When Should you Offer an Upgrade: Online Upgrading Mechanisms for Resource Allocation

Patrick Jaillet

Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, jaillet@mit.edu

Chara Podimata

Sloan School of Management, Massachusetts Institute of Technology, podimata@mit.edu

Andrew Vakhutinsky

Oracle Lab, andrew.vakhutinsky@oracle.com

Zijie Zhou

Operations Research Center, Massachusetts Institute of Technology, Cambridge, MA, 02139, zhou98@mit.edu

In this work, we study an upgrading scheme for online resource allocation problems. We work in a sequential setting, where at each round a request for a resource arrives and the decision-maker has to decide whether to accept it (and thus, offer the resource) or reject it. The resources are ordered in terms of their value. If the decision-maker decides to accept the request, they can offer an *upgrade-for-a-fee* to the next more valuable resource. This fee is dynamically decided based on the currently available resources. After the upgrade-for-a-fee option is presented to the requester, they can either accept it, get upgraded, and pay the additional fee, or reject it and maintain their originally allocated resource.

We take the perspective of the decision-maker and wish to design upgrading mechanisms in a way that simultaneously maximizes revenue and minimizes underutilization of resources. Both of these desiderata are encapsulated in a notion of *regret* that we define, and according to which we measure our algorithms' performance. We present a fast algorithm that achieves $O(\log T)$ regret. Finally, we implemented our algorithm utilizing data akin to those observed in the hospitality industry and estimated our upgrading mechanism would increase the annual revenue by over 17%.

1. Introduction

In online resource allocation, we model the interplay between a *supplier* that allocates resources and *requesters* that arrive sequentially and request resources. In general, the resources can be of variable value both to the supplier and the requester, i.e., some are more premium than others. The supplier is tasked with making real-time, irrevocable decisions concerning the method and pricing of the resource to be allocated, despite lacking complete foresight into future demand. Once allocated, the requester provides some payment to the supplier in return. Online resource allocation models have been studied for many different settings such as airline seat allocation (Bertsimas and De Boer 2005, Zhang and Cooper 2009) and hotel room allocation (Ivanov and Zhechev 2012, Ivanov 2014); their impact has been tremendous both in the academic and the industry world.

An important challenge in online resource allocation occurs when a supplier does not have information about the future realized demand from the requesters and the availability of the resources remains constant. For example, the number of seats on an airplane or rooms in a hotel is fixed, while demand is subject to variability. Demand can surge during certain seasons and shrink during others, but the supplier lacks the flexibility to adjust its resource availability accordingly. As a consequence, there are cases when premium resources can be underutilized. For example, according to Travel (2020), approximately 80% of hotel rooms in the United States remain vacant. Addressing this inefficiency calls for the development of advanced resource allocation algorithms, ensuring that even during periods of low demand for premium resources, they are optimally utilized.

One popular idea for ensuring that premium resources do not go underutilized is to employ some *upgrading* scheme. Existing models in the literature (Shumsky and Zhang 2009, Yu et al. 2015,

Cui et al. 2023) primarily assume that suppliers can only offer upgrades to requesters at no extra charge. This assumption (although helpful in simplifying the mathematical modeling) diverges from what happens in reality, where upgrades come at an *extra cost*. For instance, when booking a hotel room through an online platform, it is common for a requester to encounter offers for upgrading to a premium room at an additional fee, typically highlighted with prompts such as "Would you like to upgrade to a premium room for an additional $50?" This reflects the reality that upgrades are not always complimentary, and determining the optimal pricing for such upgrades is crucial for effective revenue management.

1.1.   Our Contributions and Techniques

Model. We subscribe to the agenda of using upgrades as a way of using premium resources properly and introduce an online upgrading mechanism *for a fee*. At a high level, our model is as follows. Resources of interest (e.g., airline/railway seats, hotel rooms, ride-hailing services, gaming equipment etc.) exhibit an ascending order in both quality and cost. The online upgrading mechanism operates by sequentially processing resource requests. Decisions to accept or reject these requests, along with potential upgrade plans, are made based on factors like arrival rates, remaining availabilities, and associated costs. An upgrade plan entails offering the customer an option to move to the next superior resource category at an additional price. Customers can either accept the upgrade, incurring the extra cost for better resources, or continue using their initially chosen resource at the original price. Our formal model is given in Section 2.

Upgrading Mechanism for Online Resource Allocation. Our main contribution is the development of a mechanism that we call DynUp (stands for "Dynamic Upgrading") for identifying *when* and *for how much* a supplier should offer an upgrade. For ease of exposition, we first explain the mechanism DynUp-2 and its analysis for the simpler case of 2 resource types in Section 3. We use DynUp-2 as a building block for our general mechanism DynUp-n (presented in Section 4) which works for n resource types. The performance of our algorithms is measured in terms of *regret* to an in-hindsight optimum benchmark. Roughly, our algorithms achieve the following guarantees.

THEOREM 1 **(Informal)**. *For online resource allocation with 2 resource types, DynUp-2 incurs regret $O(\log T)$. For online resource allocation with n resource types, DynUp-n incurs regret $O(n \log T)$.*

Before we delve into the high level description of our mechanism for the 2-type resource case, we first need to explain the benchmark we are comparing against. At first sight, it may seem tricky to establish a reliable benchmark for the online upgrading problem since at every period we not only have to decide whether to accept or reject a request but also decide on an upgrade price if we do decide to accept. Importantly, this decision at a period t crucially affects the space of allowable decisions for future periods; this is because if we do offer an upgrade price that is ultimately accepted, then we are taking one unit of the premium resource away from future rounds. To navigate this challenge, we propose a "hybrid programming" approach which corresponds to an upper bound of the hindsight optimal expected revenue. The "hybrid programming" interpolates between a *deterministic* and a *hindsight optimum* counterpart, which draw intuition from approaches traditionally found in price- and quantity- based revenue management respectively. As it will become clear moving forward, we use the "hybrid programming" in the development of algorithms DynUp-2 and DynUp-n. Detailed explanations of the hybrid programming and the benchmark for 2 and n resource types can be found in Sections 3.1 and 4.1 respectively.

At a high level, Algorithm DynUp-2 determines the upgrading price at each time period t by solving the aforementioned hybrid programming with the remaining availability of each resource and remaining expected demand for each type of resource at t. Although the hybrid programming is a non-convex optimization problem, we show that we can obtain a closed-form solution to it. To understand the $O(\log T)$ regret of DynUp-2, we need to study how the upgrade pricing decisions of DynUp-2 evolve over time. We distinguish between two cases. First, when the initial resource

availability exceeds the expected demand, the stochastic process of DynUp-2's pricing decisions forms a martingale with an expected value change of 0. Therefore, by comparing this martingale process at each period and the optimal decision variable of "hybrid programming", from a known result of Jasin and Kumar (2012), we can obtain that the regret can be upper bounded by $O(\log T)$. Second, when the resource availability falls short of expected demand, DynUp-2's dynamic pricing decision is not a martingale. To address this bottleneck we construct a complicated stochastic process which can be shown to be a martingale, closely mirroring the pattern of the actual dynamic pricing decision made by DynUp-2 so that the distance of this martingale and the dynamic pricing decision made by DynUp-2 at each period can be upper bounded by $O(1/T^2)$.

We then extend our findings to the scenario involving n types of resources. Algorithm DynUp-n is implemented by incorporating Algorithm DynUp-2 between consecutive resource types, i.e., between type i and type $i+1$, for all $i \in [n-1]$. In Section 4, we show that the online upgrading problem with n types of resources can be fully decomposed to $n-1$ "parallel" problems with 2 types of resources, leading to a regret bound of $O(n \log T)$. Practically, the number of resource types does not increase with time T, making n a fixed parameter. So, under the assumption that n is constant, Algorithm DynUp-n achieves a regret bound of $O(\log T)$, demonstrating its efficiency and scalability for managing resources across various types.

Empirical Study in Hotel Management. Our empirical investigation involved a real dataset from the hospitality industry. This dataset encompassed 13,155 requests for rooms spanning one year. Applying Algorithm DynUp-n to this real-world data led to a significant increase in annual revenue by 17%. Notably, in the high-demand month of November, the implementation of Algorithm DynUp-n resulted in a revenue increase exceeding 26%.

1.2. Related Work

Price-based Revenue Management. Gallego and Van Ryzin (1994, 1997) introduced the price-based revenue management model. This model processes requests sequentially, with the central decision revolving around the pricing strategy. Typically, higher pricing leads to reduced demand. The authors provided an optimal static pricing technique that achieves an $O(\sqrt{T})$ regret. Subsequently, Jasin and Kumar (2012), Jasin (2014) developed a dynamic pricing strategy that entails solving a deterministic program for each incoming request, demonstrating that this approach results in an $O(\log T)$ regret. Further advancing this field, Wang and Wang (2022) employed an innovative technique to establish that the dynamic pricing method attains an $O(1)$ regret. Additionally, there are variations to these models, such as Aydin and Ziya (2008), which explored dynamic pricing in the context of product bundling.

While our online upgrading process shares certain similarities with price-based revenue management, there are two key distinctions to consider. First, from a modeling standpoint, our approach requires a decision not just on pricing but also on whether to accept each request. This dual consideration positions our model as a hybrid of quantity-based (Reiman and Wang 2008, Bumpensanti and Wang 2020, Sun et al. 2020) and price-based models. Second, from technical standpoint, all literature on price-based revenue management assumes a concave revenue function, leading to an online convex optimization problem. In contrast, our model, incorporating both quantity and price variables, is characterized by a non-concave revenue function.

Upgrading Models. The domain of upgrading models has seen various contributions. Shumsky and Zhang (2009) proposed a dynamic model where upgrades are complimentary, allowing each request to be upgraded only to the immediately superior resource category. McCaffrey and Walczak (2016) employed dynamic programming to effectively resolve the issue of free upgrading in a two-type resource scenario. Extending this concept, Yu et al. (2015) introduced a multi-step free upgrading framework, permitting requests to be upgraded to any higher-tier resource without cost. Cui et al. (2023) further expanded this model by transitioning from reactive to proactive upgrading. These studies, however, diverge significantly from our work as they are predicated on the assumption of costless upgrading, aligning them more with quantity-based revenue management. In

contrast, Gallego and Stefanescu (2009) introduced the concept of "upsell" as a non-complimentary upgrade in their research as an extension. They focused on static upselling strategies but did not provide any theoretical performance guarantees for this method.

2. Model

We consider a setting where a supplier wants to allocate n distinct resources to interested individuals (aka *requesters*). Each resource has a different type $i \in [n]$. For each resource type $i \in [n]$, there is an associated availability c_i and a cost r_i. We assume that resources are ordered in terms of their cost to the requesters and that higher costs are associated with higher value for all the requesters, i.e., $r_1 < r_2 < \ldots < r_n$ and the n^{th} resource type is the most premium. On the request side, requests arise sequentially across a discrete time horizon spanning T periods. In any given period, at most one requester may show up, and request precisely one unit of some resource. The request arrival rate for resource type i is known and denoted by λ_i, and λ_0 represents the rate of no request. We normalize the arrival rates such that $\sum_{i=0}^{n} \lambda_i = 1$. Consequently, the probability of a request for resource type i within a specific period is λ_i, and the probability for no request is λ_0.

When a request for a resource of type i emerges, the supplier is faced with the decision to either accept or reject this request. If the request is rejected, there are no resources allocated. On the other hand, if the request is accepted, the supplier offers an upgrade plan[1]: would the requester be willing to pay an additional cost, $u_i \in [0, r_{i+1} - r_i]$, in exchange for the superior type $i+1$ resource[2]? If this plan is accepted by the requester, the supplier assigns one unit of the type $i+1$ resource, leading the requester to pay a combined cost of $r_i + u_i$. Conversely, a rejection results in the allocation of the type i resource, with a corresponding cost of r_i. We posit that any request for a resource of type i has a $f_i(u_i)$ likelihood of accepting the upgrade for an additional cost of u_i. The overall objective for the supplier is to *maximize the revenue* obtained given the resources and their availabilities, and the designated upgrading offer mechanism. We use *regret* to measure the performance of the online upgrading mechanism, where the regret is defined as the hindsight optimal revenue $\mathbb{E}[W^{\text{OPT}}]$ minus the expected cumulative revenue generated by the mechanism:

$$R(T) = \mathbb{E}[W^{\text{OPT}} - W^{\pi}]$$

where W^{π} denotes the revenue generated by a mechanism π.

