A Nonparametric Self-Adjusting Control for Joint Learning and Optimization of Multi-Product Pricing with Finite Resource Capacity

Qi (George) Chen

Naveen Jindal School of Management, University of Texas at Dallas, Richardson, TX, USA georgeqc@utdallas.edu

Stefanus Jasin, Izak Duenyas

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

We study a multi-period network revenue management problem where a seller sells multiple products, made from multiple resources with finite capacity, in an environment where the underlying demand function is a priori unknown (in the nonparametric sense). The objective of the seller is to simultaneously learn the unknown demand function and dynamically price his products to minimize the expected revenue loss. For the problem where the number of selling periods and initial capacity are scaled by k > 0, it is known that the expected revenue loss of any non-anticipating pricing policy is $\Omega(\sqrt{k})$. However, there is a considerable gap between this theoretical lower bound and the performance bound of the best known heuristic control in the literature. In this paper, we propose a *Nonparametric Self-adjusting Control* and show that its expected revenue loss is $\mathcal{O}(k^{1/2+\epsilon} \log k)$ for any arbitrarily small $\epsilon > 0$, provided that the underlying demand function is sufficiently smooth. This is the tightest bound of its kind for the problem setting that we consider in this paper and it significantly improves the performance bound of existing heuristic controls in the literature; in addition, our intermediate results on the large deviation bounds for spline estimation and nonparametric stability analysis of constrained optimization are of independent interest and are potentially useful for other applications.

Key words: Revenue management; learning; self-adjusting control; spline approximation; asymptotic analysis
 MSC2000 subject classification: Primary: 90C59; secondary: 62G20, 90B50

OR/MS subject classification: Primary: probability, stochastic model applications; secondary: statistics, estimation

History:

1. Introduction. Revenue management (RM), which was first implemented in the 1960s by legacy airline companies to maintain their edge in the competitive airline market, has recently become widespread in many industries such as hospitality, fashion goods, and car rentals ([29], [27]). The sellers in these industries face the common challenge of allocating a *fixed* capacity of perishable resources (e.g., seats in a jet, rooms in a hotel, etc.) to satisfy *volatile* demand of products or services. If the seller fails to satisfy demand appropriately, a considerable amount of profit is at stake either due to the zero salvage value of unused capacity or the loss of potential revenue. (For example, in the airline industry, it is known that the benefit of using RM is roughly comparable to the airline's annual margin, which is about 4-5% of total revenue [29].) Given this, RM is aimed at helping the sellers to make optimal decisions such that the right product is sold to the right customer at the right time and at the right price. One type of operational leverage often employed by the sellers is *dynamic pricing*: By adjusting the prices over time, the seller can effectively control the rate at which demand arrives so he can better match demand with available resources.

Despite its known benefits [29], the efficacy of dynamic pricing hinges upon the seller's knowledge of market's response to different prices, i.e., the underlying demand function. Unfortunately, in most (if not all) real-life applications, this underlying demand function is not easily accessible to the sellers. Although many sellers have adopted sophisticated statistical methods, the estimated demand function is inevitably subject to estimation error, which in turn affects the quality of the sellers' pricing decisions. The negative impact of inaccurate demand estimation is further magnified in practice because typical RM industries tend to have a large sales volume; thus, even small errors can lead to a significant revenue loss in absolute term. Given this limitation, one pressing issue faced by RM practitioners is how to dynamically price their products when the underlying demand function is unknown a priori. This paper studies joint learning and pricing problem in a general network RM setting with *multiple* products and *multiple* capacitated resources for the nonparametric demand case. (By nonparametric, we mean the case where the seller does not even know the functional form of demand. This is in contrast to the so-called *parametric* case where the seller a priori knows the form of demand function (e.g., linear, exponential, logit, etc.) and he only needs to estimate the unknown parameters (e.g., the intercept and the slope of a linear demand function).) In this paper, we construct a heuristic control that is not only easy to implement for large-scale problems but also has a provable analytical performance bound. Our bound significantly improves the performance bound of existing heuristic controls in the literature.

Literature review. A large body of RM literature has investigated the canonical dynamic pricing problem where the seller knows the underlying demand function. The prevailing view is that, even in the case where learning is not in play, computing an optimal control is already challenging to do. This is so because the common technique for solving sequential decision problems, the so-called Dynamic Programming (DP), suffers from the well-known curse of dimensionality. This curse of dimensionality is exacerbated in many RM industries because the sellers typically have to manage the price of at least *thousands* of products on a daily basis. To illustrate, a typical major US airline operates more than a thousand flights daily, each of which has more than ten different booking classes that are characterized by different combinations of service level and purchase restriction. Since passengers book tickets in advance, the airline needs to price not only the tickets for the same-day flights but also those with departure dates several months in the future. All these factors put together can easily translate into a daily pricing decision for *millions* of itineraries. Due to this challenge, instead of finding the optimal pricing control, a considerable body of existing literature has focused on developing computationally implementable heuristic controls with provably good performance guarantee. (See Bitran and Caldentey [9] and Elmaghraby and Keskinocak [15] for a comprehensive review of the literature.)

Within the canonical RM literature, some works have focused on developing heuristic controls based on the solution of a deterministic pricing problem, i.e., the deterministic counterpart of the original stochastic control problem, which is computationally much easier to solve than the DP. This approach was first proposed by [18]. They develop a static price control by first solving a convex optimization problem at the beginning of the selling season and then using its optimal solution as static price throughout the selling season, subject to available resources. Although the proposed heuristic control is easy to implement, its drawback is obvious: It ignores the observed demand realizations, which leaves room for further improvement. One intuitively appealing idea that has been studied in the literature involves frequent *re-optimization* of the deterministic pricing problem throughout the selling season. Maglaras and Meissner [26] show that the re-optimization (in asymptotic sense). However, it is not immediately clear from their analysis alone whether reoptimization actually guarantees a better performance (and if so, by how much). A recent work by Jasin [21] answers this question in the affirmative by showing that RSC *does* significantly improve the performance of static price control (again, in asymptotic sense). While existing literature has shown frequent re-optimization to be beneficial, its implementation can be very time-consuming especially when applied to large-scale problems that often arise in practice; this has motivated the development for computationally much easier yet equally effective heuristic controls. For example, motivated by the optimal structure of the diffusion control problem of the continuous-time dynamic pricing problem, Atar and Reiman [1] develop a re-optimization-free bridge pricing control that guarantees the same asymptotic performance as RSC. An equally effective heuristic control is also obtained in Jasin [21]. Motivated by the structure of the re-optimized prices under RSC, Jasin [21] proposes a real-time control, called *Linear Rate Correction* (LRC), that has a similar structure as the bridge control and does not require any re-optimization at all. To be precise, LRC only requires a single optimization at the beginning of the selling season and automatically adjusts the price according to a pre-specified update rule throughout the remaining selling season. Inspired by the strong performance of bridge pricing control and LRC in the setting with known demand function, in this paper, we construct a nonparametric self-adjusting control akin to LRC and show that its asymptotic performance is very close to the theoretical lower bound on the performance of any feasible pricing control in the setting with unknown demand function. Below, we discuss the literature on joint learning and pricing.

There is a growing literature that studies joint demand learning and pricing problem. Most existing works have combined a particular statistical learning procedure (e.g., Maximum Likelihood, Least Squares, etc.) with a certain dynamic pricing control (most notably, the static price control). A central highlight in this literature is the trade-off between the cost of learning the demand function (exploration) and the reward of using the "optimal" price computed using the estimated demand function (exploitation). The longer the time the seller spends on learning the demand function, the less opportunity there is to exploit the knowledge of the newly estimated demand function. On the flip side, if the exploration time is too short, it will result in a poor estimation, which yields highly sub-optimal prices. What is the best performance that any non-anticipating pricing control can achieve in the setting with unknown demand function? Suppose that we scale the length of the selling season and the initial resource capacity by a factor of k > 0. (The constant k can be interpreted as the size of the problem. See $\S2.5$ for more discussions on this.) One way to measure the performance of a feasible control is to study the order of expected revenue loss which is defined as the order (with respect to k) of the gap between the total expected revenue earned under this control and a well-established deterministic upper bound. (See $\S2.4$ for more details on this performance metric.) It is widely known in the literature that the expected revenue loss of any feasible pricing control in general is $\Omega(\sqrt{k})$ (e.g., Besbes and Zeevi [6], Broder and Rusmevichientong [11], Keskin and Zeevi [23]). For the case of uncapacitated RM, where there is no limit on the number of resources that can be used, this lower bound has been repeatedly shown to be tight (e.g., Broder and Rusmevichientong [11], Keskin and Zeevi [23]). As for the case of capacitated RM, most existing literature has primarily focused on the setting of a singleproduct and single-resource RM (often called single-leg RM due to the early application of RM in airline industry). Besbes and Zeevi [6] is among the first to investigate this problem under both parametric and nonparametric cases. Their proposed heuristic control for the parametric case yields an expected revenue loss of $\mathcal{O}(k^{2/3}\log^{0.5}k)$ whereas their proposed heuristic control for the nonparametric case guarantees an expected revenue loss of $\mathcal{O}(k^{3/4} \log^{0.5} k)$. This suggests that there is a considerable gap between the performance of parametric and nonparametric approaches. Recent works by Wang et al. [30] and Lei et al. [25] have managed to significantly shrink this gap; they develop sophisticated nonparametric heuristic controls that guarantee a $\mathcal{O}(\sqrt{k}\log^{4.5}k)$ and $\mathcal{O}(\sqrt{k})$ expected revenue loss, respectively. Thus, for the case of capacitated RM in single-leg setting, existing works in the literature have managed to not only completely close the gap between the performance of parametric and nonparametric approaches, at least in the asymptotic sense, but also show that the theoretical lower bound of $\Omega(\sqrt{k})$ is indeed tight.

The general network RM problem with multiple products and multiple limited resources is more difficult to analyze than the single-leg RM. (In §4, we will explain why the proofs and the arguments in the uncapacitated setting cannot be applied to the capacitated setting. Moreover, the arguments in Wang et al. [30] and Lei et al. [25] for the single-leg capacitated RM also cannot be applied to the more general network RM setting. This is so because both Wang et al. [30] and Lei et al. [25] heavily exploit the special structure in the single-leg RM. Unfortunately, no analogous structure is known for the network RM.) To the best of our knowledge, the only paper that addresses the joint learning and pricing problem for general network RM is Besbes and Zeevi [7]. They consider the nonparametric case only and show that the performance bound of their proposed heuristic control (i.e., so-called Algorithm 2 in their paper) is $\mathcal{O}(k^{(n+2)/(n+3)}\log^{0.5}k)$, where n is the number of products. Note that the fraction (n+2)/(n+3) in the bound highlights the curse of dimensionality for network RM since the performance bound quickly deteriorates as the number of products nincreases. If, however, the true demand function is sufficiently smooth (e.g., infinitely differentiable), this bound can be reduced; they propose another nonparametric heuristic control (Algorithm 3) that guarantees a $\mathcal{O}(k^{2/3+\epsilon}\log^{0.5}k)$ expected revenue loss for some $\epsilon > 0$ that can be arbitrarily small. As one can see, there is still a considerable gap between the lower bound of $\Omega(\sqrt{k})$ and the performance bound of $\mathcal{O}(k^{2/3+\epsilon}\log^{0.5}k)$. Our proposed heuristic control in this paper significantly reduces this gap from $k^{2/3+\epsilon}$ to $k^{1/2+\epsilon}$ (up to logarithmic terms).

We would like to note here that all the results discussed above for the joint learning and pricing problem are derived for the setting where the seller can use a continuum of prices drawn from a certain convex and compact set. This distinction is crucial as the complexity of the problem changes as we switch from a continuum setting to the setting where the set of feasible prices is finite. Besbes and Zeevi [7] have also considered this finite set setting and proposed a heuristic control (Algorithm 1) with a performance bound of $\mathcal{O}(k^{2/3}\log^{0.5}k)$. A recent work by Ferreira et al. [16] improves this bound to $\mathcal{O}(\sqrt{k}\log k \log \log k)$ by using a Thompson sampling-based heuristic control. Note that although the best known performance bound that we are aware of for the network RM with finite price set setting is already close to \sqrt{k} , it is not clear that the theoretical lower bound for this setting is still $\Omega(\sqrt{k})$. In fact, a recent work by Flajolet and Jaillet [17] show that the UCB-Simplex control they propose attains a logarithmic performance bound in a simpler setting where there is only one capacity constraint (besides time). This seems to suggest that the lower bound of the general network RM with finite price set setting could be as small as logarithmic. Finally, we want to point out that, in order to derive the order of performance bound discussed above, the heuristic controls developed for finite price setting are typically compared with a revenue upper bound benchmark under the finite price setting while the heuristic controls developed for the continuum price setting are compared with a larger revenue upper bound benchmark under the continuum setting. Since the two upper bound benchmarks are different, with the former being smaller than the latter, the heuristic controls developed under the two settings are not easily comparable by simply looking at the order of their performance bounds as a function of k. Moreover, one also cannot simply extend existing heuristic controls developed in finite price setting to continuum price setting since the performance bounds of these heuristic controls (e.g., [16]) deteriorate quickly as the number of the feasible prices increases. This is so because these heuristic controls do not exploit existing relationship between expected demand value at different price points and need to learn the expected demand value at most price points separately, which is not very efficient if the number of price points is large.

More broadly, the joint learning and pricing problem is closely related to the literature on *bandit* problems (e.g., [24], [2], etc.). This stream of literature had not considered the inter-temporal constraints on actions over time (such as the capacity constraint in the network RM setting) until only very recently (e.g., [4], [5], [14], etc.). Badanidiyuru et al. [4] is among the first to consider the so-called *bandit with knapsack* (BwK) problem. In BwK, a decision maker has a fixed amount

of resources and needs to accumulate rewards by sequentially selecting from a finite set of arms whose reward distributions are unknown. Pulling each arm stochastically depletes those resources according to an unknown consumption distribution of that arm, and the decision maker stops collecting rewards when he runs out of resources. Note that the general RM with finite price set and unknown demand fits into the BwK framework by treating each feasible price vector as a bandit and viewing time as a resource with deterministic depletion rate. Badanidiyuru et al. [5] show that the performance lower bound of BwK is $\Omega(\sqrt{k})$ and propose two heuristic controls that match this lower bound up to logarithmic factors. While the network RM with finite number of feasible prices is a special case of the BwK problem, we would like to point out that the bounds derived in Badanidiyuru et al. [5] cannot be directly compared to the bounds derived in all the RM papers discussed above since the quantification orders used in evaluating the asymptotic bounds are different. To be precise, Badanidiyuru et al. [5] allow the underlying demand distribution to vary in k while existing RM works assume that the underlying demand distribution does not vary as k scales; hence, both the lower and upper bounds in Badanidiyuru et al. [5] have weakly larger asymptotic order. (This phenomena is not unique to Badanidiyuru et al. [5] alone. For example, the gap in performance bounds due to the use of different quantification orders also arises in the traditional bandit setting such as in the logarithmic bounds of Auer et al. [2] versus the square-root bounds of Auer et al. [3]—see page 50 in Auer et al. [3] for more discussions.) As an extension, Badanidiyuru et al. 5 also consider the case where the decision maker has a continuum of arms (this corresponds to the setting we study where the seller has a continuum of feasible prices) and derive a revenue loss bound in the order of $\mathcal{O}(k^{2/3})$. While we study a similar setting, our revenue loss bound is sharper due to two reasons: (1) following the paradigm in RM literature, we do not allow the underlying demand distribution to vary as k scales; (2) our result relies on smoothness assumptions on the underlying demand function whereas Badanidiyuru et al. [5] does not require this assumption. Hence, while the setting in Badanidiyuru et al. [5] is closely related to our setting, the bounds they derive are not directly comparable to ours. Combes et al. [14] also study a related setting where the seller has a finite number of arms and the seller is faced with sample-path dependent budget constraints on how many times each arm can be pulled. The constraints they consider are separable in the sense that pulling an arm only affects its own budget but not the budgets of other arms. Our setting is different from theirs along at least two dimensions: First, we consider a continuum of feasible actions rather than a finite number of actions; secondly, the resource constraints in our setting are not separable because the consumption of a product could affect the capacity/inventory levels of multiple resources.

Apart from the RM and joint learning and optimization literature, our work is also closely related to Spline Regression in the statistics literature. A typical problem in statistics is to estimate the mean response as a function of some input variables. (The demand learning aspect of our problem is one such problem: our goal is to estimate the mean demand of each product as a function of the prices of all products.) Spline Regression generates an estimate in the form of a linear combination of spline basis functions (originally studied in Applied Mathematics to approximate deterministic functions) and uses Least Squares criterion to compute the corresponding coefficients. (See Gyorfi et al. [20] for more details on Spline Regression.) To the best of our knowledge, most existing literature on Spline Regression is mostly concerned with the estimation accuracy of the response function; thus, the typical convergence result for Spline Regression is limited to only the estimation error of the response function itself. In our problem, estimation and optimization are intertwined and it is crucial to understand how the estimation error affects the subsequent optimization. This requires results on bounds of the error between higher order partial derivatives of the response function and its spline estimate. To derive these bounds, deviating from the Spline Regression approach, we generate our spline estimate by using a specific linear operator (i.e., \mathcal{L} defined in Step 3 of Technical Details for Spline Approximation part (b) in $\S3.1$ instead of using Least Squares criterion. (We call our estimation method *Spline Estimation* to differentiate our approach from Spline Regression. For more details on Spline Estimation, see §3.1. We want to emphasize here that although we choose Spline Estimation in combination with our optimization, this is only for the purpose of mathematical analysis. In general, we suspect that the seller can also use other estimation schemes such as Spline Regression, local polynomial approximation, etc. in lieu of Spline Estimation in our proposed heuristic control and still enjoys a strong performance.) This linear operator, also known as a *quasi-interpolant* in Schumaker [28], is originally devised to analyze the error of using spline functions to approximate a *deterministic* function. We generalize the analysis of spline approximation for deterministic functions in Schumaker [28] to estimating response functions using spline functions. We manage to not only derive large deviation bounds of the estimation error of the demand function itself but also its partial derivatives, which are very useful in analyzing our proposed heuristic control.

