Appendices

Appendix A: Limitations of Balance and Primal-dual for Stochastic Usage Durations

Recall that in case of non-reusable resources, the Balance algorithm combined with primal-dual analysis leads to the best possible $(1-1/e)$ guarantee in a variety of settings. Through simple examples we now demonstrate some of the challenges with applying these ideas to the more general case of reusable resources. These examples also illustrate the ability of our new algorithm and analysis approach to address uncertainty in reusability.

A.1. Performance of Balance for Two-point usage distributions

We consider settings where for every resource $i \in I$, the usage durations are stochastic but take only one of two values - a value $d_i > 0$ or $+\infty$. This gives us a simple and natural generalization of the setting of non-reusable resources as well as the setting of deterministic usage durations. We show that Balance has a competitive ratio strictly less than $(1-1/e)$ in this setting. This indicates that quantities such as the time interval between arrivals and the probability of matched units returning before the next arrival may play an important role.

We start with a simple example that demonstrates the difference between Balance and RBA.

Example A.1 Consider an instance with two resources, labeled 1 and 2. Resources have the same reward and a large capacity n. We use a two-point usage distribution for both resources, with support $\{1, +\infty\}$ and probability of return 0.5. Consider a sequence of $4n$ arrivals as follows: At time 0 we have a burst of $2n$ arrivals all within a very short amount of time $\epsilon \to 0$. These arrivals can only be matched to resource 1. This is followed by n regularly spaced arrivals at time epochs $\{2, 4, \dots, 2n\}$, each of which can be matched to either resource. Finally, we have a burst of n arrivals that can only be matched to resource 2, at time $2n+2$. For large n, clairvoyant can match $\sim 3n$ arrivals with high probability (w.h.p.).

Comparing Balance with RBA: Balance will match to the resource with highest fraction of remaining capacity. So it matches the first n arrivals $t \in [n]$ (half of the first burst), to resource 1. With high probability (w.h.p.), nearly $n/2$ units of resource 1 return by time 2 and the rest never return. W.h.p., Balance matches most of n the arrivals $t \in [2n]\backslash[n]$ occurring between time 2 and $2n$ to resource 2 and thus, can only match half of the final burst of n arrivals $t \in [3n]\backslash[2n]$ at time $2n+2$. Balance fails to recognize that due to reusability the $n/2$ remaining units of resource 1 could all be matched to the second set of n arrivals and, in this way, the "effective" remaining capacity of resource 1 is n.

Consider the decisions of RBA in this instance. RBA coincides with Balance over the first $2n$ arrivals. After the first burst, w.h.p., the highest available unit of resource 1 in RBA is no lower than $n - O(\log n)$. In fact, w.h.p., RBA manages to match a constant fraction ($\sim 1/3$) of the n spread out arrivals to resource 1, successfully gauging "effective" inventory of the resource. Overall, RBA matches $\sim n/6$ more arrivals than Balance. In general, the metric $z_i(t)$ is very sensitive to reusability, i.e., it tends to have a high value when arrivals are spaced out and units return "frequently enough", and when this is not the case it acts closer to Balance and protects resources with low inventory.

What about fluid reusability? Does the framework of turning realization dependent algorithms into realization independent ones by means of using a fluid guide (see Section 3), addresses the above issue with Balance? Notice that even when we consider the fluid versions of usage distributions in Example A.1, the actions of Balance do not change. In fact, by modifying Example A.1, we establish an upper bound of 0.626 on the competitive ratio of Balance.

In Example A.1, let r_t denote the reduced price computed by Balance for the resource matched to arrival $t \in [4n]$.

Example A.2 We consider the instance in Example A.1 and augment it as follows. For every t, we have a new low reward non-reusable resource i_t with price $\max\{0, \frac{r_t}{(1-1/e)} - \delta\}$, for some small $\delta > 0$. Resource i_t does not have an edge to any arrival except t. For every t, let i_t have large capacity and usage duration of $+\infty$. We also augment the arrival sequence in Example A.1 with additional $n/2$ arrivals at the end (time $2n+2$), making a total of $1.5n$ arrivals in the final burst and $4.5n$ arrivals in total. The final $n/2$ arrivals have an edge to resource 2 only.

THEOREM 1 (RESTATED) *Consider the family of instances given in Example A.2. Suppose we allow fractional matching and experience fluid versions of usage distributions such that, if an ϵ amount of resource i is matched to t then 0.5ϵ returns at time $t+1$ and the other half never returns. Then, the total revenue of fractional Balance algorithm is $< 0.626\,\text{OPT}$ ($< (1-1/e)\,\text{OPT}$).*

Proof. Notice that, the matching output by Balance does not change with the addition of the low reward resources. Further, Balance does not match the additional $n/2$ arrivals added in the final burst. Therefore, the total reward of Balance is $2.5n$. We now compute a lower bound on the clairvoyant. Let OPT match the first n arrivals (from 1 to n) to the low reward resources. This gives a reward of at least

$$\frac{\sum_{i=1}^{n} 1 - e^{-i/n}}{1-1/e} - O(n\delta),$$

here $\frac{1-e^{-t/n}}{1-1/e} - \delta$ is the revenue from matching arrival t to i_t. Let OPT match the next $2n$ arrivals (from $n+1$ to $3n$) to resource 1 (as much as possible). Using concentration bounds for large n, we have that OPT matches $2n - o(n)$ of the arrivals to resource 1 w.h.p.. In the final burst of arrivals, OPT matches first $n/2$ arrivals (from $3n+1$ to $3.5n$) to low reward resources and final n arrivals (from $3.5n$ to $4.5n$) to resource 2. This gives additional reward of at least $\frac{\sum_{i=0}^{0.5n-1} 1-e^{-0.5+i/n}}{1-1/e} - O(n\delta) + n$. The total value of OPT is at least,

$$\begin{aligned}
&\frac{\sum_{i=1}^{n} 1 - e^{-i/n}}{1-1/e} + 2n + \frac{\sum_{i=0}^{0.5n-1} 1 - e^{-0.5+i/n}}{1-1/e} + n - O(n\delta), \\
&\geq n\left(3 + \frac{1/e}{1-1/e} + \frac{0.5 - (1/\sqrt{e} - 1/e)}{1-1/e} - O(\delta)\right), \\
&\overset{\delta \to 0}{\geq} \frac{1}{0.626} \times 2.5\, n.
\end{aligned}$$

□

A.2. Comparison for Exponential usage distributions

We consider a different example below to further demonstrate RBA's ability to adapt to arrival sequence and usage distribution, without explicitly using the distribution.

Example A.3 Consider a setting with two resources that have rewards $r_1 = 1$ and $r_2 = 2$ and large starting capacities n. The usage time distribution of both resources is exponential with rate μ (mean $1/\mu$). We receive an arrival sequence where the first $n-1$ arrivals come in a very short span of time $\epsilon \to 0$ and are only interested in a unit of resource 2. Suppose that the next arrival, call it t_0, is 1 unit of time later and can be matched to either resource.

Comparing Greedy, Balance, and RBA: All algorithms match the first $n-1$ arrivals to resource 2. Greedy matches t_0 to resource 2 for every value of μ. On the other hand, Balance will be quite risk-averse and protect resource 2 from being matched unless $\mu > 0.5$ (roughly). RBA, on the other hand, responds quite nimbly to μ. It will protect resource 2 when $\mu \to 0$ i.e., the 1 unit time interval is insignificant for inventory to replenish. For every non-infinitesimal value of μ, the highest available unit of resource 2 will have index at least $n - O(\log n)$ with high probability and therefore, RBA will act greedily and match the arrival to resource 2. Figure 1 summarizes these differences between the algorithms.

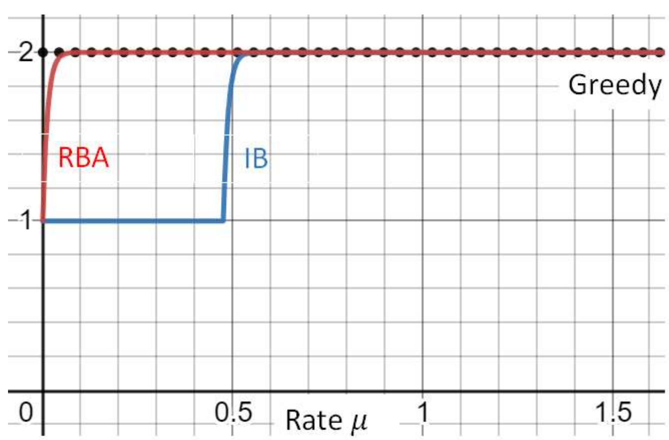

Figure 1 Comparison of Greedy, Balance, and RBA on Example A.3. Numbers on the y axis denote the resource matched to t_0 w.h.p..

A.3. Variants of Balance

Examples A.1 and A.3 hint at the following "switching" behavior: In phases of arrivals where resources return "frequently" relative to the arriving demand, it is better to be greedy. On the other hand, when arrivals occur in a large batch and resources can not return in time, we should follow Balance to protect diminishing resources. Since we have no information about future arrivals it is not clear how an online algorithm can make this "switch" optimally. While RBA manages to perform this switch quite nimbly, the following natural variants of Balance fail to do so.

(i) An algorithm that uses distributional knowledge to deduce when some items are not going to return and refreshes the maximum capacity instead of just the remaining capacity. In the context of Example A.1, this algorithm would realize after 1 unit of time that the units of resource 1 that have not returned, will never return. Subsequently, it computes a new maximum capacity of $n/2$ for resource 1 at time 1, and treats the resource as if it were at full capacity. It is not hard to see that in general this ends up converging to the greedy algorithm (when the return probabilities approach 0 for instance).

(ii) An algorithm that anticipates that items are going to return in the future and considers a more optimistic remaining capacity level. In Example A.1, this algorithm would deduce that there are no further items returning after time 1. Therefore, it makes the same decisions as Balance on the instance in Example A.1.

A.4. Challenges with Standard Primal-dual Analysis

From an analysis standpoint, it has previously been observed that using the primal-dual framework of Devanur et al. (2013) and Buchbinder et al. (2007), a cardinal technique of analysis in case of non-reusable resources, presents non-trivial challenges and typical dual fitting arguments do not seem to work (Rusmevichientong et al. 2020, Gong et al. 2022). For a concrete demonstration, consider the LP upper bound on clairvoyant and its dual,

$$\begin{aligned} \textbf{Primal} \quad & \max \sum_{(i,t)\in E} r_i y_{it} \\ & s.t. \ \sum_{t=1}^{\tau} [1 - F_i(a(\tau) - a(t))] y_{it} \leq c_i \quad \forall \tau \in \{1, \dots, T\}, \forall i \in I \\ & \sum_{i\in I} y_{it} \leq 1 \quad \forall t \in T \\ & 0 \leq y_{it} \leq 1 \quad \forall t \in T, \quad i \in I \end{aligned} \tag{19}$$

$$\begin{aligned} \textbf{Dual} \quad & \min \sum_{t} \lambda_t + \sum_{(i,t)\in E} c_i \theta_{it} \\ & s.t. \ \lambda_t + \sum_{\tau | a(\tau) \geq a(t)} [1 - F_i(a(\tau) - a(t))] \, \theta_{i\tau} \geq r_i \quad \forall (i,t) \in E \\ & \lambda_t, \theta_{it} \geq 0 \quad \forall t \in T, i \in I \end{aligned} \tag{20}$$

Given an online algorithm ALG (with expected revenue ALG), if one can find a dual fitting $\lambda_t \geq 0$ and $\theta_{i\tau} \geq 0$ such that,

(i) $\lambda_t + \sum_{\tau | a(\tau) \geq a(t)} [1 - F_i(a(\tau) - a(t))] \, \theta_{i\tau} \geq \alpha \, r_i, \, \forall (i,t) \in E.$

(ii) $\text{ALG} \geq \sum_t \lambda_t + \sum_{(i,t)\in E} c_i \theta_{it}.$

Then, by weak duality ALG is α–competitive. For this certificate, we demonstrate that the standard dual fitting approach for finding a feasible dual solution cannot be used to obtain tight guarantees for stochastic (or even fluid) reusability.

In the standard procedure for defining a candidate solution to the dual, we initialize all dual variables to 0 and then update the values in tandem with the matching decisions made by the online algorithm (Buchbinder et al. 2007, Devanur et al. 2013). For simplicity, consider a deterministic online algorithm ALG. When ALG matches arrival t to resource i, the reward r_i is split into two parts and the dual variables are updated as follows,

$$\lambda_t = r_i \alpha_{it}; \quad \sum_{\tau \in T} \theta_{it} = \sum_{\tau \in T} \theta_{it} + \frac{r_i}{c_i}(1 - \alpha_{it}),$$

here $\alpha_{it} \in [0,1]$ decides how the gain of r_i is split between λ_t and $\sum_{\tau \in T} \theta_{it}$. This procedure ensures that at the end of the planning horizon we have,

$$\sum_{t\in T} \lambda_t + \sum_{i \in I} \sum_{t \in T} c_i \theta_{it} = \text{ALG}.$$

The values of splitting parameters α_{it} need to be carefully tuned to obtain the desired performance guarantee. We show that when usage durations are stochastic there is, in general, no setting of the splitting parameters that can be used to certify the actual performance of Balance.

For family of instances given in Example A.2, to show that Balance has total reward αOPT using this approach, we need to find splitting parameters that satisfy the following system of inequalities,

$$\begin{aligned}
&\lambda_t + \sum_{\tau \in [2n]\backslash[t-1]} \theta_{1,\tau} + \sum_{\tau \in [3n]\backslash[2n]} 0.5\theta_{1,\tau} \geq \alpha \quad \forall t \in [2n] \\
&\lambda_t + \theta_{i,t} + \sum_{\tau \in \{t+1,\cdots,(3+1.5(i-1))n\}} 0.5\theta_{i,\tau} \geq \alpha \quad \forall i \in \{1,2\}, \, t \in \{2n+1,\cdots,3n\} \\
&\lambda_t + \sum_{\tau \in \{t,\cdots,4.5n\}} \theta_{2,\tau} \geq \alpha \quad \forall t \in \{3n+1,\cdots,4.5n\} \\
&\lambda_t + \theta_{i_t} \geq \alpha r_{i_t} \quad \forall t \in [4.5n] \\
&\lambda_t, \theta_{it} \geq 0,
\end{aligned}$$

here we use a single variable θ_{i_t} for low reward resource i_t since resource i_t does not have an edge to any arrival other than t. Note that for $t \in ([2n]\backslash[n]) \cup (\{3.5n,\cdots,4.5n\})$, the low reward resources have price 0 and can be ignored. Now, under a fluid version of reusability, Balance does not match any arrival to low reward resources (with non-zero price). Based on the standard dual fitting procedure, we set,

$$\theta_{i_t} = 0 \quad \forall t \in [4.5n].$$

Thus, $\lambda_t \geq \alpha r_{i_t}$ for all $t \in [4.5n]$. Similarly, we have $\lambda_t = 0$ for $t \in [2n]\setminus[n]$ and $t \geq 3.5n+1$. Now, consider a simplified and reduced system of inequalities,

$$\begin{aligned}
\theta_{1,2n} + \sum_{\tau \in [3n]\setminus[2n]} 0.5\theta_{1,\tau} &\geq \alpha, \\
\theta_{1,t} + \sum_{\tau \in [3n]\setminus[t]} 0.5\theta_{1,\tau} &\geq \alpha - \lambda_t \quad \forall t \in \{2n+1, \cdots, 3n\}, \\
\theta_{2,4.5n} &\geq \alpha, \\
\lambda_t &\geq \alpha r_{i_t} \quad \forall t \in [4.5n], \\
\sum_{t \in [4.5n]} \lambda_t + n \sum_{i \in \{1,2\}, t \in [4.5n]} \theta_{it} &= 2.5n.
\end{aligned}$$

For $\alpha = 0.626$, it can be verified that this system does not have a feasible solution, i.e., there is no setting of the splitting parameters α_{it} that gives a feasible solution to the inequalities above.

Appendix B: Validity of Generalized Certificate

LEMMA 3. *Given an online algorithm* ALG*, non-negative values* $\{\lambda_t(\omega,\nu)\}_{t,\omega,\nu}$ *and* $\{\theta_i\}_i$ *such that conditions* (2) *and* (3) *hold, we have*

$$\text{ALG} \geq \frac{\min_{i \in I} \alpha_i}{\beta} \text{OPT}.$$

Proof. We start by summing both sides of constraint (3) over $i \in I$,

$$\begin{aligned}
\sum_{i \in I} \alpha_i \, \text{OPT}_i \leq \sum_i \Big(\theta_i + \mathbb{E}_{\boldsymbol{\omega},\boldsymbol{\nu}}\Big[\sum_{t \in O(\omega,i)} \lambda_t(\omega,\nu)\Big]\Big) &= \sum_{i \in I} \theta_i + \mathbb{E}_{\boldsymbol{\omega},\boldsymbol{\nu}}\Big[\sum_{i \in I} \sum_{t \in O(\omega,i)} \lambda_t(\omega,\nu)\Big] \\
&\overset{(a)}{\leq} \sum_{i \in I} \theta_i + \mathbb{E}_{\boldsymbol{\omega},\boldsymbol{\nu}}\Big[\sum_{t \in T} \lambda_t(\omega,\nu)\Big] \\
&\overset{(b)}{=} \sum_{i \in I} \theta_i + \sum_{t \in T} \mathbb{E}_{\boldsymbol{\omega},\boldsymbol{\nu}}\Big[\lambda_t(\omega,\nu)\Big] \\
&\leq \beta \text{ALG}.
\end{aligned}$$

Inequality (a) follows from the fact that $\{O(\omega,i)\}_{i \in I}$ is collection of disjoint subsets of T and also that $\lambda_t(\omega,\nu) \geq 0 \ \forall t \in T$. Equality (b) follows by exchanging the order of the sum and expectation. The final inequality follows from (2). □

B.1. Tightness of LP free Certificate

Let ω denote a sample path w.r.t. all the randomness in OPT. Let ν denote a sample path w.r.t. all the randomness in ALG. The LP free certificate is given by the following inequalities,

$$\begin{aligned}
\sum_{t \in T} \mathbb{E}_{\boldsymbol{\omega},\boldsymbol{\nu}}[\lambda_t(\omega,\nu)] + \sum_{i \in I} \theta_i &\leq \beta \text{ALG} \\
\mathbb{E}_{\boldsymbol{\omega},\boldsymbol{\nu}}\left[\sum_{t \in O(\omega,i)} \lambda_t(\omega,\nu)\right] + \theta_i &\geq \alpha \text{OPT}_i \quad \forall i \in I, \\
\lambda_t, \, \theta_i &\geq 0.
\end{aligned}$$

We view the linear system above as an LP with a trivial objective of minimizing 0 (a constant). The dual of this LP is,

$$\begin{aligned} \max \quad & -\beta \text{ALG}\, y + \alpha \sum_{i \in I} \text{OPT}_i\, x_i \\ \text{s.t.} \quad & 0 \le x_i \le y \quad \forall i \in I, \\ & y \ge 0, \end{aligned}$$

here we use the fact that for any given arrival t and sample paths ω and ν, there is at most one set $O(\omega, i)$ that contains t. The optimal value of this LP is $\max_{y \ge 0} \left[(-\beta \text{ALG} + \alpha \sum_{i \in I} \text{OPT}_i)\, y\right]$. Thus, from strong duality we have that $\text{ALG} \ge \frac{\alpha}{\beta}\text{OPT}$, if and only if our LP-free system has a feasible solution. This implies that our certificate is tight, i.e., if ALG has a competitive ratio guarantee of $\gamma \in (0,1]$, then there exists a feasible solution to our linear system with $\alpha_i = 1\, \forall i \in I$ and $\beta = 1/\gamma$.

Appendix C: Missing Details from Proof of Theorem 2

C.1. Properties of $(F, \boldsymbol{\sigma}, \boldsymbol{p})$ Random Process

LEMMA 5. *Given a distribution F, arrival set $\boldsymbol{\sigma} = \{\sigma_1, \cdots, \sigma_T\}$, and probability sequences $\boldsymbol{p}_1 = (p_{11}, \cdots, p_{1T})$ and $\boldsymbol{p}_2 = (p_{21}, \cdots, p_{2T})$ such that, $p_{1t} \le p_{2t}$ for every $t \in [T]$, we have,*

$$r(F_i, \boldsymbol{\sigma}, \boldsymbol{p}_1) \le r(F_i, \boldsymbol{\sigma}, \boldsymbol{p}_2).$$

Proof. Suppose the resource is available at some arrival $\sigma_t \in \boldsymbol{\sigma}$. Recall that independent of all other randomness in the random process, w.p. p_t we switch the resource to in-use at σ_t and w.p. $1 - p_t$ the resource stays available till at least the next arrival. Consider an alternative random process where given a set $\boldsymbol{\sigma}$ and probability sequence $\boldsymbol{p}$, we first sample a random subset $\boldsymbol{\sigma}_{\boldsymbol{p}}$ of $\boldsymbol{\sigma}$ as follows: independently for each arrival $\sigma_t \in \boldsymbol{\sigma}$, include the arrival in the subset w.p. p_t. Taking expectation over this random sampling we claim that,

$$\mathbb{E}[r(F_i, \boldsymbol{\sigma}_{\boldsymbol{p}})] = r(F_i, \boldsymbol{\sigma}, \boldsymbol{p}).$$

This is a direct implication of the fact that an available resource is switched to unavailable independent of other events.

