j_1 in Algorithm 2 which implies $\hat{I}_i \geq \hat{I}_{j_1}$.

- *Case 2 ($j_1, j_2 < i$):*

- *Case 2a ($q_{j_1} < q_{j_2}$ or $q_{j_2} < q_{j_1}$):* Without loss of generality, we assume $q_{j_1} < q_{j_2}$ ($q_{j_2} < q_{j_1}$ is similar):

- *Case 2a1 (interval j_1 is removed before interval j_2):* In this case, consider the moment when Algorithm 2 removes j_1 due to interval q_{j_1} . By the construction of Algorithm 2, since j_1 is removed before j_2 , it implies $|\mathcal{N}_{q_{j_1}}^{\mathcal{W}}| > |\mathcal{N}_{q_{j_2}}^{\mathcal{W}}|$. Thus, $\hat{I}_{j_1}$ is strictly larger than $\hat{I}_{j_2}$, a contradiction.

- *Case 2a2 (interval j_2 is removed before interval j_1):* In this case, consider the moment when Algorithm 2 removes j_2 due to interval q_{j_2} . Notice that interval j_2 also covers position $a_{q_{j_1}}$, i.e., $j_2 \in \mathcal{N}_{q_{j_1}}$. By the construction of Algorithm 2, since j_2 is removed before j_1 , it implies $|\mathcal{N}_{q_{j_1}}^{\mathcal{W}}| \leq |\mathcal{N}_{q_{j_2}}^{\mathcal{W}}|$. Thus, once Algorithm 2 removes j_2 , $|\mathcal{N}_{q_{j_1}}^{\mathcal{W}/\{j_2\}}|$ (which is an upper bound of $\hat{I}_{j_1}$) becomes strictly smaller than $|\mathcal{N}_{q_{j_2}}^{\mathcal{W}}|$ (which is equal to $\hat{I}_{j_2}$), a contradiction.

- *Case 2b ($q_{j_1} = q_{j_2}$):* Without loss of generality, we assume interval j_1 is removed before interval j_2 , and the case where interval j_2 is removed before interval j_1 is symmetric. Let $q^* = q_{j_1} = q_{j_2}$. Consider

the moment when Algorithm 2 removes j_1 due to q^* . By the construction of Algorithm 2, $\hat{I}_{j_1} = |\mathcal{N}_{q^*}^{\mathcal{W}}|$ and $\hat{I}_{j_2} \leq |\mathcal{N}_{q^*}^{\mathcal{W}/\{j_1\}}| = |\mathcal{N}_{q^*}^{\mathcal{W}}| - 1$. Thus, $\hat{I}_{j_1} > \hat{I}_{j_2}$, which is a contradiction.

Proof of claim (b). For any interval i with $\hat{I}_i > 1$, consider the next removed interval j among $\mathcal{N}_{q_i}$ after Algorithm 2 removes interval i . It is sufficient to show $\hat{I}_j = \hat{I}_i - 1$, which implies that Algorithm 2 sets $p_i = j$ when it removes interval j . Since j is the next interval removed among $\mathcal{N}_{q_i}$ after interval i , we have $\hat{I}_i - 1 \leq \hat{I}_j \leq \hat{I}_i$. To show $\hat{I}_j < \hat{I}_i$, we use the proof by contradiction. Suppose $\hat{I}_j = \hat{I}_i$. If $i \in \mathcal{N}_{q_j}$, then at the moment when Algorithm 2 removes i , we have $\hat{I}_i = |\mathcal{N}_{q_i}^{\mathcal{W}}| \geq |\mathcal{N}_{q_j}^{\mathcal{W}}| \geq \hat{I}_j + 1 = \hat{I}_i + 1$, which is a contradiction. If $i \notin \mathcal{N}_{q_j}$, since $i < j$, we have mapping $q_i < q_j$ and at the moment when Algorithm 2 removes i , $|\mathcal{N}_{q_i}^{\mathcal{W}}| = \hat{I}_i = \hat{I}_j \leq |\mathcal{N}_{q_j}^{\mathcal{W}}|$, a contradiction. $\square$

5. Conclusion

Summary & takeaways. We studied real-time assortment planning under both endogenous and exogenous inventory shocks. We introduced an extension of the inventory balancing algorithm for this problem, called Batched Inventory Balancing (BIB). This new algorithm keeps track of the exogenous inventory shock of each product separately by smartly batching units of this product into buckets in an adaptive fashion. It then penalizes the price of different units of the product using a concave penalty function with the (normalized) inventory of the bucket of the unit as input. We analyzed the BIB algorithm with general concave penalty functions. The analysis in the prior work for online assortment (without inventory shocks) breaks as inventory levels are not monotone and the function does not satisfy a certain differential equation. To handle this challenge, we considered an LP benchmark that depends on the execution of the BIB algorithm, designed a batch-based randomized dual-fitting primal-dual analysis, and most interestingly proposed a reduction from the correct randomized dual construction in our problem to another algorithmic yet combinatorial problem. This combinatorial problem, called Interval Assignment Problem (IAP), searches for a certain re-assignment of the bucket inventory levels, so that the new inventory levels are always monotone and can be used in a primal-dual analysis as before. We concluded by showing how to solve the IAP in polynomial time. We then show how to put all the pieces together and obtain a universal bound comparable to [Golrezaei et al. \(2014\)](#), but with both endogenous and exogenous inventory shock.

Open problems & directions. There are several promising technical directions for future research stemming from our work, including: extending the model to stochastic or heterogeneous return durations; incorporating inventory holding or restocking costs; handling negative inventory shocks; studying joint inventory–assortment control; and exploring application-specific operational replenishment policies. A more detailed discussion of these open problems is provided in Section [EC.5](#).