Our contributions and the organization of the paper. Our contributions in this paper can be summarized in the following two points:

1. We develop a nonparametric control called Nonparametric Self-adjusting Control (NSC) that can be applied to the general network RM setting with multiple products and multiple limited resources. We show that if the underlying demand function is sufficiently smooth, the expected revenue loss of NSC is $\mathcal{O}(k^{1/2+\epsilon}\log k)$ for some $\epsilon > 0$ that can be arbitrarily small (see Theorem 1 in §4). This is the *tightest* bound of its kind for the setting that we are considering (i.e., the continuum price setting): It significantly improves the $\mathcal{O}(k^{2/3+\epsilon}\log^{0.5}k)$ bound of Besbes and Zeevi [7] and is only slightly worse than the theoretical lower bound of $\Omega(\sqrt{k})$. In a nutshell, NSC is a combination of four elements: (1) Spline Estimation of the underlying demand function, (2) linear approximation of the estimated demand function, (3) quadratic approximation of the estimated revenue function, and (4) approximate self-adjusting control akin to the one developed in Jasin [21]. In this paper, we show that although each of these elements introduces its own error (and some of them are *not* even unbiased), under properly selected tuning parameters, the cumulative impact of these errors is asymptotically only slightly larger than $\Theta(\sqrt{k})$. This makes it possible to prove the strong performance guarantee of NSC. Note that, per our discussions above, the first element (i.e., using Spline Estimation to estimate the underlying demand function) is only for the purpose of the analysis; in practice, the seller can still use other estimation schemes in combination with the other three elements. Although the proper selection of the tuning parameters will depend on the specific estimation scheme being used, we suspect that the resulting heuristic control still enjoys a similar strong performance guarantee.

2. In addition to contributing to the RM literature, our intermediate results in this paper also contribute to the more general statistics and optimization literature. For the nonparametric estimation, we generalize the analysis of spline approximation for deterministic functions in Schumaker [28] to the setting with noisy observations and derive large deviation bounds for the estimation error of the function itself and its higher order partial derivatives (see Lemma 1 in §3). These bounds seem to be new—although spline functions have been used in statistics, we are not aware of existing large deviation bound for higher order partial derivatives of the estimate—and are particularly useful for our analysis because the resulting spline estimate is ultimately used in the subsequent optimization phase in our heuristic control. Moreover, the bound for partial derivatives also facilitates the stability analysis of the optimal solution. Aside from the statistical error bound, for the analysis of NSC, we also need to derive a nonparametric Lipschitz-type stability result for the optimal solution of a perturbed optimization problem (see Lemma 2 in §3). This result also seems to be new—although *parametric* stability analysis of optimization problem has been intensively studied in the literature (see Bonnans and Shapiro [10]), nonparametric stability analysis is very rare. As of the writing of this paper, we are not aware of any existing result on

179.R2

7

nonparametric stability analysis that can be directly used for our purpose. Aside from its use in the analysis of NSC, our stability result is of independent interest and is potentially applicable to other optimization problems.

The reminder of this paper is organized as follows. We first formulate the problem in §2. Our nonparametric approach is discussed in §3-§5. In particular, §3 provides some preliminaries on spline approximation and nonparametric stability analysis; §4 describes the proposed NSC and its performance bound (Theorem 1); §5 provides the proof of Theorem 1. Finally, we conclude the paper in §6. Unless otherwise noted, all the extra details of the proofs can be found in the appendix at the end of this paper.

2. Problem formulation. In this section we describe the problem setting, modeling assumptions, and the asymptotic regime.

2.1. Notation. The following notation will be used throughout the paper. (Additional notation will be introduced when necessary.) We denote by \mathbb{R} , \mathbb{R}_+ , and \mathbb{R}_{++} the set of real, nonnegative real, and positive real numbers respectively. For column vectors $a = (a_1; \ldots; a_n) \in \mathbb{R}^n$, $b = (b_1; \ldots; b_n) \in \mathbb{R}^n$, we denote by $a \succeq b$ if $a_i \ge b_i$ for all i, and by $a \succ b$ if $a_i > b_i$ for all i. Similarly, we denote by \mathbb{Z} , \mathbb{Z}_+ , and \mathbb{Z}_{++} the set of integers, non-negative integers, and positive integers respectively. For any $a, b \in \mathbb{Z}$ with $a \le b$, let $[a, b] := \{a, a+1, \ldots, b-1, b\}$. We denote by \cdot the inner product of two vectors and by \otimes the tensor product of sets or function spaces. We use a prime to denote the transpose of a vector or a matrix, an I to denote an identity matrix with a proper dimension, and an \mathbf{e} to denote a vector of ones with a proper dimension. Following the standard notation, for any real matrix $M = [M_{ij}] \in \mathbb{R}^{m \times n}$, we use $||M||_1 := \max_{1 \le j \le n} \sum_{i=1}^m |M_{ij}|, ||M||_2 :=$ the largest eigenvalue of M'M, $||M||_{\infty} := \max_{1 \le i \le m} \sum_{j=1}^n |M_{ij}|)$, and $||M||_F := \sqrt{\sum_{i=1}^m \sum_{j=1}^n |M_{ij}|^2}$ to denote the induced 1-norm, induced 2-norm, induced ∞ -norm, and Frobenius norm of M respectively. (Note that $||M||_1 = ||M'||_{\infty}$.) For any function $f : \mathcal{X} \to \mathcal{Y}$, we denote by $||f(.)||_{\infty} := \sup_{x \in \mathcal{X}} ||f(x)||_{\infty}$ the infinity-norm of f. We use ∇ to denote the usual derivative operator and a subscript ∇ means that the derivative is applied to all variables.) If $f : \mathbb{R}^n \to \mathbb{R}$, then $\nabla_x f = (\frac{\partial f}{\partial x_1}; \ldots; \frac{\partial f}{\partial x_n})$; if, on the other hand, $f = (f_1; \ldots; f_n) : \mathbb{R}^n \to \mathbb{R}^n$, then

$$\nabla_x f = \begin{bmatrix} \frac{\partial f_1}{\partial x_1} & \cdots & \frac{\partial f_n}{\partial x_1} \\ \vdots & \ddots & \vdots \\ \frac{\partial f_1}{\partial x_n} & \cdots & \frac{\partial f_n}{\partial x_n} \end{bmatrix}$$

We denote by $C^{s}(S)$ the set of functions whose first s^{th} order partial derivatives are continuous on its domain S, and by $P^{s}([a,b])$ the set of single variate polynomial functions with degree s on an interval $[a,b] \subseteq \mathbb{R}$, e.g., $P^{1}([0,1])$ is the set of all linear functions on the interval [0,1].

2.2. The model. We consider the setting of a monopolist selling his products to incoming customers during a finite selling season, aiming to maximize his total expected revenue. There are n types of products, each of which is made up of a combination of a subset of m types of resources. For example, in the airline setting, a product refers to a multi-flight itinerary and a resource refers to a seat in a jet of a single-leg flight; in the hotel setting, a product refers to a multi-day stay and a resource refers to a one-night stay at a particular room. We denote by $A = [A_{ij}] \in \mathbb{R}^{m \times n}$ the resource consumption matrix, which characterizes the types and amounts of resources needed by each product (i.e., a single unit of product j requires A_{ij} units of resource i). Without loss of generality, we assume that the matrix A has full row rank. (If this is not the case, then we

can first apply the standard row elimination procedure to delete the redundant rows. See Jasin [21].) We denote by $C \in \mathbb{R}^m$ the vector of initial capacity levels of all resources at the beginning of the selling season. Since, in many industries (e.g., hotels and airlines), replenishment of resources during the selling season is either too costly or simply not feasible, following the standard model in the literature [19], we will assume that the seller has no opportunity to procure additional units of resources during the selling season. In addition, we also assume without loss of generality that the remaining resources at the end of the selling season have zero salvage value.

The selling season is divided into T discrete periods, indexed by t = 1, 2, ..., T. At the beginning of period t, the seller first decides the price $p_t = (p_{t,1}; \ldots; p_{t,n})$ for his products, where p_t is chosen from a convex and compact set $\mathcal{P} = \bigotimes_{l=1}^{n} [p_l, \bar{p}_l] \subseteq \mathbb{R}^n$ of feasible price vectors. The posted price p_t , in turn, induces a demand, or sale, for one of the products with a certain probability. Here, we implicitly assume that at most one sale for one product occurs in each period. (We have made this assumption and chosen to focus on discrete time model to simplify the presentation of the analysis. Our analysis can be easily extended to either a discrete-time model with bounded demand arrivals in each period or continuous-time model with compound Poisson process (see Chen et al. [13] for more details).) Let $\Delta^{n-1} := \{(x_1; \ldots; x_n) \in \mathbb{R}^n | \sum_{i=1}^n x_i \leq 1, \text{ and } x_i \geq 0 \text{ for all } i\}$ denote the standard (n-1)-simplex. Let $\lambda^*(.): \mathcal{P} \to \Delta^{n-1}$ denote the induced demand rate or purchase probability vector; we also call $\lambda^*(.)$ the underlying demand function. Contrary to most existing RM literature where it is assumed that the seller knows $\lambda^*(.)$ a priori, in this paper, we simply assume that this function can be estimated using statistical learning procedures. Let $\Lambda_{\lambda^*} := \{\lambda^*(p) : p \in \mathcal{P}\}$ denote the convex and compact set of feasible demand rates and let $D_t(p_t) = (D_{t,1}(p_t); \ldots; D_{t,n}(p_t))$ denote the vector of realized demand in period t under price p_t . It should be noted that, although demands for different products in the same period are not necessarily independent, demands over different periods are assumed to be independent (i.e., D_t only depends on the posted price p_t in period t). By definition, we have $D_t(p_t) \in \mathcal{D} := \{D : \sum_{j=1}^n D_j \leq 1, D_j \in \{0,1\} \text{ for all } j\}$ and $\mathbb{E}[D_t(p_t)] = \lambda^*(p_t)$. This allows us to write $D_t(p_t) = \lambda^*(p_t) + \Delta_t(p_t)$, where $\Delta_t(p_t)$ is a zero-mean random vector. For notational simplicity, whenever it is clear from the context which price p_t is being used, we will simply write $D_t(p_t)$ and $\Delta_t(p_t)$ as D_t and Δ_t respectively. The sequence $\{\Delta_t\}_{t=1}^T$ will play an important role in our analysis later. Define the revenue function $r^*(p) := p \cdot \lambda^*(p)$ to be the one-period expected revenue that the seller can earn under price p. It is typically assumed in the literature that $\lambda^*(.)$ is invertible (see the regularity assumptions below). We can then write $r^*(p) = p \cdot \lambda^*(p) = p^*(\lambda) \cdot \lambda = r^*_{\lambda}(\lambda)$ to emphasize the dependency of revenue on demand rate instead of on price. We make the following regularity assumptions on $\lambda^*(.)$, $r^*(.)$ and $r^*_{\lambda}(.)$:

REGULARITY ASSUMPTIONS. There exists positive constants \bar{r} , $\underline{v} < \bar{v}$ such that:

R1. $\lambda^*(.): \mathcal{P} \to \Lambda_{\lambda^*}$ is in $\mathsf{C}^2(\mathcal{P})$ with Lipschitz continuous second order partial derivatives, and it has an inverse function $p^*(.): \Lambda_{\lambda^*} \to \mathcal{P}$ that is in $\mathsf{C}^2(\Lambda_{\lambda^*})$;

R2. There exists a set of turn-off prices $p_j^{\infty} \in [\underline{p}_j, \overline{p}_j]$ for j = 1, ..., n such that for any $p = (p_1; \ldots; p_n), p_j = p_j^{\infty}$ implies that $\lambda_j^*(p) = 0$.

R3. $||r_{\lambda}^{*}(.)||_{\infty} \leq \bar{r}, r_{\lambda}^{*}(.)$ is strongly concave, and all the eigenvalues of $\nabla^{2}r_{\lambda}^{*}(\lambda)$ are between $-\bar{v}$ and $-\underline{v}$ for all $\lambda \in \Lambda_{\lambda^{*}}$.

Assumption R1 is fairly natural and is easily satisfied by many popular demand functions such as linear, logit, and exponential functions. Assumption R2 is common in the literature. (See Besbes and Zeevi [6] and Wang et al. [30] for similar assumptions.) Its purpose is to allow the seller to effectively shut down the demand for any product whenever needed, e.g., in the case of stockout. (The existence of such turn-off prices follows naturally when we consider truncated demand functions. It is also possible to consider an unbounded set of feasible prices instead of the compact set we assume above, with a potentially infinite turn-off price; in such setting, our results still hold.) As for Assumption R3, the boundedness of $r_{\lambda}^*(.)$ follows from the compactness of Λ_{λ^*} and the continuity of $r_{\lambda}^{*}(.)$. The strong concavity of $r_{\lambda}^{*}(.)$ is a standard assumption in the literature and is satisfied by many commonly used demand functions such as linear, exponential, and logit functions. It should be noted that although some of these functions, such as logit, do not naturally correspond to a concave revenue function when viewed as a function of p, they are nevertheless concave when viewed as a function of λ . This highlights the benefit of treating revenue as a function of demand rate instead of as a function of price.

In addition, following Besbes and Zeevi [7] and the literature on nonparametric estimation, we will assume that the function $\lambda^*(.)$ has a certain level of *smoothness*. Let \bar{s} denote the largest integer such that $\left|\partial^{a_1,\ldots,a_n}\lambda_j^*(p)/\partial p_1^{a_1}\ldots\partial p_n^{a_n}\right|$ is uniformly bounded for all $j \in [1,n]$ and $0 \le a_1,\ldots,a_n \le \bar{s}$. We call \bar{s} the *smoothness index*. We make the following smoothness assumptions:

NONPARAMETRIC FUNCTION SMOOTHNESS ASSUMPTIONS.

N1. $\bar{s} \ge 2$.

N2. There exists a constant W > 0 such that for all $j \in \overline{[1,n]}$ and $p \in \mathcal{P}$ and integers $0 \leq a_1, \ldots, a_n \leq \overline{s}$, we have $\left| \frac{\partial^{a_1, \ldots, a_n} \lambda_j^*(p)}{\partial p_1^{a_1} \ldots \partial p_n^{a_n}} \right| \leq W$.

The above assumptions are fairly mild and are satisfied by most commonly used demand functions, including linear, polynomial with higher degree, logit, and exponential with a bounded domain of feasible prices. We note that very similar assumptions are also made in Besbes and Zeevi [7]. More broadly, this type of smoothness assumptions are commonly made in the nonparametric estimation literature in statistics (see, for example, [20]). The smoothness index \bar{s} indicates the level of difficulty in estimating the corresponding demand function: The larger the value of \bar{s} , the smoother the demand function, and the easier it is to estimate its shape because its value cannot have a drastic local change.

2.3. Admissible controls and the induced probability measures. Let $D_{1:t} := (D_1, D_2, \dots, D_t)$ and $p_{1:t} := (p_1, p_2, \dots, p_t)$ denote respectively the observed vectors of demand and price realizations up to and including period t. Let \mathcal{H}_t denote the σ -field generated by $D_{1:t}$ and $p_{1:t}$. We define a control π as a sequence of functions $\pi = (\pi_1, \pi_2, \dots, \pi_T)$, where π_t is a \mathcal{H}_{t-1} -measurable mapping that maps the history $D_{1:t-1}$ and $p_{1:t-1}$ to a distribution of price vectors on $\bigotimes_{j=1}^n [\underline{p}_j, \overline{p}_j]$, and the price to be used in period t under policy π is drawn from this distribution $\pi_t(\mathcal{H}_{t-1})$. This class of controls is referred to as non-anticipating controls because the decision in each period depends only on the accumulated information up to the beginning of the period. Under control π , the seller sets the price in period t equal to $p_t^{\pi} = \pi_t(D_{1:t-1}, p_{1:t-1})$. Let Π denote the set of all admissible controls. That is,

$$\Pi := \left\{ \pi : \sum_{t=1}^{T} AD_t(p_t^{\pi}) \preceq C \text{ a.s., and } p_t^{\pi} = \pi_t(\mathcal{H}_{t-1}) \right\}.$$

Note that, even though the seller does not know the underlying demand function, the existence of turn-off prices $p_1^{\infty}, \ldots, p_n^{\infty}$ guarantees that the capacity constraints can be satisfied if the seller applies p_j^{∞} for product j as soon as the remaining capacity is not sufficient to produce one more unit of product j. Let $\mathbb{P}_{\lambda^*,t}^{\pi}$ denote the induced probability measure under an admissible control $\pi \in \Pi$, i.e.,

$$\mathbb{P}_{\lambda^*,t}^{\pi}(d_{1:t}) = \mathbb{P}_{\lambda^*,t}^{\pi}(D_{1:t} = d_{1:t}) = \prod_{s=1}^t \left[\left(1 - \sum_{j=1}^n \lambda_j^*(p_s^{\pi}) \right)^{\left(1 - \sum_{j=1}^n d_{s,j}\right)} \prod_{j=1}^n \lambda_j^*(p_s^{\pi})^{d_{s,j}} \right],$$

where $p_s^{\pi} = \pi_s(d_{1:s-1}, p_{1:s-1})$ and $d_s = [d_{s,j}] \in \mathcal{D}$ for all $s = 1, \ldots, t$. (By definition, the term $1 - \sum_{j=1}^n \lambda_j^*(p_s^{\pi})$ can be interpreted as the probability of no-purchase in period s under price p_s^{π} .)

For notational simplicity, we will simply write $\mathbb{P}_{\lambda^*,T}^{\pi}$ as $\mathbb{P}_{\lambda^*}^{\pi}$ and denote by $\mathbb{E}_{\lambda^*}^{\pi}$ the expectation with respect to the probability measure $\mathbb{P}_{\lambda^*}^{\pi}$. The total expected revenue under $\pi \in \Pi$ is $R^{\pi} = \mathbb{E}_{\lambda^*}^{\pi} [\sum_{t=1}^{T} p_t^{\pi} \cdot D_t(p_t^{\pi})]$.

2.4. The deterministic formulation and performance metric. The following optimization is the deterministic analog of the original stochastic pricing problem:

$$J^{D} := \max_{p_{t} \in \mathcal{P}} \left\{ \sum_{t=1}^{T} r^{*}(p_{t}) : \sum_{t=1}^{T} A\lambda^{*}(p_{t}) \preceq C \right\},$$

or equivalently,
$$J^{D} := \max_{\lambda_{t} \in \Lambda_{\lambda^{*}}} \left\{ \sum_{t=1}^{T} r^{*}_{\lambda}(\lambda_{t}) : \sum_{t=1}^{T} A\lambda_{t} \preceq C \right\}.$$

By assumption R3, the second optimization above is a convex program and can be efficiently solved. (To avoid triviality, we assume that both optimizations are feasible.) It can be shown that J^D is in fact an upper bound for the total expected revenue under any admissible control. That is, $R^{\pi} \leq J^D$ for all $\pi \in \Pi$. (See Besbes and Zeevi [7] for more details.) This allows us to use J^D as a benchmark to quantify the performance of any admissible pricing control. In this paper, we follow the standard convention and define the *expected revenue loss* of an admissible control $\pi \in \Pi$ as $\rho^{\pi} := J^D - R^{\pi}$. Since $r^*_{\lambda}(.)$ is strongly concave, by Jensen's inequality, it can be shown that the optimal solutions of J^D are static, i.e., $p_t = p^D$ and $\lambda_t = \lambda^D$ for all t, where p^D and λ^D can be obtained by solving the following "one-period" optimizations, respectively:

$$(\mathbf{P}) \qquad r^{D} := \max_{p \in \mathcal{P}} \left\{ r^{*}(p) : A\lambda^{*}(p) \preceq \frac{C}{T} \right\},$$

and,
$$(\mathbf{P}_{\lambda}) \qquad r^{D} := \max_{\lambda \in \Lambda_{\lambda^{*}}} \left\{ r^{*}_{\lambda}(\lambda) : A\lambda \preceq \frac{C}{T} \right\}.$$

Note that $p^D = p^*(\lambda^D)$ and $Tr^D = J^D$. Moreover, the optimal dual variables that correspond to the capacity constraints in **P** are the same as the optimal dual variables that correspond to the capacity constraints in \mathbf{P}_{λ} ; we denote these dual variables as μ^D . Let $\mathsf{Ball}(x,r)$ denote a closed Euclidean ball centered at x with radius r. We state our fourth regularity assumption below:

R4. (INTERIOR ASSUMPTION) There exists $\phi > 0$ such that $\mathsf{Ball}(p^D, \phi) \subseteq \mathcal{P}$.