Now, consider random processes $(F_i, \boldsymbol{\sigma}, \boldsymbol{p}_1)$ and $(F_i, \boldsymbol{\sigma}, \boldsymbol{p}_2)$ as given in the statement of the lemma. We establish the main claim by using the alternative viewpoint defined above to couple the two processes. To be more precise, we couple the subset sampling stage in the two processes. First, sample a random subset $\boldsymbol{\sigma}_{\boldsymbol{p}_1}$ by including each arrival σ_t with corresponding probability p_{1t}. Next, sample subset $\boldsymbol{\sigma}_{\boldsymbol{p}_2 \backslash \boldsymbol{p}_1}$ of $\boldsymbol{\sigma}$ by independently including arrival σ_t w.p. $\frac{1}{1-p_{1t}}(p_{2t} - p_{1t})$, for every $\sigma_t \in \boldsymbol{\sigma}$. Finally, let $\boldsymbol{\sigma}_{\boldsymbol{p}_2} = \boldsymbol{\sigma}_{\boldsymbol{p}_1} \cup \boldsymbol{\sigma}_{\boldsymbol{p}_2 \backslash \boldsymbol{p}_1}$. Since for every t, $\sigma_t \in \boldsymbol{\sigma}_{\boldsymbol{p}_2}$ with probability p_{2t} independent of other points in $\boldsymbol{\sigma}_{\boldsymbol{p}_2}$, we have,

$$\mathbb{E}[r(F_i, \boldsymbol{\sigma}_{\boldsymbol{p}_2})] = r(F_i, \boldsymbol{\sigma}, \boldsymbol{p}_2) \text{ and } \mathbb{E}[r(F_i, \boldsymbol{\sigma}_{\boldsymbol{p}_1})] = r(F_i, \boldsymbol{\sigma}, \boldsymbol{p}_1).$$

Observe that $\boldsymbol{\sigma}_{\boldsymbol{p}_1} \subseteq \boldsymbol{\sigma}_{\boldsymbol{p}_2}$ on every sample path. Thus, to finish the proof it suffices to argue that $r(F,\boldsymbol{\sigma}_1) \leq r(F,\boldsymbol{\sigma}_2)$ if $\boldsymbol{\sigma}_1 \subseteq \boldsymbol{\sigma}_2$. Consider $(F,\boldsymbol{\sigma}_1)$ and $(F,\boldsymbol{\sigma}_2)$ random processes and the straightforward coupling of usage durations where we have a list of i.i.d. samples drawn according to distribution F and each process independently parses this list in order, moving to the next sample whenever the current sample is used and never skipping samples. On any coupled path, the number of transitions made on arrivals in $\boldsymbol{\sigma}_2$ is lower bounded by the number of transitions made on arrivals in $\boldsymbol{\sigma}_1$, giving us the desired. □

LEMMA 6. *Given an $(F,\boldsymbol{\sigma},\boldsymbol{p})$ random process, let $\boldsymbol{\sigma}' \subset \boldsymbol{\sigma}$ be a subset of arrivals where the resource is unavailable with probability 1. Then, at every arrival $\sigma_t \in \boldsymbol{\sigma}$, the probability that the resource is available at σ_t is identical in $(F,\boldsymbol{\sigma},\boldsymbol{p})$ and $(F,\boldsymbol{\sigma},\boldsymbol{p} \vee \mathbf{1}_{\boldsymbol{\sigma}'})$.*

Proof. It suffices to show the lemma for a subset $\boldsymbol{\sigma}'$ consisting of a single arrival. The result for general $\boldsymbol{\sigma}'$ then follows by repeated application. Now, let σ_t denote an arbitrary arrival in $\boldsymbol{\sigma}$ such that w.p. 1, the resource is in-use when σ_t arrives. Observe that we can change the probability p_t associated with σ_t arbitrarily, but this does not change the probability of resource being available at σ_t. In particular, w.p. 1, the resource is in-use when σ_t arrives in the $(F,\boldsymbol{\sigma},\boldsymbol{p} \vee \mathbf{1}_t)$ random process as well. Consequently, the probabilities at other arrivals are unchanged in going from $(F,\boldsymbol{\sigma},\boldsymbol{p})$ to $(F,\boldsymbol{\sigma},\boldsymbol{p} \vee \mathbf{1}_t)$. □

LEMMA 7. *The following statements are true for every $(F,\boldsymbol{\sigma},\boldsymbol{p})$ random process:*

(i) For every $\sigma_t \in \boldsymbol{\sigma}$, the probability that the resource is available at σ_t equals the fraction of the resource available at σ_t in the fluid $(F,\boldsymbol{\sigma},\boldsymbol{p})$ process.

(ii) The expected reward $r(F,\boldsymbol{\sigma},\boldsymbol{p})$, equals the total reward in the fluid $(F,\boldsymbol{\sigma},\boldsymbol{p})$ process.

Proof. Note that, statement (ii) of the lemma is a direct consequence of statement (i). The proof of (i) hinges on the fact that in the $(F,\boldsymbol{\sigma},\boldsymbol{p})$ random process, the duration of every state transition is independent of past randomness. Using this we write a recursive equation for the probability of reward at every arrival. Let $\eta(\sigma_t)$ denote the probability that the resource is available when $\sigma_t \in \boldsymbol{\sigma}$ arrives. We have,

$$\eta(\sigma_t) = \eta(\sigma_{t-1})\big(1 - p_{t-1}\big) + \sum_{\tau=1}^{t-1} \eta(\sigma_\tau) p_\tau \big(F(\sigma_t - \sigma_\tau) - F(\sigma_{t-1} - \sigma_\tau)\big),$$

where $\eta(\sigma_1) = 1$. By forward induction, it is easy to see that this set of equations has a unique solution. Now, consider the fluid $(F,\boldsymbol{\sigma},\boldsymbol{p})$ process and let $\eta'(\sigma_t)$ denote the fraction of resource available at σ_t in the fluid process. Clearly, $\eta'(\sigma_1) = 1$ and we have,

$$\eta'(\sigma_t) = \eta'(\sigma_{t-1})\big(1 - p_{t-1}\big) + \sum_{\tau=1}^{t-1} \eta'(\sigma_\tau) p_\tau \big(F(\sigma_t - \sigma_\tau) - F(\sigma_{t-1} - \sigma_\tau)\big).$$

Thus,

$$\eta(\sigma_t) = \eta'(\sigma_t) \quad \forall t \leq T.$$

The expected reward from a match occurring at σ_t in the $(F,\boldsymbol{\sigma},\boldsymbol{p})$ random process is $p_t \eta(\sigma_t)$, which, using the equality above, is equal to the fraction of resource consumed at σ_t in the fluid process. □

C.2. Applying Chernoff and Missing Pieces of Theorem 2

LEMMA 12. *Given integer $\tau > 0$, real value $c > 0$, independent indicator random variables $\mathbb{1}(t)$ for $t \in [\tau]$ and $\delta = \sqrt{\frac{\log c}{c}}$ such that, $\sum_{t=1}^{\tau} \mathbb{E}[\mathbb{1}(t)] \le \frac{c}{1+\delta}$. We have,*

$$\mathbb{P}\left(\sum_{t=1}^{\tau} \mathbb{1}(t) \ge c\right) \le \frac{1}{\sqrt{c}}.$$

Proof. Let $\mu = \frac{c}{1+\delta}$. Let the condition on total mean hold with equality, i.e.,

$$\sum_{t=1}^{\tau} \mathbb{E}[\mathbb{1}(t)] = \mu$$

This is w.l.o.g., as we can always add some number of dummy independent binary random variables to make the condition hold with equality. Applying the standard multiplicative Chernoff bound for independent Bernoulli random variables $\{\mathbb{1}(t)\}_{t\in[\tau]}$, we have

$$\mathbb{P}\Big(\sum_{t=1}^{\tau} \mathbb{1}(t) \ge (1+\delta)\mu = c\Big) \le e^{-\frac{\mu\delta^2}{2+\delta}} < \frac{1}{\sqrt{c}},$$

here $\frac{\mu\delta^2}{2+\delta} = \frac{\log c}{1+2\delta+\delta^2} \ge 0.5\log c$ for $c \ge 1$. □

Appendix D: From Matching to Multi-unit Assortments

In this section we generalize the $(1-1/e)$ result for the following model.

Online Assortments with Multi-unit demand: Customer t requires up to $b_{it} \ge 0$ units of resource $i \in I$. Given an assortment S, the customer chooses at most one resource from S with probabilities given by choice model ϕ_t. Let $y_i(t)$ denote the number of units of resource i when t arrives. Selection of resource i results in $\min\{y_i(t), b_{it}\}$ units of resource i being used for an independently drawn duration $d \sim F_i$, and a reward $\min\{y_i(t), b_{it}\} r_i$ (results hold even if each of the b_{it} units is used for an independent random duration). The assortment S that we offer must belong to a downward closed feasible set $\mathcal{F}_t$. Choice model ϕ_t and quantities b_{it} are revealed when t arrives.

Similar to the online matching problem, we focus on the large capacity regime. Due to multi-unit (budgeted) allocations this is more accurately the large budget to bid ratio regime, where the parameter,

$$\gamma := \min_{i\in I,\, t\in T} \frac{c_i}{b_{it}}, \text{ approaches } +\infty.$$

We compare online algorithms against a clairvoyant algorithm that knows the choice models and quantities b_{it} for all arrivals in advance but makes assortment decisions in order of the arrival sequence and observes (i) realizations of customer choice after showing the assortment and (ii) realizations of usage duration when used units return (same as an online algorithm). Further, in case of assortments we make the standard assumptions (Golrezaei et al. (2014), Rusmevichientong et al. (2020), Gong et al. (2022)) that for every arrival $t \in T$, choice model ϕ_t satisfies the weak substitution property, i.e.,

$$\phi_t(S,i) \geq \phi_t(S \cup \{j\}, i), \quad \forall i, j \notin S, \forall t \in T. \tag{21}$$

We also assume access to an assortment optimization oracle that takes a choice model ϕ and set of feasible solutions $\mathcal{F}$ as input and outputs a feasible revenue maximizing assortment More generally, an α approximate oracle is also acceptable and in this case the competitive ratio guarantee is $(1-1/e)\,\alpha$.

Since we now need to think in terms of sets of resources offered to arrivals, a relatively straightforward way to generalize G-ALG will be to fractionally “match” every arrival to a collection of revenue maximizing assortments/sets, consuming constituent resources in a fluid fashion in accordance with the choice probabilities.

ALGORITHM 4: Assort G-ALG

Output: For every arrival t, collection of assortments and probabilities $\{A(\eta,t), u(\eta,t)\}_\eta$;

Let $g(t) = e^{-t}$, and initialize $Y(k_i) = 1$ for every $i \in I, k_i \in [c_i]$;

for *every new arrival* t **do**

For every $k_i \in [c_i]$ and $t \geq 2$, update values

$$Y(k_i) = Y(k_i) + \sum_{\tau=1}^{t-1} \Big(F_i\big(a(t)-a(\tau)\big) - F_i\big(a(t-1)-a(\tau)\big)\Big) y(k_i,\tau)$$

`// Fluid update of returning capacity`

Initialize $S_t = \{i \mid (i,t) \in E\}$, values $\eta = 0$, and $y(k_i,t) = 0$ for all $i \in S_t, k_i \in [c_i]$;

while $\eta < 1$ *and* $S_t \neq \emptyset$ **do**

for $i \in S_t$ **do**

if $Y(k_i) = 0$ for every $k_i \in [c_i]$ **then** remove i from S_t;

else $z_i = \arg\max_{k_i \in [c_i]} \{k_i \mid Y(k_i) > 0\}$; `// Highest available unit`

end

$A(\eta,t) = \arg\max_{S \in S_t} \sum_{i \in S} b_{it} r_i \phi_t(S,i) \left(1 - g\left(\frac{z_i}{c_i}\right)\right)$ `// Optimal assortment with RBA`

$u(\eta,t) = \min\left\{1-\eta, \min_{i \in A(\eta,t)} \frac{Y(z_i)}{b_{it}\phi_t(A(\eta,t),i)}\right\}$ `// Fractional assortment`

Update $\eta \to \eta + u(\eta,t)$;

for $i \in A(\eta,t)$ **do**

Update $y(z_i,t) \to y(z_i,t) + u(\eta,t)\, b_{it}\, \phi_t(A(\eta,t),i)$; $\quad Y(z_i) \to Y(z_i) - y(z_i,t)$;

`// Inventory update after fluid customer choice`

end

end

end

Description of Assort G-ALG (Algorithm 4): Observe that the stochasticity due to choice has been converted to its fluid version. Specifically, arrival $t \in T$ is fractionally "matched" to assortments $A(1,t), \cdots, A(m,t)$ for some $m \geq 0$. The weight/fraction of assortment $A(j,t)$ is given by $u(j,t) > 0$ and we have, $\sum_{j=1}^{m} u(j,t) \leq 1$. The amount of resource i (fluidly) consumed as a result of this is given by $\sum_{A(j,t)\ni i} u(j,t)\, \phi\big(A(j,t), i\big)$. The collection of assortments is found by computing the revenue maximizing assortment with reduced prices computed according to RBA rule, as in the case of matching. The values $y(k_i,t)$ in the algorithm correspond to the total fraction of unit k_i that is fluidly chosen by arrival t. The weights $u(j,t)$ are chosen to ensure that $y(k_i,t)$ does not exceed the fraction of k_i that was available when t arrived. We assume w.l.o.g. that the oracle that outputs revenue maximizing assortments never includes resources with zero probability of being chosen in the assortment. Interestingly, the performance guarantee of the relaxed online algorithm Assort G-ALG, depends only on $c_{\min} = \min_{i\in I} c_i$.

Lemma D1 *For every instance of the online budgeted assortment problem we have,*

$$\textit{Assort}\ \text{G-ALG} \geq (1-1/e)\, e^{-\frac{1}{c_{\min}}}\, \text{OPT}.$$

Proof. Note that the sample path ω now also includes the randomness due to customer choice. Let $O(\omega,i)$ denote the set of all arrivals on sample path ω in OPT where some units of i are *chosen*. Since each arrival chooses (possibly multiple units of) *at most one* resource, we interpret $O(\omega,i)$ as the set of arrivals that choose i. Let $b(\omega,i,t)$ denote the number of units of i chosen in OPT at arrival $t \in O(\omega,i)$. Let $z_i(t^+)$ be the highest index unit of resource i that has a non-zero fraction available in Assort G-ALG at t^+. We use the generalized certification with sample path based variables $\lambda_t(\omega)$. Given a sample path ω and resource $i \in I$, we set,

$$\lambda_t(\omega) = b(\omega,i,t) r_i \left(1 - g\Big(\frac{z_i(t^+)}{c_i}\Big)\right), \tag{22}$$

for every arrival $t \in O(\omega,i)$. Let $\lambda_t = \mathbb{E}_{\boldsymbol{\omega}}[\lambda_t(\omega)]$. Recall that $y(k_i,t)$ is the total fraction of unit k_i that is fluidly chosen by arrival t. By definition of Assort G-ALG, we have

$$\lambda_t \leq \sum_{i\in I,\, k_i\in[c_i]} y(k_i,t)\, r_i \left(1 - g\left(\frac{k_i}{c_i}\right)\right),$$

i.e., λ_t is at most the expected total reward at t in Assort G-ALG calculated with units at their reduced price. Setting θ_i as before (in the proof of Lemma 4), i.e.,

$$\theta_i = c_i \left(e^{\frac{1}{c_i}} - 1\right) r_i \sum_t \sum_{k\in[c_i]} y(k_i,t) g\Big(\frac{k}{c_i}\Big),$$

we have that condition (2) of the certificate is satisfied with $\beta = e^{1/c_{\min}}$.

The remaining analysis now mimics the proof of Lemma 4. We generalize the basic setup to demonstrate this formally. Let $\mathbb{1}(\neg k, t^+)$ indicate that no fraction of unit k is available in Assort G-ALG right after t has been matched. Recall that $\Delta g(k) = g\big(\frac{k-1}{c_i}\big) - g\big(\frac{k}{c_i}\big)$. By definition of λ_t (see (22)), we have

$$\mathbb{E}_{\boldsymbol{\omega}}\Big[\sum_{t\in O(\omega,i)} \lambda_t(\omega)\Big] \geq (1-1/e)\text{OPT}_i - r_i\mathbb{E}_{\boldsymbol{\omega}}\Big[\sum_{t\in O(\omega,i)} b(\omega,i,t)\sum_{k\in[c_i]}\Delta g(k)\mathbb{1}(\neg k,t^+)\Big].$$

Fix an arbitrary unit k_O of i and let $O(\omega, k_O)$ denote the set of arrivals on sample path ω in OPT where k_O is one of the chosen units of i. Note that $O(\omega, k_O)$ is a subset of $O(\omega, i)$. Using the decomposition above, it suffices to show that,

$$\mathbb{E}_{\boldsymbol{\omega}}\Big[\sum_{t\in O(\omega,k_O)}\sum_{k\in[c_i]}\Delta g(k)\mathbb{1}(\neg k,t^+)\Big] \leq \frac{1}{c_i r_i}\theta_i.$$

The proof of this inequality follows Lemma 4 verbatim. Crucially, we have the following inequalities that complete the proof,

$$\mathbb{E}_{\boldsymbol{\omega}}\Big[\sum_{t\in O(\omega,k_O)}\mathbb{1}(\neg k,t^+)\Big] \leq r(F_i,\boldsymbol{s}(k)) \leq r(F_i,\boldsymbol{T},\boldsymbol{p}(k)),$$

where $\boldsymbol{s}(k), \boldsymbol{T}$, and $\boldsymbol{p}(k)$ are as defined in Lemma 4, i.e., $\boldsymbol{s}(k)$ is the ordered set of all arrivals t (arrival times $a(t)$ to be precise) such that $\mathbb{1}(\neg k, t^+) = 1$ in Assort G-ALG. $\boldsymbol{T}$ is the ordered set of all arrivals. Finally, probabilities $p(k,t) \in \boldsymbol{p}(k)$ are defined as follows: $p(k,t) = 0$ if $y(k,t) = 0$, otherwise $p(k,t) = \frac{y(k,t)}{\eta(k,t)}$, where $\eta(k,t)$ is the fraction of k available in Assort G-ALG when t arrives. □

The main new challenge in turning Assort G-ALG into Sample Assort G-ALG is that we must deal with scenarios where Assort G-ALG directs some mass towards a set A but only a subset of resources in A are available in Sample G-ALG. Recall that in case of matching, if the randomly chosen resource is unavailable we simply leave t unmatched. We could consider a similar approach here whereby if any unit of sampled set A is unavailable then we do not offer A. However, this will not preserve the overall revenue in expectation as the probability of every resource in A being available simultaneously is likely to be small. If it were acceptable to offer an assortment with items that are not available in Sample G-ALG then we could also offer the set A as is. The underlying assumption in such a case is that if the arrival chooses an unavailable item then we earn no reward and the arrival simply departs (called static substitution in Ma et al. (2021)). However, in many applications it may not be possible or desirable to offer an assortment where some items are unavailable.

An alternative approach is to offer the subset S of A that is available in Sample G-ALG at t. However, this can affect the choice probability for resources $i \in S \cap A$ in non-trivial ways, and thus, affect future availability of resources in a way that is challenging to control. In other words, the concentration bounds that show Sample G-ALG has the same performance as G-ALG for large inventory, will not apply here. Consequently, we need to find a way to display some subsets of A such that, (i) the overall probability of any given resource being allocated is no larger than if we offered A itself and (ii) we do not rely on multiple resources in A being available simultaneously. The main novelty of our approach to tackle this problem will be to switch our perspective from sets of resources back to individual resources. Specifically, for each resource we find the overall probability that the resource is chosen by a given arrival and then use these probabilities as our guideline, i.e., given the subset $S \subseteq A$ of resources that is available, we find a new collection of assortments so that for every available resource, the overall probability of the resource being chosen matches this probability in the original collection of assortments in Assort G-ALG. The main idea here is a probability matching, made non trivial by the fact that we are restricted to choice probabilities given by the choice model. We find an iterative polytime algorithm (Algorithm 3) that ensures that the probability of a resource being chosen by an arrival in Sample G-ALG, matches that in Assort G-ALG.

Recall that $\gamma := \min_{(i,t)\in E} \frac{c_i}{b_{it}}$. We are interested in the case where $\gamma \to +\infty$. Note that we assume knowledge of a lower bound on γ in Sample Assort G-ALG. Overloading notation, we denote this lower bound also as γ.

ALGORITHM 5: Sample Assort G-ALG

Initialize capacities $y_i(0) = c_i$ and let $\delta = \sqrt{\frac{\log \gamma}{\gamma}}$;

for *every new arrival* t **do**

  Update capacities $\{y_i(t)\}_{i\in I}$ for resources with returning units and let $S_t = \{i \mid y_i(t) \geq b_{it}\}$;

  Get collection of assortments $\{A(\eta,t), u(\eta,t)\}_\eta$ from Assort G-ALG;

  Randomly sample collection η w.p. $u(\eta,t)$;

  For sampled η, let $\hat{\mathcal{A}}, \hat{\mathcal{U}} =$ Probability Match $\Big(A(\eta,t) \cap S_t, \{\frac{1}{1+\delta}\phi_t(A(\eta,t),s)\}_{s\in A(\eta,t)\cap S_t}\Big)$;

  Randomly sample asssortment $\hat{A}_j \in \hat{\mathcal{A}}$ with distribution $\hat{\mathcal{U}}$;

  `// Assortment may be empty with non-zero probability`

  Offer $\hat{A}_j$ to t and update capacity after t chooses;

end

LEMMA 6. *(Lemma 16 restated) Consider a choice model $\phi : 2^N \times N \to [0,1]$ satisfying the weak substitution property (see* (21)*), an assortment $A \subseteq N$ belonging to a downward closed feasible set $\mathcal{F}$, a subset $S \subseteq A$ and target probabilities p_s such that, $p_s \leq \phi(A,s)$ for every $s \in S$. There exists a collection $\mathcal{A} = \{A_, \cdots, A_m\}$ of $m = |S|$ assortments along with weights $\big(u_i\big) \in [0,1]^m$, such that the following properties are satisfied:*

(i) For every $i \in [m]$, $A_i \subseteq S$ and thus, $A_i \in \mathcal{F}$.

(ii) Sum of weights, $\sum_{i\in[m]} u_i \leq 1$.

(iii) For every $s \in S$, $\sum_{A_i \ni s} u_i\, \phi(A_i,s) = p_s$.

Algorithm 3 computes such a collection $\mathcal{A}$ along with weights (u_i) in $O(m^2)$ time.

Proof of Lemma 16. We give a constructive proof that outlines Algorithm 3 introduced earlier. Let,

$$q_s^0 = \phi(S,s) \text{ and } \zeta_s^0 = \frac{p_s}{q_s^0} \text{ for every } s \in S.$$

Observe that $q_s^0 \geq \phi(A,s) \geq p_s$, due to substitutability. Thus, $\zeta_s^0 \leq 1$ for every $s \in S$.