2.1. Model Assumptions and Notations

We assume that the likelihood functions $f_i(\cdot)$ are known both by the supplier and the requesters. These functions are bijective, continuous, and monotone-decreasing. Further, let $f_i(0) = 1$, capturing the fact that no individual would decline a free upgrade. The inverse function of $f_i(\cdot)$ is denoted by $p_i(\cdot)$ (e.g., if $v_i = f_i(u_i)$, then $u_i = p_i(v_i)$). In other words, function $p_i(v)$ returns the upgrade price u such that the probability that u is accepted for the upgrading from resource i to $i+1$ is $f_i(u)$. Furthermore, the functions $p_i(\cdot)$ are also bijective, continuous, and monotone decreasing. Since $f_i(0) = 1$, then $p_i(1) = 0$. The expected revenue arising from the upgrade are captured by $R_i(v_i) = v_i p_i(v_i) = u_i f_i(u_i)$. We assume that the functions $R_i(\cdot)$ are bounded and quasi-concave. The assumption of quasi-concavity aligns with common functional forms of $p_i(\cdot)$, such as linearly and exponentially decreasing functions, which inherently lead to quasi-concave $R_i(\cdot)$. We also assume that $R_i(\cdot)$ has a finite maximizer, denoted by $v_i^\star$, which lies in the interval $[f_i(r_{i+1} - r_i), 1]$.

---

[1] If $i = n$ or the type $i+1$ resource has availability less than 1, no upgrade plan will be offered.

[2] If we accept the type i request, but type i resource has availability less than 1, we can only upgrade this request to the type $i+1$ resource with $u_i = 0$.

3.   Algorithm for Two Types of Resources

In this section, we focus on the setting with two distinct resource types where only one of them can be upgraded to the other. We assume that the availability of the premium resource (type 2) exceeds its expected demand, namely $c_2 > \lambda_2 T$. Otherwise, the problem becomes trivial as there would be no need for upgrading. We first describe the optimal stochastic and deterministic formulations of the online upgrading mechanism to introduce the upper bound performance for any upgrading algorithm (Section 3.1). Second, we present our algorithm, and its regret analysis in Section 3.2.

3.1.   Benchmark Revenue Upper Bound

In each period $t \in [T]$, let $\Lambda_i^{(t)}$ be the binary variable that represents whether the type i request arrives in period t, and $y_i^{(t)}$ be the binary variable that indicates whether the supplier accepts the type i request. Let $D_i^{(t)}$ represent the quantity of type i resources consumed in period t. For example, at period t, if a type 1 request arrives and accepts the upgrading plan, then no type 1 resource is used and one unit of type 2 resource is consumed, and hence $D_1^{(t)} = 0$ and $D_2^{(t)} = 1$. Note that $D_i^{(t)}$ is a random variable, influenced by the upgrading cost u_i. Hence, the optimal cumulative revenue can be expressed as the following stochastic programming:

$$W^{\text{OPT}} = \max_{\mathbf{y_1}, \mathbf{y_2}, \mathbf{u_1}} \sum_{t \in [T]} \mathbb{E}\left[\Lambda_1^{(t)} y_1^{(t)}\left(r_1 D_1^{(t)} + (r_1 + u_1^{(t)})D_2^{(t)}\right) + \Lambda_2^{(t)} y_2^{(t)} r_2 D_2^{(t)}\right]$$

$$\text{s.t.} \sum_{t \in [T]} \Lambda_1^{(t)} y_1^{(t)} D_1^{(t)} \le c_1$$

$$\sum_{t \in [T]} (\Lambda_1^{(t)} y_1^{(t)} + \Lambda_2^{(t)} y_2^{(t)})D_2^{(t)} \le c_2, \tag{SP}$$

where the expectation is taken on the randomness of the arrival process, and the constraints must hold almost surely. In (SP), we maximize the sum of the expected revenue over all periods $t \in [T]$. In each period t, if a type 1 request shows up ($\Lambda_1^{(t)} = 1$) and is accepted ($y_1^{(t)} = 1$), the expected revenue generated is $r_1 D_1^{(t)}$ (the request rejects the upgrading plan) plus $(r_1 + u_1^{(t)})D_2^{(t)}$ (the request accepts the upgrading plan). If a type 2 request shows up ($\Lambda_2^{(t)} = 1$) and is accepted ($y_2^{(t)} = 1$), the expected revenue generated is $r_2 D_2^{(t)}$. The first constraint guarantees that the sum of the type 1 resource consumed over all periods does not exceed c_1, and the second constraint ensures that the sum of the type 2 resource consumed (type 1 request accepts the upgrading plan or type 2 request is accepted) over all periods does not surpass c_2.

Observe that SP only outlines the decision-making process under the condition that both resource types remain available; it does not account for scenarios where either resource is exhausted. For example, in cases where type 2 resource runs out, there is no upgrading option. Under these circumstances, the approach mandates rejecting all type 2 requests, while accepting those from type 1, provided there is remaining availability in type 1 resources. To solve the (SP) we can use the following dynamic program:

$$W_t(c_1^{(t)}, c_2^{(t)}) = \max_{y_1^{(t)}, y_2^{(t)}, u_1^{(t)}} \mathbb{E}\left[\Lambda_1^{(t)} y_1^{(t)}\left(r_1 D_1^{(t)} + (r_1 + u_1^{(t)})D_2^{(t)}\right) + \Lambda_2^{(t)} y_2^{(t)} r_2 D_2^{(t)}\right.$$

$$\left. + W_{t+1}(c_1^{(t)} - \Lambda_1^{(t)} y_1^{(t)} D_1^{(t)}, c_2^{(t)} - (\Lambda_1^{(t)} y_1^{(t)} + \Lambda_2^{(t)} y_2^{(t)})D_2^{(t)})\right].$$

where $c_1^{(t)}, c_2^{(t)}$ are the remaining availability for type 1, 2 resource respectively at time t. However, the curse of dimensionality presents a significant challenge to computing the exact solution of the dynamic program. To counteract this problem, a prevalent approach in the context of price-based revenue management is to formulate a *deterministic* counterpart to the stochastic program defined in (SP) (Gallego and Van Ryzin 1994, 1997). This is achieved by substituting the upgrade price random variable $v_1^{(t)} = f_1(u_1^{(t)})$ with a fixed variable v_1. However, contrary to price-based revenue

management, in our problem, there also exists a *quantity-based* variable $y_i^{(t)}$ capturing whether to accept or reject the arriving type i request and we need to account for it. A well-known formulation that can be treated as an upper bound of the (SP) and contains quantity-based variables is the *hindsight optimum formulation* (Reiman and Wang 2008, Bumpensanti and Wang 2020), where we define Y_i as the total number of type i requests to accept, and define Λ_i as the aggregate number of arrivals for each type i. Here, Λ_i is a binomial random variable, namely $\Lambda_i \sim \text{Bin}(T, \lambda_i)$, and in the hindsight formulation, one assumes that the value of Λ_i is known in advance.

Putting everything together, we propose a *hybrid programming* that integrates both the *hindsight* and *deterministic* formulations, where Λ_i is the hindsight variable and v_1 is the deterministic variable. This formulation will henceforth be referred to as "hybrid programming" (HP).

$$
\begin{aligned}
w^{\text{H}} = \max_{Y_1, v_1} \Lambda_2 r_2 &+ Y_1 \left((r_1 + p(v_1))v_1 + r_1(1 - v_1) \right) \\
\text{s.t. } & Y_1 \leq \Lambda_1 \\
& Y_1 v_1 \leq c_2 - \Lambda_2 \\
& Y_1 (1 - v_1) \leq c_1.
\end{aligned}
\tag{1}
$$

We provide some more details about the optimization problem in (1). The objective function captures the total revenue. Since we always accept type 2 requests when there is availability, the revenue generated by type 2 requests is $\Lambda_2 r_2$. For type 1 requests, we accept Y_1 of them and each request will generate an expected revenue of $(r_1 + p(v_1))v_1$ (accept the upgrading plan) plus $r_1(1 - v_1)$ (reject the upgrading plan). The first constraint ensures that the number of type 1 requests accepted is no larger than the total number of type 1 arrivals. The second constraint guarantees that the total number of upgraded type 1 requests is less than or equal to the availability of type 2 resource minus the number of type 2 requests. The third constraint suggests that the total number of type 1 requests who reject the upgrading plan is less than or equal to the availability of type 1 resource. As $(r_1 + p(v_1))v_1 + r_1(1 - v_1)$ is always positive, the optimization problem (1) is equivalent to the following hybrid of hindsight and deterministic programming:

$$
w^{\text{H}} = \max_{v_1} \Lambda_2 r_2 + \min \left\{ \Lambda_1, \frac{c_2 - \Lambda_2}{v_1}, \frac{c_1}{1 - v_1} \right\} (R_1(v_1) + r_1).
\tag{HP}
$$

Let w^{H} be the optimal objective value; note that w^{H} is a random variable that depends on Λ_1 and Λ_2. We define $W^{\text{H}} = \mathbb{E}[w^{\text{H}}]$.

Recall that (SP) only captures the decision-process only up to the point where a resource runs out. Consequently, W^{H} should be regarded as an approximate upper bound for the total revenue achievable prior to the depletion of any resources. In our framework, decision-making ends either at the exhaustion of both resources or upon surpassing the designated time horizon, denoted as $[0, T]$. Accordingly, it becomes necessary to split the analysis based on the order of resource depletion, with comprehensive elaborations provided in the following theorem.

THEOREM 2. *Define three sets* $\mathcal{V}_1$, $\mathcal{V}_2$, $\mathcal{V}_3$ *such that* $\mathcal{V}_1 = \left\{ v_1 : \Lambda_1 = \min\{\Lambda_1, \frac{c_2 - \Lambda_2}{v_1}, \frac{c_1}{1-v_1}\} \right\}$, $\mathcal{V}_2 = \left\{ v_1 : \frac{c_2 - \Lambda_2}{v_1} = \min\{\Lambda_1, \frac{c_2 - \Lambda_2}{v_1}, \frac{c_1}{1-v_1}\} \right\}$, *and* $\mathcal{V}_3 = \left\{ v_1 : \frac{c_1}{1-v_1} = \min\{\Lambda_1, \frac{c_2 - \Lambda_2}{v_1}, \frac{c_1}{1-v_1}\} \right\}$.

For $v_1 \in \mathcal{V}_1$:
$$
w^{\text{U}_1} = \max_{v_1 \in \mathcal{V}_1} \Lambda_2 r_2 + \Lambda_1 (R_1(v_1) + r_1)
$$

For $v_1 \in \mathcal{V}_2$:
$$
w^{\text{U}_2} = \max_{v_1 \in \mathcal{V}_2} \Lambda_2 r_2 + \frac{c_2 - \Lambda_2}{v_1} R_1(v_1) + r_1 \min\{\Lambda_1, c_1 + c_2 - \Lambda_2\}
$$

For $v_1 \in \mathcal{V}_3$:
$$
w^{\text{U}_3} = \max_{v_1 \in \mathcal{V}_3} \Lambda_2 r_2 + \frac{c_1}{1 - v_1} R_1(v_1) + r_1 \min\{\Lambda_1, c_1 + c_2 - \Lambda_2\}.
$$

Let w^U *be such that:*
$$
w^{\text{U}} = \max\{w^{\text{U}_1}, w^{\text{U}_2}, w^{\text{U}_3}\}.
\tag{Upper-HP}
$$

Algorithm 1: Optimal Upgrade Price

Input: time period t; remaining availability $c_1^{(t)}$, $c_2^{(t)}$; arrival rate λ_1, λ_2.
Output: Optimal upgrade price at time t.
if $c_1^{(t)} + c_2^{(t)} < (\lambda_1 + \lambda_2)(T - t + 1)$ **then**

   Let upgrade plan $v_1^{(t)} = \min \left\{ \max \left\{ \dfrac{c_2^{(t)} - \lambda_2(T-t+1)}{c_1^{(t)} + c_2^{(t)} - \lambda_2(T-t+1)}, 0 \right\}, 1 \right\}$

else

   Let upgrade plan $v_1^{(t)} = \max \left\{ \min \left\{ v_1^{\star}, \dfrac{c_2^{(t)}}{\lambda_1(T-t+1)} - \dfrac{\lambda_2}{\lambda_1} \right\}, 1 - \dfrac{c_1^{(t)}}{\lambda_1(T-t+1)} \right\}$

return $\min \left\{ \max \left\{ p_1(v_1^{(t)}), 0 \right\}, 1 \right\}$.

Then, given any Λ_1 and Λ_2, the optimal solution to Upper-HP is the same as the one to HP. Furthermore, $W^{\mathrm{U}} \geq \mathbb{E}[W^{\mathrm{OPT}}]$, where $W^{\mathrm{U}} = \mathbb{E}[w^{\mathrm{U}}]$.

The proof of Theorem 2 can be found in Appendix EC.1. To demonstrate that Equations (Upper-HP) and (HP) share the same optimal solution, we only need to calculate the optimal solution in three distinct cases: when v_1 is inside $\mathcal{V}_1$, or $\mathcal{V}_2$, or $\mathcal{V}_3$. To establish that w^{U} is the maximum potential revenue achievable by any online algorithm, we find that, in each case, upon fixing the hindsight variables, Λ_i, the optimization problem with respect to v_1 is convex. Then, by Gallego and Van Ryzin (1994), the statement immediately follows.