Assumption R4 is sufficiently mild. Intuitively, it states that the static price should neither be too low that it attracts too much demand nor too high that it induces no demand. A similar interior assumption has also been made in Jasin [21] and Chen et al. [12].

2.5. Asymptotic setting. Following the standard convention in the literature (e.g., Besbes and Zeevi [6] and Wang et al. [30]), in this paper, we will consider a sequence of increasing problems where the length of the selling season and the initial resource capacity are scaled by a factor of k > 0. To be precise, in the k^{th} problem, the length of the selling season and the initial capacity are given by kT and kC, respectively. (One can interpret k as the size of the problem. For example, in single-leg setting, C = 1 and k = 50 could correspond to a small jet with capacity 50 seats and k = 500 could correspond to a large jet with capacity 500 seats.) The optimal solutions for **P** and **P**_{λ} in the k^{th} problem are still p^D and λ^D ; the optimal dual solution corresponding to the capacity constraints in **P** and **P**_{λ} is still μ^D . But, the deterministic upper bound becomes $J^D(k) = kTr^D = kJ^D$. Let $\rho^{\pi}(k)$ denote the expected revenue loss under an admissible control $\pi \in \Pi$ for the problem with scaling factor k. We are primarily interested in identifying the order of $\rho^{\pi}(k)$ for large k. (Intuitively, one would expect that a better heuristic control will have an expected revenue loss that grows more slowly with respect to k.) The following notations will be used throughout the remainder of the paper. For any two functions $f: \mathbb{Z}_{++} \to \mathbb{R}$ and $g: \mathbb{Z}_{++} \to \mathbb{R}_+$, we write $f(k) = \Omega(g(k))$ if there exists M > 0 independent of k such that $f(k) \ge Mg(k)$. Similarly, we write $f(k) = \mathcal{O}(g(k))$ if there exists K > 0 independent of k such that $f(k) \le Kg(k)$.

3. Supporting technical results. In this section, we present some technical results on spline estimation and nonparametric stability analysis of a perturbed optimization problem, and we will also introduce a quadratic programming approximation of **P**. We will use these technical results in the analysis of NSC in §4. As noted in §1, many of these results are of independent interest and can potentially be used in different application areas.

3.1. Spline approximation. We first describe the problem of approximating a deterministic function from noiseless observations using spline approximation and then we will discuss the problem of estimating a function from noisy observations. Spline functions have been widely used in engineering to approximate complicated functions, and their popularity is primarily due to their flexibility in effectively approximating complex curve shapes [28]. This flexibility lies in the piecewise nature of spline functions—a spline function is constructed by attaching piecewise polynomial functions with a certain degree, and the coefficients of these polynomials are computed in such a way that a sufficiently high degree of smoothness is ensured in the places where the polynomials are connected. More formally, for all $l \in [1, n]$, let $\underline{p}_l = x_{l,0} < x_{l,1} \cdots < x_{l,d} < x_{l,d+1} = \overline{p}_l$ be a partition that divides $[p_l, \bar{p}_l]$ into d+1 sub-intervals of equal length where $d \in \mathbb{Z}_{++}$. Let $\mathcal{G} := \bigotimes_{l=1}^n \mathcal{G}_l$ denote a set of grid points, where $\mathcal{G}_l = \{x_{l,i}\}_{i=0}^{d+1}$. We define the function space of tensor-product polynomial splines of order $(s; \ldots; s) \in \mathbb{R}^n$ with a set of grid points \mathcal{G} as $\mathsf{S}(\mathcal{G}, s) := \bigotimes_{l=1}^n \mathsf{S}_l(\mathcal{G}_l, s)$, where $\mathsf{S}_l(\mathcal{G}_l,s) := \{f \in \mathsf{C}^{s-2}([p_l,\bar{p}_l]) : f \text{ is a single-variate polynomial of degree } s-1 \text{ on each sub-interval}\}$ $[x_{l,i}, x_{l,i+1})$, for all $i \in [0, d-1]$ and $[x_{l,d}, x_{l,d+1}]$. One of the key questions that spline approximation theory addresses is the following: Given an arbitrary function f that satisfies N1-N2, find a spline function $q^* \in \mathsf{S}(\mathcal{G}, s)$ that approximates f well. Among the various approaches, one of the most popular approximations is using the so-called *tensor-product B-Spline basis functions* [28]. This approach is based on the key observation that $S(\mathcal{G}, s)$ is actually a linear space of dimension $(s+d)^n$. This means that there exists a set of $(s+d)^n$ basis functions (this set is not necessarily unique) such that any function in $S(\mathcal{G}, s)$ can be represented by a linear combination of these basis functions. We choose to use tensor-product B-Spline basis functions, denoted by $\{N_{i_1,\ldots,i_n}(x_1,\ldots,x_n)\}_{i_1=1,\ldots,i_n=1}^{s+d,\ldots,s+d}$ as the set of basis functions. These functions are defined formally in the Technical Details part (a) below. Given the basis functions, for any spline function $g \in \mathsf{S}(\mathcal{G}, s)$, there exists a set of coefficients $\{c_{i_1,\ldots,i_n}\}_{i_1=1,\ldots,i_n=1}^{s+d}$ such that $g(x) = \sum_{i_1=1}^{s+d} \cdots \sum_{i_n=1}^{s+d} c_{i_1,\ldots,i_n} N_{i_1,\ldots,i_n}(x)$ for all $x \in \mathcal{P}$. Therefore, the problem of finding g^* is reduced to the problem of computing the coefficients for representing g^* , which we address below in the Technical Details part (b). For a more comprehensive discussion of this approach, see Schumaker [28].

Technical Details for Spline Approximation: The B-Spline Approach

(a) Tensor-product B-Spline Basis Functions.

Step 1: For each $l \in \overline{[1,n]}$, define an *extended partition* $\mathcal{G}_l^e := \{y_{l,i}\}_{i=1}^{2s+d}$, where

$$y_{l,1} = \dots = y_{l,s} = x_{l,0}, \ y_{l,s+1} = x_{l,1}, \dots, y_{l,s+d} = x_{l,d}, \ y_{l,s+d+1} = \dots = y_{l,2s+d} = x_{l,d+1}$$

Step 2: For $i_l \in \overline{[1, s+d]}, l \in \overline{[1, n]}$, define the *tensor-product B-Spline basis function* as $N_{i_1, \dots, i_n}(x_1, \dots, x_n) = \prod_{l=1}^n N_{l, i_l}^s(x_l)$, where

$$N_{l,i}^{s}(x_{l}) = \begin{cases} (-1)^{s} (y_{l,i+s} - y_{l,i}) [y_{l,i}, \dots, y_{l,i+s}] (x_{l} - y)_{+}^{s-1}, \text{ if } y_{l,i} \le x_{l} < y_{l,i+s} \\ 0, & \text{otherwise} \end{cases}$$

for all $x_l \in [p_l, \bar{p}_l]$ for all $l \in [1, n]$ and for all $i \in [1, s+d]$, where $(x_l - y)_+ = \max\{0, x_l - y\}$, and $[t_1, \ldots, t_{r+1}]\bar{f}(y)$ is the r^{th} order divided difference of a single variate real function f over the points $t_1 < t_2 < \cdots < t_r < t_{r+1}$ defined as follows (see Definition 2.49 and Theorem 2.50 in Schumaker [28] for more discussion):

$$[t_1, \dots, t_{r+1}]f(y) := \sum_{i=1}^{r+1} \frac{f(t_i)}{\prod_{j=1, j \neq i}^{r+1} (t_i - t_j)}$$

(b) Calculating the Linear Coefficients.

Step 1: For $l \in \overline{[1,n]}, i \in \overline{[1,s+d]}$, let

$$\tau_{l,i,j} = y_{l,i} + (y_{l,i+s} - y_{l,i})\frac{j-1}{s-1} \quad \text{and} \quad \beta_{l,i,j} = \sum_{v=1}^{j} \frac{(-1)^{v-1}}{(s-1)!} \phi_{l,i,s}^{(s-v)}(0) \psi_{l,i,j}^{(v-1)}(0), \quad \text{for } j \in \overline{[1,s]},$$

where $\phi_{l,i,s}(t) = \prod_{r=1}^{s-1} (t - y_{l,i+r}), \ \psi_{l,i,j}(t) = \prod_{r=1}^{j-1} (t - \tau_{l,i,r}), \ \psi_{l,i,1}(t) \equiv 1.$ **Step 2:** For any $f \in C^0(\mathcal{P})$, define a set of linear functionals $\{\gamma_{l,i} : C^0([p_l, \bar{p}_l]) \to \mathbb{R}\}_{l=1,i=1}^{n,s+d}$ as:

$$\gamma_{l,i}f := \sum_{j=1}^{s} \beta_{l,i,j}[\tau_{l,i,1}, \dots, \tau_{l,i,j}]f = \sum_{j=1}^{s} \beta_{l,i,j} \sum_{r=1}^{j} \frac{f(x_1, \dots, x_{l-1}, \tau_{l,i,r}, x_{l+1}, \dots, x_n)}{\prod_{s=1, s \neq r}^{j} (\tau_{l,i,r} - \tau_{l,i,s})}$$

where f is viewed as a single variate function of x_l here, and the second equality follows by Theorem 2.50 in Schumaker [28] (note that for any given l and i, $\tau_{l,i,1}, \ldots, \tau_{l,i,s}$ are pairwise distinct). Define another set of linear functionals $\{\gamma_{i_1,\ldots,i_n}\}_{i_1=1,\ldots,i_n=1}^{s+d,\ldots,s+d}$ such that

$$\gamma_{i_1,\ldots,i_n}f=\gamma_{1,i_1}\circ\gamma_{2,i_2}\cdots\circ\gamma_{n,i_n}f,$$

where γ_{l,i_l} is understood as being applied to f as a function of x_l . By the construction of γ_{l,i_l} , basic algebra yields:

$$\gamma_{i_1,\dots,i_n} f = \sum_{j_1=1}^s \sum_{r_1=1}^{j_1} \cdots \sum_{j_n=1}^s \sum_{r_n=1}^{j_n} \frac{f(\tau_{1,i_1,r_1},\dots,\tau_{n,i_n,r_n}) \prod_{l=1}^n \beta_{l,i_l,j_l}}{\prod_{l=1}^n \prod_{l=1}^{j_l} \prod_{l=1,s_l \neq r_l}^{j_l} (\tau_{l,i_l,r_l} - \tau_{l,i_l,s_l})}$$

Step 3: Define a linear operator $\mathcal{L}_l : \mathsf{C}^0([\underline{p}_l, \bar{p}_l]) \to \mathsf{S}_l(\mathcal{G}_l, s)$ as $\mathcal{L}_l f(x_l) = \sum_{i=1}^{s+d} (\gamma_{l,i} f) N_{l,i}^s(x_l)$, for all $l \in [1, n]$. Similarly, define a linear operator $\mathcal{L} : \mathsf{C}^0(\mathcal{P}) \to \mathsf{S}(\mathcal{G}, s)$ as

$$\mathcal{L}f(x_1, \dots, x_n) = \sum_{i_1=1}^{s+d} \cdots \sum_{i_n=1}^{s+d} (\gamma_{i_1, \dots, i_n} f) N_{i_1, \dots, i_n}(x_1, \dots, x_n)$$

Note that $\mathcal{L} = \mathcal{L}_1 \circ \mathcal{L}_2 \circ \cdots \circ \mathcal{L}_n$, where this composition of linear operators is understood as \mathcal{L}_l being applied to a function of x_l .

Step 4: Set $g^* = \mathcal{L}f$.

Spline approximation with noisy observations. We now discuss the estimation of demand function $\lambda^*(.)$ using spline approximation under noisy observations. Let $\tilde{\mathcal{G}} := \{(\tau_{1,i_1,j_1}; \ldots; \tau_{n,i_n,j_n}): i_l \in \overline{[1,s+d]}, j_l \in \overline{[1,s]} \text{ for all } l \in \overline{[1,n]}\}$. Note that the constants $\{\gamma_{i_1,\ldots,i_n}\lambda_j^*\}_{i_1=1,\ldots,i_n=1}^{s+d,\ldots,s+d}$ depend on $\lambda_j^*(.)$ only via $\lambda_j^*(p), p \in \tilde{\mathcal{G}}$. So, if the seller could observe the demand rate of product j under prices in $\tilde{\mathcal{G}}$, he could construct an approximation of $\lambda_j^*(.)$ using a linear combination of tensor-product B-splines. In our problem, the seller cannot observe $\lambda_j^*(p)$ for $p \in \tilde{\mathcal{G}}$, but only its noisy observation $D_j(p) = \lambda_j^*(p) + \Delta_j$. To address this, we use empirical mean as a surrogate for $\lambda_j^*(p)$ and propose the following *Spline Estimation* algorithm to estimate the demand.

Spline Estimation

Input Parameters: L_0, n, s ; Tuning Parameter: dAlgorithm:

- Step 1: Estimate $\lambda^*(p)$ at points $p \in \tilde{\mathcal{G}}$. For each $p \in \tilde{\mathcal{G}}$ a. Apply price $p \ L_0$ times
 - b. Let $\lambda(p)$ be the sample mean of the L_0 observations.

Step 2: Construct spline approximation.

a. For all $j \in \overline{[1,n]}$ and $i_l \in \overline{[1,s+d]}$, $l \in \overline{[1,n]}$, calculate coefficients c_{i_1,\ldots,i_n}^j as:

$$c_{i_1,\dots,i_n}^j = \sum_{j_1=1}^s \sum_{r_1=1}^{j_1} \cdots \sum_{j_n=1}^s \sum_{r_n=1}^{j_n} \frac{\tilde{\lambda}_j(\tau_{1,i_1,r_1},\dots,\tau_{n,i_n,r_n}) \prod_{l=1}^n \beta_{l,i_l,j_l}}{\prod_{l=1}^n \prod_{s_l=1,s_l \neq r_l}^{j_l} (\tau_{l,i_l,r_l} - \tau_{l,i_l,s_l})}$$

b. Construct a tensor-product spline function $\tilde{\lambda}(p) = (\tilde{\lambda}_1(p); \ldots; \tilde{\lambda}_n(p))$, where

$$\tilde{\lambda}_j(p) = \sum_{i_1=1}^{s+d} \cdots \sum_{i_n=1}^{s+d} c_{i_1,\dots,i_n}^j N_{i_1,\dots,i_n}(p).$$

Note that the algorithm above conducts $\tilde{L}_0 := L_0(s+d)^n s^n$ samples. The idea of *Spline Estimation* is as follows. In Step 1, we apply each $p \in \tilde{\mathcal{G}}$ as many as L_0 times and calculate its empirical mean $\tilde{\lambda}(p)$. In Step 2, we approximate the underlying demand function $\lambda^*(.)$ using a spline function. In particular, we use a modified version of B-Spline approach by replacing the actual function value $\lambda^*(p)$ ($p \in \tilde{\mathcal{G}}$) with its empirical mean $\tilde{\lambda}(p)$. Note that our estimation approach is different from the so-called *Spline Regression* (see Gyorfi et al. [20]). While Spline Regression uses Least Squares to compute the linear coefficients for each of the spline basis function, we use the empirical means at sample points and a specific linear operator (originally devised and analyzed in the deterministic approximation theory of spline functions, see Schumaker [28]) to compute the linear coefficients. We choose to use Spline Estimation in our heuristic instead of Spline Regression because it allows us to use existing results on Spline Approximation Theory to derive the large deviation bounds for Spline Estimation in Lemma 1. We suspect that similar results also hold for Spline Regression. Let $a \wedge b = \min\{a, b\}$. The following lemma shows how well $\tilde{\lambda}(.)$ approximates $\lambda^*(.)$.

LEMMA 1. Set $d = \lceil (L_0^{1/2}/\log k)^{1/(s+n/2)} \rceil$ and let $L_0 \ge \log^3 k$ be a positive integer that may depend on k. Suppose that $s \ge 2$. There exist positive constants Ψ_r for each $r \in [0, (s-2) \land \overline{s}]$ and K independent of $k \ge 1$ such that for all $j \in [1,n]$ and $r_l \in \mathbb{Z}_+, l \in [1,n]$ satisfying $\sum_{l=1}^n r_l = r$ we have:

$$\mathbb{P}^{\pi}\left(\left\|\left|\frac{\partial^{r}(\lambda_{j}^{*}(.)-\tilde{\lambda}_{j}(.))}{\partial p_{1}^{r_{1}}\ldots\partial p_{n}^{r_{n}}}\right\|\right\|_{\infty} \geq \Psi_{r}\left(\frac{\log k}{\sqrt{L_{0}}}\right)^{\frac{s\wedge\bar{s}-r}{s+n/2}}\right) \leq K\exp(-\log^{2}k).$$
(1)

Proof: See Appendix A. \Box

The condition $L_0 \geq \log^3 k$ implies $(\log k/\sqrt{L_0})^{\frac{s/\bar{s}-r}{s+n/2}} \to 0$ as $k \to \infty$. This means that the difference between the r^{th} order partial derivatives of the underlying demand function and the spline approximation is uniformly small with a high probability for large k. When r = 0, our bound becomes a large deviation bound for the function estimate itself and is similar to the known bound for Spline Regression. (Suppose we set $s = \bar{s}$ and re-write our bound in a Hoeffding-type form, see Remark 1 below. Integrating the right hand side with respect to x over \mathbb{R}_+ , we obtain the ∞ -risk of Spline Estimation which is of order $(1/\sqrt{\tilde{L}_0})^{2\bar{s}/(2\bar{s}+n)}$. This is to be compared with the well-known 2-risk of Spline Regression which is of order $(\sqrt{\log \tilde{L}_0/\tilde{L}_0})^{2\bar{s}/(2\bar{s}+n)}$, see Corollary 15.1 in Gyorfi et al. [20].) We want to stress that the large deviation bound for the function estimate itself is *not* sufficient for our purpose. Specifically, we need additional large deviation bounds for the first and second order partial derivatives of the estimated demand function, as in Lemma 1, in order to conduct a stability analysis of the deterministic optimization problem (**P**) in our analysis later.

