

Online Resource Allocation under Partially Predictable Demand

Dawsen Hwang

Google, Chicago, IL 60607, dawsen@gmail.com

Patrick Jaillet

Department of Electrical Engineering and Computer Science, Laboratory for Information and Decision Systems, Operations Research Center, Massachusetts Institute of Technology, Cambridge, MA 02139, jaillet@mit.edu

Vahideh Manshadi

Yale School of Management, New Haven, CT 06511, vahideh.manshadi@yale.edu

For online resource allocation problems, we propose a new demand arrival model where the sequence of arrivals contains both an adversarial component and a stochastic one. Our model requires no demand forecasting; however, due to the presence of the stochastic component, we can partially predict future demand as the sequence of arrivals unfolds. Under the proposed model, we study the problem of the online allocation of a single resource to two types of customers, and design online algorithms that outperform existing ones. Our algorithms are adjustable to the relative size of the stochastic component, and our analysis reveals that as the portion of the stochastic component grows, the loss due to making online decisions decreases. This highlights the value of (even partial) predictability in online resource allocation. We impose no conditions on how the resource capacity scales with the maximum number of customers. However, we show that using an adaptive algorithm—which makes online decisions based on observed data—is particularly beneficial when capacity scales linearly with the number of customers. Our work serves as a first step in bridging the long-standing gap between the two well-studied approaches to the design and analysis of online algorithms based on (1) adversarial models and (2) stochastic ones. Using novel algorithm design, we demonstrate that even if the arrival sequence contains an adversarial component, we can take advantage of the limited information that the data reveals to improve allocation decisions. We also study the classical secretary problem under our proposed arrival model, and we show that randomizing over multiple stopping rules may increase the probability of success.

Key words: online resource allocation, competitive analysis, analysis of algorithms

1. Introduction

E-commerce platforms host markets for perishable resources in various industry sectors ranging from airlines to hotels to internet advertising. In these markets, demand realizes sequentially, and the firms need to make online (irrevocable) decisions regarding how (and at what price) to allocate resources to arriving demand without precise knowledge of future demand. The success of any online

allocation algorithm depends crucially on a firm’s ability to predict future demand. If demand can be predicted, then under some conditions on the amount of available resources, making online decisions incurs little loss (as shown in Agrawal et al. (2014), among others). However, in many markets, demand cannot be perfectly predicted due to unpredictable components such as traffic spikes and strategy changes by competitors. The emergence of sharing-economy platforms such as Airbnb, which can scale supply at negligible cost and on short notice (Zervas et al. 2016), has significantly added to unpredictable variability in demand even for products that are not new (e.g., existing hotels).

In such cases firms can take a worst-case approach and assume that demand is controlled by an imaginary adversary and thus is unpredictable. Such an approach, however, usually results in online policies that are too conservative (as studied in Ball and Queyranne (2009) and others). Instead, firms may wish to employ online policies based on models that assume the future demand can partially be predicted, avoiding being too conservative while not being reliant on fully accurate predictions. This paper aims to investigate to what extent the above goal is achievable. We propose a new demand model, called *partially predictable*, that contains both adversarial (thus unpredictable) and stochastic (predictable) components. We design novel algorithms to demonstrate that even though demand is assumed to include an unpredictable component, firms can make use of the limited information that the data reveals and improve upon the completely conservative approach.

We study a basic online allocation problem of a single resource with an arbitrary capacity to a sequence of customers, each of which belongs to one of two types. Each customer demands one unit of the resource. If the resource is allocated, the firm earns a type-dependent revenue. Type-1 and -2 customers generate revenue of 1 and $a < 1$, respectively. Our demand model takes a parameter $0 < p < 1$ and works as follows. An unknown number of customers of each type will be revealed to the firm in unknown order. Both the number and the order of customers are assumed to be controlled by an imaginary adversary. However, a fraction p of randomly chosen customers does *not follow* this prescribed order and instead arrives at uniformly random times. This group of customers represents the stochastic component of the demand that is mixed with the adversarial element. Although we cannot identify which customers belong to the stochastic group, we can still *partially* predict future demand, because this group is almost uniformly spread over the time horizon. Therefore parameter p determines the level of predictability of demand.

From a practical point of view, our demand model requires *no forecast* for the number of customers of each type prior to arrival; instead, it assumes a rather mild “regularity” in the arrival pattern: a fraction p of customers of each type is spread throughout the time horizon. We motivate this through a simple example. Suppose an airline launches a new flight route for which it has no demand forecast. However, using historical data on customer booking behavior, the airline

knows that there is heterogeneity in booking behavior of customers, namely, the time they request a booking varies across customers of each type. Such heterogeneity results in the gradual arrival of a portion of customers of each type. For example, CWT (2016) illustrates a significant disparity in the advanced booking behavior of business travelers based on their age, gender, and travel frequency. Therefore, the airline can reasonably assume that demand from business travelers (who correspond to type-1 in our model) is, to some degree, spread over the sale horizon.

From a theoretical point of view, our demand model aims to address the limitations of the main two approaches that have been taken so far in the literature: (1) adversarial models and (2) stochastic ones.¹ Under the adversarial modeling approach, the sequence of arrivals is assumed to be completely unpredictable. The online algorithms developed for these models aim to perform well in the worst-case scenario, often resulting in very conservative bounds (see Ball and Queyranne (2009) for the single-resource revenue management problem and Mehta et al. (2007) and Buchbinder and Naor (2009) for online allocation problems in internet advertising). On the other hand, the stochastic modeling approach assumes that demand follows an unknown distribution (Kleinberg 2005, Devanur and Hayes 2009, Agrawal et al. 2014).² In this case we can predict future demand after observing a small fraction of it. For example, after observing the first 10% of the demand, if we observe that 15% of customers are of type-1, we can predict that roughly 15% of the remaining customers are also of type-1. The limitation of such an approach is that it cannot model variability across time. In some cases, real data does not confirm the stochastic structure presumed in these models, as shown in Wang et al. (2006) and Shamsi et al. (2014). In fact, as discussed in Mirrokni et al. (2012) and Esfandiari et al. (2015), large online markets (such as internet advertising systems) often use modified versions of these algorithms to make them less reliant on accurate demand prediction. Our model provides a middle ground between the aforementioned approaches by assuming that the arrival sequence contains both an adversarial component and a stochastic one.

For the above problem, we design two online algorithms (a non-adaptive and an adaptive one³) that perform well in the partially predictable model. We use the metric of competitive ratio, which is commonly used to evaluate the performance of online algorithms. Competitive ratio is the worst-case ratio between the revenue of the online scheme to that of a clairvoyant solution (see Definition 1). The competitive ratio of our algorithms is parameterized by p , and for both algorithms the ratio increases with p : *as the relative size of the stochastic component grows, the*

¹ A few papers consider arrival models outside these two categories. We carefully review them and compare them with our model in Section 2.

² In fact these papers assume a more general model, the random order model, that we discuss in Section 2.

³ We call an algorithm “adaptive” if it makes decisions based on the sequence of arrivals it has observed so far.

loss due to making online decisions decreases. We further show that using an adaptive algorithm is particularly beneficial when the capacity scales linearly with the maximum number of customers. Our algorithms are easily adjustable with respect to parameter p . Therefore, if a firm wishes to use different levels of predictability for different products, then it can use the same algorithm with different parameters p .

In designing of our algorithms, we keep track of the number of accepted customers of each type, and we decide whether to accept an arriving type-2 customer by comparing the number of already accepted type-2 customers with optimally designed *dynamic* thresholds.⁴ Our non-adaptive algorithm strikes a balance between “smoothly” allocating the inventory over time (by not accepting many type-2 customers toward the beginning) and not protecting too much inventory for potential late-arriving type-1 customers (see Algorithm 1 and Theorem 1). Our adaptive algorithm frequently recomputes upper bounds on the number of future customers of each type based on observed data and uses these upper bounds to ensure that we protect enough inventory for future type-1 customers. We show that such an adaptive policy significantly improves the performance guarantee when the initial inventory is large relative to the maximum number of customers (see Algorithm 2 and Theorem 2). Both algorithms could reject a type-2 customer early on but accept another type-2 customer later. This is consistent with practice. For example, in online airline booking systems, lower fare classes can open up after being closed out previously (Cheapair 2016).

From a methodological standpoint, an analysis of the competitive ratio of our algorithms presents many new technical challenges arising from the fact that our arrival model contains both an adversarial and a stochastic component. Our analysis crucially relies on a concentration result that we establish for our arrival model (see Lemma 1) as well as fairly intricate case analyses for both algorithms. Further, to prove the lower bound on the competitive ratio of our adaptive algorithm, we construct a novel *factor-revealing* nonlinear mathematical program (see MP1 and Section 5.2).

The two extreme cases of our model where all or none of the customers belong to the adversarial group (i.e., $p = 0$ and $p = 1$) reduce to the adversarial and stochastic modeling approaches that have been mainly studied in the literature thus far (for instance, Ball and Queyranne (2009) study the former model and Agrawal et al. (2014) study the latter). Our algorithms recover the known performance guarantees for these two extreme cases. For the regime in between (i.e., when $0 < p < 1$), we show that our algorithms achieve competitive ratios better than what can be achieved by any of the algorithms designed for these extreme cases (or even any combination of them). This highlights the need to design new algorithms when departing from traditional arrival models.

We also study the classic secretary problem under our partially predictable arrival model. The secretary problem, a stopping time problem, corresponds to the setting in which we initially have

⁴ We always accept a type-1 customer if there is remaining inventory.

one unit of inventory; each customer is of a different type, and we wish to maximize the probability of allocating the inventory to the type generating the highest revenue. We show that, unlike the classic setting (which corresponds to $p = 1$ in our model), the celebrated deterministic stopping rule policy based on a deterministic observation period is no longer optimal. Due to the presence of the adversarial component, randomizing over the length of the observation period may result in improvement (see Algorithm 3, Theorem 3, and Proposition 3).

We conclude this section by highlighting our motivations and contributions. For many applications, demand arrival processes are inherently prone to contain unpredictable components that even advanced information technologies cannot mitigate. An allocation policy whose design is based on stochastic modeling cannot incorporate the presence of such unpredictable components. At the same time, taking a worst-case adversarial approach usually leads to allocation policies that are too conservative. We introduce the *first arrival model* that contains *both adversarial (thus unpredictable) and stochastic* components. Through novel algorithm design, we show that (1) we can take advantage of even limited available information (due to the presence of the stochastic component) to improve a firm’s revenue when compared to algorithms that take a worst-case approach and that (2) there is an unavoidable loss due to the presence of an adversarial component, which emphasizes *the value of stochastic information and predictability* in online resource allocation.

The rest of the paper is organized as follows. In Section 2, we review the related literature and highlight the differences between the current paper and previous work. In Section 3, we formally introduce our demand arrival model and our performance metric, and prove a consequential concentration result for the arrival process. Sections 4 and 5 are dedicated to description and analysis of our two algorithms. In Section 6, we present upper bounds on the performance of any online algorithm, and we compare the performance of our algorithms with that of existing ones. Section 7 studies the secretary problem under our new arrival model. In Section 8, we conclude by outlining several directions for future research. For the sake of brevity, we include proofs of only selected results in the main text. Detailed proofs of the remaining statements are deferred to clearly marked appendices.

2. Literature Review

Online allocation problems have broad applications in revenue management, internet advertising, scheduling appointments (through web applications) in health care, just to name a few. Thus it has been studied in various forms in operations research and management, as well as computer science. As discussed in the introduction, the approach taken in modeling the arrival process is the first consequential step in studying these problems. Therefore, in this literature review, we

categorize related streams of research by modeling approach rather than by the particular problem formulation and application.

First, we note that the single-leg revenue management (RM) problem and its generalizations have been extensively studied using frameworks other than online resource allocation problems and competitive analysis. Earlier papers assumed *low-before-high* models (where all low-fare demand realizes before high-fare demand) with known demand distributions (Belobaba 1987, 1989, Brumelle and McGill 1993, Littlewood 2005) or assumed the arrival process is known, and formulated the problem as a Markov decision problem (Lee and Hersh 1993, Lautenbacher and Stidham Jr. 1999). We refer the reader to Talluri and Van Ryzin (2006) for a comprehensive review of RM literature. Further, many recent papers in revenue management study dynamic pricing when the underlying price-sensitive demand process is unknown. See, for example, seminal work by Besbes and Zeevi (2009) and Araman and Caldentey (2009). For the sake of brevity, we will not review these streams of work.

Adversarial models: Ball and Queyranne (2009) studied the single-leg revenue management problem under an adversarial model and showed that in the two-fare case the optimal competitive ratio is $\frac{1}{2-a}$ where $a < 1$ is the ratio of two fares. As discussed in the introduction, our model reduces to that of Ball and Queyranne (2009) for $p = 0$. In this special case, our non-adaptive algorithm reduces to the threshold policy of Ball and Queyranne (2009) and recovers the same performance guarantee. However, when $0 < p < 1$, we show that for a certain class of instances our algorithms perform better than that of Ball and Queyranne (2009) (see Subsection 6.2), indicating the need for new algorithms for our new arrival model.

Several papers studied the adwords problem under the adversarial model (Mehta et al. 2007, Buchbinder and Naor 2009). This problem concerns allocating ad impressions to budget-constrained advertisers. As mentioned in Mehta et al. (2007), even though the optimal competitive ratio under an adversarial model is $1 - 1/e$, one would expect to do better when statistical information is available. Later, Mirrokni et al. (2012) showed that it is impossible to design an algorithm with a near-optimal competitive ratio under both adversarial and random arrival models. Such an impossibility result affirms the need for new modeling approaches to serve as a middle ground between these two models. Our paper takes a step in this direction.

Stationary stochastic models: A general form of these models is the *random order model*, which assumes that the sequence of arrivals is a random permutation of an arbitrary sequence (Kleinberg 2005, Devanur and Hayes 2009, Agrawal et al. 2014). In such a model, after observing a small fraction of the input, one can predict pattern of future demand. This intuition is used to develop primal- and dual-based online algorithms that achieve near-optimal revenue, under

appropriate conditions on the relative amount of available resources to allocate. These algorithms rely heavily on learning from observed data, either once (Devanur and Hayes 2009) or repeatedly (Kleinberg 2005, Agrawal et al. 2014, Kesselheim et al. 2014). As discussed in the introduction, arrival patterns could experience high variability across time, limiting the performance of these algorithms in practice (Mirrokni et al. 2012, Esfandiari et al. 2015). We note that assuming i.i.d. arrivals with known or unknown distributions also falls into this category of modeling approaches. Several revenue management papers provided asymptotic analysis of linear programming-based (LP-based) approaches for such settings; see Talluri and Ryzin (1998), Cooper (2002) and Jasin (2015).

Our model reduces to a special case of the model of Agrawal et al. (2014) for $p = 1$, and like their algorithm, ours also achieves near-optimal revenue when $p = 1$. However, when $0 < p < 1$, we show, in Subsection 6.2, that for a certain class of instances our algorithms perform better than that of Agrawal et al. (2014).

Nonstationary stochastic models: Motivated by advanced service reservation and scheduling, Wang and Truong (2015) and Stein et al. (2016) studied online allocation problems where demand arrival follows a *known* nonhomogeneous Poisson process. For such settings, they developed online algorithms with constant competitive ratios. Further, Ciocan and Farias (2012) considered another interesting setting where the (unknown) arrival process belongs to a broad class of stochastic processes. They proved a constant factor guarantee for the case where arrival rates are uniform. Our modeling strategy differs from both approaches by assuming that $(1 - p)$ fraction of the input is adversarial. Even for the stochastic component, we assume no prior knowledge of the distribution. However, we limit the adversary’s power by assuming that these two components are *mixed*. Also, we note that the aforementioned papers studied more general allocation problems in settings like network revenue management.

Other models: Several earlier papers also acknowledged and addressed the limitation of both the adversarial and random order (or stochastic) models using various approaches. Mahdian et al. (2007) and Mirrokni et al. (2012) considered allocation problems where the demand can *either* be perfectly estimated or adversarial. They designed and analyzed algorithms that have good performance guarantees in both worst-case and average-case scenarios. Unlike these works, our demand model contains *both* stochastic and adversarial components at the same time, and we design algorithms that take advantage of partial predictability.

Another approach to address unpredictable patterns in demand is to use robust stochastic optimization (Ben-Tal and Nemirovski 2002, Bertsimas et al. 2004). These papers aim to optimize allocations when the demand belongs to a class of distributions (or uncertainty set). This approach

limits the adversary’s power by restricting the class of demand distributions. Here, we take a different approach. We do not limit the class of distribution that the adversary can choose from; instead, we assume that a fraction p of the demand will not follow the adversary.

Lan et al. (2008) also took a robust approach, studying the single-leg multi-fare class revenue management problem in a very interesting setting, where the only prior knowledge about demand is the lower and upper bounds on the number of customers from each fare class. Lan et al. (2008) used fixed upper and lower bounds to develop optimal static policies in the form of nested booking limits, and also showed that dynamically adjusting these policies can improve the competitive ratio. Unlike their work, we do not assume prior knowledge of lower and upper bounds on the number of customers from each class. Instead, in our model, we learn the bounds as the sequence unfolds.

Shamsi et al. (2014) used a real data set from display advertising at AOL/Advertising.com to show that arrival patterns do not satisfy the crucial property implied by assuming a random order model for demand. In particular, they showed that the dual prices of the offline allocation problem at different times can vary significantly. They used a risk minimization framework to devise allocation rules that outperform existing algorithms when applied to AOL data. Even though the results are practically promising, the paper provides no performance guarantee, nor does it offer insights on how to model traffic in practice.

Further, Esfandiari et al. (2015) also considered a hybrid arrival model where the input comprises known stochastic i.i.d. demand and an unknown number of arrivals that are chosen by an adversary (which is motivated by traffic spikes). They do not assume any knowledge of the traffic spikes, but the performance guarantee of their algorithm is parameterized by λ , roughly the fraction of the revenue in the optimal solution that is obtained from the stochastic (predictable) part of the demand. Parameter λ plays a similar role as parameter p in our model, in that it controls the adversary’s power. However, the underlying arrival processes in these two models differ considerably and cannot be directly compared. In particular, we do not assume any prior knowledge of the stochastic component; instead we partially predict it. However, we do assume that the adversary determines only the initial order of arrivals (i.e., before knowing which customer will eventually follow its order).

Our work is also closely related to the literature on the secretary problem. In the original formulation of the problem, n secretaries with unique values arrive in uniformly random order; the goal is to maximize the probability of hiring the most valuable secretary. The optimal solution to this problem is an observation-selection policy: observe the first n/e secretaries, then select the first one whose value exceeds that of the best of the previously observed secretaries (Lindley 1961, Dynkin 1963, Freeman 1983, Ferguson 1989). Recently, Kesselheim et al. (2015) relaxed the assumption of uniformly random order, and analyzed the performance of the above policy under certain classes

of nonuniform distribution over permutations. Here, we study the secretary problem in our new arrival model (i.e., only a p fraction of secretaries arrive in uniformly random order) and show that a deterministic observation period is not optimal.

3. Model and Preliminaries

A firm is endowed with b (identical) units of a product to sell over $n \geq 3$ periods, where $n \geq b$. In each period, at most one customer arrives demanding one unit of the product; customers belong to two types depending on the revenue they generate. Type-1 and type-2 customers generate revenue of 1 and $0 < a < 1$, respectively. Upon the arrival of a customer, the firm observes the type of the customer and must make an irrevocable decision to accept this customer and allocate one unit, or to reject this customer. If a firm accepts a type-1 (type-2) customer, it will earn \$1 (\$ a). Our goal is to devise online allocation algorithms that maximize the firm's revenue. We evaluate the performance of an algorithm by comparing it to the optimum offline solution (i.e., the clairvoyant solution).

Before proceeding with the model, we introduce a few notations and briefly discuss the structure of the problem. We represent each customer by the value of revenue she generates if accepted, and the sequence of arrival by $\vec{v} = (v_1, v_2, \dots, v_n)$, where $v_i \in \{0, a, 1\}$; $v_i = 0$ implies that no customer arrives at period i . We denote the number of type-1 (type-2) customers in the entire sequence as n_1 (n_2). Note that the optimum offline solution that we denote by $OPT(\vec{v})$ has the following simple structure: accept all the type-1 customers, and if $n_1 < b$, then accept $\min\{n_2, b - n_1\}$ type-2 customers. Therefore,

$$OPT(\vec{v}) = \min\{b, n_1\} + a \min\{n_2, (b - n_1)^+\}, \quad (1)$$

where $(x)^+ \triangleq \max\{x, 0\}$, and we use the symbol “ \triangleq ” for definitions. At each period, a reasonable online algorithm will accept an arriving type-1 customer if there is inventory left. Thus the main challenge for an online algorithm is to decide whether to accept/reject an arriving type-2 customer facing the following natural trade-off: accepting a type-2 customer may result in rejecting a potential future type-1 customer due to limited inventory; on the other hand, rejecting a type-2 customer may lead to having unused inventory at the end. Therefore, any good online algorithm needs to strike a balance between accepting *too few* and *too many* type-2 customers. We denote by $ALG(\vec{v})$ the revenue obtained by an online algorithm.

Next we introduce our *partially predictable* demand arrival model that works as follows. The adversary determines an initial sequence which we denote by $\vec{v}_I = (v_{I,1}, v_{I,2}, \dots, v_{I,n})$, where $v_{I,j} \in \{0, a, 1\}$, for $1 \leq j \leq n$. However, a subset of customers will not follow this order. We call this subset the *stochastic* group, which we denote by \mathcal{S} . Each customer joins the stochastic group independently




Figure 1 Illustration of the customer arrival model

and with the same probability p . Other customers are in the *adversarial* group denoted by \mathcal{A} . Customers in the stochastic group are permuted uniformly at random among themselves. Formally, a permutation $\sigma_{\mathcal{S}} : \mathcal{S} \rightarrow \mathcal{S}$ is chosen uniformly at random and determines the order of arrivals among the stochastic group. In the resulting overall arriving sequence, the adversarial group follows the adversarial sequence according to \vec{v}_I , and those in the stochastic group follow the random order given by $\sigma_{\mathcal{S}}$. Given \vec{v}_I , we denote the *random* customer arrival sequence by $\vec{V} = (V_1, V_2, \dots, V_n)$, and the *realization* of it by $\vec{v} = (v_1, v_2, \dots, v_n)$.

The example presented in Figure 1 illustrates the arrival process. The top row (gray nodes) shows the initial sequence (\vec{v}_I). The middle row shows which customers belong to the stochastic group (the black nodes) and which belong to the adversarial group (the white ones). The bottom row shows both the permutation $\sigma_{\mathcal{S}}$ and the actual arrival sequence. In this example, $\mathcal{S} = \{2, 5, 6, 8\}$, and $\sigma_{\mathcal{S}}(2) = 6, \sigma_{\mathcal{S}}(5) = 2, \sigma_{\mathcal{S}}(6) = 5$, and $\sigma_{\mathcal{S}}(8) = 8$.

Note that the extreme cases $p = 0$ and $p = 1$ correspond to the adversarial and random order models that have been studied before (e.g., Ball and Queyranne (2009) and Agrawal et al. (2014), respectively). Hereafter, we assume that $0 < p < 1$. For a given $p \in (0, 1)$, at any time over the horizon, we can use the number of past observed type-1 (type-2) customers to obtain bounds on the number of customers of each type to be expected over the rest of the horizon. This idea is formalized later in Subsection 3.2 along with further analysis of our model.

Having described the arrival process, we now define the competitive ratio of an online algorithm under the proposed partially predictable model as follows:

DEFINITION 1. An online algorithm is c -competitive in the proposed partially predictable model if for any adversarial instance \vec{v}_I ,

$$\mathbb{E} \left[\text{ALG}(\vec{V}) \right] \geq c \text{OPT}(\vec{v}_I),$$

where the expectation is taken over which customers belong to the stochastic group (i.e., subset \mathcal{S}), the choice of the random permutation $\sigma_{\mathcal{S}}$, and any possibly randomized decisions of the online algorithm.

Note that $OPT(\vec{V}) = OPT(\vec{v}_I)$ and thus, in the above definition, $\mathbb{E} [ALG(\vec{V})] \geq cOPT(\vec{v}_I)$ is equivalent to $\mathbb{E} [ALG(\vec{V})] \geq c\mathbb{E} [OPT(\vec{V})]$.

In Sections 4 and 5, we present two online algorithms that perform well in the proposed partially predictable model for various ranges of b and n . Before introducing our online algorithms, in the following subsections we introduce a series of notations used throughout the paper and state a consequential concentration result that will allow us to partially predict future demand using past observed data.

3.1. Notational Conventions

Throughout the paper, we use uppercase letters for random variables and lowercase ones for realizations. We have already used this convention in defining \vec{V} vs. \vec{v} . We normalize the time horizon to 1, and represent time steps by $\lambda = 1/n, 2/n, \dots, 1$. First, we introduce notations related to the random customer arrival sequence \vec{V} . At any time step λ , for $j = 1, 2$, the number of type- j customers *to be observed* by the online algorithms up to time λ is denoted by $O_j(\lambda)$. Further, we denote by $O_j^S(\lambda)$ the number of type- j customers in the stochastic group that arrive up to time λ in \vec{V} . Note that the online algorithm cannot distinguish between customers in the stochastic group and customers in the adversarial group. Therefore, the online algorithm does *not* observe the realizations of $O_j^S(\lambda)$. We denote realizations of $O_j(\lambda)$ and $O_j^S(\lambda)$ by $o_j(\lambda)$ and $o_j^S(\lambda)$, respectively.

Next, we introduce notations related to the initial adversarial sequence \vec{v}_I . As discussed earlier, we denote the total number of type- j customers in \vec{v}_I by n_j . In addition, given the sequence \vec{v}_I , we denote the total number of type- j customers among the first λn customers by $\eta_j(\lambda)$. Note that both n_j and $\eta_j(\lambda)$ are deterministic. Also, we define $\tilde{o}_j(\lambda) \triangleq (1-p)\eta_j(\lambda) + p\lambda n_j$ and $\tilde{o}_j^S(\lambda) \triangleq p\lambda n_j$ which will serve as deterministic approximations for $O_j(\lambda)$ and $O_j^S(\lambda)$, respectively (see Lemma 1 and the subsequent discussion for motivation of this definition).

Here we return to the example in Figure 1 and review the notations. Suppose $\lambda = 5/8$ and $p = 0.5$; in this example, looking at the bottom row that shows the sequence \vec{v} , we have: $o_1(5/8) = 3$, $o_1^S(5/8) = 1$, which are realizations of random variables $O_1(5/8)$ and $O_1^S(5/8)$, respectively. Looking at the top row that shows sequence \vec{v}_I , we have: $n_1 = 4$, $\eta_1(5/8) = 3$, $\tilde{o}_1(5/8) = 0.5 \times 3 + 0.5 \times 4 \times (5/8) = 2.75$, and $\tilde{o}_1^S(5/8) = 0.5 \times 4 \times (5/8) = 1.25$ that are all deterministic quantities. Similarly, for type-2 customers, $o_2(5/8) = 2$, $o_2^S(5/8) = 1$, $n_2 = 2$, $\eta_2(5/8) = 1$, $\tilde{o}_2(5/8) = 0.5 \times 1 + 0.5 \times 2 \times (5/8) = 1.125$, and $\tilde{o}_2^S(5/8) = 0.5 \times 2 \times (5/8) = 0.625$.