Theorem 2 describes hybrid programming formulations (Upper-HP) that effectively capture the order of resource depletion. Specifically, $w^{\mathrm{U_1}}$ captures the case where no resource will run out because the number of type 1 arrivals is not enough. $w^{\mathrm{U_2}}$ captures the case where the type 2 resource runs out first, and after that, the supplier cannot give an upgrading option to the remaining type 1 requests. In this case, the remaining type 1 requests can only generate a revenue of r_1 each. $w^{\mathrm{U_3}}$ describes the case where the type 1 resource first runs out, and after that, the supplier can only give free upgrades to the remaining type 1 requests. In that case, the remaining type 1 requests can also only generate a revenue of r_1 each. Finally, Theorem 2 states that the optimal solution to (HP) coincides with the one to (Upper-HP), namely $v_1^{(\mathrm{H})} = \mathrm{argmax}_{v_1} w^{\mathrm{H}} = \mathrm{argmax}_{v_1} w^{\mathrm{U}}$.

Going back to our benchmark revenue $\mathbb{E}[W^{\mathrm{OPT}}]$, by Theorem 2, W^{U} is an upper bound for $\mathbb{E}[W^{\mathrm{OPT}}]$, which implies that $R(T) \leq \mathbb{E}[W^{\mathrm{U}} - W^{\pi}]$, where the expectation is taken on the arrival process. Henceforth, we will use $\mathbb{E}[W^{\mathrm{U}} - W^{\pi}]$ as a stronger performance metric.

3.2. Algorithm Description and Regret Analysis

In this section, we introduce a fast algorithm (Algorithm DynUp-2) for the supplier to make the decision at each time period: either accept or reject the demand, and if they accept, what the upgrade price will be. Formally, we prove the following theorem.

THEOREM 3. *The regret of Algorithm DynUp-2 is bounded by $O(\log T)$.*

At a high level, in Algorithm DynUp-2, we split the time horizon into two segments. The first segment, $[1, \tau]$, starts from the beginning and extends up to the moment τ, defined as the first instance when the type 1 resource is completely depleted, i.e., $\tau = \inf\{t : c_1^{(t)} = 0\}$. The second segment, $[\tau, T]$, covers the remaining portion of the time horizon. In the first segment $[1, \tau]$, requests are accepted on a first-come first-serve basis (i.e., demand is rejected only if there is insufficient availability). The upgrade price is determined by Algorithm 1 in each time step t. In the second segment $[\tau, T]$, as the type 1 resource is depleted, the problem becomes a single-leg revenue management problem: there is one type of resource serving two types of customers with revenue for the supplier r_1 and r_2 respectively. In this case, the last step of DynUp-2 borrows the state-of-the-art algorithm for single-leg revenue management from Vera and Banerjee (2021). Let $W^{\mathrm{U}}(t_1, t_2)$ be

---

Algorithm 2: DynUp-2

Input: time horizon T; initial availability $c_1^{(1)}$, $c_2^{(1)}$; arrival rate λ_1, λ_2.

for $t \in \{1, 2, ..., T\}$ **do**

  Observe demand of type i. If $i = 2$, we accept the request if and only if $c_2^{(t)} \geq 1$.

  if $i = 1$ *and* $c_1^{(t)} \geq 1$ **then**

   Input t, $c_1^{(t)}$, $c_2^{(t)}$, λ_1, and λ_2 to Algorithm 1, and denote the output as $u_1^{(t)}$.

   if $c_2^{(t)} \geq 1$ **then**

    Accept the request and give an upgrade plan with additional price $u_1^{(t)}$.

   else

    Accept the request without an upgrade plan.

  else if $i = 1$ *and* $c_1^{(t)} = 0$ **then**

   Accept the request and give a free upgrade if and only if

   $\left(c_2^{(t)} - \lambda_2(T - t + 1) \right)^+ > \frac{1}{2}\lambda_1(T - t + 1)$

---

the upper bound of the revenue between $[t_1, t_2]$, and $W^\pi(t_1, t_2)$ the total revenue generated by Algorithm DynUp-2 between $[t_1, t_2]$. From the definition of regret:

$$
\begin{aligned}
R(T) &= \mathbb{E}[W^{\mathrm{U}}(1, T) - W^\pi(1, T)] \\
&= \mathbb{E}[W^{\mathrm{U}}(1, \tau - 1) - W^\pi(1, \tau - 1)] + \mathbb{E}[W^{\mathrm{U}}(\tau, T) - W^\pi(\tau, T)]
\end{aligned}
\tag{2}
$$

Then, we bound the two expectations in Equation (2) respectively. First, we suggest a lemma to upper bound $\mathbb{E}[W^{\mathrm{U}}(1, \tau - 1) - W^\pi(1, \tau - 1)]$ as follows:

LEMMA 1. *The regret of the first time segment is upper bounded as:*
$$
\mathbb{E}[W^{\mathrm{U}}(1, \tau - 1) - W^\pi(1, \tau - 1)] \leq \mathbb{E}\left[\sum_{t \in [\tau - 1]} \left(R(v_1^{(\mathrm{H})}) - R(v_1^{(t)}) \right) \right].
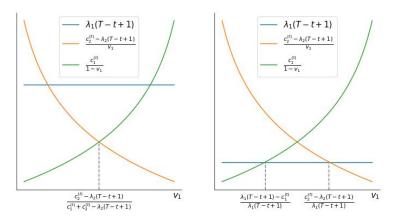$$

The proof can be found in Appendix EC.2. Next, we study the relationship between $v_1^{(\mathrm{H})}$ and $v_1^{(t)}$. Observe that $v_1^{(t)}$ is a stochastic process, which corresponds to the optimal solution to (Upper-HP), taking into account both the remaining availability of each resource, $c_i^{(t)}$, and the anticipated demand for each category, $\lambda_i(T - t + 1)$, $i \in \{1, 2\}$. According to Theorem 2, the optimal solutions for (Upper-HP) and (HP) are equivalent, and our focus shifts primarily to solving (HP). Notice that (HP) is a non-convex optimization problem at each period t, the following lemma shows the closed-form solutions of (HP) when $t = 1$.

LEMMA 2. *If* $c_1^{(1)} + c_2^{(1)} > (\lambda_1 + \lambda_2)T$, *we have that:*

$$
v_1^{(1)} = \max\left\{ \min\left\{ v_1^\star, \frac{c_2^{(1)} - \lambda_2 T}{\lambda_1 T} \right\}, \frac{\lambda_1 T - c_1^{(1)}}{\lambda_1 T} \right\} \quad and \quad v_1^{(\mathrm{H})} = \max\left\{ \min\left\{ v_1^\star, \frac{c_2^{(1)} - \Lambda_2}{\Lambda_1} \right\}, \frac{\Lambda_1 - c_1^{(1)}}{\Lambda_1} \right\}.
$$

If $c_1^{(1)} + c_2^{(1)} \leq (\lambda_1 + \lambda_2)T$, *we have that:*
$$
v_1^{(1)} = \frac{c_2^{(1)} - \lambda_2 T}{c_1^{(1)} + c_2^{(1)} - \lambda_2 T} \quad and \quad v_1^{(\mathrm{H})} = \frac{c_2^{(1)} - \Lambda_2}{c_1^{(1)} + c_2^{(1)} - \Lambda_2}.
$$

Before providing the full proof, we first give some intuition. In the first case where $c_1^{(t)} + c_2^{(t)} < (\lambda_1 + \lambda_2)(T - t + 1)$, the aggregate availability is less than the cumulative estimated expected future arrivals; the left diagram in Fig. 1 shows that $\min\left\{ \lambda_1(T - t + 1), \frac{c_2^{(t)} - \lambda_2(T - t + 1)}{v_1}, \frac{c_1^{(t)}}{1 - v_1} \right\}$ cannot equal $\lambda_1(T - t + 1)$. If $\frac{c_2^{(t)} - \lambda_2(T - t + 1)}{v_1}$ is the minimum, Lemma EC.1 from Appendix EC.1 implies that W^{H}

Figure 1 **Auxiliary image for the proof sketch. Left: the case where** $c_1^{(t)} + c_2^{(t)} < (\lambda_1 + \lambda_2)(T - t + 1)$. **Right: the case where** $c_1^{(t)} + c_2^{(t)} \geq (\lambda_1 + \lambda_2)(T - t + 1)$

is decreasing with respect to v_1. Conversely, if $\frac{c_1^{(t)}}{1-v_1}$ is the minimal value, the same lemma suggests that W^{H} is increasing with respect to v_1. Consequently, at time period t, the optimal solution to (HP) is at the intersection of the functions $\frac{c_2^{(t)}-\lambda_2(T-t+1)}{v_1}$ and $\frac{c_1^{(t)}}{1-v_1}$, yielding $v_1^{(\mathrm{H})} = \frac{c_2^{(t)}-\lambda_2(T-t+1)}{c_1^{(t)}+c_2^{(t)}-\lambda_2(T-t+1)}$.

For the second case where $c_1^{(t)} + c_2^{(t)} \geq (\lambda_1 + \lambda_2)(T - t + 1)$, the right diagram of Fig. 1 shows the relationship between $\lambda_1(T-t+1)$, $\frac{c_2^{(t)}-\lambda_2(T-t+1)}{v_1}$, and $\frac{c_1^{(t)}}{1-v_1}$. By Lemma EC.1, the monotonicity of W^{H} in each function shows that the optimal solution $v_1^{(\mathrm{H})}$ is within

$$\mathcal{I} = \left\{ v_1 : \lambda_1(T-t+1) = \min\left\{ \lambda_1(T-t+1), \frac{c_2^{(t)}-\lambda_2(T-t+1)}{v_1}, \frac{c_1^{(t)}}{1-v_1} \right\} \right\}.$$

In this context, (HP) becomes a convex optimization problem: $\max_{v_1 \in \mathcal{I}} R_1(v_1)$. $R_1(v_1)$ has a unique maximizer $v_1^\star$, so we set $v_1^{(t)}$ as the projection of $v_1^\star$ to $\mathcal{I}$. The formal proof follows.

Proof of Lemma 2 If $c_1^{(1)} + c_2^{(1)} > (\lambda_1 + \lambda_2)T$, Lemma EC.1 implies that $v_1^{(1)} \in \left[\frac{\lambda_1 T - c_1^{(1)}}{\lambda_1 T}, \frac{c_2^{(1)}-\lambda_2 T}{\lambda_1 T} \right]$. Let the unique maximizer of $R_1(v_1)$ be $v_1^\star$. If $v_1^\star \in \left[\frac{\lambda_1 T - c_1^{(1)}}{\lambda_1 T}, \frac{c_2^{(1)}-\lambda_2 T}{\lambda_1 T} \right]$, then

$$v_1^{(1)} = v_1^\star = \operatorname*{argmax}_{v_1 \in \left[\frac{\lambda_1 T - c_1^{(1)}}{\lambda_1 T}, \frac{c_2^{(1)}-\lambda_2 T}{\lambda_1 T} \right]} R_1(v_1).$$

If $v_1^\star > \frac{c_2^{(1)}-\lambda_2 T}{\lambda_1 T}$, by Lemma EC.1, we have $v_1^{(1)} = \frac{c_2^{(1)}-\lambda_2 T}{\lambda_1 T}$. Similarly, if $v_1^\star < \frac{\lambda_1 T - c_1^{(1)}}{\lambda_1 T}$, by Lemma EC.1, we have $v_1^{(1)} = \frac{\lambda_1 T - c_1^{(1)}}{\lambda_1 T}$. By the same approach, we can also get the expression of $v^{(\mathrm{H})}$.

If $c_1^{(1)} + c_2^{(1)} \leq (\lambda_1 + \lambda_2)T$, by Theorem 2, we have in this case, $W^{\mathrm{U}} = \max\{W^{\mathrm{U}_2}, W^{\mathrm{U}_3}\}$. The max can be attained if and only if $W^{\mathrm{U}_2} = W^{\mathrm{U}_3}$, which implies that

$$\frac{c_2 - \lambda_2 T}{v_1^{(1)}} = \frac{c_1}{1 - v_1^{(1)}}.$$

Because the left hand side is monotone decreasing in $v_1^{(1)}$ and right hand side is monotone increasing in $v_1^{(1)}$, there is only one solution, which is

$$v_1^{(1)} = \frac{c_2^{(1)} - \lambda_2 T}{c_1^{(1)} + c_2^{(1)} - \lambda_2 T}.$$

By the same approach, we can also get the expression for $v^{(\mathrm{H})}$. $\square$

Lemma 2 provides the closed form of the decision of Algorithm DynUp-2 in the first period, $v_1^{(1)}$, and the optimal solution to (HP), $v_1^{(\mathrm{H})}$. The following lemma shows their distance and the proof can be found in Appendix EC.2.

LEMMA 3. $\mathbb{E}[v_1^{(1)} - v^{(\mathrm{H})}] = O(1/T)$.

Given that the initial value of the stochastic process $v_1^{(t)}$ is close to $v^{(\mathrm{H})}$, we discuss how $v_1^{(t)}$ evolves. This requires to split the analysis into two cases based on the remaining availability and time.