References

- Gagan Aggarwal, Gagan Goel, Chinmay Karande, and Aranyak Mehta. Online vertex-weighted bipartite matching and single-bid budgeted allocations. In *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms*, pages 1253–1264. SIAM, 2011.
- Jackie Baek and Will Ma. Bifurcating constraints to improve approximation ratios for network revenue management with reusable resources. *Operations Research*, 70(4):2226–2236, 2022.
- Fernando Bernstein, A Gürhan Kök, and Lei Xie. Dynamic assortment customization with limited inventories. *Manufacturing & Service Operations Management*, 17(4):538–553, 2015.
- Benjamin E Birnbaum and Claire Mathieu. On-line bipartite matching made simple. *Sigact News*, 39(1):80–87, 2008.
- Agostino Bruzzone, Antonio Cimino, Rafael Diaz, and Francesco Longo. Inventory control with products returns: a state of the art overview. In *Proceedings of the 2010 Spring Simulation Multiconference*, pages 1–7, 2010.
- Jinzhì Bu, Huanan Zhang, and Stefanus Jasin. Asymptotic optimality of simple replenishment policies for a lost-sales inventory system with delivery lead time and purchase returns. *Available at SSRN*, 2023.
- Carri W Chan and Vivek F Farias. Stochastic depletion problems: Effective myopic policies for a class of dynamic optimization problems. *Mathematics of Operations Research*, 34(2):333–350, 2009.
- Marisa P de Brito and Rommert Dekker. Modelling product returns in inventory control—exploring the validity of general assumptions. *International Journal of Production Economics*, 81:225–241, 2003.
- Steven DeLong, Alireza Farhadi, Rad Niazadeh, Balasubramanian Sivan, and Rajan Udawani. Online bipartite matching with reusable resources. *Mathematics of Operations Research*, 49(3):1825–1854, 2024.
- Nikhil R Devanur, Kamal Jain, and Robert D Kleinberg. Randomized primal-dual analysis of ranking for online bipartite matching. In *Proceedings of the twenty-fourth annual ACM-SIAM symposium on Discrete algorithms*, pages 101–107. Society for Industrial and Applied Mathematics, 2013.
- Farbod Ekbatabi, Yiding Feng, and Rad Niazadeh. Online resource allocation with buyback: Optimal algorithms via primal-dual. In *Proceedings of the 24th ACM Conference on Economics and Computation*, pages 583–583, 2023.
- Yiding Feng and Rad Niazadeh. Batching and optimal multistage bipartite allocations. *Management Science*, 71(5):4108–4130, 2025.
- Yiding Feng, Rad Niazadeh, and Amin Saberi. Near-optimal bayesian online assortment of reusable resources. *Operations Research*, 72(5):1861–1873, 2024.
- Negin Golrezaei, Hamid Nazerzadeh, and Paat Rusmevichientong. Real-time optimization of personalized assortments. *Management Science*, 60(6):1532–1551, 2014.
- Xiao-Yue Gong, Vineet Goyal, Garud N Iyengar, David Simchi-Levi, Rajan Udawani, and Shuangyu Wang. Online assortment optimization with reusable resources. *Management Science*, 2021.
- Vineet Goyal, Garud Iyengar, and Rajan Udawani. Asymptotically optimal competitive ratio for online allocation of reusable resources. *Operations Research*, 2025.
- Zhiyi Huang, Zhihao Gavin Tang, Xiaowei Wu, and Yuhao Zhang. Online vertex-weighted bipartite matching: Beating $1-1/e$ with random arrivals. *ACM Transactions on Algorithms (TALG)*, 15(3):1–15, 2019.
- Ganesh Janakiraman and Robin O Roundy. Lost-sales problems with stochastic lead times: Convexity results for base-stock policies. *Operations Research*, 52(5):795–803, 2004.
- Billy Jin and Will Ma. Online bipartite matching with advice: Tight robustness-consistency tradeoffs for the two-stage model. *Advances in Neural Information Processing Systems*, 35:14555–14567, 2022.
- Billy Jin and David P Williamson. Improved analysis of ranking for online vertex-weighted bipartite matching in the random order model. In *International Conference on Web and Internet Economics*, pages 207–225. Springer, 2021.

- Bala Kalyanasundaram and Kirk R Pruhs. An optimal deterministic algorithm for online b-matching. *Theoretical Computer Science*, 233(1-2):319–325, 2000.
- Suho Kang, Ziyang Liu, and Rajan Udhwani. A black-box approach for exogenous replenishment in online resource allocation. In *International Conference on Integer Programming and Combinatorial Optimization*, pages 326–340. Springer, 2025.
- Richard M Karp, Umesh V Vazirani, and Vijay V Vazirani. An optimal algorithm for on-line bipartite matching. In *Proceedings of the twenty-second annual ACM symposium on Theory of computing*, pages 352–358. ACM, 1990.
- Alys Jiaxin Liang, Stefanus Jasin, and Joline Uichanco. Combining a smart pricing policy with a simple replenishment policy: Managing uncertainties in the presence of stochastic purchase returns. *Available at SSRN 3563847*, 2020.
- Alys Jiaxin Liang, Stefanus Jasin, and Xiuli Chao. Assortment and inventory planning under dynamic (stockout-based) substitution in the presence of customer returns: A fluid analysis. *Available at SSRN 4468430*, 2023.
- Irene Lo, Vahideh Manshadi, Scott Rodilitz, and Ali Shameli. Commitment on volunteer crowdsourcing platforms: Implications for growth and engagement. *Manufacturing & Service Operations Management*, 26(5):1787–1805, 2024.
- Will Ma and David Simchi-Levi. Algorithms for online matching, assortment, and pricing with tight weight-dependent competitive ratios. *Operations Research*, 68(6):1787–1803, 2020.
- Vahideh Manshadi and Scott Rodilitz. Online policies for efficient volunteer crowdsourcing. *Management Science*, 68(9):6572–6590, 2022.
- Vahideh Manshadi, Scott Rodilitz, Daniela Saban, and Akshaya Suresh. Online algorithms for matching platforms with multichannel traffic. *Management Science*, 2024.
- Aranyak Mehta. Online matching and ad allocation. *Foundations and Trends® in Theoretical Computer Science*, 8(4):265–368, 2013.
- Aranyak Mehta and Debmalaya Panigrahi. Online matching with stochastic rewards. In *2012 IEEE 53rd Annual Symposium on Foundations of Computer Science*, pages 728–737. IEEE, 2012.
- Aranyak Mehta, Amin Saberi, Umesh Vazirani, and Vijay Vazirani. Adwords and generalized online matching. *Journal of the ACM (JACM)*, 54(5):22–es, 2007.
- Sentao Miao. Managing the inventory with product return: An approximation algorithm. *Available at SSRN*, 2023.
- Kumar Muthuraman, Sridhar Seshadri, and Qi Wu. Inventory management with stochastic lead times. *Mathematics of Operations Research*, 40(2):302–327, 2015.
- Alvin E Roth. *Online and matching-based market design*. Cambridge University Press, 2023.
- Paat Rusmevichientong, Mika Sumida, and Huseyin Topaloglu. Dynamic assortment optimization for reusable products with random usage durations. *Management Science*, 66(7):2820–2844, 2020.
- David Simchi-Levi, Zeyu Zheng, and Feng Zhu. On greedy-like policies in online matching with reusable network resources and decaying rewards. *Management Science*, 2025.
- Vincent P Simpson. Optimum solution structure for a repairable inventory problem. *Operations research*, 26(2):270–281, 1978.
- Jing-Sheng Song, Geert-Jan Van Houtum, and Jan A Van Mieghem. Capacity and inventory management: Review, trends, and projections. *Manufacturing & Service Operations Management*, 22(1):36–46, 2020.
- Rajan Udhwani. Adwords with unknown budgets and beyond. *Management Science*, 71(2):1009–1026, 2025.
- Vijay V Vazirani. Towards a practical, budget-oblivious algorithm for the adwords problem under small bids. In *43rd IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science*, page 21, 2023.

This page is intentionally blank. The e-companion (EC) starts next page.

EC.1. Further Related Work

Online assortment has an extensively growing literature recently, both under the prior-free/adversarial and the Bayesian settings. For non-reusable products, [Bernstein et al. \(2015\)](#) study a model with two products, i.i.d. consumer types and Poisson arrivals. [Chan and Farias \(2009\)](#) study a model with non-stationary consumer types. [Golrezaei et al. \(2014\)](#) propose inventory balancing algorithms for assortment—inspired by the seminal work of [Mehta et al. \(2007\)](#) for online ad allocation—and analyze their performance guarantee in the adversarial setting using a primal-dual approach. This analysis is later improved and generalized to the case where there are finite steps in the penalty function [Ma and Simchi-Levi \(2020\)](#). For online assortment of reusable products, [Rusmevichientong et al. \(2020\)](#); [Baek and Ma \(2022\)](#); [Feng et al. \(2024\)](#) study the Bayesian setting under various modeling assumptions. [Gong et al. \(2021\)](#); [Goyal et al. \(2025\)](#) study the adversarial setting with stationary rental fees and rental duration distributions. [Simchi-Levi et al. \(2025\)](#) introduce a more general framework in the adversarial setting with decaying rental fees and rental duration distributions.