For convenience of reference, in Table 1 we present a summary of the defined notations.

Finally, to avoid carrying cumbersome expressions in the statement of our results for second-order quantities (e.g., bounds on approximation errors), we use the following approximation notations.

Table 1 Notations

\vec{v}_I	$\vec{v}_I = (v_{I,1}, v_{I,2}, \dots, v_{I,n})$, initial customer sequence
\mathcal{S}	subset of customers in the stochastic group
\mathcal{A}	subset of customers in the adversarial group
\vec{V}	$\vec{V} = (V_1, V_2, \dots, V_n)$, random customer arrival sequence
\vec{v}	$\vec{v} = (v_1, v_2, \dots, v_n)$, a realization of \vec{V} (what online algorithm actually observes)
n_j	number of type- j , $j = 1, 2$, customers in \vec{v}_I (which is the same as in \vec{V})
λ	normalized time: $\lambda = 1/n, \dots, 1$
$O_j(\lambda)$	random number of type- j customers arriving up to time λ
$o_j(\lambda)$	a realization of $O_j(\lambda)$
$O_j^{\mathcal{S}}(\lambda)$	random number of type- j customers in \mathcal{S} arriving up to time λ
$o_j^{\mathcal{S}}(\lambda)$	a realization of $O_j^{\mathcal{S}}(\lambda)$
$\eta_j(\lambda)$	number of type- j , $j = 1, 2$, customers among the first λn ones in \vec{v}_I
$\tilde{o}_j(\lambda)$	$(1-p)\eta_j(\lambda) + p\lambda n_j$ (a deterministic approximation of $O_j(\lambda)$)
$\tilde{o}_j^{\mathcal{S}}(\lambda)$	$p\lambda n_j$ (a deterministic approximation of $O_j^{\mathcal{S}}(\lambda)$)

DEFINITION 2. Suppose $f, g : \mathcal{X} \rightarrow \mathbb{R}$ are two functions defined on set \mathcal{X} . We use the notation $f = O(g)$ if there exists a constant k such that $f(x) < kg(x)$ for all $x \in \mathcal{X}$.

DEFINITION 3. Suppose $f, g : \mathbb{N} \rightarrow \mathbb{R}$ are two functions defined on natural numbers. We use the notation $f = o(g)$ if $\lim_{n \rightarrow \infty} \frac{f(n)}{g(n)} = 0$, and the notation $f = \omega(g)$ if $\lim_{n \rightarrow \infty} \left| \frac{f(n)}{g(n)} \right| = \infty$.

3.2. Estimating Future Demand

At time $\lambda < 1$, upon observing $o_j(\lambda)$, $j = 1, 2$ (but not n_j and $\eta_j(\lambda)$), we wish to estimate future demand, or equivalently the total demand n_j . To make such an estimation, we establish the following concentration result:

LEMMA 1. Define constants $\alpha \triangleq 10 + 2\sqrt{6}$, $\bar{\epsilon} \triangleq 1/24$, and $k \triangleq 16$. For any $\epsilon \in [\frac{1}{n}, \bar{\epsilon}]$, with probability at least $1 - \epsilon$, all the following statements hold:

- If $n_1 \geq \frac{k}{p^2} \log n$, then for all $\lambda \in \{0, 1/n, 2/n, \dots, 1\}$,

$$|O_1(\lambda) - \tilde{o}_1(\lambda)| < \alpha \sqrt{n_1 \log n}, \text{ and} \quad (2a)$$

$$|O_1(\lambda) + O_2(\lambda) - (\tilde{o}_1(\lambda) + \tilde{o}_2(\lambda))| < \alpha \sqrt{(n_1 + n_2) \log n} \quad (2b)$$

- If $n_2 \geq \frac{k}{p^2} \log n$, then for all $\lambda \in \{0, 1/n, 2/n, \dots, 1\}$,

$$|O_2(\lambda) - \tilde{o}_2(\lambda)| < \alpha \sqrt{n_2 \log n}, \text{ and} \quad (3a)$$

$$|O_2^{\mathcal{S}}(\lambda) - \tilde{o}_2^{\mathcal{S}}(\lambda)| < \alpha \sqrt{n_2 \log n}. \quad (3b)$$

The lemma is proved in Appendix EC.1. Given that there are two layers of randomization (selection of subset \mathcal{S} and the random permutation), proving the above concentration results requires a fairly delicate analysis that builds upon several existing concentration bounds. Because proving

concentration results is not the main focus of our work, we will not outline the proof in the main text, and refer the interested reader to Appendix EC.1.⁵ Here we focus on the following two questions: (i) what is our motivation for using deterministic approximations $\tilde{o}_j(\lambda)$ and $\tilde{o}_j^S(\lambda)$? and (ii) how do such approximations help us to estimate n_j ?

To answer the first question, let us count the number of type- j customers in $O_j(\lambda)$ that belong to the stochastic and adversarial groups separately. We start with the stochastic group. Roughly, a total of pn_j type- j customers belong to the stochastic group, and a λ fraction of them arrive by time λ , because these customers are spread almost uniformly over the entire time horizon. As a result, there are approximately $pn_j\lambda$ type- j customers from \mathcal{S} arriving up to time λ . Now we move on to the adversarial group: there are a total of $\eta_j(\lambda)$ of type- j customers in the first λn customers in \vec{v}_I . Since with probability $1 - p$ each of them will be in the adversarial group, the total number of type- j customers from the adversarial group arriving up to time λ is approximately $(1 - p)\eta_j(\lambda)$. Combining these two approximate counting arguments gives us:

$$O_j(\lambda) \approx (1 - p)\eta_j(\lambda) + p\lambda n_j = \tilde{o}_j(\lambda). \quad (4)$$

A similar argument shows that $O_j^S(\lambda) \approx p\lambda n_j = \tilde{o}_j^S(\lambda)$. Lemma 1 confirms that these approximations hold with high probability. Lemma 1 also it provides upper bounds on the corresponding approximation errors. Further, we note that $\tilde{o}_j(\lambda) \neq \mathbb{E}[O_j(\lambda)]$, as shown in Appendix EC.1.1. However, the difference between the two is very small and vanishing in n . Given that $\tilde{o}_j(\lambda)$ provides a very intuitive deterministic approximation for random variable $O_j(\lambda)$ and admits a simple closed-form expression, we use it instead of the $\mathbb{E}[O_j(\lambda)]$.

Now, let us answer the second posed question. There are simple relations between n_j and $\eta_j(\lambda)$ such as $n_j \geq \eta_j(\lambda)$ and $\eta_j(\lambda) + (1 - \lambda)n \geq n_j$.⁶ Combining these with our deterministic approximations leads us to compute upper bounds on the total number of customers as established in Lemma 9.

Finally, based on Lemma 1, we partition the sample space of arriving sequences into two subsets, \mathcal{E} and its complement $\bar{\mathcal{E}}$, and define event \mathcal{E} as follows:

DEFINITION 4. Event \mathcal{E} occurs if the realized arrival sequence \vec{v} satisfies all the conditions of Lemma 1, i.e.,

⁵ We present the values of the constants, defined in the statement of the lemma, only to clarify that they exist and do not depend on n ; however, they are not optimized.

⁶ The first inequality follows from definition. The second one also follows from definition and from the observation that the number of type- j customers arriving between λ and 1 cannot be more than the number of remaining time steps, i.e., $(1 - \lambda)n$.

- If $n_1 \geq \frac{k}{p^2} \log n$, then for all $\lambda \in \{0, 1/n, 2/n, \dots, 1\}$,

$$|o_1(\lambda) - \tilde{o}_1(\lambda)| < \alpha \sqrt{n_1 \log n} \quad \text{and} \quad |o_1(\lambda) + o_2(\lambda) - (\tilde{o}_1(\lambda) + \tilde{o}_2(\lambda))| < \alpha \sqrt{(n_1 + n_2) \log n},$$

- If $n_2 \geq \frac{k}{p^2} \log n$, then for all $\lambda \in \{0, 1/n, 2/n, \dots, 1\}$,

$$|o_2(\lambda) - \tilde{o}_2(\lambda)| < \alpha \sqrt{n_2 \log n} \quad \text{and} \quad |o_2^S(\lambda) - \tilde{o}_2^S(\lambda)| < \alpha \sqrt{n_2 \log n}.$$

Lemma 1 confirms that event \mathcal{E} occurs *with high probability*. In all our analyses, we use the above definition to focus on the event that the deterministic approximations (i.e., $\tilde{o}_j(\lambda)$) are in fact “very close” to the observed sequence. This greatly helps us simplify the analysis and its presentation.

4. A Non-Adaptive Algorithm

In this section, we present and analyze our first online algorithm for the resource allocation problem and the demand model described in Section 3. First, in Section 4.1, we describe the algorithm. Then, in Section 4.2, we present the analysis of its competitive ratio.

4.1. The Algorithm

Our first algorithm is a non-adaptive online algorithm that uses predetermined dynamic thresholds to accept or reject customers. This algorithm combines some ideas from the primal algorithm of Kesselheim et al. (2014) and the threshold algorithm of Ball and Queyranne (2009) to generate maximal revenue from both the stochastic and adversarial components of the demand.

In particular, our non-adaptive algorithm makes use of the fact that customers from the stochastic group are uniformly spread over the entire horizon. Therefore, at least a fraction p of the inventory should be allocated at a roughly constant rate. To this end, we define an *evolving threshold* that works as follows: at any time λ , accept a type-2 customer if the total number of accepted customers by this rule does not exceed $\lfloor \lambda p b \rfloor$.

However, the arrival pattern of the other $1 - p$ fraction can take any arbitrary form. In particular, if the adversary puts many type-2 customers at the very beginning of the time horizon but none toward the end, then we may reject too many type-2 customers early on. To prevent this loss, we keep another quota for a type-2 customer rejected by the evolving threshold. We only reject that customer if the number of such type-2 customers accepted so far exceeds the *fixed* threshold of $\theta \triangleq \frac{1-p}{2-a}$. When $p = 0$, this is the same threshold as in Ball and Queyranne (2009).

The formal definition of our algorithm is presented in Algorithm 1. Note that q_1 , $q_{2,e}$, and $q_{2,f}$ respectively represent counters for the number of accepted type-1 customers, the number of type-2 customers accepted by the evolving threshold, and the number of type-2 customers accepted by the fixed threshold.

Algorithm 1 Online Non-adaptive Algorithm (ALG_1)

1. Initialize $q_1, q_{2,e}, q_{2,f} \leftarrow 0$, and define $\theta \triangleq \frac{1-p}{2-a}$.
2. Repeat for time $\lambda = 1/n, 2/n, \dots, 1$, accept customer $i = \lambda n$ arriving at time λ if there is remaining inventory and one of the following conditions holds:
 - (a) $v_i = 1$; update $q_1 \leftarrow q_1 + 1$.
 - (b) **Evolving threshold rule:** $v_i = a$ and $q_1 + q_{2,e} < \lfloor \lambda p b \rfloor$; update $q_{2,e} \leftarrow q_{2,e} + 1$.
 - (c) **Fixed threshold rule:** $v_i = a$ and $q_{2,f} < \lfloor \theta b \rfloor$; update $q_{2,f} \leftarrow q_{2,f} + 1$.

We prioritize the evolving threshold rule if both of the last two conditions are satisfied.

4.2. Competitive Analysis

In this subsection, we analyze the competitive ratio of Algorithm 1. Our main result is the following theorem:

THEOREM 1. *For $p \in (0, 1)$, the competitive ratio of Algorithm 1 is at least $p + \frac{1-p}{2-a} - O\left(\frac{1}{a(1-p)^p} \sqrt{\frac{\log n}{b}}\right)$ in the partially predictable model.*

Before proceeding to the proof of the above theorem, we make the following remarks:

REMARK 1. Our competitive analysis of Algorithm 1 is tight (up to an $O\left(\sqrt{\frac{\log n}{b}}\right)$ term). In particular, for the following instance, Algorithm 1 can attain only a $p + \frac{1-p}{2-a}$ fraction of the optimum offline solution: Suppose $b = n$ and all customers are of type-2. The revenue of the optimum offline algorithm is ab . On the other hand, if we employ Algorithm 1, at the end we will have $q_1 = 0$, $q_{2,e} \leq pb$ and $q_{2,f} \leq \theta b$. This results in a competitive ratio of at most $p + \theta = p + \frac{1-p}{2-a}$.

REMARK 2. In Subsection 6.1, we prove that no online algorithm can have a competitive ratio larger than $p + \frac{1-p}{2-a} + o(1)$ when $b = o(\sqrt{n})$. On the other hand, Theorem 1 indicates that Algorithm 1 achieves a competitive ratio of $p + \frac{1-p}{2-a} - o(1)$ when $b = \omega(\log n)$. Combining the two results implies that for fixed a and p , Algorithm 1 achieves the best possible competitive ratio (up to an $o(1)$ term) in the regime where conditions $b = \omega(\log n)$ and $b = o(\sqrt{n})$ hold simultaneously.

REMARK 3. Note that even though $p + \frac{1-p}{2-a}$ is the convex combination of the competitive ratios of Ball and Queyranne (2009) and of Agrawal et al. (2014), it cannot be achieved by simply randomizing between these two algorithms. Suppose we flip a biased coin; with probability p , we follow the algorithm of Agrawal et al. (2014) (or any other algorithms designed for a random order model such as Kesselheim et al. (2014)) and with probability $(1-p)$ we follow the fixed threshold algorithm of Ball and Queyranne (2009). In Subsection 6.2 we show that for a certain class of instances, such a randomized algorithm does not generate $p + \frac{1-p}{2-a}$ fraction of the optimum offline solution.

Proof of Theorem 1: We start the proof by making the following observation: Theorem 1 is non-trivial only if $\sqrt{\frac{\log n}{b}}$ is small enough, such that the approximation term $O(\cdot)$ is negligible. Therefore, without loss of generality, we can restrict attention to the case where $\sqrt{\frac{\log n}{b}}$ is small. In particular, recalling that we defined constant $\bar{\epsilon} = 1/24$ in Lemma 1, if $\frac{1}{a(1-p)p} \sqrt{\frac{\log n}{b}} \geq \bar{\epsilon}$, then $O\left(\frac{1}{a(1-p)p} \sqrt{\frac{\log n}{b}}\right)$ becomes $O(1)$ and Theorem 1 becomes trivial. Therefore, without loss of generality, we assume $\frac{1}{a(1-p)p} \sqrt{\frac{\log n}{b}} < \bar{\epsilon}$, or equivalently,

$$b > \frac{1}{\bar{\epsilon}^2} \frac{\log n}{a^2(1-p)^2 p^2}. \quad (5)$$

We denote the random revenue generated by Algorithm 1 by $ALG_1(\vec{V})$. To analyze $\mathbb{E} [ALG_1(\vec{V})]$ we condition it on the event \mathcal{E} . Thus we have:

$$\frac{\mathbb{E} [ALG_1(\vec{V})]}{OPT(\vec{v}_I)} \geq \frac{\mathbb{E} [ALG_1(\vec{V})|\mathcal{E}] \mathbb{P}(\mathcal{E})}{OPT(\vec{v}_I)}.$$

Define $\epsilon \triangleq \frac{1}{a(1-p)p} \sqrt{\frac{\log n}{b}}$. For b that satisfies condition (5), and assuming that $n \geq 3$, we have $\frac{1}{n} \leq \epsilon \leq \bar{\epsilon}$. Therefore, we can apply Lemma 1 to get:

$$\frac{\mathbb{E} [ALG_1(\vec{V})]}{OPT(\vec{v}_I)} \geq \frac{\mathbb{E} [ALG_1(\vec{V})|\mathcal{E}] \mathbb{P}(\mathcal{E})}{OPT(\vec{v}_I)} \geq \frac{\mathbb{E} [ALG_1(\vec{V})|\mathcal{E}]}{OPT(\vec{v}_I)} (1 - \epsilon).$$

This will allow us to focus on the realizations that belong to event \mathcal{E} . In the main part of the proof we show that, for any realization \vec{v} belonging to event \mathcal{E} ,

$$\frac{ALG_1(\vec{v})}{OPT(\vec{v}_I)} \geq p + \frac{1-p}{2-a} - O(\epsilon).$$

Fixing a realization \vec{v} that belongs to event \mathcal{E} , we define $q_1(\lambda)$, $q_{2,e}(\lambda)$, and $q_{2,f}(\lambda)$ to be the values of counters q_1 , $q_{2,e}$, and $q_{2,f}$ right after the algorithm determines whether to accept the customer arriving at time λ . Further, we define $\Delta \triangleq \alpha \sqrt{b \log n}$ (constant α is defined in Lemma 1). To analyze the competitive ratio we analyze three cases separately.

Case (i): $n_1 \geq \frac{k}{p^2} \log n$, and Algorithm 1 exhausts the inventory.

Note that when $n_1 \geq \frac{k}{p^2} \log n$, we can apply the concentration result (2a) from Lemma 1. When Algorithm 1 exhausts the inventory it is possible that the algorithm accepts *too many* type-2 customers, which results in rejecting type-1 customers and losing revenue. We control for this loss

by establishing the following upper bound on the number of type-2 customers accepted by the evolving threshold.⁷ In particular, we have the following lemma:

LEMMA 2. *Under event \mathcal{E} , if $n_1 \geq \frac{k}{p^2} \log n$, then*

$$q_{2,e}(1) \leq p(b - n_1)^+ + \Delta.$$

Proof: We assume, without loss of generality, that $n_1 \leq b$. Otherwise, we construct a modified adversarial instance, denoted by $\vec{v}_{I,M}$, as follows: keep an arbitrary subset of type-1 customers with size b in \vec{v}_I (before the random permutation), and remove the remaining type-1 customers (e.g., set their revenue to be 0). For the same realization of the stochastic group and random permutation, we claim that at any time $\lambda \in \{1/n, \dots, 1\}$, the number of type-2 customers accepted through the evolving threshold rule in the original instance is not larger than that in the modified one. This holds because $o_1(\lambda, \vec{v}) \geq o_1(\lambda, \vec{v}_M)$, where the second argument is added to $o_1(\cdot, \cdot)$ to indicate the corresponding instance. Note that because the algorithm accepts all type-1 customers, this implies $q_1(\lambda, \vec{v}) \geq q_1(\lambda, \vec{v}_M)$, which proves our claim (i.e., $q_{2,e}(\lambda, \vec{v}) \leq q_{2,e}(\lambda, \vec{v}_M)$). Thus, without loss of generality, we assume $n_1 \leq b$. Further, note that because of condition (5), we have $n_1(\vec{v}_M) = b \geq \frac{k}{p^2} \log n$.⁸ Thus we are still in Case (i) for the modified instance.

If no type-2 customer is accepted by the evolving threshold, then $q_{2,e}(1) = 0$ and the proof is complete. Otherwise, let $\bar{\lambda} \leq 1$ be the last time that a type-2 customer is accepted by the evolving rule. Then we have

$$\begin{aligned} q_{2,e}(1) &= q_{2,e}(\bar{\lambda}) \leq \bar{\lambda}pb - o_1(\bar{\lambda}) && \text{(Evolving threshold rule)} \\ &\leq \bar{\lambda}pb - (\bar{\lambda}pn_1 + (1-p)\eta_1(\bar{\lambda}) - \Delta) && \text{((2a))} \\ &\leq p(b - n_1) + \Delta. && (\eta_1(\bar{\lambda}) \geq 0, n_1 \leq b, \text{ and } \bar{\lambda} \leq 1) \end{aligned}$$

The reason for each inequality appears in the same line. We remark that in the second inequality, we crucially use the concentration result of Lemma 1. \square

Using Lemma 2, we prove, in Appendix EC.2, the following lemma that gives a lower bound on the competitive ratio for Case (i):

LEMMA 3. *Under event \mathcal{E} , if $n_1 \geq \frac{k}{p^2} \log n$ and $q_1(1) + q_{2,e}(1) + q_{2,f}(1) = b$, then $\frac{ALG_1(\vec{v})}{OPT(\vec{v})} \geq p + \frac{1-p}{2-a} - \frac{(1-a)\Delta}{ab}$.*

⁷ Note that we already have an upper bound on the number of type-2 customers accepted by the fixed threshold: $q_{2,f}(1) \leq \theta b$.

⁸ This follows from Condition (5) and the fact that $\frac{1}{a^2(1-p)^2} > 1$ and by definition (given in Lemma 1) $\frac{1}{\epsilon^2} \geq k$ which imply $\frac{1}{\epsilon^2} \frac{1}{a^2(1-p)^2} \geq k$.

Case (ii): $n_1 \geq \frac{k}{p^2} \log n$, and **Algorithm 1 does not exhaust the inventory.**

In this case, all type-1 customers are accepted. Therefore, the ratio between $ALG_1(\vec{v})$ and $OPT(\vec{v})$ can be expressed as:

$$\frac{ALG_1(\vec{v})}{OPT(\vec{v})} = \frac{n_1 + a[q_{2,e}(1) + q_{2,f}(1)]}{n_1 + a \min\{n_2, (b - n_1)\}}.$$

The only ‘‘mistake’’ that the algorithm may make is to reject too many type-2 customers. The following lemma establishes a lower bound on the number of accepted type-2 customers:

LEMMA 4. *Under event \mathcal{E} , if $n_1 \geq \frac{k}{p^2} \log n$ and $q_1(1) + q_{2,e}(1) + q_{2,f}(1) < b$, then one of the following conditions holds:*

- (a) $q_{2,e}(1) + q_{2,f}(1) = n_2$,
- (b) $q_{2,f}(1) = \lfloor \theta b \rfloor$ and $n_1 > bp - 3\Delta$, or
- (c) $q_{2,f}(1) = \lfloor \theta b \rfloor$, $n_1 \leq bp - 3\Delta$, and $q_{2,e}(1) \geq (p(n_1 + n_2) - n_1 - 5\Delta)^+$.

Proof: First note that $q_{2,f}(1) < \lfloor \theta b \rfloor$ means that Algorithm 1 never rejects a type-2 customer. This implies that $q_{2,e}(1) + q_{2,f}(1) = n_2$, i.e., condition (a) holds. Now suppose $q_{2,f}(1) = \lfloor \theta b \rfloor$. If $n_1 > bp - 3\Delta$, then condition (b) holds. The most interesting case is when $q_{2,f}(1) = \lfloor \theta b \rfloor$, and $n_1 \leq bp - 3\Delta$. In the following, we show that in this case, condition (c) will hold.

In this case, without loss of generality, we can assume $n_1 + n_2 \leq b$. Otherwise, we construct an alternative adversarial instance, denoted by $\vec{v}_{I,A}$, as follows: keep an arbitrary subset of type-2 customers with size $b - n_1$ in \vec{v}_I (before the random permutation), and remove the remaining type-2 customers (e.g., set their revenue to be 0). With the same realization of the stochastic group and random permutation, we claim that:

$$q_{2,e}(\lambda, \vec{v}) \geq q_{2,e}(\lambda, \vec{v}_A), \quad \lambda \in \{0, 1/n, \dots, 1\}. \quad (6)$$

To show (6), we use induction. The base case, corresponding to taking $\lambda = 0$, is trivial. Suppose (6) holds for $\lambda - 1/n$. We show it will hold for λ as well. At time λ , if $q_{2,e}(\lambda, \vec{v}_A) = q_{2,e}(\lambda - 1/n, \vec{v}_A)$, then (6) holds because $q_{2,e}(\lambda, \vec{v}) \geq q_{2,e}(\lambda - 1/n, \vec{v})$. Otherwise, $q_{2,e}(\lambda, \vec{v}_A) = q_{2,e}(\lambda - 1/n, \vec{v}_A) + 1$. This implies that a type-2 customer arrives at time λ in \vec{v}_A , and thus also in \vec{v} . If $q_{2,e}(\lambda, \vec{v}) = q_{2,e}(\lambda - 1/n, \vec{v}) + 1$, then (6) again holds. Otherwise, under customer arrival sequence \vec{v} , we do not accept the type-2 customer at time λ by the evolving threshold rule, which means that $o_1(\lambda, \vec{v}) + q_{2,e}(\lambda, \vec{v}) = \lfloor \lambda pb \rfloor$. Because $o_1(\lambda, \vec{v}_A) + q_{2,e}(\lambda, \vec{v}_A) \leq \lfloor \lambda pb \rfloor$, and $o_1(\lambda, \vec{v}) = o_1(\lambda, \vec{v}_A)$, we can conclude that (6) holds in the last case as well. This concludes the induction. Thus, without loss of generality, we assume $n_1 + n_2 \leq b$.

To prove that condition (c) holds when $q_{2,f}(1) = \lfloor \theta b \rfloor$ and $n_1 \leq bp - 3\Delta$, we make two important observations: (i) In this case, the number of type-2 customers is large enough to apply the concentration results of (3b). In particular, we have:

$$n_2 \geq \theta b \geq \frac{k \log n}{p^2} \quad (7)$$

where the last inequality holds because of (5), and definitions of $\theta = \frac{1-p}{2-a}$ and k (defined in Lemma 1). (ii) The number of type-1 customers is so small that after a certain time the evolving threshold accepts a sufficient number of type-2 customers that ensures condition (c) holds. In particular, define

$$\bar{\lambda} \triangleq \frac{1}{n} \lceil \frac{n(n_1(1-p) + 3\Delta)}{p(b-n_1)} \rceil.$$

Note that $\bar{\lambda} \leq 1$ when $n_1 \leq bp - 3\Delta$. For any $\lambda \geq \bar{\lambda}$, we have:

$$\begin{aligned} o_1(\lambda) + o_2^S(\lambda) - o_2^S(\bar{\lambda}) &\leq \lambda p n_1 + (1-p)\eta_1(\lambda) + \Delta + \lambda p n_2 + \Delta - (\bar{\lambda} p n_2 - \Delta) && ((2a), (3b)) \\ &\leq \lambda p n_1 + (1-p)n_1 + (\lambda - \bar{\lambda})p n_2 + 3\Delta && (\eta_1(\lambda) \leq n_1) \\ &= \bar{\lambda} p n_1 + (1-p)n_1 + (\lambda - \bar{\lambda})p(n_1 + n_2) + 3\Delta \\ &\leq \bar{\lambda} p n_1 + (1-p)n_1 + (\lambda - \bar{\lambda})p b + 3\Delta && (n_1 + n_2 \leq b) \\ &\leq \lambda p b. && (\text{definition of } \bar{\lambda}) \end{aligned}$$

Note that because $o_1(\lambda) + o_2^S(\lambda) - o_2^S(\bar{\lambda})$ is an integer, the above inequality also implies

$$o_1(\lambda) + o_2^S(\lambda) - o_2^S(\bar{\lambda}) \leq \lfloor \lambda p b \rfloor \quad \text{for all } \lambda \geq \bar{\lambda}. \quad (8)$$

Further, the above inequality implies that for $\lambda \geq \bar{\lambda}$, there is a gap between $o_1(\lambda)$ and the evolving threshold $\lfloor \lambda p b \rfloor$, which in turn implies that the evolving threshold will accept type-2 customers. Next, for $\lambda \geq \bar{\lambda}$, we establish a lower bound on the number of type-2 customers that the evolving threshold accepts. In particular, we show that

$$q_{2,e}(\lambda) \geq o_2^S(\lambda) - o_2^S(\bar{\lambda}) \quad \text{for all } \lambda \geq \bar{\lambda}. \quad (9)$$

We show (9) by induction. The base case $\lambda = \bar{\lambda}$ is trivial. Suppose (9) holds for $\lambda - 1/n \geq \bar{\lambda}$. We show it will also hold for λ : If the arriving customer is not a type-2 customer belonging to the stochastic group, then $o_2^S(\lambda) = o_2^S(\lambda - 1/n)$; but $q_{2,e}(\lambda) \geq q_{2,e}(\lambda - 1/n)$, and thus (9) holds. Otherwise, we have $o_2^S(\lambda) = o_2^S(\lambda - 1/n) + 1$. Now if this customer is accepted by the evolving threshold rule, then both sides of (9) are increased by one, and thus inequality (9) still holds. Otherwise, if the customer is not accepted, it implies we have reached the threshold. Therefore

$$q_{2,e}(\lambda) = \lfloor \lambda pb \rfloor - o_1(\lambda). \quad (10)$$

Now we utilize the gap between $\lfloor \lambda pb \rfloor$ and $o_1(\lambda)$ that we established above in (8). Combining (10) and (8) proves that (9) holds in this case as well. This completes the induction, and thus the proof of (9).