Let s_1 be an element in S with the smallest value $\zeta_{s_1}^0$. Let $A_1 = S$ be the first set added to collection $\mathcal{A}$ with $u_1 = \zeta_{s_1}^0$, so that $u_1\phi(A_1,s_1) = p_{s_1}$. We will ensure that all subsequent sets added to $\mathcal{A}$ do not include the element s_1 and this will guarantee condition (iii) for element s_1. Next, define the set $S^1 = S\backslash\{s_1\}$. Let,

$$q_s^1 = \phi(S^1,s) \geq q_s^0 \text{ and } \zeta_s^1 = \frac{p_s - u_1 q_s^0}{q_s^1} \text{ for every } s \in S^1.$$

Observe that $\zeta_s^1 \in [0,1]$ for every $s \in S^1$. Let s_2 denote the element with the smallest value $\zeta_{s_2}^1$, out of all elements in S^1. If $\zeta_{s_2}^1 = 0$ we stop, otherwise we now add the second set $A_2 = S^1$ to the collection with $u_2 = \zeta_{s_2}^1$. Inductively, after i iterations of this process, we have added i nested sets $A_i \subset A_{i-1} \subset \cdots \subset A_1$ to the collection and have the remaining set $S^i = A_i\backslash\{s_i\}$ of $|A| - i$ elements. Define values,

$$q_s^i = \phi(S^i,s) \text{ and } \zeta_s^i = \frac{p_s - \sum_{k=1}^{i} u_k q_s^{k-1}}{q_s^i} \text{ for every } s \in S^i.$$

Let $s_{i+1} \in S^i$ be the element with the smallest value $\zeta_{s_{i+1}}^i$. If $\zeta_{s_{i+1}}^i > 0$, we add the set $A_{i+1} = S^i$ to the collection with $u_{i+1} = \zeta_{s_{i+1}}^i$ and continue.

Clearly, this process terminates in at most $m = |S|$ steps, resulting in a collection of size at most m. Each step involves updating the set of remaining elements, computing the new values $\zeta_s^{(\cdot)}$ and finding the minimum of these values. Thus every iteration requires at most $O(m)$ time and the overall algorithm takes at most $O(m^2)$ time. Due to the nested nature of the sets and downward closedness of $\mathcal{F}$, condition (i) is satisfied for every set added to the collection. It is easy to verify that condition (iii) is satisfied for every element by induction. We established the base case for element s_1 in the first iteration. Suppose that the property holds for all elements $s_1, s_2, \cdots, s_{i-1}$. Then, by the definition of u_i we have for element s_i,

$$\sum_{j=1}^{i} u_j \phi(A_j, s_i) = \big(p_s - \sum_{j=1}^{i-1} u_j q_{s_i}^{j-1}\big) + \sum_{j=1}^{i-1} u_j q_{s_i}^{j-1} = p_{s_i}.$$

Since s_i is excluded from all future sets added to the collection, this completes the induction for (iii). Finally, to prove property (ii) it suffices to show that,

$$u_m \leq \zeta_{s_{m-1}}^0 - \sum_{i=1}^{m-1} u_i,$$

as this immediately implies, $\sum_{i\in[m]} u_i \leq \zeta_{s_m}^0 \leq 1$. The desired inequality follows by substituting u_m and using the following facts: (i) $q_{s_m}^j$ is non-decreasing in j due to substitutability, (ii) $u_i \geq 0$ for every $i \in [m]$ since we perform iteration i only if $u_i = \zeta_{s_i}^{i-1} > 0$. Therefore,

$$u_m = \frac{p_{s_m} - \sum_{i=1}^{m-1} u_i q_{s_m}^{i-1}}{q_{s_m}^{m-1}} \leq \frac{p_{s_m} - q_{s_m}^0 \sum_{i=1}^{m-1} u_i}{q_{s_m}^0} = \zeta_{s_m}^0 - \sum_{i=1}^{m-1} u_i.$$

□

The probability matching algorithm be executed more efficiently for the commonly used MNL choice model. Using properties of MNL it suffices to sort the resources in order of values ζ_s^0 in the beginning and this ordering does not change as we remove more and more elements. Each iteration only takes $O(1)$ time and so the process has runtime dominated by sorting a set of size $m = |S|$, i.e., $O(m \log m)$.

Proof of Theorem 15. Recall that $\gamma = \min_{(i,t)\in E} \frac{c_i}{b_{it}}$, and $y(k_i, t)$ is the total fraction of unit k_i that is fluidly chosen by arrival t in Assort G-ALG. The proof rests simply on showing that for every $i \in I$ and $t \in T$, at least c_i/γ $(\geq b_{it})$ units of i are available at t w.p. at least $1 - 1/\gamma$. Conditioned on this, from Lemma 16 we have that in Sample Assort G-ALG, i is offered to and chosen by arrival t w.p. $\frac{1}{b_{it}(1+\delta)} \sum_{k\in[c_i]} y(k,t)$. This implies a lower bound of $\frac{1-1/\gamma}{1+\delta} \sum_{k\in[c_i]} y(k,t)$ on the expected reward from t choosing i, completing the proof.

Let $x_{it} := \frac{1}{b_{it}(1+\delta)} \sum_{k\in[c_i]} y(k,t)$. To show that i is available at t w.h.p., we first draw out some hidden independence in the events of concern. Recall that if less than b_{it} units of i are available at t then Sample Assort G-ALG does not offer i in any (randomized) assortment to t. Otherwise, Probability Matching (Algorithm 3) ensures that i is *offered and chosen* by arrival t w.p. exactly x_{it}. Now, consider the following alternative process at every arrival,

1. Given collection $\{A(\eta,t), u(\eta,t)\}_\eta$ from Assort G-ALG, sample assortment $A(\eta,t)$ independently w.p. $\frac{1}{1+\delta} u(\eta,t)$.
2. Offer $A(\eta,t)$ and it customer chooses resource i with insufficient inventory, reject the customer request.

We refer to this alternative process as the *static* process. Observe that the static process is probabilistically identical to Sample Assort G-ALG, which first checks the inventory of resources and then offers an assortment (after running probability matching). Thus, it suffices to show that in the static process, for every $i \in I$ and $t \in T$, at least c_i/γ units of i are available at t w.p. at least $1 - 1/\gamma$.

Now, consider an arbitrary resource i in the static process and let $\mathbb{1}(i \to t)$ indicate the event that i is offered to and chosen by t. W.l.o.g., we independently pre-sample usage durations for every possible match and let $\mathbb{1}(d_t > a(\tau) - a(t))$ indicate that the duration of usage pre-sampled for (a potential) match of i to arrival t is at least $a(\tau) - a(t)$. Let $\gamma_i = \frac{c_i}{\max_{t'\in T} b_{it'}}$. Observe that $\gamma_i \leq \gamma \ \forall i \in I$. The static process never fails to satisfy customer request for resource i if,

$$\sum_{\tau=1}^{t} b_{i\tau} \mathbb{1}(i \to \tau) \mathbb{1}(d_\tau > a(t) - a(\tau)) \leq c_i(1 - 1/\gamma_i) \quad \forall t \in T.$$

Define binary (not necessarily Bernoulli) random variables $X_\tau = \frac{b_{i\tau}}{\max_{t'\in T} b_{it'}} \mathbb{1}(i \to \tau)\mathbb{1}(d_\tau > a(t) - a(\tau))$ for all $\tau \leq t - 1$. Random variables X_τ are independent of each other as the assortment sampled, customer choice, and the duration of usage are all independently sampled at each arrival (in the static process). From Assort G-ALG, we have the following upper bound on the total expectation,

$$\mu := \mathbb{E}\Big[\sum_{\tau=1}^{t} X_\tau\Big] = \sum_{\tau=1}^{t} \frac{b_{i\tau}}{\max_{t'\in T} b_{it'}} x_{i\tau}\big(1 - F_i(a(t) - a(\tau))\big) \leq \frac{\gamma_i}{1+\delta}.$$

Applying the generalized Chernoff bound as stated in Lemma D2, completes the proof. □

Lemma D2 *Given integer $\tau > 0$, real value $\gamma > 0$, independent binary random variables $X_t \in \{0, x_t\}$ with $x_t \in (0,1]$ $\forall t \in [\tau]$ and $\sum_{t=1}^{\tau} \mathbb{E}[X_t] \leq \frac{\gamma}{1+\delta}$, where $\delta = \sqrt{\frac{\log \gamma}{\gamma}}$. We have,*

$$\mathbb{P}\left(\sum_{t=1}^{\tau} X_t > \gamma - 1\right) \leq \frac{1}{\sqrt{\gamma}}.$$

Proof. Consider Bernoulli random variables Y_t such that $E[Y_t] = x_t p_t$ $\forall t \in [\tau]$ and $\sum_{t=1}^{\tau} E[Y_t] \leq \frac{\gamma}{1+\delta}$. From Lemma 12, we have

$$\mathbb{P}\left(\sum_{t=1}^{\tau} Y_t > \gamma - 1\right) \leq \frac{1}{\sqrt{\gamma}}.$$

The key step in deriving the Chernoff bound that underlies this inequality is the following upper bound on the moment generating function of Y_t (Goemans 2015). For $s \geq 0$,

$$E[e^{sY_t}] = (x_t p_t)\, e^s + (1 - x_t p_t) = 1 + x_t p_t\, (e^s - 1) \leq e^{x_t p_t\, (e^s - 1)}. \tag{23}$$

Using this upper bound along with Markov's inequality and independence of random variables, gives the desired bound. Thus, to show the desired bound for independent (non-Bernoulli) random variables $X_t \in \{0, x_t\}$ (where $E[X_t] = x_t p_t$), it suffices to establish the upper bound on moment generating functions given by (23). For $s \geq 0$, we have

$$E[e^{sX_t}] = p_t\, e^{s x_t} + (1 - p_t) = 1 + p_t\, (e^{s x_t} - 1).$$

Unlike (23), here we have the term x_t in the exponent. However, for $s \geq 0$, and x_t, $p_t \in [0,1]$, we have, $p_t\, (e^{s x_t} - 1) \leq e^{p_t x_t\, (e^s - 1)} - 1$, giving us the same upper bound as (23). This completes the proof. □

Appendix E: Missing Details for RBA

E.1. Boundedness of Common IFR Distributions

Given,

$$L_i(\epsilon) = \max_{x \geq 0} \frac{F_i\big(x + F_i^{-1}(\epsilon)\big) - F_i(x)}{\epsilon},$$

we show that RBA is asymptotically $(1 - 1/e)$–competitive if $L_i(\epsilon)\epsilon \to$ as $\epsilon \to 0$. In particular, if $\epsilon L_i(\epsilon)$ decreases to 0 as strongly as $O(\epsilon^{\eta})$ for some $\eta > 0$, then we have a polynomial convergence rate of $\tilde{O}(c_i^{-\frac{\eta}{1+\eta}})$. We show that this holds for many common IFR families (Chapter 2 of Barlow and Proschan (1996)). An approximation for $L(\epsilon)\epsilon$ that eases calculations is given by $\{F^{-1}(\epsilon) \max_{x \geq 0} f(x)\}$ and $F(x) \to O(1) x f(x)$ for small x.

- **For exponential, uniform and IFR families with non-increasing density:** it is easy to see that $L_i(\epsilon) = 1$ and thus, $\eta = 1$. This implies a $\tilde{O}(c_i^{-0.5})$ convergence.

- **Weibull distributions:** This family is characterized by two non-negative parameters λ, k with c.d.f. given by, $F(x) = 1 - e^{(-x/\lambda)^k}$ for $x \geq 0$. The family is IFR only for $k \geq 1$. It is easy to see that,

$$L(\epsilon)\epsilon = O(1)\epsilon^{1/k}.$$

Hence, for any finite λ, k we have $\tilde{O}(c_i^{-\frac{1}{1+k}})$ convergence to $(1-1/e)$.

- **Gamma distributions:** This family is given by non-negative parameters k, θ such that the c.d.f. is $F(x) = \frac{1}{\Gamma(k)}\gamma(k, x/\theta)$. Here Γ, γ are the upper and lower incomplete gamma functions resp.. The family is IFR only for $k \geq 1$. It can be shown that for small ϵ, $L(\epsilon)\epsilon \to O(1)\epsilon^{1/k}$. Thus, we have $\tilde{O}(c_i^{-\frac{1}{1+k}})$ convergence for any finite set of parameter values.

- **Modified Extreme Value Distributions:** This family is characterized by the density function, $f(x) = \frac{1}{\lambda}e^{-\frac{e^x-1}{\lambda}+x}$ where parameter $\lambda > 0$. For small ϵ we have that $L(\epsilon)\epsilon \to O(1)\epsilon$, for any finite λ. Giving a a $\tilde{O}(c_i^{-0.5})$ convergence.

- **Truncated Normal:** For truncated normal distribution with finite mean μ, finite variance σ^2, and support $[0, b]$ where b could $\to \infty$, it is easy to see that $L(\epsilon) = O(1)$, leading to a $\tilde{O}(c_i^{-0.5})$ convergence. Note that if the support is given by $[a, b]$ for some $a > 0$, then we do not have convergence as $L(\epsilon)\epsilon \to O(1)$.

E.2. Detailed Overview of Analysis of RBA

We devote this section to setting up the overall framework for analyzing RBA and postpone detailed proofs to Appendix E.3. We start by introducing (and recalling) important notation. We use ω to denote the sample path of usage durations in OPT and $O(\omega, i)$ to denote the set of all arrivals matched to i in OPT on sample path ω. Similarly, ν denotes a sample path of usage durations in RBA. Recall that RBA tracks the highest available unit for each resource, denoted as $z_i(t)$ for resource i at arrival t. In fact, since RBA is realization dependent $z_i(t)$ is more accurately written as $z_i(\nu, t)$. Given set of resources S_t with an edge to t, RBA matches according to the following simple rule,

$$\arg\max_{i \in S_t} r_i \left(1 - g\left(\frac{z_i(t)}{c_i}\right)\right).$$

Let $D(t)$ (technically $D(\nu, t)$) denote the resource matched to t by RBA on sample path ν. Let $z_{D(t)}$ denote the highest available unit of resource $D(t)$ at t. Finally, recall that $\Delta g(k) = g(\frac{k-1}{c_i}) - g(\frac{k}{c_i}) = e^{-\frac{k}{c_i}}(e^{\frac{1}{c_i}} - 1)$.

Recall from (17), we use the following candidate solution for the LP free certificate,

$$\theta_i = r_i\,\mathbb{E}_{\boldsymbol{\nu}}\Big[\sum_{t|D(t)=i} g\Big(\frac{z_i(t)}{c_i}\Big)\Big],\ \forall i \in I \text{ and } \lambda_t = \mathbb{E}_{\boldsymbol{\nu}}\Big[r_{D(t)}\Big(1 - g\Big(\frac{z_{D(t)}}{c_{D(t)}}\Big)\Big)\Big],\ \forall t \in T.$$

Similar to the analysis of G-ALG in Section 3.2, we defined λ_t and θ_i as deterministic quantities. Since λ_t is independent of ω and ν, it suffices to use the simplified system of inequalities given by (6) and (7). The candidate solution above satisfies condition (6) by definition, with $\beta = 1$. To prove conditions (7) are satisfied with $\alpha_i = (1-1/e)$ for every $i \in I$, we start by lower bounding the term $\mathbb{E}_{\boldsymbol{\omega}}[\sum_{t \in O(\omega, i)} \lambda_t]$. The following lemma brings out the key source of difficulty. A proof of the lemma is in Appendix E.3.

LEMMA 14. *Given λ_t as defined by* (17)*, we have for every resource i,*

$$\mathbb{E}_{\boldsymbol{\omega}}\Big[\sum_{t\in O(\omega,i)}\lambda_t\Big]\geq (1-1/e)\mathrm{OPT}_i - r_i\mathbb{E}_{\boldsymbol{\omega}}\Big[\sum_{t\in O(\omega,i)}\sum_{k=1}^{c_i}\Delta g(k)\mathbb{P}_{\boldsymbol{\nu}}\big[k>z_i(t)\big]\Big]. \tag{24}$$

Consider a resource $i\in I$ and let k denote a unit of i. Contrast decomposition (24) with its counterpart given in Lemma 9. Deterministic binary valued quantities $\mathbb{1}(\neg k, t^+)$, that signified unavailability of a unit k after arrival t is matched in G-ALG, have been replaced by (non-binary) probabilities $\mathbb{P}_{\boldsymbol{\nu}}[k>z_i(t)]$. For any given arrival t, probability $\mathbb{P}_{\boldsymbol{\nu}}[k>z_i(t)]$ represents the likelihood that unit k and all higher units of i are unavailable in RBA when t arrives. This dependence on other units, combined with the sensitivity of $z_i(t)$ to small changes on the sample path, makes it challenging to bound these probabilities in a meaningful way. To address these challenges we introduce two new ingredients. The first ingredient simplifies the stochastic dependencies by introducing a conditional version of the probability $\mathbb{P}_{\boldsymbol{\nu}}[k>z_i(t)]$. This ingredient uses special structure that is only available in settings beyond online matching (see Appendix E.6 for further discussion). The second ingredient builds on the first one and addresses the non-binary nature of these probabilities.

Ingredient 1: Conditioning. Recall that to analyze G-ALG, we captured the allocation of each individual unit via a natural $(F,\boldsymbol{\sigma},\boldsymbol{p})$ random process. In an $(F,\boldsymbol{\sigma})$ random process, at every $\sigma_t\in\boldsymbol{\sigma}$ the only event of relevance is the availability of the item at σ_t (as probability $p_t=1$). If the item is available, then it is matched to σ_t regardless of the usage durations realized before σ_t. In order to naturally capture the actions of RBA on individual units through a random process, a similar property must hold in RBA. To be more specific, consider a single unit k of i in RBA and fix the randomness associated with all units and resources except k. Suppose we have an arrival t with edge to i and two distinct sample paths over usage durations of k prior to the arrival of t. We are also given that k is available at t on both sample paths. Clearly, RBA makes a deterministic decision on each sample path. Is it possible that RBA matches t to k on one sample path but not the other?

Interestingly, in the case of matching we can show that this is impossible, i.e., RBA always makes the same decision on both sample paths in such instances. We show that this is not the case in general. In particular, when we include the aspect of customer choice or budgeted allocations, reusability can interact with these elements in an undesirable way (see example in Section E.6). For online matching we leverage this property to show that conditioned on the randomness of all other resources and units, the allocation of k in RBA is characterized by an $(F_i,\boldsymbol{\sigma})$ random process.

Let us formalize the discussion above. First, since usage durations are independently sampled we generate ν as follows: For each edge, we independently sample a usage duration for every unit of the resource that the edge is incident on. RBA only sees the samples that correspond to units that are matched and does not see the other samples. Thus, sample path ν includes a set of c_i usage durations for every edge (i,t). If RBA matches t to unit k on ν then, the usage duration annotated for k is realized out of the c_i samples for edge (i,t). Let ν_k be the collection of all samples annotated for unit k on path ν. Thus, ν_k is a sample path of usage durations of k. Let ν_{-k} denote the sample path of all units and resources except unit k of i. We use $\mathbb{P}_{\boldsymbol{\nu_k}}$ and $\mathbb{E}_{\boldsymbol{\nu_k}}$ to denote probability and expectation over randomness in sample paths ν_k. Define indicators,

$$\mathbb{1}(k,t) = \begin{cases} 1 & \text{if unit } k \text{ is available when } t \text{ arrives,} \\ 0 & \text{otherwise.} \end{cases}$$

Now, consider the conditional version of probability $\mathbb{P}_{\boldsymbol{\nu}}[k > z_i(t)]$ in (24),

$$\mathbb{P}_{\boldsymbol{\nu_k}}[k > z_i(t) \mid \nu_{-k}] = \mathbb{E}_{\boldsymbol{\nu_k}}[1 - \mathbb{1}(k,t) \mid \nu_{-k}] \times \prod_{k'>k}^{c_i} [(1 - \mathbb{1}(k',t)) \mid \nu_{-k}]. \tag{25}$$

For every $k' \neq k$, the conditional indicator $\mathbb{1}(k',t) \mid \nu_{-k}$ is deterministic. $\mathbb{E}_{\boldsymbol{\nu_k}}[\mathbb{1}(k,t) \mid \nu_{-k}]$ is the likelihood of k being available at t conditioned on fixing the usage durations of all other units. We show that it suffices to condition the expectation only on a certain subset of units as opposed to all units. This subset consists of all units that strictly precede k in the ordering defined below.

Order of Units: Unit k_i of resource i precedes unit k_j of resource j, denoted $k_i \succ k_j$, iff $r_i(1-g(k_i/c_i)) > r_j(1-g(k_j/c_j))$, i.e., RBA would prefer to match k_i rather than k_j if both were available. In case of a tie we let the unit with the lower resource index precede and let RBA follow the same tie breaking rule. Observe that this ordering is transitive.

Let $D(t)$ denote the unit matched to t by RBA. Given edge (i,t) and a unit k of i, let $\mathbb{1}_\nu[k \succeq D(t)]$ be the indicator for the event $k \succeq D(t)$ in RBA. Notice that if $\mathbb{1}_\nu[k \succeq D(t)] = 1$, then on sample path ν, all units with an edge to t that precede k are unavailable when t arrives. If k is available at t and $k \succeq D(t)$, then we have $D(t) = k$.

Lemma E3 *For every edge (i,t) and unit k of i, the random variable $\mathbb{1}_\nu[k \succeq D(t)]$ is independent of the usage duration of k and all units that succeed k.*

See Appendix E.3 for the proof. Let k^+ denote all units that strictly precede k and let ν_{k^+} denote the sample path of usage durations for these units. As a consequence of Lemma E3, we have the following properties for RBA.

Corollary E4 *Given arrival t with edge to i and a unit $k \in [c_i]$, we have $\mathbb{P}_{\boldsymbol{\nu_k}}[k > z_i(t) \mid \nu_{k^+}] = \mathbb{P}_{\boldsymbol{\nu_k}}[k > z_i(t) \mid \nu_{-k}]$.*

The corollary follows by using Lemma E3 to transform the RHS of equation (25). See Appendix E.3 for the proof.

Corollary E5 *For every unit k, given a sample path ν_{k^+} we have an ordered set of arrivals $\boldsymbol{\sigma}(\nu_{k^+}) = \{\sigma_1, \cdots, \sigma_e\}$ with $\sigma_1 < \cdots < \sigma_e$ such that for any arrival t,*

- *If $t \notin \boldsymbol{\sigma}(\nu_{k+})$, then conditioned on ν_{k+}, the probability that k is matched to t is 0.*
- *If $t \in \boldsymbol{\sigma}(\nu_{k+})$, then conditioned on ν_{k+}, k is matched to t w.p. 1 if it is available.*

Therefore, conditioned on ν_{k+}, the $(F_i, \boldsymbol{\sigma}(\nu_{k+}))$ random process fully characterizes the matches of k in RBA.

See Appendix E.3 for the proof. As an immediate consequence of Corollary E5, $r(F_i, \boldsymbol{\sigma}(\nu_{k+}))$ gives the expected number of times k is matched in RBA conditioned on sample sample path ν_{k+}. We remind the reader that technically the set $\boldsymbol{\sigma}$ in an $(F, \boldsymbol{\sigma})$ random process is a set of arrival times. We are using arrivals and arrival times interchangeably for convenience.