[Vazirani \(2023\)](#); [Udwani \(2025\)](#) study a variant of AdWords problem (cf. [Mehta et al., 2007](#); [Mehta, 2013](#)), where the budget (aka., capacity) of each offline node is initially unknown to the online algorithms, and is only revealed to the algorithms immediately after it is exhausted. They present the competitive ratio analysis for Perturbed Greedy algorithm and Generalized Perturbed Greedy algorithm. [Manshadi et al. \(2024\)](#) introduce a new variant of the online matching with stochastic reward (cf. [Mehta and Panigrahi, 2012](#)). In their model, a fixed fraction of online nodes are single-minded (i.e., each of them has a single incident edge). They show the non-optimality of the Inventory Balancing (IB) in this model, and propose a new modification – Adaptive Capacity (AC) – with improved competitive ratio. In all [Vazirani \(2023\)](#); [Manshadi et al. \(2024\)](#); [Udwani \(2025\)](#), the online algorithms suffer some specific forms of uncertainty about the capacities of offline nodes. However, our model is fundamentally different from them and thus not directly comparable.

Our results are also comparable to the literature on online bipartite matching (with vertex arrival) and its extensions. [DeLong et al. \(2024\)](#) is the most related, where the authors consider the adversarial arrival setting without replenishment under small capacity (in particular, the capacity equals one) and restricts to problem instances with identical duration across offline nodes (aka products) and time. The paper proposes integral matching algorithms that achieve competitive ratios slightly better than $1/2$, but still not optimal. In contrast, our paper allows for heterogeneous duration across offline nodes (aka products) and achieves an asymptotically optimal competitive ratio (rather than looking at the small capacity regime). We refer the reader to related surveys and books (cf. [Mehta, 2013](#); [Roth, 2023](#)) for a comprehensive study of the online bipartite matching literature.

As mentioned earlier, a major source of exogenous shocks to inventory is customer returns, either from the original channel or from other channels. The impact of customer returns on inventories and assortments has been studied in the literature, e.g. recent work of [Liang et al. \(2023, 2020\)](#); [Miao \(2023\)](#). See [Bruzzone et al. \(2010\)](#) for an overview of the older literature on inventory control with returns. In addition, we

interpret endogenous shocks as replenishment of sold units of the product when the retailer employs a simple discrete-period restocking policy with a fixed lead time. This is in fact a simplification and there is a long literature in supply chain and inventory theory on various forms of replenishment policies; for example, see [Bruzzone et al. \(2010\)](#); [Bu et al. \(2023\)](#) for a literature review. Simple restocking with fixed lead time is a special case of stochastic lead time, which has been studied in the literature, e.g. [Janakiraman and Roundy \(2004\)](#); [Muthuraman et al. \(2015\)](#). Also, a more practical restocking policy is base-stock with a lower-threshold ([Simpson, 1978](#); [Song et al., 2020](#)), e.g., restocking only when α fraction or k number of units is sold (instead of each time a unit is sold). We leave further investigations on these models for future research. Recently, [Kang et al. \(2025\)](#) extended the framework developed in this paper to more general settings (e.g., endogenous inventory shocks with i.i.d. stochastic durations). Their work builds directly on our core ideas and demonstrates a general reduction framework that lifts a broad class of algorithms—and their accompanying theoretical guarantees—from the setting without exogenous shocks to one that accommodates arbitrary exogenous inventory shocks.

EC.2. Discussion on Practical Relevance

The online assortment planning model with endogenous and exogenous shocks studied in this paper is rather general and rooted in various applications. In this section, we discuss several such applications in detail and explain the micro-foundations behind both endogenous and exogenous shocks in each context.

Online platforms with rental or reusable products (e.g., cloud computing, car sharing, and hospitality services). Rental-based platforms such as AWS (cloud computing), Zipcar (car sharing), and Airbnb (short-term rentals) are natural instances of online assortment with inventory shocks. When a user selects a server instance, vehicle, or home listing, that unit immediately becomes unavailable for a period equal to the rental duration. This leads to an *endogenous shock*, as the temporary depletion of inventory is triggered directly by the platform’s own allocation decision. Once the usage period ends, the same unit returns to the available inventory, replenishing supply. Meanwhile, *exogenous shocks* occur when new resources unexpectedly enter the system—for example, when a new host lists a property on Airbnb, when additional vehicles are relocated to a specific city, or when unused cloud servers are released from internal operations and become available to external customers. The assortment planner must dynamically adjust which items to display to incoming users, accounting for both the predictable restocking of rented items and the uncertain inflow of new supply.

Dynamic resource allocation in public-service and humanitarian operations. In social-impact and public-sector contexts—such as refugee resettlement, housing allocation, or hospital bed management—resources are inherently reusable. A resource (e.g., a case manager, a housing unit, or a hospital bed) becomes available again once a prior case or patient exits, which creates an *endogenous shock*. At the same time, *exogenous shocks* occur when additional capacity arrives due to new government funding, emergency shelters, or partner institutions joining the network. The central planner’s task is to decide, in real time, which

clients or cases to match with available resources given uncertain future arrivals of both clients and capacity. This is effectively an assortment problem over a dynamically fluctuating inventory, directly mapping to our framework with both endogenous and exogenous shocks.

Volunteer nudging and matching platforms. Many volunteer coordination platforms (e.g., Catchafire or Food Rescue U.S.) face an online assignment problem where *volunteers* serve as reusable resources and *jobs or requests* arrive sequentially over time. When the platform nudges a volunteer to accept a task, that volunteer becomes temporarily unavailable for the task duration, producing an *endogenous shock*. After completing the task, she becomes available again to accept new opportunities. At the same time, *exogenous shocks* occur as new volunteers join, existing volunteers extend their commitment hours, or additional partner organizations contribute workforce capacity. The platform’s decision at each arrival corresponds to an online “assortment” of volunteers or tasks—deciding which subset of volunteers to nudge for each new job while ensuring balanced utilization and avoiding fatigue. Here, endogenous shocks capture temporary unavailability caused by engagement decisions, whereas exogenous shocks capture new volunteer inflows driven by external outreach or seasonal events. The platform’s goal of maximizing completed jobs subject to dynamic volunteer availability aligns precisely with our model’s structure.

Inventory-sharing and circular-economy systems. In sharing-economy or circular supply systems, resources such as tools, clothing, or electronics are repeatedly lent out, repaired, or refurbished, creating *endogenous shocks* that replenish inventory. Each time a borrowed or refurbished unit completes its cycle, it re-enters the pool of available items. In addition, *exogenous shocks* occur through donation campaigns, buy-back programs, or asset transfers from partner organizations. For instance, in a shared-equipment platform like Fat Llama¹⁵, a camera lent to a user for three days returns afterward (endogenous), while new listings or donated items are added unpredictably (exogenous). The platform must dynamically decide which items to display to maximize engagement while balancing short-term utilization against uncertain restocking and inflow rates.

Freelance and gig-economy marketplaces. In marketplaces such as Upwork, Uber, or DoorDash, workers or service providers act as reusable capacity. When a driver accepts a trip or a freelancer takes on a project, she becomes unavailable for a stochastic service duration—an *endogenous shock* arising directly from the platform’s allocation. After completion, the provider re-enters the pool, replenishing supply. Simultaneously, the platform experiences *exogenous shocks* as new drivers sign on, existing workers log back in after a break, or temporary policy changes (e.g., surge incentives, weather effects) alter availability. The platform’s control resembles an assortment decision: which jobs to display or assign to which subset of available providers to optimize earnings and reliability. The interplay between reusable labor capacity (endogenous restocking) and unpredictable supply inflows (exogenous additions) is naturally captured by our model.

