



LBS Research Online

Q Chen, S Jasin and I Duenyas

Real-Time Dynamic Pricing with Minimal and Flexible Price Adjustment

Article

This version is available in the LBS Research Online repository: <https://lbsresearch.london.edu/id/eprint/1020/>

Chen, Q, Jasin, S and Duenyas, I

(2016)

Real-Time Dynamic Pricing with Minimal and Flexible Price Adjustment.

Management Science, 62 (8). pp. 2437-2455. ISSN 0025-1909

DOI: <https://doi.org/10.1287/mnsc.2015.2238>

INFORMS (Institute for Operations Research and Management Sciences)

<http://pubsonline.informs.org/doi/abs/10.1287/mnsc...>

Users may download and/or print one copy of any article(s) in LBS Research Online for purposes of research and/or private study. Further distribution of the material, or use for any commercial gain, is not permitted.

Real-Time Dynamic Pricing with Minimal and Flexible Price Adjustment

Qi (George) Chen, Stefanus Jasin, Izak Duenyas

Stephen M. Ross School of Business
University of Michigan, Ann Arbor, MI 48109 USA
georgeqc, sjasin, duenyas@umich.edu

We study a standard dynamic pricing problem where the seller (a monopolist) possesses a finite amount of inventories and attempts to sell the products during a finite selling season. Despite the potential benefits of dynamic pricing, many sellers still adopt a static pricing policy due to (1) the complexity of frequent re-optimizations, (2) the negative perception of excessive price adjustments, and (3) the lack of flexibility caused by existing business constraints. In this paper, we develop a family of pricing heuristics that can be used to address all these challenges. Our heuristic is computationally easy to implement; it requires only a single optimization at the beginning of the selling season and automatically adjusts the prices over time. Moreover, to guarantee a strong revenue performance, the heuristic only needs to adjust the prices of a *small* number of products and do so infrequently. This property helps the seller focus his effort on the prices of the most important products instead of all products. In addition, in the case where not all products are equally admissible to price adjustment (due to existing business constraints such as contractual agreement, strategic product positioning, etc.), our heuristic can immediately substitute the price adjustment of the original products with the price adjustment of similar products and maintain an equivalent revenue performance. This property provides the seller with extra flexibility in managing his prices.

Key words: dynamic pricing; revenue management; heuristic; asymptotic analysis.

History: .

1. Introduction

Nowadays, Revenue Management (RM) practice has become very prevalent in many industries such as airlines, hospitality, fashion, ground transportation, and many others (Talluri and van Ryzin 2005, chap.10). In a typical RM setting, the seller possesses a finite amount of inventories and attempts to maximize his revenue by selling a collection of products during a finite selling season. Often times, replenishment of inventory is not viable during the selling season and the leftovers have little salvage value (e.g., empty hotel rooms). There are two types of RM commonly found in practice: *quantity*-based RM and *price*-based RM. In the first category, prices are fixed over the selling season and the focus is on making a dynamic resource allocation. As for the second category, prices become the key decision variables and the seller adjusts his prices as often as he wishes and sells all products until stock-out. Although the two types of RM are not mutually exclusive, market context and the seller's value proposition may dictate which of the two is more appropriate. In

this paper we are primarily interested in price-based RM. (For a review of quantity-based RM, see Talluri and van Ryzin (2005, chap.2).)

Pricing is, without doubt, one of the most important decisions that affect the seller's profitability. According to a study by McKinsey & Company, "Pricing right is the fastest and most effective way for managers to increase profits" (Marn et al. 2003). The study argues that a 1% price increase in a typical S&P 1500 company would generate an 8% increase in operating profit, an impact which is almost 50% greater than that of a 1% reduction in variable cost and more than three times greater than that of a 1% increase in volume. Perhaps more strikingly, an annual report of the operating profit for airlines and rental car companies in the US during 2009 reveals that a 1% increase in average price improved total operating profit by up to 67% and 30%, respectively (Sen 2013). (Although a 67% improvement in profit is arguably rather unusual, a moderate 8% – 25% increase via dynamic pricing is not uncommon (Sahay 2007).) And yet, despite its apparent benefit, dynamic pricing still poses several serious challenges. First, the *complexity* of the required large-scale optimization leads to prohibitive computational burden. To illustrate, a typical major US airline operates thousands of flights daily and posts fares several months into the future. Accounting for the number of different booking classes per flight, this can easily translate into daily pricing decisions for *millions* of itineraries. Hotel industry is no exception. Koushik et al. (2012) reports that a single run of price optimization at the InterContinental Hotels Group (excluding the estimation time) takes about four hours to complete. Similarly, Pekgun et al. (2013) also reveals that it takes about six hours for the Carlson Rezidor Hotel Group to complete its price optimization *once*. Given the increased competition in many industries where the prices of some products are now being adjusted even *hourly* (Rigby et al. 2012), this begs the question whether there exists a scalable pricing heuristic which can be easily implemented in real-time.

Second, dynamic pricing typically involves *frequent* price adjustments of *many* products, which may not be desirable for the firms. For one thing, even when full-scale dynamic pricing tools are readily available, the seller may want to intentionally avoid excessive price adjustments due to brand positioning and customer relationship considerations. Widely accepted as it is in the airline industry, dynamic pricing suffers a considerable setback in some other industries due to negative customers' perception. For example, in hotel industry, the most common criticism of dynamic pricing is that it treats customers unequally and unfairly (Ramasastry 2005), and lab experiments confirm the unfairness perception of price discrimination (Haws and Bearden 2006). Aside from the customers' perception issue, frequent price adjustments of many products may also not be feasible due to existing *business constraints*, i.e., the seller may not have the flexibility to adjust the prices of some products because of existing regulations and contractual agreement. For example, hotels often face customers from the so-called *negotiated segment* and provide *fixed* corporate rates for

large travel buyers such as IBM and HP (Koushik et al. 2012). Thus, hotels are practically forced to provide a fixed price that cannot be adjusted over time to the negotiated segment while at the same time are free to dynamically adjust the prices for other customer segments. This situation is not unique to hotel industry alone. The practice of *selective* dynamic pricing, which combines dynamic pricing of some products with fixed pricing of other products, is not uncommon and can be found in many industries (e.g., with the exception of Sears, Amazon.com, and Kmart, most retailers only change their prices daily on less than 10% of their assortments (Rigby et al. 2012)). And yet, despite its common practice, we are not aware of any work in the academic literature that rigorously analyzes the feasibility and effectiveness of such approach.

The preceding discussions lead to several important research questions: (1) Can we construct a pricing heuristic that is easy to implement and does not require frequent price adjustments? (2) Can we adjust the price of only a *small* number of products in order to mitigate customers' negative perception while at the same time maintaining a decent revenue performance? If such minimal price adjustment is possible, (3) how should we pick the set of products whose prices are to be updated? Is there a simple rule that can be used as a guidance? Moreover, in the case where the seller's business constraints disallow him to dynamically adjust the prices of some products, (4) can he still maintain an equivalent revenue performance by dynamically adjusting the price of other products? If yes, which *other* products should be used? In this paper we address all these questions. In particular, we will construct a family of real-time heuristics which, depending on the firm's need, can be used to address any of the aforementioned issues.

Static price control and re-optimization. There is a rich operations management (OM) literature on dynamic pricing. (See Bitran and Caldentey (2003) and Elmaghraby and Keskinocak (2003) for overviews.) In the RM context, motivated by the well-known curse of dimensionality of Dynamic Program (DP), many existing works have focused on the construction of easy-to-use heuristics. There are two popular approaches that can be found in the literature. The first is based on the so-called *Approximate Dynamic Programming* (ADP). Some works along this line are Erdelyi and Topaloglu (2011) and Kunnumkal and Topaloglu (2010). The second approach, which is closer to our work in this paper, is based on solving a *deterministic* analog of the original stochastic problem. One of the seminal works on this approach is Gallego and van Ryzin (1997). The trade-off between the two approaches is obvious. On the one hand, the sophisticated ADP requires more computational power than the deterministic approach. On the other hand, while the former yields an "adaptive" price sequence, which depends on sales realization, the latter only results in a deterministic (static) price. The good news is that static price control is *asymptotically* optimal (Gallego and van Ryzin 1997). This may partly explain its decent performance, hence its wide

adoption, in many industries. Yet, a considerable amount of revenue is still lost. As noted earlier, the main drawback of static pricing is that it completely ignores the observed demand realizations and the remaining inventory levels. One potential way of utilizing this progressively revealed information is to periodically *re-optimize* the aforementioned deterministic optimization. The impact of re-optimization in quantity-based RM has been extensively studied in the literature (e.g., see Chen and de Mello (2010), Reiman and Wang (2008), Secomandi (2008), Ciocan and Farias (2012), Jasin and Kumar (2012, 2013)). As for price-based RM, Maglaras and Meissner (2006) is the first to show that re-optimizing static price control guarantees at least the same asymptotic performance as static price without re-optimization. Thus, although re-optimization does not necessarily result in a monotonically increasing revenue, it cannot severely degrade revenue either. This is in contrast to the potentially negative impact of re-optimization in quantity-based RM (Jasin and Kumar 2013). Chen and Farias (2013) analyze the impact of re-optimization in the presence of imperfect forecast for a single product RM. They show that a combination of re-optimization and re-estimation yields a significant improvement in revenue. The paper that is perhaps closest to ours is Jasin (2014). The author provides a tighter bound for the expected revenue loss of the re-optimized static price control studied in Maglaras and Meissner (2006). This confirms the theoretical benefit of re-optimization for a very general class of multi-product and multi-resource RM. In addition, the author also proposes a simple pricing heuristic that can be implemented in real-time. (See Section 4 for further discussions on this.) A parallel but independent work by Atar and Reiman (2012) studies a continuous time version of the same problem and shows that the problem can be reduced to a diffusion control problem whose optimal solution is a Brownian bridge. The Brownian bridge structure motivates them to develop a diffusion-scale dynamic pricing heuristic that has similar error correction terms as the simple heuristic developed in Jasin (2014).

Although re-optimization is intuitively appealing and enjoys a good theoretical guarantee, unfortunately, it is not always practically feasible. As previously discussed, even a single optimization of a large-scale real problem instance can take hours to complete (Pekgun et al. 2013). This obviously serves as a bottleneck for the number of re-optimizations that can be implemented in one day. A recent work by Golrezaei et al. (2014) in the context of assortment optimization also highlights the same issue. The problem being re-optimized in their setting is a linear program, which is considered by many as one of the most tractable family of optimization problems. And yet, their simulation shows that the running time of frequent re-optimizations can be 800 times larger than that of a single optimization. While the resulting time-lag due to re-optimization may not be too detrimental for brick-and-mortar stores who update their prices less frequently, it is clearly less feasible for online retailers with more frequent price adjustments. In such settings, any proposed control must ideally be implementable in real-time without unnecessarily invoking large-scale re-optimization.

The proposed heuristic. In this paper, we introduce a new family of dynamic pricing heuristics, which we call *Linear Price Correction* (LPC). LPC only requires a single deterministic optimization at the beginning of the selling season and can be implemented in real-time. In addition, LPC only needs to adjust the price of a small number of products, admits a general asynchronous update schedule, and allows update substitution among “similar” products. Needless to say, it is also possible to couple LPC with occasional re-optimizations to further improve its performance. All these properties taken together allow the seller to enjoy the benefit of dynamic pricing while at the same time reducing the computational burden of re-optimization and mitigating the negative effect of frequent price changes on customers’ perception.

The remainder of the paper is organized as follows. Section 2 describes the problem setting and the asymptotic approach we take to analyze the performance of any dynamic pricing heuristic. The proposed heuristic LPC is formally introduced in Section 3 where we also discuss its minimal and asynchronous price adjustment properties which allow LPC to achieve good performance by adjusting the prices of only a small number of products and do so infrequently. In Section 4, we show the flexibility of LPC in choosing the prices of which products to adjust by demonstrating how to achieve equivalent revenue performances by adjusting prices of different sets of products that are “equivalent”. Section 5 uses numerical experiments to show the strong performance of LPC and its modifications, and to illustrate the managerial insights drawn from previous sections. Finally, Section 6 concludes. The proof of Theorem 1 can be found in the Online Supplement and the proofs of other results are deferred to Appendix A and B.

2. Problem Formulation

We consider a multi-period and multi-product pricing problem where the seller sells a catalog of n products (indexed by j), each of which is made up of a combination of m types of resources (indexed by i) whose initial inventory levels are given by $C \in \mathbb{R}^m$. As is usually the case, the number of products is much larger than the number of resources. We introduce a matrix $A = [A_{ij}]$, commonly known as the *consumption matrix*, whose element A_{ij} indicates the amount of resource i required by one unit of product j . Without loss of generality, we assume that the rows of A are linearly independent. The selling season is finite and divided into T periods. At the beginning of period t , the seller posts the price $p_t = (p_{t,j})$. The price then induces a demand $D_t(p_t) = (D_{t,j}(p_t))$ with rate $\lambda(p_t) = \mathbf{E}[D_t(p_t)]$. As is common in the literature, we allow at most one customer arrival per period. Hence, the function $\lambda(p_t)$ can also be interpreted as the arrival probability in period t . Let $r(p_t) := p'_t \lambda(p_t)$ denote the revenue rate in period t , where p'_t indicates the transpose of p_t . Let Ω_p and Ω_λ denote the convex set of feasible prices and demand rates, respectively. We make the following assumptions:

(A1) The demand function $\lambda(p_t) : \Omega_p \rightarrow \Omega_\lambda$ is invertible, twice differentiable, monotonically decreasing in its individual argument, and bounded from above by $\bar{\lambda}$.

(A2) The revenue function $r(p_t) = p_t' \lambda(p_t) = \lambda_t' p(\lambda_t) = r_t(\lambda_t)$ is continuous, strictly jointly concave in λ_t , and bounded from above by \bar{r} .

(A3) For each product j , there exists a turn-off price p_j^∞ such that if $\{p^k\}$ is any price sequence satisfying $p_j^k \rightarrow p_j^\infty$, then we have $\lambda_j(p^k) \rightarrow 0$.

(A4) The absolute eigenvalues of $\nabla^2 \lambda_j(p_t)$ and $\nabla^2 r(p_t)$ are bounded from above by \bar{v} .

Assumptions (A1) - (A3) are similar to the standard regularity conditions in Gallego and van Ryzin (1997). (A1) is a mild assumption to ensure basic analytical properties of the demand rate. (A2) follows from the invertibility assumption in (A1) and is needed to guarantee that the function $r(\cdot)$ has a unique, bounded optimizer. The revenue functions under a vast class of demand models such as linear and logit demand satisfy these assumptions. As for (A3), the existence of turn-off prices allow us to effectively shut down the demand for any product whenever desirable. (A4) is easily satisfied in general, especially for compact Ω_p . The constants $\bar{\lambda}$, \bar{r} and \bar{v} are independent of t .

The RM pricing problem. The optimal stochastic pricing problem can be written as:

$$(SPP): \quad J^{Stoc} = \max_{\pi \in \Pi_p} \mathbf{E} \left[\sum_{t=1}^T (p_t^\pi)' D_t(p_t^\pi) \right] \quad \text{s.t.} \quad A \left[\sum_{t=1}^T D_t(p_t^\pi) \right] \leq C,$$

where Π_p is the set of all non-anticipating pricing policies and the constraints must hold almost surely. Alternatively, by the invertibility of demand function, we can also use $\{\lambda_t\}$ as the decision variables and replace p_t and $D_t(p_t)$ with $p_t(\lambda_t)$ and $D_t(\lambda_t)$ respectively. We then replace the random variables in SPP by their mean and obtain a more tractable deterministic formulation below.

$$(DPP): \quad J^{Det} = \max \sum_{t=1}^T r(\lambda_t) \quad \text{s.t.} \quad \sum_{t=1}^T A \lambda_t \leq C \quad \text{and} \quad \lambda_t \in \Omega_\lambda, \quad \forall t.$$

Let $\{\lambda_t^D\}$ denote the unique optimal solution to DPP. Correspondingly, we define $p_t^D := p(\lambda_t^D)$. Since demand is time-homogeneous, it can be shown that $\lambda_t^D = \lambda_1^D := \lambda^D$ and $p_t^D = p_1^D := p^D$ for all t . This explains the name *static* pricing. For analytical tractability, we will assume that λ^D lies in the interior of Ω_λ . We formally state this assumption below.

(A5) There exist strictly positive constants ϕ_L and ϕ_U such that $[\lambda^D - \phi_L \mathbf{e}, \lambda^D + \phi_U \mathbf{e}] \subseteq \Omega_\lambda$.

Assumption (A5) essentially says that *all* products matter. It implies the optimal deterministic price is neither so low that it induces too many requests nor so high that it completely shuts down the demand of some products. As a practical rule of thumb, if some products are not profitable (i.e.

$\lambda_j^D = 0$ for some j), they can be discarded from the catalog and we can re-run the optimization. This helps the seller to focus on the products that matter. Hence, (A5) is *not* restrictive at all.

Performance measure and asymptotic regime. Ideally, we would like to define *revenue loss* of any control π as the difference between the revenue earned under the optimal pricing policy and the revenue earned under the control. Since the former is not easy to compute, we resort to using an upper bound as an approximation. It is known that $J^{Stoc} \leq J^{Det}$. (This is a standard result in the literature and is an immediate consequence of Jensen’s inequality. We omit its proof.) Let R_π denote total revenue earned under heuristic π throughout the selling season. The expected revenue loss of heuristic π is then defined as: $RL_\pi = J^{Det} - \mathbf{E}[R_\pi]$. Following Gallego and van Ryzin (1997), in this paper we consider a sequence of increasing problems parameterized by $\theta > 0$. To be precise, in the θ^{th} problem, we scale both the length of selling season and the initial inventory levels by a factor of θ while keeping all the other parameters unchanged. If we let $T(\theta)$ and $C(\theta)$ denote the length of the selling season and initial inventory levels in the θ^{th} problem, respectively, then $T(\theta) = \theta T$ and $C(\theta) = \theta C$. One may interpret the parameter θ as the *scale*, or *relative size*, of the problem. (If C is normalized to 1, then θ has an immediate interpretation as the size of initial inventory levels. Alternatively, if T is normalized to 1, the scale θ can be interpreted as the size of potential demands.) Notationwise, we will simply attach (θ) as a reference to the θ^{th} problem. Observe that the optimal solution of the scaled deterministic problem is the same as the optimal solution of the unscaled one (i.e., $\lambda_t^D(\theta) = \lambda^D$ and $p_t^D(\theta) = p^D$), so we have $J^{Det}(\theta) = \theta J^{Det}$.

3. Minimal and Asynchronous Price Adjustments

In this section, we will develop a pricing heuristic that adjusts the prices of only a small number of products and admits a general asynchronous update schedule. We show that our heuristic guarantees a strong asymptotic performance despite the fact that it only adjusts the prices of a small number of products. This has an obvious managerial significance. For example, at Chicago O’Hare airport, United Airlines operates more than forty routes to and from the North East and another thirty or so routes to and from the West Coast and the Mountain Area (see www.united.com). Assuming one fare class per flight, the company needs to price approximately $40 \times 30 = 1,200$ itineraries from the North East to the West Coast and the Mountain Area that make one stop at O’Hare airport. Our result suggests that United only needs to dynamically price $40 + 30 = 70$ itineraries instead of 1,200. Moreover, the price of these 70 itineraries can be adjusted asynchronously instead of simultaneously.

To introduce our heuristic, we start with a notion of a *base*. (This is the set of products whose prices are to be adjusted under the heuristic. We will allow more adjustable prices in Section 4.) A subset of products \mathcal{B} is said to be a *base* if (1) it contains exactly m products and (2) the products

in \mathcal{B} span the resource space, meaning the columns of matrix $A\nabla\lambda(p^D)$ that correspond to the products in \mathcal{B} (by the same index) span the whole \mathbb{R}^m . Note that, since the rows of A are linearly independent and $\nabla\lambda(p^D)$ is invertible, the rank of $A\nabla\lambda(p^D)$ is m . So, there always exists at least one base. Let H be a real n by m matrix satisfying $AH = I$, where I is an m by m identity matrix. We call H a *projection* matrix and say that a projection matrix H *selects the base* \mathcal{B} if the rows of $\nabla p(\lambda^D)H$ (by the same index) that correspond to the products *not* in \mathcal{B} are all zero vectors. As will be evident shortly, a proper choice of matrix H is important to ensure that only the prices of the base products are dynamically adjusted while the prices of the non-base products are never changed. The following lemma establishes the existence of a projection matrix for any given base.