Corollary E6 *For every arrival t with an edge to unit k of i, conditioned on ν_{k+} and also on the availability of unit k at t, the event $D(t) = k$ is independent of the (past) usage durations of k.*

The corollary follows from the fact that conditioned on ν_{k+}, the $(F_i, \boldsymbol{\sigma}(\nu_{k+}))$ random process fully characterizes the matches of k in RBA.

Ingredient 2: Covering. If the probabilities $\mathbb{P}_{\boldsymbol{\nu}}[k > z_i(t) \mid \nu_{k+}]$ were all either 1 or 0, the rest of the analysis could proceed along the same lines as the analysis of G-ALG. Of course, in general these probabilities have non-binary values and are hard to meaningfully bound. To address this challenge we introduce the ingredient of *covering*. For intuition behind the idea, consider an instance where RBA matches all n units of a resource to arrivals in an infinitesimal interval $[0, \epsilon]$. Each unit returns after 1 unit of time with probability 0.5 and never returns with probability 0.5. Let $\mathbb{1}(k, 2)$ indicate that unit k of the resource is available at time 2. We have,

$$\mathbb{P}_{\boldsymbol{\nu}}[\mathbb{1}(k,2)] = \frac{1}{2} \text{ and therefore, } \mathbb{P}_{\boldsymbol{\nu}}[k > z_1(2)] < \frac{1}{2^k}.$$

This implies that $z_1(2) \geq n - O(\log n)$ w.h.p.. Therefore, the fact that all units are available at t with a constant probability implies that the probability $\mathbb{P}_{\boldsymbol{\nu}}[k > z_1(2)]$ decreases geometrically as k increases, resulting in a small value ($O(\log n)$) for the negative term in (24). Inspired by this, for every k and sample path ν_{k+}, we classify all arrivals t into two groups. Roughly speaking, the first group consists of arrivals where k is available with sufficiently high probability. These are called *uncovered* arrivals. We show that the overall contribution to (24) from this group of arrivals, summed appropriately over all units k, is well approximated by the expectation of a geometric random variable. This generalizes the bound that we observed in the example above. The remaining *covered* arrivals, are the ones where the probability that k is available is low. We show that the contribution from these terms to (24) is effectively canceled out by θ_i , i.e., these arrivals are *covered* by θ_i.

To describe this formally, for values $\epsilon_i \in (0,1]$, let $\mathcal{X}_k(\nu_{k+}, \epsilon_i, t) \in \{0,1\}$ denote the *covering function* that performs this classification.

Fix unit k, a value $\epsilon_i \in (0,1]$, and condition on sample path ν_{k+}. Given an arrival t, we say that t is uncovered and set $\mathcal{X}_k(\nu_{k+}, \epsilon_i, t) = 0$ if,

$$\mathbb{P}_{\boldsymbol{\nu_k}}[\mathbb{1}(k,t) = 1 \mid \nu_{k+}] \geq \epsilon_i,$$

i.e., at an uncovered arrival we have a lower bound of ϵ_i on the probability of k being available. Equivalently, this condition imposes a lower bound of ϵ_i on the probability of k being free at time $a(t)$ in an $(F_i, \boldsymbol{\sigma}(\nu_{k+}))$ random process.

Notice that as a direct consequence of the above definition, for every covered arrival t we have $\mathcal{X}_k(\nu_{k+}, \epsilon_i, t) = 1$ and $\mathbb{P}_{\boldsymbol{\nu_k}}[\mathbb{1}(k,t) = 1 \mid \nu_{k+}] < \epsilon_i$. For ease of notation we omit the input ϵ_i in covering functions and use the abbreviation,

$$\mathcal{X}_k(\nu_{k+}, t),$$

with the understanding that ϵ_i is present and will be set later to optimize the result. Note that a given arrival t may be covered w.r.t. one unit k of i but uncovered w.r.t. another unit k' of i.

Combining the Ingredients: The following lemma uses the notions of covering and conditioning to further decompose and upper bound (24). Proof of the lemma is deferred to Appendix E.3.

Lemma E7 *For every resource i,*

$$\mathbb{E}_{\boldsymbol{\omega}}\Big[\sum_{t \in O(\omega,i)} \sum_{k=1}^{c_i} \Delta g(k) \mathbb{P}_{\boldsymbol{\nu}}\big[k > z_i(t)\big]\Big] \leq \mathbb{E}_{\boldsymbol{\omega}}\Big[\sum_{t \in O(\omega,i)} \Big(\sum_{k=1}^{c_i} \Delta g(k) \mathbb{E}_{\boldsymbol{\nu_{k+}}}\big[\mathcal{X}_k(\nu_{k+}, t)\big]\Big)\Big] + \frac{2}{r_i \epsilon_i c_i} \mathrm{OPT}_i.$$

Remarks: The term $\frac{2}{r_i \epsilon_i c_i}\mathrm{OPT}_i$ on the RHS is an upper bound on the contribution of the terms arising due to uncovered arrivals. As we discussed when motivating the second ingredient, this bound is a result of geometrically decreasing probability of multiple uncovered units being jointly unavailable at t. The formal proof of this bound relies crucially on the property that the usage durations of units $k' \prec k$ do not affect the matching decision of RBA for unit k (the first ingredient).

The other term on the RHS of inequality in Lemma E7 captures the contribution from covered arrivals. As we informally stated earlier, we shall upper bound this term by θ_i, i.e.,

$$\mathbb{E}_{\boldsymbol{\omega}}\Big[\sum_{t \in O(\omega,i)} \Big(\sum_{k=1}^{c_i} \Delta g(k) \mathbb{E}_{\boldsymbol{\nu_{k+}}}\big[\mathcal{X}_k(\nu_{k+}, t)\big]\Big)\Big] \leq \frac{1 + o(1)}{r_i} \theta_i, \tag{26}$$

where $o(1)$ represents a term that goes to 0 as $c_i \to +\infty$. An obstacle to proving this inequality is the dependence on set $O(\omega, i)$, which is determined by actions of OPT. At a high level, we tackle this difficulty in the same way as the proof of Lemma 4. We untangle the dependence on OPT in the LHS of (26) by upper bounding it as the expected reward of an $(F, \boldsymbol{\sigma})$ random process. From Corollary E5 we have a way of lower bounding θ_i by the expectation of a (different) random processes. Consequently, showing (26) boils down to proving a new perturbation property of $(F, \boldsymbol{\sigma})$ random processes.

Upper bounding the LHS in (26)**:** Consider the set $\boldsymbol{s}(\nu_{k^+})$ of all arrivals t that are covered i.e.,

$$\boldsymbol{s}(\nu_{k^+}) = \{t \in T \mid \mathcal{X}_k(\nu_{k^+}, t) = 1\}.$$

The following lemma shows that the quantity, $r\big(F_i, \boldsymbol{s}(\nu_{k^+})\big)$ is an upper bound on the contribution from covered arrivals to (24).

Lemma E8 *For any resource i, unit k, and path ν_{k^+}, we have,*

$$\mathbb{E}_{\boldsymbol{\omega}}\Big[\sum_{t \in O(\omega,i)} \mathcal{X}_k(\nu_{k^+}, t)\Big] \le c_i r\big(F_i, \boldsymbol{s}(\nu_{k^+})\big).$$

Consequently,

$$\mathbb{E}_{\boldsymbol{\omega}}\Big[\sum_{t \in O(\omega,i)} \Big(\sum_{k=1}^{c_i} \Delta g(k) \mathbb{E}_{\boldsymbol{\nu}_{\boldsymbol{k}^+}}\left[\mathcal{X}_k(\nu_{k^+}, t)\right]\Big)\Big] \le \Big(1 + \frac{2}{c_i}\Big) \sum_{k=1}^{c_i} \Big(g\Big(\frac{k}{c_i}\Big) \cdot \mathbb{E}_{\boldsymbol{\nu}_{\boldsymbol{k}^+}}\left[r\big(F_i, \boldsymbol{s}(\nu_{k^+})\big)\right]\Big).$$

The proof of this lemma is included in Appendix E.3 and closely mimics the two list coupling argument used to establish the inequality $\mathbb{E}_{\boldsymbol{\omega}}\Big[\sum_{t \in O(\omega, k_O)} \mathbb{1}(\neg k, t^+)\Big] \le r(F_i, \boldsymbol{s}(k))$, in the analysis of G-ALG. To finish the overall analysis of RBA, we need the following new property for random processes.

Proposition 1 (Perturbation property) *For every resource i, every unit k of i, every path ν_{k^+}, and parameter $\epsilon_i = \frac{1}{o(c_i)}$, we have*

$$\boxed{r\big(F_i, \boldsymbol{s}(\nu_{k^+})\big) \le \big(1 + \kappa_i(\epsilon_i, c_i)\big) r\big(F_i, \boldsymbol{\sigma}(\nu_{k^+})\big),} \tag{27}$$

for some non-negative function κ_i that approaches 0 for $c_i \to +\infty$.

Recall that given set of arrivals T, subset $\boldsymbol{\sigma}(\nu_{k^+})$, and value ϵ_i, the set $\boldsymbol{s}(\nu_{k^+})$ of covered arrivals is fully determined by the $\big(F_i, \boldsymbol{\sigma}(\nu_{k^+})\big)$ random process. Thus, inequality (27) is purely a statement about an explicitly defined random process. Note, the requirement that $\epsilon_i = \frac{1}{o(c_i)}$ or $\frac{1}{\epsilon_i c_i} \to 0$ for $c_i \to +\infty$, is clearly a necessity due to the term $\frac{2}{r_i \epsilon_i c_i}\mathrm{OPT}_i$ in Lemma E7. In general, the difficulty in proving (27) is that we require ϵ_i to be suitably large and also desire κ_i to be small. Larger ϵ_i leads to more arrivals $\boldsymbol{s}(\nu_{k^+}) \backslash \boldsymbol{\sigma}(\nu_{k^+})$, leading to a (possibly) larger value of $r\big(F_i, \boldsymbol{s}(\nu_{k^+})\big)$. This tension between the two parameter values is made explicit through the competitive ratio guarantee in the next lemma.

Lemma E9 *Given, values θ_i and λ_t set according to* (17)*. For every usage distribution such that Proposition 1 holds, we have*

$$\theta_i + \mathbb{E}_{\boldsymbol{\omega}}\Big[\sum_{t \in O(\omega,i)} \lambda_t\Big] \ge \alpha_i \mathrm{OPT}_i, \quad \forall i \in I,$$

with values $\alpha_i = \frac{\left(1 - 1/e - \frac{2}{\epsilon_i c_i}\right)}{\left(1 + \kappa_i(\epsilon_i, c_i)\right)(1 + 2/c_i)}$. If we can choose ϵ_i, κ_i such that, $\frac{1}{\epsilon_i c_i} \to 0$ and $\kappa_i(\epsilon_i, c_i) \to 0$ for $c_i \to +\infty$, then RBA *is asymptotically $(1 - 1/e)$–competitive, with the rate of convergence given by,*

$$O\Big(\min_i \Big\{\frac{1}{\epsilon_i c_i} + \kappa_i(\epsilon_i, c_i)\Big\}\Big).$$

Proof. For convenience, we refer to $\kappa_i(\epsilon_i, c_i)$ simply as κ_i respectively. By definition of θ_i in (17),

$$\begin{aligned}\frac{1}{r_i}\theta_i &= \mathbb{E}_{\boldsymbol{\nu}}\Big[\sum_{t|D(t)=i} g\Big(\frac{z_i(t)}{c_i}\Big)\Big],\\ &= \sum_{k=1}^{c_i} g\Big(\frac{k}{c_i}\Big) \sum_{t|D(t)=i,\, z_i(t)=k} 1,\\ &= \sum_{k=1}^{c_i} g\Big(\frac{k}{c_i}\Big)\mathbb{E}_{\boldsymbol{\nu_{k+}}}\big[r\big(F_i, \boldsymbol{\sigma}(\nu_{k+})\big)\big],\end{aligned}$$

where the last equality follows from Corollary (E5). Combining this with Lemma E7, Lemma E8 and Proposition 1, we have,

$$\mathbb{E}_{\boldsymbol{\omega}}\Big[\sum_{t\in O(\omega,i)}\sum_{k=1}^{c_i}\Delta g(k)\mathbb{P}_{\boldsymbol{\nu}}\big[k > z_i(t)\big]\Big] \leq \Big(1+\frac{2}{c_i}\Big)(1+\kappa_i)\,\frac{1}{r_i}\theta_i + \frac{2}{r_i\epsilon_i c_i}\mathrm{OPT}_i.$$

Substituting this in Lemma 14, we get,

$$\mathbb{E}_{\boldsymbol{\omega}}\Big[\sum_{t\in O(\omega,i)}\lambda_t\Big] \geq \Big(1-1/e-\frac{2}{\epsilon_i c_i}\Big)\mathrm{OPT}_i - \Big(1+\frac{2}{c_i}\Big)(1+\kappa_i)\,\theta_i.$$

□

The final step in proving Theorem 13 is to establish the Perturbation property (27) for the various families of usage distributions discussed in Section 4. In Appendix E.4, we prove (27) for these families of distributions, starting with the simplest case of two-point distributions. These proofs rely on using the knowledge of distribution F_i to further characterize the set $\boldsymbol{s}(\nu_{k+})$, followed by carefully constructed coupling arguments that also critically use the structure of F_i.

Note that to prove (27), we will typically show a stronger statement by considering a set $\boldsymbol{S}(\nu_{k+})$ that is a superset of $\boldsymbol{s}(\nu_{k+})$ and show that,

$$r\big(F_i, \boldsymbol{S}(\nu_{k+})\big) \leq \big(1+\kappa_i(\epsilon_i, c_i)\big)r\big(F_i, \boldsymbol{\sigma}(\nu_{k+})\big),$$

for a suitably small κ_i. Inequality (27) then follows by using the monotonicity property (Lemma 5). This will give us freedom to consider sets $\boldsymbol{S}(\nu_{k+})$ that are easier to characterize than $\boldsymbol{s}(\nu_{k+})$. In absence of any assumptions about the distribution F_i, it is not formally clear to us if (27) still holds, though we conjecture that it does.

E.3. Proofs of Individual Components

LEMMA 14. *Given λ_t as defined by* (17)*, we have for every resource i,*

$$\mathbb{E}_{\boldsymbol{\omega}}\Big[\sum_{t\in O(\omega,i)}\lambda_t\Big] \geq (1-1/e)\mathrm{OPT}_i - r_i\mathbb{E}_{\boldsymbol{\omega}}\Big[\sum_{t\in O(\omega,i)}\sum_{k=1}^{c_i}\Delta g(k)\mathbb{P}_{\boldsymbol{\nu}}\big[k > z_i(t)\big]\Big].$$

Proof. We start with the LHS,

$$\begin{aligned}
\mathbb{E}_{\boldsymbol{\omega}}\Big[\sum_{t\in O(\omega,i)}\lambda_t\Big] &= \mathbb{E}_{\boldsymbol{\omega},\boldsymbol{\nu}}\Big[\sum_{t\in O(\omega,i)} r_{D(t)}[1-g\Big(\frac{z_{D(t)}}{c_{D(t)}}\Big)]\Big],\\
&\overset{(a)}{\geq} \mathbb{E}_{\boldsymbol{\omega},\boldsymbol{\nu}}\Big[\sum_{t\in O(\omega,i)} r_i[1-g\Big(\frac{z_i(t)}{c_i}\Big)]\Big],\\
&\overset{(b)}{=} r_i\mathbb{E}_{\boldsymbol{\omega},\boldsymbol{\nu}}\Big[\sum_{t\in O(\omega,i)}\Big(1-1/e-\sum_{k=z_i(t)+1}^{c_i}\Delta g(k)\Big)\Big],\\
&= (1-1/e)\mathrm{OPT}_i - r_i\mathbb{E}_{\boldsymbol{\omega},\boldsymbol{\nu}}\Big[\sum_{t\in O(\omega,i)}\sum_{k=z_i(t)+1}^{c_i}\Delta g(k)\Big],
\end{aligned} \tag{28}$$

where (a) follows by the fact that RBA matches every arrival to the resource that has maximum reduced price. Equation (b) follows from the inequality,

$$1-g\Big(\frac{z_i(t)}{c_i}\Big) = (1-1/e) - \sum_{k=z_i(t)+1}^{c_i}\Delta g(k),$$

obtained by setting $k_i = z_i(t)$ in (9) (which states one of the properties of function $\Delta g(\cdot)$). Let $\mathbb{1}_\nu(k > z_i(t))$ indicate the event that $k > z_i(t)$. We rewrite the $\mathbb{E}_{\boldsymbol{\omega},\boldsymbol{\nu}}[\cdot]$ term in (28) as follows,

$$\begin{aligned}
\mathbb{E}_{\boldsymbol{\omega},\boldsymbol{\nu}}\Big[\sum_{t\in O(\omega,i)}\sum_{k=z_i(t)+1}^{c_i}\Delta g(k)\Big] &= \mathbb{E}_{\boldsymbol{\omega}}\Big[\sum_{t\in O(\omega,i)}\mathbb{E}_{\boldsymbol{\nu}}\Big[\sum_{k=z_i(t)+1}^{c_i}\Delta g(k)\Big]\Big],\\
&= \mathbb{E}_{\boldsymbol{\omega}}\Big[\sum_{t\in O(\omega,i)}\mathbb{E}_{\boldsymbol{\nu}}\Big[\sum_{k=1}^{c_i}\Delta g(k)\,\mathbb{1}_\nu(k > z_i(t))\Big]\Big],\\
&= \mathbb{E}_{\boldsymbol{\omega}}\Big[\sum_{t\in O(\omega,i)}\sum_{k=1}^{c_i}\Delta g(k)\mathbb{P}_{\boldsymbol{\nu}}\big[k > z_i(t)\big]\Big],
\end{aligned} \tag{29}$$

here the final equality follows by linearity of expectation. Combining (28) and (29) completes the proof.
□

LEMMA E3 (RESTATED). *Let $D(t)$ denote the resource matched to t by* RBA. *Given arrival t with edge to i let $\mathbb{1}_\nu[k \succeq D(t)]$ denote the event $k \succeq D(t)$ in* RBA. *Then, the event $\mathbb{1}_\nu[k \succeq D(t)]$ is independent of the usage durations of k and all units that succeed k.*

Proof. Suppose that the statement is false for some unit k. Then, there exists sample paths ν_1 and ν_2 such that $\mathbb{1}_{\nu_1}[k \succeq D(t)] = 1$ and $\mathbb{1}_{\nu_2}[k \succeq D(t)] = 0$ for some arrival t and ν_1 and ν_2 agree on the usage durations of all units except unit ℓ, for some $\ell \prec k$. For every unit $k' \succeq k$, the usage durations of k' are the same on both ν_1 and ν_2. Without loss of generality,

$$\mathbb{1}_{\nu_1}[k' \succeq D(\tau)] = \mathbb{1}_{\nu_2}[k' \succeq D(\tau)],$$

for every $k' \succeq k$ and every arrival τ prior to t. Thus, the matching decisions over units $k' \succeq k$ are identical prior to arrival t and when t arrives, the availability status of units $k' \succeq k$ is the same on both sample paths.

Now by definition of RBA, if t is matched to a unit $k' \succeq k$ on one path then it will be matched to the same unit on the other path, contradiction. □

COROLLARY E4. *Given arrival t with edge to i and a unit $k \in [c_i]$, we have $\mathbb{P}_{\boldsymbol{\nu_k}}[k > z_i(t) \mid \nu_{k+}] = \mathbb{P}_{\boldsymbol{\nu_k}}[k > z_i(t) \mid \nu_{-k}]$.*

Proof. Equality (25) states that,

$$\mathbb{P}_{\boldsymbol{\nu_k}}[k > z_i(t) \mid \nu_{-k}] = \mathbb{E}_{\boldsymbol{\nu_k}}[1 - \mathbb{1}(k,t) \mid \nu_{-k}] \times \prod_{k'>k}^{c_i} [(1 - \mathbb{1}(k',t)) \mid \nu_{-k}],$$

where $\mathbb{1}(k',t) \mid \nu_{-k}$ is deterministic for every $k' \neq k$. In fact, $\mathbb{1}(k',t) \mid \nu_{-k} = \mathbb{1}(k',t) \mid \nu_{k+} \ \forall k' > k$. From Lemma E3, we have, $\mathbb{E}_{\boldsymbol{\nu_k}}[\mathbb{1}(k,t) \mid \nu_{-k}] = \mathbb{E}_{\boldsymbol{\nu_k}}[\mathbb{1}(k,t) \mid \nu_{k+}]$. Thus,

$$\mathbb{P}_{\boldsymbol{\nu_k}}[k > z_i(t) \mid \nu_{-k}] = \mathbb{E}_{\boldsymbol{\nu_k}}[1 - \mathbb{1}(k,t) \mid \nu_{k+}] \times \prod_{k'>k}^{c_i} [(1 - \mathbb{1}(k',t)) \mid \nu_{k+}] = \mathbb{P}_{\boldsymbol{\nu_k}}[k > z_i(t) \mid \nu_{k+}].$$

□

COROLLARY E5. *For every unit k, given a sample path ν_{k+} we have an ordered set of arrivals $\boldsymbol{\sigma}(\nu_{k+}) = \{\sigma_1, \cdots, \sigma_e\}$ with $\sigma_1 < \cdots < \sigma_e$ such that for any arrival t,*

- *If $t \notin \boldsymbol{\sigma}(\nu_{k+})$, then conditioned on ν_{k+}, the probability that k is matched to t is 0.*
- *If $t \in \boldsymbol{\sigma}(\nu_{k+})$, then conditioned on ν_{k+}, k is matched to t w.p. 1 if it is available.*

Therefore, conditioned on ν_{k+}, the $(F_i, \boldsymbol{\sigma}(\nu_{k+}))$ random process fully characterizes the matches of k in RBA.

Proof. From Lemma E3 we have that the matching decisions over units k^+ (that strictly precede k) do not depend on usage durations of unit k and its successors. We condition on ν_{k+} and this fixes the state of units k^+ at every arrival. Let $\boldsymbol{\sigma}(\nu_{k+})$ denote the ordered set of arrivals that have an edge to k and are not matched to a unit preceding k. If arrival t has an edge to k and $t \notin \boldsymbol{\sigma}(\nu_{k+})$, then $D(t) \succ k$ and k is not matched to t. If $t \in \boldsymbol{\sigma}(\nu_{k+})$, then $k \succeq D(t)$ by definition and k is matched to t if it is available at t. The availability of k at $t \in \boldsymbol{\sigma}(\nu_{k+})$ depends only on the usage durations of k prior to t, which are sampled independently according to distribution F_i. Thus, the $(F_i, \boldsymbol{\sigma}(\nu_{k+}))$ random process fully characterizes the matches of k in RBA.