¹⁵ <https://fatllama.com/us>

Omni-channel and online retail with returns and restocking. In omni-channel retailing, inventories “come and go” through multiple operational processes. When an online sale is made, the retailer may automatically place a replenishment order with a supplier, which arrives after a lead time—an *endogenous shock* that replenishes the same product based on the previous allocation. Meanwhile, *exogenous shocks* arise when additional units are returned from other sales channels (e.g., physical stores or marketplaces), or when inventory records are corrected due to system discrepancies or shipping cancellations. For example, a pair of shoes sold online might be restocked from the distribution center after three days (endogenous), while at the same time another pair may be returned to the warehouse due to a customer refund (exogenous). The retailer’s online assortment (which items to feature or promote) must thus remain robust to both predictable restocking cycles and unpredictable inflows, motivating assortment algorithms that adapt to both types of inventory shocks.¹⁶

EC.3. Numerical Experiments

In this section, we provide numerical justifications for the performance of our BIB algorithm on synthetic instances.

Policies. In our numerical experiments, we compare the revenue of five different policies:

1. *Batched Inventory Balancing* (BIB): this is the family of policies (Algorithm 1) parameterized by penalty function Ψ and batch-size threshold γ proposed in Section 3. It generalizes the classic Inventory Balancing (IB) when there are inventory shocks. It achieves the asymptotic optimal competitive ratio of $1 - \frac{1}{e} \approx 0.632$ under exponential penalty function $\Psi(x) = \frac{e^{1-x} - e}{1-e}$ and batch-size threshold $\gamma = \sqrt{c_0}$ (Section 3.3).
2. *Static-Capacity Inventory Balancing* (SCIB): this is the family of policies parameterized by penalty function Ψ . It generalizes the classic IB when there are inventory shocks. Specifically, for each time period t and product i , the marginal revenue of all its units is discounted by the same factor of $\Psi(\min\{I_{i,t}/c_i, 1\})$. Here $I_{i,t}$ is the inventory level of product i at time period t . See Algorithm 3 for its formal description.
3. *Dynamic-Capacity Inventory Balancing* (DCIB): this is the family of policies parameterized by penalty function Ψ . It generalizes the classic IB when there are inventory shocks. Specifically, for each time period t and product i , the marginal revenue of all its units is discounted by the same factor of $\Psi(I_{i,t}/(c_i + \sum_{t' \in [t]} \xi_{i,t'}))$. Here $I_{i,t}$ is the inventory level of product i at time period t , and $\xi_{i,t}$ is the new units of product i added due to exogenous shocks. See Algorithm 4 for its formal description.

¹⁶ Notably, lead times are not fixed in practical applications, but for the purpose of theoretical tractability we consider a simplified model with fixed lead times (i.e., fixed duration). We explore the change in durations in our numerical simulations. See Section EC.3.2. As it turns out, our algorithm is robust to moderate fluctuations in the duration.

4. *Unit-Specific Inventory Balancing* (USIB): this is the family of policies parameterized by penalty function Ψ . Each new unit of each product added due to exogenous shocks are treated as a new type of products. This is equivalent to the IBB algorithm (Algorithm 1) with batch size threshold $\gamma = 1$ (which automatically does not distinguish between “charging” and “ready” batches as well). It generalizes the classic IB when there are inventory shocks.
5. *Myopic Greedy* (GREED): this is the policy that offers the (myopic) optimal assortment given the available products at each time. This policy is a special case of IB when $\Psi(x) = \mathbb{1}\{x > 0\}$. Its competitive ratio is at most 0.5.

Algorithm 3: STATIC-CAPACITY INVENTORY BALANCING

Input: penalty function Ψ

```

1 for each time period  $t \in [T]$  do
    /* consumer  $t$  with choice model  $\phi_t$  arrives. */
    /*  $\xi_{i,t}$  new units of product  $i$  are added due to exogenous shock. */
2   update inventory level  $\{I_{i,t}\}_{i \in [n]}$  of all products at time period  $t$ .
3   solve  $S_t \leftarrow \arg \max_{S \in \mathcal{S}} \sum_{i \in S} r_i \Psi(\min\{I_{i,t}/c_i, 1\}) \cdot \phi_t(S, i)$ .
4   allocate assortment  $S_t$  to consumer  $t$ .
```

Algorithm 4: DYNAMIC-CAPACITY INVENTORY BALANCING

Input: penalty function Ψ

```

1 for each time period  $t \in [T]$  do
    /* consumer  $t$  with choice model  $\phi_t$  arrives. */
    /*  $\xi_{i,t}$  new units of product  $i$  are added due to exogenous shock. */
2   update inventory level  $\{I_{i,t}\}_{i \in [n]}$  of all products at time period  $t$ .
3   solve  $S_t \leftarrow \arg \max_{S \in \mathcal{S}} \sum_{i \in S} r_i \Psi(I_{i,t}/(c_i + \sum_{t' \in [t]} \xi_{i,t'})) \cdot \phi_t(S, i)$ .
4   allocate assortment  $S_t$  to consumer  $t$ .
```

EC.3.1. Numerical Performance under Stylized Instances

In this section, we first numerically illustrate that three natural modifications of IB — Static-Capacity Inventory Balancing (SCIB), Dynamic-Capacity Inventory Balancing (DCIB) and Unit-Specific Inventory Balancing (USIB) — cannot achieve the optimal asymptotic competitive ratio $(1 - 1/e)$ even if we consider online bipartite matching instances (i.e., the special case of online assortment where choice model $\phi(S, i) \in \{0, 1\}$ and $\phi(S, i) = 1$ only if $S = \{i\}$) with large initial inventories and no endogenous inventory shock

(i.e., $d_i = \infty$ for every product i). Then we examine the numerical performance of different policies on these constructed instances.

EC.3.1.1. Competitive ratio upper bound of 0.552 for SCIB. Our constructed instance $G(N, r, s)$ is parameterized by $N \in \mathbb{N}$ and $r, s \in [0, 1]$. In this instance, there are $N + 1$ product subgroups $\{U_\ell\}_{\ell \in [0:N]}$, indexed by $[0 : N]$. Each product subgroup U_ℓ with $\ell \in [0 : N]$ contains n_ℓ products, and each product i in subgroup U_ℓ has identical per unit price $r_i \triangleq r^\ell$ and identical initial inventory $c_i \triangleq c$. Here c is a large enough constant. On the other side, there are $2N + 1$ consumer subgroups $\{V_k\}_{k \in [0:2N]}$ indexed by $[0 : 2N]$. Each consumer subgroup V_k with $k \in [0 : 2N]$ contains m_k consumers, who arrive sequentially.

Consumer subgroup V_0 has size $m_0 \triangleq c \cdot n_0$. For each $i \in [n_0]$ and each time period $t \in [(i-1)c + 1 : ic]$, a new consumer from subgroup V_0 who accepts every product from $[n_0 - i + 1] \subseteq U_0$ (i.e., the first $n_0 - i + 1$ products of U_0) arrives.