LEMMA 1. *For any base \mathcal{B} , there exists a unique projection matrix H that selects it.*

The heuristic. Fix a base \mathcal{B} and assume without loss of generality that $\mathcal{B} = \{1, \dots, m\}$. For each $j \in \mathcal{B}$, define $\gamma_j = \{t_l^j : 1 \leq l \leq K_j\}$ to be the updating schedule for product j . (An updating schedule can be viewed as a business constraint that prescribes when the price of a given product is adjustable.) In particular, the l^{th} updating time is denoted by t_l^j and the number of updates is K_j . For convenience, we will write $t_0^j = 1$ and $t_{K_j+1}^j = T + 1$. Let $k_t^j = \max\{k : t_k^j \leq t\}$ denote the number of price updates for product j by time t . This setting is very general: *We allow the price of each product in the base to be updated asynchronously (i.e., independently of the other products in the base).* Let H be a projection matrix that selects \mathcal{B} . For any set $\mathcal{A} \subseteq \{1, \dots, n\}$, let $E^{\mathcal{A}}$ denote an n by n diagonal matrix with $E_{ii}^{\mathcal{A}} = 1$ if $i \in \mathcal{A}$ and 0 otherwise. (This matrix helps select a set of rows of another matrix when it is left-multiplied, e.g., $E^j \nabla p(\lambda^D)H$ is a matrix whose j^{th} row is the same as the j^{th} row of $\nabla p(\lambda^D)H$ and all its other rows are zeros.) Define $\Delta_t(p_t) := D_t(p_t) - \mathbf{E}[D_t(p_t)] = D_t(p_t) - \lambda(p_t)$ and $\tilde{\Delta}_l^j := \sum_{s=t_{l-1}^j}^{t_l^j-1} \Delta_s(p_s)$, $l = 1, \dots, K_j + 1$. The term $\Delta_t(p_t)$ can be interpreted as demand error during period t and the term $\tilde{\Delta}_l^j$ can be interpreted as *cumulative* demand errors between two subsequent updating times for product j . (For brevity, whenever there is no confusion, we will often suppress notational dependency on p_t and simply write Δ_t , D_t , and λ_t .) Let C_t denote the remaining inventory levels at the end of period t . The definition of our heuristic is given below.

Linear Price Correction (LPC)

1. During period 1, set $p_1 = p^D$.
2. At the beginning of period $t > 1$, do:

- a. First compute $\hat{p}_t = p^D - \sum_{j=1}^m E^j \nabla p(\lambda^D)H \left[\sum_{l=1}^{k_t^j} \frac{A\tilde{\Delta}_l^j}{T - t_l^j + 1} \right]$.

b. Set the price according to the following rule:

- (1) If $C_{t-1} \geq A^j$ for all j , and $\hat{p}_s \in \Omega_p$ for all $s \leq t$, set $p_{t,j} = \hat{p}_{t,j}$;
- (2) Otherwise, set $p_{t,j} = p_j^\infty$.

The idea behind LPC is to use static price p^D as baseline prices and apply real-time adjustment to only the prices of m products in the chosen base. The proposed adjustment has an intuitive interpretation: If past demand realization is higher than expected (i.e., the term $\tilde{\Delta}$'s are positive), then LPC immediately increases future prices; if, on the other hand, past demand realization is lower than expected, then LPC immediately decreases future prices. To see that the given update formula only adjusts prices of base products, define $\tilde{\xi}_l^j \mathbf{e}_j := E^j \nabla p(\lambda^D) H A \tilde{\Delta}_l^j$ and $\xi_s^j \mathbf{e}_j := E^j \nabla p(\lambda^D) H A \Delta_s$, where \mathbf{e}_j is a vector with proper size whose j^{th} element equals one and any of its other elements equals zero. Note that we can write \hat{p}_t as:

$$\begin{bmatrix} \hat{p}_{t,1} \\ \vdots \\ \hat{p}_{t,m} \\ \hat{p}_{t,m+1} \\ \vdots \\ \hat{p}_{t,n} \end{bmatrix} = \begin{bmatrix} p_1^D - \sum_{l=1}^{k_t^1} \frac{\tilde{\xi}_l^1}{T-t_l^1+1} \\ \vdots \\ p_m^D - \sum_{l=1}^{k_t^m} \frac{\tilde{\xi}_l^m}{T-t_l^m+1} \\ p_{m+1}^D \\ \vdots \\ p_n^D \end{bmatrix}.$$

Obviously, only the prices of the first m products are adjusted. Moreover, for each $j \in \mathcal{B}$, if the current period t is such that $t_{l-1}^j < t < t_l^j$ for some l , then $p_{t,j} = p_{t-1,j}$. So, the price of product $j \in \mathcal{B}$ in the periods between two subsequent updating times does not change. To help the reader better understand the mechanism of this pricing heuristic, we give an example below.

EXAMPLE 1. Consider a network RM with 3 products and 2 resources. Without loss of generality, we assume that $\mathcal{B} = \{1, 2\}$ is a base. Suppose that $\gamma_1 = \{2, 5, \dots\}$ and $\gamma_2 = \{4, 5, \dots\}$ (i.e., we want to adjust the price of product 1 in periods 2, 5, etc. and the price of product 2 in periods 4, 5, etc.). Assuming no stock-out, the price formula for the first five periods, are given by:

$$\begin{bmatrix} p_{1,1} \\ p_{1,2} \\ p_{1,3} \end{bmatrix} = \begin{bmatrix} p_1^D \\ p_2^D \\ p_3^D \end{bmatrix}, \quad \begin{bmatrix} p_{2,1} \\ p_{2,2} \\ p_{2,3} \end{bmatrix} = \begin{bmatrix} p_1^D - \frac{\xi_1^1}{T-1} \\ p_2^D \\ p_3^D \end{bmatrix}, \quad \begin{bmatrix} p_{3,1} \\ p_{3,2} \\ p_{3,3} \end{bmatrix} = \begin{bmatrix} p_1^D - \frac{\xi_1^1}{T-1} \\ p_2^D \\ p_3^D \end{bmatrix},$$

$$\begin{bmatrix} p_{4,1} \\ p_{4,2} \\ p_{4,3} \end{bmatrix} = \begin{bmatrix} p_1^D - \frac{\xi_1^1}{T-1} \\ p_2^D - \frac{\xi_1^2 + \xi_2^2 + \xi_3^2}{T-3} \\ p_3^D \end{bmatrix}, \quad \text{and} \quad \begin{bmatrix} p_{5,1} \\ p_{5,2} \\ p_{5,3} \end{bmatrix} = \begin{bmatrix} p_1^D - \left(\frac{\xi_1^1}{T-1} + \frac{\xi_2^1 + \xi_3^1 + \xi_4^1}{T-4} \right) \\ p_2^D - \left(\frac{\xi_1^2 + \xi_2^2 + \xi_3^2}{T-3} + \frac{\xi_4^2}{T-4} \right) \\ p_3^D \end{bmatrix}.$$

General performance bound. We will now discuss the performance of LPC. We first provide a general bound that can be applied to arbitrary updating schedule and then we discuss its implication for several specific schedules. For the sake of generality, we will allow the choice of updating schedule to also depend on θ , i.e., $\gamma_j(\theta) = \{t_l^j(\theta) : 1 \leq l \leq K_j(\theta)\}$, $j \in \mathcal{B}$. Let $R_{H, \gamma_{\mathcal{B}}}(\theta)$ denote the total revenue earned under LPC with projection matrix H and updating schedules $\gamma_{\mathcal{B}} := \{\gamma_j(\theta)\}_{j \in \mathcal{B}}$. Let $\|\cdot\|_2$ denote the usual spectral norm of a matrix, i.e., $\|X\|_2^2$ equals the maximum eigenvalue of $X'X$. We state our result below.

THEOREM 1. *There exist positive constants Ψ and $\bar{\Psi}$ independent of $\theta \geq 1$, the projection matrix H that selects \mathcal{B} , and the choice of updating schedules $\{\gamma_j(\theta)\}_{j \in \mathcal{B}}$ such that*

$$\begin{aligned} J^{Det}(\theta) - \mathbf{E}[R_{H, \gamma_{\mathcal{B}}}(\theta)] &\leq \Psi + \bar{\Psi} \sum_{j \in \mathcal{B}} \sum_{t=1}^{T(\theta)-1} \min \left\{ 1, \|\nabla p(\lambda^D)HA\|_2^2 U_1^j(T(\theta), t) \right\} \\ &\quad + \bar{\Psi} \sum_{j \in \mathcal{B}} \sum_{t=1}^{T(\theta)-1} \min \left\{ 1, \|\nabla p(\lambda^D)HA\|_2^2 U_2^j(T(\theta), t) \right\}, \end{aligned}$$

where the terms $U_1^j(T, t)$ and $U_2^j(T, t)$ are defined as

$$U_1^j(T, t) = \frac{t - t_{k_t^j}^j + 1}{(T-t)^2} + \sum_{l=1}^{k_t^j} \frac{t_l^j - t_{l-1}^j}{(T-t_l^j + 1)^2} \quad \text{and} \quad U_2^j(T, t) = \frac{1}{T-t} \sum_{s=1}^t \sum_{l=1}^{k_s^j} \frac{t_l^j - t_{l-1}^j}{(T-t_l^j + 1)^2}.$$

We want to stress: The above bound is *very* general. It characterizes the performance of LPC for *any* given base and *any* given updating schedule¹, either synchronous or asynchronous. (The implications of Theorem 1 for specific schedules will be discussed below.) Note that the bound is *separable* over the products in the base. This suggests that the seller cannot compensate the lack of updating of one product in the base by applying more frequent updates to the remaining product(s) in the base. If there exist multiple feasible bases, the bound in Theorem 1 suggests that we use the base \mathcal{B} and the corresponding projection matrix H that minimizes $\|\nabla p(\lambda^D)HA\|_2$. Although, in general, it is not possible to explicitly characterize the “optimal” base products chosen by this selection rule, it turns out that we can provide a very intuitive characterization of the “optimal” base product for the case of single-resource RM.

LEMMA 2. *Suppose that $m = 1$. Among all projection matrices that select a base, the projection matrix H^* that achieves the smallest $\|\nabla p(\lambda^D)HA\|_2$ selects the base that consists of product $j^* = \arg \max_{j=1, \dots, n} |(A \nabla \lambda(p^D))_j|$.*

¹ In the setting of quantity-based RM, Jasin and Kumar (2012) also provide a bound for revenue loss which depends on a general choice of updating schedule. However, they assume that the admission control for *all* products must be *simultaneously* updated at the same time. In contrast, LPC allows *each* product to have its own updating schedule. This level of generality, together with the non-linearity of the objective function and capacity constraints, introduces non-trivial analytical subtleties which do not previously exist in the analysis of Jasin and Kumar (2012).

The intuition of the above lemma is most easily explained if we consider a special case of single-resource RM with $A = [1, \dots, 1]$ and separable demands (i.e., $\lambda_j(p)$ only depends on p_j). In this setting, $A\nabla\lambda(p^D)$ becomes a row vector whose j^{th} element equals the demand sensitivity of product j with respect to its own price, $\lambda'_j(p_j^D)$. Thus, under LPC, the optimal projection matrix selects the most price-sensitive product into the base. This can be intuitively explained as follows: Among all products, product j^* needs the smallest price perturbation to correct the same demand error. Since we are using the deterministic model as our performance benchmark, ideally, we would want to have a price trajectory that stays as close as possible to the baseline price p^D . This can be achieved by adjusting the product that requires the smallest perturbation. As for the more general case of single-resource RM with general demand and general capacity consumption matrix A , a similar intuition also holds: We want to pick the product whose price adjustment has the largest impact on capacity consumption.

Special updating schedules. We will now apply the result of Theorem 1 to derive an explicit performance bound for several special updating schedules that only adjust the prices of base products and draw some managerial insights. We start with the most commonly used update schedule where prices are being adjusted periodically according to some frequencies.

COROLLARY 1. (*h*-PERIODIC SCHEDULE) *Fix $h(\theta) \geq 1$ and define $t_l^j(\theta) = t_l(\theta) = lh(\theta) + 1$ for all $j \in \mathcal{B}$. There exist positive constants Ψ , $\hat{\Psi}$, and $\bar{\Psi}$ independent of $\theta \geq 1$ and $h(\theta) \geq 1$ such that the expected revenue loss of LPC is bounded by $\Psi + \hat{\Psi}\sqrt{h(\theta)} + \bar{\Psi}\log^2\theta$.*

Two comments are in order. First, if $h(\theta) = T(\theta)$, then the periodic schedule reduces to static pricing and the revenue loss is $O(\sqrt{\theta})$. This bound is consistent with the result in Gallego and van Ryzin (1997). If, on the other hand, $h(\theta) = 1$, the revenue loss is reduced to $O(\log^2\theta)$. Since LPC requires only one optimization followed by simple price updates, it provides a significant improvement² over static pricing with negligible computational effort. Second, although Corollary 1 assumes a synchronous schedule, it is not difficult to derive a bound for an asynchronous periodic update schedule because the bound is separable in individual product. For example, one plausible asynchronous schedule would be to adjust the prices of base products on weekly basis, but on different days of the week. The asymptotic performance bound will remain the same as in Corollary 1. One caveat of periodic schedule is that, in order to reduce the revenue loss to $O(\log^2\theta)$, a very frequent updates of the prices of all base products (roughly $\Theta(\theta)$ times) is required. But, per

² Since θ represents the size of the problem, the percentage revenue loss under LPC is approximately $\frac{\log^2\theta}{\theta} \times 100\%$ whereas the percentage revenue loss under static pricing is about $\frac{\sqrt{\theta}}{\theta} \times 100\%$. Numerically, for a problem instance with initial inventory levels equal to 100, as in a typical airplane with 100 seats, our experiments in Section 6 show a 2% improvement in revenue, which is quite significant for typical RM applications.

our discussions in Section 1, this may not be practically feasible – or even if it is, it may not be strategically desirable due to customers’ perception issue. To address this, below we propose two schedules that still guarantee $O(\log^2 \theta)$ revenue loss albeit with much fewer price updates.

COROLLARY 2. (α -POWER SCHEDULE) Fix $\alpha \geq 1$. For all $j \in \mathcal{B}$, let $t_0^j(\theta) = t_0(\theta) = 1$ and define $t_l^j(\theta) = t_l(\theta) = \left\lceil T(\theta) - \sum_{s=1}^{K(\theta)-l+1} s^\alpha \right\rceil$ for $1 \leq l \leq K(\theta)$, where $K(\theta) := \{k : \sum_{s=1}^k s^\alpha < T(\theta), \sum_{s=1}^{k+1} s^\alpha \geq T(\theta)\}$. Then $K(\theta) \leq ((\alpha + 1)T(\theta))^{1/(\alpha+1)}$ and there exist positive constants Ψ and $\bar{\Psi}$ independent of $\theta \geq 1$ such that the expected revenue loss of LPC is bounded by $\Psi + \bar{\Psi} \log^2 \theta$.

COROLLARY 3. (β -GEOMETRIC SCHEDULE) Fix $\beta > 1$. For all $j \in \mathcal{B}$, let $t_0^j(\theta) = t_0(\theta) = 1$, and for $l \geq 1$, iteratively define $t_l^j(\theta) = t_l(\theta) = \left\lceil \frac{(\beta-1)T(\theta) + t_{l-1}(\theta)}{\beta} \right\rceil$ as long as $t_{l-1}(\theta) < T(\theta)$. Let $K(\theta)$ be such that $t_{K(\theta)}^j(\theta) = T(\theta)$. Then, $K(\theta) \leq 1 + \log_\beta T(\theta)$, and there exist positive constants Ψ and $\bar{\Psi}$ independent of $\theta \geq 1$ such that the expected revenue loss of LPC is bounded by $\Psi + \bar{\Psi} \log^2 \theta$.

Corollaries 2 and 3 offer two interesting insights. First, by carefully choosing the update times, we can use a small number of updates (only about $\theta^{\frac{1}{\alpha+1}}$ updates with power schedule and $\log_\beta \theta$ updates with geometric schedule) to guarantee a $O(\log^2 \theta)$ revenue loss.³ Second, for both schedules, most of the updates happen near the end of the selling season. This implies that the crucial moments for dynamic pricing is near the end of the selling season instead of at the beginning, which suggests that the seller can perhaps apply static price at the beginning of the season and only switch to dynamic pricing later. Needless to say, although Corollaries 2 and 3 assume synchronous schedules, it is also possible to use asynchronous schedules. For example, the prices of some base products can be updated using power schedule and the prices of other base products can be updated using geometric schedule. Again, since the bound in Theorem 1 is separable over the products in the base, the $O(\log^2 \theta)$ bound still holds.

The impact of adjusting the prices of fewer, or more, than m products. Since adjusting the price of all products may not be desirable, or even feasible, it is important that we understand the impact of restricting the number of adjustable products on revenue. Corollaries 1-3 partially answer this question by showing a surprising result that adjusting the prices of only m products (in the base) is sufficient to guarantee a $O(\log^2 \theta)$ revenue loss.⁴ This is a powerful result because,

³ Our simulations show that the non-asymptotic performance of *1-Power* schedule is almost the same as that of *1-Periodic* schedule. This is very impressive since when $\theta = 500$, *1-Power* needs 44 adjustments while *1-Periodic* requires 500 adjustments. For larger θ , the difference is even bigger.

⁴ Since we only have m resources, it seems “intuitive” that we should be able to perform well by adjusting the prices of only m products. However, since adjusting the prices of only m products also affects the demands for the other $n - m$ products whose prices are not adjusted, it is not immediately clear what impact this would have on revenue. Our result is different from the so-called *action-space reduction* discussed in pg. 220 of Talluri and van Ryzin (2005). Under

in most RM applications, the number of resources m is typically much smaller than the number of products n . In particular, it provides an important managerial insight that the seller does not need to aggressively adjust the prices of all products to benefit from dynamic pricing. The result on minimal price adjustment, however, leads to two interesting questions. First, can we still guarantee the $O(\log^2 \theta)$ revenue loss by adjusting the prices of fewer than m products? The answer is unfortunately negative and the revenue loss under such scenario is of order $\sqrt{\theta}$ in general. To understand why this is so, consider the case where demands are separable and $A = I$ is an m by m identity matrix. Since this corresponds to an aggregate of m independent problems (e.g., m independent one-stop flights), if we only dynamically adjust the price of $m' < m$ products, then we are effectively applying static price control to the remaining $m - m'$ problems, which we already know has $\Theta(\sqrt{\theta})$ revenue loss in general (Jasin 2014). Second, what is the incremental benefit of adjusting the prices of more than m products? To answer this, we again consider the case of a single-resource RM. (By minimal price adjustment property, we already know that we only need to adjust the price of *one* product to guarantee a significant improvement over static pricing. The question is whether adjusting the prices of more products has a significant impact on performance.) Let $b = (A \nabla \lambda(p^D))'$ and denote by $b_{(i)}$ the i^{th} largest element (in absolute value) of b . For $k \geq 1$, let Π_k denote the set of all non-anticipating pricing policies that adjust the price of at most k products in each period. (If the price of product j is not adjusted in period t under $\pi \in \Pi_k$, then $p_{t,j}^\pi = p_{t-1,j}^\pi$.) Then we have,

THEOREM 2. *Suppose that $m = 1$. There exist positive constants Ψ and $\bar{\Psi}$ independent of $\theta \geq 1$ and $1 \leq k \leq n$ such that*

$$\min_{\pi \in \Pi_k} \{J^{Det}(\theta) - \mathbf{E}[R_\pi(\theta)]\} \leq \Psi + \frac{\bar{\Psi}}{\sum_{i=1}^k b_{(i)}^2} \log^2 \theta.$$

The above performance bound suggests that the incremental benefit of adjusting the price of an additional product decreases as the number of the adjustable products increases. To see this, suppose that $A = [1, \dots, 1]$ and demands are separable and identical across different products with $\lambda_j(\cdot) = \lambda_1(\cdot)$ for all j . This implies $p_j^D = p_1^D$ for all j and $b_{(i)} = \lambda_1'(p_1^D)$ for all i . Then, the bound in Theorem 2 is of order $\frac{\log^2 \theta}{k}$. Since the function $1/k$ drops quickly for small k and slowly for large k , this suggests that it is not necessary for the seller to adjust the prices of too many products to get most of the potential revenue. (See Section 5 for numerical evidence of this observation in the

the action-space reduction scenario, we first compute the optimal *aggregate* decision variable and then *disaggregate* this variable to recover the optimal price for each product. However, there is no guarantee that this disaggregation will result in the adjustment of only the prices of m products. In contrast, under our scenario, the prices of $n - m$ products are never changed.

multi-resource case. Our results show that the revenue improvement of adjusting the price of m products over static pricing is about 80 – 90% of the revenue improvement of adjusting the price of all n products, in most cases. Moreover, in terms of revenue loss, while adjusting the price of m products reduces the revenue loss of static pricing by about 1 – 1.2%, adjusting the price of n products only further reduces the revenue loss by an additional 0.1 – 0.2% in most cases. (See Table 2 in Appendix C.) Given that the average margins in RM industries are typically very small, only about 3% (Irvine 2014), this highlights the practical significance of minimal adjustments for real-world implementation.) In particular, if the seller wishes to adjust the prices of more than m products to further increase revenue, then s/he only needs to consider adjusting the prices of a few more products instead of all.