3.2.1. Case 1: $c_1^{(t)} + c_2^{(t)} > (\lambda_1 + \lambda_2)(T - t + 1)$
In this case, by Lemma 2, we can obtain

$$v_1^{(t)} = \max\left\{\min\left\{v_1^\star, \frac{c_2^{(t)}}{\lambda_1(T-t+1)} - \frac{\lambda_2}{\lambda_1}\right\}, 1 - \frac{c_1^{(t)}}{\lambda_1(T-t+1)}\right\}.$$

We have $v_1^{(1)} = \max\left\{\min\left\{v_1^\star, \frac{c_2^{(1)} - \lambda_2 T}{\lambda_1 T}\right\}, \frac{\lambda_1 T - c_1^{(1)}}{\lambda_1 T}\right\}$. If $v_1^\star$ is in the middle of the interval $\left[\frac{\lambda_1 T - c_1^{(1)}}{\lambda_1 T}, \frac{c_2^{(1)} - \lambda_2 T}{\lambda_1 T}\right]$, then $v_1^{(1)} = v_1^\star$. For each period t, $v_1^{(t)}$ is the projection of $v_1^\star$ into $\left[1 - \frac{c_1^{(t)}}{\lambda_1(T-t+1)}, \frac{c_2^{(t)}}{\lambda_1(T-t+1)} - \frac{\lambda_2}{\lambda_1}\right]$. As the value of the lower and upper bound of the interval is slightly changing for each t, we have most likely, $v_1^\star$ is always in this interval, and $v_1^{(t)} = v_1^\star = v_1^{(1)}$. Therefore, in this case, we almost have zero loss.

However, if $v_1^\star$ is outside the interval $\left[\frac{\lambda_1 T - c_1^{(1)}}{\lambda_1 T}, \frac{c_2^{(1)} - \lambda_2 T}{\lambda_1 T}\right]$, for example, $v_1^\star > \frac{c_2^{(1)} - \lambda_2 T}{\lambda_1 T}$, then $v_1^{(1)} = \frac{c_2^{(1)} - \lambda_2 T}{\lambda_1 T}$. Therefore, in each period t, if the upper bound moves left, we will have $v_1^{(t)} < v_1^{(1)}$, which result in some loss. The next lemma constructs a martingale to characterize the movement of the upper bound of the interval.

LEMMA 4. *Let $\triangle_2(s)$ be the amount of type 2 resources assigned in period s. Suppose that $v_1^{(1)} = \frac{c_2^{(1)} - \lambda_2 T}{\lambda_1 T}$. Construct a stochastic process $\epsilon_t^U = v_1^{(1)} - \left(\frac{c_2^{(t)}}{\lambda_1(T-t+1)} - \frac{\lambda_2}{\lambda_1}\right)$. Then, we have*

$$\epsilon_t^U = \frac{1}{\lambda_1} \sum_{s=1}^{t-1} \frac{\triangle_2(s) - \mathbb{E}[\triangle_2(s)]}{T - s + 1},$$

and ϵ_t^U is a martingale.

Proof of Lemma 4   As $v_1^{(1)} = \frac{c_2^{(1)} - \lambda_2 T}{\lambda_1 T}$, we have $c_2^{(1)} = \lambda_1 T v_1^{(1)} + \lambda_2 T$. If $t = 2$, we have $\frac{c_2^{(1)}}{\lambda_1(T-2+1)} - \frac{\lambda_2}{\lambda_1} = \frac{c_2^{(1)} - \triangle_2(1)}{\lambda_1(T-1)} - \frac{\lambda_2}{\lambda_1}$. Then, we can obtain

$$\epsilon_2^U = \frac{1}{\lambda_1}\left(\frac{c_2^{(1)}}{T} - \frac{c_2^{(1)} - \triangle_2(1)}{T-1}\right) = \frac{1}{\lambda_1} \frac{\triangle_2(1)T - c_2^{(1)}}{T(T-1)}$$

$$= \frac{1}{\lambda_1} \frac{\triangle_2(1)T - (\lambda_1 T v_1^{(1)} + \lambda_2 T))}{T(T-1)} = \frac{1}{\lambda_1} \frac{\triangle_2(1) - \mathbb{E}[\triangle_2(1)]}{T-1},$$

where the last step is because $\triangle_2(1) = 1$ if and only if premium request arrives or basic request arrives and accepted the upgrading plan. Therefore, the expectation of $\triangle_2(1)$ is $\lambda_1 v_1^{(1)} + \lambda_2$.

Next, by math induction, we have

$$\epsilon_t^U = \frac{1}{\lambda_1} \sum_{s=1}^{t-1} \frac{\triangle_2(s) - \mathbb{E}[\triangle_2(s)]}{T - s + 1}.$$

Because at each t expected increment $\mathbb{E}\left[\triangle_2(t) - \mathbb{E}[\triangle_2(t)]\right] = 0$, we have that ϵ_t^U is a martingale. $\square$

Similarly, if $v_1^\star < \frac{c_2^{(1)} - \lambda_2 T}{\lambda_1 T}$, then $v_1^{(1)} = \frac{\lambda_1 T - c_1^{(1)}}{\lambda_1 T}$. The next lemma constructs a martingale to characterize the movement of $\frac{\lambda_1(T-t+1) - c_1^{(t)}}{\lambda_1(T-t+1)}$, and the proof can be found in Appendix EC.2.

LEMMA 5. *Let $\triangle_1(s)$ be the amount of type 1 resources assigned in period s. Suppose that $v_1^{(1)} = \frac{\lambda_1 T - c_1^{(1)}}{\lambda_1 T}$. Construct a stochastic process $\epsilon_t^L = v_1^{(1)} - \frac{\lambda_1(T-t+1) - c_1^{(t)}}{\lambda_1(T-t+1)}$. Then, we have*

$$\epsilon_t^L = \frac{1}{\lambda_1} \sum_{s=1}^{t-1} \frac{\mathbb{E}[\triangle_1(s) - \triangle_1(s)]}{T - s + 1},$$

and ϵ_t^L is a martingale.

Lemmas 4 and 5 show that in this case, $v_1^{(1)} - v_1^{(t)}$ is a martingale with expected increment 0.

3.2.2. Case 2: $c_1^{(t)} + c_2^{(t)} \leq (\lambda_1 + \lambda_2)(T - t + 1)$
In this case, Lemma 2 implies that

$$v_1^{(t)} = 1 - \frac{c_1^{(t)}}{c_1^{(t)} + c_2^{(t)} - \lambda_2(T - t + 1)}.$$

However, different with case 1, it is easy to see that $v_1^{(1)} - v_1^{(t)}$ in this case is **not** a martingale. The next lemma constructs a complicated stochastic process α_t which satisfies two key properties: (i) it is a martingale with expected increment 0, (ii) the difference between α_t and ϵ_t can be bounded by $O(1/T^2)$ for $t \in [1, \tau]$.

LEMMA 6. *Define $\epsilon_t = v_1^{(1)} - v_1^{(t)}$ for any $t \in [T]$. For any constant $\gamma \in (0,1)$, on the time horizon $[1, \gamma T]$, there exists a martingale α_t and a constant ζ, such that $|\alpha_t - \epsilon_t| \leq \zeta/T^2$.*

Proof of Lemma 6 First, when $c_1^{(t)} + c_2^{(t)} = (\lambda_1 + \lambda_2)(T - t + 1)$, we have

$$1 - \frac{c_1^{(t)}}{c_1^{(t)} + c_2^{(t)} - \lambda_2(T - t + 1)} = \frac{c_2^{(t)}}{\lambda_1(T - t + 1)} - \frac{\lambda_2}{\lambda_1} = 1 - \frac{c_1^{(t)}}{\lambda_1(T - t + 1)}.$$

This implies that the phase transition during the decision-making process is smooth and continuous. By Lemmas 4 and 5, we have when $c_1^{(t)} + c_2^{(t)} > (\lambda_1 + \lambda_2)(T - t + 1)$, ϵ_t is a martingale. Therefore, we only need to take $\alpha_t = \epsilon_t$.

Then, we focus on the case where $v_1^{(t)} = 1 - \frac{c_1^{(t)}}{c_1^{(t)} + c_2^{(t)} - \lambda_2(T - t + 1)}$. We proceed by induction. Take $\alpha_1 = 0$. For $t = 2$,

- With probability $1 - \lambda_1 - \lambda_2$, no request arrives, and we have:

$$\begin{aligned}
v_1^{(1)} - v_1^{(2)} &= -\left(\frac{c_1^{(1)}}{c_1^{(1)} + c_2^{(1)} - \lambda_2 T} - \frac{c_1^{(1)}}{c_1^{(1)} + c_2^{(1)} - \lambda_2 T + \lambda_2} \right) \\
&= -c_1^{(1)} \lambda_2 \frac{1}{(c_1^{(1)} + c_2^{(1)} - \lambda_2 T)(c_1^{(1)} + c_2^{(1)} - \lambda_2 T + \lambda_2)} \\
&= -c_1^{(1)} \lambda_2 \frac{1}{(c_1^{(1)} + c_2^{(1)} - \lambda_2 T)^2} + O(1/T^3).
\end{aligned}$$

- With probability $\lambda_1(1 - v_1^{(1)})$, a type 1 request arrives and the upgrade is not accepted. Then:

$$
\begin{aligned}
v_1^{(1)} - v_1^{(2)} &= -\left(\frac{c_1^{(1)}}{c_1^{(1)} + c_2^{(1)} - \lambda_2 T} - \frac{c_1^{(1)} - 1}{c_1^{(1)} - 1 + c_2^{(1)} - \lambda_2 T + \lambda_2} \right) \\
&= -\frac{c_2^{(1)} - \lambda_2 T + \lambda_2 c_1^{(1)}}{(c_1^{(1)} + c_2^{(1)} - \lambda_2 T)(c_1^{(1)} - 1 + c_2^{(1)} - \lambda_2 T + \lambda_2)} \\
&= -c_1^{(1)} \left(\frac{v_1^{(1)}}{1 - v_1^{(1)}} + \lambda_2 \right) \frac{1}{(c_1^{(1)} + c_2^{(1)} - \lambda_2 T)(c_1^{(1)} - 1 + c_2^{(1)} - \lambda_2 T + \lambda_2)} \\
&= -c_1^{(1)} \left(\frac{v_1^{(1)}}{1 - v_1^{(1)}} + \lambda_2 \right) \frac{1}{(c_1^{(1)} + c_2^{(1)} - \lambda_2 T)^2} + O(1/T^3),
\end{aligned}
$$

where the third equality is because

$$
\begin{aligned}
c_2^{(1)} - \lambda_2 T + \lambda_2 c_1^{(1)} &= \left(\frac{c_2^{(1)} - \lambda_2 T}{c_1^{(1)}} + \lambda_2 \right) c_1^{(1)} = \left(\frac{c_2^{(1)} - \lambda_2 T}{c_1^{(1)} + c_2^{(1)} - \lambda_2 T} \frac{c_1^{(1)} + c_2^{(1)} - \lambda_2 T}{c_1^{(1)}} + \lambda_2 \right) c_1^{(1)} \\
&= \left(\frac{v_1^{(1)}}{1 - v_1^{(1)}} + \lambda_2 \right) c_1^{(1)}.
\end{aligned}
$$

- With probability $\lambda_2 + \lambda_1 v_1^{(1)}$, a type 2 request arrives or a type 1 request arrives and accepts the upgrading plan, we have

$$
\begin{aligned}
v_1^{(1)} - v_1^{(2)} &= \frac{c_1^{(1)}}{c_1^{(1)} + c_2^{(1)} - 1 - \lambda_2 T + \lambda_2} - \frac{c_1^{(1)}}{c_1^{(1)} + c_2^{(1)} - \lambda_2 T} \\
&= c_1^{(1)}(1 - \lambda_2) \frac{1}{(c_1^{(1)} + c_2^{(1)} - \lambda_2 T)(c_1^{(1)} + c_2^{(1)} - 1 - \lambda_2 T + \lambda_2)} \\
&= c_1^{(1)}(1 - \lambda_2) \frac{1}{(c_1^{(1)} + c_2^{(1)} - \lambda_2 T)^2} + O(1/T^3).
\end{aligned}
$$

Next, we define α_2 as the value of $v_1^{(1)} - v_1^{(2)}$ without the $O(1/T^3)$ term:

$$
\alpha_2 = \begin{cases}
-c_1^{(1)} \lambda_2 \frac{1}{(c_1^{(1)} + c_2^{(1)} - \lambda_2 T)^2} & \text{with probability } 1 - \lambda_1 - \lambda_2 \\
-c_1^{(1)} \left(\frac{v_1^{(1)}}{1 - v_1^{(1)}} + \lambda_2 \right) \frac{1}{(c_1^{(1)} + c_2^{(1)} - \lambda_2 T)^2} & \text{with probability } \lambda_1(1 - v_1^{(1)}) \\
c_1^{(1)}(1 - \lambda_2) \frac{1}{(c_1^{(1)} + c_2^{(1)} - \lambda_2 T)^2} & \text{with probability } \lambda_2 + \lambda_1 v_1^{(1)}
\end{cases}
$$