All remaining consumer subgroups $\{V_k\}_{k \in [2N]}$ and products subgroups $\{U_k\}_{k \in [N]}$ are constructed recursively based on the execution of Dynamic-Capacity Inventory Balancing.¹⁷ Specifically, for each $\ell \in [N]$,

- Consider time period $t_\ell \triangleq \sum_{k \in [0:2(\ell-1)]} m_k$, i.e., the time period when the first consumer from consumer subgroup $V_{2\ell-1}$ arrives. Let x_{i,t_ℓ} be the number of available units of product $i \in U_{\ell-1}$ at the beginning of time period t_ℓ in the algorithm. For each product $i \in U_{\ell-1}$, there are $\xi_{i,t_\ell} \triangleq \max\{s \cdot c - x_{i,t_\ell}, 0\}$ new units added due to the exogenous inventory shocks. As a sanity check, for each product $i \in U_{\ell-1}$ with $\xi_{i,t_\ell} > 0$, its inventory level becomes sc and has a discount factor of $\Psi(s)$ in the algorithm. With slight abuse of notation, let $\xi_\ell \triangleq \sum_{i \in U_{\ell-1}} \xi_{i,t_\ell}$ be the total number of new units added for products from product subgroup $U_{\ell-1}$.
- Product subgroup U_ℓ has size $n_\ell \triangleq \frac{\xi_\ell}{c \cdot (1 - \Psi^{-1}(\frac{\Psi(s)}{r}))}$, where Ψ^{-1} is the inverse function of penalty function Ψ .¹⁸
- Consumer subgroup $V_{2\ell-1}$ have size $m_{2\ell-1} \triangleq \xi_\ell$. For each new unit added for products from product subgroup $U_{\ell-1}$ at time period t_ℓ due to exogenous shocks, there is a consumer in consumer subgroup $V_{2\ell-1}$ who accepts product $i \in U_{\ell-1}$ of this unit and one product $i' \in U_\ell$. Here product i' is constructed such that for each product $i' \in U_\ell$, there are exactly $\frac{\xi_\ell}{n_\ell}$ consumers from $V_{2\ell-1}$ accepts this product. Under this construction, for each consumer in $V_{2\ell-1}$, suppose she accepts product $i \in U_{\ell-1}$ and product $i' \in U_\ell$, the algorithm allocates product $i' \in U_\ell$ to her and collects price r^ℓ , since the discounted price of product $i' \in U_\ell$ is larger than the discounted price of product $i \in U_{\ell-1}$ due to the construction of

¹⁷ Since the algorithm is deterministic, it suffices for us to construct instance $G(N, r, s)$ adaptively, based on the execution of the algorithm.

¹⁸ In the construction of instance $G(N, r, s)$, we impose the requirement that $\frac{\Psi(s)}{r} \in [0, 1]$ for parameters $r, s \in [0, 1]$. Note that under this requirement, together with the assumption that penalty function Ψ is exponential function, the inverse function Ψ^{-1} and n_ℓ are well-defined.

n_ℓ . More specifically, note that the inventory level of product i' when serving this consumer is at least $c - \frac{\xi_\ell}{n_\ell}$. Thus, the discounted price of product i' is at least

$$r_{i'} \Psi \left(1 - \frac{\xi_\ell}{c \cdot n_\ell} \right) = r^\ell \Psi \left(1 - \frac{\xi_\ell}{c \cdot n_\ell} \right) = r^\ell \frac{\Psi(s)}{r} = r^{\ell-1} \Psi(s) = r_i \Psi(s)$$

where the first and last equalities hold due to the construction of $r_i, r_{i'}$ from subgroups $U_{\ell-1}, U_\ell$; and the second equality holds due to the construction of n_ℓ . Because of this allocation in the algorithm, after serving all consumers from consumer subgroup $V_{2\ell-1}$, the inventory level of each product from product subgroup U_ℓ is the same, which is equal to $c - \frac{\xi_\ell}{n_\ell}$.

- Consumer subgroup $V_{2\ell}$ has size $m_{2\ell} \triangleq c \cdot n_\ell$. For the i -th consumer in consumer subgroup $V_{2\ell}$, she accepts the first $n_\ell - \lfloor \frac{i}{c} \rfloor$ products of product subgroup U_ℓ .

The clairvoyant optimum policy assigns a perfect matching between consumer subgroup $V_{2\ell}$ and product subgroup U_ℓ for every $\ell \in [0 : N]$, and assigns every consumer from consumer subgroup $V_{2\ell-1}$ with a unit of her acceptable product from product subgroup $U_{\ell-1}$ that is added at time period t_ℓ due to exogenous shocks for every $\ell \in [N]$. Hence, the total revenue of the clairvoyant optimum policy is $\sum_{\ell \in [0:N]} r^\ell \cdot c \cdot n_\ell + \sum_{\ell \in [N]} r^{\ell-1} \cdot \xi_\ell$.

On the other end, SCIB is sub-optimal because of the revenue loss from two parts. First, it only assigns an imperfect matching between consumer subgroup $V_{2\ell}$ and product subgroup U_ℓ for every $\ell \in [0 : N]$. Moreover, for every $\ell \in [N]$ and every consumer from consumer subgroup $V_{2\ell-1}$, the algorithm allocates a unit of her acceptable product from product subgroup U_ℓ .

Based on the aforementioned instance construction, we numerically evaluate the performance SCIB with exponential penalty function $\Psi(x) = \frac{e^{1-x}-e}{1-e}$ on instance $G(20, 0.5, 0.32)$ with $n_0 = 500$, sufficiently large c and obtain a competitive ratio upper bound of 0.552.

EC.3.1.2. Competitive ratio upper bound of 0.53 for DCIB. Our constructed instance $\hat{G}(N, r, s)$ is parameterized by $N \in \mathbb{N}$ and $r, s \in [0, 1]$. The construction of $\hat{G}(N, r, s)$ is almost the same as $G(N, r, s)$ above, expect that we set $\xi_{i,t_\ell} \triangleq \frac{s \cdot c - x_{i,t_\ell}}{1-s}$ for every $\ell \in [N]$ and $i \in U_{\ell-1}$. Under this modification of $\{\xi_{i,t_\ell}\}$, we still ensure that for every $\ell \in [N]$ and every consumer from consumer subgroup $V_{2\ell-1}$, DCIB allocates a unit of her acceptable product from product subgroup U_ℓ , while the clairvoyant optimum policy allocates a unit of her acceptable product from product subgroup $U_{\ell-1}$.

Based on the aforementioned instance construction, we numerically evaluate the performance DCIB with exponential penalty function $\Psi(x) = \frac{e^{1-x}-e}{1-e}$ on instance $\hat{G}(20, 0.5, 0.3)$ with $n_0 = 500$, sufficiently large c and obtain a competitive ratio upper bound of 0.53.

EC.3.1.3. Competitive ratio upper bound of 0.5 for USIB. Consider an instance $\bar{G}(\epsilon)$ with 2 products. Product 1 has price per unit $r_1 \triangleq 1$ and a large enough initial inventory $c_1 \triangleq c$. Product 2 has price per unit $r_2 \triangleq 1 - \epsilon$ and infinite initial inventory $c_2 \triangleq \infty$. There are $T \triangleq c + 2c^2$ time periods. At each time period $t \in \{c+1, c+3, \dots, T-1\}$, a new unit of product 1 is added due to exogenous inventory shock.

For every time period $t \in [c]$, a consumer who only accepts product 1 arrives. For every time period $t \in \{c+1, c+3, \dots, T-1\}$, a consumer who accepts both products 1 and 2 arrives. For every time period $t \in \{c+2, c+4, \dots, T\}$, a consumer who only accepts product 1 arrives.

The clairvoyant optimum policy allocates a unit of product 1 to every consumer $t \in [t] \cup \{c+2, c+4, \dots, T\}$ and collect per unit price $r_1 = 1$. Meanwhile, it allocates a unit of product 2 to every consumer $t \in \{c+1, c+3, \dots, T-1\}$ and collect per unit price $r_2 = 1 - \epsilon$. Hence, the total revenue of the clairvoyant optimum policy is $T - \frac{T-c}{2}\epsilon = c + 2c^2 - c^2\epsilon$.