4. Equivalent Performance via Adjusting the Prices of Other Products

Corollaries 2 and 3 in the previous section provide an important managerial insight: Managers need to update the prices of only a small subset of their products, and do so sufficiently rarely, to guarantee a strong revenue performance. Those results, however, assume that only the prices of the same m products are updated throughout the selling season. Can we do better? For example, why should we update the price of one product ten times and the other products not at all if a major concern of some practitioners is that customers get upset by frequent price changes? Can we reduce the number of price updates per product by somehow *distributing* the required adjustments across different products over different time periods (e.g., one price update per product for ten different products instead of ten price updates for one product)? Also, what if the seller dictates that the price of some products should not, or cannot, be changed either due to existing business constraints or contractual agreements? Can we somehow re-assign the scheduled update for these products to other “similar products”? As discussed in Section 1, although these questions have significant practical relevance and are faced by many sellers, we are not aware of any existing work in the literature addressing these issues. In this section, we will discuss a generalization of LPC that partially addresses these issues. Our proposed heuristic provides important practical insights on how to do *equivalent* pricing via adjusting the prices of similar products. To illustrate the basic idea, we start with two examples.

EXAMPLE 2. Consider a single flight RM with n types of ticket. We assume that each ticket only requires one seat and demands are separable. Note that $\nabla\lambda(p^D)$ is a diagonal matrix. As Corollary 3 indicates, it is sufficient to adjust the price of only one type of ticket $\Theta(\log_2 \theta)$ times to obtain $O(\log^2 \theta)$ revenue loss. If we evenly distribute these adjustments to all n tickets, the number of price updates per ticket is about $\lceil (\log_2 \theta)/n \rceil$. It turns out that this still guarantees $O(\log^2 \theta)$ revenue loss. Thus, dynamically adjusting one type of ticket $\Theta(\log_2 \theta)$ times is equivalent

to dynamically adjusting n types of tickets $\Theta((\log_2 \theta)/n)$ times for each. This has an important managerial implication. As an illustration, consider economy seats. There are usually about 13 different fare classes for economy seats. Since a typical US passenger flight has fewer than 500 seats and $\log_2(500) = 8.96$, by our previous arguments, we can either adjust the price of one fare class nine times or the price of *any* nine fare classes once during the selling season.

EXAMPLE 3. Consider a network RM problem with 3 resources and 6 products and suppose that

$$A\nabla\lambda(p^D) = \begin{bmatrix} -1 & 0 & 0 & 0 & 0 & -2 \\ 0 & -1 & 0 & -2 & -2 & 0 \\ 0 & 0 & -1 & -2 & 0 & -1 \end{bmatrix}.$$

Obviously, $\mathcal{B} = \{1, 2, 3\}$ forms a base. Suppose that the previously prescribed schedule for \mathcal{B} is $\gamma_1 = \{2, 3, 5\}$, $\gamma_2 = \{3, 4, 5\}$, and $\gamma_3 = \{4, 6\}$. Unlike in the previous example where we can arbitrarily pick any nine products, here, the choice of “similar products” is more subtle. A new set of products is similar to the original set of products if its corresponding columns (by the same index) in $A\nabla\lambda(p^D)$ can linearly represent the columns in $A\nabla\lambda(p^D)$ that correspond to the original set of products. In our example, this means that we can replace updating $\{2, 3\}$ in period 4 with $\{4, 5\}$, or replace updating $\{3\}$ in period 6 with $\{4, 5\}$. We cannot directly replace the price adjustment of product 3 in period 4 with product 4 because column 4 is not parallel to column 3. But, since product 2 will be adjusted in period 4 under both the original schedule and the new schedule, we can achieve an equivalent revenue by bundling the price adjustment of product 2 and 3 in period 4 and substituting it with the price adjustment of $\{2, 4\}$.

Equivalent pricing control. We now formally state the idea behind the preceding examples. For clarity, we assume that $\mathcal{B} = \{1, \dots, m\}$ is a base and H is a projection matrix that selects \mathcal{B} . Let $\gamma_{\mathcal{B}} := \{\gamma_j(\theta)\}_{j=1}^m$ denote the *existing* updating schedule for base products. We will show in this section that, for any *equivalent schedule* of $\gamma_{\mathcal{B}}$ (to be formally defined below), we can construct a pricing heuristic that guarantees the same asymptotic performance as LPC under $\gamma_{\mathcal{B}}$. In other words, if the seller wants to modify the current price updating schedules to a new one for strategic considerations, then we can provide a new pricing control that guarantees an equivalent performance as long as the new updating schedule is equivalent to the current updating schedule. Before introducing equivalent schedule, we first introduce the concept of *equivalent set*: A set of products $\mathcal{G} \subseteq \{1, \dots, n\}$ is said to be *equivalent* to the set $\mathcal{S} \subseteq \mathcal{B}$ (mathematically, we write: $\mathcal{G} \sim_{\mathcal{B}} \mathcal{S}$) if the columns in $A\nabla\lambda(p^D)$ that correspond to the products in \mathcal{S} can be written as a linear combination of the columns in $A\nabla\lambda(p^D)$ that correspond to products in \mathcal{G} . (Note that, by our definition, $\mathcal{G} \sim_{\mathcal{B}} \mathcal{S}$ does not imply $\mathcal{S} \sim_{\mathcal{B}} \mathcal{G}$.) Let $\mathcal{S}_t \subseteq \mathcal{B}$ be a subset of products that are adjusted in period t under $\gamma_{\mathcal{B}}$. Let \mathcal{G}_t be one of the (possibly) many sets that are equivalent to \mathcal{S}_t . We say that

a price updating schedule γ is an *equivalent schedule* of $\gamma_{\mathcal{B}}$ if in each period t only products in \mathcal{G}_t are adjusted under γ . Let $\mathbf{\Gamma}(\gamma_{\mathcal{B}})$ denote the set of all equivalent schedules of $\gamma_{\mathcal{B}}$. We now define an equivalent pricing control for any $\gamma \in \mathbf{\Gamma}(\gamma_{\mathcal{B}})$. Let $\mathcal{G}_t \sim_{\mathcal{B}} \mathcal{S}_t$ and denote by S_t and G_t the submatrices of $A\nabla\lambda(p^D)$ whose columns correspond to the products in \mathcal{S}_t and \mathcal{G}_t , respectively. By definition of equivalent set, there exists a $|\mathcal{G}_t|$ by $|\mathcal{S}_t|$ matrix Y_t such that $S_t = G_t Y_t$. For any such \mathcal{G}_t , \mathcal{S}_t and Y_t , we can construct a unique n by n matrix $Q_t = Q(Y_t, \mathcal{G}_t, \mathcal{S}_t)$ as follows: its submatrix with rows and columns not in $\mathcal{G}_t \cup \mathcal{S}_t$ equals an identity matrix, its submatrix with rows in \mathcal{G}_t and columns in \mathcal{S}_t equals Y_t , and any of its other elements equals 0. We call $Q(Y_t, \mathcal{G}_t, \mathcal{S}_t)$ a *transformation matrix* because, from its construction, it uses the matrix Y_t to transform the price adjustment for products in \mathcal{S}_t into price adjustment for products in \mathcal{G}_t . The following lemma provides some important properties of $Q(Y_t, \mathcal{G}_t, \mathcal{S}_t)$.

LEMMA 3. *For any $\mathcal{G}_t \sim_{\mathcal{B}} \mathcal{S}_t$ and any Y_t such that $S_t = G_t Y_t$, let $Q_t = Q(Y_t, \mathcal{G}_t, \mathcal{S}_t)$. Then, we have the following:*

- (1) $A\nabla\lambda(p^D)Q_t E^{\mathcal{B}} = A\nabla\lambda(p^D)E^{\mathcal{B}} = A\nabla\lambda(p^D)E^{\mathcal{G}_t \cup (\mathcal{B} - \mathcal{S}_t)} Q_t$;
- (2) *There exists a projection matrix H_t such that $\nabla p(\lambda^D)H_t = Q_t \nabla p(\lambda^D)H$ and the rows in $\nabla p(\lambda^D)H_t$ that correspond to products not in $\mathcal{G}_t \cup (\mathcal{B} - \mathcal{S}_t)$ are zeros;*
- (3) *The rows in $Q_t E^{\mathcal{S}_t} \nabla p(\lambda^D)H$ that correspond to products not in \mathcal{G}_t are zeros.*

Define $\mathcal{Q}_t(\gamma) := \arg \min_Q \{\|Q\|_2 : Q = Q(Y, \mathcal{G}_t, \mathcal{S}_t), S_t = G_t Y\}$ in each period t . (This optimization problem turns out to be a convex optimization with linear constraints and can be efficiently solved off-line.) We are now ready to introduce the concept of *equivalent pricing*. Let γ be an equivalent schedule of the existing schedule $\gamma_{\mathcal{B}}$. Then, a pricing control π with schedule γ is said to be *equivalent* to an existing LPC with updating schedule $\gamma_{\mathcal{B}}$ if, in Step 2a in the definition of LPC, it uses the following update formula:

$$\hat{p}_t = p^D - \sum_{j=1}^m \sum_{l=1}^{k_t^j} Q_{t_l^j} E^j \nabla p(\lambda^D) H \frac{A \tilde{\Delta}_l^j}{T - t_l^j + 1}$$

for some $Q_t \in \mathcal{Q}_t(\gamma)$ ⁵ in each period t . (In light of part (3) of Lemma 3, the above update formula guarantees that only adjustable products under γ are adjusted in each period.)

EXAMPLE 2 (CONT'D). Consider again the single flight problem described in Example 2. Suppose that $n = 3$ and assume, without loss of generality, that $\mathcal{B} = \{1\}$ with the corresponding projection

⁵ Note that, given $\gamma \in \mathbf{\Gamma}(\gamma_{\mathcal{B}})$, $\mathcal{Q}_t(\gamma)$ may not be a singleton. However, as can be seen in the proof of Theorem 3, the performance bound of an equivalent pricing control under γ depends on Q_t only via its spectral norm $\|Q_t\|_2$. In particular, the smaller the norm, the smaller the revenue loss bound. This observation motivates our definition of $\mathcal{Q}_t(\gamma)$ where $\|Q_t\|_2$ is minimized. Since all matrices in $\mathcal{Q}_t(\gamma)$ have the same spectral norm, our performance bound does not depend on the particular selection of Q_t within $\mathcal{Q}_t(\gamma)$.

matrix $H = (1, 0, 0)'$. Suppose that the seller originally plans to periodically adjust the price of only product 1 at the beginning of every period using the following update formula:

$$\begin{bmatrix} p_{t,1} \\ p_{t,2} \\ p_{t,3} \end{bmatrix} = p^D - \sum_{s=1}^{t-1} \nabla p(\lambda^D) H \frac{A \Delta_s}{T-s} = \begin{bmatrix} p_1^D - \sum_{s=1}^{t-1} p_1'(\lambda_1^D) \frac{\Delta_s}{T-s} \\ p_2^D \\ p_3^D \end{bmatrix}.$$

To develop an equivalent pricing control, which alternates among the three products such that the price of only one product is being adjusted in every period, we construct a sequence of transformation matrices $\{Q_{t_l}\}$ for each update time t_l as follows. Let Q^1 be a 3 by 3 identity matrix. For $j \in \{2, 3\}$, denote by Q^j the transformation matrix that transform the price adjustment of product 1 into price adjustment of product j . In particular, by the construction of transformation matrix

$$Q^2 = \begin{bmatrix} 0 & 0 & 0 \\ \frac{p_2'(\lambda_2^D)}{p_1'(\lambda_1^D)} & 0 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad Q^3 = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ \frac{p_3'(\lambda_3^D)}{p_1'(\lambda_1^D)} & 0 & 0 \end{bmatrix}.$$

For all l satisfying $l \equiv j \pmod{3}$, set $Q_{t_l} = Q^j$. The resulting equivalent pricing control is then given by $\hat{p}_t = p^D - \sum_{s=1}^{t-1} Q_s \nabla p(\lambda^D) H \frac{A \Delta_s}{T-s}$. Assuming no stock-out, the explicit formulae of the price of all three products for the first five periods are:

$$\begin{bmatrix} p_{1,1} \\ p_{1,2} \\ p_{1,3} \end{bmatrix} = \begin{bmatrix} p_1^D \\ p_2^D \\ p_3^D \end{bmatrix}, \quad \begin{bmatrix} p_{2,1} \\ p_{2,2} \\ p_{2,3} \end{bmatrix} = \begin{bmatrix} p_1^D - p_1'(\lambda_1^D) \frac{\Delta_1}{T-1} \\ p_2^D \\ p_3^D \end{bmatrix}, \quad \begin{bmatrix} p_{3,1} \\ p_{3,2} \\ p_{3,3} \end{bmatrix} = \begin{bmatrix} p_1^D - p_1'(\lambda_1^D) \frac{\Delta_1}{T-1} \\ p_2^D - p_2'(\lambda_2^D) \frac{\Delta_2}{T-2} \\ p_3^D \end{bmatrix},$$

$$\begin{bmatrix} p_{4,1} \\ p_{4,2} \\ p_{4,3} \end{bmatrix} = \begin{bmatrix} p_1^D - p_1'(\lambda_1^D) \frac{\Delta_1}{T-1} \\ p_2^D - p_2'(\lambda_2^D) \frac{\Delta_2}{T-2} \\ p_3^D - p_3'(\lambda_3^D) \frac{\Delta_3}{T-3} \end{bmatrix}, \quad \text{and} \quad \begin{bmatrix} p_{5,1} \\ p_{5,2} \\ p_{5,3} \end{bmatrix} = \begin{bmatrix} p_1^D - p_1'(\lambda_1^D) \frac{\Delta_1}{T-1} - p_1'(\lambda_1^D) \frac{\Delta_4}{T-4} \\ p_2^D - p_2'(\lambda_2^D) \frac{\Delta_2}{T-2} \\ p_3^D - p_3'(\lambda_3^D) \frac{\Delta_3}{T-3} \end{bmatrix}.$$

Thus, in this example, we have shown how to adjust the prices of three products $T/3$ times each instead of adjusting the price of one product T times using equivalent pricing.

Performance result. For any updating schedule $\gamma \in \Gamma(\gamma_B)$, let $\mathbf{Q} \in \mathcal{Q}(\gamma) := \{\{Q_t\}_{t=1}^T : Q_t \in \mathcal{Q}_t(\gamma)\}$ denote a sequence of transformation matrices that correspond to γ and let $R_{H, \gamma_B, \gamma}^{\mathbf{Q}}$ denote the resulting revenue. The following theorem provides a uniform performance bound for equivalent pricing control under any updating schedule γ that is equivalent to γ_B .

THEOREM 3. *There exist positive constants Ψ and $\bar{\Psi}$ independent of $\theta \geq 1$, the projection matrix H that selects \mathcal{B} , and the choice of updating schedules $\gamma_{\mathcal{B}}$ such that*

$$\begin{aligned} \sup_{\gamma \in \Gamma(\gamma_{\mathcal{B}})} \sup_{\mathbf{Q} \in \mathcal{Q}(\gamma)} \{J^{Det}(\theta) - \mathbf{E}[R_{H, \gamma_{\mathcal{B}}, \gamma}^{\mathbf{Q}}(\theta)]\} &\leq \Psi + \bar{\Psi} \sum_{j \in \mathcal{B}} \sum_{t=1}^{T(\theta)-1} \min\{1, \|\nabla p(\lambda^D)HA\|_2^2 U_1^j(T(\theta), t)\} \\ &+ \bar{\Psi} \sum_{j \in \mathcal{B}} \sum_{t=1}^{T(\theta)-1} \min\{1, \|\nabla p(\lambda^D)HA\|_2^2 U_2^j(T(\theta), t)\}, \end{aligned}$$

where the terms $U_1^j(T, t)$ and $U_2^j(T, t)$ are defined as in Theorem 1.

Observe that the bound in Theorem 3 is similar to the bound in Theorem 1. This shows that, for any schedule γ that is equivalent to the base schedule $\gamma_{\mathcal{B}}$, the seller can use equivalent pricing to guarantee the same asymptotic performance as the LPC under the base schedule $\gamma_{\mathcal{B}}$. This result provides the seller with an extra flexibility to manage his prices.

LPC with synchronous price adjustment of more than m products. Although the LPC discussed in Section 3 allows for arbitrary asynchronous price adjustment, it is restricted to adjust the price of *exactly* m products. Generalizing LPC to the case of *arbitrary* asynchronous price adjustment of more than m products is not a trivial task and beyond the scope of this paper. It is, however, possible to use equivalent pricing to develop a version of LPC that *synchronously* adjusts the prices of $k \geq m$ products. To illustrate how to use equivalent pricing to do synchronous price adjustment for $k \geq m$ products, consider the LPC discussed in Section 3 where the base is \mathcal{B} and $\gamma_j(\theta) = \gamma_1(\theta)$ for all $j \in \mathcal{B}$. Let \mathcal{G} denote a set of $k \geq m$ products that span the resource space (i.e., the set of products whose corresponding columns (by the same index) in $A\nabla\lambda(p^D)$ span \mathbb{R}^m). Since $\mathcal{G} \sim_{\mathcal{B}} \mathcal{B}$, we can construct a transformation matrix Q as described above and apply equivalent pricing with $Q_t = Q$ for all t . The resulting price update formula is given by

$$\hat{p}_t = p^D - \sum_{l=1}^{k_t^1} Q \nabla p(\lambda^D) H \frac{A \tilde{\Delta}_l^1}{T - t_l^1 + 1} = p^D - \sum_{l=1}^{k_t^1} \nabla p(\lambda^D) \tilde{H} \frac{A \tilde{\Delta}_l^1}{T - t_l^1 + 1},$$

where the second equality follows from the second part of Lemma 3 with \tilde{H} being a projection matrix such that the rows in $\nabla p(\lambda^D) \tilde{H}$ that correspond to products *not* in \mathcal{G} are zeros. Note that such pricing control has a practical implication: It provides the seller with an extra flexibility to trade off the negative impact of excessive price adjustment with the incremental improvement in revenue due to adjusting the price of more products. (See Theorem 2 and numerical experiments in Section 5 for further discussions.)

The difference between LPC and LRC. As briefly mentioned in Section 1, Jasin (2014) has developed a dynamic pricing heuristic which he calls *Linear Rate Correction* (LRC), and it adjusts

the price in period t using the update formula $\hat{p}_t = p\left(\lambda^D - H \sum_{s=1}^{t-1} \frac{A\Delta_s}{T-s}\right)$, where H is a projection matrix. To see the difference between LPC and LRC, first, note that, since $p(\cdot)$ is not always separable, the prices of *all* n products under LRC must be *simultaneously* updated at the same time. (Even if the projection matrix H is chosen to select a certain base, there is no guarantee that LRC will adjust the price of only the products in the base.) Thus, minimal price adjustment of only m products is, in general, not possible with LRC. Second, since $p(\cdot)$ is not always separable, there is no analog of the general LPC update formula for LRC. This means that neither asynchronous update nor equivalent pricing is possible with LRC, which may limit the applicability of LRC for real-world implementation (e.g., due to existing business constraints). Indeed, aside from the fact that LRC and LPC are examples of *linear control*⁶, they are close only in the special case where the prices of all products are updated at the same time (e.g., the synchronous *1-Periodic* schedule). In that special case, the update formula of LPC can be viewed as a linearization of the update formula of LRC. (The generic asynchronous LPC, however, is *not* a linearization of any form of LRC.)