REMARK 1 (INTERPRETING (1) AS A HOEFFDING-TYPE ERROR BOUND). Hoeffding-type error bounds commonly appear in statistical estimations. Informally, they relate a measure of estimation error (e.g., 2-norm of the parameter estimation error in parametric models) with the number of samples L_0 in the following way:

$$\mathbb{P}(\operatorname{Error} \geq x) \leq C_1 \exp(-C_2 L_0 x^2),$$

for some constants C_1 and C_2 that are independent of x and L_0 . Note that, in Hoeffding-type of inequality, the right hand side converges to zero as x tends to zero and the variable x shows up as a quadratic term in the exponent. In contrast, when we write (1) into a similar form, we obtain

$$\mathbb{P}^{\pi}\left(\left\|\left|\frac{\partial^{r}(\lambda_{j}^{*}(.)-\tilde{\lambda}_{j}(.))}{\partial p_{1}^{r_{1}}\ldots\partial p_{n}^{r_{n}}}\right\|_{\infty}\geq x\right)\leq K\exp(-\bar{\Psi}_{r}L_{0}x^{\frac{2s+n}{s\wedge\bar{s}-r}}),$$

where $\bar{\Psi}_r = \Psi_r^{-\frac{2s+n}{s\wedge s-r}}$. Due to the well-known curse of dimensionality in nonparametric function estimation, the right hand side of our inequality does *not* tend to zero as fast as a typical Hoeffding-type inequality (i.e., $x^2 > x^{\frac{2s+n}{s\wedge s-r}}$ when x is small). Moreover, the convergence rate on the right hand side also depends on model parameters. In particular, it is decreasing in n and r, and is increasing in \bar{s} . This makes intuitive sense: as the problem dimension n increases, estimation becomes more difficult; as the order of derivative r decreases or as the smoothness index \bar{s} increases, the underlying demand function (or the partial derivative of the underlying demand function) becomes smoother and is easier to estimate. The convergence rate is increasing in s when $s \leq \bar{s}$ because higher s allows more flexibility in spline approximation. Interestingly, when $s > \bar{s}$, the convergence rate actually decreases in s. This is possibly due to the fact that, when s is higher than smoothness index \bar{s} , the extra flexibility introduces unnecessary complexity (i.e., redundant linear coefficients to be estimated), which leads to more sampling.

3.2. Stability analysis. In this subsection, we first present a nonparametric stability result for a class of optimization problems, and then apply this result to the perturbation analysis of our deterministic optimization **P**. Consider the following non-linear optimization problems:

(**NP**)
$$\max_{x \in \mathcal{X}} \left\{ f(x) : Ug(x) \preceq V \right\} \text{ and } (\tilde{\mathbf{NP}}(\delta)) \qquad \max_{x \in \mathcal{X}} \left\{ \tilde{f}(x) : U\tilde{g}(x) \preceq V - \delta \right\}.$$

where \mathcal{X} is a convex compact subset of \mathbb{R}^n , $f: \mathcal{X} \to \mathbb{R}$ and $g: \mathcal{X} \to \mathbb{R}^n$ are both twice continuously differentiable functions, $\tilde{f}: \mathcal{X} \to \mathbb{R}$ and $\tilde{g}: \mathcal{X} \to \mathbb{R}^n$ are continuously differentiable approximations

15

of f and g, $\delta \in \mathbb{R}^m$, $V \in \mathbb{R}^m$, and U is an m by n non-negative matrix that has full row rank. Let x^* and \tilde{x}_{δ} denote the optimal solution of **NP** and $\tilde{NP}(\delta)$ respectively (i.e., if they are feasible). We state a useful stability result.

PROPOSITION 1. Suppose that the following conditions hold:

(i) g(.) has a twice continuously differentiable inverse function $g^{-1}(.): \mathcal{Y} \to \mathcal{X}$ where $\mathcal{Y} := g(\mathcal{X})$ is a convex compact subset of \mathbb{R}^n ;

(ii) $f(g^{-1}(.)): \mathcal{Y} \to \mathbb{R}^n$ is strongly concave;

(*iii*) **NP** is feasible;

(iv) x^* is in the interior of \mathcal{X} .

Then, there exist $\bar{\delta} > 0$ and K > 0 such that for all δ , $\tilde{f}(.)$ and $\tilde{g}(.)$ satisfying $||Ug(.) - U\tilde{g}(.) + \delta||_{\infty} \leq \bar{\delta}$, $\tilde{\mathbf{NP}}(\delta)$ is feasible and

$$||x^* - \tilde{x}_{\delta}||_2 \leq K \left(||(\nabla f(.) - \nabla \tilde{f}(.))'||_{\infty} + ||g(.) - \tilde{g}(.)||_{\infty} + ||\delta||_{\infty} \right).$$

Proof: See Appendix B. \Box

The above result can be viewed as a Lipschitz-type stability result for a family of nonparametric optimization problems. Per our discussions in §1, although stability analysis of parametric optimization problems has been intensively studied in the literature (e.g., Bonnans and Shapiro [10]), stability results for nonparametric optimization problems are very rare. (See Remark 2 for a brief discussion on the relationship between Proposition 1 and existing results on parametric stability analysis.) In our case, since the original unperturbed optimization can be transformed into a convex optimization, we can use a convexity argument to establish Proposition 1.

We now apply Proposition 1 to our deterministic optimization problem **P**. Using the spline approximate $\tilde{\lambda}(p)$ derived in §3.1, we can formulate an approximate optimization of **P** as follows:

$$(\tilde{\mathbf{P}}) \quad \tilde{r}^D := \max_{p \in \mathcal{P}} \left\{ \tilde{r}(p) : \ A \tilde{\lambda}(p) \preceq \frac{C}{T} \right\}$$

where $\tilde{r}(p) = p \cdot \tilde{\lambda}(p)$. Let \tilde{p}^D denote an optimal solution of $\tilde{\mathbf{P}}$ if it is feasible and let $\tilde{\lambda}^D = \tilde{\lambda}(\tilde{p}^D)$. The following lemma follows directly from Proposition 1 and provides a characterization of \tilde{p}^D .

LEMMA 2. Suppose that $s \geq 3$. Then, there exist constants $\bar{\delta} > 0$ and K > 0 such that if $||\lambda^*(.) - \tilde{\lambda}(.)||_{\infty} \leq \bar{\delta}$, $\tilde{\mathbf{P}}$ is feasible and $||p^D - \tilde{p}^D||_2 \leq K(||\lambda^*(.) - \tilde{\lambda}(.)||_{\infty} + ||(\nabla \lambda^*(.) - \nabla \tilde{\lambda}(.))'||_{\infty})$.

Proof: See Appendix C. \Box

Lemma 2 tells us that if demand estimation error is small, $\tilde{\mathbf{P}}$ is feasible and its optimal solution \tilde{p}^D lies in close proximity of p^D . This observation is crucial for our analysis later.

REMARK 2 (ON PROPOSITION 1 AND EXISTING PARAMETRIC STABILITY RESULT). The existing Lipschitz stability result of the optimal solution of a parameterized optimization problem (e.g., Theorem 5.53 part (a) in Bonnans and Shapiro [10]) can be viewed as a special case of our nonparametric Lipschitz-type stability result in Proposition 1. Let $\mathcal{U} \subseteq \mathbb{R}^q, q \in \mathbb{Z}_{++}$, be a compact parameter set. Suppose that the objective functions f and \tilde{f} come from a family of parameterized functions $\{f(.;u)\}_{u\in\mathcal{U}}$ where $f(.) = f(.;u_0)$ and $\tilde{f}(.) = f(.;v)$ for $u_0, v \in \mathcal{U}$. Also, suppose that the constraint functions g and \tilde{g} come from a family of parameterized functions $\{g(.;u)\}_{u\in\mathcal{U}}$ where $g(.) = g(.;u_0)$ and $\tilde{g}(.) = g(.;v)$ for $u_0, v \in \mathcal{U}$. For simplicity, assume $\delta = \mathbf{0}$. In the perturbation analysis of parametric optimization problems, f(.;.) and g(.;.) are typically assumed to be twice continuously differentiable, which means that $||(\nabla f(.;u_0) - \nabla f(.;v))'||_{\infty} = \mathcal{O}(||u_0 - v||_{\infty})$ and $||g(.;u_0) - g(.:v)||_{\infty} = \mathcal{O}(||u_0 - v||_{\infty})$. Applying Proposition 1 to this setting immediately yields Theorem 5.53 part (a) in Bonnans and Shapiro [10]. **3.3.** An approximate quadratic program. In this subsection, we introduce a quadratic program approximation of **P**. (This will be useful when we discuss our heuristic in §4.) The idea is simple: We approximate the objective of **P** with a quadratic function and its constraints with linear functions. Our objective here is to show that if the parameters of the quadratic and linear functions are correctly chosen, the resulting quadratic program will have the same optimal solution as **P** and it will possess some very useful stability properties. To begin with, we first linearize the constraints of **P**. Since the capacity constraints form an affine transformation of the demand function, we will simply linearize the demand function. For any $a \in \mathbb{R}^n, B \in \mathbb{R}^{n \times n}$, let B_1, \ldots, B_n be the columns of *B* and define $\theta_i = (a; B_1; \ldots; B_n) \in \mathbb{R}^{n^{2+n}}$, where the subscript ι stands for *linear demand*. We denote a linear demand function by $\lambda(p; \theta_i) = a + B'p$. Next, we discuss a quadratic approximation for the objective of **P**. For any $E \in \mathbb{R}, F \in \mathbb{R}^n, G \in \mathbb{R}^{n \times n}$, let G_1, \ldots, G_n denote the columns of *G* and define $\theta_o = (E; F; G_1; \ldots; G_n) \in \mathbb{R}^{n^{2+n+1}}$ where the subscript o stands for *objective*. We denote the resulting quadratic function by $q(p; \theta_o) = E + F'p + \frac{1}{2}p'Gp$. Finally, let $\theta = (\theta_o; \theta_\iota) \in \mathbb{R}^{2n^2+2n+1}$. For any $\theta \in \mathbb{R}^{2n^2+2n+1}, \delta \in \mathbb{R}^m$, we can define a quadratic program $\mathbf{QP}(\theta; \delta)$ as follows:

$$(\mathbf{QP}(\theta; \delta)) \qquad \max_{p \in \mathcal{P}} \left\{ q(p; \theta_o) : A\lambda(p; \theta_\iota) \preceq \frac{C}{T} - \delta \right\}.$$

If we choose the parameters θ and δ carefully, $\mathbf{QP}(\theta; \delta)$ can be a very good approximation of \mathbf{P} . Specifically, let $\theta_{\iota}^* = (a^*; B_1^*; \ldots; B_n^*)$, where $B^* := \nabla \lambda^*(p^D)$ and $a^* := \lambda^D - (B^*)'p^D$. Define an n by n symmetric matrix $H^* := B^* \nabla^2 r_{\lambda}^*(\lambda^D)(B^*)' - B^* - (B^*)'$. Then, one can verify that

$$H_{ij}^* = -(u_{ij}^*)'(B^*)^{-1}\lambda^D, \text{ where } u_{ij}^* = \left[\frac{\partial^2 \lambda_1^*(p^D)}{\partial p_i \partial p_j}; \dots; \frac{\partial^2 \lambda_n^*(p^D)}{\partial p_i \partial p_j}\right].$$
(2)

(See Appendix D for derivation.) Let $\theta_o^* = (E^*; F^*; G_1^*; \ldots; G_n^*)$ where

$$E^* := \frac{1}{2} (p^D)' H^* p^D, \ F^* := a^* - H^* p^D, \ G^* := B^* + (B^*)' + H^*,$$

and let $\theta^* := (\theta_o^*; \theta_\iota^*)$. Note that $\mathbf{QP}(\theta^*; \mathbf{0})$ is a very intuitive approximation of \mathbf{P} since the function $\lambda(p; \theta_\iota^*) = a^* + (B^*)'p = \lambda^D + (B^*)'(p - p^D)$ can be viewed as a linearization of $\lambda^*(.)$ at p^D . (Since $\nabla \lambda^*(p^D)$ is invertible as implied by R1 and R4, we can write $p(\lambda; \theta_\iota^*) = p^D + ((B^*)')^{-1}(\lambda - \lambda^D)$ as the inverse demand function). Note also that the gradients of the objective function and the constraints in $\mathbf{QP}(\theta^*; \mathbf{0})$ at p^D coincide with those in \mathbf{P} . By Karush-Kuhn-Tucker (KKT) optimality conditions, it can be shown that the optimal solution of $\mathbf{QP}(\theta^*; \mathbf{0})$ is the same as the optimal solution of \mathbf{P} . We formally state these results in Lemma 3 below. Let $p_{\delta}^D(\theta)$ and $\mu_{\delta}^D(\theta) = \lambda(p_{\delta}^D(\theta); \theta_{\iota})$.

LEMMA 3. There exist constants $\kappa > 0$, $\omega > 0$ and $\overline{\delta} > 0$ such that, for all $\theta_{\iota} \in \mathsf{Ball}(\theta_{\iota}^*, \overline{\delta})$, $\theta_o \in \mathsf{Ball}(\theta_o^*, \overline{\delta})$ and $\delta \in \mathsf{Ball}(\mathbf{0}, \overline{\delta})$, the following results hold:

(a) B is invertible and $||(B')^{-1}||_2 \leq \omega$;

 $\begin{array}{l} (b) \ \ For \ all \ p \in \mathcal{P} \ \ and \ \ for \ \ all \ i, j \in \overline{[1,n]}, \ ||\lambda(p;\theta_{\iota}) - \lambda(p;\theta_{\iota}^{*})||_{2} \leq \omega ||\theta_{\iota} - \theta_{\iota}^{*}||_{2} \ \ and \ |\frac{\partial \lambda_{j}}{\partial p_{i}}(p;\theta_{\iota}) - \frac{\partial \lambda_{j}}{\partial p_{i}}(p;\theta_{\iota})||_{2} \leq \omega ||\theta_{\iota} - \theta_{\iota}^{*}||_{2} \ \ and \ |\frac{\partial \lambda_{j}}{\partial p_{i}}(p;\theta_{\iota}) - p(\lambda';\theta_{\iota})||_{2} \leq \omega ||\lambda - \lambda'||_{2}; \\ (c) \ \ For \ \ all \ \lambda, \lambda' \in \lambda(\mathcal{P};\theta_{\iota}), \ ||p(\lambda;\theta_{\iota}) - p(\lambda';\theta_{\iota})||_{2} \leq \omega ||\lambda - \lambda'||_{2}; \\ (d) \ \ q(p(.;\theta_{\iota});\theta_{o}) \ \ is \ strongly \ concave. \\ (e) \ \ p^{D} = p_{0}^{D}(\theta^{*}), \ \lambda^{D} = \lambda_{0}^{D}(\theta^{*}), \ \mu^{D} = \mu_{0}^{D}(\theta^{*}); \\ (f) \ \ \mathbf{QP}(\theta;\delta) \ \ is \ feasible \ and \ has \ a \ unique \ optimal \ solution. \ Moreover, \ p_{\delta}^{D}(\theta) \in \mathsf{Ball}(p_{0}^{D}(\theta^{*}), \phi/2), \end{array}$

(f) $\mathbf{QP}(\theta; \delta)$ is feasible and has a unique optimal solution. Moreover, $p_{\delta}^{D}(\theta) \in \mathsf{Ball}(p_{0}^{D}(\theta^{*}), \phi/2)$, $\mathsf{Ball}(p_{\delta}^{D}(\theta), \phi/2) \subseteq \mathcal{P}, ||p_{0}^{D}(\theta^{*}) - p_{\delta}^{D}(\theta)||_{2} \le \kappa(||\theta^{*} - \theta||_{2} + ||\delta||_{2}), ||\lambda_{0}^{D}(\theta^{*}) - \lambda_{\delta}^{D}(\theta)||_{2} \le \kappa(||\theta^{*} - \theta||_{2} + ||\delta||_{2}), and the constraints of <math>\mathbf{QP}(\theta; \delta)$ that correspond to the rows $\{i : \mu_{0,i}^{D}(\theta^{*}) > 0\}$ are binding. Note that Lemma 3 part (f) not only establishes Lipschitz continuity of the optimal solution, but also provides additional results regarding the properties of the capacity constraints at the optimal solution. These play an important role in deriving a sharp performance bound of our heuristic.

REMARK 3 (ON THE QUADRATIC REVENUE FUNCTION APPROXIMATION). Note that, as the equation below shows, $q(p; \theta_o^*)$ can be viewed as the revenue function under the approximate linear demand function plus an additional correction term:

$$\begin{split} q(p;\theta_o^*) &= \frac{1}{2} (p^D)' H^* p^D + p'(a^* - H^* p^D) + \frac{1}{2} p'(B^* + (B^*)' + H^*) p \\ &= p'(a^* + (B^*)' p) + \frac{1}{2} (p - p^D)' H^* (p - p^D) \\ &= r(p;\theta_\iota^*) + \frac{1}{2} (p - p^D)' H^* (p - p^D), \end{split}$$

where $r(p; \theta_{\iota}^*) := p \cdot \lambda(p; \theta_{\iota}^*)$ is the natural revenue function under the approximate linear demand function. We add a correction term above in order to ensure that if we do a change of variables $p = p(\lambda; \theta_{\iota}^*)$ to change the pricing decision to demand rate decision, the resulting objective function is actually the second order Taylor's expansion of r_{λ}^* :

$$\begin{aligned} q(p(\lambda;\theta_{\iota}^{*});\theta_{o}^{*}) &= r(p(\lambda;\theta_{\iota}^{*});\theta_{\iota}^{*}) + \frac{1}{2}(p(\lambda;\theta_{\iota}^{*}) - p^{D})'H^{*}(p(\lambda;\theta_{\iota}^{*}) - p^{D}) \\ &= \lambda'(p^{D} + ((B^{*})')^{-1}(\lambda - \lambda^{D})) + \frac{1}{2}(\lambda - \lambda^{D})'(B^{*})^{-1}H^{*}((B^{*})')^{-1}(\lambda - \lambda^{D}) \\ &= \lambda'(p^{D} + ((B^{*})')^{-1}(\lambda - \lambda^{D})) - (\lambda - \lambda^{D})'((B^{*})')^{-1}(\lambda - \lambda^{D}) + \frac{1}{2}(\lambda - \lambda^{D})'\nabla^{2}r_{\lambda}^{*}(\lambda^{D})(\lambda - \lambda^{D}) \\ &= \lambda'p^{D} + (\lambda - \lambda^{D})'(B^{*})^{-1}\lambda^{D} + \frac{1}{2}(\lambda - \lambda^{D})'\nabla^{2}r_{\lambda}^{*}(\lambda^{D})(\lambda - \lambda^{D}) \\ &= r_{\lambda}^{*}(\lambda^{D}) + (\lambda - \lambda^{D})'(p^{D} + (B^{*})^{-1}\lambda^{D}) + \frac{1}{2}(\lambda - \lambda^{D})'\nabla^{2}r_{\lambda}^{*}(\lambda^{D})(\lambda - \lambda^{D}) \\ &= r_{\lambda}^{*}(\lambda^{D}) + \nabla r_{\lambda}^{*}(\lambda^{D})'(\lambda - \lambda^{D}) + \frac{1}{2}(\lambda - \lambda^{D})'\nabla^{2}r_{\lambda}^{*}(\lambda^{D})(\lambda - \lambda^{D}). \end{aligned}$$

Hence, in light of R3, $q(p(\lambda; \theta_{\iota}^*); \theta_{o}^*)$ is strongly concave in λ . This observation is important because it allows us to use the general result in Proposition 1 to derive perturbation result for the optimal primal and dual solutions of $\mathbf{QP}(\theta; \delta)$ (see condition (ii) in Proposition 1).