On the other end, USIB allocates a unit of product 1 to every consumer $t \in [t] \cup \{c+1, c+3, \dots, T-1\}$, and allocates nothing to every consumer $t \in \{c+2, c+4, \dots, T\}$. Consequently, the total revenue of the algorithm is $c + \frac{T-c}{2} = c + c^2$.

Putting two pieces together, the competitive ratio of the algorithm on this constructed instance is $\frac{c+c^2}{c+2c^2-c^2\epsilon}$ which converges to $\frac{1}{2}$ as ϵ goes to zero and c goes to infinity.

EC.3.1.4. Comparison between policies on constructed instances. Finally, we examine all policies on the aforementioned instances $G(N, r, s)$, $\hat{G}(N, r, s)$, $\bar{G}(\epsilon)$, which numerically upper bound the competitive ratio of SCIB, DCIB, USIB, respectively (Sections EC.3.1.1 to EC.3.1.3). BIB, SCIB, DCIB and USIB all use exponential penalty function $\Psi(x) = \frac{e^{1-x}-e}{1-e}$, and the batch-size threshold γ in BIB is set as $\gamma = 10$.

We record the revenue of different policies in Table EC.1. It can be seen that in all three instances, all four IB policies outperform the myopic greedy policy (GREED). In particular, the revenue gap between BIB and GREED are 15.3%, 12.0%, 103.1% for three instances, respectively. It can also be observed that BIB has the most robust performance for all three instances. In instances $G(10, 0.5, 0.32)$ and $\hat{G}(10, 0.5, 0.3)$, BIB outperforms SCIB and DCIB significantly, and is close to USIB (which is equivalent to BIB with $\gamma = 1$). In instance $\bar{G}(0.1)$, BIB outperforms USIB significantly, and is close to SCIB and DCIB.

Table EC.1 Comparing the average revenue of different policies (Section EC.3.1). Instances $G(N, r, s)$, $\hat{G}(N, r, s)$, $\bar{G}(\epsilon)$ are constructed to upper bound the competitive ratio of SCIB, DCIB, USIB, respectively. For instances $G(N, r, s)$ and $\hat{G}(N, r, s)$, we set $n_0 = 100$ and $c = 50$. For instance $\bar{G}(\epsilon)$, we set $c = 50$.

	$G(10, 0.5, 0.32)$	$\hat{G}(10, 0.5, 0.3)$	$\bar{G}(0.1)$
BIB	5607.72	6295.95	4570.50
SCIB	5103.57	5765.26	4800.00
DCIB	5103.57	5515.90	4800.00
USIB	5848.45	6658.80	2250.00
GREED	4864.29	5622.82	2250.00

EC.3.2. Numerical Performance under Random Instances

In this section, we examine the numerical performance of different policies on the following randomly generated instances.

Experimental setup. We generate the following assortment instances. There are six products indexed by $[6]$ and $T = 3000$ time periods. Each product $i \in [6]$ has an identical initial inventory $c_0 = 30$ and per unit price r_i drawn uniformly between 10 and 25. Without loss of generality, we reindex products such that r_i is decreasing in $i \in [6]$. Each product $i \in [6]$ has identical endogenous inventory shock with lead time $d_i = T/3$ and exogenous inventory shock $\xi_{i,t}$ drawn i.i.d. from geometric distribution $\text{Geo}(0.98)$ for every time period $t \in [T]$.

There are six types of consumers indexed by $1 : 6$. For each consumer type $j \in [6]$, it is associated with a “consideration set” $[j]$ and each product $i \in [j]$ is selected based on a Multinomial Logit (MNL) choice model. Specifically, if assortment $S \subseteq [6]$ is offered, consumers of type $j \in [6]$ selects each product $i \in S$ with probability $\phi^{(j)}(S, i) = \alpha_i^{(j)} / (\alpha_0^{(j)} + \sum_{i' \in S} \alpha_{i'}^{(j)})$, where $\alpha_i^{(j)} = 0$ if $i > j$ (i.e., she does not select a product outside of her consideration set), and $\alpha_i^{(j)}$ is drawn uniformly between 0.9 and 1.1 if $i \in [j]$. Finally, the outside option (i.e., selecting nothing) is picked by constant probability 0.1 by setting $\alpha_0^{(j)}$ so that $\alpha_0^{(j)} / (\alpha_0^{(j)} + \sum_{i' \in S} \alpha_{i'}^{(j)}) = 0.1$.

In each time period $t \in [T]$, a new consumer is sampled from the following stochastic process similar to [Rusmevichientong et al. \(2020\)](#). We divide the whole time horizon into 3 equal-length phases, each with 6 equal-length chunks. Thus, there are in total 18 chunks with equal length $\tau = T/18$. Let $\tau^{(j)} = (6 - j)\tau + 1$ for every $j \in [6]$. At each time period $t \in [T]$, let $t' = t \bmod (T/3) \in [T/3]$, a consumer of type $j \in [6]$ arrives with probability proportional to $e^{-0.001\kappa|t' - \tau^{(j)}|}$. Here, κ is a parameter that controls the arrival pattern of the consumer types. Specifically, $\kappa = 0$ corresponds to i.i.d. arrival of types from the uniform distribution; and the arrival becomes more heterogeneous with convergence towards a deterministic arrival in descending order of types in every phase as κ increases.

We consider four scenarios corresponding to $\kappa \in [4]$. For each scenario, we consider running policies over independent sample paths of the input (including exogenous inventory shocks, type realizations, and consumer choices) by doing 20 iterations of the Monte-Carlo simulation. Batched Inventory Balancing, Static-Capacity Inventory Balancing, Dynamic-Capacity Inventory Balancing and Unit-Specific Inventory Balancing all use exponential penalty function $\Psi(x) = \frac{e^{1-x} - e}{1-e}$, and the batch-size threshold γ in Batched Inventory Balancing is set as $\gamma = 10$.

In addition to the aforementioned experimental setup, we also consider two further variants, both of which extend beyond the model analyzed in our theoretical results. In the first variant, we allow for *negative exogenous inventory shocks*. (Such shocks can be interpreted as unexpected damage to currently available product units.) Specifically, after drawing $\xi_{i,t}$ from the geometric distribution $\text{Geo}(0.98)$, we flip its sign (i.e., set $\xi_{i,t} \leftarrow -\xi_{i,t}$) with probability 0.2. In the second variant, we allow for *i.i.d. stochastic endogenous inventory*

Table EC.2 Comparing the average revenue of different policies (Section EC.3.2) in the original setup.

	$\kappa = 0$	$\kappa = 1$	$\kappa = 2$	$\kappa = 3$
BIB	27836.91	32308.41	29972.79	23121.82
SCIB	27747.12	32197.73	29775.70	22914.32
DCIB	27852.76	32316.99	30101.17	23170.69
USIB	27017.10	31385.22	29070.33	22572.46
GREED	25999.49	30184.01	27926.82	21502.78

shocks. Specifically, for each consumer and her selected product unit, we draw a random duration independently from the geometric distribution $\text{Geo}(3/T)$, which has an expected value of $T/3$.

Results. We record the revenues of all 20 iterations of the Monte-Carlo simulation for each experiment setup.

For the original setup, we use “box and whisker” plots to demonstrate the median revenues and the resulting confidence intervals (Figure EC.1) and compute the average revenue as an estimate for the expected revenue of each policy (Table EC.2). It is clear from both Figure EC.1 and Table EC.2 that all four IB policies outperform the myopic greedy policy and the largest revenue gap is 7.1%, 7.0%, 7.8%, 7.8% for all four scenarios $\kappa \in [4]$, respectively. Among all four IB policies, Unit-Specific Inventory Balancing has the lowest revenue, Static-Capacity Inventory Balancing has the second lowest revenue. Batched Inventory Balancing and Dynamic-Capacity Inventory Balancing have almost the same best revenue.