Then, we calculate the expectation of α_2:

$$
\mathbb{E}[\alpha_2] = \frac{c_1^{(1)}}{(c_1^{(1)} + c_2^{(1)} - \lambda_2 T)^2} \left(-\lambda_2(1 - \lambda_1 - \lambda_2) - \lambda_1(1 - v_1^{(1)}) + (\lambda_2 + \lambda_1 v_1^{(1)})(1 - \lambda_2) \right) = 0.
$$

Therefore, we define the stochastic process

$$
\alpha_t = \begin{cases}
-c_1^{(t)} \lambda_2 \frac{1}{(c_1^{(t)} + c_2^{(t)} - \lambda_2 (T - t + 1))^2} & \text{with probability } 1 - \lambda_1 - \lambda_2 \\
-c_1^{(t)} \left(\frac{v_1^{(t)}}{1 - v_1^{(t)}} + \lambda_2 \right) \frac{1}{(c_1^{(t)} + c_2^{(t)} - \lambda_2 (T - t + 1))^2} & \text{with probability } \lambda_1(1 - v_1^{(t)}) \\
c_1^{(t)}(1 - \lambda_2) \frac{1}{(c_1^{(t)} + c_2^{(t)} - \lambda_2 (T - t + 1))^2} & \text{with probability } \lambda_2 + \lambda_1 v_1^{(t)}
\end{cases}
$$

By math induction, we have α_t is a martingale with expectation 0. In addition, between α_t and ϵ_t, there is an error term of $O(\frac{1}{(T-t+1)^3})$. On time horizon $[1, \gamma T]$, we can use a union bound to upper bound the gap between α_t and ϵ_t: for any $t \in [1, \gamma T]$,

$$|\alpha_t - \epsilon_t| = \sum_{t=1}^{\gamma T} O\left(\frac{1}{(T-t+1)^3}\right) \leq \frac{\zeta}{T^2}$$

for some constant $\zeta > 0$. $\quad\square$

Lemma 6 shows that we can always find a martingale to approximate the gap between our decision variable $v_1^{(t)}$ and the optimal variable $v_1^{(1)}$ for any $t \in [1, \gamma T]$, where $\gamma \in (0, 1)$. Recall that our goal is to approximate this gap for any $t \in [1, \tau]$. Therefore, we capture the value of the stopping time τ by the following lemma.

LEMMA 7. *For any arbitrarily small constant $\eta > 0$, with probability at least $1 - O(1/T)$,*

$$\left(\frac{c_1^{(1)} + c_2^{(1)} - \lambda_2 T}{\lambda_1 T} - \eta\right) T \leq \tau \leq \left(\frac{c_1^{(1)} + c_2^{(1)} - \lambda_2 T}{\lambda_1 T} + \eta\right) T$$

Proof of Lemma 7 Let $\triangle_1(s)$ be the total number of type 1 resources assigned in period s, and we have $\triangle_1(s) \sim \mathrm{Ber}\left(\lambda_1(1 - v_1^{(s)})\right)$. Then, we have

$$\mathbb{P}\left(\tau \geq \left(\frac{c_1^{(1)} + c_2^{(1)} - \lambda_2 T}{\lambda_1 T} - \eta\right)T\right) = \mathbb{P}\left(\sum_{s=1}^{\left(\frac{c_1^{(1)} + c_2^{(1)} - \lambda_2 T}{\lambda_1 T} - \eta\right)T} \triangle_1(s) < c_1^{(1)}\right)$$

$$= \mathbb{P}\left(\sum_{s=1}^{\left(\frac{c_1^{(1)} + c_2^{(1)} - \lambda_2 T}{\lambda_1 T} - \eta\right)T} \left(\triangle_1(s) - \lambda_1 \frac{c_1^{(1)}}{c_1^{(1)} + c_2^{(1)} - \lambda_2 T}\right) < \eta T\right).$$

By Lemma 6, we have $\mathbb{E}[\triangle_2(s)] = \lambda_1 \frac{c_1^{(1)}}{c_1^{(1)} + c_2^{(1)} - \lambda_2 T} + O(1/T^2)$. Therefore, by Chernoff's inequality:

$$\mathbb{P}\left(\tau \geq \left(\frac{c_1^{(1)} + c_2^{(1)} - \lambda_2 T}{\lambda_1 T} - \eta\right)T\right) < \frac{1}{T}.$$

Similarly, with a symmetric statement, we can also get

$$\mathbb{P}\left(\tau \leq \left(\frac{c_1^{(1)} + c_2^{(1)} - \lambda_2 T}{\lambda_1 T} + \eta\right)T\right) < \frac{1}{T}.$$

$\square$

Lemma 7 shows that $\tau = \gamma T$ for some constant $\gamma \in (0, 1)$. Therefore, by Lemma 6, the error of the constructed martingale α_t can be bounded by $O(\frac{1}{T^2})$.

Combined the results from *Case 1* and *Case 2*, we are ready to upper bound the value of $\mathbb{E}\left[\sum_{t \in [\tau-1]} \left(R(v_1^{(\mathrm{H})}) - R(v_1^{(t)})\right)\right]$.

LEMMA 8.

$$\mathbb{E}\left[\sum_{t \in [\tau-1]} \left(R(v_1^{(\mathrm{H})}) - R(v_1^{(t)})\right)\right] = O(\log T).$$

The proof, detailed in Appendix EC.2, utilizes the Taylor series expansion of $R(v_1^{(t)})$. The underlying intuition draws from Lemmas 4, 5, and 6, which establish the existence of a martingale closely approximating the stochastic process $v_1^{(t)}$. Given that $R(\cdot)$ is bounded, continuous, and concave, Jasin and Kumar (2012), Jasin (2014) suggest that the expected difference $\mathbb{E}\left[\sum_{t\in[\tau-1]}\left(R(v_1^{(\text{H})})-R(v_1^{(t)})\right)\right]$ scales logarithmically with T, denoted as $O(\log T)$.

Finally, we bound $\mathbb{E}[W^{\text{U}}(\tau,T)-W^{\pi}(\tau,T)]$ by the proposition below.

PROPOSITION 1. *The regret for the second time segment is:* $\mathbb{E}[W^{\text{U}}(\tau,T)-W^{\pi}(\tau,T)]=O(1)$.

Proof of Proposition 1 In the horizon $[\tau,T]$, we only have type 2 resources. In this case, when a type 1 request arrives, we can only upgrade them for free. Therefore, accepting a type 1, 2 request can generate a revenue of r_1, r_2 respectively. This turns the problem to be the classical quantity-based single leg revenue management problem under stochastic arrival process. By Vera and Banerjee (2021), the regret is $O(1)$. $\square$

Combining Lemma 8 and Proposition 1 in Equation (2), we can compute the regret.

$$R(T) = O(\log T) + O(1) = O(\log T)$$

4. Algorithm for Multiple Types of Resources

In this section, we focus on scenarios involving n distinct types of resources, where n is a given parameter. Similarly to the steps in the previous section, we first describe the upper bound of the performance of any online mechanism in Section 4.1, and subsequently, we present and analyze our algorithm in Section 4.2.

4.1. Upper Bound on Performance of Any Online Algorithm

In Section 3.1, we introduced the hybrid programming (Upper-HP), whose optimal objective value establishes an upper bound across all algorithms. Nevertheless, the complexity of (Upper-HP) arises from the necessity to discuss the sequence in which resources become exhausted. Considering the presence of n distinct resource types in this context, it is impracticable to exhaustively enumerate every potential sequence of depletion for each resource type. To address this, we define $\Lambda = [\Lambda_1, \Lambda_2, \ldots, \Lambda_n]$ as a vector of random variables, each representing the total number of arrivals for each resource type. Furthermore, $\mathbf{v} = [v_1, v_2, \ldots, v_{n-1}]$ is a vector of decision variables. Given $\Lambda_i, i \in [n]$, and a fixed upgrade probability v_i for each type $i \in [n-1]$, we let the total revenue be $W^{\text{n}}(\Lambda, \mathbf{v})$. Therefore, the hybrid formulation can be articulated as:

$$w^{\text{U}} = \max_{\mathbf{v}} W^{\text{n}}(\Lambda, \mathbf{v}) \qquad \text{(n-type HP)}$$

Since w^{U} is a random variable depending on the value of Λ, we define $W^{\text{U}} = \mathbb{E}[w^{\text{U}}]$ as the expected hindsight total revenue. The hybrid programming formulation (n-type HP) finds the optimal fixed upgrading probabilities for each category, with the objective of maximizing the total revenue under Λ_i arrivals of type i requests. Specifically, for type n resources, which are classified as the most superior resource, we should accept all such requests. This acceptance results in a residual availability of $(c_n - \Lambda_n)^+$ resources of type n, which are then exclusively allocated for potential upgrades of requesters classified under type $n-1$. Given the arrival of Λ_{n-1} requesters of type $n-1$, the optimal upgrading probability, denoted as v_{n-1}, is computed utilizing Equation (Upper-HP). With a little bit abuse of notation, we denote w^{U} as $w^{\text{U}}(c_1, c_2, \Lambda_1, \Lambda_2)$, explicitly indicating its dependence on the availability of resources and number of requesters. Consequently, the optimal decision variable $v_{n-1}^{(\text{H})}$ is determined as $v_{n-1}^{(\text{H})} = \text{argmax}_{v_{n-1}} w^{\text{U}}(c_{n-1}, (c_n - \Lambda_n)^+, \Lambda_{n-1}, 0)$. Subsequently, for type $n-2$ requests, the remaining resources of type $n-1$ are quantified as $\left(c_{n-1} - \frac{\Lambda_{n-1}}{1-v_{n-1}^{(\text{H})}}\right)^+$, leading to the derivation of the optimal solution for v_{n-2}, which is formulated as $v_{n-2}^{(\text{H})} = \text{argmax}_{v_{n-2}} w^{\text{U}}(c_{n-2}, \left(c_{n-1} - \frac{\Lambda_{n-1}}{1-v_{n-1}^{(\text{H})}}\right)^+, \Lambda_{n-2}, 0)$. This methodology is iteratively applied across the spectrum of types, from type $n-1$ to type 1. The subsequent theorem encapsulates the derived upper bound of this analysis where the proof can be found in Appendix EC.3.

Algorithm 3: DynUp-n

Input: time horizon T; initial availability vector $\mathbf{c} = [c_1^{(1)}, c_2^{(1)}, \ldots, c_n^{(1)}]$; arrival rate vector $\lambda = [\lambda_1, \lambda_2, \ldots, \lambda_n]$; optimal solution vector to a defined optimization problem: $\mathbf{v}^{(\text{H})} = [v_1^{(\text{H})}, v_2^{(\text{H})}, \ldots, v_{n-1}^{(\text{H})}, 0]$.

for $i \in \{1, 2, \ldots, n-1\}$ **do**

    Apply DynUp-2 for type i and type $i+1$ requests with input T, $c_i^{(1)}$, $\left(c_{i+1}^{(1)} - \frac{\lambda_{i+1}T}{1-v_{i+1}^{(\text{H})}}\right)^+$, λ_i,

    0.

THEOREM 4. *The optimal solution to* (n-type HP) *is:* $v_{n-1}^{(\text{H})} = argmax_{v_{n-1}} w^{\text{U}}(c_{n-1}, (c_n - \Lambda_n)^+, \Lambda_{n-1}, 0)$; *and for* $i \in \{n-2, n-3, \ldots, 1\}$, $v_i^{(\text{H})} = argmax_{v_i} w^{\text{U}}(c_i, \left(c_{i+1} - \frac{\Lambda_{i+1}}{1-v_{i+1}^{(\text{H})}}\right)^+, \Lambda_i, 0)$. *Moreover,* W^{U} *is an upper bound of the total revenue among all online algorithms, where* $W^{\text{U}} = \mathbb{E}[w^{\text{U}}]$.

4.2. Algorithm Description and Regret Analysis

In this section, we present Algorithm DynUp-n, designed for n resource types. By applying Theorem 4, we can obtain the closed form of optimal solution vector to (n-type HP): $\mathbf{v}^{(\text{H})} = [v_1^{(\text{H})}, v_2^{(\text{H})}, \ldots, v_{n-1}^{(\text{H})}]$. Next, for each type i resource, we protect $\left(c_{i+1}^{(1)} - \frac{\lambda_{i+1}T}{1-v_{i+1}^{(\text{H})}}\right)^+$ for type $i-1$ requesters to upgrade. Thereafter, we implement Algorithm 1 to dynamically decide the upgrade pricing between each adjacent pair of request types. The next theorem summarizes the result.

THEOREM 5. *The regret of Algorithm DynUp-n is bounded by* $O(n \log T)$.

The proof can be found in Appendix EC.3. The total revenue, w^{U}, can be decomposed into the sum of total revenue accrued between each consecutive pair of resources. Theorem 4 shows that the mechanism by which DynUp-n allocates protection to each category i resource for the potential upgrade demands of category $i-1$ is the same as (n-type HP). By Theorem 3, the regret between each adjacent pair of resource types is $O(\log T)$. This leads to an overall regret of $O(n \log T)$.
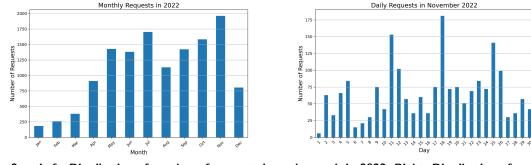
5. Empirical Study for Hotel Upgrading

In this section, we present an empirical study based on real data from the hospitality industry.