5. Numerical Experiments

In this section, we run several experiments to illustrate the theoretical results in Sections 3 and 4 as well as to highlight the applicability of our heuristic in practice and its managerial implications. For our simulations, we use a multinomial logit demand with 10 products and 4 resources. (See Appendix C for detail.) We use $T = 1$ and $C_i = 0.1$ for each resource i . Note that, per our definition, the actual number of selling periods and initial inventory levels are given by θT and θC , respectively. For example, $\theta = 1,000$ corresponds to a problem instance with 1,000 selling periods and initial inventory levels equal to 100. We compare the expected revenue loss under different heuristics for a wide range of θ 's. In particular, since typical RM firms sell about 100-1,000 inventories per season (e.g., mid-size airplanes have about 100-500 seats and large-size hotels can easily have more than 1,000 rooms), we use θ ranging from 500-10,000.

We denote by *Static* the static price control developed in Gallego and van Ryzin (1997), and by *LRC* the linear rate control developed in Jasin (2014). As for our heuristics, we denote by *LPC-k* the LPC that simultaneously adjusts the prices of $k \geq m$ products in every period. (Recall that

⁶ Linear control has been widely studied in engineering (Ben-Tal et al. 2009) and finance (Calafiore 2009), and has only been recently studied in operations management (Bertsimas et al. 2010, Atar and Reiman 2012, Jasin 2014). In general, a linear control assumes the form of a baseline control plus a linear combination of past system perturbations. (This explains the forms of LRC and LPC.) While most existing literature on linear control focuses on finding a way to compute the optimal control parameters, our work explicitly constructs a particular form of linear control, which has certain desirable properties, and proposes a particular choice of parameters values that yields a strong performance guarantee. Needless to say, once the form is assumed, it may be possible to apply standard techniques in the literature to optimize the parameters of LPC. However, this is beyond the scope of this paper.

to ensure LPC adjusts at most k prices, we only need to find a proper transformation matrix. We select the transformation matrix following the proposed guideline in Section 4.) Correspondingly, we use $RSC-k$ to denote the heuristic that adjusts the prices of the same k products as in $LPC-k$ via exact re-optimization of DPP in every period, with an additional constraint that the prices of the unadjustable products remain the same as the static price. In addition to the said heuristics, we also test two simple modifications of $LPC-k$ that only adjust the same k prices and can improve the *non-asymptotic* performance of the vanilla $LPC-k$. The first one is a projection-based LPC where, in each period, we apply LPC update formula followed by a projection into $[(1 - \alpha\%)p^D, (1 + \alpha\%)p^D]$; we denote the resulting heuristic by $Pro\alpha-k$. If α is small, $Pro\alpha-k$ is very similar to static price control; if α is large, $Pro\alpha-k$ is very similar to $LPC-k$. Per our discussions in Section 3, since we are using static price as our benchmark, we would ideally like to have a heuristic whose price trajectory stays as close as possible to the static price. However, since demands are random, we must also allow some room for price adjustments to account for demand variability. This motivates the use of projection as a way to control the intensity of price fluctuation. The second modification of $LPC-k$ is a re-optimization-based LPC, denoted by $Hyb\beta-k$, where we re-optimize DPP at the first β updating times of the *2-Geometric* schedule and apply LPC in the remaining periods.

Experiment 1: Performance of LPC. Figure 1 illustrates the performance of $LPC-10$ and other existing heuristics. Consistent with our asymptotic results, $LPC-10$ performs much better than *Static*.⁷ Figure 1 also shows that $LPC-10$ performs slightly worse than LRC and $RSC-10$, which is not surprising because both LRC and $RSC-10$ are known to have a slightly stronger performance guarantee of $O(\log \theta)$ than LPC (Jasin 2014). We want to stress that although $RSC-10$ performs very well, it is also very time-consuming (see Table 1). In contrast, $LPC-10$ is computationally very fast. Admittedly, there is still a revenue gap between the “ideal but not implementable” $RSC-10$ and $LPC-10$. The question is whether there is a cheap way to improve the performance of $LPC-10$ without resorting to heavy frequent re-optimizations. It turns out that we can significantly narrow the gap between $RSC-10$ and $LPC-10$ by simple modifications of $LPC-10$. The first plot in Figure 2 shows that $Pro30-10$, which enforces the prices of LPC to fluctuate within a 30% band around the static price, can reduce the revenue loss gap by almost a half. This tells us that a simple projection can have a significant impact on revenue. (In general, we can also use product-dependent

⁷ It is interesting to note that not all linear price controls are guaranteed to perform well. For example, under 1-Periodic schedule, one intuitively appealing linear price control is $\hat{p}_t = p^D - \sum_{j=1}^m E^j \nabla p(\lambda^D) H \sum_{s=1}^{t-1} A \Delta_s$. Similar to LPC, this heuristic also adjusts prices to compensate for randomness in demand realizations. But, in contrast to LPC, this heuristic adjusts the price in a myopic manner; it attempts to fully correct the errors made in the previous period in the next period. Although this heuristic appears reasonable at first sight, our numerical experiments suggest that it is not even asymptotically optimal. This highlights that developing a linear price control that has strong performance is not a trivial task.

Figure 1 Revenue loss under different heuristics.

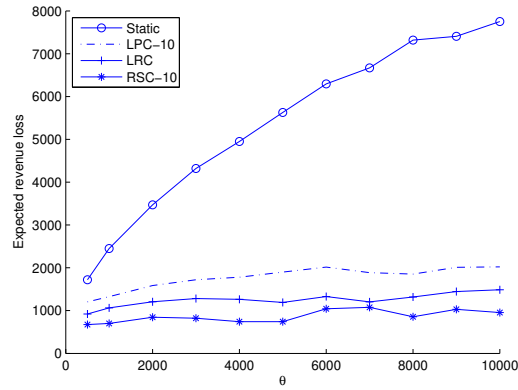
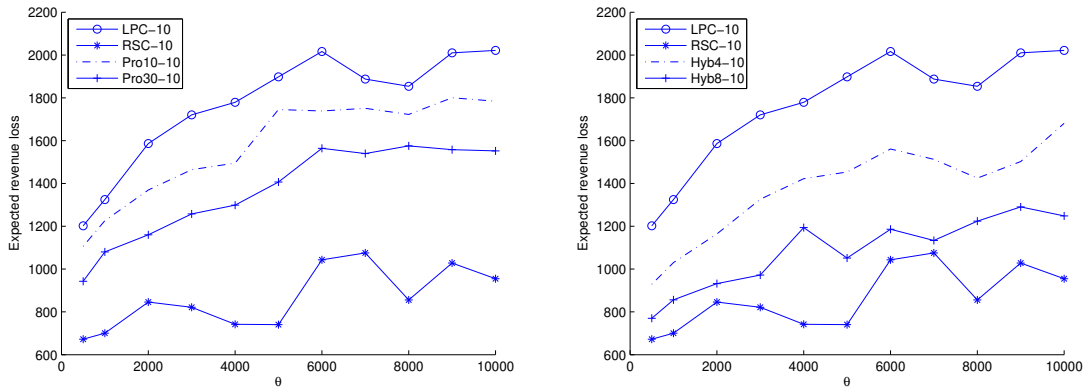


Figure 2 Improving LPC-10 using projection and occasional re-optimizations.

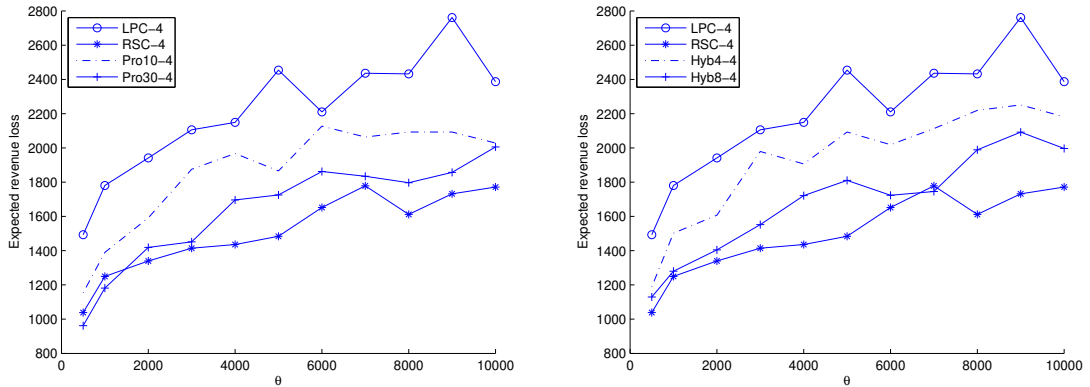


α parameters and optimize them by running an off-line Monte-Carlo optimization.) The second plot in Figure 2 further shows that *Hyb8-10*, which combines LPC with *only* 8 optimizations, can reduce the revenue loss gap by more than 75%. This is fairly impressive considering the fact that, even for small $\theta = 500$, *RSC-10* already requires 500 re-optimizations. It highlights the versatility of LPC for practical implementation; in particular, we can use LPC in combination with occasional re-optimizations in the case where frequent re-optimizations is clearly not feasible.

Table 1 Typical running time (in milliseconds) for a single simulation for selected heuristics.

θ	RSC-10	LPC-10	Hyb8-10
500	8305.0	13.3	209.7
5000	87552.4	86.2	212.3

Experiment 2: Minimal price adjustment. In this experiment, we test the minimal adjustment property discussed in Section 3. The plots in Figure 3 show the comparison between *LPC-4* and *RSC-4*, as well as the two types of modified LPC with the same projection matrix as *LPC-4*. All these heuristics adjust the prices of the same $m = 4$ products. (Note that *LRC* cannot be included

Figure 3 Improving LPC-4 using projection and occasional re-optimizations.

in this comparison because it cannot adjust prices of fewer than $n = 10$ products.) Similar to experiment 1, while $RSC-4$ performs very well, it requires a lot of re-optimizations, which may not be feasible in practice. The two simple modifications of $LPC-4$, $Pro30-4$ and $Hyb8-4$, which are computationally much cheaper, can attain a similar performance as $RSC-4$.

At the end of Section 3, we discussed the impact of increasing the number of adjustable products on revenue performance. Figure 4 illustrates our theoretical results. (See also Table 2 in Appendix C.) The first plot in Figure 4 shows that, in comparison to *Static* that adjusts no prices at all, allowing $m = 4$ adjustable products yields a significant reduction in revenue loss. This is due to the minimal adjustment property of LPC . Beyond the initial four products, although allowing more adjustable products further decreases the revenue loss, its incremental benefit becomes much smaller. In particular, the plot shows that the impact of allowing two additional adjustable products (see the gap between $LPC-4$ and $LPC-6$) captures almost half of the benefit of allowing six more adjustable products (see the gap between $LPC-4$ and $LPC-10$). We observe the same phenomenon in the second plot in Figure 4 for $Hyb8$ heuristics. This suggests that the managerial insights drawn from Theorem 2 still hold in network setting: If the seller wishes to adjust the prices of more than m products to increase revenue, then adjusting a few more products is sufficient to capture pretty much all the potential benefit of adjusting all products.

Experiment 3: Equivalent pricing with business constraints. In this experiment, we study a case where the seller has additional constraints on when and what prices to adjust. We assume that (1) the prices of products 5, 8 and 9 cannot be adjusted, (2) the prices of products 2, 3, 4 can *only* be adjusted in the second half of the selling season, and (3) the prices of products 6, 7, 10 can only be adjusted in the first half of the selling season. These are plausible constraints motivated by practical applications. For example, products 5, 8 and 9 can be viewed as corporate rate rooms that cannot be adjusted over time. Products 2-4 and 6, 7, 10 can be viewed as special

Figure 4 Revenue impact of the number of adjustable products for LPC and Hyb8.

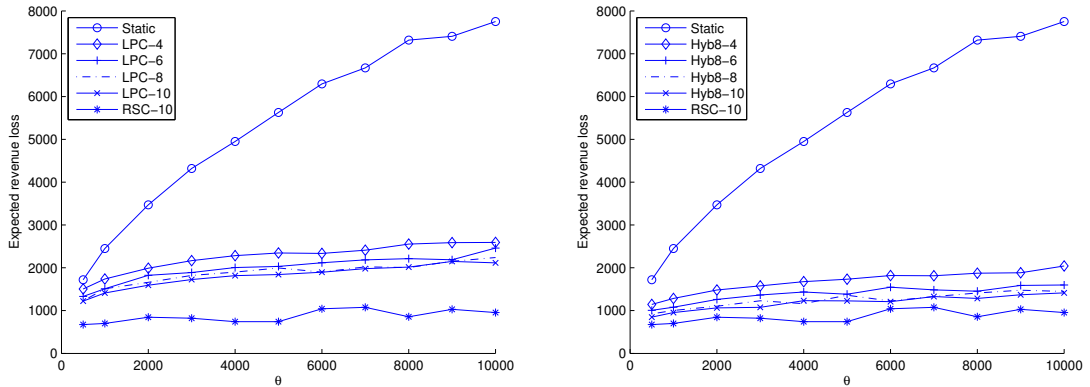
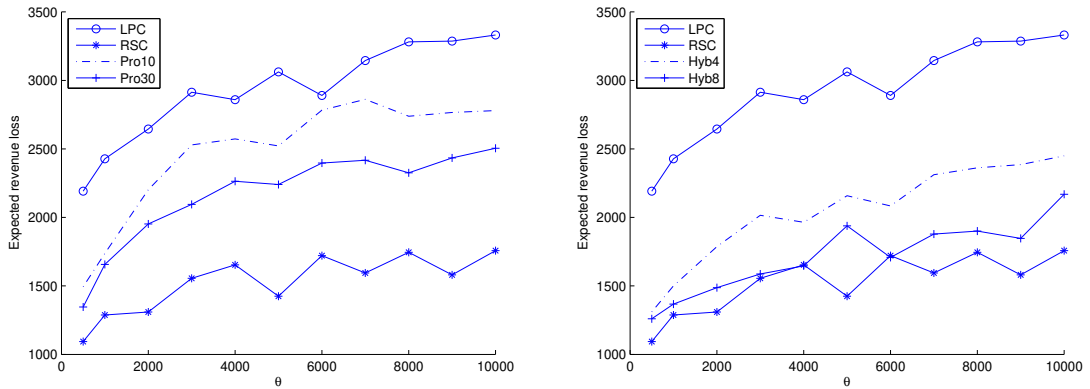


Figure 5 Improving LPC using projection and occasional re-optimizations.



rate rooms for certain events (e.g., conference) whose prices cannot be adjusted in a certain time window. Based on our discussions in Section 4, LPC can be automatically adapted to this setting via equivalent pricing with an original base of $\mathcal{B} = \{1, 2, 3, 4\}$; we denote this heuristic simply as *LPC*. Similar to previous experiments, we can apply re-optimized static price control with the additional constraints that certain prices cannot be adjusted in particular periods; we denote the resulting heuristic simply as *RSC*. It is also possible to use the modified LPC, which we denote as *Pro α* and *Hyb k* , accordingly. Figure 5 shows that simple modifications of LPC, which is computationally easy, can attain a similar performance as *RSC* which requires frequent re-optimizations and may not be implementable in practice. This highlights the versatility of LPC for practical implementation in the presence of business constraints.

6. Closing Remarks

In this paper, we consider a standard dynamic pricing problem and propose a new family of pricing heuristics, which we call LPC. We show that LPC provides a strong improvement over static pricing: The revenue loss is reduced from $O(\sqrt{\theta})$ to $O(\log^2 \theta)$. In addition, it also has desirable features

that can be used to address practical concerns. First, LPC only requires a single optimization and can be implemented in real-time, which makes it useful for solving large-scale problems where other computationally intensive heuristics are not viable. Second, LPC guarantees a strong revenue performance by adjusting the price of a few “important” products infrequently. This helps address the issue of acceptability of dynamic pricing in the eyes of customers due to excessive price adjustments. Third, LPC allows the seller to maintain an equivalent revenue performance via adjusting the prices of other products. This not only can be used to further reduce the number of required price changes per product, but also provides an extra flexibility for the sellers to manage his prices in the presence of various business constraints. Our simulation results show that LPC not only has a good theoretical performance but also works well numerically. Furthermore, its performance can be further improved by simple modifications such as projection and occasional re-optimizations. To conclude, we believe that our work provides novel managerial insights that make dynamic pricing more applicable and practically appealing for real-world implementation.

Acknowledgments

The authors thank Noah Gans and Assaf Zeevi (the department editors), the anonymous associate editor, and the three anonymous referees for their comments and suggestions that have improved this paper.

Appendix A: Proofs of Section 3

A.1. Proof of Lemma 1

Throughout, we use superscript j and subscript i to indicate the j^{th} column and the i^{th} row of a matrix respectively. Define $\bar{A} := A\nabla\lambda(p^D)$. By definition, a base must span the resource space which has rank m , so it must contain at least m products. Without loss of generality, suppose that $\mathcal{B} = \{1, 2, \dots, m\}$. The matrix $[\bar{A}^1 \bar{A}^2 \dots \bar{A}^m]$ is invertible and we can define its inverse $\bar{U} = [\bar{A}^1 \bar{A}^2 \dots \bar{A}^m]^{-1}$. We now construct an n by m matrix U as follows: $U_i = \bar{U}_i$ for $i = 1, \dots, m$ and $U_i = 0$ otherwise. Observe that $A\nabla\lambda(p^D)U = \bar{A}U = I$. Let $H = \nabla\lambda(p^D)U$. Since only the first m rows of $\nabla p(\lambda^D)H = U$ are non-zeros and $\mathcal{B} = \{1, 2, \dots, m\}$, we conclude that H selects \mathcal{B} . To show the uniqueness of H , we use contradiction. Suppose not, then we have at least two n by m matrices $H \neq \tilde{H}$ that select \mathcal{B} . Let $U = \nabla p(\lambda^D)H, \tilde{U} = \nabla p(\lambda^D)\tilde{H}$. Since $\nabla p(\lambda^D)$ is full rank and $H - \tilde{H} \neq 0$, we conclude that $U - \tilde{U} \neq 0$. Since the last $n - m$ rows of U and \tilde{U} are all zero vectors, we conclude that $U_i \neq \tilde{U}_i$ for some $1 \leq i \leq m$ which contradicts with the uniqueness of the inverse of \bar{A} .

A.2. Proof of Lemma 2

We will prove a more general result of picking the best k prices. For any $v \in \mathbb{R}^n$ define $\|v\|_0 := |\{i : v_i \neq 0\}|$. Let $a = A'$, $x = \nabla p(\lambda^D)H$, $b = (A\nabla\lambda(p^D))'$. Since $m = 1$, a, x, b are all vectors in \mathbb{R}^n . The optimization problem $\min_H \{\|\nabla p(\lambda^D)HA\|_2 : AH = 1, \|\nabla p(\lambda^D)H\|_0 \leq k\}$ is equivalent to $\min_x \{\|xa'\|_2^2 : b'x = 1, \|x\|_0 \leq k\}$. Since $xa'ax'$ is a rank one matrix, its maximum eigenvalue is just its trace. So $\|xa'\|_2^2 = \text{tr}(xa'ax') = \text{tr}(x'xa'a) = \|x\|_2^2 \|a\|_2^2$. Note also that the equality constraint is equivalent to $\|b\|_2 \|x\|_2 \cos(b, x) = 1$, where $\cos(b, x)$ is the cosine of the angle between vectors b and x . Therefore, as long as $\|x\|_2 = 1/(\|b\|_2 \cos(b, x))$, the equality constraint can be satisfied. So the problem becomes $\min_x \{\|a\|_2^2 \|b\|_2^{-2} \cos^{-2}(b, x) : \|x\|_0 \leq k\}$. Let $b_{(i)}$ denote the i^{th} largest element in absolute value in b , then the optimal solution x^* is parallel with a vector b^k which has the exact same elements as b in the k largest elements in absolute values but zeros in other elements. The optimal objective value is $\|a\|_2^2 (\sum_{i=1}^k b_{(i)}^2)^{-1}$.