4. Main result. We are now ready to describe Nonparametric Self-adjusting Control (NSC) and discuss its asymptotic performance; the proof of this result is given in §5. NSC consists of an exploration procedure and an exploitation procedure. The exploration procedure uses the Spline Estimation algorithm discussed in §3.1 to construct a spline approximation $\tilde{\lambda}(.)$ of the underlying demand function $\lambda^*(.)$. This function $\tilde{\lambda}(.)$ is then used to construct a linear function $\lambda(.;\hat{\theta}_{\iota})$ that closely approximates $\lambda(.;\theta_{\iota}^*)$ in the neighborhood of p^D and a quadratic program that closely approximates **P**. During the exploitation phase, we use the optimal solution of the approximate quadratic program as baseline control and automatically adjust the price according to a predetermined price update rule. Further detail will be provided below. Recall that \tilde{L}_0 is the duration of the Spline Estimation algorithm. Let C_t denote the remaining capacity at the end of period t. Let $\hat{\theta} := (\hat{\theta}_o; \hat{\theta}_\iota)$, where $\hat{\theta}_\iota := (\hat{a}; \hat{B}_1; \ldots; \hat{B}_n), \hat{\theta}_o := (\hat{E}; \hat{F}; \hat{G}_1; \ldots; \hat{G}_n)$ for

$$\hat{B} := \nabla \tilde{\lambda}(\tilde{p}^D), \ \hat{a} := \tilde{\lambda} - \hat{B}'\tilde{p}^D, \ \hat{E} := \frac{1}{2}(\tilde{p}^D)'\hat{H}\tilde{p}^D, \ \hat{F} := \hat{a} - \hat{H}\tilde{p}^D,$$
$$\hat{G} := \hat{B} + \hat{B}' + \hat{H}, \ \text{and} \ \hat{H} = [\hat{H}_{ij}] \ \text{where} \ \hat{H}_{ij} := -\hat{u}'_{ij}\hat{B}^{-1}\tilde{\lambda}^D \ \text{and} \ \hat{u}_{ij} := \left[\frac{\partial^2 \tilde{\lambda}_1(\tilde{p}^D)}{\partial p_i \partial p_j}; \dots; \frac{\partial^2 \tilde{\lambda}_n(\tilde{p}^D)}{\partial p_i \partial p_j}\right].$$

(From §3.2, \tilde{p}^D is an optimal solution of $\tilde{\mathbf{P}}$.) Our proposed NSC heuristic is given below.

Nonparametric Self-adjusting Control (NSC) Tuning Parameters: d, L_0 Input parameters: n, s, Stage 1 (Exploration Phase 1 - Spline Estimation) a. For t = 1 to $L_0 \wedge T$: - If $C_{t-1} \prec A_j$ for some $j = 1, \ldots, n$, set $p_{t,j} = p_j^{\infty}$ for all $j = 1, \ldots, n$. - Otherwise, follow Step 1 in *Spline Estimation* algorithm. b. At the end of period $L_0 \wedge T$, do: - If $\tilde{L}_0 \geq T$, terminate NSC. - If $L_0 < T$ and $C_{\tilde{L}_0} \prec A_j$ for some $j = 1, \ldots, n$: - For all $t > \tilde{L}_0$, set $p_{t,j} = p_j^\infty$ for all $j = 1, \dots, n$. - Terminate NSC. - If $\tilde{L}_0 < T$ and $C_{\tilde{L}_0} \succeq A_j$ for all $j = 1, \ldots, n$: - Follow Step 2 in *Spline Estimation* algorithm to get $\lambda(.)$. - Go to Stage 2 below. Stage 2 (Exploration Phase 2 - Function Approximation) a. Solve $\tilde{\mathbf{P}}$ and obtain the optimizer \tilde{p}^D . b. Let $\delta := C/T - C_{\tilde{L}_0}/(T - \tilde{L}_0)$. c. Compute $\hat{a}, \hat{B}, \hat{E}, \hat{F}, \hat{G}, \hat{H}$ and $\hat{\theta} = (\hat{\theta}_o; \hat{\theta}_\iota)$. - If \hat{B} is invertible, go to Stage 2(d) below. - Otherwise, for $t = L_0 + 1$ to T: - If $C_{t-1} \succeq A_j$ for $j = 1, \ldots, n$, apply $p_t = \tilde{p}^D$. - Otherwise, for product j = 1 to n, do: - If $C_{t-1} \prec A_j$, set $p_{t,j} = p_j^{\infty}$. - Otherwise, set $p_{t,j} = p_{t-1,j}$. d. Solve $\mathbf{QP}(\theta; \delta)$ for its static price $p_{\delta}^{D}(\hat{\theta})$. Stage 3 (Exploitation) For $t = L_0 + 1$ to T: - Compute: $\hat{p}_t = p_{\delta}^D(\hat{\theta}) - \nabla_{\lambda} p(\lambda_{\delta}^D(\hat{\theta}); \hat{\theta}_{\iota}) \cdot \sum_{s=\tilde{L}_0+1}^{t-1} \frac{\tilde{\Delta}_s}{T-s}$, where $\tilde{\Delta}_t := D_t - \lambda(p_t; \hat{\theta}_{\iota})$. - If $\hat{p}_t \in \mathcal{P}$ and $C_{t-1} \succeq A_j$ for $j = 1, \dots, n$, apply $p_t = \hat{p}_t$. - Otherwise, for product j = 1 to n, do:

- If $C_{t-1} \prec A_j$, set $p_{t,j} = p_j^{\infty}$.
- Otherwise, set $p_{t,j} = p_{t-1,j}$.

We now explain the main ideas behind NSC. The exploitation part (Stage 3) of NSC is motivated by LRC heuristic developed in Jasin [21], which (roughly) uses

$$p_t = p^* \left(\lambda^D - \sum_{s=1}^{t-1} \frac{\Delta_s}{T-s} \right), \text{ where } \Delta_t = D_t(p_t) - \lambda^*(p_t)$$

and has a strong performance guarantee in the setting of known demand function. In our setting, the demand function $\lambda^*(.)$ is unknown (hence, the inverse demand function $p^*(.)$ is also unknown) and the sequence $\{\Delta_t\}_{t=1}^T$ is not observable. If we still wish to use LRC, an intuitive fix is to replace $\lambda^*(.)$ and $\{\Delta_t\}_{t=1}^T$ with their best estimates. This motivates the use of Spline Estimation in Stage 1 to compute an approximate demand function $\tilde{\lambda}(.)$. However, although $\tilde{\lambda}(.)$ can approximate $\lambda^*(.)$ well by tapping into the smoothness of $\lambda^*(.)$, the piece-wise nature of spline functions and the shape of the spline basis functions imply that $\tilde{\lambda}(.)$ may not be invertible, i.e., $\tilde{\lambda}(.)$ may not

admit a well-defined inverse demand function. But, this is crucial since LRC uses $p^*(.)$ to adjust the prices. This motivates us to use demand linearization in Stage 2. The objective of Stage 2 is to construct a linear function that closely approximates the linearization of the true demand function $\lambda^*(.)$ around p^D and construct a quadratic program that closely approximates **P** around its optimal solution p^{D} . We choose to use linear approximation of the demand function and quadratic approximation of the revenue function because, by Lemma 3 part (e), the optimal solution of the constructed approximate quadratic program coincides with the optimal solution of \mathbf{P} if the parameters are chosen to be $(\theta^*, \mathbf{0})$ (see §3.3 for more discussions). Although θ^* is unknown to the seller, we can utilize the spline approximation $\lambda(.)$ to construct parameters $\hat{\theta}$ that closely approximate θ^* . To see why this is so, note that if L_0 is carefully selected, the spline estimation procedure yields a spline function $\lambda(.)$ that closely approximates $\lambda^*(.)$, together with its first and second order partial derivatives (by Lemma 1), with a very high probability; this in turn indicates that any optimizer \tilde{p}^D of $\tilde{\mathbf{P}}$ lies in a close proximity of p^D (by Lemma 2). Since θ^* (resp. $\hat{\theta}$) can essentially be viewed as a function of p^D (resp. \tilde{p}^D), $\lambda^*(p^D)$ (resp. $\tilde{\lambda}(\tilde{p}^D)$) and its first and second order derivatives evaluated at p^D (resp. \tilde{p}^D), this suggests that $\hat{\theta} = (\hat{\theta}_{\alpha}; \hat{\theta}_{t})$ is a good approximation of $\theta^* = (\theta^*_{\alpha}; \theta^*_{t})$. It is worth stressing that Spline Estimation is crucial for determining reasonably good linear demand and quadratic revenue function approximations. As mentioned above, among all possible approximate linear demand functions, only those that are linearized at a point close to p^D are effective. (Similarly for the revenue functions.) To find a point that is close to p^D (i.e., \tilde{p}^D in our NSC) via optimizing the approximate deterministic pricing problem, we need to use Spline Estimation to get an approximate function that *uniformly* approximates the underlying demand function well.

Finally, after obtaining $\lambda(.;\hat{\theta}_{\iota})$, we replace $p^*(.)$ and Δ_t in LRC with $p(.;\hat{\theta}_{\iota})$ and $\tilde{\Delta}_t$. This leads to the price update formula in Stage 3:

$$\hat{p}_t = p_{\delta}^D(\hat{\theta}) - \nabla_{\lambda} p(\lambda_{\delta}^D(\hat{\theta}); \hat{\theta}_\iota) \cdot \sum_{s=\tilde{L}_0+1}^{t-1} \frac{\tilde{\Delta}_s}{T-s}$$

It is natural to expect that if $\lambda(.;\hat{\theta})$ approximates $\lambda^*(.)$ well, then NSC should retain the strong performance of LRC, this intuition is not immediately obvious and requires a mathematical justification. Note that, in addition to demand randomness, there are at least three sources of errors that affect the performance of NSC: (1) errors from functional estimation (i.e., due to estimating $\lambda^*(.)$ with $\tilde{\lambda}(.)$), (2) errors from function approximation (i.e., due to demand linearization and quadratic approximation), and (3) errors from systematic biases due to the terms $\{\tilde{\Delta}_t\}_{t=1}^T$. (In LRC, the perturbation term $\sum_{s=1}^{t-1} \Delta_s/(T-s)$ is unbiased because $\mathbb{E}_{\lambda^*}^{\pi}[\Delta_t] = 0$. In contrast, in NSC, the perturbation term $\sum_{s=\tilde{L}_0+1}^{t-1} \tilde{\Delta}_s/(T-s)$ is biased because $\mathbb{E}_{\lambda^*}^{\pi}[\Delta_t] \neq 0$. This means that we are systematically introducing new biases in each period. It is not a priori clear what kind of impact these biases will have on revenue performance.) Thus, despite the strong performance of LRC in the known demand function setting, it is not a priori clear whether self-adjusting alone, without re-optimizations and without re-estimations during Stage 3, is sufficient to reduce the impact of these errors on expected revenue loss. Interestingly, the following result states that the performance of NSC is close to the best achievable (asymptotic) performance bound.

THEOREM 1. Suppose that we use $s \ge 4$, $L_0 = \lceil (kT)^{(s+n/2)/(2s+n-2)} (\log(kT))^{(2s+n-4)/(2s+n-2)} \rceil$ and $d = \lceil (L_0^{1/2}/\log(kT))^{1/(s+n/2)} \rceil$. There exists a constant $M_1 > 0$ independent of k > 3 such that for all $s \ge 4$, we have

$$\rho^{NSC}(k) \le M_1 k^{\frac{1}{2} + \epsilon(n, s, \bar{s})} \log k, \text{ where } \epsilon(n, s, \bar{s}) = \frac{1}{2} \left(\frac{2s - 2(s \wedge \bar{s}) + n + 2}{2s + n - 2} \right).$$

Some comments are in order. First, unlike the heuristic control proposed in Besbes and Zeevi [7], which requires the knowledge of \bar{s} , NSC does *not* require the knowledge of the smoothness index \bar{s} . This is practically appealing because it is usually difficult to guess the smoothness index of a function when the function itself is unknown. Second, since most commonly used demand functions such as polynomial with arbitrary degree, logit, and exponential are infinitely differentiable (i.e., \bar{s} can be arbitrarily large), for any fixed $\epsilon > 0$, we can select integers $s > (n+2)/(4\epsilon) - (n-2)/2$ such that the performance under NSC is $\mathcal{O}(k^{1/2+\epsilon}\log k)$. Theoretically, this means that the asymptotic performance of NSC is very close to the best achievable performance lower bound of $\Omega(\sqrt{k})$. Third, despite the systematic biases it introduces, self-adjusting control in Stage 3 (surprisingly) plays a vital role in guaranteeing the stated performance bound; specifically, compared to a static pricing control, self-adjusting control has the ability to reduce the negative impact of systematic biases on revenue. To illustrate, consider the case where \bar{s} is arbitrarily large. Suppose that we only apply static price $p_t = p_{\delta}^D(\hat{\theta})$ throughout Stage 3, subject to capacity constraints. Then, under the optimally tuned L_0 and s, one can show that the resulting expected revenue loss is $\mathcal{O}(k^{2/3+\epsilon}\log k)$, which is significantly worse than the bound in Theorem 1. This underscores the importance of selfadjusting price update in reducing the expected revenue loss from $\mathcal{O}(k^{2/3+\epsilon}\log k)$ to $\mathcal{O}(k^{1/2+\epsilon}\log k)$. (Note that NSC's ability to reduce the negative impact of systematic biases on revenue is also observed in the re-optimized PAC heuristic developed in Jasin [22] for the quantity-based network RM problem with unknown demand arrival rates. The re-optimizations in Jasin [22] create a negative feedback mechanism which is similar to the self-adjusting pricing rule studied here.) Finally, to further validate the theoretical result in Theorem 1, we conduct a simple numerical study with two types of products and two types of resources. Table 1 shows that NSC performs well: For problems with a wide range of k, its relative revenue loss (i.e., $\rho^{\pi}(k)/J^{D}(k)$) is about 3 - 8% lower than the relative revenue loss of Algorithm 3 in Besbes and Zeevi [7]. To implement NSC for large-scale problems, the main computational burden lies in solving the nonlinear optimization **P** because $\lambda(p)$ is stitched together by many (not necessarily concave) multinomial function. (In fact, Algorithm 3 in Besbes and Zeevi [7] also suffers from this computational complexity. Moreover, we would also like to point out that, in contrast to local polynomial approximation used in Algorithm 3 in Besbes and Zeevi (2012), our spline approximation is globally differentiable and is more amenable to optimizations.) Thus, for problems with many different types of products and resources, one may want to optimize an approximation of $\tilde{\mathbf{P}}$ that is computationally more tractable. The question of which approximation should be used is an important and practically relevant one; however, it is beyond the scope of the current paper and we leave it for future research pursuit.

REMARK 4 (ON THE ANALYSIS OF UNCAPACITATED VS. CAPACITATED RM). Per our discussions in §1, most existing literature on joint learning and pricing focus on the setting of uncapacitated RM where there is no limit on the number of resources that can be used. In such setting, it has been repeatedly shown in the literature that the $\Omega(\sqrt{k})$ lower bound is actually tight for both single product and multiple products settings (see [6], [11], [23]). The presence of capacity constraints makes the problem significantly more challenging. To see this, note that, if we mis-calculate p^D by ϵ (i.e., we use $\tilde{p}^D = p^D + \epsilon$), by the strong concavity of $r_{\lambda}^*(.)$ and Lipschitz property of demand, $r^*(p^D) - r^*(\tilde{p}^D)$ is approximately on the order of ϵ^2 in the uncapacitated setting (because $\nabla r^*(p^D) = \mathbf{0}$ due to p^D being the unconstrained optimizer of $r^*(p)$). Thus, the expected revenue loss during T periods is $\mathcal{O}(T\epsilon^2)$. In contrast, in the capacitated setting, $\nabla r^*(p^D) \neq \mathbf{0}$ in general. This means that $r^*(p^D) - r^*(\tilde{p}^D)$ is on the order of ϵ , which implies that the expected revenue loss during T periods is $\mathcal{O}(T\epsilon)$. This is the reason why the analysis in uncapacitated RM is not directly applicable to capacitated RM.

REMARK 5 (APPLYING NSC IN DETERMINISTIC DEMAND ARRIVAL CASE). Although NSC is designed for the stochastic demand case, it can be readily adapted and applied in the deterministic

	Algorithm 3 in Boshog and Zoovi [7]			NSC		
	Algorithm 5 m Despes and Zeevi [7]			NSC		
k	Revenue loss	Stdev	Relative revenue $loss(\%)$	Revenue loss	Stdev	Relative revenue $loss(\%)$
500	5331	53	52.6	4681	11	46.2
1000	10490	114	51.8	8823	46	43.5
2000	20307	214	50.1	17320	85	42.7
3000	30074	300	49.5	26647	240	43.8
4000	39167	392	48.3	34421	289	42.5
5000	48922	466	48.3	40796	307	40.3
6000	57804	522	47.5	49578	477	40.8
7000	65459	594	46.1	57310	605	40.4
8000	74990	688	46.3	62319	640	38.4
9000	82891	721	45.4	70797	800	38.8
10000	89703	734	44.3	75179	814	37.1
100000	500623	2972	24.7	426665	8173	21.1
1000000	3173343	17856	15.7	1829065	40502	9.0
10000000	20342474	68912	10.0	8105010	26139	4.0

TABLE 1. Revenue loss of Algorithm 3 in Besbes and Zeevi [7] and NSC

In this numerical example, we set n = 2, m = 2, A = [1, 1; 0, 2], C = [0.1; 0.1]. The true demand function is a logit function, and $[\lambda_1(p_1, p_2); \lambda_2(p_1, p_2)] = (1 + \exp(0.4 - 0.015p_1) + \exp(0.8 - 0.02p_2))^{-1} [\exp(0.4 - 0.015p_1); \exp(0.8 - 0.02p_2)]$. For ease of performance comparison, we use s = 4 for both Algorithm 3 in Besbes and Zeevi [7] and NSC. We vary k from 500 (a capacity level of 50 for each resource) to 10000000 (a capacity level of 1000000 for each resource) and run 1000 trials for each k. The fourth and the seventh columns correspond to the relative revenue loss for the corresponding heuristic π defined as $\rho^{\pi}(k)/J^{D}(k)$. (The Matlab code of our numerical example can be found at https://sites.google.com/a/umich.edu/georgeqc/research.)

demand case as well. In this case, there is no random noise in demand observations, so one can simply set $L_0 = 1$ in the Spline Estimation subroutine. The other tuning parameter d needs to be adjusted accordingly. Specifically, given $L_0 = 1$, for any $s \ge 4$ and d, the estimation error of the demand function and its first order partial derivatives are in the order of $\epsilon := \mathcal{O}(d^{-(s \land \bar{s} - 1)})$ by a similar analysis as in Step 1 in the proof of Lemma 1. The expected revenue loss during the exploration stages is in the order of the number of prices being tested, i.e., $\mathcal{O}(d^n)$, while the expected revenue loss during the exploitation stage is $\mathcal{O}(\epsilon^2 k)$. Hence, the expected revenue loss throughout the selling season is $\mathcal{O}(\epsilon^{-\frac{n}{s \land \bar{s} - 1}} + \epsilon^2 k)$, which is minimized at $\epsilon = k^{-\frac{s \land \bar{s} - 1}{2(s \land \bar{s}) + n - 2}}$. Thus, by setting $d = k^{\frac{1}{2(s \land \bar{s}) + n - 2}}$, the performance bound of NSC for deterministic demand is $\mathcal{O}(k^{\frac{2(s \land \bar{s}) + n - 2}})$. This means that when the demand function is sufficiently smooth (i.e., $\bar{s} = \infty$), for any $\epsilon > 0$, we can choose s large enough so that the performance of NSC in the deterministic demand setting is $\mathcal{O}(k^{\epsilon})$. This highlights the fact that stochastic and deterministic demand cases have different complexities.