Data Description. The dataset contains 13,155 upgrade offers recorded over one year. The left histogram of Figure 2 illustrates the monthly distribution of these offers. The peak in the volume of the offers is observed in November. Consequently, in this section, we use the data corresponding to November as an example for a detailed examination and implementation of the proposed algorithm[3]. The right histogram of Figure 2 describes the daily distribution of requests in November.

The dataset for each reservation request includes the following attributes:
- *price*: initial payment by the customer.
- *from*: requested room category.
- *total*: cost of the next higher room category.
- *to*: next higher room category.
- *upgrade*: upgrade price presented to the customer.
- *decision*: the customer accepts or rejects the upgrade offer.

It is important to note that room prices, even within the same category, are subject to dynamic fluctuation. To effectively measure the relative cost of an upgrade, the dataset employs the upgrade price proportion, calculated as *upgrade proportion = upgrade/(total − price)*. In addition, the existing upgrading strategy does not incorporate information about the remaining capacities of each room category. Instead, it only considers factors related to pricing and demand. Therefore, for

---

[3] We also implement the algorithm for all days in all the months.

Figure 2 **Left: Distribution of number of requests in each month in 2022. Right: Distribution of a number of requests in each day in November 2022.**

every request to room category i, the existing system provides roughly the same *upgrade proportion* for upgrading to the category $i + 1$.

Simulation Setup. The dataset encompasses three room categories with inventories of 60, 30, and 2, respectively, within the examined booking channel. To derive functions $f_1(\cdot)$ and $f_2(\cdot)$ mapping the *upgrade proportion* to the probability of accepting the upgrade plan for type 1 and type 2 requests, a regression analysis is conducted. Based on data visualization, we propose a parametric exponential decay family of functions, $\mathcal{F}$, consisting of all functions of the form $f(v) = e^{-av^b}$. Subsequent regression analysis identifies the functions with the minimal mean square error, resulting in $f_1(v_1) = e^{-4.4853v_1^{0.9889}}$ and $f_2(v_2) = e^{-2.33v_2}$, which best fit the observed data.

We assume that each request corresponds to a single-night occupancy, which confines the decision-making process to a daily basis. Each day, we randomly generate 100 permutations of the arrival order to represent arriving instances. The benchmark algorithm employs the upgrade price found in the data and follows a first-come, first-served protocol, continuing until either all requests have arrived or the inventory is full. To implement Algorithm DynUp-n, we first approximate the arrival rates, denoted as λ_i, $i \in [3]$. We set a horizon T to simulate daily reservations maintained by hotels and add random noise to λ, accounting for cancellation in guest attendance or arrivals without reservations: $\lambda_i \sim \mathrm{Unif}(\Lambda_i - \sqrt{\Lambda_i}, \Lambda_i + \sqrt{\Lambda_i})/T$. Subsequently, Algorithm DynUp-n is applied to the 100 samples and we compare its performance to the benchmark using the same sample instances.

Results. The left diagram of Figure 3 provides a comparative evaluation of daily revenues generated during November by both the benchmark and Algorithm DynUp-n. Despite noticeable fluctuations in daily demand, DynUp-n consistently outperforms the benchmark. Particularly on days characterized by higher demand, DynUp-n achieves a significant increase in total revenue, exceeding 26%. The right panel of Figure 3 and Table 1 illustrate the monthly revenue trends throughout the year, with improvements of over 15% highlighted in red. Over the entire period, the aggregate revenue generated by Algorithm DynUp-n exceeded the benchmark by approximately 17%.

Table 1   Monthly Summary in 2022

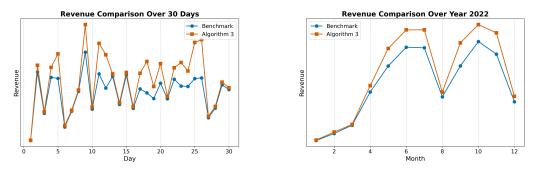| Month	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Total
Improvement	4.49%	7.84%	2.98%	10.61%	20.56%	16.75%	17.24%	8.55%	27.22%	15.20%	26.82%	10.79%	17.07%

Figure 3 **Left: Daily total revenue of Benchmark and Algorithm DynUp-n in November 2022. Right: Monthly total revenue of Benchmark and Algorithm DynUp-n in 2022.**

6. Conclusion

In this paper, we studied upgrading mechanisms for a fee for general resource allocation problems. Our proposed mechanism (Algorithm DynUp-n) achieves regret $O(\log T)$ against a "hybrid programming" benchmark that we analyze. Algorithm DynUp-n stands out not only for its theoretical regret bound but also for its performance in real-world hotel management scenarios.

There are a lot of avenues for future research. For instance, accurately estimating the functions $f_i(\cdot)$ presents a challenge in practical settings. This raises the question of whether effective blind online upgrading can be implemented without prior knowledge of $f_i(\cdot)$. Specifically, it necessitates developing strategies for upgrade pricing to make online learning of $f_i(\cdot)$ possible without incurring significant losses during the learning phase. Additionally, exploring other factors, such as the impact of customer loyalty on upgrading costs, is essential. Determining optimal strategies for discounting upgrade fees to reward loyalty requires careful consideration. Moreover, ensuring fairness in upgrading costs poses a significant challenge. Given that our dynamic upgrading mechanism assigns different fees to each request, incorporating fairness constraints necessitates adjustments to the algorithm's structure. Addressing these questions will enhance the efficacy of the online upgrading model and mechanism.

References

Goker Aydin and Serhan Ziya. 2008. Pricing promotional products under upselling. *Manufacturing & Service Operations Management* 10, 3 (2008), 360–376.

Dimitris Bertsimas and Sanne De Boer. 2005. Simulation-based booking limits for airline revenue management. *Operations Research* 53, 1 (2005), 90–106.

Pornpawee Bumpensanti and He Wang. 2020. A re-solving heuristic with uniformly bounded loss for network revenue management. *Management Science* 66, 7 (2020), 2993–3009.

Zheng Cui, Daniel Zhuoyu Long, and Lijian Lu. 2023. Dynamic Resource Allocation with Proactive Upgrades: The Value and Design of Fulfillment Flexibility. *HKUST Business School Research Paper* 2023-101 (2023).

Guillermo Gallego and Catalina Stefanescu. 2009. Upgrades, upsells and pricing in revenue management. *Available at SSRN 1334341* (2009).

Guillermo Gallego and Garrett Van Ryzin. 1994. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management science* 40, 8 (1994), 999–1020.

Guillermo Gallego and Garrett Van Ryzin. 1997. A multiproduct dynamic pricing problem and its applications to network yield management. *Operations research* 45, 1 (1997), 24–41.

Stanislav Ivanov. 2014. *Hotel revenue management: From theory to practice.* Zangador.

Stanislav Ivanov and Vladimir Zhechev. 2012. Hotel revenue management–a critical literature review. *Tourism: an international interdisciplinary journal* 60, 2 (2012), 175–197.

Stefanus Jasin. 2014. Reoptimization and self-adjusting price control for network revenue management. *Operations Research* 62, 5 (2014), 1168–1178.

Stefanus Jasin and Sunil Kumar. 2012. A re-solving heuristic with bounded revenue loss for network revenue management with customer choice. *Mathematics of Operations Research* 37, 2 (2012), 313–345.

David McCaffrey and Darius Walczak. 2016. Optimal dynamic upgrading in revenue management. *Production and Operations Management* 25, 11 (2016), 1855–1865.

Martin I Reiman and Qiong Wang. 2008. An asymptotically optimal policy for a quantity-based network revenue management problem. *Mathematics of Operations Research* 33, 2 (2008), 257–282.

Robert A Shumsky and Fuqiang Zhang. 2009. Dynamic capacity management with substitution. *Operations research* 57, 3 (2009), 671–684.

Rui Sun, Xinshang Wang, and Zijie Zhou. 2020. Near-optimal primal-dual algorithms for quantity-based network revenue management. *arXiv preprint arXiv:2011.06327* (2020).

CNN Travel. 2020. Nearly 80% of hotel rooms in the US are empty, according to new data. *https://www.cnn.com/2020/04/08/us/hotel-rooms-industry-coronavirus-trnd/index.html* (2020).

Alberto Vera and Siddhartha Banerjee. 2021. The bayesian prophet: A low-regret framework for online decision making. *Management Science* 67, 3 (2021), 1368–1391.

Yining Wang and He Wang. 2022. Constant regret resolving heuristics for price-based revenue management. *Operations Research* 70, 6 (2022), 3538–3557.

Yueshan Yu, Xin Chen, and Fuqiang Zhang. 2015. Dynamic capacity management with general upgrading. *Operations Research* 63, 6 (2015), 1372–1389.

Dan Zhang and William L Cooper. 2009. Pricing substitutable flights in airline revenue management. *European Journal of Operational Research* 197, 3 (2009), 848–861.

Online Appendix

EC.1. Supplemental Materials for Section 3.1

LEMMA EC.1. *Fix any value of $\Lambda_1, \Lambda_2 \geq 0$, $\Lambda_2 r_2 + \frac{c_2 - \Lambda_2}{v_1} R_1(v_1) + r_1 \min\left\{\Lambda_1, c_1 + c_2 - \Lambda_2\right\}$ is a monotone decreasing function in $v_1 \in [0,1]$. $\Lambda_2 r_2 + \frac{c_1}{1-v_1} R_1(v_1) + r_1 \min\left\{\Lambda_1, c_1 + c_2 - \Lambda_2\right\}$ is a monotone increasing function in $v_1 \in [0,1]$.*

Proof of Lemma EC.1 First, to check the monotonicity of $\Lambda_2 r_2 + \frac{c_2 - \Lambda_2}{v_1} R_1(v_1) + r_1 \min\left\{\Lambda_1, c_1 + c_2 - \Lambda_2\right\}$, we only need to show that $\frac{R_1(v_1)}{v_1}$ is decreasing. This is obvious because $\frac{R_1(v_1)}{v_1} = p(v_1)$, which is defined as a monotone decreasing function.

Second, to check the monotonicity of $\Lambda_2 r_2 + \frac{c_1}{1-v_1} R_1(v_1) + r_1 \min\left\{\Lambda_1, c_1 + c_2 - \Lambda_2\right\}$, we only need to show that $\frac{R_1(v_1)}{v_1}$ is increasing in v_1. As $R(\cdot)$ is concave, we have $(v_1 p(v_1))'' < 0$, which implies that $p''(v_1) < \frac{-2p'(v_1)}{v_1}$. Then, we have

$$\left(\frac{R(v_1)}{1-v_1}\right)' = \frac{p(v_1) + v_1(1-v_1)p'(v_1)}{(1-v_1)^2}.$$

Observe that

$$\left(p(v_1) + v_1(1-v_1)p'(v_1)\right)' = 2(1-v_1)p'(v_1) + p''(v_1)(v_1 - v_1^2) < 2(1-v_1)p'(v_1) + \frac{-2p'(v_1)}{v_1}(v_1 - v_1^2) < 0,$$

which implies that $p(v_1) + v_1(1-v_1)p'(v_1)$ is a decreasing function. Therefore,

$$\begin{aligned}\left(\frac{R(v_1)}{1-v_1}\right)' &= \frac{p(v_1) + v_1(1-v_1)p'(v_1)}{(1-v_1)^2} \\ &\geq \frac{p(1) + 1 \cdot (1-1)p'(1)}{(1-v_1)^2} = 0,\end{aligned}$$

and we can obtain that $\frac{R(v_1)}{1-v_1}$ is increasing. □

EC.1.1. Proof of Theorem 2

Proof of Theorem 2 We split the proof into two parts, where in part 1, we show that the optimal solution to Upper-HP is the same as the one to HP, and in part 2, we prove that $W^{\mathrm{U}} \geq \mathbb{E}[W^{\mathrm{OPT}}]$.

Part 1

Let $v_1^{(\mathrm{H})} = \mathrm{argmax}_{v_1} \Lambda_2 r_2 + \min\left\{\Lambda_1, \frac{c_2 - \Lambda_2}{v_1}, \frac{c_1}{1-v_1}\right\}\left(R_1(v_1) + r_1\right)$. Then, we show that $v_1^{(\mathrm{H})}$ is also the optimal solution to w^{U}.

Define $v_1^\star(\mathrm{U}_1) = \mathrm{argmax}_{v_1 \in \mathcal{V}_1} w^{\mathrm{U}_1}$, $v_1^\star(\mathrm{U}_2) = \mathrm{argmax}_{v_1 \in \mathcal{V}_2} w^{\mathrm{U}_2}$, and $v_1^\star(\mathrm{U}_3) = \mathrm{argmax}_{v_1 \in \mathcal{V}_3} w^{\mathrm{U}_3}$. By Lemma EC.1, we have $v_1^\star(\mathrm{U}_2) = \inf \mathcal{V}_2$ and $v_1^\star(\mathrm{U}_3) = \sup \mathcal{V}_3$. Then, we split the discussion into two cases.