□

LEMMA E7. *For every resource i,*

$$\mathbb{E}_{\boldsymbol{\omega}}\Big[\sum_{t \in O(\omega,i)} \sum_{k=1}^{c_i} \Delta g(k) \mathbb{P}_{\boldsymbol{\nu}}\big[k > z_i(t)\big]\Big] \leq \mathbb{E}_{\boldsymbol{\omega}}\Big[\sum_{t \in O(\omega,i)} \Big(\sum_{k=1}^{c_i} \Delta g(k) \mathbb{E}_{\boldsymbol{\nu_{k+}}}\big[\mathcal{X}_k(\nu_{k+}, t)\big]\Big)\Big] + \frac{2}{r_i \epsilon_i c_i} \mathrm{OPT}_i.$$

Proof. We start with the observation, $\mathbb{P}_{\boldsymbol{\nu}}\big[k > z_i(t)\big] = \mathbb{E}_{\boldsymbol{\nu}_{\boldsymbol{k}+}}\big[\mathbb{P}_{\boldsymbol{\nu_k}}[k > z_i(t) \mid \nu_{k+}]\big]$. Using the classification of each arrival t afforded by the covering function, we get,

$$\begin{aligned}\mathbb{P}_{\boldsymbol{\nu_k}}[k > z_i(t) \mid \nu_{k+}] &= \mathcal{X}_k(\nu_{k+}, t)\ \mathbb{P}_{\boldsymbol{\nu_k}}[k > z_i(t) \mid \nu_{k+}] \\ &\quad + (1 - \mathcal{X}_k(\nu_{k+}, t))\ \mathbb{P}_{\boldsymbol{\nu_k}}[k > z_i(t) \mid \nu_{k+}] \\ &\le \mathcal{X}_k(\nu_{k+}, t) + (1 - \mathcal{X}_k(\nu_{k+}, t)) \cdot \mathbb{P}_{\boldsymbol{\nu_k}}[k > z_i(t) \mid \nu_{k+}], \end{aligned} \tag{30}$$

here we use the upper bound of 1 on probability $\mathbb{P}_{\boldsymbol{\nu_k}}[k > z_i(t) \mid \nu_{k+}]$ for the scenario where t is covered[12]. Now, by substitution,

$$\begin{aligned}&\mathbb{E}_{\boldsymbol{\omega}}\Big[\sum_{t \in O(\omega,i)} \sum_{k=1}^{c_i} \Delta g(k) \mathbb{P}_{\boldsymbol{\nu}}\big[k > z_i(t)\big]\Big] \le \\ &\mathbb{E}_{\boldsymbol{\omega}}\Big[\sum_{t \in O(\omega,i)} \Big(\sum_{k=1}^{c_i} \Delta g(k) \mathbb{E}_{\boldsymbol{\nu}_{\boldsymbol{k}+}}\big[\mathcal{X}_k(\nu_{k+}, t)\big]\Big)\Big] + \\ &\mathbb{E}_{\boldsymbol{\omega}}\Big[\sum_{t \in O(\omega,i)} \Big(\sum_{k=1}^{c_i} \Delta g(k) \mathbb{E}_{\boldsymbol{\nu}_{\boldsymbol{k}+}}\big[(1 - \mathcal{X}_k(\nu_{k+}, t)) \mathbb{P}_{\boldsymbol{\nu_k}}[k > z_i(t) \mid \nu_{k+}]\big]\Big)\Big].\end{aligned}$$

The first term in the inequality is as desired. The next lemma (Lemma E10) shows that the second term is further upper bounded by $\frac{1}{r_i \epsilon_i c_i}\mathrm{OPT}_i$. □

Let us recall the definition of the covering function $\mathcal{X}_k$. Fix a resource i and a value $\epsilon_i \in (0,1]$. Consider unit k of i and condition on sample path ν_{k+}. Given an arrival t, we say that t is uncovered and set $\mathcal{X}_k(\nu_{k+}, \epsilon_i, t) = 0$ if,

$$\mathbb{P}_{\boldsymbol{\nu_k}}[\mathbb{1}(k,t) = 1 \mid \nu_{k+}] \ge \epsilon_i,$$

i.e., at an uncovered arrival we have a lower bound of ϵ_i on the probability of k being available. Due to the importance of parameter ϵ_i in the next lemma, we use the full form of the covering function without omitting ϵ_i.

Lemma E10 *For any resource i and arrival t, we have,*

$$\sum_{k=1}^{c_i} \Delta g(k) \mathbb{E}_{\boldsymbol{\nu}_{\boldsymbol{k}+}}\big[(1 - \mathcal{X}_k(\nu_{k+}, \epsilon_i, t))\ \mathbb{P}_{\boldsymbol{\nu_k}}[k > z_i(t) \mid \nu_{k+}]\big] \le \frac{2}{\epsilon_i c_i}$$

Consequently,

$$\mathbb{E}_{\boldsymbol{\omega}}\Big[\sum_{t \in O(\omega,i)} \Big(\sum_{k=1}^{c_i} \Delta g(k) \mathbb{E}_{\boldsymbol{\nu}_{\boldsymbol{k}+}}\big[(1 - \mathcal{X}_k(\nu_{k+}, \epsilon_i, t)) \mathbb{P}_{\boldsymbol{\nu_k}}[k > z_i(t) \mid \nu_{k+}]\big]\Big)\Big] \le \frac{2}{r_i \epsilon_i c_i}\mathrm{OPT}_i.$$

The upper bound can be tightened to 0 for $\epsilon_i = 1$.

[12] As an aside from the proof, one might wonder if this is too loose an upper bound. Conditioned on sample path ν_{k+}, in general it can be shown that if t is covered in k then it is also covered in all units of i preceding k, conditioned on corresponding partial sample paths that are consistent with ν_{k+}. So the bound is reasonably tight.

Proof. Let us focus on the first part of the lemma. If $\epsilon_i = 1$, then for every unit k and sample path ν_{k^+}, conditioned on ν_{k^+}, we have,

$$1 - \mathcal{X}_k(\nu_{k^+}, 1, t) = 1 \quad \Rightarrow \quad \mathbb{P}_{\boldsymbol{\nu_k}}[\mathbb{1}(k,t) = 1 \mid \nu_{k^+}] = 1 \quad \Rightarrow \quad \mathbb{P}_{\boldsymbol{\nu_k}}[z_i(t) \geq k \mid \nu_{k^+}] = 1.$$

Thus, $(1 - \mathcal{X}_k(\nu_{k^+}, 1, t)) \times \mathbb{P}_{\boldsymbol{\nu_k}}[k > z_i(t) \mid \nu_{k^+}] = 0$, and we are done. So, let $\epsilon_i < 1$. Using $e^x \leq 1 + x + x^2$ for $x \in [0,1]$, we have that,

$$\Delta g(k) \leq (e^{\frac{1}{c_i}} - 1) \leq 1/c_i + 1/c_i^2 \leq 2/c_i, \text{ for } c_i \geq 1.$$

Therefore, to prove the first part of the lemma, it suffices to show that for every $\epsilon_i < 1$,

$$\sum_{k=1}^{c_i} \mathbb{E}_{\boldsymbol{\nu_{k^+}}} \Big[(1 - \mathcal{X}_k(\nu_{k^+}, \epsilon_i, t))\, \mathbb{P}_{\boldsymbol{\nu_k}}[k > z_i(t) \mid \nu_{k^+}] \Big] \leq \frac{1}{\epsilon_i}. \tag{31}$$

The proof relies on establishing that the LHS is upper bounded by the expectation of a geometric r.v. with success probability ϵ_i.

Fix an arbitrary arrival t and consider the following iterative process for generating sample paths ν. We sample the usage durations of units one by one in decreasing order of precedence over units. Suppose that as we generate ν, we also grow an array $\mathbf{b}(\nu)$ with binary values. $\mathbf{b}(\nu)$ is initially an empty array. At the step where usage durations of all units k^+ that strictly precede unit k of i have been sampled, the covering function $\mathcal{X}_k(\nu_{k^+}, \epsilon_i, t)$ is well defined. If t is not covered w.r.t. k on the sample path ν_{k^+} generated so far, then we add an entry to the array $\mathbf{b}$. Next, we sample the durations of k, i.e., generate ν_k. This determines the availability of k at arrival t. We set the new array entry to 1 if k is available at t and set the entry to 0 otherwise. Thereafter, we move to the next unit dictated by the precedence order. We repeat the same procedure whenever we reach a unit of resource i. For all other units, we simply generate a sample path of its usage durations and move to the next unit in order. At the end of the process (after visiting all units in this manner), we add an entry to the end of the array with value 1 so that there is always an entry with value 1 in the array.

Let $M(\nu)$ denote the first entry in the array that has value 1. We show that the LHS of (31) is upper bounded by the expected value of $M(\nu)$. Then, we prove that $\mathbb{E}_{\boldsymbol{\nu}}[M(\nu)] \leq \frac{1}{\epsilon_i}$. Let $b_k(\nu)$ denote the array entry (if any) generated at unit k. If no array entry is generated at k then we set $b_k(\nu) = \emptyset$. If an entry is generated at k, then $b_k(\nu)$ is a boolean random variable which takes value 1 on sample paths where k is available at t. Note that the value of $b_k(\cdot)$ depends only on ν_{k^+} and ν_k, and is independent of the usage durations of units that succeed k. Let $\mathbb{1}(b_k(\nu) \neq \emptyset)$ indicate the event that a new entry is added to the array when we generate the usage durations of unit k. By definition,

$$1 - \mathcal{X}_k(\nu_{k^+}, \epsilon_i, t) = \mathbb{1}(b_k(\nu) \neq \emptyset).$$

Let $\ell(\nu)$ denote the first unit of i where the corresponding array entry has value 1, i.e.,

$$b_{\ell(\nu)}(\nu) = 1 \quad \text{and} \quad b_k(\nu) \in \{0, \emptyset\} \; \forall k \in \{\ell(\nu)+1, \cdots, c_i\}.$$

If $\ell(\nu) \leq k$ then there exists a unit k' of i that is available at t and that precedes k. In this scenario, we have, $k \leq z_i(t)$. Therefore,

$$\mathbb{P}_{\boldsymbol{\nu_k}}[k > z_i(t) \mid \nu_{k+}] \leq \mathbb{P}_{\boldsymbol{\nu_k}}[\ell(\nu) > k \mid \nu_{k+}].$$

Using the observations above, we have the following upper bound on the LHS of (31),

$$\begin{aligned}
\sum_{k=1}^{c_i} \mathbb{E}_{\boldsymbol{\nu_{k+}}}\Big[(1 - \mathcal{X}_k(\nu_{k+}, \epsilon_i, t))\, \mathbb{P}_{\boldsymbol{\nu_k}}[k > z_i(t) \mid \nu_{k+}]\Big] &\leq \sum_{k=1}^{c_i} \mathbb{E}_{\boldsymbol{\nu_{k+}}}\big[\mathbb{1}(b_k(\nu) \neq \emptyset)\; \mathbb{P}_{\boldsymbol{\nu_k}}[\ell(\nu) > k \mid \nu_{k+}]\big], \\
&= \sum_{k=1}^{c_i} \mathbb{P}_{\boldsymbol{\nu}}\big[(b_k(\nu) \neq \emptyset) \;\wedge\; (\ell(\nu) > k)\big], \\
&\leq \sum_{k=1}^{c_i} \mathbb{P}_{\boldsymbol{\nu}}\big[(b_k(\nu) \neq \emptyset) \;\wedge\; (\ell(\nu) \geq k)\big], \\
&= \mathbb{E}_{\boldsymbol{\nu}}\left[\sum_{k=1}^{c_i} \mathbb{1}\left(b_k(\nu) \neq \emptyset \;\wedge\; \ell(\nu) \geq k\right)\right], \\
&= \mathbb{E}_{\boldsymbol{\nu}}[M(\nu)].
\end{aligned}$$

It remains to show that $\mathbb{E}_{\boldsymbol{\nu}}[M(\nu)] \leq \frac{1}{\epsilon_i}$. Since the last array entry is set to 1 we have $M(\nu) \in [c_i + 1]$ on every sample path ν. Recall that for any unit k of i, conditioned on ν_{k+}, a new entry is generated only if t is not covered w.r.t. unit k, i.e., $\mathcal{X}_k(\nu_{k+}, \epsilon_i, t) = 0$. By definition, $\mathbb{P}_{\boldsymbol{\nu_k}}[\mathbb{1}(k,t) = 1 \mid \nu_{k+}] \geq \epsilon_i$. When all units of i have been visited, we add a final entry to the array with value 1. Thus, conditioned on the previous array entries, any new entry in the array has value 1 w.p. at least ϵ_i and we get,

$$\mathbb{P}_{\boldsymbol{\nu}}[M(\nu) = j \mid M(\nu) > j-1] \;\geq \begin{cases} \epsilon_i & \forall j \in [c_i], \\ 1 & \text{for } j = c_i + 1. \end{cases}$$

Using this we show that the expected number of entries until we see an entry with value 1 is at most $\frac{1}{\epsilon_i}$.

$$\begin{aligned}
\mathbb{E}_{\boldsymbol{\nu}}[M(\nu)] &= \sum_{j=0}^{c_i} \mathbb{P}_{\boldsymbol{\nu}}[M(\nu) > j] \\
&= 1 + \sum_{j=1}^{c_i} \mathbb{P}_{\boldsymbol{\nu}}[M(\nu) > j-1]\; \mathbb{P}_{\boldsymbol{\nu}}[M(\nu) > j \mid M(\nu) > j-1] \\
&\leq 1 + (1 - \epsilon_i) \sum_{j=1}^{c_i} \mathbb{P}_{\boldsymbol{\nu}}[M(\nu) > j-1], \\
&\leq 1 + (1 - \epsilon_i)\, \mathbb{E}_{\boldsymbol{\nu}}[M(\nu)] \\
&\leq \frac{1}{\epsilon_i}.
\end{aligned}$$

To show the second part of the lemma, we observe that (31) holds for every arrival t. Thus, for any set S of arrivals, we have,

$$\sum_{t\in S}\sum_{k=1}^{c_i}\Delta g(k)\mathbb{E}_{\boldsymbol{\nu}_{\boldsymbol{k}+}}\left[(1-\mathcal{X}_k(\nu_{k^+},t))\,\mathbb{P}_{\boldsymbol{\nu}_{\boldsymbol{k}}}[k>z_i(t)\mid \nu_{k^+}]\right]\leq \frac{2}{\epsilon_i c_i}|S|.$$

Plugging in $S=O((\omega,i)$ in the inequality above completes the proof. □

LEMMA E8. *For any resource i, unit k, and path ν_{k^+}, we have,*

$$\mathbb{E}_{\boldsymbol{\omega}}\Big[\sum_{t\in O(\omega,i)}\mathcal{X}_k(\nu_{k^+},t)\Big]\leq c_i r\big(F_i,\boldsymbol{s}(\nu_{k^+})\big).$$

Consequently,

$$\mathbb{E}_{\boldsymbol{\omega}}\Big[\sum_{t\in O(\omega,i)}\Big(\sum_{k=1}^{c_i}\Delta g(k)\mathbb{E}_{\boldsymbol{\nu}_{\boldsymbol{k}+}}\left[\mathcal{X}_k(\nu_{k^+},t)\right]\Big)\Big]\leq\Big(1+\frac{2}{c_i}\Big)\sum_{k=1}^{c_i}\Big(g\Big(\frac{k}{c_i}\Big)\,\mathbb{E}_{\boldsymbol{\nu}_{\boldsymbol{k}+}}\left[r\big(F_i,\boldsymbol{s}(\nu_{k^+})\big)\right]\Big).$$

Proof. Fix an arbitrary unit k_O of i in OPT and let $O(\omega,k_O)$ denote the set of arrivals matched to this unit on sample path ω in OPT. It suffices to show that,

$$\mathbb{E}_{\boldsymbol{\omega}}\Big[\sum_{t\in O(\omega,k_O)}\mathcal{X}_k(\nu_{k^+},t)\Big]\leq r\big(F_i,\boldsymbol{s}(\nu_{k^+})\big). \tag{32}$$

Now, recall that the set $\boldsymbol{s}(\nu_{k^+})$ is defined as the set of all covered arrivals on sample path ν_{k^+}, i.e., arrivals t for which $\mathcal{X}_k(\nu_{k^+},t)=1$. Therefore,

$$\mathbb{E}_{\boldsymbol{\omega}}\Big[\sum_{t\in O(\omega,k_O)}\mathcal{X}_k(\nu_{k^+},t)\Big]=\mathbb{E}_{\boldsymbol{\omega}}\Big[\big|O(\omega,k_O)\cap\boldsymbol{s}(\nu_{k^+})\big|\Big].$$

Now the proof of (32) follows by a straightforward application of the two list coupling argument used to prove (15) (Lemma 4). To prove the corollary statement, we change the order of summation in the LHS,

$$\mathbb{E}_{\boldsymbol{\omega}}\Big[\sum_{t|O(t)=i}\Big(\sum_{k=1}^{c_i}\Delta g(k)\mathbb{E}_{\boldsymbol{\nu}_{\boldsymbol{k}+}}\left[\mathcal{X}_k(\nu_{k^+},t)\right]\Big)\Big]=\mathbb{E}_{\boldsymbol{\omega}}\Big[\sum_{k=1}^{c_i}\Delta g(k)\sum_{t|O(t)=i}\mathbb{E}_{\boldsymbol{\nu}_{\boldsymbol{k}+}}\left[\mathcal{X}_k(\nu_{k^+},t)\right]\Big].$$

Using $\Big(1-\frac{1}{c_i}\Big)\Delta g(k)\leq\frac{1}{c_i}g(\frac{k}{c_i})$ completes the proof. □

E.4. Proof of Perturbation Property (Proposition 1) for Various Families

E.4.1. $\{d_i,+\infty\}$ Distributions. We consider the case where the usage duration for resource i takes a finite value d_i w.p. p_i and with remaining probability $1-p_i$, the unit is never returned. This generalizes both non-reusable resources as well as deterministic reusability. We show inequality (27) for this family of usage distributions. Then using Lemma E9 we have that for any set of values $\{p_i\}_{i\in I}$, RBA is $(1-1/e)$–competitive with convergence rate $O(\frac{1}{\sqrt{c_{min}}})$. Focusing on a single resource i with two-point distribution, we show more strongly that regardless of the usage distributions of other resources, condition (3) holds for

i with $\alpha_i = 1-1/e - O(\frac{1}{\sqrt{c_i}})$. We can also sharpen the convergence rate to $O\big(\frac{\log c_i}{c_i}\big)$, as promised in Table 2. This requires a subtler analysis that will take us away from the outline developed previously. Thus, we include a proof with sharper convergence rate of $O\big(\frac{\log c_i}{c_i}\big)$ separately in Appendix E.5.

Now, to prove (27), fix an arbitrary resource i, unit k and path ν_{k^+}. For convenience we treat ϵ_i, κ_i as parameters instead of functions. Their relationship with each other and with parameter c_i will be determined towards the end when we optimize the convergence rate. First, we claim the following.

Lemma E11 *Any two arrivals t_1, t_2 in the set $\boldsymbol{\sigma}(\nu_{k^+})$ are such that $|a(t_1) - a(t_2)| \geq d_i$.*

Proof. If we are given an ordered set $\boldsymbol{\sigma}(\nu_{k^+})$ where this is not true, then consider the earliest pair of contiguous arrivals in $\boldsymbol{\sigma}(\nu_{k^+})$ where this is false and note that the probability of k being matched to the later of the two arrivals is 0. So we can remove this arrival from the set w.l.o.g. Repeating this gives a set with the desired property. □

The next lemma gives a superset on the set $\boldsymbol{s}(\nu_{k^+})$. For a positive value $\epsilon_i \in (0,1]$, let $l_0 \geq 1$ denote the largest integer such that $p_i^{l_0-1} \geq \epsilon_i$. The value of ϵ_i, which is also the probability lower bound used in defining the covering function, will be suitably chosen later in order to optimize the convergence rate. If $\boldsymbol{\sigma}(\nu_{k^+})$ contains less than l_0 elements, let $\sigma_{l_0} = T$.

Lemma E12 *Define, $\boldsymbol{S}(\nu_{k^+}) = \{t \mid \exists \sigma_j$ s.t. $a(t) - a(\sigma_j) \in [0, d_i)$ or $t \geq \sigma_{l_0}\}$. Then $\boldsymbol{s}(\nu_{k^+}) \subseteq \boldsymbol{S}(\nu_{k^+})$.*

Proof. For any arrival $t \notin \boldsymbol{s}(\nu_{k^+})$, clearly $t < \sigma_{l_0}$. Now, if k has usage duration d_i for at least the first $l_0 - 1$ uses, then we claim that k is available at t. This would imply that that the probability of k being available at t is at least $p_i^{l_0-1}$, as desired. To see the claim, note that since k can be matched at most $l_0 - 1$ times prior to any $t < \sigma_{l_0}$, and a finite duration for the first $l_0 - 1$ uses of k implies that k is in use precisely during the intervals $\cup_{j=1}^{l_0-1}(a(\sigma_j), a(\sigma_j)+d_i)$, which do not contain t by assumption. Therefore, k is available at t when its first $l_0 - 1$ durations are finite. □

It remains to show that, $r\big(F_i, \boldsymbol{S}(\nu_{k^+})\big) \leq \big(1+\kappa_i\big) r\big(F_i, \boldsymbol{\sigma}(\nu_{k^+})\big)$.

Proof of (27). The proof follows by considering the simple coupling where we sample the same number l_f of finite usage durations, for both random processes in question. When the value $l_f \leq l_0 - 1$, the number of matches on set $\boldsymbol{S}(\nu_{k^+})$ equals the number of matches on $\boldsymbol{\sigma}(\nu_{k^+})$. For $l_f > l_0 - 1$ we have a higher reward on the LHS, however we also have the following,

$$r\big(F_i, \boldsymbol{S}(\nu_{k^+})\big) \leq r\big(F_i, \boldsymbol{\sigma}(\nu_{k^+})\big) + p_i^{l_0} r\big(F_i, \boldsymbol{S}(\nu_{k^+})\big).$$

Applying Lemma 5 then gives us (27), with $\kappa_i = \frac{\epsilon_i}{1-\epsilon_i}$. □

Convergence rate: The rate $O\big(\frac{1}{\epsilon_i c_i} + \frac{\epsilon_i}{1-\epsilon_i}\big)$ with $\epsilon_i \to 0$ for large c_i, is optimal for $\epsilon_i = \frac{1}{\sqrt{c_i}}$. A sharper rate of $O\big(\frac{\log c_i}{c_i}\big)$ is shown in Appendix E.5.