A.3. Proof of Corollary 1

We compute, part by part, the bound in Theorem 1 under periodic price update schedule. Without loss of generality, we assume that $T = 1$. For notational clarity, we suppress the dependence on θ whenever there is no confusion. We start with the summation over $U_1^j(\theta, t)$. First of all, we have $\sum_{j=1}^m \sum_{t=1}^{\theta-1} \min \left\{ 1, \|\nabla p(\lambda^D)HA\|_2^2 U_1(\theta, t) \right\} \leq \max\{1, \|\nabla p(\lambda^D)HA\|_2^2\} \sum_{j=1}^m \sum_{t=1}^{\theta-1} \min \{1, U_1^j(\theta, t)\}$. We bound the summation after the inequality as follows: $\sum_{j=1}^m \sum_{t=1}^{\theta-1} \min \{1, U_1^j(\theta, t)\} = m \sum_{t=1}^{\theta-1} \min \left\{ 1, \frac{t-hk_t}{(\theta-t)^2} + \sum_{l=1}^{k_t} \frac{t_l-t_{l-1}}{(\theta-t_l+1)^2} \right\} \leq m \sum_{t=1}^{\theta-1} \min \left\{ 1, \frac{t-hk_t}{(\theta-t)^2} + \frac{t_{k_t}-t_{k_t-1}}{(\theta-t_{k_t}+1)^2} + \int_1^{t_{k_t}} \frac{1}{(\theta-x+1)^2} dx \right\} \leq m \sum_{t=1}^{\theta-1} \min \left\{ 1, \frac{2h}{(\theta-t)^2} + \frac{1}{\theta-t} \right\}$. The first equality follows since we update the price of the m products at the same time. The first inequality is the integral approximation and the last inequality follows from the fact that $0 \leq t - k_t h \leq h$. Now define $t^* = \lceil \theta - \sqrt{h} \rceil$. We make further approximation of the inequality above by breaking down the summation over t into two parts, before and after t^* : $\sum_{j=1}^m \sum_{t=1}^{\theta-1} \min \{1, U_1(\theta, t)\} \leq m \left[\int_1^{t^*} \frac{2h}{(\theta-x)^2} dx + \int_1^{t^*} \frac{1}{\theta-x} dx + \theta - t^* \right] \leq m \left(\frac{2h}{\theta-t^*} + \log \left(\frac{\theta-1}{\theta-t^*} \right) + \theta - t^* \right) \leq m \left(1 + 3\sqrt{h} + \log \theta \right)$ where the first inequality follows from the integration approximation and the third inequality follows from the fact that $1 \leq \sqrt{h} \leq \theta - t^* \leq \sqrt{h} + 1$. Now we compute the summation over $U_2^j(\theta, t)$. Similarly, it suffices to bound the following: $\sum_{j=1}^m \sum_{t=1}^{\theta-1} \min \{1, U_2^j(\theta, t)\} \leq m \sum_{t=1}^{\theta-1} \min \left\{ 1, \frac{1}{\theta-t} \sum_{s=1}^t \left(\frac{h}{(\theta-s)^2} + \frac{1}{\theta-s} \right) \right\}$. Again, we break the summation into two parts and use integral approximation: $\sum_{j=1}^m \sum_{t=1}^{\theta-1} \min \{1, U_2(\theta, t)\} \leq m \left[\sum_{t=1}^{t^*-1} \frac{1}{\theta-t} \int_1^{t+1} \left(\frac{h}{(\theta-x)^2} + \frac{1}{\theta-x} \right) dx + \theta - t^* \right] \leq m \left[\sum_{t=1}^{t^*-1} \left(\frac{h}{(\theta-t-1)^2} + \frac{1}{\theta-t-1} \log \left(\frac{\theta-1}{\theta-t-1} \right) \right) + \theta - t^* \right] \leq m \left[\frac{h}{(\theta-t^*)^2} + \int_1^{t^*-1} \frac{h}{(\theta-x-1)^2} dx + \sum_{t=1}^{t^*-1} \frac{1}{\theta-t-1} \log \left(\frac{\theta-1}{\theta-t-1} \right) + \theta - t^* \right] \leq m \left(2 + 2\sqrt{h} + \log \theta + \log^2 \theta \right)$, where the last inequality holds because $\sum_{t=1}^{t^*-1} \frac{1}{\theta-t-1} \log \left(\frac{\theta-1}{\theta-t-1} \right) \leq \frac{\log \left(\frac{\theta-1}{\theta-t^*} \right)}{\theta-t^*} + \int_1^{t^*-1} \frac{\log \left(\frac{\theta-1}{\theta-t-1} \right)}{\theta-t-1} dt \leq \log \theta + \log^2 \theta$.

A.4. Proof of Corollary 2

We assume without loss of generality that $T = 1$ and suppress the dependence on θ for brevity. Note that $K(\theta)$ is well-defined since $\sum_{s=1}^k s^\alpha$ is strictly increasing in k and is unbounded as $k \rightarrow \infty$ for all $\alpha \geq 1$. Since $\theta > \sum_{s=1}^K s^\alpha \geq K^{\alpha+1}/(\alpha+1)$, we have $K \leq ((\alpha+1)\theta)^{1/(\alpha+1)}$. We now analyze the performance bound. We first derive bound for the summation over $U_1^j(\theta, t)$. Similar to the proof of Corollary 1, it suffices to bound the following: $\sum_{j=1}^m \sum_{t=1}^{\theta-1} \min \{1, U_1^j(\theta, t)\}$. By definition, for $1 \leq l \leq K$, we have $\theta - t_l + 1 \geq \sum_{s=1}^{K-l+1} s^\alpha \geq \frac{(K-l+1)^{\alpha+1}}{\alpha+1} \geq \frac{(K-l+2)^{\alpha+1}}{2^{\alpha+1}(\alpha+1)}$. In addition, we also have that for $2 \leq l \leq K$, $t_l - t_{l-1} \leq (K-l+2)^\alpha + 1 \leq 2(K-l+2)^\alpha$, and for $l = 1$, $t_1 - t_0 \leq \theta + 1 - \sum_{s=1}^K s^\alpha - 1 \leq (K+1)^\alpha \leq 2(K+1)^\alpha$. Then, for $t < \theta - 1$, since $k_t < K$, we have $U_1^j(\theta, t) \leq \sum_{l=1}^{k_t+1} \frac{t_l-t_{l-1}}{(\theta-t_l+1)^2} \leq \sum_{l=1}^{k_t+1} \frac{2(K-l+2)^\alpha (\alpha+1)^2 2^{2\alpha+2}}{(K-l+2)^{2\alpha+2}} = (\alpha+1)^2 2^{2\alpha+3} \sum_{l=1}^{k_t+1} \frac{1}{(K-l+2)^{\alpha+2}}$. Hence, $\sum_{j=1}^m \sum_{t=1}^{\theta-1} \min \{1, U_1^j(\theta, t)\} \leq \sum_{j=1}^m \left(1 + \sum_{t=1}^{\theta-2} U_1^j(\theta, t) \right) \leq m + m \sum_{l=1}^K 2(K-l+2)^\alpha \sum_{s=1}^l \frac{(\alpha+1)^2 2^{2\alpha+3}}{(K-s+2)^{\alpha+2}} \leq m + m \sum_{l=1}^K 2(K-l+2)^\alpha \int_1^{l+1} \frac{(\alpha+1)^2 2^{2\alpha+3}}{(K-s+2)^{\alpha+2}} ds \leq m + m(\alpha+1) 2^{2\alpha+4} \sum_{l=1}^K \frac{(K-l+2)^\alpha}{(K-l+1)^{\alpha+1}} \leq m + m(\alpha+1) 2^{3\alpha+4} \sum_{l=1}^K \frac{1}{(K-l+1)} \leq m + m(\alpha+1) 2^{3\alpha+4} \log K$. Since $K \leq ((\alpha+1)\theta)^{\frac{1}{\alpha+1}}$, $\sum_{j=1}^m \sum_{t=1}^{\theta-1} \min \{1, U_1^j(\theta, t)\} \leq m(1 + 2^{3\alpha+4} \log(\alpha+1) + 2^{3\alpha+4} \log \theta)$. As for the summation over $U_2^j(\theta, t)$, we have $\sum_{j=1}^m \sum_{t=1}^{\theta-1} \min \{1, U_2^j(\theta, t)\} \leq m + \sum_{j=1}^m \sum_{t=1}^{\theta-2} U_2^j(\theta, t) \leq m + \sum_{j=1}^m \sum_{t=1}^{\theta-2} \frac{1}{\theta-t} \sum_{s=1}^t U_1^j(\theta, s) \leq m + \sum_{j=1}^m \sum_{t=1}^{\theta-2} \frac{1}{\theta-t} \sum_{s=1}^{\theta-2} U_1^j(\theta, s) \leq m(1 + 2^{3\alpha+4} \log(\alpha+1) \log \theta + 2^{3\alpha+4} \log^2 \theta)$.

A.5. Proof of Corollary 3

We assume without loss of generality that $T = 1$ and suppress the dependence on θ for brevity. We first show that $K \leq 1 + \log_\beta \theta$. Note that since $\{t_l\}$ are strictly increasing integers, so K is well defined and by definition of t_l we have $t_{K-1} \leq \theta - 1$. By definition, we have $t_l \geq \lceil (\beta-1)\theta + t_{l-1} \rceil / \beta$, so $\theta - t_l \leq (\theta - t_{l-1}) / \beta \leq \theta / \beta^l$. Therefore, $\theta - 1 \geq t_{K-1} > \theta - \theta / \beta^{K-1}$ which implies that $K \leq 1 + \log_\beta \theta$. We now analyze the performance bound. By definition, we have $t_l \leq \lceil (\beta-1)\theta + t_{l-1} \rceil / \beta + 1$, so we have the following useful bound which will be used a couple of times later: for $l \leq K$,
 (*) $\frac{t_l-t_{l-1}}{\theta-t_l+1} \leq \frac{\lceil (\beta-1)\theta + t_{l-1} \rceil / \beta + 1 - t_{l-1}}{\theta - \lceil (\beta-1)\theta + t_{l-1} \rceil / \beta + 1} = \frac{(\beta-1)(\theta-t_{l-1}+1)+1}{\theta-t_{l-1}} \leq 2\beta - 1$. We derive an upper bound for the summation

over $U_1^j(\theta, t)$ first. Similar to the proof of Corollary 1, it suffices to bound the following: $\sum_{j=1}^m \sum_{t=1}^{\theta-1} \min\{1, U_1^j(\theta, t)\} \leq m \sum_{t=1}^{\theta-1} \left(\frac{t-t_{k_t}+1}{(\theta-t)^2} + \sum_{l=1}^{k_t} \frac{2\beta-1}{\theta-t_l+1} \right) \leq m \sum_{t=1}^{\theta-1} \left(\frac{t_{k_t+1}-t_{k_t}+1}{(\theta-t)(\theta-t_{k_t+1}+1)} + \sum_{l=1}^{k_t} \frac{2\beta-1}{\theta-t_l+1} \right) \leq m \sum_{t=1}^{\theta-1} \left(\frac{2\beta-1}{(\theta-t)} + \sum_{l=1}^{k_t} \frac{2\beta-1}{\theta-t_l+1} \right) \leq m(2\beta-1) \left(\log \theta + \sum_{t=1}^{\theta-1} \sum_{l=1}^{k_t} \frac{1}{\theta-t_l+1} \right)$, where the first and the third inequalities follow from (\star) . Note that $\sum_{t=1}^{\theta-1} \sum_{l=1}^{k_t} \frac{1}{\theta-t_l+1} = \sum_{j=0}^{K-1} \sum_{t=t_j}^{t_{j+1}-1} \sum_{l=1}^j \frac{1}{\theta-t_l+1} = \sum_{j=0}^{K-1} \sum_{l=1}^j \frac{t_{j+1}-t_j}{\theta-t_l+1} = \sum_{j=1}^{K-1} \frac{\theta-t_j}{\theta-t_j+1} \leq K-1 \leq \log_{\beta} \theta$. Hence, we have $\sum_{j=1}^m \sum_{t=1}^{\theta-1} \min\{1, U_1^j(\theta, t)\} \leq m(2\beta-1) (\log \theta + \log_{\beta} \theta)$. Now we approximate the summation over $U_2^j(\theta, t)$ as follows: $\sum_{j=1}^m \sum_{t=1}^{\theta-1} \min\{1, U_2^j(\theta, t)\} \leq m \sum_{t=1}^{\theta-1} \frac{1}{\theta-t} \sum_{s=1}^{\theta-1} \sum_{l=1}^{k_s} \frac{2\beta-1}{\theta-t_l+1} \leq m(2\beta-1) \log \theta \log_{\beta} \theta$.

A.6. Proof of Theorem 2

We use a slight modification of LPC with synchronous 1-Periodic Schedule as follows: follow the LPC heuristic but uses $\hat{p}_t = p^D - \nabla p(\lambda^D) H \sum_{s=1}^{t-1} \frac{A \Delta_s}{T-s}$, where H is a projection matrix. Call this heuristic π_H . Pick an H that satisfies $\|\nabla p(\lambda^D) H\|_0 \leq k$. Then we have $\pi_H \in \Pi_k$. Following a similar argument as Theorem 1, there exist positive constants Ψ and $\hat{\Psi}$ such that $J^{Det} - \mathbf{E}[R_{\pi_H}(\theta)] \leq \Psi + \hat{\Psi} \|\nabla p(\lambda^D) H A\|_2^2 \log^2 \theta$. By the proof of Lemma 2, if we minimize $\|\nabla p(\lambda^D) H A\|_2$ subject to $AH = 1$ and $\|\nabla p(\lambda^D) H\|_0 \leq k$, the optimal projection matrix H^* attains $\|\nabla p(\lambda^D) H^* A\|_2^2 = \|a\|_2^2 (\sum_{i=1}^k b_{(i)}^2)^{-1}$. Therefore, $\min_{\pi \in \Pi_k} \{J^{Det} - \mathbf{E}[R_{\pi}(\theta)]\} \leq J^{Det} - \mathbf{E}[R_{\pi_{H^*}}(\theta)] \leq \Psi + \bar{\Psi} (\sum_{i=1}^k b_{(i)}^2)^{-1} \log^2 \theta$, where $\bar{\Psi} = \hat{\Psi} \|a\|_2^2$.

Appendix B: Proofs of Section 4

B.1. Proof of Lemma 3

By the construction of $Q(Y_t, \mathcal{G}_t, \mathcal{S}_t)$, it is straightforward to verify that $A \nabla \lambda(p^D) Q_t E^{\mathcal{B}} = A \nabla \lambda(p^D) E^{\mathcal{B}} = A \nabla \lambda(p^D) E^{\mathcal{G}_t \cup (\mathcal{B} - \mathcal{S}_t)} Q_t$ holds. (See Figure 6 for an illustration.) This proves (1). For (2), construct $H_t := \nabla \lambda(p^D) Q_t \nabla p(\lambda^D) H$. Note that H_t is a projection matrix since $A H_t = A \nabla \lambda(p^D) Q_t \nabla p(\lambda^D) H = A \nabla \lambda(p^D) Q_t E^{\mathcal{B}} \nabla p(\lambda^D) H = A \nabla \lambda(p^D) E^{\mathcal{B}} \nabla p(\lambda^D) H = A \nabla \lambda(p^D) \nabla p(\lambda^D) H = I$ where the second and the fourth equality follows by the fact that only the first m rows of $\nabla p(\lambda^D) H$ are nonzero. Note also that $\nabla p(\lambda^D) H_t = Q_t \nabla p(\lambda^D) H$. So, to verify that rows in $\nabla p(\lambda^D) H_t$ that correspond to products not in $\mathcal{G}_t \cup (\mathcal{B} - \mathcal{S}_t)$ are zeros, we only need to verify it for $Q_t \nabla p(\lambda^D) H$. For any $j \notin \mathcal{G}_t \cup (\mathcal{B} - \mathcal{S}_t)$, either (a) $j \in \mathcal{S}_t$ and $j \notin \mathcal{G}_t$, or (b) $j \notin \mathcal{B} \cup \mathcal{G}_t$. In case (a), the result holds since the j^{th} row of Q_t is a zero vector. In case (b), the only nonzero element in row j of Q_t is the j^{th} element, but the j^{th} row of $\nabla p(\lambda^D) H$ is a zero vector. This proves (2). Finally, since the only nonzero elements in $Q_t E^{\mathcal{S}_t}$ are in the submatrix consisting of rows in \mathcal{G}_t and columns in \mathcal{S}_t , we conclude that the rows in $Q_t E^{\mathcal{S}_t} \nabla p(\lambda^D) H$ that correspond to products not in \mathcal{G}_t are zero vectors. This completes the proof of (3).

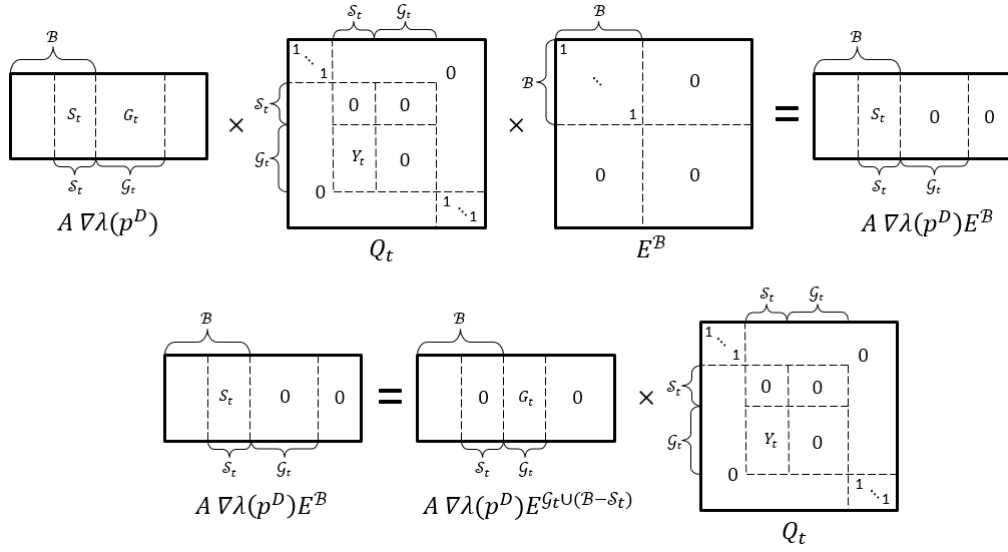
B.2. Proof Sketch of Theorem 3

The proof of Theorem 3 follows the same outline of the proof for Theorem 1 (see the Online Supplement) with three nontrivial twists.

1) Resource Correction Equivalence. We first show that in terms of error correction, equivalent pricing is “equivalent” to LPC. In particular, let $\tilde{\epsilon}_t = \sum_{j=1}^m \sum_{l=1}^{k_t^j} Q_{t_l^j} E^j \nabla p(\lambda^D) H \frac{A \Delta_l^j}{T-t_l^j+1}$. For simplicity, disregard the second order term of Taylor expansion of λ_t , then we have exactly the same capacity error below as (2) in the proof of Theorem 1: $A \lambda_t - A \lambda^D = -A \nabla \lambda(p^D) \tilde{\epsilon}_t = -\sum_{j=1}^m \sum_{l=1}^{k_t^j} A \nabla \lambda(p^D) Q_{t_l^j} E^j \nabla p(\lambda^D) H \frac{A \Delta_l^j}{T-t_l^j+1} = -\sum_{j=1}^m \sum_{l=1}^{k_t^j} A \nabla \lambda(p^D) E^j \nabla p(\lambda^D) H \frac{A \Delta_l^j}{T-t_l^j+1} = -M^{-1} \sum_{j=1}^m \sum_{l=1}^{k_t^j} \frac{\tilde{\xi}_l^j \mathbf{e}_j}{T-t_l^j+1}$, where the third equality follows by Lemma 3 part (1).