We complete the proof of the lemma by using (9) with $\lambda = 1$, to have the following lower bound:

$$\begin{aligned} q_{2,e}(1) &\geq o_2^S(1) - o_2^S(\bar{\lambda}) && ((9)) \\ &\geq pn_2 - \Delta - (\bar{\lambda}pn_2 + \Delta) && ((3b)) \\ &\geq pn_2 - (n_1(1-p) + 3\Delta) - 2\Delta && (b - n_1 \geq n_2) \\ &= p(n_1 + n_2) - n_1 - 5\Delta. \end{aligned}$$

□

Using Lemma 4, we prove, in Appendix EC.2, the following lemma that gives a lower bound on the competitive ratio for Case (ii):

LEMMA 5. *Under event \mathcal{E} , if $n_1 \geq \frac{k}{p^2} \log n$ and $q_1(1) + q_{2,e}(1) + q_{2,f}(1) < b$, then*

$$\frac{ALG_1(\vec{v})}{OPT(\vec{v})} \geq p + \frac{1-p}{2-a} - \frac{5\Delta}{\theta b}.$$

Case (iii): $n_1 < \frac{k}{p^2} \log n$.

The competitive ratio analysis for Case (iii) is fairly similar to that for Case (ii). It follows from the next two lemmas. The proofs are deferred to Appendix EC.2.

LEMMA 6. *Under event \mathcal{E} , if $n_1 < \frac{k}{p^2} \log n$, then one of the following conditions holds:*

- (a) $q_1(1) + q_{2,e}(1) + q_{2,f}(1) = b$,
- (b) $q_1(1) = n_1$ and $q_{2,e}(1) + q_{2,f}(1) = n_2$, or
- (c) $q_1(1) = n_1$, $q_{2,f}(1) = \lfloor \theta b \rfloor$ and $q_{2,e}(1) \geq pn_2 - \frac{k}{p^2} \log n - 4\Delta$.

Using Lemma 6, the following lemma (proven in Appendix EC.2) gives a lower bound on the competitive ratio for Case(iii):

LEMMA 7. *Under event \mathcal{E} , if $n_1 < \frac{k}{p^2} \log n$, then*

$$\frac{ALG_1(\vec{v})}{OPT(\vec{v})} = \frac{n_1 + a [q_{2,e}(1) + q_{2,f}(1)]}{n_1 + a \min\{n_2, (b - n_1)\}} \geq \min \left\{ p + \frac{1-p}{2-a} - \frac{\frac{k}{p^2} \log n}{ab}, p + \frac{1-p}{2-a} - \frac{\frac{k}{p^2} \log n + 4\Delta}{\theta b} \right\}.$$

Using Lemmas 3, 5, and 7, we have lower bounds on the competitive ratio of Algorithm 1 for all possible cases. We complete the proof of the theorem by the following lemma (proven in Appendix EC.2) that ensures that the error terms in Lemmas 3, 5, and 7 are $O(\epsilon)$.

LEMMA 8. *The error terms in Lemmas 3, 5, and 7 are $O(\epsilon)$, i.e., we have: (a) $\frac{(1-a)\Delta}{ab} = O(\epsilon)$, (b) $\frac{5\Delta}{\theta b} = O(\epsilon)$, (c) $\frac{\frac{k}{p^2} \log n}{ab} = O(\epsilon)$, and (d) $\frac{\frac{k}{p^2} \log n + 4\Delta}{\theta b} = O(\epsilon)$.*

□

5. The Adaptive Algorithm

In the design of Algorithm 1, we used the observation that in the partially predictable model, the demand has a stochastic component that is uniformly spread over the entire horizon. This observation motivated us to define the evolving threshold rule. We remark that in Algorithm 1 neither the evolving threshold rule nor the fixed threshold rule adapts to the observed data, which makes Algorithm 1 a non-adaptive algorithm. As noted in Remark 3, when the initial inventory b is small compared to the horizon n , the competitive ratio of Algorithm 1, $p + \frac{1-p}{2-a}$, is in fact the best possible, and it can be achieved with our non-adaptive algorithm. Therefore, in this regime, adapting to the data, i.e., setting thresholds based on the observed data, would not improve the performance. More precisely, when $b = o(\sqrt{n})$ the inventory is so small compared to the time horizon that there may not be enough time to effectively adapt to the observed data. The adversary can mislead us to allocate all the inventory before we can observe a sufficient portion of the data. However, as b becomes larger, we will have more chance to observe and adapt to the data before allocating a significant part of the inventory. In this section, in fact, we design an adaptive algorithm that achieves a better competitive ratio for large enough b (relative to n). In Section 5.1, we first present the ideas behind our adaptive algorithm along with its formal description. Then, in Section 5.2, we analyze the competitive ratio of our algorithm.

5.1. The Algorithm

In this section, we describe our adaptive algorithm, denoted by $ALG_{2,c}$, which takes $c \in [0, 1]$ as a parameter. For a certain range of c , we show that $ALG_{2,c}$ attains a competitive ratio of c (up to an error term); however, if c becomes too large (for example if $c = 1$), then $ALG_{2,c}$ no longer guarantees a c fraction of the optimum offline solution. We call this algorithm adaptive because it makes decisions based on the sequence of arrivals it has observed so far. In particular, this algorithm repeatedly computes upper bounds on the total number of type-1/-2 customers based on the observed data and uses these upper bounds to decide whether to accept an arriving type-2 customer or not. Before proceeding with the algorithm, we first introduce two functions $u_1(\lambda)$ and

$u_{1,2}(\lambda)$ that will prove useful in constructing the aforementioned upper bounds. In particular we define:

$$u_1(\lambda) \triangleq \begin{cases} b & \text{if } \lambda < \delta \text{ (not enough data observed).} \\ \min \left\{ \frac{o_1(\lambda)}{\lambda p}, \frac{o_1(\lambda) + (1-\lambda)(1-p)n}{1-p+\lambda p} \right\} & \text{if } \lambda \geq \delta. \end{cases}$$

$$u_{1,2}(\lambda) \triangleq \begin{cases} b & \text{if } \lambda < \delta \text{ (not enough data observed).} \\ \min \left\{ \frac{o_1(\lambda) + o_2(\lambda)}{\lambda p}, \frac{o_1(\lambda) + o_2(\lambda) + (1-\lambda)(1-p)n}{1-p+\lambda p} \right\} & \text{if } \lambda \geq \delta, \end{cases}$$

where $\delta \triangleq \frac{(1-c)b}{(1-a)n}$. Note that $u_1(\lambda)$ and $u_{1,2}(\lambda)$ are functions of the observed data $o_1(\lambda)$ and $o_2(\lambda)$. In the following lemma, we show how $u_1(\lambda)$ and $u_{1,2}(\lambda)$ provide upper bounds on n_1 and $n_1 + n_2$ when the realized sequence, \vec{v} , belongs to event \mathcal{E} and the number of type-1 customers as well as the initial inventory b is large enough (as specified in the lemma's statement). Recall that we defined Δ to be $\alpha\sqrt{b \log n}$, where constant α itself is defined in Lemma 1.

LEMMA 9. *Under event \mathcal{E} , suppose $n_1 \geq \frac{k}{p^2} \log n$ and $b > \left(\frac{1}{\bar{\epsilon}} \frac{n\sqrt{\log n}}{(1-c)^2 a p^{3/2}} \right)^{\frac{2}{3}}$, where constants k and $\bar{\epsilon}$ are defined in Lemma 1. Then for all $\lambda \in \{1/n, 2/n, \dots, 1\}$,*

$$u_1(\lambda) \geq \min \left\{ b, n_1 - \frac{2\Delta}{\delta p} \right\}, \text{ and} \quad (11a)$$

$$u_{1,2}(\lambda) \geq \min \left\{ b, n_1 + n_2 - \frac{2\Delta}{\delta p} \right\}. \quad (11b)$$

Lemma 9 is proven in Appendix EC.3.

Having defined $u_1(\lambda)$ and $u_{1,2}(\lambda)$, now we describe how the adaptive algorithm determines whether to accept an arriving type-2 customer when there is remaining inventory. In the following, $q_j(\lambda)$, $j = 1, 2$, represents the number of type- j customers accepted by the algorithm up to time λ (for a particular realization \vec{v}). Suppose the arriving customer at time λ is of type-2. If $u_{1,2}(\lambda) < b$, then we accept the customer, because (11b) implies that the total number of type-1 and type-2 customers will not exceed b (neglecting the error term), and thus we will have extra inventory at the end. On the other hand, if $u_{1,2}(\lambda) \geq b$, we may want to reject this customer to reserve inventory for a future type-1 customer. The decision of whether to accept the customer is based on the following two observations:

Observation 1 *If $u_1(\lambda) \geq n_1$, then*

$$OPT(\vec{v}) \leq \min\{n_1, b\} + a(b - n_1)^+ = (1-a) \min\{n_1, b\} + ab \leq \min\{u_1(\lambda), b\}(1-a) + ab.$$

Observation 2 *If we accept the current type-2 customer, then the maximum revenue we can get is $(b - (q_2(\lambda - 1/n) + 1)) + a(q_2(\lambda - 1/n) + 1)$.*

To have a competitive ratio of at least c , Observations 1 and 2 motivate us to accept the type-2 customer only if

$$\frac{(b - (q_2(\lambda - 1/n) + 1)) + a(q_2(\lambda - 1/n) + 1)}{\min\{u_1(\lambda), b\}(1 - a) + ab} \geq c. \quad (12)$$

After rearranging terms, we get the following threshold for accepting the type-2 customer:

$$q_2(\lambda - 1/n) + 1 \leq \frac{1 - c}{1 - a}b + c(b - u_1(\lambda))^+. \quad (13)$$

Thus when $u_{1,2}(\lambda) \geq b$, we use Condition (13) to accept/reject a type-2 customer. For notational convenience, we define $\phi \triangleq \frac{1-c}{1-a}$. We point out the right-hand side of (13) may not be an integer; thus, in our algorithm, we use a slightly modified version of it, defined as follows:

$$q_2(\lambda - 1/n) \leq \lfloor \frac{1 - c}{1 - a}b + c(b - u_1(\lambda))^+ \rfloor. \quad (14)$$

Note that by the definition of the threshold given in (14), we always accept the first $\lfloor \phi b \rfloor$ type-2 customers. The formal definition of our algorithm is presented in Algorithm 2. In Algorithm 2, q_j represents the counter for the number of accepted customers of type- j so far.

Algorithm 2 Online Adaptive Algorithm ($ALG_{2,c}$)

1. Initialize $q_1, q_2 \leftarrow 0$, and define $\phi \triangleq \frac{1-c}{1-a}$, and $\delta \triangleq \frac{\phi b}{n}$.
2. Repeat for time $\lambda = 1/n, 2/n, \dots, 1$:

(a) Calculate functions $u_1(\lambda)$ and $u_{1,2}(\lambda)$ (to construct upper bounds for n_1 and $n_1 + n_2$):

$$u_1(\lambda) \triangleq \begin{cases} b & \text{if } \lambda < \delta \text{ (not enough data observed).} \\ \min \left\{ \frac{o_1(\lambda)}{\lambda^p}, \frac{o_1(\lambda) + (1-\lambda)(1-p)n}{1-p+\lambda p} \right\} & \text{if } \lambda \geq \delta. \end{cases}$$

$$u_{1,2}(\lambda) \triangleq \begin{cases} b & \text{if } \lambda < \delta \text{ (not enough data observed).} \\ \min \left\{ \frac{o_1(\lambda) + o_2(\lambda)}{\lambda^p}, \frac{o_1(\lambda) + o_2(\lambda) + (1-\lambda)(1-p)n}{1-p+\lambda p} \right\} & \text{if } \lambda \geq \delta. \end{cases}$$

(b) Accept customer $i = \lambda n$ arriving at time λ if there is remaining inventory and one of the following conditions holds:

- i. $v_i = 1$; update $q_1 \leftarrow q_1 + 1$.
- ii. $v_i = a$ and $u_{1,2}(\lambda) < b$; update $q_2 \leftarrow q_2 + 1$.
- iii. $v_i = a$ and $q_2 \leq \lfloor \phi b + c(b - u_1(\lambda))^+ \rfloor$; update $q_2 \leftarrow q_2 + 1$.

We prioritize the second condition if both the second and the third ones hold.

Before we analyze the algorithm, we highlight two key properties of threshold $\lfloor \phi b + c(b - u_1(\lambda))^+ \rfloor$: (i) The threshold is decreasing in $u_1(\lambda)$; the smaller $u_1(\lambda)$ is, the less inventory we reserve for future type-1 customers. (ii) The threshold is decreasing in c as well (the right-hand side

of (14) can be expressed as $\lfloor \frac{1}{1-a}b - c(\frac{b}{1-a} - (b - u_1(\lambda))^+) \rfloor$. When c is too large, we may reject too many type-2 customers, which in turn hurts the revenue in a certain class of instances. Said another way, note that Inequality (14) only gives a “necessary” condition for achieving c -competitiveness. We identify the sufficient condition for c -competitiveness by solving the *factor-revealing* mathematical program presented in (MP1). We will explain the construction of this program in the analysis of the competitive ratio (in Section 5.2). On a high level, we construct the feasible region such that it contains any valid instance that can violate the c -competitiveness; by minimizing over c , we find the smallest value of c for which the feasible region is not empty.

	$\underset{(l, n_1, n_2, \eta_1, \eta_2, c)}{\text{Minimize } c}$	(MP1)						
subject to	$c \geq \frac{a(n_2 - \tilde{o}_2 + \frac{b}{1-a}) + n_1}{a \min\{n_1 + n_2, b\} + (1-a)n_1 + \frac{a^2b}{1-a} + a \min\{\tilde{u}_1, b\}}$	(15a)						
	$\tilde{u}_{1,2} \geq b$	(15b)						
$l \leq 1$	(15c)	$\eta_1 + \eta_2 \leq ln$	(15d)	$\eta_1 \leq n_1$	(15e)	$\eta_2 \leq n_2$	(15f)	
$n_1 \leq b$	(15g)	$n_1 + n_2 \leq n$	(15h)	$n_1 + n_2 \leq \eta_1 + \eta_2 + (1-l)n$				(15i)
where $\tilde{o}_1 \triangleq (1-p)\eta_1 + pn_1l$, $\tilde{o}_2 \triangleq (1-p)\eta_2 + pn_2l$, $\tilde{u}_1 \triangleq \min\left\{\frac{\tilde{o}_1}{lp}, \frac{\tilde{o}_1 + (1-l)(1-p)n}{(1-p+lp)}\right\}$, and $\tilde{u}_{1,2} \triangleq \min\left\{\frac{\tilde{o}_1 + \tilde{o}_2}{lp}, \frac{\tilde{o}_1 + \tilde{o}_2 + (1-l)(1-p)n}{(1-p+lp)}\right\}$.								

Before we analyze Algorithm 2, we also evaluate the solution of (MP1). Denote the optimal objective value of (MP1) by c^* . As will be stated in Theorem 2, ALG_{2,c^*} achieves a competitive ratio of c^* (minus an error term). First, we solve (MP1) numerically for the regime where $b = \kappa n$ (where $0 < \kappa \leq 1$ is a constant), and show that if $b/n > 0.5$, then Algorithm 2 achieves a better competitive ratio than Algorithm 1.

In Figure 2, we fix $a = 0.5, 0.7$, and plot c^* for $p = 0.05, 0.1, \dots, 0.95$ for three cases of $b/n = 0.9, 0.7$, and 0.5 . Figure 2 leads us to make the following observation: The competitive ratio of ALG_{2,c^*} is at least that of ALG_1 , and it is significantly larger when (i) p is small and (ii) b/n is large. This observation highlights the power of adapting to the data, even though it contains an adversarial component: Consider $a = 0.7$, $b = 0.7n$ and $p = 0.2$; this means that 80% of the demand belongs to the adversarial group. Our adaptive algorithm guarantees 10% more revenue than the non-adaptive algorithm does.

In addition, we note that as the initial inventory b becomes larger (for a fixed time horizon n), the adversary’s power naturally declines. Thus one would expect that a “smart” algorithm achieves

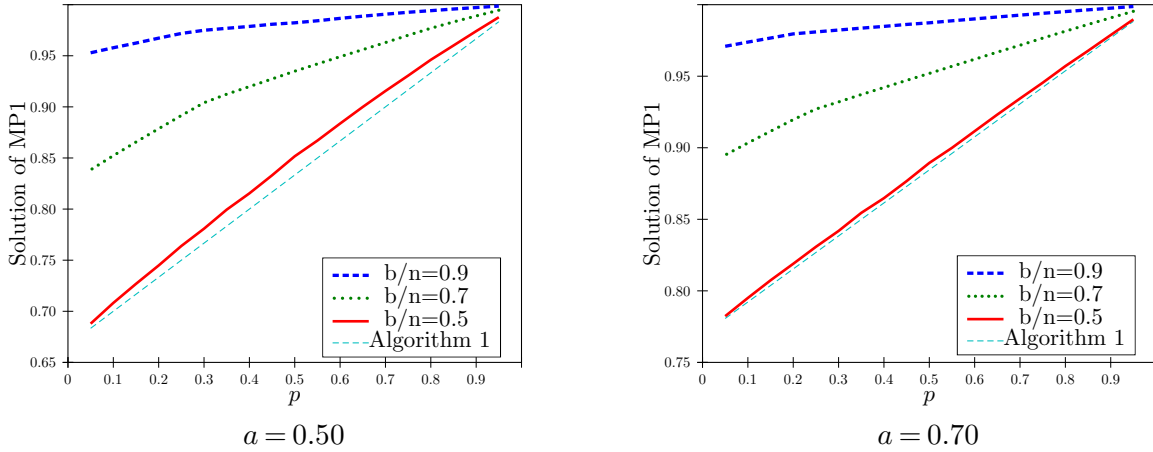


Figure 2 Solution of (MP1), c^* , vs. p for $a = 0.50$ and 0.70

a higher competitive ratio. Our adaptive algorithm indeed attains a higher competitive ratio as the initial inventory increases. In contrast, the competitive ratio of our non-adaptive algorithm remains the same.

We conclude our study of (MP1) by establishing a lower bound on its optimum solution. The following proposition states that c^* is at least $p + \frac{1-p}{2-a}$, which is the competitive ratio of Algorithm 1 (ignoring the error term).

PROPOSITION 1. *For any $b \leq n$, we have: $c^* \geq p + \frac{1-p}{2-a}$. Further, if $b = n$, then $c^* = 1$.*

5.2. Competitive Analysis

In this section we analyze the competitive ratio of Algorithm 2 and prove the following theorem:

THEOREM 2. *For $p \in (0, 1)$, let c^* be the optimal objective value of (MP1). For any $c \leq c^*$ such that $c < 1$, $ALG_{2,c}$ is $c - O\left(\frac{1}{(1-c)^2 ap^{3/2}} \sqrt{\frac{n^2 \log n}{b^3}}\right)$ competitive in the partially predictable model.*

The above theorem implies that if $c^* < 1$, then ALG_{2,c^*} is $c^* - O\left(\frac{1}{(1-c^*)^2 ap^{3/2}} \sqrt{\frac{n^2 \log n}{b^3}}\right)$ competitive. However, the same does not hold when $c^* = 1$. For this special case, we have the following corollary of Theorem 2:

COROLLARY 1. *When $c^* = 1$, for $c = 1 - \sqrt[3]{\frac{1}{ap^{3/2}} \sqrt{\frac{n^2 \log n}{b^3}}}$ the competitive ratio of $ALG_{2,c}$ is $1 - O\left(\sqrt[3]{\frac{1}{ap^{3/2}} \sqrt{\frac{n^2 \log n}{b^3}}}\right)$.*

REMARK 4. Theorem 2 combined with Proposition 1 shows that in the asymptotic regime (where n and b both grow), if the scaling factor $\sqrt{\frac{n^2 \log n}{b^3}}$ (which appears in the error term of the competitive ratio) is vanishing (i.e., order of $o(1)$), then our adaptive algorithm outperforms our non-adaptive one. For instance, the aforementioned condition holds if $b = \kappa n$ where $0 < \kappa \leq 1$ is a constant.

Proof of Theorem 2: Similar to the proof of Theorem 1, we start by making the observation that Theorem 2 is nontrivial only if $\sqrt{\frac{n^2 \log n}{b^3}}$ is small enough such that the approximation term $O(\cdot)$ is negligible. Therefore, without loss of generality, we can restrict attention to the case where $\sqrt{\frac{n^2 \log n}{b^3}}$ is small. In particular, if $\frac{1}{(1-c)^2 ap^{3/2}} \sqrt{\frac{n^2 \log n}{b^3}} \geq \bar{\epsilon}$, then $O\left(\frac{1}{(1-c)^2 ap^{3/2}} \sqrt{\frac{n^2 \log n}{b^3}}\right)$ becomes $O(1)$ and Theorem 2 becomes trivial (recall that constant $\bar{\epsilon} = 1/24$ is defined in Lemma 1). Therefore, without loss of generality, we assume $\frac{1}{(1-c)^2 ap^{3/2}} \sqrt{\frac{n^2 \log n}{b^3}} < \bar{\epsilon}$, or equivalently,

$$b^{\frac{3}{2}} > \frac{1}{\bar{\epsilon}} \frac{n \sqrt{\log n}}{(1-c)^2 ap^{3/2}}. \quad (16)$$

We remark that we impose the same condition on b in Lemma 9. We denote the random revenue generated by Algorithm 2 by $ALG_{2,c}(\vec{V})$. Similar to the proof of Theorem 1, we define an appropriate ϵ that allows us to focus on the realizations that belong to event \mathcal{E} . In particular, let $\epsilon = \frac{1}{(1-c)^2 ap^{3/2}} \sqrt{\frac{n^2 \log n}{b^3}}$. For b that satisfies condition (16), and assuming that $n \geq 3$, we have $\frac{1}{n} \leq \epsilon \leq \bar{\epsilon}$. Therefore, we can apply Lemma 1 to get:

$$\frac{\mathbb{E} \left[ALG_{2,c}(\vec{V}) \right]}{OPT(\vec{v}_I)} \geq \frac{\mathbb{E} \left[ALG_{2,c}(\vec{V}) | \mathcal{E} \right] \mathbb{P}(\mathcal{E})}{OPT(\vec{v}_I)} \geq \frac{\mathbb{E} \left[ALG_{2,c}(\vec{V}) | \mathcal{E} \right]}{OPT(\vec{v}_I)} (1 - \epsilon).$$

In the main part of the proof, we show that for any realization \vec{v} belonging to event \mathcal{E} ,

$$\frac{ALG_{2,c}(\vec{v})}{OPT(\vec{v}_I)} \geq c - O(\epsilon).$$

To analyze the competitive ratio we analyze three cases separately.

Case (i): $n_1 \geq \frac{k}{p^2} \log n$, and Algorithm 2 exhausts the inventory.

When $n_1 \geq \frac{k}{p^2} \log n$, we can apply (2a) from Lemma 1 and Lemma 9. Because Algorithm 2 exhausts the inventory, we know that $n_1 + n_2 \geq b$. Now we have either (a) $n_1 + n_2 - \frac{2\Delta}{\delta p} \leq b$ or (b) $n_1 + n_2 - \frac{2\Delta}{\delta p} > b$. If (a) happens, then (according to Lemma 9) we may have $u_{1,2}(\lambda) < b$, which may result in accepting a type-2 customer through the second condition that we should have rejected. However, in this case, we also have a tight upper bound on the optimum offline solution. As shown in the proof of Lemma 11—which analyzes the competitive ratio of the two cases (a) and (b) separately—such a bound allows us to establish the desired lower bound on the competitive ratio. Case (b) is the more interesting case, which accepts type-2 customers through the third condition of Algorithm 2. It is possible that the algorithm accepts *too many* type-2 customers through this condition, resulting in rejecting type-1 customers, and thus in revenue loss. In the following

lemma, we control for this loss by establishing an upper bound on the number of accepted type-2 customers. The proof of the lemma, which uses similar ideas to those in Lemma 2, is deferred to Appendix EC.3.

LEMMA 10. *Under event \mathcal{E} , if $n_1 \geq \frac{k}{p^2} \log n$, then one of the following conditions holds:*

- (a) $n_1 + n_2 - \frac{2\Delta}{\delta p} \leq b$, or
- (b) $n_1 + n_2 - \frac{2\Delta}{\delta p} > b$ and $q_2(1) \leq \frac{1-c}{1-a}b + c(b - n_1)^+ + c\frac{2\Delta}{\delta p} + 1$.

Using Lemma 10 and the discussion before the lemma, in Appendix EC.3, we prove the following lemma, which gives a lower bound on the competitive ratio for Case (i):

LEMMA 11. *Under event \mathcal{E} , if $n_1 \geq \frac{k}{p^2} \log n$ and $q_1(1) + q_2(1) = b$, then*

$$\frac{ALG_{2,c}(\vec{v})}{OPT(\vec{v})} \geq c - \frac{3\Delta}{ab\delta p}.$$

Case (ii): $n_1 \geq \frac{k}{p^2} \log n$, and Algorithm 2 does not exhaust the inventory.