REMARK 6 (ON OUR DEMAND LINEARIZATION APPROACH). Although the estimated spline function is not used in the exploitation stage once the function approximations (i.e., linear demand approximation and quadratic revenue approximation) are conducted, we would like to re-iterate that Spline Estimation is crucial in NSC as it ensures that the demand function is linearized at a point that is sufficiently close to p^D so that the resulting function approximations are reasonably good. Other demand linearization approaches have been proposed in the literature as well. For example, in the single product without capacity constraint setting, by using the simple structure of the optimal price $\lambda(p^*) + p^*\lambda'(p^*) = 0$, Besbes and Zeevi [8] propose a simpler and more direct demand linearization approach that works well in their setting. This approach is unlikely to work in our multiple products with multiple capacity constraints setting because the optimal solution p^D does not permit the same simple structure anymore; instead, it is characterized by the KKT conditions (i.e., one needs to compare a combinatorial number of KKT points to find p^D). REMARK 7 (FINDING THE TUNING PARAMETERS FOR IMPLEMENTATION). To implement NSC in practice, a seller needs to find the tuning parameters L_0 and d which depend on k; but k is an asymptotic scaling factor and does not have a physical meaning in the original problem. How should a seller implement NSC? Note that L_0 and d only depend on k via kT which can be interpreted as the aggregated potential market size, i.e., the total number of customers who would potentially purchase a product the seller offers during the whole selling season. Estimating this aggregated number is not very difficult; a firm may be able to get a ballpark estimate from its own experience in the industry, or from market analysis research conducted by consulting companies. Then, the seller can divide the whole selling season into kT decision periods, compute L_0 and d accordingly, and apply NSC. Note that the estimate of kT does not need to be perfect; as long as this estimate is of the same order as the true aggregated market size, the corresponding tuning parameters will have the appropriate order of magnitude.

5. Proof of Theorem 1. In this section, we provide a complete proof of Theorem 1. We first discuss an outline of the proof, together with the key ideas and key lemmas, in §5.1 and then we fill in the remaining details in §5.2 - 5.4. Throughout this section, we fix $\pi = \text{NSC}$ and assume that T = 1 (this is without lost of generality).

5.1. Key ideas and outline of the proof. The proof of Theorem 1 uses a combination of large deviation arguments, stability analysis, and stopping time arguments. Below, we divide the proof into three parts.

$\underline{Part 1}$

In this part, we argue that, if k is large, $||\theta^* - \hat{\theta}||_2$ is small with a very high probability. This result allows us to use the perturbation result in Lemma 3 when analyzing the revenue loss later (in Part 3). Let $\epsilon(L_0) = (\log k/\sqrt{L_0})^{((s\wedge\bar{s})-2)/(s+n/2)}$ and define $\mathcal{E} := \{||\theta^* - \hat{\theta}||_2 \leq M_2 \epsilon(L_0)\}$, where M_2 is as defined in Lemma 4 below.

LEMMA 4. There exist constants $M_2, M_3 > 0$ independent of $k \ge 1$ and L_0 such that if $L_0 \ge \log^3 k$ and $s \ge 4$, then $\mathbb{P}^{\pi}_{\lambda^*}(||\theta^* - \hat{\theta}||_2 > M_2\epsilon(L_0)) \le M_3/k$.

The complete proof of Lemma 4 is given in §5.2. Here, we simply provide some basic intuition behind the proof. The proof uses Lemma 1 and Lemma 2. In particular, recall that Lemma 1 indicates that, with a very high probability, the spline function $\tilde{\lambda}(.)$ closely approximates the underlying demand function $\lambda^*(.)$ both in terms of the function value and its first and second order partial derivatives when $s \ge 4$ and $\bar{s} \ge 2$. This result together with the nonparametric perturbation result in Lemma 2 establishes that \tilde{p}^D is very close to p^D with very high probability. But then, the first and second order derivatives of $\tilde{\lambda}(.)$ evaluated at \tilde{p}^D also closely approximate those of $\lambda^*(.)$ evaluated at p^D . Note that by construction, θ^* (resp. $\hat{\theta}$) can essentially be viewed as a function of p^D (resp. \tilde{p}^D), $\lambda^*(p^D)$ (resp. $\tilde{\lambda}(\tilde{p}^D)$) and its first and second order derivatives evaluated at p^D (resp. \tilde{p}^D). Hence, $\hat{\theta}$ closely approximates θ^* with a very high probability.

In the remainder of this part, we first discuss some observations that follow from Lemma 4 and then use these observations to define a constant Ω_1 , which will be used in Parts 2 and 3. (As will be clear later, the problem behaves nicely and the prerequisite of Lemma 3 is satisfied when $k \geq \Omega_1$. This sets the stage for the analysis in Parts 2 and 3.) Four observations are in order. First, note that, as k becomes large, the probability of \mathcal{E} tends to one. Second, $\epsilon(L_0) \to 0$ as $k \to \infty$ under the condition that $L_0 > \log^3 k$, and this condition is satisfied for all sufficiently large k because our selection of L_0 implies

$$k^{\frac{s+n/2}{2s+n-2}} (\log k)^{\frac{2s+n-4}{2s+n-2}} \le L_0 \le 2k^{\frac{s+n/2}{2s+n-2}} (\log k)^{\frac{2s+n-4}{2s+n-2}}, \tag{4}$$

and $\log^3 k$ is smaller than the left hand side of (4) for large k. In light of Lemma 4, this observation means that if k is large, θ^* and $\hat{\theta}$ can be arbitrarily close with a very high probability. By (4), we have the following bounds for $\epsilon(L_0)$ as well:

$$2^{-\frac{s\wedge\bar{s}-2}{2s+n}} \left(\frac{\log k}{\sqrt{k}}\right)^{\frac{s\wedge\bar{s}-2}{2s+n-2}} \le \epsilon(L_0) \le \left(\frac{\log k}{\sqrt{k}}\right)^{\frac{s\wedge\bar{s}-2}{2s+n-2}}.$$
(5)

s+n

Third, by definition of L_0 and our choice of d in Theorem 1, we can bound

$$\tilde{L}_0 = s^n (s+d)^n L_0 \le s^n (s+1)^n d^n L_0 \le s^n (s+1)^n 2^{n+2} k^{\frac{s+n}{2s+n-2}} (\log k)^{\frac{2(s-2)}{2s+n-2}},$$
(6)

where the second inequality follows because, by (4) and the definition of d in Theorem 1,

$$d^{n}L_{0} \leq \left[2(\sqrt{L_{0}}/\log k)^{\frac{1}{s+n/2}}\right]^{n}L_{0} = 2^{n}(\log k)^{-\frac{n}{s+n/2}}L_{0}^{\frac{1}{s+n/2}}$$
$$\leq 2^{n}(\log k)^{-\frac{n}{s+n/2}}2^{\frac{s+n}{s+n/2}}k^{\frac{s+n}{2s+n-2}}(\log k)^{\frac{2s+n-4}{2s+n-2}}(\log k)^{\frac{2s+n-4}{2s+n-2}}$$
$$\leq 2^{n+2}k^{\frac{s+n}{2s+n-2}}(\log k)^{\frac{2s+n-4}{s+n/2}-\frac{n}{s+n/2}}=2^{n+2}k^{\frac{s+n}{2s+n-2}}(\log k)^{\frac{2(s-2)}{2s+n-2}}.$$

Note that (6) implies that $L_0/k \to 0$ as $k \to \infty$. So, there exists a constant $\Omega_0 > 0$ such that for all $k > \Omega_0$, we have $\tilde{L}_0 \le k/2$. Fourth, there exists a constant $M_4 > 0$ independent of $k \ge \Omega_0$ such that for all $k \ge \Omega_0$,

$$\begin{aligned} ||\delta||_{2} &= \left| \left| C - \frac{C_{\tilde{L}_{0}}}{k - \tilde{L}_{0}} \right| \right|_{2} = \left| \left| \frac{(kC - \tilde{L}_{0}C) - (kC - A\sum_{s=1}^{\tilde{L}_{0}}D_{s})}{k - \tilde{L}_{0}} \right| \right|_{2} = \left| \left| \frac{A\sum_{s=1}^{\tilde{L}_{0}}D_{s} - \tilde{L}_{0}C}{k - \tilde{L}_{0}} \right| \right|_{2} \\ &\leq 2\left(||A\mathbf{e}||_{2} + ||C||_{2} \right) \frac{\tilde{L}_{0}}{k} \leq 2\left(||A\mathbf{e}||_{2} + ||C||_{2} \right) s^{n}(s+1)^{n} 2^{n+2} (\log k/\sqrt{k})^{\frac{2(s-2)}{2s+n-2}} \\ &\leq 2\left(||A\mathbf{e}||_{2} + ||C||_{2} \right) s^{n}(s+1)^{n} 2^{n+2} (\log k/\sqrt{k})^{\frac{2(s/\tilde{s}-2)}{2s+n-2}} \leq M_{4} \epsilon(L_{0})^{2} \end{aligned}$$

where the first inequality follows because $L_0 \leq k/2$ for $k \geq \Omega_0$ and we have at most one arrival per period, the second inequality follows from (6), and the last inequality follows from (5).

Let δ be defined as in Lemma 3. Putting together the four observations above, we conclude that there exists a constant $\Omega_1 > 0$ independent of k such that for all $k \ge \Omega_1$ the following holds:

$$\mathbb{P}^{\pi}_{\lambda^*}(\mathcal{E}) \ge 1 - M_3/k \ge 1/2; \tag{7}$$

Conditioning on
$$\mathcal{E}, ||\theta^* - \hat{\theta}||_2 \le M_2 \epsilon(L_0) \le \bar{\delta};$$
 (8)

$$||\delta||_2 \le M_4 \epsilon (L_0)^2 \le \bar{\delta}. \tag{9}$$

Inequality (7) indicates that we only need to focus on the revenue loss on the event \mathcal{E} . Inequalities (8) and (9) are crucial; they ensure that, for $k \geq \Omega_1$, conditioning on \mathcal{E} , the prerequisite of Lemma 3 is satisfied and the perturbation bounds therein can be used to analyze the performance of NSC.

$\underline{Part 2}$

In this part, we define a stopping time τ and analyze its properties. This will be crucial for our analysis in Part 3. In particular, it helps us to quantify the amount of revenue loss under NSC during the exploitation phase. (Stopping time argument is also used in Jasin [21]. However, unlike the arguments in Jasin [21], which assume *known demand function*, here we also need to deal with estimation errors, approximation errors, and systematic biases.) We first define τ and state its properties in Lemmas 5 and 6. For clarity, we delay the complete proof of these two lemmas in §5.3 and only discuss the intuition here. Let τ be the minimum of k and the first time $t \geq \tilde{L}_0 + 1$ such that the following condition (†) is violated:

(†)
$$\psi > S(t)$$
, where $\psi := \sqrt{\epsilon(L_0)}$, $S(k) := \infty$ and $\forall t \in \overline{[1, k-1]}$,
 $S(t) := \left| \left| \sum_{s=\tilde{L}_0+1}^t \frac{\tilde{\Delta}_s}{k-s} \right| \right|_2 + \frac{1}{k-t}$.

(Recall that $\hat{\Delta}_s = \Delta_s + \lambda^*(p_s) - \lambda(p_s; \hat{\theta}_\iota)$.) The purpose of condition (†) is to guarantee that \hat{p}_t is not too far way from $p_{\delta}^D(\hat{\theta})$ (see the pricing formula in Stage 3 of NSC) before τ and the cumulative deviation of the actual demand realizations from the target average demand is not too large before τ . Let Ω_1 and \mathcal{E} be as defined in Part 1. We state two lemmas.

LEMMA 5. Suppose that $L_0 \ge \log^3 k$. There exists a constant $\Omega_2 > \Omega_1$ independent of $k \ge 1$ such that for all $k \ge \Omega_2$ and all sample paths on \mathcal{E} , $\hat{p}_t \in \mathcal{P}$ (i.e., $p_t = \hat{p}_t$) and $C_t \succeq A_j$ for all $j \in [1, n]$ and $t \in [\tilde{L}_0 + 1, \tau - 1]$.

LEMMA 6. There exists a constant $M_5 > 0$ independent of $k \ge 1$ such that, for all $k \ge \Omega_2$, we have $\mathbb{E}_{\lambda^*}^{\pi}[k-\tau|\mathcal{E}] \le M_5(\epsilon(L_0)^2k + \epsilon(L_0)^{-1}\log k + \epsilon(L_0)^{-2}).$

Lemma 5 essentially says that, when k is sufficiently large, everything behaves "nicely" before the stopping time τ on \mathcal{E} . As will be clear in Part 3, this enables us to explicitly characterize the cumulative revenue loss under NSC before τ . After τ , NSC may end up charging the turn-off prices (i.e., due to stock-out) and the characterization of p_t becomes less tractable. Fortunately, Lemma 6 tells us that $\mathbb{E}^{\pi}_{\lambda^*}[k-\tau]$ is small for large k (i.e., τ is large). So, by regularity condition R3, we can simply bound the per period revenue loss after τ with \bar{r} .

The complete proof of Lemma 6 is deferred to §5.3. For now, we provide the main intuition and highlight how our argument differs from that in Jasin [21]. Note that, since $\mathbb{E}_{\lambda^*}^{\pi}[k-\tau|\mathcal{E}] = \sum_{t=1}^{k-1} \mathbb{P}_{\lambda^*}^{\pi}(\tau \leq t|\mathcal{E})$, the proof of Lemma 6 boils down to computing a bound (for each t) for the conditional probability $\mathbb{P}_{\lambda^*}^{\pi}(\tau \leq t|\mathcal{E})$. Roughly speaking, this is equivalent to analyzing the probability that S(s) is smaller than the threshold ψ for $\tilde{L}_0 + 1 \leq s \leq t$. Note that S(t) can be bounded as follows:

$$S(t) \leq \underbrace{\left\| \sum_{s=\tilde{L}_{0}+1}^{t} \frac{\Delta_{s}}{k-s} \right\|_{2}}_{\text{random noise}} + \underbrace{\left\| \sum_{s=\tilde{L}_{0}+1}^{t} \frac{\lambda^{*}(p_{s}) - \lambda(p_{s};\hat{\theta}_{\iota})}{k-s} \right\|_{2}}_{\text{systematic biases}} + \frac{1}{k-t},$$

where the random noise comes from the stochasticity of demand and the systematic bias comes from the estimation error due to Spline Estimation and demand linearization. The systematic biases term does not appear in Jasin [21]; here, it is the primary driving force of the order of $\mathbb{E}_{\lambda^*}^{\pi}[k-\tau|\mathcal{E}]$. In the proof, we use Markov's inequality and integration inequality to bound that term. Note that in order to derive a tight bound using Markov's inequality, we need to make sure that the order of $||\lambda^*(p_s) - \lambda(p_s; \hat{\theta}_i)||_2$ is small enough. It turns out that, for all $s < \tau$, the definitions of τ and ψ ensure that p_s is very close to p^D ; moreover, since $\lambda(.;\theta_{\iota}^*)$ is a good approximation of $\lambda^*(.)$ in the neighborhood of p^D , $\lambda(p_s; \theta_i^*)$ is very close to $\lambda^*(p_s)$ as well. This observation together with Lemma 4 further implies that, conditioning on \mathcal{E} , for all $s < \tau$, the order of $||\lambda^*(p_s) - \lambda(p_s; \hat{\theta}_{\iota})||_2 = \mathcal{O}(\epsilon(L_0))$ (see derivation in (44) for more details) is sufficiently small. However, for $s \ge \tau$, p_s is not guaranteed to be sufficiently close to p^D and the order of $||\lambda^*(p_s) - \lambda(p_s; \hat{\theta}_{\iota})||_2$ could be as large as $\Theta(1)$, which will blow up the Markov's bound we derive. (Although the spline estimate $\lambda(.)$ is uniformly close to $\lambda^*(.)$, its linear approximation $\lambda(.; \hat{\theta}_{\iota})$ is not always close to $\lambda^*(.)$, except for prices that are sufficiently close to p^{D} , see (44).) This means that we cannot use Markov's inequality directly on τ as it is defined in (†). The culprit here is the term S(t) which, by definition of Δ_s , includes the summation of many systematic biases terms that may turn out to be very large. To address this, we introduce another stopping time $\tilde{\tau}$ as the minimum of k and the first time $t \geq L_0 + 1$ such that the following condition $(\dagger\dagger)$ is violated:

$$\begin{array}{l} (\dagger\dagger) \qquad \psi > \tilde{S}(t), \quad \text{where } \tilde{S}(k) := \infty \text{ and } \forall t \in \overline{[1, k-1]}, \\ \tilde{S}(t) := \left| \left| \sum_{s=\tilde{L}_0+1}^t \frac{\Delta_s}{k-s} + \sum_{s=\tilde{L}_0+1}^t \frac{\lambda^*(p_s) - \lambda(p_s;\hat{\theta}_t)}{k-s} \mathbf{1}_{\{s \le \tau\}} \right| \right|_2 + \frac{1}{k-t}. \end{aligned}$$

We prove in Lemma 6 (see §5.3) that $\tilde{\tau}$ actually equals τ on every sample path, but $\tilde{\tau}$ is easier to work with because the term $\tilde{S}(t)$ in the stopping criterion only involves one systematic bias term that may be large (i.e., $(\lambda^*(p_{\tau}) - \lambda(p_{\tau}; \hat{\theta}_{\iota}))/(k - \tau)$). The desired result can then be attained by Markov's inequality and integration inequality.