Case 1: $\mathcal{V}_1 = \emptyset$. In this case, as $\frac{c_2 - \Lambda_2}{v_1}$ is decreasing in v_1 and $\frac{c_1}{1-v_1}$ is increasing in v_1, these two functions can have only one intersection. Due to $\mathcal{V}_1 = \emptyset$, we have $\inf \mathcal{V}_2 = \sup \mathcal{V}_3$, which implies that $v_1^\star(\mathrm{U}_2) = v_1^\star(\mathrm{U}_3)$. By definition, we have $w^{\mathrm{U}_2} = w^{\mathrm{U}_3} = w^{\mathrm{U}}$.

Moreover, by Lemma EC.1, if $\mathcal{V}_1 = \emptyset$, we have $v_1^{(\mathrm{H})} = \inf \mathcal{V}_2 = \sup \mathcal{V}_3 = v_1^\star(\mathrm{U}_2) = v_1^\star(\mathrm{U}_3)$. Therefore, both w^{U} and W^{H} share the same optimal solution in this case.

Case 2: $\mathcal{V}_1 \neq \emptyset$. In this case, as Λ_1 is a constant function, it can have only one intersection with $\frac{c_2-\Lambda_2}{v_1}$ and one intersection with $\frac{c_1}{1-v_1}$. Therefore, we have $v_1^\star(\mathrm{U}_2) = \inf \mathcal{V}_2 = \sup \mathcal{V}_1$ and $v_1^\star(\mathrm{U}_3) = \sup \mathcal{V}_3$. This implies that

$$w^{\mathrm{U}_2} = \Lambda_2 r_2 + \Lambda_1 R_1(v_1^\star(\mathrm{U}_2)) + r_1 \min\left\{\Lambda_1, c_1 + c_2 - \Lambda_2\right\},$$

and

$$w^{\mathrm{U}_3} = \Lambda_2 r_2 + \Lambda_1 R_1(v_1^\star(\mathrm{U}_3)) + r_1 \min\left\{\Lambda_1, c_1 + c_2 - \Lambda_2\right\}.$$

It is easy to see that

$$w^{\mathrm{U}_1} = \max_{v_1 \in \mathcal{V}_1} \Lambda_2 r_2 + \Lambda_1 \Big(R_1(v_1) + r_1\Big) \geq w^{\mathrm{U}_2},$$

and

$$w^{\mathrm{U}_1} = \max_{v_1 \in \mathcal{V}_1} \Lambda_2 r_2 + \Lambda_1 \Big(R_1(v_1) + r_1\Big) \geq w^{\mathrm{U}_3}.$$

Therefore, we have $w^{\mathrm{U}} = w^{\mathrm{U}_1}$, and its optimal solution is $v_1^\star(\mathrm{U}_1)$.

Moreover, by Lemma EC.1, if $\mathcal{V}_1 \neq \emptyset$, we have $\Lambda_1 = \min\left\{\Lambda_1, \frac{c_2-\Lambda_2}{v_1^{(\mathrm{H})}}, \frac{c_1}{1-v_1^{(\mathrm{H})}}\right\}$, and this implies that $v_1^{(\mathrm{H})} \in \mathcal{V}_1$. Therefore, $v_1^{(\mathrm{H})} = v_1^\star(\mathrm{U}_1)$, and both w^{U} and W^{H} share the same optimal solution in this case.

Part 2

Let $v_1^{(\mathrm{H})}$ be the optimal solution to (HP). From part 1, $v_1^{(\mathrm{H})}$ is also the optimal solution to (Upper-HP). Then, we split the proof into three cases:

Case 1: $\Lambda_1 = \min\left\{\Lambda_1, \frac{c_2-\Lambda_2}{v_1^{(\mathrm{H})}}, \frac{c_1}{1-v_1^{(\mathrm{H})}}\right\}$. In this case, by definition, $w^{\mathrm{U}} = W^{\mathrm{H}}$. In expectation, no resource is depleted during the time horizon $[0, T]$. As our decision process is Markovian and the objective is concave, by a well-known result (Gallego and Van Ryzin (1994) Theorem II), we have $w^{\mathrm{U}} = W^{\mathrm{H}} \geq W^{\mathrm{OPT}}$.

Case 2: $\frac{c_2-\Lambda_2}{v_1^{(\mathrm{H})}} = \min\left\{\Lambda_1, \frac{c_2-\Lambda_2}{v_1^{(\mathrm{H})}}, \frac{c_1}{1-v_1^{(\mathrm{H})}}\right\}$. In this case, the premium resource (type 2) will be depleted at certain point during the time horizon $[0, T]$. Again, by Gallego and Van Ryzin (1994) Theorem II, if the decision process stops when the premium resource is depleted, $\max_{v_1 \in \mathcal{V}_2} \Lambda_2 r_2 + \frac{c_2-\Lambda_2}{v_1}(R_1(v_1) + r_1)$ serves an upper bound to the optimal policy. However, for a given $v_1 \in \mathcal{V}_2$, there is $c_1 - \frac{c_2-\Lambda_2}{v_1}(1 - v_1)$ basic resource (type 1) remaining unused. Also, there is $\Lambda_1 - \frac{c_2-\Lambda_2}{v_1}$ type 1 demand units who will arrive in the remaining time horizon in expectation. Therefore, the upper bound formulation is:

$$\max_{v_1 \in \mathcal{V}_2} \Lambda_2 r_2 + \frac{c_2 - \Lambda_2}{v_1}(R_1(v_1) + r_1) + r_1 \min\left\{\Lambda_1 - \frac{c_2 - \Lambda_2}{v_1}, c_1 - \frac{c_2 - \Lambda_2}{v_1}(1 - v_1)\right\}$$

$$= \max_{v_1 \in \mathcal{V}_2} \Lambda_2 r_2 + \frac{c_2 - \Lambda_2}{v_1} R_1(v_1) + r_1 \min\left\{\Lambda_1, c_1 + c_2 - \Lambda_2\right\}$$

$$= w^{\mathrm{U}_2}$$

Because $w^{\mathrm{U}} \geq w^{\mathrm{U}_2}$, we have $w^{\mathrm{U}} \geq W^{\mathrm{OPT}}$.

Case 3: $\frac{c_1}{1-v_1^{(\mathrm{H})}} = \min\left\{\Lambda_1, \frac{c_2-\Lambda_2}{v_1^{(\mathrm{H})}}, \frac{c_1}{1-v_1^{(\mathrm{H})}}\right\}$. In this case, the basic resource (type 1) will be depleted at certain point during the time horizon $[0, T]$. Again, by Gallego and Van Ryzin (1994) Theorem II, if the decision process stops when the basic resource is depleted, $\max_{v_1 \in \mathcal{V}_3} \Lambda_2 r_2 + \frac{c_1}{1-v_1}(R_1(v_1) + r_1)$ serves an upper bound to the optimal policy. However, for a given $v_1 \in \mathcal{V}_3$, there is $c_2 - \Lambda_2 - \frac{c_1}{1-v_1}v_1$ premium resource (type 2) remaining unused. Also, there is $\Lambda_1 - \frac{c_1}{1-v_1}$ type 1 demand units who will arrive in the remaining time horizon in expectation. These type 1 demand units can be upgraded

to the premium resource even if the basic resource is depleted. Thus, the upper bound formulation is:

$$
\begin{aligned}
& \max_{v_1 \in \mathcal{V}_3} \Lambda_2 r_2 + \frac{c_1}{1 - v_1}\left(R_1(v_1) + r_1\right) + r_1 \min\left\{\Lambda_1 - \frac{c_1}{1 - v_1}, c_2 - \Lambda_2 - \frac{c_1}{1 - v_1}v_1\right\} \\
&= \max_{v_1 \in \mathcal{V}_3} \Lambda_2 r_2 + \frac{c_1}{1 - v_1}R_1(v_1) + r_1 \min\left\{\Lambda_1, c_1 + c_2 - \Lambda_2\right\} \\
&= w^{\mathrm{U}_3}
\end{aligned}
$$

Because $w^{\mathrm{U}} \geq w^{\mathrm{U}_3}$, we have $w^{\mathrm{U}} \geq W^{\mathrm{OPT}}$. This also implies that $W^{\mathrm{U}} \geq \mathbb{E}[W^{\mathrm{OPT}}]$. $\quad\square$

EC.2. Supplementary Materials for Section 3.2

Proof of Lemma 1 Denote $\mathcal{I}_1$ as the time index that a type 1 customer arrives before τ.

$$
\begin{aligned}
& \mathbb{E}[W^{\mathrm{U}}(1, \tau - 1) - W^{\pi}(1, \tau - 1)] \\
&= \mathbb{E}\left[\Lambda_2(1, \tau - 1)r_2 + \Lambda_1(1, \tau - 1)\left(R_1(v_1^{(\mathrm{H})}) + r_1\right) - \Lambda_2(1, \tau - 1)r_2 + \sum_{t \in \mathcal{I}_1}\left(R(v_1^{(t)}) + r_1\right)\right] \\
&\leq \mathbb{E}\left[\sum_{t=1}^{\tau - 1}\left(R(v_1^{(\mathrm{H})}) - R(v_1^{(t)})\right)\right],
\end{aligned}
$$

where the first equality is because by Wald's equality, at stopping time τ, the remaining capacity of type 2 resource is $c_2^{(\tau)} \approx \lambda_2(T - \tau + 1) > 0$, which implies that before τ, no resource is depleted.

Proof of Lemma 3 If both $v_1^{(1)}$ and $v^{(\mathrm{H})}$ are $v_1^\star$, we have $v_1^{(1)} - v^{(\mathrm{H})} = 0$. If $v_1^{(1)} = \frac{c_2^{(1)} - \lambda_2 T}{\lambda_1 T}$ and $v^{(\mathrm{H})} = \frac{c_2^{(1)} - \Lambda_2}{\Lambda_1}$, since Λ_1 and Λ_2 are independent, we have

$$
\mathbb{E}[v_1^{(1)} - v^{(\mathrm{H})}] = \mathbb{E}\left[\frac{c_2^{(1)} - \lambda_2 T}{\lambda_1 T} - \frac{c_2^{(1)} - \Lambda_2}{\Lambda_1}\right] = (c_2^{(1)} - \lambda_2 T)\left(\frac{1}{\lambda_1 T} - \mathbb{E}\left[\frac{1}{\Lambda_1}\right]\right).
$$

As $\Lambda_1 \sim \mathrm{Bin}(T, \lambda_1)$, we have

$$
\mathbb{E}\left[\frac{1}{\Lambda_1}\right] \approx \mathbb{E}\left[\frac{1}{1 + \Lambda_1}\right].
$$

Use the fact that $\mathbb{E}\left[\frac{1}{a+X}\right] = \int_0^1 t^{a-1}P_X(t)dt$, where $P_X(t)$ is the probability generating function for X, we can obtain

$$
\begin{aligned}
\mathbb{E}\left[\frac{1}{1 + \Lambda_1}\right] &= \int_0^1 t^0 P_{\Lambda_1}(t)dt \\
&= \int_0^1 (1 - \lambda_1 + \lambda_1 t)^T dt \\
&= \frac{1 - (1 - \lambda_1)^{T+1}}{(T+1)\lambda_1}
\end{aligned}
$$

Therefore, we have $\frac{1}{\lambda_1 T} - \mathbb{E}\left[\frac{1}{\Lambda_1}\right] \approx \frac{1}{\lambda_1 T} - \frac{1 - (1 - \lambda_1)^{T+1}}{(T+1)\lambda_1} = O(T^{-2})$. Similarly, if $v_1^{(1)} = \frac{\lambda_1 T - c_1^{(1)}}{\lambda_1 T}$ and $v_1^{(\mathrm{H})} = \frac{\Lambda_1 - c_1^{(1)}}{\Lambda_1}$, we can use the same method to derive that $\mathbb{E}[v_1^{(1)} - v^{(\mathrm{H})}] = O(T^{-2})$.