First note that in this case $OPT(\vec{v}) = n_1 + a \min\{b - n_1, n_2\}$. Also, in this case, we accept all type-1 customers. Therefore, $q_1(1) = n_1$. To lower-bound the competitive ratio, we need to show only that we do not reject too many type-2 customers, i.e., $q_2(1)$ is large enough. Note that if for all $\lambda \in \{1/n, 2/n, \dots, 1\}$, condition (14) holds, then all type-2 customers are accepted, and we have $q_2(1) = n_2$. This implies that $ALG_{2,c}(\vec{v}) = OPT(\vec{v})$. The more interesting case is when there exists at least one time step for which condition (14) is violated. Let l be the last time that we reject a type-2 customer. This means that at time l , we have:

$$u_{1,2}(l) \geq b, \quad (17) \quad q_2(l) \geq \frac{1-c}{1-a}b + c(b - u_1(l))^+. \quad (18)$$

This also provides the following lower bound on the number of accepted type-2 customers:

$$q_2(1) = q_2(l) + [n_2 - o_2(l)] \geq \frac{1-c}{1-a}b + c(b - u_1(l))^+ + [n_2 - o_2(\bar{\lambda})]. \quad (19)$$

Therefore, when $q_1(1) + q_2(1) < b$,

$$\frac{ALG_{2,c}(\vec{v})}{OPT(\vec{v})} \geq \frac{n_1 + a \left(\frac{1-c}{1-a}b + c(b - u_1(l))^+ + [n_2 - o_2(l)] \right)}{n_1 + a \min\{b - n_1, n_2\}}. \quad (20)$$

For a fixed c , if for all possible instances the right-hand side of (20) is greater than c , then $ALG_{2,c}$ would be c -competitive. However, if c is too large, then there will be instances for which the right-hand side of (20) will be less than c . We identify a superset of these instances by all possible combinations of $(l, n_1, n_2, \eta_1(l), \eta_2(l))$ that satisfy certain constraints to ensure they correspond to valid instances. As a reminder, $\eta_j(l)$ represents the number of type- j customers by time l in the

initial sequence (determined by the adversary, i.e., \vec{v}_I). As we describe these constraints below, it becomes clear that (1) any instance of the problem would satisfy all these constraints, and (2) these constraints correspond to the feasible region of the mathematical program in (MP1).

We start with the straightforward constraints: for every instance, $n_1 + n_2 \leq n$. Also, $\eta_1(l) \leq n_1$, and $\eta_2(l) \leq n_2$. Further, in the initial customer sequence \vec{v}_I , at time l we cannot have more than ln customers, thus $\eta_1(l) + \eta_2(l) \leq ln$. Similarly after time l , we cannot have more than $(1-l)n$ customers, and therefore $n_1 + n_2 - [\eta_1(l) + \eta_2(l)] \leq (1-l)n$. By definition of l , we have $l \leq 1$. We also add the condition $n_1 \leq b$, which is always true under the case when $q_1(1) + q_2(1) < b$. Note that these are Constraints (15c)-(15i) in (MP1), where in (MP1), with a slight abuse of notation, we simplify by substituting η_j for $\eta_j(l)$.

For a moment, suppose $o_j(l) = \tilde{o}_j(l)$. First, we remind the reader that $\tilde{o}_j(l) = (1-p)\eta_j(l) + pln_j$ is the deterministic approximation of $o_j(l)$ that we introduced in Section 3, and also is redefined in (MP1) (at the bottom). Further, note that this is just to explain the idea behind constructing (MP1). Later in the proof, we address the difference between $\tilde{o}_j(l)$ and $o_j(l)$. In this case, we have:

$$\tilde{u}_{1,2}(l) \triangleq \min \left\{ \frac{\tilde{o}_1(l) + \tilde{o}_2(l)}{lp}, \frac{\tilde{o}_1(l) + \tilde{o}_2(l) + (1-l)(1-p)n}{(1-p+lp)} \right\} = u_{1,2}(l) \geq b \quad (21)$$

where the last inequality is the same as inequality (17). Further note that rejecting a customer at time l implies that $l \geq \frac{\phi b}{n} = \delta$, and thus by definition $u_{1,2}(l) = \min \left\{ \frac{o_1(l) + o_2(l)}{lp}, \frac{o_1(l) + o_2(l) + (1-l)(1-p)n}{1-p+lp} \right\}$.⁹ Note that Inequality (21) is Constraint (15b), where in (MP1), again with a slight abuse of notation, we simplify by substituting $\tilde{u}_{1,2}$ for $\tilde{u}_{1,2}(l)$ and \tilde{o}_j for $\tilde{o}_j(l)$.

Further, the most interesting constraint, Constraint (15a), comes from condition (20). By rearranging terms, we can show that the right-hand side of (20) being smaller or equal to c is equivalent to:

$$c \geq \frac{a(n_2 - o_2(l) + \frac{b}{1-a}) + n_1}{a \min\{n_1 + n_2, b\} + (1-a)n_1 + \frac{a^2 b}{1-a} + a \min\{u_1(l), b\}} \quad (22)$$

which is Constraint (15a) after substituting $o_2(l)$ with \tilde{o}_2 and $u_1(l)$ with \tilde{u}_1 .

Overall, the above conditions define the feasible region of the math program (MP1). By minimizing c , we find the threshold for making (MP1) infeasible: Let c^* be the solution of (MP1); for any $c < c^*$, (MP1) is infeasible, and the only constraint that $(l, n_1, n_2, \eta_1(l), \eta_2(l), c)$ can violate is (15a) (same as (22)). This implies that $ALG_{2,c}$ is c -competitive.

We now go back and address the issue that $\tilde{o}_j(l)$ and $o_j(l)$ are not equal. Due to the difference between $o_j(l)$ and $\tilde{o}_j(l)$, (1) Constraint (15b) might be violated (even though (17) is satisfied)

⁹ Note that when $l < \delta$ Algorithm 2 never rejects a customer, because $q_2(l) \leq ln < \delta n = \phi b$.

and (2) violating Constraint (15a) does not imply violating (22). To address these issues, first in Lemma 12, we give a slightly modified tuple that satisfies Constraints (15b)-(15i); then, in Lemma 13, we prove that for any $c \leq c^*$, if Constraint (15a) is violated, then $\frac{ALG_{2,c}(\vec{v})}{OPT(\vec{v})} \geq c - \frac{4\Delta n}{\phi^2 b^2 p}$. The proofs of both lemmas are deferred to Appendix EC.3, and they amount to applying the concentration results of Lemma 1 and carefully analyzing the error terms. These two lemmas complete the analysis of competitive ratio in Case (ii).

LEMMA 12. *Under event \mathcal{E} , if $n_1 \geq \frac{k}{p^2} \log n$ and $q_1(1) + q_2(1) < b$, then the tuple $(l', n'_1, n'_2, \eta'_1, \eta'_2, c') \triangleq (l, n_1, n_2 + \xi, \eta_1(l), \eta_2(l) + \bar{\xi}, c)$ satisfies Constraints (15b)-(15i), where*

$$\xi \triangleq \begin{cases} 0 & \text{if } n_1 + n_2 \geq b, \\ \min \left\{ n - (n_1 + n_2), \frac{\Delta n}{\phi b p} \right\} & \text{if } n_1 + n_2 < b; \end{cases}$$

$$\bar{\xi} \triangleq \begin{cases} 0 & \text{if } n_1 + n_2 \geq b, \\ \min \{ \xi, l n - (\eta_1(l) + \eta_2(l)) \} & \text{if } n_1 + n_2 < b, \end{cases}$$

and where $\Delta = \alpha \sqrt{b \log n}$, $\phi = \frac{1-c}{1-a}$, and l is the last time that we reject a type-2 customer.

LEMMA 13. *Under event \mathcal{E} , if $n_1 \geq \frac{k}{p^2} \log n$ and $q_1(1) + q_2(1) < b$, then*

$$\frac{ALG_{2,c}(\vec{v})}{OPT(\vec{v})} \geq c - \frac{4\Delta n}{\phi^2 b^2 p}.$$

Case (iii): $n_1 < \frac{k}{p^2} \log n$.

The competitive ratio analysis for this case uses ideas similar to those in the previous two cases, and it follows from the next two lemmas. The proofs are deferred to Appendix EC.3.

LEMMA 14. *Under event \mathcal{E} , if $n_1 < \frac{k}{p^2} \log n$, then one of the following three conditions holds:*

- (a) $q_1(1) + q_2(1) = b$;
- (b) $q_1(1) = n_1$ and $q_2(1) = n_2$; or
- (c) $q_1(1) = n_1$ and $q_2(1) \geq cb$.

Using Lemma 14, in the following lemma, we establish a lower bound on the competitive ratio for Case (iii):

LEMMA 15. *Under event \mathcal{E} , if $n_1 < \frac{k}{p^2} \log n$, then $\frac{ALG_{2,c}(\vec{v})}{OPT(\vec{v})} \geq c$.*

Having Lemmas 11, 13, and 15, we have lower bounds on the competitive ratio of Algorithm 2 for all possible cases. We complete the proof of the theorem by the following lemma (proven in Appendix EC.3) that ensures that the error terms in Lemmas 11 and 13 are $O(\epsilon)$.

LEMMA 16. *The error terms in Lemmas 11 and 13 are $O(\epsilon)$, i.e., (a) $\frac{3\Delta}{ab\delta p} = O(\epsilon)$ and (b) $\frac{4\Delta n}{\phi^2 b^2 p} = O(\epsilon)$.*

□




Figure 3 Two problem instances between which online algorithms cannot distinguish at time $\frac{b}{n}$, where $b = 4$ and $n = 8$.

6. Discussion of the Model

In this section, we further study the performance of online algorithms in our demand model. First, in Section 6.1, we present an upper bound on the competitive ratio achievable by any online algorithm under our demand model when the initial inventory b is small—more precisely, $b = o(\sqrt{n})$. Next, in Section 6.2, we highlight the need for our new online algorithms by presenting a problem instance for which our algorithms outperform existing ones in our partially predictable model.

6.1. Upper Bounds

In this section, we present an upper bound on the competitive ratio of any online algorithm when $b = o(\sqrt{n})$. We start with a warm-up example that illustrates a fundamental limit of any online algorithm in the partially predictable model. Figure 3 shows two instances with $n = 8$. The bottom row shows the sequence that the online algorithm will see; as a reminder, we represent the nodes of the stochastic group as filled (even though the online algorithm cannot distinguish between the two groups of customers). Suppose $b = 4$; in the instance presented on the left, the optimum offline solution rejects all type-2 customers, and in the instance on the right, it accepts all of them. Now, by time $\lambda = b/n = 4/8$, online algorithms cannot distinguish between these two instances, and hence cannot perform as well as the optimal offline algorithm on *both* of these instances. Similar to this example, in the following proposition, we establish the upper bound by constructing two problem instances that are “difficult” for online algorithms to distinguish between up to time $\frac{b}{n}$, and show that the trade-off between accepting too many or too few type-2 customers limits the competitive ratio of any online algorithm.

PROPOSITION 2. *Under the partially predictable arrival model, and for any $p \in (0, 1)$, no online algorithm, deterministic and randomized, can achieve a competitive ratio better than $\frac{1-p}{2-a} + p + O\left(\frac{pb^2}{n}\right)$. Therefore, when $b = o(\sqrt{n})$, no online algorithm can achieve a competitive ratio better than $\frac{1-p}{2-a} + p + o(1)$.*

The details of the proof are deferred to Appendix EC.4. As explained above, the main idea of the proof is to construct two instances that are almost indistinguishable up to time $\frac{b}{n}$ to any online algorithm. In the proof we show that the following two instances \vec{v}_I and \vec{w}_I serve our purpose:

$$v_{I,j} = \begin{cases} a, & 1 \leq j \leq b, \\ 0, & b < j \leq 2b, \\ 0, & j > 2b. \end{cases} \quad w_{I,j} = \begin{cases} a, & 1 \leq j \leq b, \\ 1, & b < j \leq 2b, \\ 0, & j > 2b. \end{cases}$$

6.2. Comparison with Existing Algorithms

In this section, we show that, under our demand arrival model, there exists a class of instances for which our algorithms achieve higher revenue than algorithms designed for either the worst-case (Ball and Queyranne 2009) or the random-order model (Devanur and Hayes 2009, Agrawal et al. 2014), which respectively correspond to $p = 0$, and $p = 1$ in our model. To this end, we consider instance \vec{v}_I where

$$v_{I,j} = \begin{cases} a & \text{for } 1 \leq j \leq b, \\ 0 & \text{for } j > b. \end{cases}$$

Algorithm	Worst-Case (Ball and Queyranne (2009))	Random-Order (Idea of Agrawal et al. (2014))	Algorithm 1 (Non-Adaptive Algorithm)	Algorithm 2 (Adaptive Algorithm)
Ratio	$\frac{1}{2-a}$	at most $p + \frac{b}{n}(1-p)$	$p + \frac{1-p}{2-a} - O\left(\frac{1}{a(1-p)^p} \sqrt{\frac{\log n}{b}}\right)$	1

Table 2 Ratio between the expected revenue of different algorithms and the optimum offline solution.

Table 2 presents the ratio between the expected revenue of different online algorithms and that of the optimum offline solution. In the following, we will explain how we compute these bounds. Before that, we discuss the implications of this example. This instance class shows that, for any $p \in (0, 1)$, when $b = \omega(\sqrt{\log n})$ and $b = o(n)$ the ratio for both of our algorithms is better than existing ones. Further, note that the ratio for Algorithm 1 is in fact its competitive ratio; thus the same ratio holds for any other instance as well. This implies that the competitive ratio of our non-adaptive algorithm is higher than those of Ball and Queyranne (2009) and Agrawal et al. (2014) under the partially predictable model. Also note that for the same instance, randomizing between the algorithm of Ball and Queyranne (2009) (with probability $1-p$) and that of Agrawal et al. (2014) (with probability p) leads to a ratio of $\frac{1-p}{2-a} + p^2 + o(1)$, which is not the convex combination of the competitive ratios of these two algorithms (as also pointed out in Remark 3).

Next, we calculate the ratios listed in Table 2. The offline solution is $OPT(\vec{v}_I) = ab$. The algorithm of Ball and Queyranne (2009), proposed for the adversarial model, has a fixed threshold of $\frac{1}{2-a}b$ for accepting type-2 customers, and hence accepts $\frac{1}{2-a}b$ type-2 customers.

Next we compute the ratio for algorithms designed for the random-order model (e.g., Devanur and Hayes (2009), Agrawal et al. (2014), and Kesselheim et al. (2014)). We note that, for the sake of brevity, we present an analysis based on the idea of these papers, which is allocating

inventory at a roughly uniform rate over the entire horizon. In particular, these algorithms accept roughly λb customers at any time $\lambda \in [0, 1]$. As a result, for this instance, they accept at most b^2/n type-2 customers up to time $\lambda = b/n$. According to our model, in the arriving instance \vec{v} , there are approximately $(1 - b/n)bp$ type-2 customers arriving after time b/n . Therefore, these algorithms can accept at most $b^2/n + (1 - b/n)bp$ type-2 customers, which corresponds to a ratio of at most $p + \frac{b}{n}(1 - p)$. Note that $p + \frac{b}{n}(1 - p) < p + \frac{1-p}{2-a}$ for any $b < \frac{n}{2-a}$.

Our Algorithm 1 achieves a ratio of at least its competitive ratio as given in Theorem 1 and the ratio is tight for this instance (up to an additive error term of $O\left(\frac{1}{a(1-p)p} \sqrt{\frac{\log n}{b}}\right)$). For Algorithm 2, let $c \in (0, 1)$ be an arbitrary constant. We show that $ALG_{2,c}$ achieves the ratio of 1 because the third condition in Algorithm 2, i.e., the dynamic threshold, is never violated. To see this we compute the threshold as follows:

$$\lfloor \phi b + c(b - u_1(\lambda))^+ \rfloor = \begin{cases} \lfloor \phi b \rfloor & \lambda < \delta = \frac{\phi b}{n} \\ \lfloor \phi b + cb \rfloor & \lambda \geq \delta \end{cases}$$

where we use the fact that $u_1(\lambda) = b$ for $\lambda < \delta$, and $u_1(\lambda) = 0$ for $\lambda \geq \delta$. In both cases we have $\lfloor \phi b + c(b - u_1(\lambda))^+ \rfloor > \lambda$, which implies that the algorithm never rejects a type-2 customer because $o_2(\lambda) \leq \lambda < \lfloor \phi b + c(b - u_1(\lambda))^+ \rfloor$.

7. The Secretary Problem under Partially Predictable Demand

In this section, we study the online secretary problem under our new arrival model. In our setting, the secretary problem corresponds to having one unit of inventory, i.e., $b = 1$, and n customers, where $v_{I,j} \in \mathbb{R}^+$ for $1 \leq j \leq n$, i.e., we relax the assumption that there are only two types. The objective is to maximize the probability of selecting the highest-revenue customer in the asymptotic regime, where $n \rightarrow \infty$.

In the classical setting, the arrival sequence is assumed to be a uniformly random permutation of n customers, which corresponds to the extreme case of $p = 1$ under our partially predictable model. In this setting, it is well known that the best-possible online algorithm is the following *deterministic* algorithm (Lindley 1961, Dynkin 1963, Ferguson 1989, Freeman 1983): Observe the first $\lfloor \gamma n \rfloor$ customers, where $\gamma = \frac{1}{e}$; then accept the next one that has the highest revenue so far (if any). The success probability of this algorithm approaches $\frac{1}{e} \approx 0.37$ as $n \rightarrow \infty$. We generalize the classical setting by studying the problem under our demand model. First, we analyze the success probability of a similar class of algorithms for any $p \in (0, 1]$. Next, we show that under our demand model where $p < 1$ —i.e., in the presence of an adversarial component—this class of algorithms is not necessarily the best possible.

For any $\gamma \in (0, 1)$, we define the Observation-Selection Algorithm (OSA_γ), which works similarly to the classical algorithm described above. The formal definition of the algorithm is presented in Algorithm 3.

Algorithm 3 : Observation-Selection Algorithm (OSA_γ , $\gamma \in (0, 1)$)

1. Initialize $v_{\max} \leftarrow 0$.
 2. **Observation period:** Repeat for customer $i = 1, 2, \dots, \lfloor \gamma n \rfloor$: reject customer i and update $v_{\max} \leftarrow \max\{v_{\max}, v_i\}$.
 3. **Selection period:** Repeat for customer $i = \lfloor \gamma n \rfloor + 1, \lfloor \gamma n \rfloor + 2, \dots, n$:
 - If $v_i \geq v_{\max}$, then select customer i and stop the algorithm.
 - Otherwise, reject customer i .
-

In Appendix EC.5, we analyze the success probability of Algorithm 3 and prove the following theorem:

THEOREM 3. *Under the partially predictable model, in the limit $n \rightarrow \infty$, the success probability of OSA_γ approaches $\gamma p \log \frac{1}{\gamma p + 1 - p}$.*

By optimizing over γ , we obtain the following corollary:

COROLLARY 2. *Let $\gamma^* \in (0, 1)$ be the unique solution to*

$$\log(\gamma^* p + 1 - p) + \frac{\gamma^* p}{\gamma^* p + 1 - p} = 0;$$

then, OSA_{γ^} achieves the highest success probability among OSA_γ for all $\gamma \in (0, 1)$.*

p	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
γ^*	0.4935	0.4863	0.4784	0.4696	0.4597	0.4482	0.4348	0.4184	0.3975	0.3679
OSA_{γ^*}	0.0026	0.0105	0.0244	0.0448	0.0724	0.1081	0.1533	0.2095	0.2796	0.3679

Table 3 The optimal length of the observation period, γ^* , and the success probability of OSA_{γ^*} vs. p .

Table 3 presents the optimal length of the observation period, γ^* , and the success probability of OSA_{γ^*} for different values of p . We observe that as the size of the stochastic component increases, i.e., as p increases, the length of the observation period decreases, whereas the success probability increases.

Next, in the following proposition, we establish a lower bound on the success probability when we randomize over the length of the observation period (γ); further, we present an example that shows such randomization increases the success probability for $p < 1$. This illustrates the benefit of employing randomized algorithms in the presence of an adversarial component in the arrival sequence.

PROPOSITION 3. *Under the partially predictable model, for any $0 < \gamma_1 < \gamma_2 < 1$ and $0 < q < 1$, the randomized algorithm that runs OSA_{γ_1} with probability q and OSA_{γ_2} with probability $1 - q$ has an asymptotic success probability of at least*

$$qs_1 + (1 - q)s_2 + \min \left\{ (1 - q)p(1 - p)(1 - \gamma_2), q(1 - p)\frac{\gamma_2 - \gamma_1}{1 - \gamma_1}s_1 \right\}$$

where for $i = 1, 2$, s_i denotes the success probability of OSA_{γ_i} .

The proposition is proven in Appendix EC.5. Suppose $p = 0.5$; randomizing over $\gamma_1 = 0.427$ and $\gamma_2 = 0.69$ with $q = 0.824$ results in a success probability of at least 0.083 (utilizing the result of Proposition 3). On the other hand, the success probability of the best possible deterministic observation period OSA_γ , given in Theorem 3 and Corollary 2, is 0.072.

8. Conclusion

Online resource allocation is a central problem in the operations of numerous online platforms ranging from airline booking systems to hotel booking systems to internet advertising. Despite advances in information technology, demand arrival processes are rarely perfectly predictable. The presence of unpredictable patterns limits the performance of most allocation algorithms that rely on fully accurate prediction of future demand based on observed data. At the same time, ignoring available information and taking a completely worst-case approach usually leads to online allocation policies that are too conservative. In this paper we take a middle ground approach and introduce the first arrival model that contains both adversarial (thus unpredictable) and stochastic (predictable) components. Our demand model requires no forecast of demand; however, the stochastic component allows us to partially predict future demand as the sequence of arrivals unfolds. In our model, the relative size of the stochastic component, p , represents the level of predictability of the demand.

Under our proposed demand model, we study the basic yet fundamental problem of allocating a single resource with an arbitrary initial inventory to a sequence of customers that belong to two types, with type-1 generating higher revenue. For this problem, we design a non-adaptive algorithm as well as an adaptive one. We analyze the competitive ratios of our algorithms and show that they outperform existing ones under our proposed demand model. The first implication of our analysis is that, by employing our algorithms, we can take advantage of limited available information (due to the presence of the stochastic component) to improve the revenue of the firm compared to a fully conservative approach. Indeed, the competitive ratios of our algorithms are parameterized by p , and for both algorithms the ratio increases with p (the relative size of the stochastic component), which highlights the value of even partial predictability.

Further, we show that our adaptive algorithm—which repeatedly computes upper bounds on the total number of customers of each type based on observed data, and makes online decisions

based on those bounds—achieves a higher competitive ratio when the initial inventory b is sufficiently large. This underlines the significance of adapting to the data, even though it contains an adversarial component. Analyzing the adaptive algorithm, however, is considerably more challenging. We establish a lower bound on the competitive ratio by constructing a novel factor-revealing mathematical program.

On the other hand, when b is small (more precisely, when $b = o(\sqrt{n})$), we prove an upper bound on the competitive ratio of any deterministic or randomized online algorithm that matches the competitive ratio of our non-adaptive algorithm (up to an error term). This implies (1) our non-adaptive algorithm is the best possible in this regime, and (2) when the initial inventory is small relative to the time horizon, we may not be able to effectively adapt to observed data before allocating most of the inventory. We also have heuristic arguments—in which we do not characterize the error terms—that indicate that (1) our adaptive algorithm achieves the best possible competitive ratio in the regime where $b = \kappa n$ (where $\kappa \in (0, 1]$ is a constant) and (2) underestimating parameter p does not affect the competitive ratio of our adaptive algorithm, whereas (3) if we overestimate p by (a small amount), its competitive ratio decreases only slightly. Because making the above results rigorous will make the paper prohibitively long, these results are not included in the paper.

To illustrate the application of our model to other online allocation problems, we study the secretary problem under our demand model. We analyze the celebrated policy of selecting the highest revenue customer after an observation period with a deterministic length of γ under our new model, and find the optimum value of γ (which is parameterized by p). We further show that, in the presence of an adversarial component and unlike the classical setting, randomizing over the length of the observation period may increase the probability of selecting the highest revenue customer.

In this paper, we use a discrete time model and also assume that the arrival times of customers from the stochastic group are randomly permuted among their predetermined positions. We believe similar results can be obtained for a model where a total of n customers from the two groups (i.e., the stochastic and adversarial group) arrive according to independent Poisson processes with rates p and $1 - p$. We leave the rigorous treatment of this alternative model for future research.

Studying other online allocation problems under our new demand model is a promising direction for future research. Our consequential concentration result from Lemma 1 can be extended to any finite number of types. Further, we believe that by combining our ideas for adaptively computing bounds on the demand of each type with those of Lan et al. (2008), and utilizing the concentration results, one can generalize our algorithms to a setting with any finite number of types. Such extensions are, however, beyond the scope of this paper.

Acknowledgments

Jaillet acknowledges the research support of the Office of Naval Research grants N00014-12-1-0999 and N00014-16-1-2786. We would like to thank Lê Nguyễn Hoàng for fruitful discussions on the concentration result.

References

- Agrawal, S., Z. Wang, and Y. Ye (2014). A Dynamic Near-Optimal Algorithm for Online Linear Programming. *Operations Research* 62(4), 876–890.
- Araman, V. F. and R. Caldentey (2009). Dynamic pricing for nonperishable products with demand learning. *Operations research* 57(5), 1169–1188.
- Ball, M. O. and M. Queyranne (2009). Toward robust revenue management: Competitive analysis of online booking. *Operations Research* 57(4), 950–963.
- Belobaba, P. P. (1987). Survey paper-airline yield management an overview of seat inventory control. *Transportation Science* 21(2), 63–73.
- Belobaba, P. P. (1989). OR practice-application of a probabilistic decision model to airline seat inventory control. *Operations Research* 37(2), 183–197.
- Ben-Tal, A. and A. Nemirovski (2002). Robust optimization – methodology and applications. *Mathematical Programming* 92(3), 453–480.
- Bertsimas, D., D. Pachamanova, and M. Sim (2004, November). Robust linear optimization under general norms. *Operations Research Letters* 32(6), 510–516.
- Besbes, O. and A. Zeevi (2009). Dynamic pricing without knowing the demand function: Risk bounds and near-optimal algorithms. *Operations Research* 57(6), 1407–1420.
- Brumelle, S. L. and J. I. McGill (1993). Airline seat allocation with multiple nested fare classes. *Operations Research* 41(1), 127–137.
- Buchbinder, N. and J. Naor (2009). The design of competitive online algorithms via a primal: dual approach. *Foundations and Trends® in Theoretical Computer Science* 3(2–3), 93–263.
- Cheapair (2016). What the airlines never tell you about airfares. <https://www.cheapair.com/blog/travel-tips/what-the-airlines-never-tell-you-about-airfares/>.
- Chernoff, H. (1952). A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *The Annals of Mathematical Statistics*, 493–507.
- Ciocan, D. F. and V. Farias (2012). Model Predictive Control for Dynamic Resource Allocation. *Mathematics of Operations Research* 37(3), 501–525.
- Cooper, W. L. (2002). Asymptotic Behavior of an Allocation Policy for Revenue Management. *Operations Research* 50(4), 720–727.