For the variant with negative exogenous inventory shocks, we compute the average revenue as an estimate of the expected revenue for each policy in Table EC.3. The numerical patterns observed in the original setup largely carry over: all four IB policies consistently outperform the myopic greedy benchmark. However, the revenue gap between the IB policies and the greedy policy becomes smaller. In particular, the largest improvements are approximately 1.9%, 1.8%, 1.7%, and 2.1% for the four scenarios $\kappa \in [4]$, respectively. Among the IB policies, Unit-Specific Inventory Balancing always yields the lowest revenue, while Static-Capacity Inventory Balancing and Dynamic-Capacity Inventory Balancing perform better and, together with Batched Inventory Balancing, achieve nearly identical top performance. More specifically, Batched Inventory Balancing attains the highest revenue when $\kappa = 1, 2$, Static-Capacity Inventory Balancing is best for $\kappa = 3$, and Dynamic-Capacity Inventory Balancing is best for $\kappa = 4$.

For the variant with geometrically distributed endogenous inventory shocks, the average revenues of all policies are reported in Table EC.4. Consistent with the previous setups, all four IB policies outperform the myopic greedy benchmark. The largest revenue improvements are about 6.3%, 6.0%, 5.5%, and 4.7% for the four scenarios $\kappa \in [4]$, respectively. Among the IB policies, Unit-Specific Inventory Balancing again yields the lowest revenue, while Batched Inventory Balancing, Static-Capacity Inventory Balancing, and Dynamic-Capacity Inventory Balancing remain closely matched and achieve highly competitive performance across all scenarios, with only minor variations in which policy attains the highest revenue.




Figure EC.1 Box and whisker comparison of revenue under different policies (Section EC.3.2); Results are based on 20 iterations of Monte-Carlo simulation.

Table EC.3 Comparing the average revenue of different policies (Section EC.3.2) with negative exogenous inventory shocks.

	$\kappa = 0$	$\kappa = 1$	$\kappa = 2$	$\kappa = 3$
BIB	13061.32	18746.69	20766.04	20000.45
SCIB	13056.95	18746.49	20782.17	20002.64
DCIB	13046.72	18744.64	20765.96	20007.51
USIB	12844.23	18498.98	20542.71	19673.44
GREED	12816.43	18409.63	20435.91	19602.96

EC.4. Illustrative Examples for the Interval Assignment Problem

Table EC.4 Comparing the average revenue of different policies (Section EC.3.2) with geometrically distributed stochastic endogenous inventory shocks.

	$\kappa = 0$	$\kappa = 1$	$\kappa = 2$	$\kappa = 3$
BIB	13061.32	18746.69	20766.04	20000.45
SCIB	13056.95	18746.49	20782.17	20002.64
DCIB	13046.72	18744.64	20765.96	20007.51
USIB	12844.23	18498.98	20542.71	19673.44
GREED	12816.43	18409.63	20435.91	19602.96

In this section, we present instances of the interval assignment problem (IAP) to illuminate its combinatorial structure, motivate the need for a carefully constructed assignment, and the execution of our proposed Algorithm 2.

The failure of real consumption levels. As explained in Section 4.2 and in the proof of Theorem 1, the output of the IAP is used to define a transformed version of the consumption level (i.e., total initial inventory minus the remaining inventory at the time of allocation). This raises a natural question: does the real consumption level—defined as $|\mathcal{N}_i| = |\{j : a_i \in [a_j, b_j]\}|$ for each interval i —already satisfy the three defining properties of the IAP, namely **LOCAL-DOMINANCE**, **GLOBAL-DOMINANCE**, and **PARTITION-MONOTONICITY**? The answer is negative. The following example demonstrates that the naive assignment $\hat{I}_i = |\mathcal{N}_i|$ may violate **PARTITION-MONOTONICITY**, even though it satisfies the other two properties by construction.

EXAMPLE EC.1 (A PATH OF OVERLAPPING INTERVALS). Consider four intervals arranged in a “path” where each overlaps only with its immediate predecessor and successor:

$$[a_1, b_1] = [1, 4], \quad [a_2, b_2] = [3, 6], \quad [a_3, b_3] = [5, 8], \quad [a_4, b_4] = [7, 10].$$

Here, the left endpoints are strictly increasing ($a_1 < a_2 < a_3 < a_4$), and for each i , the coverage set $\mathcal{N}_i = \{j : a_i \in [a_j, b_j]\}$ satisfies $\mathcal{N}_1 = \{1\}$, $\mathcal{N}_2 = \{1, 2\}$, $\mathcal{N}_3 = \{2, 3\}$, and $\mathcal{N}_4 = \{3, 4\}$. An illustration is provided in Figure EC.2.

Under the naive assignment $\hat{I}_i = |\mathcal{N}_i|$, we obtain $(\hat{I}_1, \hat{I}_2, \hat{I}_3, \hat{I}_4) = (1, 2, 2, 2)$ for Example EC.1. This assignment satisfies **LOCAL-DOMINANCE** (each coverage set contains labels that dominate $\{1, \dots, |\mathcal{N}_i|\}$) and **GLOBAL-DOMINANCE** (since $\hat{I}_i = |\mathcal{N}_i|$, the inequality holds with equality for any decreasing concave penalty). However, it *violates* **PARTITION-MONOTONICITY**: there is only one interval labeled “1”, yet three intervals labeled “2”. Consequently, it is impossible to partition the intervals into groups such that each group of size m contains exactly the labels $\{1, 2, \dots, m\}$.

This failure is significant because **PARTITION-MONOTONICITY** plays a crucial role in aligning the primal and dual objective values in the proof of Theorem 1. Without it, setting the dual variables equal to the




Figure EC.2 Example EC.1: A path of four overlapping intervals. Each interval overlaps only with its immediate neighbor. Both endpoints (a_i, b_i) for each interval i are marked in red. The bold number inside each interval indicates its IAP assignment $\hat{I}_i$, and the blue dashed arrows illustrate the predecessor links $p_j = i$ that define the partition satisfying **PARTITION-MONOTONICITY**.

real consumption levels would not recover the tight approximation guarantee established by Golrezaei et al. (2014) in the special case with no endogenous shock.

In contrast, Algorithm 2 produces a valid IAP assignment for Example EC.1. Specifically, it outputs $(\hat{I}_1, \hat{I}_2, \hat{I}_3, \hat{I}_4) = (2, 1, 2, 1)$, which satisfies all three properties. It constructs this assignment in the following order: $\hat{I}_3 \leftarrow 2$, $\hat{I}_1 \leftarrow 2$, $\hat{I}_4 \leftarrow 1$, and $\hat{I}_2 \leftarrow 1$. As illustrated in Figure EC.2, the resulting labels enable a partition into two disjoint chains— $\{1, 2\}$ and $\{3, 4\}$ —each of which satisfies **PARTITION-MONOTONICITY**.

The failure of a simple greedy procedure. Example EC.1 demonstrates that the naive assignment based on real consumption levels is insufficient, thereby motivating the need for an algorithm that explicitly enforces all three properties of IAP.

A natural follow-up question is whether a simpler algorithm—beyond the full IAP construction—might suffice. One such candidate is the following *greedy* procedure: process the intervals in increasing order of their left endpoints, and for each interval, assign the *smallest* positive integer that maintains **LOCAL-DOMINANCE** with respect to all previously processed coverage sets. The rationale is that **LOCAL-DOMINANCE** only constrains intervals whose left endpoints have already been encountered, making a forward greedy choice seemingly safe.