Finally, if $v_1^{(1)} = \frac{c_2^{(1)} - \lambda_2 T}{c_1^{(1)} + c_2^{(1)} - \lambda_2 T}$ and $v_1^{(\text{H})} = \frac{c_2^{(1)} - \Lambda_2}{c_1^{(1)} + c_2^{(1)} - \Lambda_2}$, we have

$$\mathbb{E}[v_1^{(1)} - v^{(\text{H})}] = \mathbb{E}\left[1 - \frac{c_1^{(1)}}{c_1^{(1)} + c_2^{(1)} - \lambda_2 T} - \left(1 - \frac{c_1^{(1)}}{c_1^{(1)} + c_2^{(1)} - \lambda_2 T}\right)\right]$$

$$= c_1^{(1)} \mathbb{E}\left[\frac{1}{c_1^{(1)} + c_2^{(1)} - \Lambda_2} - \frac{1}{c_1^{(1)} + c_2^{(1)} - \lambda_2 T}\right].$$

As $\Lambda_2 \sim \text{Bin}(T, \lambda_2)$, by Hoeffding's inequality, with probability at least $1 - 1/T$, we have $\Lambda_2 \in [\lambda_2 T - \sqrt{T} \log T, \lambda_2 T + \sqrt{T} \log T]$. Therefore, we have

$$\mathbb{E}\left[\frac{1}{c_1^{(1)} + c_2^{(1)} - \Lambda_2} - \frac{1}{c_1^{(1)} + c_2^{(1)} - \lambda_2 T}\right] = (1 - \frac{1}{T})O(T^{-\frac{3}{2}} \log T) + \frac{1}{T} = O(\frac{1}{T}).$$

Proof of Lemma 5 As $v_1^{(1)} = \frac{\lambda_1 T - c_1^{(1)}}{\lambda_1 T}$, we have $c_1^{(1)} = \lambda_1 T(1 - v_1^{(1)})$. If $t = 2$, we have $\frac{\lambda_1(T-2+1) - c_1^{(2)}}{\lambda_1(T-2+1)} = 1 - \frac{c_1^{(1)} - \triangle_1(1)}{\lambda_1(T-1)}$. Then, we can obtain

$$\epsilon_2^L = \frac{1}{\lambda_1}\left(\frac{c_1^{(1)} - \triangle_1(1)}{T-1} - \frac{c_1^{(1)}}{T}\right) = \frac{1}{\lambda_1} \frac{\lambda_1 T(1 - v_1^{(1)}) - \triangle_1(1)T}{T(T-1)}$$

$$= \frac{1}{\lambda_1} \frac{\mathbb{E}[\triangle_1(1)] - \triangle_1(1)}{T-1},$$

where the last step is because $\triangle_1(1) = 1$ if and only if basic request arrives and rejected the upgrading plan. Therefore, the expectation of $\triangle_1(1)$ is $\lambda_1(1 - v_1^{(1)})$.

Next, by math induction, we have

$$\epsilon_t^L = \frac{1}{\lambda_1} \sum_{s=1}^{t-1} \frac{\mathbb{E}[\triangle_1(s)] - \triangle_1(s)}{T - s + 1}.$$

Because each time the expected increment $\mathbb{E}\left[\mathbb{E}[\triangle_1(t)] - \triangle_1(t)\right] = 0$, we immediately have ϵ_t^L is a martingale. $\square$

Proof of Lemma 8 Take $\zeta = \frac{1}{T^3}$, and we perturb the term $R(v_1^{(\text{H})})$ by ζ units to make R differentiable at $v_1^{(\text{H})} - \zeta$. In Lemma 3, we upper bound the expected difference of $\beta = v_1^{(\text{H})} - v_1^{(1)}$ by $O(\frac{1}{T})$. Moreover, in Lemma 6, we defined a martingale α_t which approximately captures $\epsilon_t = v_1^{(1)} - v_1^{(t)}$, and $|\alpha_t - \epsilon_t| = O(\frac{1}{T^2})$. By Taylor's expansion, we have

$$\sum_{t=1}^{\tau-1} R(v_1^{(t)}) \geq (\tau - 1)R(v_1^{(\text{H})} - \zeta) + R'(v_1^{(\text{H})} - \zeta)\sum_{t=2}^{\tau-1}(\beta + \epsilon_t - \zeta) + \frac{1}{2}R''(v_1^{(\text{H})} - \zeta)\sum_{t=2}^{\tau-1}(\beta + \epsilon_t - \zeta)^2$$

$$+ \sum_{t=1}^{\tau-1}(R(v_1^{(t)}) - \mathbb{E}[R(v_1^{(t)})|\mathcal{F}_t]).$$

Observe that $\sum_{s=1}^{t}(R(v_1^{(s)}) - \mathbb{E}[R(v_1^{(s)})|\mathcal{F}_s])$ is a martingale. By stopping time theorem, we have

$$\mathbb{E}\left[\sum_{t=1}^{\tau-1}(R(v_1^{(t)}) - \mathbb{E}[R(v_1^{(t)})|\mathcal{F}_t])\right] = 0.$$

By continuity of R, we have

$$\mathbb{E}\left[(\tau-1)R(v_1^{(\text{H})}) - (\tau-1)R(v_1^{(\text{H})}-\zeta)\right] = O(\frac{1}{T^2}).$$

As ϵ_t is close to a martingale α_t, and $|\alpha_t - \epsilon_t| = O(\frac{1}{T^2})$, we have $\mathbb{E}[\sum_{t=2}^{\tau-1}\epsilon_t] = O(\frac{1}{T})$ and by Lemma 3, $\mathbb{E}[\beta] = O(\frac{1}{T})$, we can obtain

$$-\mathbb{E}\left[R'(v_1^{(\text{H})}-\zeta)\sum_{t=2}^{\tau-1}(\beta+\epsilon_t-\zeta)\right] = O(1).$$

Finally, we also have

$$
\begin{aligned}
-\mathbb{E}[\frac{1}{2}R''(v_1^{(\text{H})}-\zeta)\sum_{t=2}^{\tau-1}(\beta+\epsilon_t-\zeta)^2] &\leq O(1) + \frac{1}{2}R''(v_1^{(\text{H})}-\zeta)\mathbb{E}\left[\sum_{t=2}^{\tau-1}(\epsilon_t)^2\right]\\
&= O(1) + \frac{1}{2}R''(v_1^{(\text{H})}-\zeta)\mathbb{E}[\sum_{t=2}^{\tau-1}\sum_{1\leq s,v\leq t}\frac{1}{(T-s)(T-v)}\epsilon_s\epsilon_v]\\
&= O(1) + \frac{1}{2}R''(v_1^{(\text{H})}-\zeta)\mathbb{E}[\sum_{t=2}^{\tau-1}\sum_{s=1}^{t-1}\frac{1}{(T-s)^2}(\epsilon_s)^2]\\
&= O(\log T).
\end{aligned}
$$

$\square$

EC.3. Supplementary Materials for Section 4

Proof of Theorem 4 Firstly, the social cost of type $i+1$ dominates the maximum upgrade social cost of type i, namely $r_{i+1} \geq r_i + u_i$, we can decompose the hindsight formulation with multiple types to n hindsight formulation with 1 type of demand and 2 types of resources. Therefore, by Theorem 2, the hybrid programming W^{U} serves as an upper bound.

Secondly, to show the closed form solution of (n-type HP), let us focus initially on type n requests. As there are Λ_n such requests and considering that $r_n > r_{n-1} > \ldots > r_1$, the optimal approach entails accepting $\min\{\Lambda_n, c_n\}$ type n requests. This results in a surplus of $(c_n - \Lambda_n)^+$ units of type n resources. These remaining resources are exclusively allocated for fulfilling upgraded type $n-1$ demands. Given that each type $n-1$ demand utilizes either a unit of type $n-1$ or type n resource, and in both cases, the total revenue surpasses that of other resources, it is necessary to accommodate as many type $n-1$ demands as feasible. Thus, by Theorem 2, the optimal upgrade probability of type $n-1$ demands is $v_{n-1}^{(\text{H})} = \text{argmax}_{v_{n-1}} W^{\text{U}}(c_{n-1}, (c_n - \Lambda_n)^+, \Lambda_{n-1}, 0)$.

Next, by setting the upgrade probability of type $n-1$ demand as $v_{n-1}^{(\text{H})}$, the surplus of type $n-1$ resources is $\left(c_{n-1} - \frac{\Lambda_{n-1}}{1-v_{n-1}^{(\text{H})}}\right)^+$, and by similar reasoning, the optimal upgrade probability of type $n-2$ demand is $v_{n-2}^{(\text{H})} = \text{argmax}_{v_{n-2}} w^{\text{U}}(c_{n-2}, \left(c_{n-1} - \frac{\Lambda_{n-1}}{1-v_{n-1}^{(\text{H})}}\right)^+, \Lambda_{n-2}, 0)$. By math induction, we have for $i \in \{n-3, n-4, \ldots, 1\}$, $v_i^{(\text{H})} = \text{argmax}_{v_i} w^{\text{U}}(c_i, \left(c_{i+1} - \frac{\Lambda_{i+1}}{1-v_{i+1}^{(\text{H})}}\right)^+, \Lambda_i, 0)$. $\square$

Proof of Theorem 5 Recall that in Theorem 2, we have defined the upper bound of total revenue W^{U}, which directly depends on c_1, c_2, Λ_1, and Λ_2. In addition, it maximizes over all possible v_1. Therefore, in this proof, we denote the total revenue of the deterministic programming of two types of resource as $w^{\text{U}}(c_1, c_2, \Lambda_1, \Lambda_2; v_1^{(\text{H})})$. That is, if $\Lambda_1 = \min\{\Lambda_1, \frac{c_2-\Lambda_2}{v_1^{(\text{H})}}, \frac{c_1}{1-v_1^{(\text{H})}}\}$,

$$w^{\text{U}}(c_1, c_2, \Lambda_1, \Lambda_2; v_1^{(\text{H})}) = \Lambda_2 r_2 + \Lambda_1 T\left(R_1(v_1^{(\text{H})}) + r_1\right),$$

if $\frac{c_2 - \Lambda_2}{v_1^{(\text{H})}} = \min\{\Lambda_1, \frac{c_2 - \Lambda_2}{v_1^{(\text{H})}}, \frac{c_1}{1 - v_1^{(\text{H})}}\}$,

$$w^{\text{U}}(c_1, c_2, \Lambda_1, \Lambda_2; v_1^{(\text{H})}) = \Lambda_2 r_2 + \frac{c_2 - \Lambda_2}{v_1^{(\text{H})}} R_1(v_1^{(\text{H})}) + r_1 \min\left\{\Lambda_1, c_1 + c_2 - \Lambda_2\right\},$$

if $\frac{c_1}{1 - v_1^{(\text{H})}} = \min\{\Lambda_1, \frac{c_2 - \Lambda_2}{v_1^{(\text{H})}}, \frac{c_1}{1 - v_1^{(\text{H})}}\}$,

$$w^{\text{U}}(c_1, c_2, \Lambda_1, \Lambda_2; v_1^{(\text{H})}) = \Lambda_2 r_2 + \frac{c_1}{1 - v_1^{(\text{H})}} R_1(v_1^{(\text{H})}) + r_1 \min\left\{\Lambda_1, c_1 + c_2 - \Lambda_2\right\}.$$

Let $\mathbf{v}^{(\text{H})} = [v_1^{(\text{H})}, v_2^{(\text{H})}, \ldots, v_{n-1}^{(\text{H})}, 0]$, where each $v_i^{(\text{H})}$ is defined in Theorem 4, we have the total revenue of (n-type HP) is

$$w^{\text{U}} = w^{\text{U}}(c_{n-1}^{(1)}, c_n^{(1)}, \Lambda_{n-1}, \Lambda_n; v_{n-1}^{(\text{H})}) + \sum_{i=1}^{n-2} w^{\text{U}}\left(c_i^{(1)}, \left(c_{i+1}^{(1)} - \frac{\Lambda_{i+1}}{1 - v_{i+1}^{(\text{H})}}\right)^+, \Lambda_i, 0; v_i^{(\text{H})}\right),$$

where we input $\left(c_{i+1}^{(1)} - \frac{\Lambda_{i+1}}{1 - v_{i+1}^{(\text{H})}}\right)^+$ is because if we take $v_{i+1}^{(\text{H})}$ as the static upgrade probability, there will be $\frac{\Lambda_{i+1}}{1 - v_{i+1}^{(\text{H})}}$ type $i+1$ demand units not upgraded.

Define W_i^π as the expected total revenue generated by Algorithm 3 on type i requests. By Theorem 3, we have,

$$\mathbb{E}\left[w^{\text{U}}(c_{n-1}^{(1)}, c_n^{(1)}, \Lambda_{n-1}, \Lambda_n; v_{n-1}^{(\text{H})}) - w_{n-1}^\pi - w_n^\pi\right] = O(\log T),$$

and for any $i \in [n-2]$,

$$\mathbb{E}\left[w^{\text{U}}\left(c_i^{(1)}, \left(c_{i+1}^{(1)} - \frac{\Lambda_{i+1}}{1 - v_{i+1}^{(\text{H})}}\right)^+, \Lambda_i, 0; v_i^{(\text{H})}\right) - W_i^\pi\right] = O(\log T).$$

Therefore, we can obtain

$\mathbb{E}\left[w^{\text{U}} - W^\pi\right]$

$= \mathbb{E}\left[w^{\text{U}}(c_{n-1}^{(1)}, c_n^{(1)}, \Lambda_{n-1}, \Lambda_n; v_{n-1}^{(\text{H})}) - W_{n-1}^\pi - W_n^\pi\right] + \sum_{i=1}^{n-2} \mathbb{E}\left[w^{\text{U}}\left(c_i^{(1)}, \left(c_{i+1}^{(1)} - \frac{\Lambda_{i+1}}{1 - v_{i+1}^{(\text{H})}}\right)^+, \Lambda_i, 0; v_i^{(\text{H})}\right) - W_i^\pi\right]$

$= (n-1)O(\log T) = O(n \log T).$

□