-
- CWT (2016). Gender differences in booking business travel. *Carson Wagonlit Travel Report*.
- Devanur, N. R. and T. P. Hayes (2009). The adwords problem: online keyword matching with budgeted bidders under random permutations. In *Proceedings of the 10th ACM conference on Electronic Commerce*, pp. 71–78.
- Dynkin, E. B. (1963). The optimum choice of the instant for stopping a markov process. In *Soviet Math. Dokl*, Volume 4.
- Esfandiari, H., N. Korula, and V. Mirrokni (2015). Online Allocation with Traffic Spikes: Mixing Adversarial and Stochastic Models. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pp. 169–186.
- Ferguson, T. S. (1989). Who solved the secretary problem? *Statistical science*, 282–289.
- Freeman, P. R. (1983). The secretary problem and its extensions: A review. *International Statistical Review/Revue Internationale de Statistique*, 189–206.
- Hush, D. and C. Scovel (2005). Concentration of the hypergeometric distribution. *Statistics & Probability Letters* 75(2), 127–132.
- Jasin, S. (2015). Performance of an LP-Based Control for Revenue Management with Unknown Demand Parameters. *Operations Research* 63(4), 909–915.
- Kesselheim, T., R. Kleinberg, and R. Niazadeh (2015). Secretary problems with non-uniform arrival order. In *Proceedings of the Forty-Seventh annual ACM on symposium on Theory of Computing*, pp. 879–888.
- Kesselheim, T., A. Tönnis, K. Radke, and B. Vöcking (2014). Primal Beats Dual on Online Packing LPs in the Random-Order Model. In *Proceedings of the 46th annual ACM symposium on Theory of Computing*, Number 1, pp. 303–312.
- Kleinberg, R. (2005). A multiple-choice secretary algorithm with applications to online auctions. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete Algorithms*, pp. 630–631.
- Lan, Y., H. Gao, M. O. Ball, and I. Karaesmen (2008). Revenue Management with Limited Demand Information. *Management Science* 54(9), 1594–1609.
- Lautenbacher, C. J. and S. Stidham Jr. (1999). The underlying Markov decision process in the single-leg airline yield-management problem. *Transportation Science* 33(2), 136–146.
- Lee, T. C. and M. Hersh (1993). A model for dynamic airline seat inventory control with multiple seat bookings. *Transportation Science* 27(3), 252–265.
- Lindley, D. V. (1961). Dynamic programming and decision theory. *Applied Statistics*, 39–51.
- Littlewood, K. (2005). Special issue papers: Forecasting and control of passenger bookings. *Journal of Revenue and Pricing Management* 4(2), 111–123.
- Mahdian, M., H. Nazerzadeh, and A. Saberi (2007). Allocating online advertisement space with unreliable estimates. In *Proceedings of the 8th ACM conference on Electronic Commerce*, pp. 288–294.

- McDiarmid, C. (1998). Concentration. In *Probabilistic methods for algorithmic discrete mathematics*, pp. 195–248. Springer.
- Mehta, A., A. Saberi, U. Vazirani, and V. Vazirani (2007). Adwords and generalized online matching. *Journal of the ACM (JACM)* 54(5), 22.
- Mirroknii, V., S. Oveis Gharan, and M. Zadimoghaddam (2012). Simultaneous approximations for adversarial and stochastic online budgeted allocation. In *Proceedings of the Twenty-Third annual ACM-SIAM symposium on Discrete Algorithms*, pp. 1690–1701. SIAM.
- Shamsi, D., M. Holtan, R. Luenberger, and Y. Ye (2014). Online allocation rules in display advertising. *arXiv preprint arXiv:1407.5710*.
- Stein, C., V.-A. Truong, and X. Wang (2016). Advance Service Reservations with Heterogeneous Customers. *Working paper*.
- Talluri, K. and G. V. Ryzin (1998, November). An analysis of bid-price controls for network revenue management. *Manage. Sci.* 44(11), 1577–1593.
- Talluri, K. T. and G. J. Van Ryzin (2006). *The theory and practice of revenue management*, Volume 68. Springer Science & Business Media.
- Wang, H., H. Xie, L. Qiu, Y. R. Yang, Y. Zhang, and A. Greenberg (2006, August). Cope: Traffic engineering in dynamic networks. *ACM SIGCOMM Computer Communication Review* 36(4), 99–110.
- Wang, X. and V.-A. Truong (2015). Online Advance Admission Scheduling for Services, with Customer Preferences. *Working paper*.
- Zervas, G., D. Proserpio, and J. W. Byers (2016). The rise of the sharing economy: Estimating the impact of airbnb on the hotel industry. *Forthcoming, Journal of Marketing Research*.

Appendix

EC.1. Proof of Lemma 1

The proof of Lemma 1 is based on the following lemma:

LEMMA EC.1. Define constants $\alpha_{EC.1} \triangleq 5 + \sqrt{6}$, $\bar{\epsilon}_{EC.1} \triangleq 1/24$, and $k_{EC.1} \triangleq 4$. If $\epsilon' \in (0, \bar{\epsilon}_{EC.1}]$ and $n_1 > \frac{k_{EC.1}}{p^2} \log\left(\frac{1}{\epsilon'}\right)$, for any $\lambda \in \{1/n, 2/n, \dots, n/n\}$, we have:

$$\mathbb{P}\left(|O_1(\lambda) - \tilde{o}_1(\lambda)| \geq \alpha_{EC.1} \sqrt{n_1 \log\left(\frac{1}{\epsilon'}\right)}\right) \leq \epsilon'.$$

To prove Lemma EC.1, we use two existing concentration bounds for random variables obtained from sampling with and without replacement. Before proceeding to the proof, we state these

concentration bounds. Using the existing concentration bounds, Lemma EC.1 is proven through a series of auxiliary corollaries (of the concentration results) and lemmas whose proofs are deferred to Section EC.1.2.

First, we use a well-know variant of the classical Chernoff bound (Chernoff (1952)) regarding the concentration of binomial random variables as given in McDiarmid (1998):

THEOREM EC.1 (McDiarmid (1998)). *Let $0 < p < 1$, let X_1, X_2, \dots, X_n be independent binary random variables, with $\mathbb{P}(X_k = 1) = p$ and $\mathbb{P}(X_k = 0) = 1 - p$ for each k , and let $S_n = \sum_{k=1}^n X_k$. Then for any $t \geq 0$,*

$$\mathbb{P}(|S_n - np| \geq nt) \leq 2e^{-2nt^2}.$$

When applying this theorem in our proof, we find it more insightful and convenient to use the following form of the above concentration result:

COROLLARY EC.1. *For any $k \geq 0$, define constants $\alpha_{EC.1,k} \triangleq 1$ and $\bar{\epsilon}_{EC.1,k} \triangleq k/2$. For $\epsilon \in (0, \bar{\epsilon}_{EC.1,k})$, under the same setting as in Theorem EC.1, we have:*

$$\mathbb{P}\left(|S_n - np| \geq \alpha_{EC.1,k} \sqrt{n \log\left(\frac{1}{\epsilon}\right)}\right) \leq k\epsilon.$$

Second, we use a concentration result for random variables drawn from the hypergeometric distribution given by Hush and Scovel (2005). Recall that hypergeometric distribution is similar to binomial distribution when sampling without replacement is performed. It is defined precisely within the following theorem.

THEOREM EC.2 (Hush and Scovel (2005)). *Let $K \sim \text{Hyper}(n_1, n, m)$ denote the hypergeometric random variable describing the process of counting how many defectives are selected when n_1 items are randomly selected without replacement from a population of n items of which m are defective. Let $\gamma \geq 2$. Then,*

$$\mathbb{P}(K - \mathbb{E}[K] > \gamma) < e^{-2\alpha_{n_1, n, m}(\gamma^2 - 1)}$$

and

$$\mathbb{P}(K - \mathbb{E}[K] < -\gamma) < e^{-2\alpha_{n_1, n, m}(\gamma^2 - 1)},$$

where

$$\alpha_{n_1, n, m} = \max\left\{\frac{1}{n_1 + 1} + \frac{1}{n - n_1 + 1}, \frac{1}{m + 1} + \frac{1}{n - m + 1}\right\}.$$

Similar to the concentration result for binomial distribution, we find it easier to use the following form of the above concentration result:

COROLLARY EC.2. For any $k \geq 0$, define constants $\alpha_{EC.2,k} \triangleq 2$ and $\bar{\epsilon}_{EC.2,k} \triangleq k/2$, $\underline{m}_{EC.2,k} \triangleq \max \left\{ \left(\log \frac{1}{\bar{\epsilon}_{EC.2,k}} \right)^{-1}, 1 \right\}$. For $\epsilon \in (0, \bar{\epsilon}_{EC.2,k})$ and $m \geq \underline{m}_{EC.2,k}$, under the same setting as in Theorem EC.2, we have:

$$\mathbb{P} \left(|K - \mathbb{E}[K]| \geq \alpha_{EC.2,k} \sqrt{m \log \left(\frac{1}{\epsilon} \right)} \right) \leq k\epsilon.$$

Proof Sketch of Lemma EC.1: Before proceeding to the proof of Lemma EC.1, we explain the idea of the proof by going back to the example of Figure 1 from Section 3. Let us consider $\lambda = 5/8$. In the following, we count the number of customers in the stochastic group and the adversarial group in $O_1(\lambda)$ separately.

We begin by counting the number of type-1 customers in the stochastic group that arrive no later than time $5/8$ in \vec{V} in Figure 1. Among the five customers arriving by time $5/8$, two of them are in the stochastic group: customers at positions 2 and 5. We aim to count the number of type-1 customers in these two positions. There are a total of four customers in the stochastic group (there are four black nodes in the middle row). Note that only one of them is type-1 ($v_{I,5} = 1$). Now we take two samples without replacement from the four customers to fill the two positions (2 and 5). Thus, given the realization of the stochastic group (the middle row), the number of type-1 customers in these two positions follows a hypergeometric distribution with parameters $(2, 4, 1)$ (which, as defined in Theorem EC.2, corresponds to taking two samples without replacement from four customers among which one is type-1). In the particular realization of Figure 1, the type-1 customer in the stochastic group is placed in position 2.

Now we count the number of type-1 customers in the adversarial group that arrive no later than time $5/8$ in \vec{V} in Figure 1. In the adversarial sequence \vec{v}_I , there are three type-1 customers (at positions 1, 3, and 5). Any of these three customers will be in the adversarial group independent of each other and with probability $(1 - p)$, and hence the number of type-1 in the adversarial group that arrive no later than time $5/8$ in \vec{V} follows the binomial distribution $\text{Bin}(3, 1 - p)$. In the particular realization of Figure 1, among the three type-1 customers arriving no later than time $5/8$ in \vec{v}_I , two of them are in the adversarial group: customers at position 1 and 3. Therefore, the number of type-1 customers in the adversarial group that arrive no later than time $5/8$ in the particular realization \vec{v} is two.

In the proof of Lemma EC.1, we use the method described in the above example to count the number of customers in $O_1(\lambda)$. For counting the number of customers in the stochastic group in $O_1(\lambda)$: (i) First we count the number of positions before time λ that belong to the stochastic group. Call this number Z . (ii) Next we count the number of type-1 customers in the stochastic group. Call the total number of customers in the stochastic group R and the number of type-1 customers in

the stochastic group R_1 . (iii) We compute the number of type-1 customers in the stochastic group that fill one of these Z positions. Call this number Z_1 . As mentioned above, this is equivalent to taking Z samples without replacement from R customers among which R_1 are type-1. The number Z_1 is the number of customers in the stochastic group in $O_1(\lambda)$. Counting the customers in the adversarial group in $O_1(\lambda)$ is relatively simple. Call this number ζ_1 . Finally, we obtain $O_1(\lambda)$ with the equation $O_1(\lambda) = Z_1 + \zeta_1$. In summary, the random variables have the following distributions:

- $R \sim \text{Bin}(n, p)$,
- $R_1 \sim \text{Bin}(n_1, p)$,
- $Z \sim \text{Bin}(\lambda n, p)$,
- $Z_1 \sim \text{Hyper}(Z, R, R_1)^{10}$,
- $\zeta_1 \sim \text{Bin}(\eta_1(\lambda), 1 - p)$.

The proof of Lemma EC.1 includes establishing concentration results for Z_1 and ζ_1 through a series of auxiliary lemmas; Lemmas EC.2, EC.3, and EC.4 are concerned with the former random variable, and Lemma EC.5 is concerned with the latter one. The proof of these lemmas is deferred to Section EC.1.2.

The first lemma focuses on analyzing R , R_1 , and Z . In particular, we use Corollary EC.1 along with the union bound to show the following:

LEMMA EC.2. *Define constants $\bar{\epsilon}_{EC.2} \triangleq 1/24$ and $\alpha_{EC.2} \triangleq 1$. For $\epsilon \in (0, \bar{\epsilon}_{EC.2}]$, with probability at least $1 - \epsilon/4$, all the following three events happen:*

$$R \in \left(np - \alpha_{EC.2} \sqrt{n \log \left(\frac{1}{\epsilon} \right)}, np + \alpha_{EC.2} \sqrt{n \log \left(\frac{1}{\epsilon} \right)} \right), \quad (\text{EC.1a})$$

$$R_1 \in \left(n_1 p - \alpha_{EC.2} \sqrt{n_1 \log \left(\frac{1}{\epsilon} \right)}, n_1 p + \alpha_{EC.2} \sqrt{n_1 \log \left(\frac{1}{\epsilon} \right)} \right), \text{ and} \quad (\text{EC.1b})$$

$$Z \in \left(\lambda np - \alpha_{EC.2} \sqrt{\lambda n \log \left(\frac{1}{\epsilon} \right)}, \lambda np + \alpha_{EC.2} \sqrt{\lambda n \log \left(\frac{1}{\epsilon} \right)} \right). \quad (\text{EC.1c})$$

We note that conditioned on R , R_1 , and Z , the expected value of Z_1 is $\frac{ZR_1}{R}$. Thus we have:

$$\mathbb{E}[Z_1] = \mathbb{E}[\mathbb{E}[Z_1 | R, R_1, Z]] = \mathbb{E}\left[\frac{ZR_1}{R}\right]. \quad (\text{EC.2})$$

The last expectation is a non-linear function of the three random variables R , R_1 , and Z . Instead of computing the expectation directly, we use the concentration bounds of (EC.1a) - (EC.1c) to show the following lemma:

¹⁰ Note that Z , R , and R_1 are not necessarily independent.

LEMMA EC.3. Define constants $\alpha_{EC.3} \triangleq 4$, $k_{EC.3} \triangleq 4$. Conditioned on the events (EC.1a)-(EC.1c), and when $n_1 > \frac{k_{EC.3}}{p^2} \log\left(\frac{1}{\epsilon}\right)$, we have:

$$\frac{ZR_1}{R} \in \left(\lambda p n_1 - \alpha_{EC.3} \sqrt{n_1 \log\left(\frac{1}{\epsilon}\right)}, \lambda p n_1 + \alpha_{EC.3} \sqrt{n_1 \log\left(\frac{1}{\epsilon}\right)} \right). \quad (\text{EC.3})$$

Lemmas EC.2 and EC.3 together imply that, for $\epsilon \in (0, \bar{\epsilon}_{EC.2}]$:

$$\mathbb{P}\left(\left|\frac{R_1 Z}{R} - p n_1 \lambda\right| \geq \alpha_{EC.3} \sqrt{n_1 \log\left(\frac{1}{\epsilon}\right)}\right) \leq \frac{\epsilon}{4} \quad (\text{EC.4})$$

Having (EC.2) and Lemma EC.3, we are ready to establish a concentration result for Z_1 . We partition the sample space of (R, R_1, Z) into two events as follows: the event where (EC.1a)-(EC.1c) hold, denoted by \mathcal{E} ; the complement event, denoted by \mathcal{E}^c .¹¹ Note that Lemma EC.2 implies that $\mathbb{P}(\mathcal{E}^c) \leq \frac{\epsilon}{4}$. Using the law of total probability, we have: for any $\tilde{\alpha} > 0$,

$$\begin{aligned} & \mathbb{P}\left(|Z_1 - \mathbb{E}[Z_1 | R, R_1, Z]| \geq \tilde{\alpha} \sqrt{n_1 \log\left(\frac{1}{\epsilon}\right)}\right) \\ &= \mathbb{P}(\mathcal{E}^c) \mathbb{P}\left(|Z_1 - \mathbb{E}[Z_1 | R, R_1, Z]| \geq \tilde{\alpha} \sqrt{n_1 \log\left(\frac{1}{\epsilon}\right)} \mid \mathcal{E}^c\right) \\ &+ \mathbb{P}(\mathcal{E}) \mathbb{P}\left(|Z_1 - \mathbb{E}[Z_1 | R, R_1, Z]| \geq \tilde{\alpha} \sqrt{n_1 \log\left(\frac{1}{\epsilon}\right)} \mid \mathcal{E}\right) \\ &\leq \frac{\epsilon}{4} \cdot 1 + 1 \cdot \mathbb{P}\left(|Z_1 - \mathbb{E}[Z_1 | R, R_1, Z]| \geq \tilde{\alpha} \sqrt{n_1 \log\left(\frac{1}{\epsilon}\right)} \mid \mathcal{E}\right). \end{aligned} \quad (\text{EC.5})$$

Using Corollary EC.2 and the definition of events (EC.1a)-(EC.1c), we show the following lemma:

LEMMA EC.4. Define constants $\bar{\epsilon}_{EC.4} \triangleq 1/24$, $\alpha_{EC.4} \triangleq \sqrt{6}$, and $k_{EC.4} \triangleq 4$. For $\epsilon \in (0, \bar{\epsilon}_{EC.4}]$, if $n_1 > \frac{k_{EC.4}}{p^2} \log\left(\frac{1}{\epsilon}\right)$, and $(R, R_1, Z) \in \mathcal{E}$, we have:

$$\mathbb{P}\left(|Z_1 - \mathbb{E}[Z_1 | R, R_1, Z]| \geq \alpha_{EC.4} \sqrt{n_1 \log\left(\frac{1}{\epsilon}\right)} \mid R, R_1, Z\right) \leq \frac{\epsilon}{4}. \quad (\text{EC.6})$$

Putting Lemma EC.4 back to (EC.5) and setting $\tilde{\alpha} = \alpha_{EC.4}$, we get:

$$\mathbb{P}\left(|Z_1 - \mathbb{E}[Z_1 | R, R_1, Z]| \geq \alpha_{EC.4} \sqrt{n_1 \log\left(\frac{1}{\epsilon}\right)}\right) \leq \frac{\epsilon}{2} \quad (\text{EC.7})$$

Finally, we have the following lemma regarding ζ_1 :

¹¹ We note that this event is only locally defined within this appendix, and it is not the same as the one defined in Definition 4.

LEMMA EC.5. Define constants $\bar{\epsilon}_{EC.5} \triangleq 1/8$ and $\alpha_{EC.5} \triangleq 1$. For $\epsilon \in (0, \bar{\epsilon}_{EC.5}]$, we have:

$$\mathbb{P} \left(|\zeta_1 - (1-p)\eta_1(\lambda)| \geq \alpha_{EC.5} \sqrt{n_1 \log \left(\frac{1}{\epsilon} \right)} \right) \leq \frac{\epsilon}{4} \quad (\text{EC.8})$$

With the lemmas above, we are ready to prove Lemma EC.1:

Proof of Lemma EC.1: First, we note that we set the constants in Lemma EC.1 and Lemmas EC.2-EC.5 such that we get $\bar{\epsilon}_{EC.1} = \min\{\bar{\epsilon}_{EC.2}, \bar{\epsilon}_{EC.4}, \bar{\epsilon}_{EC.5}\}$, $k_{EC.1} = \max\{k_{EC.4}, k_{EC.3}\}$ and $\alpha_{EC.1} = \alpha_{EC.3} + \alpha_{EC.4} + \alpha_{EC.5}$. Now we can apply the union bound on (EC.7), (EC.4), and (EC.8) and obtain: when $0 < \epsilon' \leq \bar{\epsilon}_{EC.1}$ and $n_1 > \frac{k_{EC.1}}{p^2} \log \left(\frac{1}{\epsilon'} \right)$, with probability at least $1 - \epsilon'$,

$$\begin{aligned} |Z_1 - \mathbb{E}[Z_1 | R, R_1, Z]| &< \alpha_{EC.4} \sqrt{n_1 \log \left(\frac{1}{\epsilon'} \right)}, \\ \left| \frac{R_1 Z}{R} - pn_1 \lambda \right| &< \alpha_{EC.3} \sqrt{n_1 \log \left(\frac{1}{\epsilon'} \right)}, \text{ and} \\ |\zeta_1 - (1-p)\eta_1(\lambda)| &< \alpha_{EC.5} \sqrt{n_1 \log \left(\frac{1}{\epsilon'} \right)}. \end{aligned}$$

We have $O_1(\lambda) = Z_1 + \zeta_1$, and thus, by using the triangular inequality (note that $\mathbb{E}[Z_1 | R, R_1, Z] = \frac{R_1 Z}{R}$),

$$|O_1(\lambda) - \tilde{o}_1(\lambda)| \leq |Z_1 - \mathbb{E}[Z_1 | R, R_1, Z]| + \left| \frac{R_1 Z}{R} - pn_1 \lambda \right| + |\zeta_1 - (1-p)\eta_1(\lambda)|,$$

which, according to the three inequalities above and the definition $\alpha_{EC.1} = \alpha_{EC.3} + \alpha_{EC.4} + \alpha_{EC.5}$, is smaller than $\alpha_{EC.1} \sqrt{n_1 \log \left(\frac{1}{\epsilon'} \right)}$. \square

With Lemma EC.1, we are ready to prove Lemma 1:

Proof of Lemma 1: The proof consists of two steps: First, we note that the concentration results similar to Lemma EC.1 can be obtained for the other three random variables $O_1(\lambda) + O_2(\lambda)$, $O_2(\lambda)$, and $O_2^S(\lambda)$. When applied to $O_2^S(\lambda)$, the only modification is that we do not need to consider Lemma EC.5 and (EC.8).

Second, we apply the union bound on the probability that at least one of the $4n$ events in (2a)-(3b) is violated (note that λ takes n different non-zero values: when $\lambda = 0$, $|O_1(\lambda) - \tilde{o}_1(\lambda)| = |0 - 0| = 0$ and the same holds for the other three random variables), and choose the appropriate constants.

To prove this lemma, we first note that we set the constants in Lemmas 1 and EC.1 such that we get $\alpha = 2\alpha_{EC.1}$, $\bar{\epsilon} = \bar{\epsilon}_{EC.1}$, and $k = 4k_{EC.1}$.

Using Lemma EC.1 with $\epsilon' = \frac{\epsilon}{4n}$, the lemma holds because all of the following three statements are true:

$$\text{If } \epsilon \leq \bar{\epsilon}, \text{ then } \epsilon' \leq \bar{\epsilon}_{EC.1}. \quad (\text{EC.9a})$$

$$\text{If } n_1 > \frac{k}{p^2} \log n, \text{ then } n_1 > \frac{k_{EC.1}}{p^2} \log \left(\frac{1}{\epsilon'} \right). \quad (\text{EC.9b})$$

$$\text{If } |O_1(\lambda) - \tilde{o}_1(\lambda)| \geq \alpha \sqrt{n_1 \log n}, \text{ then } |O_1(\lambda) - \tilde{o}_1(\lambda)| \geq \alpha_{EC.1} \sqrt{n_1 \log \left(\frac{1}{\epsilon'} \right)}. \quad (\text{EC.9c})$$

Because $\epsilon' < \epsilon$, and $\bar{\epsilon} = \bar{\epsilon}_{EC.1}$, (EC.9a) holds. Before proceeding to the other two conditions, we first note that $\log \left(\frac{1}{\epsilon'} \right) = \log \left(\frac{4n}{\epsilon} \right) \leq \log \left(\frac{n^3}{\epsilon} \right) \leq \log(n^4) = 4 \log n$, where we use $n \geq 2$ in the first inequality and $\epsilon \geq \frac{1}{n}$ in the second inequality. Therefore, (EC.9b) and (EC.9c) hold because $k = 4k_{EC.1}$, and $\alpha = 2\alpha_{EC.1}$. \square

EC.1.1. Further Remark on Deterministic Approximations

REMARK EC.1. In Lemma 1, we use the deterministic value $\tilde{o}_j(\lambda)$ rather than $\mathbb{E}[O_j(\lambda)]$ to estimate $O_j(\lambda)$ because $\tilde{o}_j(\lambda)$ is a very simple function of n_j and $\eta_j(\lambda)$. Here we provide an example to show that $\tilde{o}_j(\lambda)$ and $\mathbb{E}[O_j(\lambda)]$ are not necessarily the same, which explains why we do not write the term $\tilde{o}_j(\lambda)$ as $\mathbb{E}[O_j(\lambda)]$.

Let us consider an example where $n = 2$ and $\vec{v}_I = (v_{I,1}, v_{I,2}) = (1, 0)$ at $\lambda = 1/2$. First we compute $\mathbb{E}[O_1(1/2)]$. Because $O_1(1/2)$ consists of only one customer,

$$\mathbb{E}[O_1(1/2)] = \mathbb{P}(V_1 = 1).$$

We can then use the law of total probability to express the probability as

$$\begin{aligned} \mathbb{P}(V_1 = 1) &= \mathbb{P}(V_1 = 1 | 1 \notin \mathcal{S}) \mathbb{P}(1 \notin \mathcal{S}) \\ &\quad + \mathbb{P}(V_1 = 1 | 1 \in \mathcal{S}, 2 \notin \mathcal{S}) \mathbb{P}(1 \in \mathcal{S}, 2 \notin \mathcal{S}) \\ &\quad + \mathbb{P}(V_1 = 1 | 1, 2 \in \mathcal{S}) \mathbb{P}(1, 2 \in \mathcal{S}). \end{aligned}$$

Following the definitions,

$$\mathbb{E}[O_1(1/2)] = \mathbb{P}(V_1 = 1) = 1 \cdot (1 - p) + 1 \cdot p(1 - p) + \frac{1}{2} \cdot p^2 = 1 - \frac{p^2}{2}.$$

On the other hand,

$$\tilde{o}_1(1/2) = (1 - p)\eta_1\left(\frac{1}{2}\right) + p\frac{1}{2}n_1 = 1 - \frac{p}{2}.$$

Therefore, for all $p \in (0, 1)$, $\mathbb{E}[O_1(1/2)] \neq \tilde{o}_1(1/2)$.