Indeed, for Example EC.1, this greedy procedure yields the assignment $(1, 2, 1, 2)$, which satisfies **LOCAL-DOMINANCE**, **GLOBAL-DOMINANCE**, and **PARTITION-MONOTONICITY**. However, as the next example shows, this procedure does not generalize: on more complex instances, the greedy assignment may satisfy **LOCAL-DOMINANCE** but violate **PARTITION-MONOTONICITY**, and thus fail to meet the full requirements of the IAP.

EXAMPLE EC.2 (A BINARY-TREE STRUCTURE OF INTERVALS). Consider seven intervals arranged in a complete binary tree of depth 3 (i.e., three levels), defined as follows:

$$\text{Level 1 (root): } [a_1, b_1] = [1, 14],$$

$$\text{Level 2 (children of root): } [a_2, b_2] = [2, 7], \quad [a_5, b_5] = [8, 13],$$




Figure EC.3 Example EC.2: A complete binary-tree hierarchy of seven intervals. Each parent fully contains its two children. Both endpoints (a_i, b_i) for each interval i are marked in red. The bold number inside each interval indicates its IAP assignment $\hat{I}_i$, and the blue dashed arrows illustrate the predecessor links $p_j = i$ that define the partition satisfying **PARTITION-MONOTONICITY**.

Level 3 (leaves): $[a_3, b_3] = [3, 4]$, $[a_4, b_4] = [5, 6]$, $[a_6, b_6] = [9, 10]$, $[a_7, b_7] = [11, 12]$.

where the left endpoints are strictly increasing ($a_1 < a_2 < \dots < a_6 < a_7$). Each interval is fully contained within its parent, and for every i , the coverage set $\mathcal{N}_i = \{j : a_i \in [a_j, b_j]\}$ consists precisely of interval i and all its ancestors in the tree. An illustration is provided in Figure EC.3.

For Example EC.2, the aforementioned greedy procedure yields the assignment $(1, 2, 3, 3, 2, 3, 3)$,¹⁹ which satisfies **LOCAL-DOMINANCE**, **GLOBAL-DOMINANCE**, but violates **PARTITION-MONOTONICITY**: there is only one interval labeled “1”, yet two intervals labeled “2”, and four intervals labeled “3”. Consequently, it is impossible to partition the intervals into groups such that each group of size m contains exactly the labels $\{1, 2, \dots, m\}$.

In contrast, although more involved, Algorithm 2 produces a valid IAP assignment for Example EC.2. Specifically, it outputs assignment $(3, 2, 1, 1, 2, 1, 1)$, which satisfies all three properties.²⁰ It constructs this assignment in the following order: $\hat{I}_1 \leftarrow 3$, $\hat{I}_5 \leftarrow 2$, $\hat{I}_2 \leftarrow 2$, $\hat{I}_7 \leftarrow 1$, $\hat{I}_6 \leftarrow 1$, $\hat{I}_4 \leftarrow 1$, and $\hat{I}_3 \leftarrow 1$. As illustrated in Figure EC.3, the resulting labels enable a partition into disjoint chains— $\{1, 2, 3\}$, $\{5, 6\}$, $\{4\}$, and $\{7\}$ —each of which satisfies **PARTITION-MONOTONICITY**.

EC.5. Open Problems

Our analysis of online assortment under inventory shocks is built on a clean abstraction of reusability—where endogenous shocks model the automatic return of allocated units after a fixed usage duration—and exogenous shocks capture unpredictable inflows of new supply. While this model captures a wide range of

¹⁹ The real consumption levels induce the same assignment, which violates **PARTITION-MONOTONICITY**, for Example EC.2.

²⁰ In Example EC.2, the valid IAP assignment $(3, 2, 1, 1, 2, 1, 1)$ can be considered as the inventory level. However, this does not imply that the real inventory level is always valid for IAP assignment. To see this, consider a modification of Example EC.2 in which the left endpoints are reordered—specifically, set $a_3 = 1$ and $a_1 = 3$, while keeping all other left and right endpoints unchanged. In this modified example, the inventory level becomes $(1, 2, 3, 1, 2, 1, 1)$. Although this assignment satisfies **GLOBAL-DOMINANCE** and **PARTITION-MONOTONICITY**, it violates **LOCAL-DOMINANCE**: for interval 4, the coverage set $\mathcal{N}_4$ includes intervals whose assigned values are insufficient to dominate the required set $\{1, \dots, |\mathcal{N}_4|\}$.

applications, several natural extensions arise from both practical considerations and theoretical curiosity. We outline four promising directions for future work.

Negative inventory shocks and loss modeling. Our current model assumes non-negative exogenous shocks. However, in real-world settings, inventory may be unexpectedly reduced due to damage, spoilage, theft, or system errors—phenomena we model as *negative exogenous shocks*. As shown in Section EC.3, BIB exhibits strong empirical performance even when negative shocks are introduced. From a theoretical perspective, negative shocks introduce nonlinearity into the inventory dynamics (e.g., the “available inventory” is the max of zero and prior inventory minus shock), which complicates the design of a tractable LP benchmark. Developing competitive algorithms with formal guarantees under such adverse shocks—potentially by combining our batching ideas with robustness techniques from convex optimization and duality analysis—is a compelling direction for future research.²¹

Stochastic and heterogeneous return durations. Our analysis assumes deterministic and homogeneous usage durations d_i . However, in many applications (e.g., gig platforms, hospital beds, rental markets), the duration a unit remains unavailable is random and may vary across allocations. Recent work by Kang et al. (2025) shows that our batched framework can be lifted via a black-box reduction to settings with i.i.d. stochastic return times, preserving the $1 - 1/e$ competitive ratio under mild conditions. An open question is whether similar guarantees can be obtained for adversarial or highly heterogeneous duration distributions, or whether new algorithmic ideas are required.

Endogenous shocks from operational replenishment policies. In omni-channel retail and supply chain settings, endogenous inventory replenishment often stems not from usage-driven returns but from operational policies such as (s, S) -type or threshold-based restocking. In such models, a replenishment order is triggered when on-hand inventory falls below a threshold, and units arrive after a (possibly stochastic) lead time. Extending our framework to accommodate such endogenous dynamics would require reinterpreting the “return duration” as a lead time and incorporating the replenishment decision into the adversary’s power or the algorithm’s control.

Incorporating inventory holding and restocking costs. Our objective focuses solely on revenue maximization, ignoring potential costs associated with holding inventory or placing restocking orders. Extending the model to include holding costs (e.g., per-unit-per-period) or fixed restocking costs would shift the objective toward profit maximization and likely favor more conservative allocation policies. Integrating such costs into the primal-dual analysis—perhaps via modified penalty functions or cost-aware dual variables—remains an intriguing challenge.

²¹ Using a standard charging argument, the myopic greedy policy (equivalently, the Inventory Balancing algorithm with penalty function $\Psi(x) = \mathbb{1}\{x > 0\}$) can be shown to be 0.5-competitive under adversarial negative shocks.

Joint inventory-assortment control under shocks. In our model, exogenous shocks are fully adversarial and uncontrollable. However, in practice, platforms may exert partial control over inventory inflows—for example, by procuring additional units, triggering returns, or managing volunteer sign-up campaigns. A natural generalization is a *joint inventory-assortment control* problem, where the decision maker simultaneously chooses both the assortment and an inventory action (e.g., order quantity, promotional effort) at each time step. Developing competitive algorithms for this richer setting would require coupling the batched inventory balancing framework with dynamic inventory control techniques.