$\underline{Part 3}$

Finally, we analyze the revenue loss of NSC as a function of k. Here, we collect the results from Parts 1 and 2 and use standard arguments to "count" the revenue loss incurred throughout the selling season (see, for example, Jasin [21]). If $k = \mathcal{O}(1)$, the revenue loss can be bounded by a constant; if k is large, all the useful properties of τ and \mathcal{E} derived above (Lemmas 4 - 6) hold and we can use them to analyze the revenue loss of NSC. We break down the revenue loss of NSC into three parts: (i) revenue loss incurred during the exploration stage, (ii) revenue loss incurred during the exploitation stage before τ , and (iii) revenue loss incurred during the exploitation stage after τ . Since the length of the exploration stage is \tilde{L}_0 , by R3, we can bound (i) with $\tilde{L}_0\bar{r}$. As for (ii) and (iii), we derive an upper bound by conditioning on \mathcal{E} and \mathcal{E}^c . Since $\mathbb{P}^{\pi}_{\lambda^*}(\mathcal{E}^c)$ is very small for large k, the majority of the revenue loss comes from the expected revenue loss conditioning on \mathcal{E} . This means that, roughly speaking, (iii) can be bounded by $\bar{r} \mathbb{E}^{\pi}_{\lambda^*}[k - \tau | \mathcal{E}]$. The remaining work is to carefully bound (ii) conditioning on \mathcal{E} using Taylor's expansion.

Let $\Omega := \max\{\Omega_1, \Omega_2, \Omega_3\}$, where Ω_1 is as defined in Part 1, Ω_2 is as defined in Lemma 5, and Ω_3 is a constant independent of k such that $\epsilon(L_0) < 1$ for all $k \ge \Omega_3$. If $k < \Omega$, $\rho^{\pi}(k) < \bar{r}\Omega = \mathcal{O}(1)$. So, we can focus on the case $k \ge \Omega$. Let R_t^{π} denote the revenue earned in period t under policy π , and let $\hat{R}_{\lambda^*}^{\pi}(k) := \sum_{t=\bar{L}_0+1}^k R_t^{\pi}$ denote the revenue earned during the exploitation stage. For notational brevity, we will simply write $\lambda_t = \lambda^*(p_t)$. Let $\bar{\Delta}_t := R_t^{\pi} - r^*(\lambda_t)$. Note that $\{\bar{\Delta}_t\}_{t=\bar{L}_0+1}^{k-1}$ is a martingale difference sequence with respect to filtration $\{\mathcal{H}_t\}_{t=\bar{L}_0+1}^{k-1}$. Thus, by R3 and Optional Stopping Time Theorem, we have $-\mathbb{E}_{\lambda^*}^{\pi}[\sum_{t=\bar{L}_0+1}^{\tau-1}\bar{\Delta}_t] = -\mathbb{E}_{\lambda^*}^{\pi}[\sum_{t=\bar{L}_0+1}^{\tau}\bar{\Delta}_t] + \mathbb{E}_{\lambda^*}^{\pi}[\bar{\Delta}_{\tau}] = \mathbb{E}_{\lambda^*}^{\pi}[\bar{\Delta}_{\tau}] \le \bar{r}$. Therefore, for $k \ge \Omega$, $\rho^{\pi}(k)$ can be bounded as follows:

$$\rho^{\pi}(k) = \mathbb{E}_{\lambda^{*}}^{\pi} \left[\sum_{t=1}^{\tilde{L}_{0}} \left(r^{D} - R_{t}^{\pi} \right) + \sum_{t=\tilde{L}_{0}+1}^{\tau-1} \left(r^{D} - r_{\lambda}^{*}(\lambda_{t}) - \bar{\Delta}_{t} \right) + \sum_{t=\tau}^{k} \left(r^{D} - R_{t}^{\pi} \right) \right] \\
\leq \bar{r}\tilde{L}_{0} + \bar{r} + \mathbb{E}_{\lambda^{*}}^{\pi} \left[\sum_{t=\tilde{L}_{0}+1}^{\tau-1} \left(r^{D} - r_{\lambda}^{*}(\lambda_{t}) \right) + \sum_{t=\tau}^{k} \left(r^{D} - R_{t}^{\pi} \right) \right] \\
= \bar{r}(1 + \tilde{L}_{0}) + \mathbb{E}_{\lambda^{*}}^{\pi} \left[\sum_{t=\tilde{L}_{0}+1}^{\tau-1} \left(r^{D} - r_{\lambda}^{*}(\lambda_{t}) \right) + \sum_{t=\tau}^{k} \left(r^{D} - R_{t}^{\pi} \right) \right] \mathcal{E}^{c} \right] \mathbb{P}_{\lambda^{*}}^{\pi}(\mathcal{E}^{c}) \\
+ \mathbb{E}_{\lambda^{*}}^{\pi} \left[\sum_{t=\tilde{L}_{0}+1}^{\tau-1} \left(r^{D} - r_{\lambda}^{*}(\lambda_{t}) \right) + \sum_{t=\tau}^{k} \left(r^{D} - R_{t}^{\pi} \right) \right] \mathcal{E}^{d} \right] \mathbb{P}_{\lambda^{*}}^{\pi}(\mathcal{E}) \\
\leq \bar{r}(1 + \tilde{L}_{0}) + \bar{r}k\mathbb{P}_{\lambda^{*}}^{\pi}(\mathcal{E}^{c}) + \mathbb{E}_{\lambda^{*}}^{\pi} \left[\sum_{t=\tilde{L}_{0}+1}^{\tau-1} \left(r^{D} - r_{\lambda}^{*}(\lambda_{t}) \right) + \sum_{t=\tau}^{k} \left(r^{D} - R_{t}^{\pi} \right) \right] \mathcal{E}^{d} \\
\leq \bar{r}(1 + \tilde{L}_{0}) + \bar{r}k\mathbb{P}_{\lambda^{*}}^{\pi}(\mathcal{E}^{c}) + \mathbb{E}_{\lambda^{*}}^{\pi} \left[\sum_{t=\tilde{L}_{0}+1}^{\tau-1} \left(r^{D} - r_{\lambda}^{*}(\lambda_{t}) \right) \right] \mathcal{E}^{d} + \bar{r}\mathbb{E}_{\lambda^{*}}^{\pi}[k - \tau + 1 |\mathcal{E}] \\
\leq \bar{r}\left(2 + M_{3} + \tilde{L}_{0} + \mathbb{E}_{\lambda^{*}}^{\pi}[k - \tau |\mathcal{E}] \right) + \mathbb{E}_{\lambda^{*}}^{\pi} \left[\sum_{t=\tilde{L}_{0}+1}^{\tau-1} \left(r^{D} - r_{\lambda}^{*}(\lambda_{t}) \right) \right] \mathcal{E}^{d} \right], \tag{10}$$

where the last inequality follows because $k\mathbb{P}^{\pi}_{\lambda^*}(\mathcal{E}^c) \leq M_3$ by Lemma 4. To bound the second term in (10), we use Taylor's expansion. Note that, by R3,

$$r^{D} - r_{\lambda}^{*}(\lambda_{t}) = r_{\lambda}^{*}(\lambda^{D}) - r_{\lambda}^{*}(\lambda_{t}) \leq \nabla r_{\lambda}^{*}(\lambda^{D}) \cdot (\lambda^{D} - \lambda_{t}) + \frac{\overline{v}}{2} ||\lambda^{D} - \lambda_{t}||_{2}^{2}$$

We will show in §5.4 that there exist constants $M_6, M_7 > 0$ independent of $k \ge \Omega$ such that

$$\mathbb{E}_{\lambda^{*}}^{\pi} \left[\sum_{\substack{t=\tilde{L}_{0}+1\\ 2}}^{\tau-1} \nabla r_{\lambda}^{*}(\lambda^{D}) \cdot (\lambda^{D} - \lambda_{t}) \middle| \mathcal{E} \right] \leq M_{6} \left(1 + \epsilon (L_{0})^{2} k + \epsilon (L_{0})^{-1} \log k + \epsilon (L_{0})^{-2} \right); \quad (11)$$

$$\frac{\bar{v}}{2} \mathbb{E}_{\lambda^{*}}^{\pi} \left[\sum_{\substack{t=\tilde{L}_{0}+1\\ t=\tilde{L}_{0}+1}}^{\tau-1} ||\lambda^{D} - \lambda_{t}||_{2}^{2} \middle| \mathcal{E} \right] \leq M_{7} \left(\log k + \epsilon (L_{0})^{2} k \right). \quad (12)$$

Combining (5)-(6) and (10)-(12) with Lemma 6, we conclude that there exist constants $M_8, M_9 > 0$ independent of $k > \Omega$ such that for all $k > \Omega$, we have:

$$\begin{split} \rho^{\pi}(k) &\leq M_8 \left(\epsilon (L_0)^2 k + \epsilon (L_0)^{-1} \log k + \epsilon (L_0)^{-2} + \bar{r} \tilde{L}_0 \right) \\ &\leq M_8 \left[k \left(\frac{\log^2 k}{k} \right)^{\frac{s \wedge \bar{s} - 2}{2s + n - 2}} + 2^{\frac{s \wedge \bar{s} - 2}{2s + n}} \log k \left(\frac{\sqrt{k}}{\log k} \right)^{\frac{s \wedge \bar{s} - 2}{2s + n - 2}} + 2^{\frac{2(s \wedge \bar{s} - 2)}{2s + n}} \left(\frac{k}{\log^2 k} \right)^{\frac{s \wedge \bar{s} - 2}{2s + n - 2}} \right] \\ &+ M_8 \left[\bar{r} s^n (s + 1)^n 2^{n+2} k^{\frac{s + n}{2s + n - 2}} (\log k)^{\frac{2(s - 2)}{2s + n - 2}} \right] \leq M_9 k^{\frac{2s - s \wedge \bar{s} + n}{2s + n - 2}} \log k. \end{split}$$

Letting $M_1 = M_9 + \bar{r}\Omega$ completes the proof of Theorem 1.

5.2. Part 1: Proof of Lemma 4. Define $\mathcal{F} := \mathcal{F}_1 \cap \mathcal{F}_2$, where

$$\begin{aligned} \mathcal{F}_{1} &:= \left\{ ||p^{D} - \tilde{p}^{D}||_{2} \leq C_{0}\epsilon(L_{0}) \right\}, \\ \mathcal{F}_{2} &:= \left\{ \left| \left| \frac{\partial^{r}(\lambda_{j}^{*}(.) - \tilde{\lambda}_{j}(.))}{\partial p_{1}^{-1} \dots \partial p_{n}^{-n}} \right| \right|_{\infty} < C_{1}\epsilon(L_{0}), \forall j \in \overline{[1,n]}, r \in \overline{[0,2]}, r_{l} \in \mathbb{Z}_{+}, l \in \overline{[1,n]}, \sum_{l=1}^{n} r_{l} = r \right\}, \end{aligned}$$

 C_0 is a positive constant to be chosen later and $C_1 := \max\{\Psi_0, \Psi_1, \Psi_2\}$ (recall that Ψ_0, Ψ_1, Ψ_2 are the constants discussed in Lemma 1). Let $\Phi := \max\{\Phi_1, \Phi_2\}$, where $\Phi_1 > 3, \Phi_2 > 3$ are constants to be chosen later. We first derive an upper bound for $||\theta^* - \hat{\theta}||_2$ conditioning on \mathcal{F} for $k \ge \Phi$. (Unless otherwise noted, in what follows, we will simply assume that \mathcal{F} is satisfied and $k \ge \Phi$.)

By R1 (i.e., Lipschitz continuity of the second order partial derivatives of $\lambda^*(.)$), the compactness of \mathcal{P} , and the continuity of $\tilde{\lambda}(.)$ (note that $s \ge 4 > 2$ implies $\tilde{\lambda}(.) \in \mathsf{C}(\mathcal{P})$), there exists a constant $C_2 > 0$ such that, conditioning on \mathcal{F} , for all $r \in [0,2]$, $r_l \in \mathbb{Z}_+, l \in [1,n]$ satisfying $\sum_{l=1}^n r_l = r, p \in \mathcal{P}$, and $j \in [1,n]$, we have:

$$\frac{\partial^r (\lambda_j^*(p^D) - \lambda_j^*(\tilde{p}^D))}{\partial p_1^{r_1} \dots \partial p_n^{r_n}} \bigg| \le C_2 ||p^D - \tilde{p}^D||_2 \le C_2 C_0 \epsilon(L_0) \quad \text{and} \quad |\tilde{\lambda}_j(p)| \le C_2.$$
(13)

So, the following two inequalities hold:

$$\begin{aligned} ||\lambda^{*}(p^{D}) - \lambda^{*}(\tilde{p}^{D})||_{2} &\leq \sqrt{n} \, ||\lambda^{*}(p^{D}) - \lambda^{*}(\tilde{p}^{D})||_{\infty} \\ &= \sqrt{n} \max_{j=1,\dots,n} |\lambda_{j}^{*}(p^{D}) - \lambda_{j}^{*}(\tilde{p}^{D})| \leq \sqrt{n} C_{2} C_{0} \epsilon(L_{0}), \end{aligned}$$
(14)

where the last inequality follows by applying r = 0 to (13), and

$$||u_{ij}^{*} - \hat{u}_{ij}||_{2} = \sqrt{\sum_{l=1}^{n} \left| \frac{\partial^{2} \lambda_{l}^{*}(p^{D})}{\partial p_{i} \partial p_{j}} - \frac{\partial^{2} \lambda_{l}^{*}(\tilde{p}^{D})}{\partial p_{i} \partial p_{j}} + \frac{\partial^{2} \lambda_{l}^{*}(\tilde{p}^{D})}{\partial p_{i} \partial p_{j}} - \frac{\partial^{2} \tilde{\lambda}_{l}(\tilde{p}^{D})}{\partial p_{i} \partial p_{j}} \right|^{2}} \leq \sqrt{\sum_{l=1}^{n} 2 \left| \frac{\partial^{2} \lambda_{l}^{*}(p^{D})}{\partial p_{i} \partial p_{j}} - \frac{\partial^{2} \lambda_{l}^{*}(\tilde{p}^{D})}{\partial p_{i} \partial p_{j}} \right|^{2}} + \sum_{l=1}^{n} 2 \left| \left| \frac{\partial^{2} \lambda_{l}^{*}(.)}{\partial p_{i} \partial p_{j}} - \frac{\partial^{2} \tilde{\lambda}_{l}(.)}{\partial p_{i} \partial p_{j}} \right|^{2}} \right|_{\infty} \leq \sqrt{2nC_{0}^{2}C_{2}^{2}\epsilon(L_{0})^{2} + 2nC_{1}^{2}\epsilon(L_{0})^{2}} = \epsilon(L_{0})\sqrt{2nC_{0}^{2}C_{2}^{2} + 2nC_{1}^{2}}.$$
(15)

By similar arguments as above, there exists a constant $C_3 > 0$ independent of k such that:

$$||B^* - \hat{B}||_F = \sqrt{\sum_{i=1}^n \sum_{j=1}^n \left| \frac{\partial \lambda_i^*(p^D)}{\partial p_j} - \frac{\partial \lambda_i^*(\tilde{p}^D)}{\partial p_j} + \frac{\partial \lambda_i^*(\tilde{p}^D)}{\partial p_j} - \frac{\partial \tilde{\lambda}_i(\tilde{p}^D)}{\partial p_j} \right|^2} \\ \leq \sqrt{\sum_{i=1}^n \sum_{j=1}^n 2 \left| \frac{\partial \lambda_i^*(p^D)}{\partial p_j} - \frac{\partial \lambda_i^*(\tilde{p}^D)}{\partial p_j} \right|^2} + \sum_{i=1}^n \sum_{j=1}^n 2 \left| \left| \frac{\partial \lambda_i^*(.)}{\partial p_j} - \frac{\partial \tilde{\lambda}_i(.)}{\partial p_j} \right| \right|_{\infty}^2} \\ \leq \sqrt{2n^2 C_2^2 C_0^2 \epsilon (L_0)^2 + 2n^2 C_1^2 \epsilon (L_0)^2} \leq C_3 \epsilon (L_0).$$
(16)

We now derive a bound for $||H^* - \hat{H}||_2$. To do this, we need to first find a bound for $||\hat{B}^{-1}||_2$. Let $\sigma_{\max}(X)$ and $\sigma_{\min}(X)$ denote the maximum and the minimum eigenvalues of a symmetric real matrix X, respectively. Since $B^* = \nabla \lambda^*(p^D)$ is invertible, $B^*(B^*)'$ is positive definite; so, $\bar{\sigma}^* := \sigma_{\max}(B^*(B^*)') > 0$ and $\underline{\sigma}^* := \sigma_{\min}(B^*(B^*)') > 0$. Moreover, since $C_3\epsilon(L_0) \to 0$ as $k \to \infty$, by (16), there exists $\Phi_1 > 0$ such that, for all $k > \Phi_1$, $||B^* - \hat{B}||_2 \le ||B^* - \hat{B}||_F \le C_3\epsilon(L_0) \le \underline{\sigma}^*/(4\sqrt{\bar{\sigma}^*})$. Therefore, for all $v \in \mathbb{R}^n$ with $||v||_2 = 1$,

$$\begin{aligned} v'\hat{B}'\hat{B}v &= v'(\hat{B} - B^* + B^*)'(\hat{B} - B^* + B^*)v \\ &= v'(B^*)'B^*v + v'(B^*)'(\hat{B} - B^*)v + v'(\hat{B} - B^*)'B^*v + v'(\hat{B} - B^*)'(\hat{B} - B^*)v \\ &\geq \underline{\sigma}^* - 2||v||_2^2||B^*||_2||\hat{B} - B^*||_2 \geq \underline{\sigma}^* - 2\sqrt{\overline{\sigma}^*}\underline{\sigma}^*/(4\sqrt{\overline{\sigma}^*}) = \underline{\sigma}^*/2. \end{aligned}$$

This means that $\sigma_{\min}(\hat{B}'\hat{B}) \ge \underline{\sigma}^*/2 > 0$. Since $(\hat{B}'\hat{B})^{-1} = \hat{B}^{-1}(\hat{B}^{-1})'$,

$$||\hat{B}^{-1}||_2 = \sqrt{\sigma_{\max}(\hat{B}^{-1}(\hat{B}^{-1})')} = \sqrt{\sigma_{\min}(\hat{B}'\hat{B})^{-1}} \le \sqrt{2/\underline{\sigma}^*}.$$
(17)