EC.1.2. Proof of Auxiliary Corollaries and Lemmas

Proof of Corollary EC.1: In order to apply Theorem EC.1, we define t such that $2e^{-2nt^2} \leq k\epsilon$, which corresponds to

$$t \geq \sqrt{\frac{1}{2n} \log \frac{2}{k\epsilon}}.$$

By setting t to be $\sqrt{\frac{1}{2n} \log \frac{2}{k\epsilon}}$, what is remaining to prove is that we can find $\alpha_{EC.1,k}, \bar{\epsilon}_{EC.1,k}$ such that when $0 < \epsilon < \bar{\epsilon}_{EC.1,k}$,

$$n\sqrt{\frac{1}{2n} \log \frac{2}{k\epsilon}} \leq \alpha_{EC.1,k} \sqrt{n \log \left(\frac{1}{\epsilon} \right)}.$$

This can be achieved by setting $\bar{\epsilon}_{EC.1,k} = k/2$ and $\alpha_{EC.1,k} = 1$:

$$\epsilon \leq k/2 \Rightarrow \frac{2}{k\epsilon} \leq \frac{1}{\epsilon^2} \Rightarrow \log \frac{2}{k\epsilon} \leq 2 \log \frac{1}{\epsilon} \Rightarrow n\sqrt{\frac{1}{2n} \log \frac{2}{k\epsilon}} \leq \alpha_{EC.1,k} \sqrt{n \log \left(\frac{1}{\epsilon} \right)}.$$

□

Proof of Corollary EC.2: According to Theorem EC.2, when $\gamma \geq 2$,

$$\mathbb{P}(|K - \mathbb{E}[K]| > \gamma) < 2e^{-2\alpha_{n_1,n,m}(\gamma^2-1)}.$$

We first find an upper bound of the above right-hand-side probability when γ and m are large enough. When $m \geq 1$,

$$\alpha_{n_1,n,m} = \max \left\{ \frac{1}{n_1+1} + \frac{1}{n-n_1+1}, \frac{1}{m+1} + \frac{1}{n-m+1} \right\} \geq \frac{1}{m+1} \geq \frac{1}{2m}.$$

Further, when $\gamma \geq 2$, we have: $\gamma^2 - 1 \geq \gamma^2/2$. Putting these two together,

$$2e^{-2\alpha_{n_1,n,m}(\gamma^2-1)} \leq 2e^{-\frac{1}{m} \frac{\gamma^2}{2}}.$$

Therefore, if $\alpha_{EC.2,k} \sqrt{m \log \left(\frac{1}{\epsilon} \right)} \geq 2$ and $m \geq 1$,

$$\mathbb{P} \left(|K - \mathbb{E}[K]| \geq \alpha_{EC.2,k} \sqrt{m \log \left(\frac{1}{\epsilon} \right)} \right) < 2 \exp \left(-\frac{1}{m} \frac{\alpha_{EC.2,k}^2 m \log \left(\frac{1}{\epsilon} \right)}{2} \right) = 2\epsilon^{\alpha_{EC.2,k}^2/2}.$$

Thus, it is sufficient to have $\alpha_{EC.2,k} \sqrt{m \log \left(\frac{1}{\epsilon} \right)} \geq 2$, $m \geq 1$, and

$$2\epsilon^{\alpha_{EC.2,k}^2/2} \leq k\epsilon.$$

The last condition holds by setting $\alpha_{EC.2,k} = 2$ and $\bar{\epsilon}_{EC.2,k} = k/2$ ($\epsilon \leq \bar{\epsilon}_{EC.2,k} = k/2 \Rightarrow \epsilon^2 \leq k\epsilon/2$).

The first two conditions hold by defining $\underline{m}_{EC.2,k} \triangleq \max \left\{ \left(\log \frac{1}{\bar{\epsilon}_{EC.2,k}} \right)^{-1}, 1 \right\}$.

□

Proof of Lemma EC.2: Let $k = 1/12$, Corollary EC.1 implies that there exist $\bar{\epsilon}_{EC.1,k}$ and $\alpha_{EC.1,k}$ such that when $0 < \epsilon \leq \bar{\epsilon}_{EC.1,k}$,

$$\begin{aligned} & 1 - \mathbb{P} \left(R \in \left(np - \alpha_{EC.1,k} \sqrt{n \log \left(\frac{1}{\epsilon} \right)}, np + \alpha_{EC.1,k} \sqrt{n \log \left(\frac{1}{\epsilon} \right)} \right) \right) \\ &= \mathbb{P} \left(|R - np| \geq \alpha_{EC.1,k} \sqrt{n \log \left(\frac{1}{\epsilon} \right)} \right) \leq \epsilon/12. \end{aligned} \quad (\text{EC.10})$$

Defining $\bar{\epsilon}_{EC.2} \triangleq \bar{\epsilon}_{EC.1,k}$ and $\alpha_{EC.2} \triangleq \alpha_{EC.1,k}$, repeating the same for R_1 and Z , and applying the union bound imply the statement. \square

Proof of Lemma EC.3: First we define the constant $\alpha_{EC.3} \triangleq 3\alpha_{EC.2} + 2\alpha_{EC.2}^2 \sqrt{1/k_{EC.3}}$ where $k_{EC.3} \triangleq 4\alpha_{EC.2}^2$. The reason for definition of the constants becomes clear in the process of the proof. We prove the lower bound first. Because $0 < \frac{Z}{R} \leq 1$, the ratio does not increase by subtracting the same positive number from both the denominator and the numerator if the denominator remains positive after the subtraction. In particular, we subtract $\alpha_{EC.2} \sqrt{n \log \left(\frac{1}{\epsilon} \right)}$ from both the denominator and the numerator, Therefore,

$$\frac{Z}{R} > \frac{Z - \alpha_{EC.2} \sqrt{n \log \left(\frac{1}{\epsilon} \right)}}{R - \alpha_{EC.2} \sqrt{n \log \left(\frac{1}{\epsilon} \right)}}. \quad (\text{EC.11})$$

Note that $R - \alpha_{EC.2} \sqrt{n \log \left(\frac{1}{\epsilon} \right)} > 0$, because under event (EC.1a), we have:

$$R - \alpha_{EC.2} \sqrt{n \log \left(\frac{1}{\epsilon} \right)} \geq np - 2\alpha_{EC.2} \sqrt{n \log \left(\frac{1}{\epsilon} \right)}.$$

Therefore,

$$n > \frac{4\alpha_{EC.2}^2}{p^2} \log \left(\frac{1}{\epsilon} \right) \Rightarrow R - \alpha_{EC.2} \sqrt{n \log \left(\frac{1}{\epsilon} \right)} > 0. \quad (\text{EC.12})$$

The first inequality in (EC.12) holds because $n \geq n_1$, and, by assumption in the lemma, $n_1 > \frac{k_{EC.3}}{p^2} \log \left(\frac{1}{\epsilon} \right) = \frac{4\alpha_{EC.2}^2}{p^2} \log \left(\frac{1}{\epsilon} \right)$ where we use the fact that we defined $k_{EC.3} = 4\alpha_{EC.2}^2$.

Going back to (EC.11), under the events (EC.1a) and (EC.1c), we have:

$$\begin{aligned} \frac{Z}{R} &> \frac{Z - \alpha_{EC.2} \sqrt{n \log \left(\frac{1}{\epsilon} \right)}}{R - \alpha_{EC.2} \sqrt{n \log \left(\frac{1}{\epsilon} \right)}} \\ &> \frac{\lambda np - \alpha_{EC.2} \sqrt{\lambda n \log \left(\frac{1}{\epsilon} \right)} - \alpha_{EC.2} \sqrt{n \log \left(\frac{1}{\epsilon} \right)}}{np + \alpha_{EC.2} \sqrt{n \log \left(\frac{1}{\epsilon} \right)} - \alpha_{EC.2} \sqrt{n \log \left(\frac{1}{\epsilon} \right)}} \geq \lambda - \frac{2\alpha_{EC.2} \sqrt{n \log \left(\frac{1}{\epsilon} \right)}}{np}. \end{aligned} \quad (\text{EC.13})$$

Combining (EC.13) and with the lower-bound on R_1 under event (EC.1b), we get:

$$\begin{aligned} \frac{ZR_1}{R} &> \left(\lambda - \frac{2\alpha_{EC.2}}{p} \sqrt{\frac{\log\left(\frac{1}{\epsilon}\right)}{n}} \right) \left(n_1 p - \alpha_{EC.2} \sqrt{n_1 \log\left(\frac{1}{\epsilon}\right)} \right) \\ &= \lambda n_1 p - 2\alpha_{EC.2} n_1 \sqrt{\frac{\log\left(\frac{1}{\epsilon}\right)}{n}} - \lambda \alpha_{EC.2} \sqrt{n_1 \log\left(\frac{1}{\epsilon}\right)} + 2 \frac{\alpha_{EC.2}^2}{p} \sqrt{\frac{n_1}{n}} \log\left(\frac{1}{\epsilon}\right). \end{aligned}$$

Because $n_1 \leq n$, we have $n_1 \sqrt{\frac{1}{n}} \leq \sqrt{n_1}$. By definition, $\alpha_{EC.3} \geq 3\alpha_{EC.2} \geq (2 + \lambda)\alpha_{EC.2}$, and therefore, the right hand side of the above inequality is at least

$$\lambda n_1 p - \alpha_{EC.3} \sqrt{n_1 \log\left(\frac{1}{\epsilon}\right)}.$$

Hence we complete the proof for the lower bound part of Inequality (EC.3). When it comes to the upper bound, we can use the same argument as the lower bound and obtain,

$$\begin{aligned} \frac{ZR_1}{R} &< \left(\lambda + \frac{2\alpha_{EC.2}}{p} \sqrt{\frac{\log\left(\frac{1}{\epsilon}\right)}{n}} \right) \left(n_1 p + \alpha_{EC.2} \sqrt{n_1 \log\left(\frac{1}{\epsilon}\right)} \right) \\ &= \lambda n_1 p + 2\alpha_{EC.2} n_1 \sqrt{\frac{\log\left(\frac{1}{\epsilon}\right)}{n}} + \lambda \alpha_{EC.2} \sqrt{n_1 \log\left(\frac{1}{\epsilon}\right)} + 2 \frac{\alpha_{EC.2}^2}{p} \sqrt{\frac{n_1}{n}} \log\left(\frac{1}{\epsilon}\right) \\ &\leq \lambda n_1 p + 2\alpha_{EC.2} \sqrt{n_1 \log\left(\frac{1}{\epsilon}\right)} + \alpha_{EC.2} \sqrt{n_1 \log\left(\frac{1}{\epsilon}\right)} + 2 \frac{\alpha_{EC.2}^2}{p} \sqrt{n_1 \log\left(\frac{1}{\epsilon}\right)} \sqrt{\frac{\log\left(\frac{1}{\epsilon}\right)}{n}} \\ &= \lambda n_1 p + \left(3\alpha_{EC.2} + 2\alpha_{EC.2}^2 \frac{1}{p} \sqrt{\frac{\log\left(\frac{1}{\epsilon}\right)}{n}} \right) \sqrt{n_1 \log\left(\frac{1}{\epsilon}\right)}, \tag{EC.14} \end{aligned}$$

where the inequality in the third line comes from the fact that $\sqrt{\frac{n_1}{n}} \leq 1$ and $\lambda \leq 1$. Combining the definition $\alpha_{EC.3} = 3\alpha_{EC.2} + 2\alpha_{EC.2}^2 \sqrt{1/k_{EC.3}}$ and (EC.14), it is sufficient to have

$$\frac{1}{p} \sqrt{\frac{\log\left(\frac{1}{\epsilon}\right)}{n}}$$

upper bounded by the constant $\sqrt{1/k_{EC.3}}$. By the assumption of this lemma, $n_1 \geq \frac{k_{EC.3}}{p^2} \log\left(\frac{1}{\epsilon}\right)$, and thus $\frac{1}{p} \sqrt{\frac{\log\left(\frac{1}{\epsilon}\right)}{n_1}} \leq \sqrt{1/k_{EC.3}}$. Further, because $n \geq n_1$, we have: $\frac{1}{p} \sqrt{\frac{\log\left(\frac{1}{\epsilon}\right)}{n}} \leq \frac{1}{p} \sqrt{\frac{\log\left(\frac{1}{\epsilon}\right)}{n_1}} \leq \sqrt{1/k_{EC.3}}$. Thus we have:

$$3\alpha_{EC.2} + 2\alpha_{EC.2}^2 \frac{1}{p} \sqrt{\frac{\log\left(\frac{1}{\epsilon}\right)}{n}} \leq 3\alpha_{EC.2} + 2\alpha_{EC.2}^2 \sqrt{1/k_{EC.3}} \leq \alpha_{EC.3}.$$

Putting this back in (EC.14) completes the proof of the lemma. \square

Proof of Lemma EC.4: Applying Corollary EC.2 with $k = 1/4$, there exist positive real numbers $\alpha'_{EC.4} \triangleq \alpha_{EC.2,k}$, $\bar{\epsilon}_{EC.2,k}$, and $\underline{m} \triangleq \underline{m}_{EC.2,k}$ such that when $0 < \epsilon \leq \bar{\epsilon}_{EC.2,k}$ and $R_1 \geq \underline{m}$,

$$\mathbb{P} \left(|Z_1 - \mathbb{E}[Z_1|R, R_1, Z]| \geq \alpha'_{EC.4} \sqrt{R_1 \log \left(\frac{1}{\epsilon} \right)} \right) \leq \frac{\epsilon}{4}.$$

Now, we define $\bar{\epsilon}_{EC.4} \triangleq \min\{\bar{\epsilon}_{EC.2,k}, \bar{\epsilon}_{EC.2}\}$, $k_{EC.4} \triangleq \max \left\{ 4\alpha_{EC.2}^2, \frac{2\underline{m}}{\log \left(\frac{1}{\bar{\epsilon}_{EC.4}} \right)} \right\}$ and $\alpha_{EC.4} \triangleq \alpha'_{EC.4} \sqrt{\frac{3}{2}}$.

First we check the condition $R_1 \geq \underline{m}$: Because $n_1 > \frac{k_{EC.4}}{p^2} \log \left(\frac{1}{\epsilon} \right) \geq \frac{4\alpha_{EC.2}^2}{p^2} \log \left(\frac{1}{\epsilon} \right)$,

$$\alpha_{EC.2} \sqrt{n_1 \log \left(\frac{1}{\epsilon} \right)} \leq \frac{n_1 p}{2}. \quad (\text{EC.15})$$

Therefore, under event (EC.1b), we have:

$$R_1 > n_1 p - \alpha'_{EC.4} \sqrt{n_1 \log \left(\frac{1}{\epsilon} \right)} \geq \frac{1}{2} n_1 p \geq \underline{m},$$

where the first inequality follows the definition of event (EC.1b), second one uses Inequality (EC.15), and the last one uses $k_{EC.4} \geq \frac{2\underline{m}}{\log \left(\frac{1}{\epsilon} \right)}$ (which gives $n_1 \geq \frac{k_{EC.4}}{p^2} \log \left(\frac{1}{\epsilon} \right) \geq \frac{2\underline{m}}{p^2} \geq \frac{2\underline{m}}{p}$).

Next we show that because we defined $\alpha_{EC.4} = \alpha'_{EC.4} \sqrt{\frac{3}{2}}$, we have:

$$\alpha'_{EC.4} \sqrt{R_1 \log \left(\frac{1}{\epsilon} \right)} \leq \alpha_{EC.4} \sqrt{n_1 \log \left(\frac{1}{\epsilon} \right)}. \quad (\text{EC.16})$$

This again follows by the definition of $k_{EC.4}$ and event (EC.1b) and Inequality (EC.15):

$$R_1 < n_1 p + \alpha'_{EC.4} \sqrt{n_1 \log \left(\frac{1}{\epsilon} \right)} \leq \frac{3}{2} n_1 p \leq \frac{3}{2} n_1.$$

Inequality (EC.16) implies that:

$$\begin{aligned} & \mathbb{P} \left(|Z_1 - \mathbb{E}[Z_1|R, R_1, Z]| \geq \alpha_{EC.4} \sqrt{n_1 \log \left(\frac{1}{\epsilon} \right)} \right) \\ & \leq \mathbb{P} \left(|Z_1 - \mathbb{E}[Z_1|R, R_1, Z]| \geq \alpha'_{EC.4} \sqrt{R_1 \log \left(\frac{1}{\epsilon} \right)} \right) \leq \frac{\epsilon}{4}, \end{aligned}$$

which completes the proof of the lemma. \square

Proof of Lemma EC.5: Recall that ζ_1 follows the binomial distribution $\text{Bin}(\eta_1(\lambda), 1-p)$. Hence, the lemma follows straightforwardly from Corollary EC.1. \square

EC.2. Missing proofs of Subsection 4.2

Proof of Lemma 3: The revenue of Algorithm 1 in this case is $ALG_1(\vec{v}) = b - (1 - a)[q_{2,e}(1) + q_{2,f}(1)]$, which is decreasing in $q_{2,e}(1) + q_{2,f}(1)$. Note that due to the fixed threshold rule, we already have an upper bound on $q_{2,f}(1)$, i.e., $q_{2,f}(1) \leq \theta b$. As a result, using Lemma 2, $ALG_1(\vec{v}) \geq b - (1 - a)(\theta b + p(b - n_1)^+ + \Delta)$. Note that $OPT(\vec{v}) \leq b - (1 - a)(b - n_1)^+$, and thus

$$\begin{aligned} \frac{ALG_1(\vec{v})}{OPT(\vec{v})} &\geq \frac{b - (1 - a)(\theta b + p(b - n_1)^+ + \Delta)}{b - (1 - a)(b - n_1)^+} \\ &\geq \frac{b - (1 - a)(\theta b + (p + \theta)(b - n_1)^+ + \Delta)}{b - (1 - a)(b - n_1)^+} && (\theta \geq 0) \\ &= p + \frac{1 - p}{2 - a} - \frac{(1 - a)\Delta}{OPT(\vec{v})}. && (1 - (1 - a)\theta = p + (1 - p)/(2 - a)) \end{aligned}$$

The rest follows from the simple inequality $OPT(\vec{v}) \geq ab$ due to $q_1(1) + q_{2,e}(1) + q_{2,f}(1) = b$. \square

Proof of Lemma 5: We consider three cases of Lemma 4 separately.

(a) $q_{2,e}(1) + q_{2,f}(1) = n_2$:

Algorithm 1 accepts all customers and hence achieves the optimal revenue, i.e., $\frac{ALG_1(\vec{v})}{OPT(\vec{v})} = 1$.

(b) $q_{2,f}(1) = \lfloor \theta b \rfloor$ and $n_1 > bp - 3\Delta$:

$ALG_1(\vec{v}) \geq n_1 + a\theta b$ and $OPT(\vec{v}) \leq ab + n_1(1 - a)$. Therefore,

$$\frac{ALG_1(\vec{v})}{OPT(\vec{v})} \geq \frac{n_1 + a\theta b}{ab + n_1(1 - a)},$$

which is increasing in n_1 , so the ratio is minimized at $n_1 = bp - 3\Delta$, which is a special case of the last case, which we analyze next.

(c) $q_{2,f}(1) = \lfloor \theta b \rfloor$, $n_1 \leq bp - 3\Delta$, and $q_{2,e}(1) \geq (p(n_1 + n_2) - n_1 - 5\Delta)^+$: First, we remark that following the discussion in the proof of Lemma 4, we assume, without loss of generality, $n_1 + n_2 \leq b$. Note that by construction, the alternative adversarial instance has the same optimum offline solution, i.e., $OPT(\vec{v}) = OPT(\vec{v}_A)$.

$$\begin{aligned} ALG_1(\vec{v}) &\geq n_1 + a(p(n_1 + n_2) - n_1 - 5\Delta + \theta b) \\ &\geq n_1 + a(p(n_1 + n_2) - n_1 - 5\Delta + \theta(n_1 + n_2)) && (b \geq n_1 + n_2) \\ &= n_1(1 - a + pa + \theta a) + n_2(p + \theta)a - 5\Delta a \\ &\geq n_1((1 - a)(p + \theta) + a(p + \theta)) + n_2(p + \theta)a - 5\Delta a && (p + \theta \leq 1) \\ &= (p + \theta)(n_1 + n_2 a) - 5\Delta a \geq (p + \theta)OPT(\vec{v}) - 5\Delta a \\ &= \left(p + \frac{1 - p}{2 - a}\right)OPT(\vec{v}) - 5\Delta a. && (p + \theta = p + \frac{1 - p}{2 - a}) \end{aligned}$$

Since $OPT(\vec{v}) \geq q_{2,f}(1)a = \theta ba$,

$$\frac{ALG_1(\vec{v})}{OPT(\vec{v})} \geq p + \frac{1 - p}{2 - a} - \frac{5\Delta a}{\theta ba} = p + \frac{1 - p}{2 - a} - \frac{5\Delta}{\theta b}.$$

□

Proof of Lemma 6: Clearly, either condition (a) holds or $q_1(1) = n_1$. Below we consider the cases where $q_1(1) = n_1$. If $q_{2,f}(1) < \lfloor \theta b \rfloor$, then we do not reject any type-2 customer, and thus condition (b) holds. The interesting case is when $q_{2,f}(1) = \lfloor \theta b \rfloor$. We prove that condition (c) will hold. First note that in this case, we know $n_2 \geq \theta b$. As a result, (7) implies we can use the concentration result of (3b).

Following the discussion in the proof of Lemma 4, we assume, without loss of generality, $n_1 + n_2 \leq b$. Further, if we find a time $\hat{\lambda}$ for which we have:

$$o_1(\lambda) + o_2^S(\lambda) - o_2^S(\hat{\lambda}) \leq \lfloor \lambda p b \rfloor \quad \text{for all } \lambda \geq \hat{\lambda}, \quad (\text{EC.17})$$

then, using a similar induction to the one in Lemma 4, we can show:

$$q_{2,e}(\lambda) \geq o_2^S(\lambda) - o_2^S(\hat{\lambda}) \quad \text{for all } \lambda \geq \hat{\lambda}. \quad (\text{EC.18})$$

Here we find a sufficient condition on $\hat{\lambda}$ for Condition (EC.17) to hold.

$$\begin{aligned} o_1(\lambda) + o_2^S(\lambda) - o_2^S(\hat{\lambda}) &< \frac{k}{p^2} \log n + \lambda p n_2 + \Delta - (\hat{\lambda} p n_2 - \Delta) \quad (o_1(\lambda) \leq n_1 < \frac{k}{p^2} \log n, (3b)) \\ &= \frac{k}{p^2} \log n + (\lambda - \hat{\lambda}) p n_2 + 2\Delta \\ &\leq \frac{k}{p^2} \log n + (\lambda - \hat{\lambda}) p b + 2\Delta \quad (n_1 + n_2 \leq b). \end{aligned}$$

As a result, Condition (EC.17) holds if

$$\frac{k}{p^2} \log n + (\lambda - \hat{\lambda}) p b + 2\Delta \leq \lambda p b,$$

which can be achieved when

$$\hat{\lambda} \triangleq \frac{\frac{k}{p^2} \log n + 2\Delta}{p b}.$$

If $\hat{\lambda} \leq 1$, then

$$\begin{aligned} q_{2,e}(1) &\geq o_2^S(1) - o_2^S(\hat{\lambda}) && (\text{Inequality (EC.18)}) \\ &\geq p n_2 - \Delta - (\hat{\lambda} p n_2 + \Delta) && ((7) \text{ and } (3b)) \\ &\geq p n_2 - \left(\frac{k}{p^2} \log n + 2\Delta \right) - 2\Delta && \left(\hat{\lambda} \leq \frac{\frac{k}{p^2} \log n + 2\Delta}{p n_2} \right) \\ &= p n_2 - \frac{k}{p^2} \log n - 4\Delta. \end{aligned}$$

If $\hat{\lambda} > 1$, then

$$\begin{aligned} pn_2 - \frac{k}{p^2} \log n - 4\Delta &= pn_2 - \left(\frac{k}{p^2} \log n + 2\Delta\right) - 2\Delta \\ &< pn_2 - pb - 2\Delta < p(n_2 - b) < 0 && (\hat{\lambda} > 1, n_2 \leq b) \\ &\leq q_{2,e}(1). \end{aligned}$$

□

Proof of Lemma 7: Following the discussion in the proof of Lemma 4, we assume, without loss of generality, $n_1 + n_2 \leq b$. We consider three cases in Lemma 6 separately.

If case (a) in Lemma 6 happens, then $ALG_1(\vec{v}) + n_1 \geq OPT(\vec{v})$ and $OPT(\vec{v}) \geq ab$. As a result,

$$\frac{ALG_1(\vec{v})}{OPT(\vec{v})} \geq 1 - \frac{n_1}{OPT(\vec{v})} \geq 1 - \frac{\frac{k}{p^2} \log n}{ab} \geq p + \frac{1-p}{2-a} - \frac{\frac{k}{p^2} \log n}{ab}.$$

If case (b) in Lemma 6 happens, then $ALG_1(\vec{v}) = OPT(\vec{v})$, and we are done.

If case (c) happens, then

$$ALG_1(\vec{v}) \geq n_1 + \left(pn_2 - \frac{k}{p^2} \log n - 4\Delta + \theta b\right) a.$$

Because $n_1 \geq \left(p + \frac{1-p}{2-a}\right) n_1$, $pn_2 + \theta b \geq (p + \theta)n_2 = \left(p + \frac{1-p}{2-a}\right) n_2$ and $OPT(\vec{v}) = n_1 + an_2 \geq a\theta b$, we have

$$\frac{ALG_1(\vec{v})}{OPT(\vec{v})} \geq p + \frac{1-p}{2-a} - \frac{a\left(\frac{k}{p^2} \log n + 4\Delta\right)}{OPT(\vec{v})} \geq p + \frac{1-p}{2-a} - \frac{\frac{k}{p^2} \log n + 4\Delta}{\theta b}. \quad (\text{EC.19})$$

□

Proof of Lemma 8: It is easy to check (a) $\frac{(1-a)\Delta}{ab} = O\left(\frac{1}{a(1-p)p} \sqrt{\frac{\log n}{b}}\right)$ and (b) $\frac{5\Delta}{\theta b} = O\left(\frac{1}{a(1-p)p} \sqrt{\frac{\log n}{b}}\right)$. To prove (c) $\frac{\frac{k}{p^2} \log n}{ab} = O\left(\frac{1}{a(1-p)p} \sqrt{\frac{\log n}{b}}\right)$, we first note that (5) and

$$\bar{\epsilon} \leq 1. \quad (\text{EC.20})$$

implies $\log n \leq a^2 p^2 b$, and thus $\log n = \sqrt{\log n} \sqrt{\log n} \leq ap \sqrt{b \log n}$. As a result,

$$\begin{aligned} \frac{\frac{k}{p^2} \log n}{ab} &\leq \frac{\frac{k}{p} \sqrt{b \log n}}{b} && (\log n \leq ap \sqrt{b \log n}) \\ &\leq \frac{k}{a(1-p)p} \sqrt{\frac{\log n}{b}} && (0 < p < 1 \text{ and } a < 1) \\ &= O\left(\frac{1}{a(1-p)p} \sqrt{\frac{\log n}{b}}\right). \end{aligned}$$

Similarly, to prove (d) $\frac{\frac{k}{p^2} \log n + 4\Delta}{\theta b} = O\left(\frac{1}{a(1-p)p} \sqrt{\frac{\log n}{b}}\right)$, we first note that (5) and (EC.20) implies $\log n \leq bp^2$, and thus $\log n = \sqrt{\log n} \sqrt{\log n} \leq p\sqrt{b \log n}$. As a result,

$$\begin{aligned} \frac{\frac{k}{p^2} \log n + 4\Delta}{\theta b} &\leq \frac{\frac{k}{p} \sqrt{b \log n} + 4\alpha \sqrt{b \log n}}{\theta b} && (\log n \leq p\sqrt{b \log n} \text{ and } \Delta = \alpha \sqrt{b \log n}) \\ &\leq \frac{k + 4\alpha}{pa} \sqrt{\frac{\log n}{b}} && (p < 1 \text{ and } \theta > a) \\ &\leq (k + 4\alpha) \left(\frac{1}{a(1-p)p} \sqrt{\frac{\log n}{b}} \right) \\ &= O\left(\frac{1}{a(1-p)p} \sqrt{\frac{\log n}{b}} \right). \end{aligned}$$

□

EC.3. Missing proofs of Section 5

Before proceeding with the proofs, we state and prove an auxiliary lemma that establishes an upper bound on n_1 and $n_1 + n_2$ using the deterministic approximation functions $\tilde{o}_j(\cdot)$.