E.4.2. IFR Distributions. A major advantage in the $\{d_i, +\infty\}$ case was that if a unit returned after usage, the duration of usage was always d_i. In other words, we had the additional structure that a returning unit of i in RBA and in OPT was used for the exact same duration and the main question was whether the unit returned at all. More generally, it is not simply a matter of an item returning after use but also the duration of usage. In particular, the probability that an item is available for (potential) l-th use is not stated as simply as p_i^{l-1}. In this section, we address the new issues that arise for continuous IFR distributions.

For convenience we treat ϵ_i, κ_i as parameters instead of functions. Their relationship with each other and with parameter c_i will be determined towards the end when we optimize the convergence rate. We use f_i to refer to the p.d.f. and F_i to refer to the c.d.f.. Recall that the function $L_i(\epsilon)$ is defined as the maximum possible value of the ratio $\big(F_i(x+F_i^{-1}(\epsilon)) - F_i(x)\big)/\epsilon$. We will show that when the usage distribution of i is IFR and bounded in the following sense,

$$L_i(\epsilon_i)\epsilon_i \to 0 \text{ as } \epsilon_i \to 0,$$

then RBA is asymptotically $(1-1/e)$–competitive. More specifically, inequality (3) is satisfied with,

$$\alpha_i = (1-1/e) - O\left(\frac{1}{\epsilon_i c_i} + L_i(\epsilon_i)\epsilon_i \log\left(\frac{1}{L_i(\epsilon_i)\epsilon_i}\right)\right).$$

For distributions where $L_i(\epsilon_i)\epsilon_i = O(\epsilon_i^{\eta})$ for some $\eta > 0$, the optimal rate is thus, $\tilde{O}\left(c_i^{-\frac{\eta}{1+\eta}}\right)$. In Appendix E.1, we evaluate this parameter η for some commonly known families of IFR distributions. Observe that for IFR distributions that have non-increasing densities, such as exponential, uniform etc., $L_i(\epsilon_i) = 1$ and we have a resulting convergence rate of $O(\log c_i/\sqrt{c_i})$. In fact, for exponential distributions we will show a stronger convergence rate of $O(1/\sqrt{c_i})$.

To prove these claims, let us fix an arbitrary resource i, unit k, path ν_{k+}. This also fixes the set $\boldsymbol{\sigma}(\nu_{k+})$ of arrivals that k could be matched to in RBA. Since we have fixed i, for simplicity we use the abbreviations,

$$\epsilon := \epsilon_i, \quad L := L_i(\epsilon_i), \quad F := F_i$$

We assume that the density f_i is continuous and consequently, let δ_0 and δ_L be such that,

$$F(\delta_0) = \epsilon \quad \text{and} \quad F(\delta_L) = L\epsilon,$$

where note that, $L\epsilon \le 1$ by definition of L. Now, we claim that for a $\mathcal{X}_k(\nu_{k+}, \epsilon, t)$ covering, the following set of arrivals is a superset of $\boldsymbol{s}(\nu_{k+})$,

$$\boldsymbol{S}(\nu_{k+}) = \{t \mid a(t) \in [a(\sigma_j), a(\sigma_j) + \delta_0) \text{ for some } \sigma_j \in \boldsymbol{\sigma}(\nu_{k+})\}.$$

Lemma E13 $\boldsymbol{s}(\nu_{k+}) \subseteq \boldsymbol{S}(\nu_{k+})$.

Proof. Consider an arrival $t \notin \boldsymbol{S}(\nu_{k^+})$. The closest arrival preceding t in $\boldsymbol{\sigma}(\nu_{k^+})$, is at least δ_0 time before $a(t)$. Using the IFR property, we have that the probability that k switches from being in-use to free between $a(\sigma_j)$ and $a(\sigma_j)+\delta_0$ is at least $F(\delta_0)=\epsilon$. □

Given this, in order to show (27), we now aim to to show the following lemma and then apply the monotonicity Lemma 5.

Lemma E14 *Given an IFR distribution F and two sets or arrivals $\boldsymbol{S}(\nu_{k^+})$ and $\boldsymbol{\sigma}(\nu_{k^+})$, with the property that for any arrival t in $\boldsymbol{S}(\nu_{k^+})$, there exists an arrival in $\boldsymbol{\sigma}(\nu_{k^+})$ that precedes t by at most δ_0 time. We have,*

$$r\big(F,\boldsymbol{S}(\nu_{k^+})\big) \leq \left(1+O\left(F(\delta_L)\log\left(\frac{1}{F(\delta_L)}\right)\right)\right) r\big(F,\boldsymbol{\sigma}(\nu_{k^+})\big).$$

Proof. We will show this in two main steps. First, we show for a certain modified distribution F^m (F_i^m to be precise),

$$r\big(F,\boldsymbol{S}(\nu_{k^+})\big) \leq (1+2F(\delta_L))r\big(F^m,\boldsymbol{\sigma}(\nu_{k^+})\big). \tag{33}$$

In the second step, we will relate $r(F^m,A)$ to $r(F,A)$ for any set A. The modified distribution F^m needs to satisfy a property that we describe next. The exact definition of the distribution will be introduced later in the proof.

Proposition 2 *Given a c.d.f. F, we require a modified distribution F^m such that, there exists a coupling that generates samples distributed according to F^m by using i.i.d. samples from F, with the property that for every sample of F that has value $d \geq \delta_L$, the coupled sample of F^m has value at most $d-\delta_0$.*

Suppose a distribution F^m that satisfies the above property exists and let $\pi:[\delta_L,+\infty)\to\mathbb{R}^+$ be the mapping induced by the coupling, from samples of F to those of F^m. Then, to show (33), we use the following coupling between $(F,\mathcal{S}(\nu_{k^+}))$ and $(F^m,\boldsymbol{\sigma}(\nu_{k^+}))$ random processes. For the rest of the proof, let us refer to these as processes as R and R^m resp..

Consider a long list of i.i.d. samples from distribution F. Starting with a pointer P at the first sample in the list, every time we have a transition to in-use in process R, we draw the sample P is pointing to and move P to the next sample on the list. Thus, pointer P draws from the list in order without skipping over samples. To couple the two processes, we introduce another pointer on the list, P^m, for process R^m. Let us call any sample with value at least δ_L, a *large* sample. Other samples are called *small* samples. Pointer P^m starts at the first large sample in the list. Each time R^m needs a new sample, we draw the sample pointed to by P^m and pass it through the function $\pi(\cdot)$. The output is then passed to process R^m and P^m moves down the list to the next large sample. If P and P^m point to the sample sample and P moves down the list, P^m also moves down the list to the next large sample. At every arrival, we first provide a sample, if required, to

R^m and update P^m before moving to R. Thus, pointer P^m never lags behind P. Compared to the coupling used for Lemma E8, there are two important differences. First, here we do rejection sampling by ignoring small values, as P^m always skips to the next large sample. Second, we pass the value pointed to by P^m through the function $\pi(\cdot)$ before giving the sample to R^m.

We claim that P^m never skips a large sample in the list. We shall prove this by contradiction. Consider the first large sample in the list which is skipped by P^m, let its position in the list be q. Clearly, q is a sample that is used by R otherwise P^m would never skip it. Further, q cannot be the first large sample in the list, as the earliest arrival in $\boldsymbol{\sigma}(\nu_{k+})$ is also the first arrival in $\boldsymbol{S}(\nu_{k+})$. Therefore, R^m uses the first large sample at least as early as R does. More generally, given that P^m never lags P, we have that all large samples preceding q were first used in R^m. Consider the arrival $\sigma_j \in \boldsymbol{\sigma}(\nu_{k+})$, where R^m uses the large sample immediately preceding q in the list. Due to the fact that function $\pi(\cdot)$ reduces the sample value passed to R^m by at least δ_0, we claim that R^m needs the next large sample, i.e., q, before R does. To see the claim, let s_j be the arrival in $\boldsymbol{S}(\nu_{k+})$ where R uses the large sample immediately preceding q. Observe that σ_j occurs at least as early as s_j due to P^m leading P and by the assumption on q being the first sample skipped by P^m. Now, due to the decrease in sample value resulting from applying π, we have that in R^m, k returns to free state after becoming matched to σ_j, at least δ_0 time before k returns to free state after getting matched to s_j in R. Letting t_σ and t_s denote these return times resp., we have, $t_\sigma \leq t_s - \delta_0$. Suppose that the first arrival in $\boldsymbol{S}(\nu_{k+})$ after time t_s occurs at time τ_s and the first arrival in $\boldsymbol{\sigma}(\nu_{k+})$ after time t_σ occurs at time τ_σ. To establish the claim that R^m uses sample pointed by q before R does, it suffices to argue that $\tau_\sigma \leq \tau_s$. From the definition of set $\boldsymbol{S}(\nu_{k+})$, we have that either $\tau_s \in \boldsymbol{\sigma}(\nu_{k+})$ (in which case, we are done) or τ_s is preceded by an arrival in $\boldsymbol{\sigma}(\nu_{k+})$ that occurs in the interval $[\tau_s - \delta_0, \tau_s)$. Therefore, $t_\sigma \leq t_s - \delta_0 \leq \tau_s - \delta_0$, and we have that $\tau_\sigma \leq \tau_s$. So R^m requires sample q before R, which contradicts that q is the first large sample skipped by P^m. Hence, P^m does not skip any large samples. Since P^m never lags P, we also have that the total number of large samples passed to R is upper bounded by the number of samples passed to R^m. Let $r^{lg}(F, \boldsymbol{S}(\nu_{k+}))$ denote the expected number of large sample transitions from available to in-use in process R. So far we have shown that,

$$r^{lg}(F, \boldsymbol{S}(\nu_{k+})) \leq r(F^m, \boldsymbol{\sigma}(\nu_{k+})).$$

Next, we claim that,

$$r(F, \boldsymbol{S}(\nu_{k+})) \leq (1 + 2F(\delta_L)) r^{lg}(F, \boldsymbol{S}(\nu_{k+})).$$

Consider an arbitrary transition from available to in use in R. The probability that duration of this transition is large is $1 - F(\delta_L)$. Thus, the expected contribution from this transition to $r^{lg}(F, \boldsymbol{S}(\nu_{k+}))$ is $1 - F(\delta_L)$. Summing over all transitions, we have the desired. This completes the proof of (33). To finish the main proof it remains to define a modified distribution F^m such that Proposition 2 holds and compare the expected rewards of random processes $(F^m, \boldsymbol{\sigma}(\nu_{k+}))$ with $(F, \boldsymbol{\sigma}(\nu_{k+}))$. This is the focus of the next lemma. □

Lemma E15 *Given IFR distribution F and value $\epsilon \in (0,1]$ such that $\epsilon L(\epsilon) \leq 1/2$, there exists a modified distribution F^m that satisfies Proposition 2, such that for any given set of arrivals A,*

$$r(F^m, A) \leq \left(1 + O\left(F(\delta_L)\log\left(\frac{1}{F(\delta_L)}\right)\right)\right) r(F, A). \tag{34}$$

Proof. Let us start with the case of exponential distribution as a warm up. We claim that in fact, choosing $F^m = F$ suffices in this case. Inequality (34) follows directly for this choice and it remains show that Proposition 2 is satisfied, i.e., prove the existence of a coupling/mapping π. Consider the following modified density function,

$$f^m(x) = \frac{f(x+\delta_0)}{1 - F(\delta_0)}.$$

$f^m = f$ owing to the memoryless property of exponentials. Now, consider the straightforward coupling that samples from F and rejects all samples until the first large sample is obtained, which is reduced by δ_0 before it is output. Observe that this process generates samples with distribution F^m, and satisfies Proposition 2.

For other IFR distributions, it is not clear if $r(F^m, \boldsymbol{\sigma}(\nu_{k+}))$ and $r(F, \boldsymbol{\sigma}(\nu_{k+}))$ are comparable given the current definition of F^m. So we introduce a new modified distribution. Define δ_1 so that,

$$F(\delta_1) = 1 - L\epsilon.$$

Since we started with $\epsilon L \leq 1/2$ for this lemma, we have that $\delta_L \leq \delta_1$. Now, the new distribution is,

$$f^m(x) = \begin{cases} \frac{1}{1-L\epsilon} f(x) & x \in [0, \delta_1] \\ 0 & x > \delta_1 \end{cases}$$

So f^m is a truncated version of f. First, we show that Proposition 2 holds by defining a mapping π such that, for every $t \geq \delta_L$ we have $F(t) - F(\pi(t)) = L\epsilon$. By definition of L and ϵ, this would imply $t - \pi(t) \geq \delta_0$.

Consider values $x \in [0, L\epsilon)$ and non-negative $j \in \mathbb{Z}$. Let $t(x,j) = F^{-1}(x + (j+1)L\epsilon)$. Varying x and j we have that $t(x,j)$ takes all possible values in the range $[\delta_L, +\infty)$. Defining, $\pi(t(x)) = F^{-1}(x + jL\epsilon)$, we have that $F(t) - F(\pi(t)) = L\epsilon$ for every $t \geq \delta_L$, as desired. Finally, the distribution generated by sampling i.i.d. values from F and applying mapping π to all large samples (while ignoring all small samples), corresponds to the distribution F^m as defined. To show (34), consider the random variable defined as the minimum number of values drawn i.i.d. with density f^m, such that the sum of the values is at least as large as a single random value that is drawn independently with density f. Let $\hat{n}$ denote the expected value of this random variable. Observe that in order to show the main claim it suffices to show that $\hat{n} \leq (1+\gamma)$, where $\gamma = O(L\epsilon \log(1/L\epsilon))$. To show this, let μ, μ^m denote the mean of f and f^m respectively. Recall that f^m is a truncated version of f. By definition of f^m and the IFR property of f, we have that,

$$\mu \leq (1 - L\epsilon) \cdot \mu^m + L\epsilon(\delta_1 + \mu). \tag{35}$$

For the time being, we claim μ is lower bounded as follows,

$$\mu \geq O(1) \cdot \frac{\delta_1}{\log(1/L\epsilon)}. \tag{36}$$

We proceed with the proof assuming (36) holds and prove this claim later. Substituting this inequality in (35) and using the fact that $x \log(1/x) < 1$ for $x \in (0, 1]$, we get,

$$\mu^m \geq O(1) \cdot \frac{\delta_1}{\log(1/L\epsilon)}. \tag{37}$$

Now, to upper bound $\hat{n}$ consider the following coupling: Given a random sample t from f, for $t \leq \delta_1$ we consider an equivalent sample drawn from f^m. For durations $t > \delta_1$, we draw independent samples from f^m until the sum of these samples is at least t. Thus, w.p. $(1 - L\epsilon)$, we draw exactly one sample from f^m and it suffices to show that in the remaining case that occurs w.p. $L\epsilon$, we draw in expectation $O(\log(1/L\epsilon))$ samples from f^m. Now, given sample $t \geq \delta_1$, note that from (37), the expected number of samples drawn from f^m before their sum exceeds δ_1 is, $O(\log(1/L\epsilon))$. Second, using the IFR property of f, we have that the expected number of samples of f^m such that the sum of the sample values is at least $t - \delta_1$, is upper bounded by $\hat{n}$. Overall, we have the recursive inequality,

$$\hat{n} \leq (1 - L\epsilon) + L\epsilon\big(O(\log(1/L\epsilon)) + \hat{n}\big).$$

The desired upper bound on $\hat{n}$ now follows.

It only remains to show (36). Let x_1 be such that $1 - F(x_1) = 1/2$. Then, let x_2 be such that $1 - F(x_1 + x_2) = 1/4$. More generally, let $\{x_1, \cdots, x_s\}$ be the set of values such that $1 - F(\sum_{j=1}^{s} x_j) = 1/2^s$. We let s be the smallest integer greater that equal to $\log(1/L\epsilon)$. Now, the mean $\mu \geq \sum_{j=1}^{s} x_j/2^j$. Therefore, μ/δ_1 is lower bounded by $\frac{\sum_{j=1}^{s} x_j/2^j}{\sum_{j=1}^{s} x_j}$. While in general this ratio can be quite low, due to the IFR property we have $x_j \geq x_{j+1}$ for every $j \geq 1$. Consequently, the minimum value of this ratio is $\frac{O(1)}{s} = \frac{O(1)}{\log(1/L\epsilon)}$, and occurs when all values x_j are equal (which incidentally, implies memoryless property). □

Convergence rate: We have $\kappa = O\left(F(\delta_L)\log\left(\frac{1}{F(\delta_L)}\right)\right) = O\big(L\epsilon\log(1/L\epsilon)\big)$. Thus, over all rate of convergence for resource i is, $O\left(\frac{1}{\epsilon_i c_i} + L_i(\epsilon_i)\epsilon_i \log\left(\frac{1}{L_i(\epsilon_i)\epsilon_i}\right)\right)$. Suppose $L_i(\epsilon_i)\epsilon_i = O(\epsilon_i^{\eta})$ for some $\eta > 0$, then the optimal rate is, $\tilde{O}\left(c_i^{-\frac{\eta}{1+\eta}}\right)$, where the $\tilde{O}$ hides a $\log c_i$ factor.

E.4.3. IFR with Arbitrary Mass at $+\infty$. In the previous section we, showed that for bounded IFR distributions RBA is asymptotically $(1 - 1/e)$–competitive. In certain practical scenarios, it may be reasonable to expect that with some probability resource units under use may never return back to the system. We model this by allowing an arbitrary mass at infinity. Specifically, for resource i, let p_i denote the probability that a usage duration takes value drawn from distribution with c.d.f. F_i and with the remaining probability $1 - p_i$, the duration takes value $+\infty$. In this section, we show that RBA is still asymptotically $(1 - 1/e)$–competitive for such a mixture of non-increasing IFR usage distributions with arbitrary mass at infinity.

Fix, i, k, ν_{k+} and define L,ϵ, δ_0 and δ_L as before. Given set $\boldsymbol{\sigma}(\nu_{k+})$, consider the random process $(F, \boldsymbol{\sigma}(\nu_{k+}))$ and let $\mathbb{1}(t, \texttt{finite})$ indicate the event that k has not hit a $+\infty$ duration at arrival t. For a chosen value $\gamma \in (0,1]$ (finalized later to optimize convergence rate), let σ_{l_f} be the last arrival in $\boldsymbol{\sigma}(\nu_{k+})$ where,

$$\mathbb{E}[\mathbb{1}(\sigma_{l_f}, \texttt{finite})] \geq \gamma.$$

If no such arrival exists, let $\sigma_{l_f} = T$. This implies,

$$\mathbb{E}[\mathbb{1}(\sigma_j, \texttt{finite})] \geq \gamma \quad \text{for every } \sigma_j < \sigma_{l_f} \text{ in } \boldsymbol{\sigma}(\nu_{k+}), \tag{38}$$

$$\text{and} \quad \mathbb{E}[\mathbb{1}(\sigma_j, \texttt{finite})] < \gamma \quad \text{for every } t > \sigma_{l_f}. \tag{39}$$

Using this define the set of arrivals,

$$\boldsymbol{S}(\nu_{k+}) = \{t \mid t > \sigma_{l_f} \text{ or } \exists \sigma_j \in \boldsymbol{\sigma}(\nu_{k+}) \text{ s.t. } a(t) \in [a(\sigma_j), a(\sigma_j) + \delta_0)\}.$$

Lemma E16 *For a $\mathcal{X}_k(\nu_{k+}, \gamma\epsilon, t)$ covering, we have that the set of covered arrivals satisfies the following relation, $\boldsymbol{s}(\nu_{k+}) \subseteq \boldsymbol{S}(\nu_{k+})$.*

Proof. Consider an arrival $t \notin \boldsymbol{S}(\nu_{k+})$. Clearly, $t < \sigma_{l_f}$ and the closest arrival preceding t in $\boldsymbol{\sigma}(\nu_{k+})$, call it σ_t, is at least δ_0 time before $a(t)$. From (38), we have that k has not hit a duration of ∞ by the time σ_t arrives, w.p. at least γ. Conditioned on this, we have from the IFR property that k switches from being in-use to free between $a(\sigma_j)$ and $a(\sigma_j) + \delta_0$ w.p. at least $F(\delta_0) = \epsilon$. Together, this implies that the probability that k is free at t is at least $\gamma\epsilon$. □

Lemma E17 $r\big(F, \boldsymbol{S}(\nu_{k+})\big) \leq \Big(1 + O\big(L\epsilon \log(1/L\epsilon)\big)\Big)(1+2\gamma)\, r\big(F, \boldsymbol{\sigma}(\nu_{k+})\big)$.

Proof. The proof requires a combination of the two analyses so far. Consider the truncated set $\hat{\boldsymbol{S}}(\nu_{k+}) = \boldsymbol{S}(\nu_{k+}) \backslash \{t \mid t > \sigma_{l_f}\}$. First, we have that,

$$\begin{aligned} r\big(F, \boldsymbol{S}(\nu_{k+})\big) &< r\big(F, \hat{\boldsymbol{S}}(\nu_{k+})\big) + \gamma r\big(F, \boldsymbol{S}(\nu_{k+})\big), \\ r\big(F, \boldsymbol{S}(\nu_{k+})\big) &\leq (1+2\gamma) r\big(F, \hat{\boldsymbol{S}}(\nu_{k+})\big), \end{aligned}$$

where we used the fact that $\gamma \leq 1$. It suffices to therefore show that,

$$r\big(F, \hat{\boldsymbol{S}}(\nu_{k+})\big) \leq \Big(1 + O\big(L\epsilon \log(1/L\epsilon)\big)\Big) r\big(F, \boldsymbol{\sigma}(\nu_{k+})\big).$$

Let us condition on the first l transition in each process being finite and the $l+1$-th transition being ∞. Then, the resulting expected number of transitions can be compared in the same way as the case of IFR distributions given in Lemma E14, since we now draw independently from an IFR distribution for the first l durations (having conditioned on these durations being finite). This holds for arbitrary l. Taking expectation over l then gives the desired. □

Convergence rate: We have $\kappa = O\big(\gamma + L\epsilon\log(1/L\epsilon)\big)$. Moreover, the probability lower bound at uncovered arrivals is actually $\gamma\epsilon$, instead of ϵ. Thus, over all convergence rate for resource i is, $O\Bigg(\frac{1}{\epsilon_i\gamma_i c_i} + \gamma_i + L_i(\epsilon_i)\epsilon_i\log\left(\frac{1}{L_i(\epsilon_i)\epsilon_i}\right)\Bigg)$. Suppose $L_i(\epsilon_i)\epsilon_i = O(\epsilon_i^{\eta})$, then the optimal rate is, $\tilde{O}\Big(c_i^{-\frac{\eta}{1+2\eta}}\Big)$.