2) A uniform upper bound of $\|Q_t\|_2^2$. For any set $\mathcal{I} \subseteq \{1, \dots, n\}$ and any m by n matrix M , let $M^{\mathcal{I}}$ denote the submatrix of M that consists of columns $j \in \mathcal{I}$. Then, for any pair of $\mathcal{I}_1 \subseteq \{1, \dots, n\}, \mathcal{I}_2 \subseteq \{1, \dots, n\}$ we write $\mathcal{I}_1 \approx \mathcal{I}_2$ if $(A \nabla \lambda(p^D))^{\mathcal{I}_1}$ and $(A \nabla \lambda(p^D))^{\mathcal{I}_2}$ expands the same subspace of \mathbb{R}^m and $|\mathcal{I}_1| = |\mathcal{I}_2|$. Note that $\mathcal{I}_1 \approx \mathcal{I}_2$

Figure 6 Illustration of Lemma 3 part (1)



implies that there exists a unique $|\mathcal{I}_1|$ by $|\mathcal{I}_1|$ invertible matrix $Y(\mathcal{I}_1, \mathcal{I}_2)$ such that $M^{\mathcal{I}_1} = M^{\mathcal{I}_2} Y(\mathcal{I}_1, \mathcal{I}_2)$. Let $\bar{Q} := \sup\{\|Q(Y(\mathcal{I}_1, \mathcal{I}_2), \mathcal{I}_2, \mathcal{I}_1)\|_2^2 : \mathcal{I}_1 \approx \mathcal{I}_2\}$. Note that \bar{Q} is bounded because there are only finite pairs of $\mathcal{I}_1, \mathcal{I}_2$ that satisfy $\mathcal{I}_1 \approx \mathcal{I}_2$. In addition, \bar{Q} only depends on $A \nabla \lambda(p^D)$. We now claim that for any $B, \gamma_B, \gamma \in \Gamma(\gamma_B), t$ and $Q \in \mathcal{Q}_t(\gamma)$, $\|Q\|_2^2 \leq \bar{Q}$. This is because, by definition, for any set of products $S_t \subseteq B$ being adjusted in period t under γ_B , and any set of G_t being adjusted in period t under schedule $\gamma \in \Gamma(\gamma_B)$, G_t is equivalent to S_t and there exists a set $G'_t \subseteq G_t$ such that $G'_t \approx S_t$. Without loss of generality, assume G'_t corresponds to the first $|S_t|$ elements in G_t . Then construct a $|G_t|$ by $|S_t|$ matrix Y_t whose submatrix with rows in G'_t and columns in S_t equal $Y(S_t, G'_t)$ and remaining elements equal 0. Then, by optimality, we have that for any $Q_t \in \mathcal{Q}_t(\gamma)$, $\|Q_t\|_2^2 \leq \|Q(Y_t, G_t, S_t)\|_2^2 \leq \bar{Q}$.

3) Bounding $\mathbf{E}[T - \tau]$. Note that in the proof of Theorem 1, we have (E2) : $\psi > \frac{\bar{v}}{T-t} \sum_{s=1}^t \epsilon'_s \epsilon_t$ and (E3) : $\psi > \epsilon'_t \epsilon_t$. Now, because the price deviation becomes $\tilde{\epsilon}_t$, we redefine (E2) and (E3) by replacing ϵ_t by $\tilde{\epsilon}_t$. Then the rest of the argument in the proof of Theorem 1 holds except that the argument and the bound in Lemma EC.3 in the Online Supplement will be slightly different. In particular, the bounding of τ_2, τ_3 requires extra care. Let $q_{t_l^j}(j', j)$ denote the j' -th row j -th column element of the matrix $Q_{t_l^j}$. Then, the bound in STEP 2 of Lemma EC.3 in the Online Supplement becomes:

$$\begin{aligned} \Pr(\tau_2 \leq t) &= \Pr\left(\max_{v \leq t} \frac{\bar{v}}{T-v} \sum_{s=1}^v \tilde{\epsilon}'_s \tilde{\epsilon}_s \geq \psi\right) \leq \Pr\left(\frac{\bar{v}}{T-t} \sum_{s=1}^t \tilde{\epsilon}'_s \tilde{\epsilon}_s \geq \psi\right) \\ &= \Pr\left(\frac{\bar{v}}{T-t} \sum_{s=1}^t \left[\sum_{j'=1}^n \left(\sum_{j=1}^m \sum_{l=1}^{k_s^j} \frac{q_{t_l^j}(j', j) \tilde{\xi}_l^j}{T-t_l^j+1}\right)^2\right] \geq \psi\right) \leq \min\left\{1, \frac{\bar{v}}{\psi(T-t)} \sum_{s=1}^t \mathbf{E}\left[\sum_{j'=1}^n \left(\sum_{j=1}^m \sum_{l=1}^{k_s^j} \frac{q_{t_l^j}(j', j) \tilde{\xi}_l^j}{T-t_l^j+1}\right)^2\right]\right\} \end{aligned}$$

Note that we have,

$$\begin{aligned} &\mathbf{E}\left[\sum_{j'=1}^n \left(\sum_{j=1}^m \sum_{l=1}^{k_s^j} \frac{q_{t_l^j}(j', j) \tilde{\xi}_l^j}{T-t_l^j+1}\right)^2\right] \leq \mathbf{E}\left[\sum_{j'=1}^n \left(\sum_{j=1}^m \left|\sum_{l=1}^{k_s^j} \frac{q_{t_l^j}(j', j) \tilde{\xi}_l^j}{T-t_l^j+1}\right|\right)^2\right] \leq \mathbf{E}\left[\sum_{j'=1}^n m \sum_{j=1}^m \left(\sum_{l=1}^{k_s^j} \frac{q_{t_l^j}(j', j) \tilde{\xi}_l^j}{T-t_l^j+1}\right)^2\right] \\ &= m \sum_{j=1}^m \sum_{l=1}^{k_s^j} \mathbf{E}\left[\sum_{j'=1}^n \frac{q_{t_l^j}^2(j', j) (\tilde{\xi}_l^j)^2}{(T-t_l^j+1)^2}\right] = m \sum_{j=1}^m \sum_{l=1}^{k_s^j} \mathbf{E}\left[\frac{\|Q_{t_l^j} E^j \nabla p(\lambda^D) H A \tilde{\Delta}_l^j\|_2^2}{(T-t_l^j+1)^2}\right] \\ &\leq m \sum_{j=1}^m \sum_{l=1}^{k_s^j} \frac{\|Q_{t_l^j} E^j \nabla p(\lambda^D) H A\|_2^2 \mathbf{E}\left[\|\tilde{\Delta}_l^j\|_2^2\right]}{(T-t_l^j+1)^2} \leq m \sum_{j=1}^m \sum_{l=1}^{k_s^j} \bar{Q} \|\nabla p(\lambda^D) H A\|_2^2 \frac{t_l^j - t_{l-1}^j}{(T-t_l^j+1)^2} \end{aligned}$$

where the first equality follows because $\forall l \neq l', \mathbf{E}[\tilde{\xi}_l^j \tilde{\xi}_{l'}^j] = 0$ by the martingale property. With the inequality above, we conclude that: $\sum_{t=1}^{T-1} \Pr(\tau_2 \leq t) \leq \sum_{j=1}^m \sum_{t=1}^T \min \left\{ 1, \frac{m\bar{Q}}{\psi} \|\nabla p(\lambda^D) HA\|_2^2 U_2^j(T, t) \right\} \leq \max \left\{ 1, \frac{m\bar{Q}}{\psi} \right\} \sum_{j=1}^m \sum_{t=1}^T \min \{ 1, \|\nabla p(\lambda^D) HA\|_2^2 U_2^j(T, t) \}$.

We use a similar argument to modify STEP 3 in Lemma EC.3 in the Online Supplement.

$$\begin{aligned} \Pr(\tau_3 \leq t) &= \Pr \left(\max_{v \leq t} \tilde{c}'_v \tilde{e}_v \geq \psi \right) = \Pr \left(\max_{v \leq t} \left[\sum_{j'=1}^n \left(\sum_{j=1}^m \sum_{l=1}^{k_v^j} \frac{q_{t_l^j}(j', j) \tilde{\xi}_l^j}{T - t_l^j + 1} \right)^2 \right] \geq \psi \right) \\ &\leq \Pr \left(\max_{v \leq t} m \sum_{j'=1}^n \sum_{j=1}^m \left(\sum_{l=1}^{k_v^j} \frac{q_{t_l^j}(j', j) \tilde{\xi}_l^j}{T - t_l^j + 1} \right)^2 \geq \psi \right) \leq \min \left\{ 1, \frac{m}{\psi} \mathbf{E} \left[\sum_{j'=1}^n \sum_{j=1}^m \left(\sum_{l=1}^{k_t^j} \frac{q_{t_l^j}(j', j) \tilde{\xi}_l^j}{T - t_l^j + 1} \right)^2 \right] \right\} \end{aligned}$$

The above inequality implies that: $\sum_{t=1}^{T-1} \Pr(\tau_3 \leq t) \leq \sum_{j=1}^m \sum_{t=1}^{T-1} \min \left\{ 1, \frac{m^2 \bar{Q}}{\psi} \|\nabla p(\lambda^D) HA\|_2^2 U_1^j(T, t) \right\} \leq \max \left\{ 1, \frac{m^2 \bar{Q}}{\psi} \right\} \sum_{j=1}^m \sum_{t=1}^{T-1} \min \{ 1, \|\nabla p(\lambda^D) HA\|_2^2 U_1^j(T, t) \}$.

Appendix C: Simulation Parameters and Table 2.

In all the experiments, we have 10 products and 4 resources. We use a multinomial logit demand (i.e., $\lambda_{t,i} = \exp(a_i - b_i p_{t,i}) / (1 + \sum_{j=1}^n \exp(a_j - b_j p_{t,j}))$) with the following parameters:

$$\begin{aligned} a &= [0.5 \quad 0.4 \quad 0.3 \quad 0.4 \quad 0.5 \quad 0.3 \quad 0.2 \quad 0.4 \quad 0.6 \quad 0.8]', \\ b &= [0.015 \quad 0.020 \quad 0.020 \quad 0.015 \quad 0.020 \quad 0.025 \quad 0.015 \quad 0.020 \quad 0.020 \quad 0.020]', \\ A &= \begin{bmatrix} 1 & 0 & 0 & 0 & 1 & 0 & 0 & 1 & 1 & 0 \\ 0 & 1 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 0 & 1 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 1 & 1 \end{bmatrix} \quad \text{and} \quad C = \begin{bmatrix} 0.1 \\ 0.1 \\ 0.1 \\ 0.1 \end{bmatrix}. \end{aligned}$$

The following table provides revenue loss (with respect to the deterministic upper bound) and revenue improvement (with respect to the static price control) of the heuristics tested in Experiment 2.

Table 2 Comparison of revenue loss (R.L.) and revenue improvement (R.I.).

θ	% R.L. compared to revenue upper bound					% R.I. over static pricing control.				% R.I. of LPC-4	
	Static	LPC-4	LPC-6	LPC-8	LPC-10	LPC-4	LPC-6	LPC-8	LPC-10	% R.I. of LPC-10	
500	5.94%	5.19%	4.59%	4.23%	4.15%	0.79%	1.43%	1.82%	1.90%	41.6%	
1000	4.22%	3.00%	2.61%	2.60%	2.28%	1.28%	1.69%	1.69%	2.02%	63.3%	
2000	2.99%	1.72%	1.57%	1.43%	1.37%	1.32%	1.46%	1.61%	1.67%	78.6%	
3000	2.48%	1.25%	1.09%	1.05%	0.99%	1.27%	1.43%	1.47%	1.53%	82.8%	
4000	2.13%	0.98%	0.86%	0.82%	0.77%	1.18%	1.30%	1.34%	1.40%	84.1%	
5000	1.94%	0.81%	0.70%	0.69%	0.65%	1.15%	1.26%	1.28%	1.31%	88.0%	
6000	1.81%	0.67%	0.61%	0.55%	0.58%	1.16%	1.22%	1.29%	1.25%	92.5%	
7000	1.64%	0.59%	0.54%	0.50%	0.47%	1.07%	1.12%	1.16%	1.20%	89.1%	
8000	1.58%	0.55%	0.48%	0.43%	0.40%	1.04%	1.12%	1.16%	1.20%	87.2%	
9000	1.42%	0.50%	0.42%	0.41%	0.39%	0.94%	1.01%	1.02%	1.05%	89.3%	
10000	1.34%	0.45%	0.42%	0.39%	0.35%	0.90%	0.93%	0.96%	1.00%	90.0%	

References

- Atar, R., M. Reiman. 2012. Asymptotically optimal dynamic pricing for network revenue management. *Stochastic Systems* **2** 232 – 276.
- Ben-Tal, A., L. E. Ghaoui, A. Nemirovski. 2009. *Robust Optimization*. Princeton University Press.

- Bertsimas, D., D. A. Iancu, P. A. Parrilo. 2010. Optimality of affine policies in multistage robust optimization. *Math. Oper. Res.* **35** 363–394.
- Bitran, G., R. Caldentey. 2003. An overview of pricing models for revenue management. *Manufacturing Service Oper. Management* **5** 203–229.
- Calafiore, G. C. 2009. An affine control method for optimal dynamic asset allocation with transaction costs. *SIAM J. Control Optim.* **48** 2254–2274.
- Chen, L., T. Homem de Mello. 2010. Re-solving stochastic programming models for airline revenue management. *Annals of Operations Research* **177** 91–114.
- Chen, Y., V. Farias. 2013. Simple policies for dynamic pricing with imperfect forecasts. *Oper. Res.* **61** 612–624.
- Ciocan, D. F., V. Farias. 2012. Model predictive control for dynamic resource allocation. *Math. Oper. Res.* **37** 501–525.
- Elmaghraby, W., P. Keskinocak. 2003. Dynamic pricing in the presence of inventory considerations: Research overview, current practices, and future directions. *Management Sci.* **49** 1287–1309.
- Erdelyi, A., H. Topaloglu. 2011. Using decomposition methods to solve pricing problems in network revenue management. *Journal of Pricing and Revenue Management* **10** 325–343.
- Gallego, G., G. van Ryzin. 1997. A multiproduct dynamic pricing problem and its applications to network yield management. *Oper. Res.* **45** 24–41.
- Golrezaei, N., H. Nazerzadeh, P. Rusmevichientong. 2014. Real-time optimization of personalized assortments. *Management Sci.* **60** 1532–1551.
- Haws, K. L., W. O. Bearden. 2006. Dynamic pricing and consumer fairness perceptions. *Journal of Consumer Research* **33** 304 – 311.
- Irvine, D. 2014. How airlines make less than \$ 6 per passenger. posted in www.cnn.com on tuesday, june 3, 2014. .
- Jasin, S. 2014. Re-optimization and self-adjusting price control for network revenue management. *Oper. Res.* **62** 1168–1178.
- Jasin, S., S. Kumar. 2012. A re-solving heuristic with bounded revenue loss for network revenue management with customer choice. *Math. Oper. Res.* **37** 313–345.
- Jasin, S., S. Kumar. 2013. Analysis of deterministic lp-based heuristics for revenue management. *Oper. Res.* **61** 1312–1320.
- Koushik, D., J. A. Higbie, C. Eister. 2012. Retail price optimization at intercontinental hotels group. *Interface* **42** 45 – 57.
- Kunnumkal, S., H. Topaloglu. 2010. A new dynamic programming decomposition method for the network revenue management problem with customer choice behavior. *Production and Oper. Management* **19** 575–590.

- Maglaras, C., J. Meissner. 2006. Dynamic pricing strategies for multiproduct revenue management problems. *Manufacturing Service Oper. Management* **8** 136–148.
- Marn, M. V., E. V. Roegner, C. C. Zawada. 2003. The power of pricing. *The McKinsey Quarterly* **1** 26–39.
- Pekgun, P., R. P. Menich, S. Acharya, P. G. Finch, F. Deschamps, K. Mallery, J. van Sistine, K. Christianson, J. Fuller. 2013. Carlson rezidor hotel group maximizes revenue through improved demand management and price optimization. *Interface* **43** 21 – 36.
- Ramasastry, A. 2005. Web sites change prices based on customers' habits. posted in www.cnn.com on friday, june 24, 2005. .
- Reiman, M. I., Q. Wang. 2008. An asymptotically optimal policy for a quantity-based network revenue management problem. *Math. Oper. Res.* **33** 257–282.
- Rigby, D., K. Miller, J. Chernoff, S. Tager. 2012. Are consumers waiting for better deals? *Bain Retail Holiday Newsletter* **4** 1–20.
- Sahay, A. 2007. How to reap higher profits with dynamic pricing. *MIT Sloan Management Review* **48** 53–60.
- Secomandi, N. 2008. An analysis of the control-algorithm re-solving issue in inventory and revenue management. *Manufacturing Service Oper. Management* **10** 468–483.
- Sen, A. 2013. A comparison of fixed and dynamic pricing policies in revenue management. *Omega* **41** 586–597.
- Talluri, K., G. van Ryzin. 2005. *The theory and Practice of Revenue Management*. Springer.

Online Supplement: Real-Time Dynamic Pricing with Minimal and Flexible Price Adjustment

Qi (George) Chen, Stefanus Jasin, Izak Duenyas

Stephen M. Ross School of Business
University of Michigan, Ann Arbor, MI 48109 USA
georgeqc, sjasin, duenyas@umich.edu

Proof of Theorem 1.

The key to the proof lies in the definition of a stopping time $\tau(\theta)$ for each of the θ^{th} problem, which can be roughly interpreted as the time when the first stock-out of any of the resources occurs. We use a martingale argument to derive an upper bound of the expectation of the remaining length of the selling season after $\tau(\theta)$, namely $\mathbf{E}[T(\theta) - \tau(\theta)]$. The main idea of the proof is to consider the revenue loss incurred before and after $\tau(\theta)$ separately. We show that both of them are in the order of $\mathbf{E}[T(\theta) - \tau(\theta)]$. Therefore, our primary task is to obtain an upper bound of $\mathbf{E}[T(\theta) - \tau(\theta)]$. The basic outline of the proof is as follows: (1) We show that, under some conditions, the resource consumption error can be explicitly written as a function of past demand errors, namely Δ_s 's; (2) We introduce some technical conditions and define a stopping time $\tau(\theta)$. This can be roughly interpreted as the stock-out time of the first depleted resource. We then compute an upper bound for $\mathbf{E}[T(\theta) - \tau(\theta)]$; (3) We break down our analysis of revenue loss into two parts, before and after period τ . Not surprisingly, the latter is in the order of $\mathbf{E}[T(\theta) - \tau(\theta)]$, thanks to the bounded revenue assumption in (A2). The rest of the proof shows that the revenue loss before $\tau(\theta)$ is also in the order of $\mathbf{E}[T(\theta) - \tau(\theta)]$.

Without loss of generality, assume $T = 1$. Then $T(\theta) = \theta$. For notational clarity, we suppress the dependency on θ whenever there is no confusion. Fix a projection matrix H that selects \mathcal{B} . We proceed in several steps.

STEP 0

We present a well-known result in linear algebra without proof. We will use this result several times.

LEMMA EC. 1. *For any real symmetric n by n matrix S , there exists an n by n orthonormal matrix $Q \in \mathbb{R}^n \times \mathbb{R}^n$ such that $Q^{-1}SQ = \Lambda$, where $\Lambda = \text{diag}(\theta_1, \dots, \theta_n)$ is a diagonal matrix whose elements are the eigenvalues of the matrix S . In addition, for any vector $v \in \mathbb{R}^n$, we have: $v'Sv \leq \max_{1 \leq i \leq n} |\theta_i| \cdot v'v$.*

STEP 1

In this step we derive an explicit formula for resource consumption error.