LEMMA EC.6. *For $\lambda \in \{1/n, 2/n, \dots, 1\}$, we have:*

$$n_1 \leq \min \left\{ \frac{\tilde{o}_1(\lambda)}{\lambda p}, \frac{\tilde{o}_1(\lambda) + (1-\lambda)(1-p)n}{1-p+\lambda p} \right\}, \text{ and} \quad (\text{EC.21a})$$

$$n_1 + n_2 \leq \min \left\{ \frac{\tilde{o}_1(\lambda) + \tilde{o}_2(\lambda)}{\lambda p}, \frac{\tilde{o}_1(\lambda) + \tilde{o}_2(\lambda) + (1-\lambda)(1-p)n}{1-p+\lambda p} \right\}. \quad (\text{EC.21b})$$

The proof essentially follows from the definition of the deterministic approximation functions. Here we just prove (EC.21a). First note that $\eta_1(\lambda) \geq 0$, thus:

$$\tilde{o}_1(\lambda) = (1-p)\eta_1(\lambda) + p\lambda n_1 \geq p\lambda n_1 \Rightarrow n_1 \leq \frac{\tilde{o}_1(\lambda)}{\lambda p}.$$

Second note that $n_1 - \eta_1(\lambda) \leq (1-\lambda)n$, therefore,

$$\tilde{o}_1(\lambda) = (1-p)\eta_1(\lambda) + p\lambda n_1 \geq (1-p)(n_1 - (1-\lambda)n) + p\lambda n_1 \Rightarrow n_1 \leq \frac{\tilde{o}_1(\lambda) + (1-\lambda)(1-p)n}{1-p+\lambda p}.$$

Which completes the proof of (EC.21a). Proof of (EC.21b) follows similar steps. □

Proof of Lemma 9: Since Inequalities (11a) and (11b) are similar, we only present the proof of Inequality (11a). When $\lambda < \delta$, $u_1(\lambda) \triangleq b$, and thus (11a) trivially holds. The more interesting case is when $\lambda \geq \delta$. Without loss of generality, we assume $n_1 \leq b + \frac{2\Delta}{\delta p}$. Otherwise, similar to the proof of Lemma 2, we construct a modified adversarial instance with only $b + \frac{2\Delta}{\delta p}$ type-1 customers and argue that, for the same realization of the stochastic group and random permutation, $u_1(\lambda)$ is lower bounded by the one corresponding to the modified instance. Note that we can apply Inequality (2a) to this modified instance, because $b + \frac{2\Delta}{\delta p} \geq \frac{k}{p^2} \log n$ under the condition imposed on b .

Note that $\delta = \frac{\phi b}{n} = \frac{(1-c)b}{(1-a)n} \geq \frac{(1-c)b}{n}$, Condition imposed on b , and

$$\bar{\epsilon} \leq \frac{3}{2\alpha} \quad (\text{EC.22})$$

which holds by the definition of constants α and $\bar{\epsilon}$ given in Lemma 1, give us

$$\frac{\Delta}{\delta p} \leq \frac{3}{2}b. \quad (\text{EC.23})$$

Therefore, $n_1 \leq b + \frac{2\Delta}{\delta p} \leq 4b$, and thus Inequality (2a) implies $o_1(\lambda) \geq \tilde{o}_1(\lambda) - \alpha\sqrt{n_1 \log n} \geq \tilde{o}_1(\lambda) - \alpha\sqrt{4b \log n} = \tilde{o}_1(\lambda) - 2\Delta$. As a result,

$$\begin{aligned} u_1(\lambda) &\triangleq \min \left\{ \frac{o_1(\lambda)}{\lambda p}, \frac{o_1(\lambda) + (1-\lambda)(1-p)n}{1-p+\lambda p} \right\} \\ &\geq \min \left\{ \frac{\tilde{o}_1(\lambda) - 2\Delta}{\lambda p}, \frac{\tilde{o}_1(\lambda) - 2\Delta + (1-\lambda)(1-p)n}{1-p+\lambda p} \right\} \quad (o_1(\lambda) \geq \tilde{o}_1(\lambda) - 2\Delta) \\ &\geq n_1 - \max \left\{ \frac{2\Delta}{\lambda p}, \frac{2\Delta}{1-p+\lambda p} \right\} = n_1 - \frac{2\Delta}{\lambda p} \quad (\text{Lemma EC.6}) \\ &\geq n_1 - \frac{2\Delta}{\delta p}. \quad (\lambda \geq \delta) \end{aligned}$$

□

Proof of Proposition 1:

Case $b < n$: For proving this case, we can relax some of the constraints and only keep Constraints (15a), (15c), (15e), and (15f) and show that for every point in this superset of the feasibility region, $c \geq p + \frac{1-p}{2-a}$.

We first notice that, for fixed n_1 and n_2 , the right hand side of Constraint (15a) is non-increasing in each of \tilde{u}_1 and \tilde{o}_2 , and hence non-increasing in each of η_1 and η_2 . By using Constraints (15e) and (15f), we can obtain upper bounds $\tilde{o}_1 \leq (1-p+lp)n_1$ and $\tilde{o}_2 \leq (1-p+lp)n_2$. With these upper bounds and the fact $\tilde{u}_1 \leq \frac{\tilde{o}_1}{lp}$, Constraint (15a) gives:

$$c \geq \frac{a(n_2 - (1-p+pl)n_2 + \frac{b}{1-a}) + n_1}{a \min\{n_1 + n_2, b\} + (1-a)n_1 + \frac{a^2b}{1-a} + a \min \left\{ \frac{(1-p+pl)n_1}{lp}, b \right\}},$$

which, after rearranging terms, leads to

$$c \geq \frac{an_2p(1-l) + \frac{ab}{1-a} + n_1}{a \min\{n_2, b - n_1\} + n_1 + \frac{a^2b}{1-a} + a \min \left\{ \left(\frac{1-p}{lp} + 1 \right) n_1, b \right\}}. \quad (\text{EC.24})$$

We focus on lower bounding the right hand side of (EC.24). When $n_2 \geq b - n_1$, the right hand side of (EC.24) is non-decreasing in n_2 because the denominator remains the same while the numerator

is non-decreasing when n_2 increases (due to Constraint (15c)). Therefore, for the sake of obtaining a lower bound, we can assume, without loss of generality,

$$n_2 \leq b - n_1. \quad (\text{EC.25})$$

With (EC.25), the right hand side of (EC.24) can be written as

$$f_1(l) \triangleq \frac{an_2p(1-l) + \frac{ab}{1-a} + n_1}{an_2 + n_1 + \frac{a^2b}{1-a} + ab} \quad \text{if } l \leq \frac{(1-p)n_1}{p(b-n_1)}, \quad (\text{EC.26a})$$

$$f_2(l) \triangleq \frac{an_2p(1-l) + \frac{ab}{1-a} + n_1}{an_2 + n_1 + \frac{a^2b}{1-a} + a\left(\frac{1-p}{lp} + 1\right)n_1} \quad \text{if } l > \frac{(1-p)n_1}{p(b-n_1)}. \quad (\text{EC.26b})$$

We prove

$$f_1(l) \geq p + \frac{1-p}{2-a}, \quad \text{if } l \leq \frac{(1-p)n_1}{p(b-n_1)} \quad (\text{EC.27})$$

$$f_2(l) \geq p + \frac{1-p}{2-a}, \quad \text{if } l > \frac{(1-p)n_1}{p(b-n_1)} \quad (\text{EC.28})$$

separately. We start by the former: Because $f_1(l)$ is non-increasing in l , we only need to prove for the case $l = \frac{(1-p)n_1}{p(b-n_1)}$; $f_1(l)$ at $l = \frac{(1-p)n_1}{p(b-n_1)}$ can be rearranged as

$$f_1\left(\frac{(1-p)n_1}{p(b-n_1)}\right) = 1 - \frac{an_2\left(1-p + \frac{(1-p)n_1}{p(b-n_1)}p\right)}{an_2 + n_1 + \frac{ab}{1-a}},$$

which, for fixed n_1 and n_2 , is non-decreasing in b . Therefore, according to (EC.25), we only need to consider the case $b = n_1 + n_2$ (in the degenerated case $n_1 = n_2 = 0$, $f_1\left(\frac{(1-p)n_1}{p(b-n_1)}\right)$ is 1, which is greater than $p + \frac{1-p}{2-a}$, so we can assume, without loss of generality, $n_1 + n_2 > 0$), in which case, we have:

$$f_1\left(\frac{(1-p)n_1}{p(b-n_1)}\right) = 1 - \frac{a(1-p)(n_1+n_2)}{an_2 + n_1 + \frac{a(n_1+n_2)}{1-a}} \geq 1 - \frac{a(1-p)(n_1+n_2)}{an_2 + an_1 + \frac{a(n_1+n_2)}{1-a}} = 1 - \frac{1-p}{1 + \frac{1}{1-a}} = p + \frac{1-p}{2-a},$$

which completes the proof of (EC.27). Next we prove (EC.28). Due to Constraint (15c), i.e., $l \leq 1$, case (EC.26b) only happens when $\frac{(1-p)n_1}{p(b-n_1)} < 1$, or equivalently,

$$b > \frac{n_1}{p}. \quad (\text{EC.29})$$

Proving (EC.28) is trickier because both the numerator and denominator decrease in l . To address this issue, we first remark that by the definition of $f_2(l)$, inequality $f_2(l) \geq p + \frac{1-p}{2-a}$ is equivalent to

$$an_2p(1-l) + \frac{ab}{1-a} + n_1 \geq \left(p + \frac{1-p}{2-a}\right) \left(an_2 + n_1 + \frac{a^2b}{1-a} + a\left(\frac{1-p}{lp} + 1\right)n_1\right),$$

which is in turn equivalent to

$$an_2p + \frac{ab}{1-a} + n_1 \geq \left(p + \frac{1-p}{2-a}\right) \left(an_2 + n_1 + \frac{a^2b}{1-a} + a \left(\frac{1-p}{lp} + 1\right) n_1\right) + an_2pl. \quad (\text{EC.30})$$

Now the left hand side of Inequality (EC.30) is not a function of l while the right hand side of (EC.30) is a function of l of the form

$$xl + \frac{y}{l} + z \quad (\text{EC.31})$$

where x, y, z are non-negative. Clearly, the second derivative of (EC.31) (with respect to l), $\frac{2y}{l^3}$, is non-negative for $l \in \left[\frac{(1-p)n_1}{p(b-n_1)}, 1\right]$. As a result, (EC.31) is convex and is maximized at extreme values of l , which in our case is at either $l = \frac{(1-p)n_1}{p(b-n_1)}$ or $l = 1$. Therefore, we only need to prove Inequality (EC.30) at these extreme two values of l . The former case, $l = \frac{(1-p)n_1}{p(b-n_1)}$, is covered in (EC.27). Thus, we only need to prove (EC.30) for $l = 1$. When $l = 1$, (EC.30) can be rearranged as

$$\frac{ab}{1-a} + n_1 - \left(p + \frac{1-p}{2-a}\right) \left(an_2 + n_1 + \frac{a^2b}{1-a} + \frac{an_1}{p}\right) \geq 0. \quad (\text{EC.32})$$

Because the left hand side of Inequality (EC.32) is a decreasing function of n_2 , using (EC.25), we only need to prove that the inequality holds when $n_2 = b - n_1$, which is equivalent to

$$\frac{ab}{1-a} + n_1 - \left(p + \frac{1-p}{2-a}\right) \left(a(b-n_1) + n_1 + \frac{a^2b}{1-a} + \frac{an_1}{p}\right) \geq 0. \quad (\text{EC.33})$$

By separating terms associated with n_1 and b , and using

$$1 - \left(p + \frac{1-p}{2-a}\right) = \frac{(1-p)(1-a)}{2-a},$$

Inequality (EC.33) is equivalent to

$$\frac{a(1-p)}{2-a}b + \frac{(1-p)(1-a)}{2-a}n_1 - \left(p + \frac{1-p}{2-a}\right) \left(\frac{a(1-p)}{p}\right)n_1 \geq 0. \quad (\text{EC.34})$$

Using the lower bound on b given by (EC.29), i.e., $b > \frac{n_1}{p}$ and then dividing $(1-p)n_1$ on both sides (in the degenerated case where $(1-p)n_1 = 0$, both sides are 0 so we are done), Inequality (EC.34) is implied by

$$\frac{a}{(2-a)p} + \frac{1-a}{2-a} - \left(p + \frac{1-p}{2-a}\right) \left(\frac{a}{p}\right) \geq 0,$$

which holds because after canceling the two terms involving $1/p$, it is equivalent to

$$\frac{(1-a)^2}{2-a} \geq 0.$$

Therefore, we proved Inequality (EC.32), and thus (EC.28). This completes the proof of proposition for the case $b < n$.

Case $b = n$: For proving this case, we relax some of the constraints and only keep Constraints (15a), (15c), (15d) and (15h) and show that for every point in this superset of the feasibility region, $c \geq 1$. According to Constraint (15a) and using $a \min\{n_1 + n_2, b\} + (1 - a)n_1 = a \min\{n_2, b - n_1\} + n_1$, it suffices to prove

$$\frac{a(n_2 - \tilde{o}_2 + \frac{b}{1-a}) + n_1}{a \min\{n_2, b - n_1\} + n_1 + \frac{a^2b}{1-a} + a\tilde{u}_1} \geq 1. \quad (\text{EC.35})$$

or equivalently,

$$a(n_2 - \tilde{o}_2 + \frac{b}{1-a}) + n_1 \geq a \min\{n_2, b - n_1\} + n_1 + \frac{a^2b}{1-a} + a\tilde{u}_1. \quad (\text{EC.36})$$

Using $\min\{n_2, b - n_1\} \leq n_2$, subtracting $an_2 + n_1$ on both sides of (EC.36), and dividing both sides by a , Inequality (EC.36) is implied by

$$-\tilde{o}_2 + \frac{b}{1-a} \geq \frac{ab}{1-a} + \tilde{u}_1.$$

Subtracting $\frac{ab}{1-a}$ on both sides and using $\tilde{u}_1 \leq \frac{\tilde{o}_1 + (1-l)(1-p)n}{(1-p+lp)}$, the above inequality is implied by

$$b - \tilde{o}_2 \geq \frac{\tilde{o}_1 + (1-l)(1-p)n}{(1-p+lp)}.$$

Multiplying $1 - p + lp$ on both sides and using $b = n$ (which is the assumption of this case), the above inequality is equivalent to

$$ln - (1 - p + lp)\tilde{o}_2 \geq \tilde{o}_1.$$

Due to Constraint (15c), $1 - p + lp \leq 1$, and thus the inequality above is implied by

$$ln \geq \tilde{o}_2 + \tilde{o}_1,$$

or equivalently,

$$ln \geq (1 - p)\eta_2 + pn_2l + (1 - p)\eta_1 + pn_1l.$$

The above inequality follows straightforwardly from Constraints (15d) and (15h). This completes our proof of (EC.36), and consequently that of the proposition in the case $b = n$. \square

Proof of Lemma 10: The only interesting case is case (b), i.e., when $n_1 + n_2 > b + \frac{2\Delta}{\delta}$. If $q_2(1) = 0$, then we are done. Otherwise, let $\bar{\lambda}$ be the last time we accept a type-2 customer. By Lemma 9, $u_{1,2}(\bar{\lambda}) \geq \min \left\{ b, n_1 + n_2 - \frac{2\Delta}{\delta p} \right\} = b$. Therefore, according to the definition of $\bar{\lambda}$, Condition (14) must be satisfied. Thus,

$$\begin{aligned}
q_2(1) &= q_2(\bar{\lambda}) \\
&\leq \frac{1-c}{1-a}b + c(b - u_1(\bar{\lambda}))^+ + 1 && \text{(Condition (14))} \\
&\leq \frac{1-c}{1-a}b + c \left(b - \min \left\{ b, n_1 - \frac{2\Delta}{\delta p} \right\} \right)^+ + 1 && \text{(Lemma 9)} \\
&\leq \frac{1-c}{1-a}b + c(b - n_1)^+ + c\frac{2\Delta}{\delta p} + 1.
\end{aligned}$$

□

Proof of Lemma 11: We consider the two cases of Lemma 10 separately. For case (a), $n_1 + n_2 \leq b + \frac{2\Delta}{\delta p}$, we note that

$$\begin{aligned}
OPT(\vec{v}) &\leq n_1 + n_2 a \\
&\leq \left(b + \frac{2\Delta}{\delta p} - n_2 \right) + n_2 a && (n_1 + n_2 \leq b + \frac{2\Delta}{\delta p}) \\
&\leq ALG_{2,c}(\vec{v}) + \frac{2\Delta}{\delta p}. && (ALG_{2,c}(\vec{v}) \geq (b - n_2) + an_2)
\end{aligned}$$

Therefore,

$$\begin{aligned}
\frac{ALG_{2,c}(\vec{v})}{OPT(\vec{v})} &\geq \frac{ALG_{2,c}(\vec{v})}{ALG_{2,c}(\vec{v}) + \frac{2\Delta}{\delta p}} \geq \frac{ALG_{2,c}(\vec{v}) - \frac{2\Delta}{\delta p}}{ALG_{2,c}(\vec{v})} && ((ALG_{2,c}(\vec{v}))^2 \geq (ALG_{2,c}(\vec{v}))^2 - \left(\frac{2\Delta}{\delta p}\right)^2) \\
&= 1 - \frac{\frac{2\Delta}{\delta p}}{ALG_{2,c}(\vec{v})} \geq 1 - \frac{2\Delta}{ab\delta p} \geq c - \frac{2\Delta}{ab\delta p}. && (ALG_{2,c}(\vec{v}) \geq ab)
\end{aligned}$$

For case (b), $n_1 + n_2 > b + \frac{2\Delta}{\delta p}$ and $q_2(1) \leq \frac{1-c}{1-a}b + c(b - n_1)^+ + c\frac{2\Delta}{\delta p} + 1$, we have

$$\begin{aligned}
\frac{ALG_{2,c}(\vec{v})}{OPT(\vec{v})} &= \frac{b - (1-a)q_2(1)}{OPT(\vec{v})} && (q_1(1) + q_2(1) = b) \\
&\geq \frac{b - (1-a) \left(\frac{1-c}{1-a}b + c(b - n_1)^+ + c\frac{2\Delta}{\delta p} + 1 \right)}{OPT(\vec{v})} && (q_2(1) \leq \frac{1-c}{1-a}b + c(b - n_1)^+ + c\frac{2\Delta}{\delta p} + 1) \\
&= \frac{c(b - (1-a)(b - n_1)^+)}{OPT(\vec{v})} - \frac{(1-a)c\frac{2\Delta}{\delta p}}{OPT(\vec{v})} - \frac{1-a}{OPT(\vec{v})} \\
&\geq c - \frac{(1-a)c\frac{2\Delta}{\delta p}}{OPT(\vec{v})} - \frac{1-a}{OPT(\vec{v})} && (OPT(\vec{v}) \leq b - (1-a)(b - n_1)^+) \\
&\geq c - \frac{2(1-a)c\Delta}{ab\delta p} - \frac{1-a}{ab} && (OPT(\vec{v}) \geq ab) \\
&\geq c - \frac{3\Delta}{ab\delta p}. && (1-a < 1, c \leq 1, \delta \leq 1, p < 1)
\end{aligned}$$

□

Proof of Lemma 12: Let us define $\tilde{o}'_1, \tilde{o}'_2, \tilde{u}'_1$ and $\tilde{u}'_{1,2}$ to be the corresponding functions defined for the modified tuple $(l', n'_1, n'_2, \eta'_1, \eta'_2, c')$. It is easy to check that $(l', n'_1, n'_2, \eta'_1, \eta'_2, c')$ satisfies Constraints (15c)-(15i). The interesting part is to show that it satisfies Constraint (15b). When $n_1 + n_2 \geq b$, we can prove it directly from Lemma EC.6 (since $\tilde{u}'_{1,2} \geq n'_1 + n'_2 = n_1 + n_2 \geq b$).

Next we focus on the case $n_1 + n_2 < b$, and we prove $\tilde{u}'_{1,2} \geq b$ by showing $\tilde{u}'_{1,2} \geq u_{1,2}(l)$; note that we have $u_{1,2}(l) \geq b$ by Inequality (17).

Recall that we reject a customer at time l and that the threshold of rejecting a customer is at least ϕb ; thus we have $ln \geq o_2(l) \geq \phi b$. This gives

$$l \geq \frac{\phi b}{n} = \delta \quad (\text{EC.37})$$

Note that by definition for $l \geq \delta$, we have $u_{1,2}(l) = \min \left\{ \frac{o_1(l) + o_2(l)}{lp}, \frac{o_1(l) + o_2(l) + (1-l)(1-p)n}{1-p+lp} \right\}$, which is a non-decreasing function of $o_1(l) + o_2(l)$. Thus, $\tilde{u}'_{1,2} \geq u_{1,2}(l)$ is implied by $\tilde{o}'_1 + \tilde{o}'_2 \geq o_1(l) + o_2(l)$. We prove this by breaking down into two cases based on the value of ξ : Case (1) $\xi = n - (n_1 + n_2)$: we have $n'_1 + n'_2 = n$, i.e., there is no time period without a customer. Thus $\eta'_1 + \eta'_2 = ln$, and $\tilde{o}'_1 + \tilde{o}'_2 = ln \geq o_1(l) + o_2(l)$. Case (2) $\xi = \frac{\Delta n}{\phi bp}$, we have

$$\begin{aligned} \tilde{o}'_1 + \tilde{o}'_2 &= l'pn'_1 + (1-p)\eta'_1 + l'pn'_2 + (1-p)\eta'_2 \\ &\geq lpn_1 + (1-p)\eta_1 + lp(n_2 + \frac{\Delta n}{\phi bp}) + (1-p)\eta_2 && (\xi = \frac{\Delta n}{\phi bp}, \bar{\xi} \geq 0) \\ &= \tilde{o}_1(l) + \tilde{o}_2(l) + \frac{\Delta n}{\phi b}l \geq \tilde{o}_1(l) + \tilde{o}_2(l) + \Delta && ((\text{EC.37})) \\ &\geq o_1(l) + o_2(l) && ((2b)) \end{aligned}$$

□

Proof of Lemma 13: We first show that, for all $c \leq c^*$, Constraint (15a) (same as (22)) is either violated or holds with equality. First, we note that, for all real number x , the tuple $(l', n'_1, n'_2, \eta'_1, \eta'_2, x)$ satisfies Constraints (15b)-(15i) because those constraints are not related to the last element in the tuple. For all $x < c^*$, $(l', n'_1, n'_2, \eta'_1, \eta'_2, x)$ is not in the feasible set of (MP1), and hence Constraint (15a) is violated. Taking the limit $x \rightarrow c^*$, Constraint (15a) is either violated or hold with equality. This means, for $ALG_{2,c}$ (with any $c \leq c^*$),

$$c = c' \leq \frac{a(n'_2 - \tilde{o}'_2 + \frac{b}{1-a}) + n'_1}{a \min\{n'_1 + n'_2, b\} + (1-a)n'_1 + \frac{a^2 b}{1-a} + a \min\{\tilde{u}'_1, b\}}. \quad (\text{EC.38})$$

After rearranging terms (EC.38) is equivalent to

$$\frac{n'_1 + a \left(\frac{1-c}{1-a}b + c(b - \tilde{u}'_1)^+ + [n'_2 - \tilde{o}'_2] \right)}{n'_1 + a \min\{b - n'_1, n'_2\}} \geq c. \quad (n_1 \geq \frac{k}{p^2} \log n) \quad (\text{EC.39})$$

Repeating (20), recall that we have:

$$\frac{ALG_{2,c}(\vec{v})}{OPT(\vec{v})} \geq \frac{n_1 + a \left(\frac{1-c}{1-a}b + c(b - u_1(l))^+ + [n_2 - o_2(l)] \right)}{n_1 + a \min\{b - n_1, n_2\}}.$$

We want to compare the right hand side of (20) with the left hand side of (EC.39). First, we compare $(b - \tilde{u}'_1)^+$ with $(b - u_1(l))^+$ and show

$$(b - \tilde{u}'_1)^+ \leq (b - u_1(l))^+ + \xi. \quad (\text{EC.40})$$

Recall that we do not exhaust the inventory, and thus $n_1 < b$. Further $\Delta = \alpha\sqrt{b \log n}$, thus we have: $\Delta \geq \alpha\sqrt{n_1 \log n}$. According to (2a), $\delta'_1 = \tilde{o}_1(l) \geq o_1(l) - \Delta$. Combining this and using an argument similar to the proof of Lemma 9,

$$\begin{aligned} \tilde{u}'_1 &\triangleq \min \left\{ \frac{\delta'_1}{l'p}, \frac{\delta'_1 + (1-l')(1-p)n}{1-p+l'p} \right\} \\ &= \min \left\{ \frac{\tilde{o}_1(l)}{lp}, \frac{\tilde{o}_1(l) + (1-l)(1-p)n}{1-p+lp} \right\} \\ &\geq \min \left\{ \frac{o_1(l) - \Delta}{lp}, \frac{o_1(l) - \Delta + (1-l)(1-p)n}{1-p+lp} \right\} \quad (\delta'_1 \geq o_1(l) - \Delta) \\ &\geq u_1(l) - \max \left\{ \frac{\Delta}{lp}, \frac{\Delta}{1-p+lp} \right\} = u_1(l) - \frac{\Delta}{lp} \\ &\geq u_1(l) - \frac{\Delta n}{\phi bp}. \end{aligned} \quad ((\text{EC.37}))$$

Note that by the definition of ξ , the above inequality implies Inequality (EC.40). Next, we compare \tilde{o}'_2 with $o_2(l)$ and we show that

$$\tilde{o}'_2 \geq o_2(l) - 2\Delta. \quad (\text{EC.41})$$

In order to prove (EC.41), we first show that we can assume, without loss of generality, $\phi b < n_2 \leq b + \frac{2\Delta}{\delta p}$. To see this, we note that when $q_1(1) + q_2(1) < b$, $q_1(1) = n_1$. Therefore, the only ‘‘mistakes’’ that the algorithm may make is to reject too many type-2 customers. When $n_2 \leq \phi b$, we never reject a type-2 customer and so $q_2(1) = n_2$ and $ALG_{2,c}(\vec{v}) = OPT(\vec{v})$. For proving the upper bound on n_2 , i.e., $n_2 \leq b + \frac{2\Delta}{\delta p}$, we first note that, clearly, if $n_2 > b + \frac{2\Delta}{\delta p}$, decreasing n_2 to $b + \frac{2\Delta}{\delta p}$ (while fixing n_1) does not modify the optimal revenue $OPT(\vec{v})$. Using Lemma 9, we know that, when $n_2 \geq b + \frac{2\Delta}{\delta p}$, $u_{1,2}(l) \geq \min \left\{ b, n_1 + n_2 - \frac{2\Delta}{\delta p} \right\} = b$. Therefore, we accept a type-2 customer arriving at time l only if the number of type-2 customer accepted so far does not reach the dynamic threshold (i.e., the third rule in Algorithm 2) that depends only on $o_1(l)$ but not on $o_2(l)$. Given all the above, similar to the proof of Lemma 4, we can construct an alternative adversarial instance where

we reduce the number of type-2 customers to $b + \frac{2\Delta}{\delta p}$, and show that, for the same realization of the stochastic group and random permutation, the number of accepted type-2 customers in the alternative instance serves as a lower bound on its counterpart in the original instance.