E.5. Refined Bound for $\{d_i, +\infty\}$

The main bottleneck in getting a stronger convergence factor is the need to ensure large enough probability lower bound ϵ_i, for uncovered arrivals. In particular, we chose the parameter l_0 to ensure that $p_i^{l_0-1} \geq 1/\sqrt{c_i}$, and consequently bound the error term $\frac{1}{p_i^{l_0-1}c_i}$ arising out of Lemma E10. Choosing a larger l_0 would worsen the convergence rate since term designates the contribution from uncovered arrivals in OPT to (24). The separation of contributions from uncovered and covered arrivals makes the analysis tractable in general. It can also be, at least for $\{d_i, +\infty\}$ distributions, pessimistic from the point of view of convergence to the guarantee. As an extreme but illustrative example, consider a sample path in OPT where a unit k_O of i is matched only to arrivals t that occur late and are all uncovered given some unit k and path ν_{k^+} in RBA. Specifically, at each of the arrivals k_O is matched to, let the probability of k in RBA being available (conditioned on ν_{k^+}), equal $p_i^{l_0-1}$. In this case, the contributions to (24) from uncovered arrivals is significant and that part of the analysis is tight, but there is no contribution from covered arrivals in OPT. Thus, the θ_i term could be used to neutralize some of the negative terms arising out of the uncovered arrivals and this would in turn allow us to set a larger value of l_0. This observation is the key idea behind improving the analysis.

Lemma E18 *For every i, (3) is satisfied with $\alpha_i = (1-1/e) - O\Big(\frac{\log c_i}{c_i}\Big)$.*

Proof. We start by fixing also an arbitrary unit k_O in OPT and show the following inequality,

$$\frac{(1+2/c_i)}{r_i c_i}\theta_i + \mathbb{E}_{\boldsymbol{\omega}}\Big[\sum_{t|O(t)=(i,k_O)}\Big((1-1/e) - \sum_{k=1}^{c_i}\Delta g(k)\mathbb{P}_{\boldsymbol{\nu}}\big[k > z_i(t)\big]\Big)\Big]$$
$$\geq \Big[1 - 1/e - O\Big(\frac{\log c_i}{c_i}\Big)\Big]\mathbb{E}_{\boldsymbol{\omega}}\Big[\big|\{t \mid O(t) = (i,k_O)\}\big|\Big]. \tag{40}$$

The proof then follows by linearity of expectation.

To show (40), we start by conditioning on all randomness in OPT arising out of usage durations of resources and units other than unit k_O of i, as well as any intrinsic randomness in OPT. Let this partial sample path be denoted as ω_{-k_O}. This fixes the set of arrivals that k_O is matched to if durations of k_O are all finite. Let this set be $\Gamma(\omega_{-k_O}) = \{t_1(\omega_{-k_O}), t_2(\omega_{-k_O}), \cdots, t_{l_f(\omega_{-k_O})}(\omega_{-k_O})\}$. Since usage durations of k_O are sampled independently, we have that,

$$\mathbb{E}_{\boldsymbol{\omega}}\Big[\big|\{t \mid O(t) = (i,k_O)\}\big|\Big] = \mathbb{E}_{\boldsymbol{\omega_{-k_O}}}\Big[\sum_{l\geq 1} p_i^{l-1}\mathbb{1}\big(l \leq l_f(\omega_{-k_O})\big)\Big]$$

Now, define l_0 as the largest integer such that $p_i^{l_0-1} \geq \frac{\log c_i}{c_i}$. On every path ω_{-k_O}, it suffices to focus on at most $l_0 - 1$ finite durations for k_O as,

$$\sum_{l\geq 1}^{l_0} p_i^{l-1} \mathbb{1}\big(l \leq l_f(\omega_{-k_O})\big) \geq \Big(1 - \frac{\log c_i}{c_i}\Big) \sum_{l\geq 1} p_i^{l-1} \mathbb{1}\big(l \leq l_f(\omega_{-k_O})\big).$$

Thus, in order to prove (40) it suffices to show that,

$$\frac{(1+2/c_i)}{r_i c_i}\theta_i + \mathbb{E}_{\boldsymbol{\omega_{-k_O}}}\Big[\sum_{l\geq 1}^{l_0} p_i^{l-1} \mathbb{1}\big(l \leq l_f(\omega_{-k_O})\big)\Big((1-1/e) - \sum_{k=1}^{c_i} \Delta g(k) \mathbb{P}_{\boldsymbol{\nu}}\big[k > z_i(t_l(\omega_{-k_O}))\big]\Big)\Big]$$
$$\geq \Big[1 - 1/e - O\Big(\frac{\log c_i}{c_i}\Big)\Big] \mathbb{E}_{\boldsymbol{\omega_{-k_O}}}\Big[\sum_{l\geq 1}^{l_0} p_i^{l-1} \mathbb{1}\big(l \leq l_f(\omega_{-k_O})\big)\Big].$$

In fact, it suffices to show more strongly that for every ordered collection $\Gamma = \{t_1, \cdots, t_{l_f}\}$ of arrivals such that any two consecutive arrivals in the set are at least d_i time apart and $l_f \leq l_0$, we have,

$$\sum_{l\geq 1}^{l_f} p_i^{l-1} \Big(\sum_{k=1}^{c_i} \Delta g(k) \mathbb{P}_{\boldsymbol{\nu}}\big[k > z_i(t_l)\big]\Big) - \frac{(1+2/c_i)}{r_i c_i}\theta_i$$
$$= O\Big(\frac{\log c_i}{c_i}\Big) \sum_{l\geq 1}^{l_f} p_i^{l-1} = O\Big(\frac{\log c_i}{c_i}\Big) \frac{1 - p_i^{l_f}}{1 - p_i}. \tag{41}$$

Now, fix unit k, sample path ν_{k^+}, and consequently, the set $\boldsymbol{\sigma}(\nu_{k^+})$. We classify arrivals $t_l \in \Gamma$ into covered and uncovered in a new way, given by a function $\mathcal{Y}_k$. Using Lemma E11, we let any two arrivals in $\boldsymbol{\sigma}(\nu_{k^+})$ be at least d_i time apart. Then,

t_l is covered and $\mathcal{Y}_k(\nu_{k^+}, t_l) = 1$ iff there are l or more arrivals preceding t_l in $\boldsymbol{\sigma}(\nu_{k^+})$.

Using this definition and the decomposition we performed in Lemma E7, we have,

$$\sum_{l\geq 1}^{l_f} p_i^{l-1} \Big(\sum_{k=1}^{c_i} \Delta g(k) \mathbb{P}_{\boldsymbol{\nu}}\big[k > z_i(t_l)\big]\Big) \leq$$
$$\sum_{l\geq 1}^{l_f} p_i^{l-1} \Big(\sum_{k=1}^{c_i} \Delta g(k) \mathbb{E}_{\boldsymbol{\nu_{k+}}}\big[\mathcal{Y}_k(\nu_{k^+}, t_l)\big]\Big) + \tag{42}$$
$$\sum_{l\geq 1}^{l_f} p_i^{l-1} \Big(\sum_{k=1}^{c_i} \Delta g(k) \mathbb{E}_{\boldsymbol{\nu_{k+}}}\big[(1 - \mathcal{Y}_k(\nu_{k^+}, t_l)) \mathbb{P}_{\boldsymbol{\nu_k}}[k > z_i(t_l) \mid \nu_{k^+}]\big]\Big).$$

Now, for any unit k in RBA, in (42) we may interpret probabilities p_i^{l-1} as the probability that at least l durations of k are finite. Then, by definition of the coupling we have that for every path ν_{k^+},

$$\sum_{l\geq 1}^{l_f} p_i^{l-1} \mathcal{Y}_k(\nu_{k^+}, t_l) \leq r(F_i, \boldsymbol{\sigma}(\nu_{k^+})).$$

Using the same algebra as the corollary statements in Lemma E8 and Proposition 2 then gives us that,

$$\sum_{l\geq 1}^{l_f} p_i^{l-1} \Big(\sum_{k=1}^{c_i} \Delta g(k) \mathbb{E}_{\boldsymbol{\nu_{k+}}}\big[\mathcal{Y}_k(\nu_{k^+}, t_l)\big]\Big) \leq \frac{(1+2/c_i)}{r_i c_i}\theta_i.$$

So in order to prove (41), it remains to show that,

$$\sum_{l\geq 1}^{l_f} p_i^{l-1}\Big(\sum_{k=1}^{c_i}\Delta g(k)\mathbb{E}_{\boldsymbol{\nu_{k+}}}\left[(1-\mathcal{Y}_k(\nu_{k+},t_l))\mathbb{P}_{\boldsymbol{\nu_k}}[k>z_i(t_l)\mid \nu_{k+}]\right]\Big) = O\Big(\frac{\log c_i}{c_i}\Big)\frac{1-p_i^{l_f}}{1-p_i}.$$

To establish this, again fix a unit k and path ν_{k+} and note that if some arrival t_l is uncovered, we have at most $l-1$ arrivals in $\boldsymbol{\sigma}(\nu_{k+})$ preceding t_l. Therefore, at any uncovered arrival t_l,

$$\mathbb{P}_{\boldsymbol{\nu_k}}[\mathbb{1}(k,t_l)=1\mid \nu_{k+}]\geq p_i^{l-1}.$$

Then, applying Lemma E10 with $\epsilon_i = p_i^{l-1}$, we have that for any $p_i<1$,

$$\sum_{k=1}^{c_i}\Delta g(k)\mathbb{E}_{\boldsymbol{\nu_{k+}}}\left[(1-\mathcal{Y}_k(\nu_{k+},t_l))\cdot\mathbb{P}_{\boldsymbol{\nu_k}}[k>z_i(t_l)\mid \nu_{k+}]\right]\leq \frac{2}{p_i^{l-1}c_i},$$

and in case $p_i=1$, the RHS equals 0 (which gives us the desired). For $p_i<1$, observe that,

$$\sum_{l\geq 1}^{l_f} p_i^{l-1}\cdot\frac{2}{p_i^{l-1}c_i}=\frac{2l_f}{c_i}\quad\text{ and }\quad l_f\leq O(\log c_i)\frac{1-p_i^{l_f}}{1-p_i}\quad \forall l_f\leq l_0,$$

which completes the proof. □

E.6. Performance of RBA Beyond Online Matching

Recall that to analyze G-ALG, we captured the allocation of each individual unit via a natural $(F,\boldsymbol{\sigma},\boldsymbol{p})$ random process. In an $(F,\boldsymbol{\sigma})$ random process, at every $\sigma_t\in\boldsymbol{\sigma}$ the only event of relevance is the availability of the item at σ_t (as probability $p_t=1$). If the item is available, then it is matched to σ_t regardless of the usage durations realized before σ_t. In order to naturally capture the actions of RBA on individual units through a random process, a similar property must hold in RBA. To be more specific, consider a single unit k of i in RBA and fix the randomness associated with all units and resources except k. Suppose we have an arrival t with edge to i and two distinct sample paths over usage durations of k prior to the arrival of t. We are also given that k is available at t on both sample paths. Clearly, RBA makes a deterministic decision on each sample path. Is it possible that RBA matches t to k on one sample path but not the other?

Interestingly, in the case of matching we can show that this is impossible, i.e., RBA always makes the same decision on both sample paths in such instances. We show that this is not the case in general. In particular, when we include the aspect of customer choice or budgeted allocations (even one of these), reusability can interact with these elements in an undesirable way.

For the setting of online budgeted allocations (i.e., online matching with multi-unit demand), we generalize RBA as follows: given arrival t with demand $b_{it}\geq 0$ for resource $i\in I$ and set S_t denoting the set of available resources with an edge to t, we match t to,

$$\arg\max_{i\in S_t}\sum_{k=1}^{b_{it}} r_i\left(1-g\left(\frac{z_i(t,k)}{c_i}\right)\right),$$

where $z_i(t,k)$ denotes the $k-$th highest available unit of i when t arrives. In the general model of online assortments with multi-unit demand, each arrival has an associated choice model ϕ_t and we offer the following assortment to t,

$$\arg\max_{S\subseteq S_t} \sum_{i\in S} \left[r_i \phi_t(S,i) \sum_{k=1}^{b_{it}} \left(1 - g\left(\frac{z_i(t,k)}{c_i}\right)\right)\right].$$

The key ingredients of our analysis for matching namely, conditioning and covering, are not helpful in case of budgeted allocations or assortments. To understand this in more detail, consider first the fundamental ordering property for online matching and Corollary E6. In case of budgeted allocations and assortments, for arbitrary arrival t with an edge to i, whether unit k of i is matched to t can depend on the past usage duration of unit k itself. The root cause is that customers can now have different preference ordering among resources. A resource j may be less preferred than resource i by customer 1 but more preferred by customer 2. This substantially complicates the stochastic dependencies that arise out of reusability. Consider the following concrete example for budgeted allocations.

Example E.1 We are given two resources $\{1,2\}$, each with a reward of 1 per unit and capacity of 2 units (example can be generalized to a setting with arbitrary capacity). Let resource 2 be non-reusable and let the usage durations of resource 1 come from a two point distribution with support $\{0.5, 1.5\}$ and probability 0.5 of either possibility. Consider a sequence with three arrivals t_1, t_2, t_3 occurring at time $1, 2$, and 3 respectively. The first two arrivals have a bid of 2 for resource 1 and bid 1 for resource 2 i.e., both arrivals prefer resource 1 if both its units are available. Arrival t_3 has the opposite preference and requires 1 unit of resource 1 or 2 units of resource 2.

Consider the actions of RBA on this instance. Arrival t_1 is allocated 2 units of resource 1. If these units return before the second arrival then resource 1 is also matched to t_2. Overall, with non-zero probability resource 1 is matched to the first two arrivals and is still available when t_3 arrives. However, in this scenario resource 2 is matched to t_3. Now, consider a different sample path where the units of resource 1 allocated to arrival t_1 do not return by t_2. In this case, RBA allocates a unit of the non-reusable resource 2 to arrival t_2. Subsequently, arrival t_3 is allocated a unit of resource 1. Therefore, we have two sample paths where resource 1 is available to allocate to arrival t_3, however, the actions of RBA are different on these sample paths. This behavior is quite unlike the case of matching (for instance, recall Corollary E6).

Next, we construct a very similar instance for the setting of online assortments.

Example E.2 Consider three arrivals t_1, t_2, t_3 that come in that order and two resources $\{1,2\}$ with unit reward and unit capacity (example can be generalized to arbitrary capacity). Let resource 2 be non-reusable. Each arrival has a multinomial logit (MNL) choice model, i.e., probability $\phi(S,i) = \frac{v_i}{v_0 + \sum_{j\in S} v_j}$ for $i \in$

$\{1,2\}$ and any set S containing i. MNL parameters for arrivals t_1 and t_2 are as follows: $v_1 = 100$, $v_2 = 1$ and $v_0 = 0.01$. Arrival t_3 has $v_1^3 = 1$, $v_2^3 = 100$ and $v_0^3 = 0.01$. Now, consider the actions of RBA on this instance. Observe that RBA offers set $\{1,2\}$ to arrival t_1 and with probability close to 1, resource 1 is chosen by this arrival. Suppose that the probability of resource 1 returning before arrival t_2 is $p \in (0,1)$ and resource 1 returns before arrival t_3 w.p. 1. Then with probability p we offer arrival t_2 the set $\{1,2\}$ and resource 1 is chosen again w.h.p.. Subsequently, arrival t_3 will choose resource 2 w.h.p., even if resource 1 returns and is available. In other words, resource 2 is the most preferred available resource for arrival t_3 in this case.

On the other hand, consider the scenario where resource 1 does not return before arrival t_2. Arrival t_2 takes resource 2 w.h.p.. Given that resource 2 is non-reusable, arrival t_3 accepts resource 1 w.h.p.. Therefore, whether arrival t_3 accepts resource 1 depends not just on whether resource 1 is available at arrival t_3, but also on the past usage duration of resource 1 itself. If resource 1 returns before arrival t_2 then arrival t_3 does not accept resource 1, otherwise arrival t_3 accepts resource 1 w.h.p.. Note that on both sample paths, resource 1 is available at arrival t_3. This violates key properties that enable the analysis of RBA for matching (see Corollary E6).

Overall, the stochasticity in reusability interacts in a non-trivial way with the arrival dependent aspect of bids in budgeted allocations, and random choice rankings in case of assortments. Without reusability, and more specifically, without stochasticity in reusability, this interaction disappears. Indeed, as shown by Feng et al. (2019), the results from online assortments with non-reusable resources generalize naturally to the special case of reusable resources with deterministic usage durations. In summary, the ingredients that enable us to derive a general framework of analysis for RBA in case of online matching, do not apply more generally. Substituting for these ingredients in more general models appears to be challenging and is left open as an interesting technical question.

Appendix F: Faster Implementations of G-ALG

Executing G-ALG has two main bottlenecks. First, we need to update values $Y(k_i)$ by taking into account past partial matches of unit k_i. Second, we may fractionally match every arrival t to many units and in the worst case, to all $\sum_{i\in I} c_i$ units. The first issue is somewhat less limiting, as one can in practice continue to update the states during the time between any two arrivals and update states of different units in parallel. The second issue is more important and has a direct impact on the time taken to decide the match for every arrival.

F.1. From Linear to $\log$ Dependence on Capacity

The first approach to improve the runtime of G-ALG is based on the observation that we only need to find an estimate of the index of the highest available unit. In particular, we can geometrically quantize the priority index of units for every resource. The rank of each unit of resource i now takes a value $\lfloor (1+\epsilon)^j \rfloor$ for some $j \in \{0, 1, \cdots, \lfloor \log_{1+\epsilon} c_i \rfloor\}$, where parameter $\epsilon > 0$ is a design choice that trades off the runtime with performance guarantee. Larger ϵ translates to smaller runtime and larger reduction in guarantee.

Since many units of a resource may now have the same index, for the sake of computation we treat all units with the same index as a single 'unit'. The improvement in runtime is immediate. In each iteration of the while loop (except the last) we decrease the index of the highest available 'unit' of at least one resource. So for $\epsilon > 0$, there are at most $O(\frac{1}{\epsilon}\sum_{i \in I} \log c_i)$ iterations to match each arrival. To understand the impact on the performance guarantee, notice that the reduced prices, $r_i\big(1 - g(\frac{z_i(t)}{c_i})\big)$, computed for fractionally matching each arrival are off by a factor of at most $(1 \pm \epsilon)$ for every resource. To address this in the analysis, we modify the candidate solution for the certificate as follows,

$$\lambda_t = \frac{1}{1-\epsilon} \sum_{i \in I} r_i \sum_{k \in [c_i]} y(k_i, t) \left(1 - g\Big(\frac{k}{c_i}\Big)\right).$$

The value of β in condition (2) is now larger by a factor of $(1-\epsilon)^{-1}$. The rest of Lemma 4 follows as is, resulting in an asymptotic guarantee of $(1-\epsilon)(1-1/e)$.

F.2. Capacity Independent Implementation

The second approach rests on the observation that if we could ensure that at each arrival, every unit either has at least a small $\epsilon > 0$ fraction available or is completely unavailable, then the number of units any arrival is partially matched to is at most $\frac{1}{\epsilon}$. We implement this idea in G-ALG by treating each unit as *unavailable unless at least ϵ fraction of it is available*. Here ϵ is a design choice; smaller the value, closer the competitive ratio guarantee to $(1-1/e)$ and larger the runtime. In particular, the time to match t now reduces to $O\big(\frac{|I|}{\epsilon}\big)$, which is within a $\frac{1}{\epsilon}$ factor of the time taken by much simpler Balance and RBA algorithms.

The deterioration in competitive ratio guarantee is more challenging to unravel. Fix a resource i and unit k, and recall the ordered set of arrivals $\boldsymbol{s}(k)$ from the analysis of Lemma 4. Roughly speaking, this is the set of arrivals where where unit k is unavailable in G-ALG. Observe that as a consequence of treating available fraction less than ϵ as unavailable, the set $\boldsymbol{s}(k)$ is now larger. To see the implication of this, define the $(F_i, \boldsymbol{T}, \boldsymbol{p}(k))$ process to capture actions of G-ALG, in the same way as Lemma 4. A key step in Lemma 4 involves showing that the expectation $r(F_i, \boldsymbol{T}, \boldsymbol{p}(k) \cup \mathbf{1}_{\boldsymbol{s}(k)})$, is the same as $r(F_i, \boldsymbol{T}, \boldsymbol{p}(k))$. However, as the set $\boldsymbol{s}(k)$ is now larger, this equality does not hold. Therefore, for a suitably defined non-negative valued function κ, we aim to show the weaker statement,

$$r\big(F_i, \boldsymbol{T}, \boldsymbol{p}(k) \cup \mathbf{1}_{\boldsymbol{s}(k)}\big) \leq \big(1 + \kappa(\epsilon)\big)\, r\big(F_i, \boldsymbol{T}, \boldsymbol{p}(k)\big), \tag{43}$$

If true, this would establish $(1-1/e)\big(1+\kappa(\epsilon)\big)^{-1}$–competitiveness (asymptotically) for G-ALG. Ideally, would like to show inequality (43) with a function κ that takes values as small as possible for every ϵ. Recall that for $\epsilon = 0$, we showed in Lemma 4 that inequality (43) holds with $\kappa(\epsilon) = 0$. However, for non-zero but small ϵ, it is not clear if inequality (43) holds with a small value $\kappa(\epsilon)$ in general.

Interestingly, this inequality has a remarkably strong connection to Proposition 1 in the analysis of RBA. In some sense, it is equivalent to Proposition 1. More concretely, one can show the inequality (43) for the families of usage distributions where we establish validity of Proposition 1 in this paper. For instance, when distribution F_i is exponential, inequality (43) holds with the linear function $\kappa(\epsilon) = 2\epsilon$, leading to a $O(n/\epsilon)$ algorithm with asymptotic guarantee $(1+\epsilon)^{-1}(1-1/e)$. More generally, we have $k(\epsilon) = \epsilon^{\eta}$ and a $(1+\epsilon^{\eta})^{-1}(1-1/e)$ guarantee for bounded IFR distributions, where η is as defined in case of RBA (Appendix E.1) Note that while the guarantees for RBA hold only for online matching, the guarantee for this modified version of G-ALG holds for budgeted allocation as well as assortment.