Define $\delta_s := A\Delta_s$, $\tilde{\delta}_i^j := A\tilde{\Delta}_i^j$, and $\epsilon_t := \sum_{j=1}^m E^j \nabla p(\lambda^D) H \sum_{l=1}^{k_t^j} \tilde{\delta}_l^j / (\theta - t_l^j + 1)$. (We follow the convention that if the lower limit of a summation is bigger than the higher limit, then the sum is zero.) By Taylor's expansion,

$$\lambda_t = \lambda^D - \nabla \lambda(p^D) \epsilon_t + \frac{1}{2} \epsilon_t' \nabla^2 \lambda(\eta_t) \epsilon_t, \quad \eta_t \in [p^D, p^D - \epsilon_t], \quad (1)$$

where, by a slight abuse of the notation, we use

$$\epsilon_t' \nabla^2 \lambda(\eta_t) \epsilon_t := \begin{bmatrix} \epsilon_t' \nabla^2 \lambda_1(\eta_t) \epsilon_t \\ \vdots \\ \epsilon_t' \nabla^2 \lambda_n(\eta_t) \epsilon_t \end{bmatrix},$$

and $\nabla^2 \lambda_j$ is the Hessian matrix of $\lambda_j(p_t)$. (Formula (1) holds if λ_t lies in the interior of Ω_λ . We will address this in STEP 2.) Since H is the projection matrix that selects \mathcal{B} . By definition, there exists an invertible m by m matrix M such that $\nabla p(\lambda^D)H = [M' \mathbf{0}']'$ and $A \nabla \lambda(p^D) = [M^{-1} | \dots]$, where the latter holds because $A \nabla \lambda(p^D) \nabla p(\lambda^D)H = AH = I$. Define M^j to be a square matrix whose j^{th} row is the same as M while the other rows are zeros. By definition, $M^j \tilde{\delta}_t^j = \tilde{\xi}_t^j \mathbf{e}_j$, where \mathbf{e}_j is a column vector with a proper size whose j^{th} element is one and the others are zeros. We can write ϵ_t as:

$$\epsilon_t = \begin{bmatrix} \sum_{j=1}^m \sum_{l=1}^{k_t^j} \frac{M^j \tilde{\delta}_t^j}{\theta - t_l^j + 1} \\ \mathbf{0} \end{bmatrix} = \begin{bmatrix} \sum_{j=1}^m \sum_{l=1}^{k_t^j} \frac{\tilde{\xi}_t^j \mathbf{e}_j}{\theta - t_l^j + 1} \\ \mathbf{0} \end{bmatrix}.$$

Because $A \nabla \lambda(p^D) = [M^{-1} | \dots]$, we have $A \nabla \lambda(p^D) E^j \nabla p(\lambda^D) H = M^{-1} M^j$ which allows us to write the following identity as long as λ_t lies in the interior of Ω_λ :

$$\begin{aligned} A \lambda_t - A \lambda^D &= -A \nabla \lambda(p^D) \sum_{j=1}^m E^j \nabla p(\lambda^D) H \sum_{l=1}^{k_t^j} \frac{\tilde{\delta}_t^j}{\theta - t_l^j + 1} + \frac{1}{2} A \epsilon_t' \nabla^2 \lambda(\eta_t) \epsilon_t \\ &= -M^{-1} \sum_{j=1}^m \sum_{l=1}^{k_t^j} \frac{M^j \tilde{\delta}_t^j}{\theta - t_l^j + 1} + \frac{1}{2} A \epsilon_t' \nabla^2 \lambda(\eta_t) \epsilon_t = -M^{-1} \sum_{j=1}^m \sum_{l=1}^{k_t^j} \frac{\tilde{\xi}_t^j \mathbf{e}_j}{\theta - t_l^j + 1} + \frac{1}{2} A \epsilon_t' \nabla^2 \lambda(\eta_t) \epsilon_t. \end{aligned} \quad (2)$$

STEP 2

We define a stopping time τ and give an upper bound for $\mathbf{E}[\theta - \tau]$. Recall that in (A4), we assume that the absolute values of the eigenvalues of the matrices $\nabla^2 \lambda_j, j = 1, \dots, n$ are bounded above by \bar{v} . Let $\lambda_L = \lambda^D - \phi_L \mathbf{e}$ and $\psi = \min \{ \psi', \psi'^2 \}$, where $\psi' = \min \left\{ \frac{\min \{ \phi_L, \phi_U \}}{\max \{ \bar{v}, 2 \cdot \|\nabla \lambda(p^D)\|_\infty \}}, \frac{\min \{ A \lambda_L \}}{\max \{ \|A \mathbf{e}\|_\infty, 2 \cdot \|M^{-1}\|_\infty \}} \right\}$. One can directly verify that $\psi > 0$. Define a stopping time τ to be the minimum of θ and the first time when any of the following conditions is violated.

$$\begin{aligned} \text{(E1)} \quad & \psi > \frac{1}{\theta - t} \left| \sum_{s=1}^t \left(\xi_s^j - \sum_{l=1}^{k_s^j} \frac{\tilde{\xi}_l^j}{\theta - t_l^j + 1} \right) \right|, \forall j = 1, \dots, m; \\ \text{(E2)} \quad & \psi > \frac{\bar{v}}{\theta - t} \sum_{s=1}^t \left\| \sum_{j=1}^m \sum_{l=1}^{k_s^j} \frac{\tilde{\xi}_l^j \mathbf{e}_j}{\theta - t_l^j + 1} \right\|_2^2; \\ \text{(E3)} \quad & \psi > \left\| \sum_{j=1}^m \sum_{l=1}^{k_t^j} \frac{\tilde{\xi}_l^j \mathbf{e}_j}{\theta - t_l^j + 1} \right\|_2^2. \end{aligned}$$

The three conditions listed above are somewhat technical and not easy to interpret. However, they are just stronger conditions of the two conditions below which have more obvious meaning.

$$\begin{aligned} \text{(E1}^*) \quad & \lambda_s \in [\lambda^D - \phi_L \mathbf{e}, \lambda^D - \phi_U \mathbf{e}] \subseteq \Omega_\lambda, \forall s \leq t; \\ \text{(E2}^*) \quad & C_t > 0, \end{aligned}$$

where C_t denotes the remaining inventory at the end of period t . The first condition states that all the target demand rates under LPC up to period t (including t) are feasible, so are the corresponding prices. The second condition states that no stock-out happens by the end of period t . Per our discussion in STEP 1, (E1^{*}) ensures the validity of expression (1) and (2). In addition, (E2^{*}) ensures that all the demand requests up to period t are satisfied, so the dynamics of the resource consumption can be fully expressed by the demand error Δ_s 's. Hence, under (E1^{*}) and (E2^{*}), we can track the inventory levels by explicitly quantifying them using past demand errors. (We emphasize that the purpose of (E1)-(E3) is simply for analytical tractability.) The following lemma reveals the connection between (E1)-(E3) and (E1^{*})-(E2^{*}).

LEMMA EC. 2. *We have: (E1)-(E3) \Rightarrow (E1^{*})-(E2^{*}). In other words, (E1^{*})-(E2^{*}) hold when $t < \tau$.*

The next lemma provides an upper bound of $\mathbf{E}[\theta - \tau]$ as a function of updating schedule $\{\gamma_j\}_{j=1}^m$.

LEMMA EC. 3. Let $U_1^j(T, t)$ and $U_2^j(T, t)$ be as defined in Theorem 1. Then, there exists a constant $\bar{\Psi}$, independent of θ and the choice of the projection matrix H such that:

$$\mathbf{E}[\theta - \tau(\theta)] \leq \bar{\Psi} \sum_{j=1}^m \sum_{t=1}^{\theta-1} \left(\min \left\{ 1, \left\| \nabla p(\lambda^D) H A \right\|_2^2 U_1^j(\theta, t) \right\} + \min \left\{ 1, \left\| \nabla p(\lambda^D) H A \right\|_2^2 U_2^j(\theta, t) \right\} \right).$$

Although the two lemmas above are crucial and their proofs are quite subtle, we defer the details for now and focus on the main thread of the proof.

STEP 3

We analyse the revenue loss incurred by LPC. Let $R_t(p_t)$ denote the revenue collected in period t under the posted price p_t . So, $R_{H, \gamma_B} = \sum_{t=1}^{\theta} R_t(p_t)$. Define $\bar{\Delta}_t := R_t(p_t) - \mathbf{E}[R_t(p_t) | \mathcal{F}_t] = R_t(p_t) - r(p_t)$. Since p^D is the optimal solution to DPP, $J^{Det} = p^{D'} \lambda(p^D) = r(p^D)$. This yields

$$\begin{aligned} J^{Det} - \mathbf{E}[R_{H, \gamma_B}] &= J^{Det} - \mathbf{E} \left[\sum_{t=1}^{\theta} R_t(p_t) \right] = \mathbf{E} \left[\sum_{t=1}^{\tau-1} (r(p^D) - R_t(p_t)) \right] + \mathbf{E} \left[\sum_{t=\tau}^{\theta} (r(p^D) - R_t(p_t)) \right] \\ &\leq \mathbf{E} \left[\sum_{t=1}^{\tau-1} (r(p^D) - R_t(p_t)) \right] + \mathbf{E} \left[\sum_{t=\tau}^{\theta} r(p^D) \right] \leq \mathbf{E} \left[\sum_{t=1}^{\tau-1} (r(p^D) - R_t(p_t)) \right] + \bar{r} \mathbf{E}[\theta - \tau + 1]. \end{aligned}$$

For $t < \tau$, by Taylor's expansion at p^D , we have $r(p_t) = r(p^D) - \nabla r(p^D) \epsilon_t + \frac{1}{2} \epsilon_t' \nabla^2 r(\rho_t) \epsilon_t$ for some $\rho_t \in [p^D, p^D - \epsilon_t]$. So, the first term after the last inequality above can be bounded as follows:

$$\begin{aligned} \mathbf{E} \left[\sum_{t=1}^{\tau-1} (r(p^D) - R_t(p_t)) \right] &= \mathbf{E} \left[\sum_{t=1}^{\tau-1} (r(p^D) - r(p_t) - \bar{\Delta}_t) \right] = \mathbf{E} \left[\sum_{t=1}^{\tau-1} \left(\nabla r(p^D) \epsilon_t - \frac{1}{2} \epsilon_t' \nabla^2 r(\rho_t) \epsilon_t - \bar{\Delta}_t \right) \right] \\ &= \mathbf{E} \left[\sum_{t=1}^{\tau-1} \nabla r(\lambda^D) \nabla \lambda(p^D) \epsilon_t \right] - \frac{1}{2} \mathbf{E} \left[\sum_{t=1}^{\tau-1} \epsilon_t' \nabla^2 r(\rho_t) \epsilon_t \right] - \mathbf{E} \left[\sum_{t=1}^{\tau-1} \bar{\Delta}_t \right] \\ &\leq \mathbf{E} \left[\sum_{t=1}^{\tau-1} \nabla r(\lambda^D) \nabla \lambda(p^D) \epsilon_t \right] - \frac{1}{2} \mathbf{E} \left[\sum_{t=1}^{\tau-1} \epsilon_t' \nabla^2 r(\rho_t) \epsilon_t \right] - \mathbf{E} \left[\sum_{t=1}^{\tau} \bar{\Delta}_t \right] + \bar{r}, \end{aligned}$$

where the third equality holds by the chain rule $\nabla r(\lambda^D) \nabla \lambda(p^D) = \nabla r(p^D)$ and the last inequality follows because $\mathbf{E}[\bar{\Delta}_\tau] \leq \mathbf{E}[R_\tau(p_\tau)] = \mathbf{E}[\mathbf{E}[R_\tau(p_\tau) | \mathcal{F}_\tau]] \leq \bar{r}$. Note that $\{\bar{\Delta}_t\}_{t=1}^{\theta}$ is a martingale with respect to the natural filtration and τ is bounded, so $\mathbf{E}[\sum_{t=1}^{\tau} \bar{\Delta}_t] = 0$ by the optional stopping theorem. Therefore, we only need to derive upper bounds for the first two terms above, which will be the primary focus of STEP 4 and 5.

STEP 4

We derive an upper bound for $\mathbf{E}[\sum_{t=1}^{\tau-1} \nabla r(\lambda^D) \nabla \lambda(p^D) \epsilon_t]$. Let π and μ denote the duals associated with the inventory constraints and the constraints $\lambda_t \in \Omega_\lambda$ of DPP respectively. Note that neither depends on θ . By assumption (A5), the optimal solution of DPP is interior. As a result of Karush-Kuhn-Tucker (KKT) optimality condition, we have $\nabla r(\lambda^D) = \pi' A$ (note that $\mu = 0$ by complementary slackness). Thus, $\mathbf{E}[\sum_{t=1}^{\tau-1} \nabla r(\lambda^D) \nabla \lambda(p^D) \epsilon_t] = \mathbf{E}[\sum_{t=1}^{\tau-1} \pi' A \nabla \lambda(p^D) \epsilon_t]$. By definition of ϵ_t and $A \nabla \lambda(p^D) E^j \nabla p(\lambda^D) H = M^{-1} M^j$ (see STEP 1), we can write

$$\begin{aligned} \mathbf{E} \left[\sum_{t=1}^{\tau-1} \nabla r(\lambda^D) \nabla \lambda(p^D) \epsilon_t \right] &= \mathbf{E} \left[\sum_{t=1}^{\tau-1} \pi' M^{-1} \sum_{j=1}^m \sum_{l=1}^{k_t^j} \frac{M^j \tilde{\delta}_l^j}{\theta - t_l^j + 1} \right] = \pi' M^{-1} \sum_{j=1}^m M^j \mathbf{E} \left[\sum_{t=1}^{\tau-1} \sum_{l=1}^{k_t^j} \frac{\tilde{\delta}_l^j}{\theta - t_l^j + 1} \right] \\ &= \pi' M^{-1} \sum_{j=1}^m M^j \mathbf{E} \left[\sum_{l=1}^{k_{\tau-1}^j} \left(1 - \frac{\theta - \tau + 1}{\theta - t_l^j + 1} \right) \tilde{\delta}_l^j \right]. \end{aligned} \quad (3)$$

The last term (3) can be further broken down into two parts as follows:

$$\begin{aligned}
& \pi' M^{-1} \sum_{j=1}^m M^j \mathbf{E} \left[\sum_{l=1}^{k_{\tau-1}^j} \left(1 - \frac{\theta - \tau + 1}{\theta - t_l^j + 1} \right) \tilde{\delta}_l^j + (1-1) \cdot \left(\delta_{t_{k_{\tau-1}^j}} + \cdots + \delta_{\tau-1} \right) \right] \\
&= \pi' M^{-1} \sum_{j=1}^m M^j \mathbf{E} \left[\left(\delta_{t_{k_{\tau-1}^j}} + \cdots + \delta_{\tau-1} + \sum_{l=1}^{k_{\tau-1}^j} \tilde{\delta}_l^j \right) - \left(\delta_{t_{k_{\tau-1}^j}} + \cdots + \delta_{\tau-1} + \sum_{l=1}^{k_{\tau-1}^j} \frac{\theta - \tau + 1}{\theta - t_l^j + 1} \tilde{\delta}_l^j \right) \right] \\
&\leq \pi' M^{-1} \sum_{j=1}^m M^j \mathbf{E} \left[\delta_{t_{k_{\tau-1}^j}} + \cdots + \delta_{\tau-1} + \sum_{l=1}^{k_{\tau-1}^j} \tilde{\delta}_l^j \right] + \pi' \mathbf{E} \left[\left| M^{-1} \sum_{j=1}^m M^j \left(\delta_{t_{k_{\tau-1}^j}} + \cdots + \delta_{\tau-1} + \sum_{l=1}^{k_{\tau-1}^j} \frac{\theta - \tau + 1}{\theta - t_l^j + 1} \tilde{\delta}_l^j \right) \right| \right]. \tag{4}
\end{aligned}$$

Since $\sum_{j=1}^m M^j = M$, by definition of δ_s and $\tilde{\delta}_l^j$ (see STEP 1), we can write

$$\pi' M^{-1} \sum_{j=1}^m M^j \mathbf{E} \left[\delta_{t_{k_{\tau-1}^j}} + \cdots + \delta_{\tau-1} + \sum_{l=1}^{k_{\tau-1}^j} \tilde{\delta}_l^j \right] = \pi' M^{-1} \sum_{j=1}^m M^j \mathbf{E} \left[\sum_{s=1}^{\tau-1} \delta_s \right] = \pi' \mathbf{E} \left[\sum_{s=1}^{\tau-1} \delta_s \right].$$

Observing that $\{\sum_{s=1}^t \Delta_s\}_{t=1}^\theta$ is a martingale and τ is bounded, $\mathbf{E}[\sum_{s=1}^\tau \Delta_s] = 0$ by optional stopping theorem. Also, the elements in π and A are all nonnegative. This implies that $\pi' A \mathbf{E}[\Delta_\tau] = \pi' A \mathbf{E}[\mathbf{E}[\Delta_\tau | \tau]] \leq \pi' A \bar{\lambda} \mathbf{e}$. Thus, the first term in (4) can be bounded by

$$\pi' \mathbf{E} \left[\sum_{s=1}^{\tau-1} \delta_s \right] = \pi' A \mathbf{E} \left[\sum_{s=1}^{\tau-1} \Delta_s \right] \leq \pi' A \mathbf{E} \left[\sum_{s=1}^\tau \Delta_s \right] + \pi' A \bar{\lambda} \mathbf{e} = \pi' A \bar{\lambda} \mathbf{e}. \tag{5}$$

As for the second term in (4), we have the following:

$$\begin{aligned}
& \pi' \mathbf{E} \left[\left| M^{-1} \sum_{j=1}^m M^j \left(\delta_{t_{k_{\tau-1}^j}} + \cdots + \delta_{\tau-1} + \sum_{l=1}^{k_{\tau-1}^j} \frac{\theta - \tau + 1}{\theta - t_l^j + 1} \tilde{\delta}_l^j \right) \right| \right] \\
&= \pi' \mathbf{E} \left[\left| M^{-1} \sum_{j=1}^m \left(\xi_{t_{k_{\tau-1}^j}}^j + \cdots + \xi_{\tau-1}^j + \sum_{l=1}^{k_{\tau-1}^j} \frac{\theta - \tau + 1}{\theta - t_l^j + 1} \tilde{\xi}_l^j \right) \mathbf{e}_j \right| \right] \\
&\leq \pi' \mathbf{E} \left[\left\| M^{-1} \right\|_\infty \max_{j=1, \dots, m} \left| \xi_{t_{k_{\tau-1}^j}}^j + \cdots + \xi_{\tau-1}^j + \sum_{l=1}^{k_{\tau-1}^j} \frac{(\theta - \tau + 1) \tilde{\xi}_l^j}{\theta - t_l^j + 1} \right| \right] \mathbf{e} \\
&= \pi' \mathbf{E} \left[\left\| M^{-1} \right\|_\infty \max_{j=1, \dots, m} \left| \sum_{s=1}^{\tau-1} \left(\xi_s^j - \sum_{l=1}^{k_s^j} \frac{\tilde{\xi}_l^j}{\theta - t_l^j + 1} \right) \right| \right] \mathbf{e} \\
&\leq \pi' \mathbf{E} \left[\frac{A \lambda_L}{2} (\theta - \tau + 1) \right] \leq \frac{\pi' A \bar{\lambda} \mathbf{e}}{2} \mathbf{E} [\theta - \tau + 1], \tag{6}
\end{aligned}$$

where the last equality holds because

$$\sum_{s=1}^{\tau-1} \left(\xi_s^j - \sum_{l=1}^{k_s^j} \frac{\tilde{\xi}_l^j}{\theta - t_l^j + 1} \right) = \xi_{t_{k_{\tau-1}^j}}^j + \cdots + \xi_{\tau-1}^j + \sum_{l=1}^{k_{\tau-1}^j} \frac{(\theta - \tau + 1) \tilde{\xi}_l^j}{\theta - t_l^j + 1}, \tag{7}$$

and the second to the last inequality results from the definition of ψ , the condition (E1) used to define τ , and the fact that $\min\{A \lambda_L\} \mathbf{e} \leq A \lambda_L$. Combining (5) – (6), we get: $\mathbf{E}[\sum_{t=1}^{\tau-1} \nabla r(\lambda^D) \nabla \lambda(p^D) \epsilon_t] \leq \pi' A \bar{\lambda} \mathbf{e} + \frac{1}{2} \pi' A \bar{\lambda} \mathbf{E}[\theta - \tau + 1]$.

STEP 5

We now derive an upper bound for $-\frac{1}{2}\mathbf{E}\left[\sum_{t=1}^{\tau-1}\epsilon'_t\nabla^2r_t(\rho_t)\epsilon_t\right]$ as follows:

$$-\frac{1}{2}\mathbf{E}\left[\sum_{t=1}^{\tau-1}\epsilon'_t\nabla^2r_t(\rho_t)\epsilon_t\right]\leq\mathbf{E}\left[\sum_{t=1}^{\tau-1}\epsilon'_t\nabla^2r_t(\rho_t)\epsilon_t\right]\leq\bar{v}\mathbf{E}\left[\sum_{t=1}^{\tau-1}\left\|\sum_{j=1}^m\sum_{l=1}^{k_t^j}\frac{\tilde{\xi}_l^j\mathbf{e}_j}{\theta-t_l^j+1}\right\|_2^2\right]\leq\psi\mathbf{E}[\theta-\tau+1],$$

where the second inequality follows from Lemma EC.1 and assumption (A5), and the last inequality follows from condition (E2) in the definition of τ .

STEP 6

Putting together results in STEP 1 - 5 proves Theorem 1. We only need to prove Lemma EC.2 and EC.3 which we do below.