Next we show that condition $n_2 \geq \frac{k \log n}{p^2}$ is satisfied which implies we can apply concentration result (3a) from Lemma 1. Because $n_2 > \phi b$, it suffices to show:

$$\phi b \geq \frac{k}{p^2} \log n. \quad (\text{EC.42})$$

Inequality (16) and

$$\bar{\epsilon} \leq \frac{1}{\sqrt{k}} \quad (\text{EC.43})$$

which holds by the definition of constants k and $\bar{\epsilon}$ given in Lemma 1, implies

$$\sqrt{b} = \frac{b^{\frac{3}{2}}}{b} \geq \frac{b^{\frac{3}{2}}}{n} > \frac{1}{\bar{\epsilon}} \frac{\sqrt{\log n}}{(1-c)^2 a^2 p^{3/2}} \geq \sqrt{k} \frac{\sqrt{\log n}}{p\sqrt{1-c}}.$$

This, together with $\phi = \frac{1-c}{1-a} \geq 1-c$, proves (EC.42). Thus, we can apply (3a). Further note that (EC.23) implies that $n_2 \leq b + \frac{2\Delta}{\delta p} \leq 4b$. Finally note that by definition $\xi \geq 0$, $\xi' \geq 0$. Putting all these together, we have:

$$\tilde{o}'_2 \geq \tilde{o}_2(l) \geq o_2(l) - \alpha \sqrt{n_2 \log n} \geq o_2(l) - \alpha \sqrt{4b \log n} = o_2(l) - 2\Delta.$$

This proves (EC.41). Having proved (EC.40) and (EC.41), at last, we complete the proof as follows:

$$\begin{aligned} c &\leq \frac{n'_1 + a \left(\frac{1-c}{1-a} b + c(b - \tilde{u}'_1)^+ + [n'_2 - \tilde{o}'_2] \right)}{n'_1 + a \min\{b - n'_1, n'_2\}} && (\text{EC.39}) \\ &\leq \frac{n_1 + a \left(\frac{1-c}{1-a} b + c \left[(b - u_1(l))^+ + \frac{\Delta n}{\phi b p} \right] + [n_2 + \xi - o_2(l) + 2\Delta] \right)}{n_1 + a \min\{b - n_1, n_2\}} && ((\text{EC.40}), (\text{EC.41}), n'_2 = n_2 + \xi \geq n_2) \\ &\leq \frac{ALG_{2,c}(\vec{v})}{OPT(\vec{v})} + \frac{a \left(c \frac{\Delta n}{\phi b p} + \xi + 2\Delta \right)}{OPT(\vec{v})} && ((20)) \\ &\leq \frac{ALG_{2,c}(\vec{v})}{OPT(\vec{v})} + \frac{4a\Delta n}{a\phi^2 b^2 p} = \frac{ALG_{2,c}(\vec{v})}{OPT(\vec{v})} + \frac{4\Delta n}{\phi^2 b^2 p} && (n_2 > \phi b, \Delta \leq \frac{\Delta n}{\phi b p}, \xi \leq \frac{\Delta n}{\phi b p}, OPT \geq a\phi b). \end{aligned}$$

□

Proof of Lemma 14: Note that if we are not in case (a), i.e., $q_1(1) + q_2(1) < b$, then $q_1(1) = n_1$. Now either $q_2(1) = n_2$, i.e., we are in case (b), or $q_2(1) < n_2$. Therefore, what is remaining is to show that if $q_1(1) + q_2(1) < b$ and $q_2(1) < n_2$, then $q_2(1) \geq cb$, i.e., we are in case (c).

Let $\bar{\lambda}$ be the last time that a customer is rejected. Then, similar to earlier discussion, Inequality (EC.37) is satisfied. Therefore,

$$\begin{aligned} u_1(\bar{\lambda}) &= \min \left\{ \frac{o_1(\bar{\lambda})}{\bar{\lambda}p}, \frac{o_1(\bar{\lambda}) + (1 - \bar{\lambda})(1 - p)n}{1 - p + \bar{\lambda}p} \right\} & (\bar{\lambda} \geq \delta) \\ &\leq \frac{o_1(\bar{\lambda})}{\bar{\lambda}p} \\ &\leq \frac{n_1 n}{\phi b p} & ((\text{EC.37}) \text{ and } o_1(\bar{\lambda}) \leq n_1) \\ &< \frac{n \frac{k}{p^2} \log n}{\phi b p} = \frac{kn \log n}{\phi b p^3}. & (n_1 < \frac{k}{p^2} \log n) \end{aligned}$$

As a result, we have:

$$q_2(1) \geq \phi b + c(b - u_1(\bar{\lambda}))^+ > (\phi + c)b - c \frac{kn \log n}{\phi b p^3},$$

In order to complete the proof, it suffices to show that

$$q_2(1) > (\phi + c)b - c \frac{kn \log n}{\phi b p^3} \geq cb, \quad (\text{EC.44})$$

The last inequality in (EC.44) holds if $b^2 \geq c \frac{kn \log n}{\phi^2 p^3}$. Thus in the following, we show $b^2 \geq c \frac{kn \log n}{\phi^2 p^3}$: Using $\phi = \frac{1-c}{1-a} \geq 1 - c$, Inequality (16), and

$$\bar{\epsilon} \leq \frac{1}{\sqrt[4]{k}}, \quad (\text{EC.45})$$

which holds by the definitions of constants k and $\bar{\epsilon}$ given in Lemma 1, we have

$$b^2 = \frac{\left(b^{\frac{3}{2}}\right)^2}{b} \geq \frac{\left(b^{\frac{3}{2}}\right)^2}{n} > \frac{1}{\bar{\epsilon}^4 (1-c)^4 a^2 p^3} \geq c \frac{kn \log n}{\phi^2 p^3}.$$

This proves $b^2 \geq c \frac{kn \log n}{\phi^2 p^3}$, and thus $q_2(1) \geq cb$. This completes the proof of the lemma. \square

Proof of Lemma 15: We consider each case of Lemma 14 separately. For case (a), $q_1(1) + q_2(1) = b$, since $n_1 < \frac{k}{p^2} \log n$, it is easy to see that

$$\frac{ALG_{2,c}(\vec{v})}{OPT(\vec{v})} \geq \frac{ab}{ab + \frac{k}{p^2} \log n} \geq \frac{ab - \frac{k}{p^2} \log n}{ab} = 1 - \frac{k \log n}{ab p^2},$$

which is at least c if $b \geq \frac{k \log n}{a(1-c)p^2}$. Inequality (16) and (EC.43) imply

$$\sqrt{b} = \frac{b^{\frac{3}{2}}}{b} \geq \frac{b^{\frac{3}{2}}}{n} > \frac{1}{\bar{\epsilon}} \frac{\sqrt{\log n}}{(1-c)^2 a^2 p^{3/2}} \geq \sqrt{k} \frac{\sqrt{\log n}}{p \sqrt{a(1-c)}},$$

and thus $b \geq \frac{k \log n}{a(1-c)p^2}$; therefore we have: $\frac{ALG_{2,c}(\vec{v})}{OPT(\vec{v})} \geq c$.

In cases (b) and (c), $q_2(1) \geq \min\{n_2, cb\}$; thus we have

$$ALG_{2,c}(\vec{v}) \geq n_1 + c(\min\{n_2, b\})a \geq c(n_1 + \min\{n_2, b\}a) \geq cOPT(\vec{v}).$$

\square

Proof of Lemma 16: Both follow from definition. \square

Proof of Corollary 1: Theorem 2 with $c = 1 - \sqrt[3]{\frac{1}{ap^{3/2}} \sqrt{\frac{n^2 \log n}{b^3}}}$ proves the corollary. \square

EC.4. Missing proofs of Section 6

Proof of Proposition 2: We prove that the competitive ratio of any online algorithm is at most $p + \frac{1-p}{2-a} + 3 \left(\frac{pb^2}{n} \right)$. Note that when $\frac{pb^2}{n} > 1/2$, $p + \frac{1-p}{2-a} + 3 \left(\frac{pb^2}{n} \right)$ is greater than 1 and hence the upper bound trivially holds. Thus in the following, we assume, without loss of generality, $\frac{pb^2}{n} \leq 1/2$.

We consider two adversarial instances \vec{v}_I and \vec{w}_I defined as

$$v_{I,j} = \begin{cases} a, & 1 \leq j \leq b, \\ 0, & b < j \leq 2b, \\ 0, & j > 2b. \end{cases} \quad w_{I,j} = \begin{cases} a, & 1 \leq j \leq b, \\ 1, & b < j \leq 2b, \\ 0, & j > 2b. \end{cases}$$

Let \mathcal{U} denote the event in which in the arrival sequence, none of the first b arrivals belongs to positions $[b+1, 2b]$ in the *initial* customer sequence, i.e., for all $i \in [b]$, we have: $i \notin \mathcal{S}$ or $\sigma_S^{-1}(i) \notin [b+1, 2b]$, where we use the following definition: For $x, y \in \mathbb{N}$ and $x < y$, $[x, y] \triangleq \{x, x+1, \dots, y\}$. Further, we define $[y] \triangleq [1, y]$. Note that under event \mathcal{U} , no online algorithms can distinguish whether the initial sequence is \vec{v}_I or \vec{w}_I up to time b/n . We first compute the probability of event \mathcal{U} as follows:

$$\begin{aligned} \mathbb{P}(\mathcal{U}) &= \mathbb{P}(\text{for all } i \in [b] : i \notin \mathcal{S} \text{ or } \sigma_S^{-1}(i) \notin [b+1, 2b]) \\ &\geq 1 - \sum_{i \in [b]} \mathbb{P}(i \in \mathcal{S} \text{ and } \sigma_S^{-1}(i) \in [b+1, 2b]) && \text{(Union bound)} \\ &\geq 1 - \frac{pb^2}{n}, && \text{(EC.46)} \end{aligned}$$

where the last inequality holds because of the following inequality (which we prove next): For all $i \neq j$,

$$\mathbb{P}(i \in \mathcal{S} \text{ and } \sigma_S^{-1}(i) = j) \leq \frac{p}{n}. \quad \text{(EC.47)}$$

To prove (EC.47), we first note for any i , we have $p = \mathbb{P}(i \in \mathcal{S}) = \sum_{j=1}^n \mathbb{P}(i \in \mathcal{S} \text{ and } \sigma_S^{-1}(i) = j)$. Second, denoting R the random variable corresponding to the size of the stochastic group, we have $\mathbb{P}(\sigma_S^{-1}(i) = i | i \in \mathcal{S}, R) = \frac{1}{R} \geq \frac{1}{n}$, and thus $\mathbb{P}(\sigma_S^{-1}(i) = i | i \in \mathcal{S}) \geq \frac{1}{n}$. Therefore,

$$\begin{aligned} \sum_{j \neq i} \mathbb{P}(i \in \mathcal{S} \text{ and } \sigma_S^{-1}(i) = j) &= p - \mathbb{P}(i \in \mathcal{S} \text{ and } \sigma_S^{-1}(i) = i) \\ &= p - \mathbb{P}(\sigma_S^{-1}(i) = i | i \in \mathcal{S}) \mathbb{P}(i \in \mathcal{S}) \leq p - \frac{p}{n} = \frac{(n-1)p}{n}. \end{aligned}$$

By symmetry, for each $j \neq i$, $\mathbb{P}(i \in \mathcal{S} \text{ and } \sigma_S^{-1}(i) = j) \leq \frac{p}{n}$, which proves (EC.47). This completes our proof of inequality (EC.46).

Under the event \mathcal{U} , in both problem instances, the revenue of each customer accepted up to time b/n is a . Conditioned on event \mathcal{U} , we denote q_2 the expected number of accepted type-2 customers

up to time b/n under either problem instances—recall that under event \mathcal{U} , up to time b/n , the online algorithm cannot distinguish the two.

We now proceed to find an upper bound on the expected revenue of any online algorithm under the two problem instances. We start by \vec{w}_I :

$$\begin{aligned} \mathbb{E} \left[\text{ALG}(\vec{W}) \right] &\leq \mathbb{E} \left[\text{ALG}(\vec{W}) \mid \mathcal{U} \right] \mathbb{P}(\mathcal{U}) + \text{OPT}(\vec{w}_I) (1 - \mathbb{P}(\mathcal{U})) \\ &\leq q_2 a + (b - q_2) + \frac{pb^2}{n} \text{OPT}(\vec{w}_I). \end{aligned}$$

Next, we proceed to establish an upper bound on the expected revenue under \vec{v}_I , by proving an upper bound on the number of type-2 customers that arrive after time b/n conditioned on the event \mathcal{U} :

$$\begin{aligned} \mathbb{E} [\{i \geq b+1 \mid V_i = a\} \mid \mathcal{U}] &= \sum_{i=b+1}^n \mathbb{P}(V_i = a \mid \mathcal{U}) \\ &\leq \frac{\sum_{i=b+1}^n \mathbb{P}(i \in \mathcal{S} \text{ and } \sigma_S^{-1}(i) \in [b])}{\mathbb{P}(\mathcal{U})} \\ &\leq \frac{(n-b)b \frac{p}{n}}{1 - \frac{pb^2}{n}} \quad (\text{Inequalities (EC.46), (EC.47)}) \\ &\leq \frac{pb}{1 - \frac{pb^2}{n}} \leq pb \left(1 + 2 \left(\frac{pb^2}{n} \right) \right), \end{aligned}$$

where we use $\frac{pb^2}{n} \leq 1/2$ in the last inequality. Note that $\mathbb{E} \left[\text{ALG}(\vec{V}) \mid \mathcal{U} \right] \leq q_2 a + a \mathbb{E} [\{i \geq b+1 \mid V_i = a\} \mid \mathcal{U}]$. As a result,

$$\begin{aligned} \mathbb{E} \left[\text{ALG}(\vec{V}) \right] &\leq \mathbb{E} \left[\text{ALG}(\vec{V}) \mid \mathcal{U} \right] \mathbb{P}(\mathcal{U}) + \text{OPT}(\vec{v}_I) (1 - \mathbb{P}(\mathcal{U})) \\ &\leq q_2 a + \left(1 + 2 \left(\frac{pb^2}{n} \right) \right) pba + \left(\frac{pb^2}{n} \right) \text{OPT}(\vec{v}_I) \\ &\leq q_2 a + pba + 3 \left(\frac{pb^2}{n} \right) \text{OPT}(\vec{v}_I). \quad (\text{OPT}(\vec{v}_I) = ba \geq pba) \end{aligned}$$

Thus, the competitive ratio is at most

$$\begin{aligned} \min \left\{ \frac{\mathbb{E} \left[\text{ALG}(\vec{V}) \right]}{\text{OPT}(\vec{v}_I)}, \frac{\mathbb{E} \left[\text{ALG}(\vec{W}) \right]}{\text{OPT}(\vec{w}_I)} \right\} &\leq \min \left\{ \frac{q_2}{b} + p + 3 \left(\frac{pb^2}{n} \right), \frac{q_2}{b} a + \left(1 - \frac{q_2}{b} + \frac{pb^2}{n} \right) \right\} \\ &\leq \min \left\{ \frac{q_2}{b} + p, \frac{q_2}{b} a + \left(1 - \frac{q_2}{b} \right) \right\} + 3 \left(\frac{pb^2}{n} \right) \\ &\leq p + \frac{1-p}{2-a} + 3 \left(\frac{pb^2}{n} \right), \end{aligned}$$

where the last inequality holds because function $g(q_2) \triangleq \min \left\{ \frac{q_2}{b} + p, \frac{q_2}{b} a + \left(1 - \frac{q_2}{b} \right) \right\}$ —defined on $q_2 \in [0, b]$ —achieves its maximum at $q_2^* = \frac{1-p}{2-a} b$, and $g(q_2^*) \leq p + \frac{1-p}{2-a}$. \square

EC.5. Missing proofs of Section 7

Proof of Theorem 3: For each integer $k \geq 2$, we denote \mathcal{F}_k the event satisfying all the following conditions:

1. All of the k highest-revenue customers are in the stochastic group.
2. The k^{th} -highest-revenue customer arrives in the observation period and the $k - 1$ customers with the highest revenue do not.
3. The highest-revenue customer arrives first among the $k - 1$ customers with the highest revenue.

Clearly, for any $k \geq 2$, \mathcal{F}_k is a success event and those events are mutually exclusive for different values of k . Furthermore, we note that, conditioned on being in the stochastic group, the probability that a customer arrives before time γ approaches γ as $n \rightarrow \infty$. This can be done by using a concentration result for random permutations similar to the one used in the proof of (3b). Further, for two customers l and \tilde{l} , conditioned on being in the stochastic group, the events that customer l arrives before time γ and customer \tilde{l} arrives after time γ are asymptotically independent. Thus, we can write the probability of event \mathcal{F}_k as:

$$\mathbb{P}(\mathcal{F}_k) = p^k \gamma (1 - \gamma)^{k-1} \frac{1}{k-1} + o(1),$$

where $o(1)$ represents a small (positive or negative) real number approaching zero as $n \rightarrow \infty$. As a result, for any fixed m , we have

$$\begin{aligned} \mathbb{P}(\text{success}) &\geq \sum_{k=2}^n \mathbb{P}(\mathcal{F}_k) \geq \sum_{k=2}^m \mathbb{P}(\mathcal{F}_k) \\ &\geq \sum_{k=2}^m \left(p^k \gamma (1 - \gamma)^{k-1} \frac{1}{k-1} - |o(1)| \right) \\ &\geq \sum_{k=2}^m \left(p^k \gamma (1 - \gamma)^{k-1} \frac{1}{k-1} \right) - m|o(1)|, \end{aligned}$$

which approaches $\sum_{k=2}^m \left(p^k \gamma (1 - \gamma)^{k-1} \frac{1}{k-1} \right)$, for fixed m , as $n \rightarrow \infty$. Since the above inequality holds for all m , we have

$$\mathbb{P}(\text{success}) \geq \lim_{n \rightarrow \infty} \sum_{k=2}^n \mathbb{P}(\mathcal{F}_k) \geq \lim_{m \rightarrow \infty} \sum_{k=2}^m p^k \gamma (1 - \gamma)^{k-1} \frac{1}{k-1} = \gamma p \log \frac{1}{\gamma p + 1 - p}. \quad (\text{EC.48})$$

Next, we show that the lower bound given in (EC.48) is tight by presenting an instance for which OSA_γ achieves a success probability of at most $\gamma p \log \frac{1}{\gamma p + 1 - p}$. Consider the following adversarial instance. The highest-revenue customer is the first customer in \vec{v}_I (As a reminder, the subscript I indicates that this is the initial sequence determined by the adversary.) For each $k = 2, 3, \dots, (1 - \gamma)n + 1$, the k^{th} -highest-revenue customer is the $(\gamma n + k - 1)^{\text{th}}$ customer in \vec{v}_I . Other customers arbitrarily fill other positions in \vec{v}_I .

For each positive integer $1 \leq l \leq (1 - \gamma)n$, we denote \mathcal{H}_l the event where all of the l highest-revenue customers are in the stochastic group but the $(l + 1)^{\text{th}}$ -highest is not. Similarly, we denote $\mathcal{H}_{(1-\gamma)n+1}$ the event where all of the $((1 - \gamma)n + 1)$ highest-revenue customers are in the stochastic group; finally, we denote \mathcal{H}_0 the event where the highest-revenue customer is in the adversarial group. Clearly, $\{\mathcal{H}_l\}_{l=0}^{(1-\gamma)n+1}$ is a partition of the sample space. In addition, for all l , $\mathbb{P}(\mathcal{H}_l) \leq p^l$. Conditioned on \mathcal{H}_0 , the highest-revenue customer arrives during the observation period, and thus the algorithm has a success probability of 0. Conditioned on \mathcal{H}_1 , the algorithm either accepts the customer with the second-highest value or does not accept any customer (if the highest-revenue customer arrives before γ), and hence again it has a success probability of 0. Further, for any $2 \leq l \leq (1 - \gamma)n$, conditioned on \mathcal{H}_l , to have a success, either one of the events $\mathcal{F}_2, \mathcal{F}_3, \dots, \mathcal{F}_l$ occurs or the highest-revenue customer arrives between time $(\gamma n + 1)$ and $(\gamma n + l - 1)$; note that conditioned on \mathcal{H}_l , the latter has a probability of $\frac{l-1}{n}$. As a result, the total success probability is at most

$$\begin{aligned}
& \sum_{l=2}^{(1-\gamma)n} \mathbb{P}(\mathcal{H}_l) \left[\mathbb{P} \left(\bigcup_{m=2}^l \mathcal{F}_m \mid \mathcal{H}_l \right) + \frac{l-1}{n} \right] + \mathbb{P}(\mathcal{H}_{(1-\gamma)n+1}) \\
& \leq \sum_{l=2}^{(1-\gamma)n} \sum_{m=2}^l \mathbb{P}(\mathcal{F}_m \cap \mathcal{H}_l) + \sum_{l=2}^{(1-\gamma)n} p^l \frac{l-1}{n} + p^{(1-\gamma)n+1} \quad (\mathbb{P}(\mathcal{H}_l) \leq p^l) \\
& \leq \sum_{m=2}^{\infty} \mathbb{P}(\mathcal{F}_m) + \sum_{l=2}^{\infty} p^l \frac{l-1}{n} + p^{(1-\gamma)n+1} \\
& = \sum_{l=2}^{\infty} \mathbb{P}(\mathcal{F}_l) + \frac{p^2}{(1-p)^2 n} + p^{(1-\gamma)n+1},
\end{aligned}$$

which converges to $\sum_{l=2}^{\infty} \mathbb{P}(\mathcal{F}_l)$ as n approaches infinity. \square

Proof of Proposition 3: The key in the proof of the proposition is that when the position of the second-highest-revenue customer in \vec{v}_I is before $\gamma_2 n$, OSA_{γ_2} has a success probability greater than s_2 ; otherwise, OSA_{γ_1} has a success probability greater than s_1 . To formalize this idea, we introduce two lemmas.

LEMMA EC.7. *If the second-highest-revenue customer is among the first $\gamma_2 n$ customers in \vec{v}_I , then OSA_{γ_2} has a success probability of at least $s_2 + p(1-p)(1-\gamma_2)$, when $n \rightarrow \infty$.*

Proof: Note that the events $\{\mathcal{F}_k\}_{k=2}^{\infty}$ defined in the proof of Theorem 3 collectively give an asymptotic success probability of s_2 . We identify another disjoint event which also results in a success. In particular, we define event $\bar{\mathcal{F}}$ that satisfies the following conditions:

1. The highest-revenue customer is in the stochastic group and arrives after time γ_2 .
2. The second-highest-revenue customer is in the adversarial group.

Note that $\bar{\mathcal{F}}$ is a success event that is disjoint from $\{\mathcal{F}_k\}_{k=2}^{\infty}$. Therefore, $\bar{\mathcal{F}}$ gives an additional success probability of $p(1-p)(1-\gamma_2) + o(1)$. \square

LEMMA EC.8. *If the second-highest-revenue customer is not among the first $\gamma_2 n$ customers in \vec{v}_I , then OSA_{γ_1} has a success probability of at least $s_1 + (1-p)\frac{\gamma_2-\gamma_1}{1-\gamma_1}s_1$, when $n \rightarrow \infty$.*

Proof: Similar to the previous lemma, we first note that the events $\{\mathcal{F}_k\}_{k=2}^\infty$ introduced in the proof of Theorem 3 collectively give an asymptotic success probability of s_1 . We identify another set of disjoint events that also results in success. In particular, for each positive integer $k \geq 2$, we define the event $\widehat{\mathcal{F}}_k$ that satisfies all of the following conditions:

1. Among the k highest-revenue customers, all except the second highest-revenue customer are in the stochastic group and arrive after time γ_1 .
2. The second-highest-revenue customer is in the adversarial group.
3. The $(k+1)^{\text{th}}$ -highest-revenue customer is in the stochastic group and arrives before time γ_1 .
4. The highest-revenue customer arrives no later than γ_2 and arrives first among the k highest-revenue customers (except for the second-highest-revenue customer).

Clearly, for any $k \geq 2$, $\widehat{\mathcal{F}}_k$ is a success event. In addition, those events are mutually exclusive for different values of k . Furthermore, $\{\widehat{\mathcal{F}}_k\}_{k=2}^\infty$ does not overlap $\{\mathcal{F}_k\}_{k \geq 2}$. Note that the probability that the highest-revenue customer arrives between time γ_1 and γ_2 , and it arrives first among the k highest-revenue customers (except for the second-highest-revenue customer) is at least $\frac{\gamma_2-\gamma_1}{k-1} + o(1)$. Therefore, $\mathbb{P}\left(\widehat{\mathcal{F}}_k\right) \geq p^k(1-p)\gamma_1(1-\gamma_1)^{k-2}(\gamma_2-\gamma_1)\frac{1}{k-1} + o(1)$. As a result, the asymptotic success probability is at least

$$\begin{aligned} s_1 + \sum_{k=2}^{\infty} \mathbb{P}\left(\widehat{\mathcal{F}}_k\right) &\geq s_1 + \sum_{k=2}^{\infty} p^k(1-p)\gamma_1(1-\gamma_1)^{k-2}(\gamma_2-\gamma_1)\frac{1}{k-1} \\ &= s_1 + (1-p)\frac{\gamma_2-\gamma_1}{1-\gamma_1}s_1. \end{aligned}$$

□

With Lemmas EC.7 and EC.8, we complete the proof of the proposition as follows: Recall that for any position of the second-highest-revenue customer in \vec{v}_I , OSA_{γ_1} has a success probability of at least s_1 and OSA_{γ_2} has a success probability of at least s_2 . As a result, using Lemma EC.7, if the second-highest-revenue customer is among the first $\gamma_2 n$ customers in \vec{v}_I , then we have a success probability of at least $qs_1 + (1-q)(s_2 + p(1-p)(1-\gamma_2))$. Similarly, using Lemma EC.8, if the second-highest-revenue customer is not among the first $\gamma_2 n$ customers in \vec{v}_I , then we have a success probability of at least $q(s_1 + (1-p)\frac{\gamma_2-\gamma_1}{1-\gamma_1}s_1) + (1-q)s_2$. As a result, for any adversarial problem instance \vec{v}_I , the asymptotic success probability is at least

$$\begin{aligned} &\min \left\{ qs_1 + (1-q)(s_2 + p(1-p)(1-\gamma_2)), q(s_1 + (1-p)\frac{\gamma_2-\gamma_1}{1-\gamma_1}s_1) + (1-q)s_2 \right\} \\ &= qs_1 + (1-q)s_2 + \min \left\{ (1-q)p(1-p)(1-\gamma_2), q(1-p)\frac{\gamma_2-\gamma_1}{1-\gamma_1}s_1 \right\}, \end{aligned}$$

which completes the proof. □