Appendix G: Challenge with Small Inventory: Connection to Stochastic Rewards

While our work provides algorithms with the best possible guarantee for reusable resources in the large inventory regime, finding an algorithm that outperforms greedy for small inventory remains open. It is worth noting that the case where all inventories are equal to 1 is the most general setting of the problem (see Proposition 1 in Gong et al. (2022)). In this section, we shed new light on the difficulty of this problem by establishing a connection between a very special case of reusability and the well studied problem of online matching with stochastic rewards.

Consider the setting where matched resources return immediately and can be re-matched to subsequent arrivals i.e., usage durations are deterministically 0. It is not surprising that the greedy solution is optimal for this instance, as the capacity of each resource is virtually unlimited. At the other end of the spectrum is the case of non-reusable resources where matched units never return. Now consider perhaps the simplest setting that captures both these extreme cases where every matched unit returns immediately with probability p, and never returns (usage duration $+\infty$) w.p. $1-p$. This setting isolates a key aspect of reusability – the stochastic nature of the problem. At first glance one might expect this setting to be straightforward given the observations for the extreme cases where $p = 0$ or $p = 1$. As it turns out, the general case is more interesting.

To make this formal, we consider the stochastic rewards problem of Mehta and Panigrahi (2012) for non-reusable resources. This problem generalizes online matching by associating a probability of success p_{it} with every edge $(i,t) \in E$. When a match is made i.e., edge is chosen, it succeeds independently with this probability. If the match fails the arrival departs but the resource is available for future rematch. The goal is to maximize the expected number of successful matches. We show the following connection.

Lemma G19 *The problem of online matching with reusability where resources have identical two point usage distributions supported on $\{0, +\infty\}$, is equivalent in the competitive ratio sense to the problem of online matching with stochastic rewards and identical edge probabilities, i.e., an α–competitive online algorithm in one setting can be translated to an α–competitive online algorithm in the other.*

Proof of the lemma is presented later in this section. For small inventory, the stochastic rewards problem is well known to be fundamentally different from classic online matching. For instance, Mehta and Panigrahi (2012) showed that when comparing against a natural LP benchmark, no online algorithm can have a guarantee better than $0.621 < (1 - 1/e)$ for stochastic rewards with small inventory, even with identical probabilities. While a recent result for stochastic rewards shows that this barrier can be circumvented by comparing directly against offline algorithms instead of LP benchmark[13], in general, the setting of *reusable resources sharply diverges and becomes much harder than stochastic rewards*. In particular, the stochastic rewards problem with heterogeneous edge probabilities admits a $1/2$–competitive result for arbitrary inventory.[14] In contrast, the corresponding generalization for reusable resources, where the probability of immediate return is arrival/edge dependent, does not admit any non-trivial competitive ratio result even for large capacity (Theorem 2 in Gong et al. (2022)).

To prove Lemma G19, we first introduce an equivalent form of stochastic rewards where the reward is deterministic and independent of the success/failure of matching. Formally, consider the stochastic rewards setting with identical success probability p for every edge and resources with unit reward. We transform this to an instance of the following problem:

Online matching with stochastic consumption: Each edge has a probability p of success. We assume that $p > 0$. If an arrival is matched to some resource, we earn a unit reward. After each match, a unit of the matched resource is used forever w.p. p, independent of other outcomes. With the remaining probability $1 - p$, we do not lose a unit of the resource. Recall, we earn a unit reward in either realization.

We can denote an instance of either of these problems simply as (G, p), where G is the graph and $p > 0$ is the edge probability. Now, consider the family of *non-anticipative algorithms* for these problems, i.e., algorithms (online or offline) that do not know the realization of any match beforehand. Every online algorithm is naturally non-anticipative. Offline algorithms such as the clairvoyant benchmark, as well the stronger fully offline benchmark that can match arrivals in an arbitrary sequence (Goyal and Udwani 2022), are both non-anticipative. Evaluating competitive ratios against non-anticipative offline algorithms, we have the following result.

Lemma G20 *The problem of online matching with stochastic consumption (with probability $p > 0$) is equivalent in the competitive ratio sense to the problem of online matching with stochastic rewards and identical edge probabilities, i.e., an α–competitive online algorithm in one setting can be translated to an α–competitive online algorithm in the other.*

[13] Goyal and Udwani (2022) give a $(1 - 1/e)$ result for instances of stochastic rewards with decomposable edge probabilities i.e., when for every edge $(i, t) \in E$ the probability can be decomposed as a product $p_{it} = p_i \times p_t$.

[14] This is achieved with greedy algorithms (Mehta et al. 2015, Golrezaei et al. 2014). Improving this is an open problem. A tight $(1 - 1/e)$ result is known for large capacity (Mehta et al. 2007, Golrezaei et al. 2014).

C onsider an instance of the stochastic rewards problem and a non-anticipative algorithm $\mathcal{A}$, that can be offline or online. Consider the alternate reward function where each time $\mathcal{A}$ makes a match we obtain a deterministic reward p regardless of the outcome of the match. Due to non-anticipativity of $\mathcal{A}$ and using the linearity of expectation, the expected total *alternative* reward of $\mathcal{A}$ is the same as its expected total reward.

Now, consider an instance (G,p) of the (online matching with stochastic) consumption problem with $p > 0$. Given algorithm $\mathcal{A}$ for the stochastic rewards problem, we can obtain an algorithm $\mathcal{A}'$ for the stochastic consumption problem by simulating $\mathcal{A}$ on a coupled instance (G,p) of the stochastic reward problem. The total expected reward of $\mathcal{A}'$ for the consumption problem is $\frac{1}{p}$ times the alternative reward of $\mathcal{A}$ on instance (G,p) for the stochastic rewards problem.

Observe that we can proceed in the reverse direction with similar arguments, i.e., given an algorithm $\mathcal{B}'$ for the consumption problem, we can construct an algorithm $\mathcal{B}$ for the stochastic rewards problem via simulating $\mathcal{B}'$ on a coupled instance of the consumption problem. The expected reward of $\mathcal{B}$ on an instance (G,p) of the stochastic rewards problem is p times the expected reward of $\mathcal{B}'$ on the instance (G,p) of the consumption problem.

Since these arguments hold for both online and offline algorithms, the constant factor of p (or $1/p$) cancels out for $p > 0$ and we have the desired competitive ratio equivalence. □

Proof of Lemma G19. It is now easy to see that the stochastic consumption problem is equivalent to the setting of online matching with reusable resources when the usage distributions for every resource is supported on $\{0, +\infty\}$, with probability of return $1-p$. To make this connection, we simply interpret unsuccessful consumption of a resource in the stochastic consumption setting as the resource returning with duration 0 in the reusable resources setting, and vice versa. Given this interpretation, we can now directly use an algorithm from one setting in the other setting with the same expected reward. □

More generally, consider the stochastic rewards problem with heterogeneous edge probabilities p_{it}. The greedy algorithm that matches each arrival to the resource with highest expected reward is $1/2$–competitive for this general problem. Now, consider the corresponding generalization in the reusable resource setting with two point usage distributions supported on $\{0, +\infty\}$ and return probability $1-p_{it}$ for edge (i,t). This problem does not admit any constant factor competitive ratio result (Theorem 2 in Gong et al. (2022)). The proof of equivalence breaks down since the expected reward of a match (i,t) in the stochastic rewards setting is now p_{it}. In contrast, the reward for a match in the reusable resource setting is 1, as before. As the ratio between these rewards is now arrival dependent (unlike the case of identical probabilities) it does not cancel out as a constant factor when evaluating the competitive ratios.

Appendix H: Miscellaneous

H.1. Impossibility for Stronger Benchmark

For online matching with reusable resources consider the offline benchmark that in addition to the arrival sequence also knows the realizations of all usage durations in advance. The following example illustrates that no non-trivial competitive ratio result is possible against this benchmark.

Consider a setting with n resources. Resources have identical reward and usage distribution. Consider a two point distribution uniformly supported on $\{0, \infty\}$. Suppose we see n^2 arrivals, each with an edge to all n resources. The expected reward of any online algorithm is at most $2n$, whereas an offline algorithm that knows the realizations of all durations in advance can w.h.p. match all n^2 arrivals as $n \to \infty$.

H.2. Clairvoyant is Deterministic

Lemma H21 *There exists a deterministic algorithm that is optimal among the class of all offline algorithms that know the entire arrival sequence but match (or decide assortments) in order of arrival and do not know realizations of stochastic elements (usage durations and customer choice) in advance.*

Proof. Given an arrival sequence the optimal algorithm is given by a dynamic program with the state space given by the number of arrivals remaining and the availability status of the resources, i.e., for each unit, whether it is currently available or in-use and how long it has been in-use for. The decision space of clairvoyant is simply the assortment decision for the current arrival in the sequence. Let $V(t, S_t)$ denote the optimal value-to-go at arrival t given that the state of resources is S_t. We use $\mathbb{1}(i, S_t)$ to indicate if a unit of i is available in state S_t and recall that $\mathcal{F}_t$ denotes the set of feasible assortments at t. Let $R(A, t+1 \mid S_t)$ denote a possible state of the resources at arrival $t+1$, given state S_t and assortment A at arrival t. Let $p(R(A, t+1 \mid S_t))$ denote the conditional probability of occurrence for this state and let $\Omega(A, t+1 \mid S_t)$ denote the set of all possible states $R(A, t+1 \mid S_t)$. Clearly,

$$V(t, S_t) = \max_{A \in \mathcal{F}_t \mid \mathbb{1}(i, S_t)=1, \forall i \in A} \left(R_t(A) + \sum_{R(A,t+1|S_t) \in \Omega(A,t+1|S_t)} p\left(R(A, t+1 \mid S_t)\right) V\left(t+1, R(A, t+1 \mid S_t)\right) \right),$$

where $R_t(A) = \sum_{i \in A} r_i \phi_t(A, i)$. At the last arrival T, the future value to go (second term in the above sum) is zero and for any given state S_T of resources at T, the optimal decision at T is simply the (deterministic) solution to a constrained assortment optimization problem. Performing a backward induction using the above equation, we have for any given set of values $V(t+1, R(A, t+1 \mid S_t))$, the optimal assortment decision at arrival t is deterministic. □

H.3. Clairvoyant Matches Fractional LP for Large Capacities

Consider the following natural LP upper bound for online matching with reusable resources (Dickerson et al. 2018, Baek and Ma 2022, Feng et al. 2019),

$$OPT(LP) = \max \sum_{(i,t) \in E} r_i y_{it}$$

$$s.t. \quad \sum_{t=1}^{\tau}[1-F_i(a(\tau)-a(t))]y_{it} \le c_i \quad \forall \tau \in \{1,\ldots,T\}, \forall i \in I$$
$$\sum_{i\in I} y_{it} \le 1 \quad \forall t \in T$$
$$0 \le y_{it} \le 1 \quad \forall t \in T, \quad i \in I \tag{44}$$

Clearly, OPT $\le$ OPT(LP) and the allocations generated by any algorithm (offline or online) can be converted into a feasible solution for the LP, regardless of c_i. Perhaps surprisingly, we show that for large c_i, the solution to this LP can be turned into a randomized clairvoyant algorithm (that does not know the realizations of usage in advance) with nearly the same expected reward, implying that the LP gives a tight asymptotic bound and moreover, all asymptotic competitive ratios shown against the clairvoyant also hold against the LP.

Theorem H22 *Let $c_{\min} = \min_{i\in I} c_i$. Then,*

$$\text{OPT}(LP)\Big(1-O\Big(\sqrt{\frac{\log c_{\min}}{c_{\min}}}\Big)\Big) \le \text{OPT} \le \text{OPT}(LP).$$

Hence, for $c_{\min} \to +\infty$, OPT $\to$ OPT(LP).

Proof. We focus on the lower bound and more strongly show that every feasible solution of the LP can be turned into an offline algorithm with nearly the same objective value. Let $\{y_{it}\}_{(i,t)\in E}$ be a feasible solution for the LP. Consider the offline algorithm that uses the LP solution as follows,

When t arrives, sample a resource i to offer, according to the distribution $\{y_{it}/(1+2\delta)\}_{i\in I}$,

where $\delta = \sqrt{\frac{\log c_{\min}}{c_{\min}}}$. Note that if the sampled resource is unavailable, the algorithm leaves t unmatched. Also w.p., $1-\sum_{i\in I} y_{it}$, the algorithm rejects t. Since the LP does not use usage durations, the offline algorithm doesn't either. The critical element to be argued is that the expected reward of this algorithm is roughly the same as the objective value for the feasible solution. This holds due to concentration bounds and the argument closely mimics the proof of Lemma 11. Finally, since this offline algorithm makes matching decisions in order of the arrival sequence and does not know realizations of usage durations in advance, its performance gives a lower bound on the performance of clairvoyant, i.e., OPT. ☐

Remark: This result generalizes naturally to the settings of online assortment and budgeted allocations. In case of assortments we use the Probability Matching algorithm from Appendix D to use the concentration bounds.

H.4. Sufficiency of Static Rewards

Lemma H23 *Given an algorithm (online or clairvoyant) for allocation that does not know the realizations of usage durations, the total expected reward of the algorithm is the same if we replace dynamic usage duration dependent rewards functions $r_i(\cdot)$ with their static (finite) expectations $r_i = \mathbb{E}_{d_i \sim F_i}[r_i(d_i)]$, for every resource.*

Proof. Consider arbitrary algorithm $\mathbb{A}$ as described in the lemma statement. Let $\mathbb{B}$ denote an algorithm that mimics the decisions of $\mathbb{A}$ but receives static rewards $r_i, \forall i \in I$ instead. Now, suppose $\mathbb{A}$ successfully allocates resource i to arrival t on some sample path $\omega(t)$ observed thus far. Then, conditioned on observing $\omega(t)$, the expected reward from this allocation is exactly r_i. More generally, using the linearity of expectation it follows that the total expected reward of $\mathbb{B}$ is the same as that of $\mathbb{A}$.

Note that for algorithms that also know the realization of usage durations in advance, the decision of allocation can depend on usage durations and conditioned on observed sample path $\omega(t)$, the expected reward from successful allocation of i to t need not be r_i. □

H.5. Upper Bound for Deterministic Arrival Dependent Usage

Lemma H24 *Suppose the usage duration is allowed to depend on the arrivals such that a resource i matched to arrival t is used for deterministic duration d_{it} (revealed when t arrives). Then there is no online algorithm with a constant competitive ratio bound when comparing against offline algorithms that know all arrivals and durations in advance.*

Proof. Using Yao's minimax, it suffices to show the bound for deterministic online algorithms over a distribution of arrival sequences. For simplicity, suppose we have a single unit of a single resource and a family of arrival sequences $A(j)$ for $j \in [n]$ (the example can be naturally extended to the setting of large capacity). The arrivals in the sequences will be nested so that all arrivals in $A(j)$ also appear in $A(j+1)$. $A(1)$ consists of a single arrival with usage duration of 1. Suppose this vertex arrives at time 0. $A(2)$ additionally consists of two more arrivals, each with usage duration of $1/2 - \epsilon$, arriving at times ϵ and $1/2$ respectively. More generally, sequence $A(j)$ is best described using a balanced binary tree where every node represents an arrival and the depth of the node determines the usage duration. Each child node has less than half the usage duration ($d/2 - \epsilon$) of its parent (d). If the parent arrives at time t, one child arrives at time $t + \epsilon$ and the other at $t + d/2$. The depth of the tree for sequence $A(j)$ is j (where depth 1 means a single node). Note that the maximum number of arrivals that can be matched in $A(j)$ is 2^{j-1}.

Let $Z = \sum_{j=1}^{n} 2^j$. Now, consider a probability distribution over $A(j)$, where probability p_j of sequence $A(j)$ occurring is $\frac{2^{n-j+1}}{Z}$. Clearly, an offline algorithm that knows the full sequence in advance can match 2^{j-1} arrivals on sequence $A(j)$ and thus, has revenue $n2^n/Z = n/2$. It is not hard to see that the best deterministic algorithm can do no better (in expectation over the random arrival sequences) than trying to match all arrivals with a certain time duration. Any such deterministic algorithm has revenue at most $Z/Z = 1$. Therefore, we have a competitive ratio upper bound of $n/2$. □

Appendix I: Numerical Experiments: Missing Details

I.1. LP Benchmark

Let p_{it} denote the probability that arrival t has an edge to i. Observe that, if t is a bursty arrival in phase k, then

$$p_{i,t} = \begin{cases} 1 & i = 2n - k + 1, \\ 0 & \text{otherwise.} \end{cases}$$

If t is a normal arrival in phase k, then

$$p_{i,t} = \begin{cases} \frac{\sum_{\ell \geq i} e^{-\kappa|\ell-(n-k+1)|}}{\sum_{\ell \geq 1} e^{-\kappa|\ell-(n-k+1)|}} & \forall i \in [n], \\ 0 & \text{otherwise.} \end{cases}$$

Clairvoyant knows the arrival sequence for each random instance. The optimal solution of the following LP is an upper bound on the expected total reward of clairvoyant.

$$\begin{aligned} \textbf{LP benchmark} \quad \max \quad & \sum_{i \in [n], t \in [2cn]} y_{i,t} \\ s.t. \quad & \sum_{t=1}^{\tau} [1 - F(a(\tau) - a(t))] y_{i,t} \leq c \quad \forall \tau \in [2cn], \forall i \in [n] \\ & \sum_{i \in [n]} y_{i,t} \leq 1 \quad \forall t \in [2cn] \\ & 0 \leq y_{i,t} \leq p_{i,t} \quad \forall t \in T[2cn], \quad i \in [n]. \end{aligned} \tag{45}$$

This LP benchmark closely resembles the standard LP relaxation of clairvoyant (see Appendix H.3). The key difference is that we impose the upper bounds $y_{i,t} \leq p_{i,t}$ (see (45)). Let $E_{i,t}$ denote the event that there is an edge between resource i and arrival t. Let $O_{i,t}$ denote the event that the clairvoyant matches t to i conditioned on the event $E_{i,t}$. To see that this LP is an upper bound on the expected performance of clairvoyant, observe that, if, for every $i \in [n]$ and $t \in [2cn]$, we set the decision variable $y_{i,t}$ as the probability that both $E_{i,t}$ and $O_{i,t}$ occur then all constraints in the LP will be satisfied and the LP objective represents the expected total reward of clairvoyant.

I.2. Other Scenarios

For Tables 4 – 6, the performance of each algorithm (ALG) in the table is reported as the ratio of the empirical average performance of ALG (based on 20×100 trials) and the optimal value of the LP benchmark. Note that the standard deviation of the reported ratios is less than 0.0001 for all algorithms except Sample G-ALG, for which it is less than 0.01.

c	F	**RBA**	**Balance**	**Greedy**	**Sample G-ALG**
5	Two-point	0.72485	0.72485	0.58894	0.5129
	Exponential	0.9901	0.9932	0.9648	0.98907
	Weibull	0.9972	0.9968	0.9732	0.91239
15	Two-point	0.79477	0.79477	0.61379	0.6192
	Exponential	0.99907	0.9996	0.97867	0.99907
	Weibull	0.9999	0.9998	0.9844	0.9999

Table 4 Average performance of online algorithms in comparison to the LP benchmark when the arrival sequence only includes the sequence of normal arrivals (T_2), i.e., the bursty arrivals are excluded. The results are for $n = 5$ and $\kappa = 1$. Observe that RBA, Balance, and greedy are all close to optimal in most scenarios and RBA and Balance have very comparable performance. For two-point usage distribution, we believe that the overall instance is relatively closer to the worst case instance for non-reusable resources and this may be why greedy performs relatively poorly.

c	F	**RBA**	**Balance**	**Greedy**	**Sample G-ALG**
5	Two-point	0.87617	0.85191	0.81745	0.75748
	Exponential	0.92708	0.92266	0.91062	0.79072
	Weibull	0.94267	0.94025	0.93380	0.81921
15	Two-point	0.90468	0.86780	0.81277	0.82592
	Exponential	0.91876	0.90630	0.90978	0.8595
	Weibull	0.94677	0.93587	0.93102	0.884263
25	Two-point	0.91813	0.87634	0.82596	0.90428
	Exponential	0.91950	0.90333	0.90711	0.91942
	Weibull	0.94641	0.93414	0.93050	0.945711

Table 5 Average performance of algorithms for $n = 5$ and $\kappa = 0$ in comparison to LP benchmark. Comparing the numbers with Table 3 for $n = 5$ and $\kappa = 1$, one can see that all algorithms perform better for $\kappa = 0$. Greedy now slightly outperforms the other algorithms in some scenarios. RBA continues to dominate Balance but the gap between them is much smaller as compared to the scenario where $\kappa = 1$.

c	F	**RBA**	**Balance**	**Greedy**	**Sample G-ALG**
5	Two-point	0.73806	0.72307	0.68854	0.6382
	Exponential	0.9675	0.9665	0.9751	0.91201
	Weibull	0.97312	0.97202	0.97742	0.95657
15	Two-point	0.74498	0.72151	0.68462	0.73217
	Exponential	0.95493	0.94877	0.9732	0.9441
	Weibull	0.96362	0.95792	0.97379	0.96991
25	Two-point	0.76864	0.74169	0.69747	0.75311
	Exponential	0.94867	0.93771	0.96965	0.94631
	Weibull	0.95822	0.94865	0.96916	0.95812

Table 6 Average performance of algorithms for $n = 20$ and $\kappa = 1$ in comparison to LP benchmark. The overall trend in this case is similar to the case of $n = 5$ and $\kappa = 0$ (see Table 5 above). For two-point usage distribution, we believe that for a larger value of n the overall instance is relatively closer to the worst case instance for non-reusable resources and this may be why greedy performs relatively poorly.

c	F	**RBA**	**Balance**	**Greedy**	**Sample G-ALG**
5	Two-point	0.84941	0.82502	0.77134	0.68801
	Exponential	0.9785	0.9776	0.9757	0.91239
	Weibull	0.97984	0.97944	0.97654	0.96112
15	Two-point	0.88137	0.84527	0.77988	0.75819
	Exponential	0.97667	0.97393	0.9735	0.95613
	Weibull	0.9784	0.97583	0.97583	0.97014
25	Two-point	0.89713	0.86034	0.82436	0.76183
	Exponential	0.92088	0.90234	0.91402	0.91015
	Weibull	0.94488	0.93114	0.93002	0.92137

Table 7 Average performance of algorithms for $n = 20$ and $\kappa = 0$ in comparison to LP benchmark. Comparing the numbers with Table 6 for $n = 20$ and $\kappa = 1$, one can see that all algorithms perform better for $\kappa = 0$. RBA continues to dominate all algorithms by a small margin. Greedy now performs comparably with Balance in some scenarios. For two-point usage distribution, we believe that for a larger value of n the overall instance is relatively closer to the worst case instance for non-reusable resources and this may be why greedy performs relatively poorly.