Proof of Lemma EC.2. We need to show that if $t < \tau$, then (E1*) and (E2*) hold. We first show that (E1*) holds:

$$|\epsilon'_t\nabla^2\lambda(\eta_t)\epsilon_t|\leq\bar{v}\mathbf{e}\epsilon'_t\epsilon_t=\bar{v}\mathbf{e}\left\|\sum_{j=1}^m\sum_{l=1}^{k_t^j}\frac{\tilde{\xi}_l^j\mathbf{e}_j}{\theta-t_l^j+1}\right\|_2^2<\min\{\phi_L,\phi_U\}\mathbf{e}.$$

The last inequality follows from (E3) in the definition of τ . In addition, we also have

$$\begin{aligned} \left\|\nabla\lambda(p^D)\epsilon_t\right\|_\infty &= \left\|\nabla\lambda(p^D)\sum_{j=1}^m\sum_{l=1}^{k_t^j}\frac{\tilde{\xi}_l^j\mathbf{e}_j}{\theta-t_l^j+1}\right\|_\infty\leq\left\|\nabla\lambda(p^D)\right\|_\infty\cdot\left\|\sum_{j=1}^m\sum_{l=1}^{k_t^j}\frac{\tilde{\xi}_l^j\mathbf{e}_j}{\theta-t_l^j+1}\right\|_\infty \\ &= \left\|\nabla\lambda(p^D)\right\|_\infty\cdot\max_{j=1,\dots,m}\left|\sum_{l=1}^{k_t^j}\frac{\tilde{\xi}_l^j}{\theta-t_l^j+1}\right|\leq\left\|\nabla\lambda(p^D)\right\|_\infty\left\|\sum_{j=1}^m\sum_{l=1}^{k_t^j}\frac{\tilde{\xi}_l^j\mathbf{e}_j}{\theta-t_l^j+1}\right\|_2<\psi'\left\|\nabla\lambda(p^D)\right\|_\infty\leq\frac{1}{2}\min\{\phi_L,\phi_U\}. \end{aligned}$$

By combining the two inequalities above with (1), we get

$$|\lambda_t-\lambda^D|\leq\left|\nabla\lambda(p^D)\epsilon_t\right|+\left|\frac{1}{2}A\epsilon'_t\nabla^2\lambda(\eta_t)\epsilon_t\right|\leq\left\|\nabla\lambda(p^D)\epsilon_t\right\|_\infty\mathbf{e}+\left|\frac{1}{2}A\epsilon'_t\nabla^2\lambda(\eta_t)\epsilon_t\right|\leq\min\{\phi_L,\phi_U\}\mathbf{e}.$$

So, (E1*) holds. We next show that (E1)-(E3) imply (E2*). Since (E1)-(E3) imply (E1*), we know formula (1) holds for all $s \leq t$. As a result, the resource consumption error formula (2) also holds. Define $C_t^D := C - \sum_{s=1}^t A\lambda^D$. Then, the remaining inventory at the end of period t satisfies

$$\begin{aligned} C_t &\geq C - \sum_{s=1}^t AD_s = C - \sum_{s=1}^t A(\Delta_s + \lambda_s + \lambda^D - \lambda^D) = C_t^D - \sum_{s=1}^t A(\Delta_s + \lambda_s - \lambda^D) \\ &= C_t^D - \sum_{s=1}^t \left(\delta_s - M^{-1} \sum_{j=1}^m \sum_{l=1}^{k_s^j} \frac{M^j \tilde{\delta}_l^j}{\theta - t_l^j + 1} + \frac{1}{2} A \epsilon'_s \nabla^2 \lambda(\eta_s) \epsilon_s \right) \\ &\geq C_t^D - \left| M^{-1} \sum_{s=1}^t \sum_{j=1}^m \left(\xi_s^j \mathbf{e}_j - \sum_{l=1}^{k_s^j} \frac{\tilde{\xi}_l^j \mathbf{e}_j}{\theta - t_l^j + 1} \right) \right| - \left| \frac{1}{2} \sum_{s=1}^t A \epsilon'_s \nabla^2 \lambda(\eta_s) \epsilon_s \right|. \end{aligned} \quad (8)$$

Because $\{\lambda^D\}$ is the optimal solution to DPP, we know that it must satisfy inventory constraint. So, $C_t^D = C - \sum_{s=1}^t A\lambda^D \geq \sum_{s=t+1}^\theta A\lambda^D$. Since we also have $\lambda^D > \lambda_L \mathbf{e}$, it must hold that $C_t^D \geq \sum_{s=t+1}^\theta A\lambda^D \geq A\lambda_L(\theta - t)$. As for the second term in (8), by (E1), we have

$$\begin{aligned} & \left| M^{-1} \sum_{s=1}^t \sum_{j=1}^m \left(\xi_s^j \mathbf{e}_j - \sum_{l=1}^{k_s^j} \frac{\tilde{\xi}_l^j \mathbf{e}_j}{\theta - t_l^j + 1} \right) \right| \leq \|M^{-1}\|_\infty \cdot \left\| \sum_{j=1}^m \sum_{s=1}^t \left(\xi_s^j \mathbf{e}_j - \sum_{l=1}^{k_s^j} \frac{\tilde{\xi}_l^j \mathbf{e}_j}{\theta - t_l^j + 1} \right) \right\|_\infty \cdot \mathbf{e} \\ & \leq \|M^{-1}\|_\infty \cdot \max_{j=1, \dots, m} \left| \sum_{s=1}^t \left(\xi_s^j - \sum_{l=1}^{k_s^j} \frac{\tilde{\xi}_l^j}{\theta - t_l^j + 1} \right) \right| \mathbf{e} < \|M^{-1}\|_\infty \psi(\theta - t) \mathbf{e} < \frac{A\lambda_L}{2}(\theta - t) \leq \frac{1}{2} C_t^D. \end{aligned}$$

For the third term in (8), the following holds by Lemma EC.1.

$$\begin{aligned} \left| \frac{1}{2} \sum_{s=1}^t A \epsilon'_s \nabla^2 \lambda(\eta_s) \epsilon_s \right| & \leq \frac{1}{2} A \sum_{s=1}^t |\epsilon'_s \nabla^2 \lambda(\eta_s) \epsilon_s| \leq \frac{1}{2} A \bar{\mathbf{v}} \mathbf{e} \sum_{s=1}^t \epsilon'_s \epsilon_s = \frac{1}{2} A \bar{\mathbf{v}} \mathbf{e} \sum_{s=1}^t \left\| \sum_{j=1}^m \sum_{l=1}^{k_s^j} \frac{\tilde{\xi}_l^j \mathbf{e}_j}{\theta - t_l^j + 1} \right\|_2^2 \\ & < \frac{1}{2} A \mathbf{e}(\theta - t) \psi \leq \frac{1}{2} \|A \mathbf{e}\|_\infty \mathbf{e}(\theta - t) \psi \leq \frac{A\lambda_L}{2}(\theta - t) \leq \frac{1}{2} C_t^D. \end{aligned}$$

Combining the two bounds above with (8), we get $C_t > 0$. So, (E2*) holds. \square

Proof of Lemma EC.3. Let τ_1^j denote the minimum of θ and the first time t such that condition (E1) is violated for j^{th} resource. Also, let denote τ_i , $i = 2, 3$, denote the minimum of θ and the first time t such that condition (Ei) is violated. Note that, by definition, $\tau = \min\{(\min_j \tau_1^j), \tau_2, \tau_3\}$. Since τ is nonnegative, $\mathbf{E}[\tau] = \sum_{t=0}^{\theta-1} \Pr(\tau > t)$. So, we can write $\mathbf{E}[\theta - \tau] = \theta - \mathbf{E}[\tau] = \sum_{t=1}^{\theta-1} \Pr(\tau \leq t)$. Since $\tau \leq t$ can only happen if either τ_1^j (for some j) or τ_2 or τ_3 gets hit by time t , by sub-additivity property of probability, we can bound: $\Pr(\tau \leq t) \leq \sum_{j=1}^m \Pr(\tau_1^j \leq t) + \Pr(\tau_2 \leq t) + \Pr(\tau_3 \leq t)$. So, it suffices to derive a bound for each component after the inequality. We do this in turn.

STEP 1

We derive an upper bound for $\Pr(\tau_1^j \leq t)$, $j = 1, \dots, m$. Fix t . For each $j = 1, \dots, m$, we define a hitting time $\tilde{\tau}_1^j$ to be the minimum of t and the first time $v \leq t$ such that $\psi \leq |S_v^j|$, where

$$S_v^j = \begin{cases} \frac{\xi_{t^j}^j + \dots + \xi_v^j}{\frac{k_v^j}{\theta - t^j} + 1} + \sum_{l=1}^{k_v^j} \frac{\tilde{\xi}_l^j}{\theta - t_l^j + 1}, & 1 \leq v \leq t^j - 1 \\ \frac{\xi_{t^j}^j + \dots + \xi_v^j}{\frac{k_v^j}{\theta - v}} + \sum_{l=1}^{k_v^j} \frac{\tilde{\xi}_l^j}{\theta - t_l^j + 1}, & t^j \leq v \leq t \end{cases}.$$

We now state a lemma which reveals the connection between τ_1^j and $\tilde{\tau}_1^j$, see STEP 4 for proof.

LEMMA EC. 4. *We have: $\Pr(\tau_1^j \leq t) \leq \Pr(\tilde{\tau}_1^j \leq t)$.*

Observe that for any given t , $\{S_v\}_{v=1}^t$ is a martingale with respect to the natural filtration $\{\mathcal{F}_v\}_{v=1}^t$. Hence, $\{|S_v|\}_{v=1}^t$ is a submartingale. By Doob's submartingale inequality and identity in (7), we have

$$\begin{aligned} \Pr(\tau_1^j \leq t) & \leq \Pr(\tilde{\tau}_1^j \leq t) = \Pr\left(\max_{v \leq t} \left| \frac{\xi_{t^j}^j + \dots + \xi_v^j}{\frac{k_v^j}{\theta - v}} + \sum_{l=1}^{k_v^j} \frac{\tilde{\xi}_l^j}{\theta - t_l^j + 1} \right| \geq \psi\right) \\ & \leq \min \left\{ 1, \frac{1}{\psi^2} \mathbf{E} \left[\left| \frac{\xi_{t^j}^j + \dots + \xi_t^j}{\frac{k_t^j}{\theta - t}} + \sum_{l=1}^{k_t^j} \frac{\tilde{\xi}_l^j}{\theta - t_l^j + 1} \right|^2 \right] \right\} = \min \left\{ 1, \frac{1}{\psi^2} \left\{ \sum_{s=t^j}^t \frac{\mathbf{E}[(\xi_s^j)^2]}{(\theta - t)^2} + \sum_{l=1}^{k_t^j} \frac{\mathbf{E}[(\tilde{\xi}_l^j)^2]}{(\theta - t_l^j + 1)^2} \right\} \right\}, \end{aligned}$$

where the last equality holds because $\mathbf{E}[\xi_s^j \xi_v^j] = 0$ for $s \neq v$, $\mathbf{E}[\tilde{\xi}_l^j \tilde{\xi}_w^j] = 0$ for $l \neq w$, and $\mathbf{E}[\xi_s^j \tilde{\xi}_l^j] = 0$ for $s \geq t^j$ and $l \leq k_t^j$. Now we want to estimate the expectations in the upper bound above.

We start with the term $(\xi_s^j)^2$. By matrix norm inequality, $(\xi_s^j)^2 \leq \sum_{i=1}^m (\xi_s^i)^2 = (MA\Delta_s)'(MA\Delta_s) \leq \|MA\|_2^2 \Delta_s' \Delta_s = \|\nabla p(\lambda^D)HA\|_2^2 \Delta_s' \Delta_s \leq \|\nabla p(\lambda^D)\|_2^2 \|H\|_2^2 \|A\|_2^2 \Delta_s' \Delta_s$. Taking expectation on both sides and using $\mathbf{E}[\Delta_t' \Delta_t] = \mathbf{Var}(\Delta_t) \leq 1$ (due to the assumption that at most one customer arrives in each period) yields $\mathbf{E}[(\xi_s^j)^2] \leq \|\nabla p(\lambda^D)\|_2^2 \|H\|_2^2 \|A\|_2^2$. By definition, $\tilde{\xi}_l^j = \sum_{s=t_{l-1}^j}^{t_l^j-1} \xi_s^j$. So we have $\mathbf{E}\left[\left(\tilde{\xi}_l^j\right)^2\right] = \sum_{s=t_{l-1}^j}^{t_l^j-1} \mathbf{E}\left[(\xi_s^j)^2\right] \leq \|\nabla p(\lambda^D)HA\|_2^2 \sum_{s=t_{l-1}^j}^{t_l^j-1} \mathbf{E}[\Delta_s' \Delta_s] \leq \|\nabla p(\lambda^D)HA\|_2^2 (t_l^j - t_{l-1}^j)$. Putting the inequalities together, we obtain that $\sum_{j=1}^m \sum_{t=1}^{\theta-1} \Pr(\tau_1^j \leq t) \leq \sum_{j=1}^m \sum_{t=1}^{\theta-1} \min\left\{1, \frac{\|\nabla p(\lambda^D)HA\|_2^2}{\psi^2} U_1^j(\theta, t)\right\}$

STEP 2

We derive an upper bound for $\Pr(\tau_2 \leq t)$. Since $\left\|\sum_{j=1}^m \sum_{l=1}^{k_s^j} \frac{\tilde{\xi}_l^j \mathbf{e}_j}{\theta - t_l^j + 1}\right\|_2^2 \geq 0$ and $\bar{v} \geq 0$, we conclude that for all $v \leq t$, $\frac{\bar{v}}{\theta-t} \sum_{s=1}^t \left\|\sum_{j=1}^m \sum_{l=1}^{k_s^j} \frac{\tilde{\xi}_l^j \mathbf{e}_j}{\theta - t_l^j + 1}\right\|_2^2 \geq \frac{\bar{v}}{\theta-v} \sum_{s=1}^v \left\|\sum_{j=1}^m \sum_{l=1}^{k_s^j} \frac{\tilde{\xi}_l^j \mathbf{e}_j}{\theta - t_l^j + 1}\right\|_2^2$. Therefore, by Markov's inequality, the following holds:

$$\begin{aligned} \Pr(\tau_2 \leq t) &= \Pr\left(\max_{v \leq t} \frac{\bar{v}}{\theta-v} \sum_{s=1}^v \left\|\sum_{j=1}^m \sum_{l=1}^{k_s^j} \frac{\tilde{\xi}_l^j \mathbf{e}_j}{\theta - t_l^j + 1}\right\|_2^2 \geq \psi\right) \\ &\leq \Pr\left(\frac{\bar{v}}{\theta-t} \sum_{s=1}^t \left\|\sum_{j=1}^m \sum_{l=1}^{k_s^j} \frac{\tilde{\xi}_l^j \mathbf{e}_j}{\theta - t_l^j + 1}\right\|_2^2 \geq \psi\right) \leq \min\left\{1, \frac{\bar{v}}{\psi(\theta-t)} \sum_{s=1}^t \mathbf{E}\left[\left\|\sum_{j=1}^m \sum_{l=1}^{k_s^j} \frac{\tilde{\xi}_l^j \mathbf{e}_j}{\theta - t_l^j + 1}\right\|_2^2\right]\right\}. \end{aligned}$$

By similar arguments as in STEP 1, we can bound

$$\mathbf{E}\left[\left\|\sum_{j=1}^m \sum_{l=1}^{k_s^j} \frac{\tilde{\xi}_l^j \mathbf{e}_j}{\theta - t_l^j + 1}\right\|_2^2\right] \leq \mathbf{E}\left[\sum_{j=1}^m \sum_{l=1}^{k_s^j} \left(\frac{\tilde{\xi}_l^j}{\theta - t_l^j + 1}\right)^2\right] \leq \sum_{j=1}^m \sum_{l=1}^{k_s^j} \frac{\mathbf{E}\left[\left(\tilde{\xi}_l^j\right)^2\right]}{(\theta - t_l^j + 1)^2} \leq \|\nabla p(\lambda^D)HA\|_2^2 \sum_{j=1}^m \sum_{l=1}^{k_s^j} \frac{t_l^j - t_{l-1}^j}{(\theta - t_l^j + 1)^2}.$$

As a result, we obtain

$$\sum_{t=1}^{\theta-1} \Pr(\tau_2 \leq t) \leq \sum_{j=1}^m \sum_{t=1}^{\theta} \min\left\{1, \frac{\bar{v} \|\nabla p(\lambda^D)HA\|_2^2}{\psi(\theta-t)} \sum_{s=1}^t \sum_{l=1}^{k_s^j} \frac{t_l^j - t_{l-1}^j}{(\theta - t_l^j + 1)^2}\right\} = \sum_{j=1}^m \sum_{t=1}^{\theta} \min\left\{1, \frac{\bar{v} \|\nabla p(\lambda^D)HA\|_2^2}{\psi} U_2^j(\theta, t)\right\}.$$

STEP 3

We derive an upper bound for $\Pr(\tau_3 \leq t)$. Observe that for all j , $\left\{\sum_{l=1}^{k_t^j} \frac{\tilde{\xi}_l^j}{\theta - t_l^j + 1}\right\}_{t=1}^{\theta}$ is a martingale. Since $\left\|\sum_{j=1}^m \sum_{l=1}^{k_t^j} \frac{\tilde{\xi}_l^j \mathbf{e}_j}{\theta - t_l^j + 1}\right\|_2^2 = \sum_{j=1}^m \left(\sum_{l=1}^{k_t^j} \frac{\tilde{\xi}_l^j}{\theta - t_l^j + 1}\right)^2$, we conclude that $\left\{\left\|\sum_{j=1}^m \sum_{l=1}^{k_t^j} \frac{\tilde{\xi}_l^j \mathbf{e}_j}{\theta - t_l^j + 1}\right\|_2^2\right\}_{t=1}^{\theta}$ is also a submartingale. So, by Doob's submartingale inequality and arguments in STEP 1, we have

$$\begin{aligned} \Pr(\tau_3 \leq t) &= \Pr\left(\max_{v \leq t} \left\|\sum_{j=1}^m \sum_{l=1}^{k_v^j} \frac{\tilde{\xi}_l^j \mathbf{e}_j}{\theta - t_l^j + 1}\right\|_2^2 \geq \psi\right) \leq \frac{1}{\psi} \mathbf{E}\left[\left\|\sum_{j=1}^m \sum_{l=1}^{k_t^j} \frac{\tilde{\xi}_l^j \mathbf{e}_j}{\theta - t_l^j + 1}\right\|_2^2\right] \\ &\leq \min\left\{1, \frac{\|\nabla p(\lambda^D)HA\|_2^2}{\psi} \sum_{j=1}^m \sum_{l=1}^{k_t^j} \frac{t_l^j - t_{l-1}^j}{(\theta - t_l^j + 1)^2}\right\}. \end{aligned}$$

As a result, the following inequality holds:

$$\sum_{t=1}^{\theta-1} \Pr(\tau_3 \leq t) \leq \sum_{j=1}^m \sum_{t=1}^{\theta-1} \min\left\{1, \frac{\|\nabla p(\lambda^D)HA\|_2^2}{\psi} \sum_{l=1}^{k_t^j} \frac{t_l^j - t_{l-1}^j}{(\theta - t_l^j + 1)^2}\right\} = \sum_{j=1}^m \sum_{t=1}^{\theta-1} \min\left\{1, \frac{\|\nabla p(\lambda^D)HA\|_2^2}{\psi} U_1^j(\theta, t)\right\}.$$

STEP 4

Putting together all the results in STEP 1-3 completes the proof of Lemma EC.3. The last thing to do is to prove Lemma EC.4 from STEP 1. We do this now.

Proof of Lemma EC.4. It suffices to show that for all $v \leq t$, if $\tau_1^j = v$ occurs, then $\tilde{\tau}_1^j \leq v$ occurs as well. By definition of S_v , this is immediately true if $t_{k_t^j}^j \leq v \leq t$. So, we only need to check the case $1 \leq v \leq t_{k_t^j}^j - 1$. Assuming $1 \leq v \leq t_{k_t^j}^j - 1$, by definition of (E1) in STEP 1, $\tau_1^j = v$ means

$$\psi \leq \left| \frac{\xi_{t_{k_t^j}^j}^j + \dots + \xi_v^j}{\theta - v} + \sum_{l=1}^{k_v^j} \frac{\tilde{\xi}_l^j}{\theta - t_l^j + 1} \right| \text{ and } \psi > \left| \sum_{l=1}^{k_v^j} \frac{\tilde{\xi}_l^j}{\theta - t_l^j + 1} \right| \text{ which imply that } \psi \leq \left| \frac{\xi_{t_{k_t^j}^j}^j + \dots + \xi_v^j}{\theta - t_{k_v^j+1}^j + 1} + \sum_{l=1}^{k_v^j} \frac{\tilde{\xi}_l^j}{\theta - t_l^j + 1} \right| = |S_v^j|.$$

So, $\tilde{\tau}_1^j \leq v$ and hence $\Pr(\tau_1^j \leq t) \leq \Pr(\tilde{\tau}_1^j \leq t)$. \square