By telescoping, we can bound

$$\begin{aligned} |H_{ij}^{*} - \hat{H}_{ij}| &= |(u_{ij}^{*})'(B^{*})^{-1}\lambda^{*}(p^{D}) - \hat{u}_{ij}'\hat{B}^{-1}\tilde{\lambda}(\tilde{p}^{D})| \\ &\leq |(u_{ij}^{*})'(B^{*})^{-1}\lambda^{*}(p^{D}) - (u_{ij}^{*})'(B^{*})^{-1}\tilde{\lambda}(\tilde{p}^{D})| \\ &+ |(u_{ij}^{*})'(B^{*})^{-1}\tilde{\lambda}(\tilde{p}^{D}) - (u_{ij}^{*})'\hat{B}^{-1}\tilde{\lambda}(\tilde{p}^{D})| + |(u_{ij}^{*})'\hat{B}^{-1}\tilde{\lambda}(\tilde{p}^{D}) - \hat{u}_{ij}'\hat{B}^{-1}\tilde{\lambda}(\tilde{p}^{D})| \\ &\leq ||u_{ij}^{*}||_{2}||(B^{*})^{-1}||_{2}||\lambda^{*}(p^{D}) - \tilde{\lambda}(\tilde{p}^{D})||_{2} \\ &+ ||u_{ij}^{*}||_{2}||(B^{*})^{-1}||_{2}||B^{*} - \hat{B}||_{2}||\hat{B}^{-1}||_{2}||\tilde{\lambda}(\tilde{p}^{D})||_{2} \\ &+ ||u_{ij}^{*} - \hat{u}_{ij}||_{2}||\hat{B}^{-1}||_{2}||\tilde{\lambda}(\tilde{p}^{D})||_{2} \end{aligned}$$
(18)

where the last inequality follows because $(B^*)^{-1} - \hat{B}^{-1} = (B^*)^{-1}(\hat{B} - B^*)\hat{B}^{-1}$. We now bound the three terms on the right hand side of (18) one by one. For the first term of (18), by (14) and the definition of \mathcal{F}_2 , we have:

$$\begin{aligned} ||u_{ij}^*||_2 ||(B^*)^{-1}||_2 ||\lambda^*(p^D) - \tilde{\lambda}(\tilde{p}^D)||_2 &\leq ||u_{ij}^*||_2 ||(B^*)^{-1}||_2 (||\lambda^*(p^D) - \lambda^*(\tilde{p}^D)||_2 + ||\lambda^*(\tilde{p}^D) - \tilde{\lambda}(\tilde{p}^D)||_2) \\ &\leq ||u_{ij}^*||_2 ||(B^*)^{-1}||_2 (\sqrt{n}C_2C_0\epsilon(L_0) + \sqrt{n}\max_{\substack{j=1,\dots,n\\ j=1,\dots,n}} \{||\lambda_j^*(.) - \tilde{\lambda}_j(.)||_\infty\}) \\ &\leq ||u_{ij}^*||_2 ||(B^*)^{-1}||_2 (\sqrt{n}C_2C_0\epsilon(L_0) + \sqrt{n}C_1\epsilon(L_0)) = \mathcal{O}(\epsilon(L_0)). \end{aligned}$$

For the second term of (18), by (13), (16) and (17),

$$||u_{ij}^*||_2||(B^*)^{-1}||_2||B^* - \hat{B}||_2||\hat{B}^{-1}||_2||\tilde{\lambda}(\tilde{p}^D)||_2 \le ||u_{ij}^*||_2||(B^*)^{-1}||_2C_3\epsilon(L_0)\sqrt{2/\underline{\sigma}^*}\sqrt{n}C_2 = \mathcal{O}(\epsilon(L_0)).$$

For the last term of (18), by (13), (15) and (17)

For the last term of (18), by (13), (15) and (17),

$$||u_{ij}^* - \hat{u}_{ij}||_2 ||\hat{B}^{-1}||_2 ||\tilde{\lambda}(\tilde{p}^D)||_2 \le \epsilon(L_0)\sqrt{2nC_0^2C_2^2 + 2nC_1^2}\sqrt{2/\underline{\sigma}^*}\sqrt{nC_2} = \mathcal{O}(\epsilon(L_0)).$$

So, $|H_{ij}^* - \hat{H}_{ij}| = \mathcal{O}(\epsilon(L_0))$ and there exists a constant $C_4 > 0$ independent of $k \ge 1$ such that

$$||H^* - \hat{H}||_2 \le ||H^* - \hat{H}||_F = \sqrt{\sum_{i=1}^n \sum_{j=1}^n |H^*_{ij} - \hat{H}_{ij}|^2} \le C_4 \epsilon(L_0).$$
(19)

Using similar arguments as above, by telescoping and (14) - (19), there exist constants $C_5, C_6, C_7, C_8 > 0$ independent of $k \ge 1$ such that on \mathcal{F} we have

$$\begin{aligned} ||a^* - \hat{a}||_2 &\leq ||\lambda^*(p^D) - \tilde{\lambda}(\tilde{p}^D)||_2 + ||(B^*)'p^D - \hat{B}'\tilde{p}^D||_2 \\ &\leq ||\lambda^*(p^D) - \lambda^*(\tilde{p}^D)||_2 + ||\lambda^*(\tilde{p}^D) - \tilde{\lambda}(\tilde{p}^D)||_2 + ||B^*||_2||p^D - \tilde{p}^D||_2 + ||\tilde{p}^D||_2||B^* - \hat{B}||_2 \\ &\leq \sqrt{n}C_2C_0\epsilon(L_0) + \sqrt{n}C_1\epsilon(L_0) + ||B^*||_2C_0\epsilon(L_0) + (\sum_{l=1}^n \bar{p}_l^2)^{1/2}C_3\epsilon(L_0) = C_5\epsilon(L_0), (20) \end{aligned}$$

$$\begin{aligned} ||E^* - \hat{E}||_2 &\leq ||\frac{1}{2}(p^D - \tilde{p}^D)'H^*p^D||_2 + ||\frac{1}{2}(\tilde{p}^D)'H^*(p^D - \tilde{p}^D)||_2 + ||\frac{1}{2}(\tilde{p}^D)'(H^* - \hat{H})\tilde{p}^D||_2 \\ &\leq (\sum_{l=1}^n \bar{p}_l^2)^{1/2} ||H^*||_2 C_0 \epsilon(L_0) + \frac{1}{2} (\sum_{l=1}^n \bar{p}_l^2) C_4 \epsilon(L_0) = C_6 \epsilon(L_0), \end{aligned}$$
(21)

$$\begin{aligned} ||F^* - \hat{F}||_2 &\leq ||a^* - \hat{a}||_2 + ||H^*||_2 ||p^D - \tilde{p}^D||_2 + ||H^* - \hat{H}||_2 ||\tilde{p}^D||_2 \\ &\leq C_5 \epsilon(L_0) + ||H^*||_2 C_0 \epsilon(L_0) + (\sum_{l=1}^n \bar{p}_l^2)^{1/2} C_4 \epsilon(L_0) = C_7 \epsilon(L_0), \end{aligned}$$
(22)

$$||G^* - \hat{G}||_F \le 2||B^* - \hat{B}||_F + ||H^* - \hat{H}||_F \le 2C_3\epsilon(L_0) + C_4\epsilon(L_0) = C_8\epsilon(L_0).$$
(23)

Putting (16) and (20) - (23) together, for all $k \ge \Phi$, we obtain:

$$||\theta^* - \hat{\theta}||_2 \le ||a^* - \hat{a}||_2 + ||B^* - \hat{B}||_F + ||E^* - \hat{E}||_2 + ||F^* - \hat{F}||_2 + ||G^* - \hat{G}||_F \le C_9 \epsilon(L_0)$$
(24)

where $C_9 := C_3 + C_5 + C_6 + C_7 + C_8$. Let $M_2 = C_9 + 1$. Since \mathcal{F} implies \mathcal{E} (i.e., because $||\theta^* - \hat{\theta}||_2 < M_2 \epsilon(L_0)$ on \mathcal{F}), we can bound:

$$\mathbb{P}^{\pi}_{\lambda^*}(\mathcal{E}^c) \leq \mathbb{P}^{\pi}_{\lambda^*}(\mathcal{F}^c) \leq \mathbb{P}^{\pi}_{\lambda^*}(\mathcal{F}^c_1) + \mathbb{P}^{\pi}_{\lambda^*}(\mathcal{F}^c_2).$$
(25)

We will now bound each $\mathbb{P}^{\pi}_{\lambda^*}(\mathcal{F}_1^c)$ and $\mathbb{P}^{\pi}_{\lambda^*}(\mathcal{F}_2^c)$ separately. Note that

$$\begin{aligned}
\mathbb{P}_{\lambda^{*}}^{\pi}(\mathcal{F}_{2}^{c}) &\leq \sum_{j=1}^{n} \sum_{r=0}^{2} \left\| \sum_{r_{l} \geq 0, \sum_{l=1}^{n} r_{l} = r} \mathbb{P}_{\lambda^{*}}^{\pi} \left(\left\| \left\| \frac{\partial^{r}(\lambda_{j}^{*}(.) - \tilde{\lambda}_{j}(.))}{\partial p_{1}^{r_{1}} \dots \partial p_{n}^{r_{n}}} \right\|_{\infty} \geq C_{1} \epsilon(L_{0}) \right) \\
&\leq \sum_{j=1}^{n} \sum_{r=0}^{2} \left\| \sum_{r_{l} \geq 0, \sum_{l=1}^{n} r_{l} = r} \mathbb{P}_{\lambda^{*}}^{\pi} \left(\left\| \left\| \frac{\partial^{r}(\lambda_{j}^{*}(.) - \tilde{\lambda}_{j}(.))}{\partial p_{1}^{r_{1}} \dots \partial p_{n}^{r_{n}}} \right\|_{\infty} \geq \Psi_{r} \left(\frac{\log k}{\sqrt{L_{0}}} \right)^{\frac{s \wedge \bar{s} - r}{s + n/2}} \right) \\
&\leq \frac{n(n+1)(n+2)}{2} \frac{K}{k},
\end{aligned} \tag{26}$$

where the first inequality follows by the definition of \mathcal{F}_2 and the union bound, the second inequality follows by the definition of C_1 and $\epsilon(L_0)$, and the last inequality follows by Lemma 1. To compute a bound for $\mathbb{P}^{\pi}_{\lambda^*}(\mathcal{F}_1^c)$, first note that

$$\mathbb{P}_{\lambda^*}^{\pi}(\mathcal{F}_1^c) = \mathbb{P}_{\lambda^*}^{\pi}(\mathcal{F}_1^c | \mathcal{F}_2) \mathbb{P}_{\lambda^*}^{\pi}(\mathcal{F}_2) + \mathbb{P}_{\lambda^*}^{\pi}(\mathcal{F}_1^c | \mathcal{F}_2^c) \mathbb{P}_{\lambda^*}^{\pi}(\mathcal{F}_2^c) \le \mathbb{P}_{\lambda^*}^{\pi}(\mathcal{F}_1^c \cap \mathcal{F}_2) + \mathbb{P}_{\lambda^*}^{\pi}(\mathcal{F}_2^c).$$
(27)

So, it suffices that we find a bound for $\mathbb{P}_{\lambda^*}^{\pi}(\mathcal{F}_1^c \cap \mathcal{F}_2)$. By Lemma 2, there exist constants $\overline{\delta}, C_{10} > 0$ independent of $k \geq 1$ such that if $||\lambda^*(.) - \widetilde{\lambda}(.)||_{\infty} \leq \overline{\delta}$, then

$$\begin{split} ||p^{D} - \tilde{p}^{D}||_{2} &\leq C_{10}(||\lambda^{*}(.) - \tilde{\lambda}(.)||_{\infty} + ||(\nabla\lambda^{*}(.) - \nabla\tilde{\lambda}(.))'||_{\infty}) \\ &= C_{10} \sup_{p \in \mathcal{P}} ||\lambda^{*}(p) - \tilde{\lambda}(p)||_{\infty} + C_{10} \sup_{p \in \mathcal{P}} ||(\nabla\lambda^{*}(p) - \nabla\tilde{\lambda}(p))'||_{\infty} \\ &= C_{10} \sup_{p \in \mathcal{P}} \max_{j=1,...,n} |\lambda^{*}_{j}(p) - \tilde{\lambda}_{j}(p)| + C_{10} \sup_{p \in \mathcal{P}} \max_{i=1,...,n} \sum_{j=1}^{n} \left| \frac{\partial\lambda^{*}_{i}(p)}{\partial p_{j}} - \frac{\partial\tilde{\lambda}_{i}(p)}{\partial p_{j}} \right| \\ &\leq C_{10} \max_{j=1,...,n} ||\lambda^{*}_{j}(.) - \tilde{\lambda}_{j}(.)||_{\infty} + C_{10} \max_{i=1,...,n} \sum_{j=1}^{n} \left| \frac{\partial\lambda^{*}_{i}(.)}{\partial p_{j}} - \frac{\partial\tilde{\lambda}_{i}(.)}{\partial p_{j}} \right| \right|_{\infty} \\ &\leq C_{10} \max_{j=1,...,n} ||\lambda^{*}_{j}(.) - \tilde{\lambda}_{j}(.)||_{\infty} + C_{10} \sum_{j=1}^{n} \max_{i=1,...,n} \left| \frac{\partial\lambda^{*}_{i}(.)}{\partial p_{j}} - \frac{\partial\tilde{\lambda}_{i}(.)}{\partial p_{j}} \right| \right|_{\infty} . \tag{28}$$

Since $C_1\epsilon(L_0) \to 0$ as $k \to \infty$, there exists a constant $\Phi_2 > 0$ such that, conditioning on \mathcal{F}_2 , for all $k \ge \Phi \ge \Phi_2$, $||\lambda^*(.) - \tilde{\lambda}(.)||_{\infty} = \max_{j=1,...,n} ||\lambda^*_j(.) - \tilde{\lambda}_j(.)||_{\infty} \le C_1\epsilon(L_0) \le \bar{\delta}$; so, (28) holds. Let $C_0 = C_{10}C_1(1+n)$. Then, for all $k \ge \Phi$,

$$\begin{split} & \mathbb{P}_{\lambda^{*}}^{\pi} \left(\mathcal{F}_{1}^{c} \cap \mathcal{F}_{2} \right) = \mathbb{P}_{\lambda^{*}}^{\pi} \left(\{ || p^{D} - \tilde{p}^{D} ||_{2} > C_{0} \epsilon(L_{0}) \} \cap \mathcal{F}_{2} \right) \\ & \leq \mathbb{P}_{\lambda^{*}}^{\pi} \left(C_{10} \max_{j=1,\dots,n} \left\{ || \lambda_{j}^{*}(.) - \tilde{\lambda}_{j}(.) ||_{\infty} \right\} + C_{10} \sum_{j=1}^{n} \max_{i=1,\dots,n} \left\{ \left| \left| \frac{\partial \lambda_{i}^{*}(.)}{\partial p_{j}} - \frac{\partial \tilde{\lambda}_{i}(.)}{\partial p_{j}} \right| \right|_{\infty} \right\} > C_{10} C_{1} (1+n) \epsilon(L_{0}) \right) \\ & \leq \mathbb{P}_{\lambda^{*}}^{\pi} \left(\max_{j=1,\dots,n} \left\{ || \lambda_{j}^{*}(.) - \tilde{\lambda}_{j}(.) ||_{\infty} \right\} > C_{1} \epsilon(L_{0}) \right) + \sum_{j=1}^{n} \mathbb{P}_{\lambda^{*}}^{\pi} \left(\max_{i=1,\dots,n} \left\{ \left| \left| \frac{\partial \lambda_{i}^{*}(.)}{\partial p_{j}} - \frac{\partial \tilde{\lambda}_{i}(.)}{\partial p_{j}} \right| \right|_{\infty} > C_{1} \epsilon(L_{0}) \right\} \right) \\ & \leq \sum_{j=1}^{n} \mathbb{P}_{\lambda^{*}}^{\pi} \left(|| \lambda_{j}^{*}(.) - \tilde{\lambda}_{j}(.) ||_{\infty} > \Psi_{0} \left(\frac{\log k}{\sqrt{L_{0}}} \right)^{\frac{s \wedge \tilde{s}}{s+n/2}} \right) + \sum_{j=1}^{n} \sum_{i=1}^{n} \mathbb{P}_{\lambda^{*}}^{\pi} \left(\left| \left| \frac{\partial \lambda_{i}^{*}(.)}{\partial p_{j}} - \frac{\partial \tilde{\lambda}_{i}(.)}{\partial p_{j}} \right| \right|_{\infty} > \Psi_{1} \left(\frac{\log k}{\sqrt{L_{0}}} \right)^{\frac{s \wedge \tilde{s}-1}{s+n/2}} \right) \\ & \leq n(n+1)K/k, \end{split}$$

where the first inequality follows from (28), the third inequality follows since $\log k/\sqrt{L_0} \leq 1$ for $L_0 \geq \log^3 k$ and $k \geq \Phi \geq 3$, the last inequality follows by Lemma 1.

Let $M_3 = \max\{\Phi, n(n+1)(n+3)K\}$. Putting all the bounds together, for $k \ge \Phi$, we have

$$\mathbb{P}_{\lambda^*}^{\pi}(\mathcal{E}^c) \le n(n+1)K/k + n(n+1)(n+2)K/k = n(n+1)(n+3)K/k \le M_3/k.$$

As for $k \leq \Phi$, by definition of M_3 , $\mathbb{P}^{\pi}_{\lambda^*}(\mathcal{E}^c) \leq 1 \leq \Phi/k \leq M_3/k$. This completes the proof. \Box

5.3. Part 2: Proofs of Lemma 5 and Lemma 6. Let $\lambda_t := \lambda^*(p_t)$ and $\hat{\lambda}_t := \lambda(p_t; \hat{\theta}_\iota)$. We prove Lemmas 5 and 6 in turn.

Proof of Lemma 5: Let $\Omega_2 = \max\{\Omega_1, K_1, K_2, K_3\}$, where Ω_1 is as defined in the last paragraph of Part 1 in §5.1 and K_1, K_2, K_3 are positive constants to be defined later. Throughout the proof, we will implicitly assume that \mathcal{E} is satisfied. We first highlight four inequalities (see (29) - (32)) that will be useful for the proof. Let $\bar{\delta}, \kappa, \omega$ be as defined in Lemma 3 and ϕ as defined in R4. First, recall that $||\hat{\theta}_{\iota} - \theta_{\iota}^*||_2 \leq ||\hat{\theta} - \theta^*||_2 \leq \bar{\delta}$ and $||\hat{\theta}_o - \theta_o^*||_2 \leq ||\hat{\theta} - \theta^*||_2 \leq \bar{\delta}$ by (8), and $||\delta||_2 \leq \bar{\delta}$ by (9). Thus, in light of part (f) of Lemma 3,

$$||p_{\mathbf{0}}^{D}(\theta^{*}) - p_{\delta}^{D}(\hat{\theta})||_{2} \le \phi/2.$$
 (29)

Second, since $\omega \psi = \omega \sqrt{\epsilon(L_0)} \to 0$ as $k \to \infty$ (recall that $\log k / \sqrt{L_0} \le 1$ for $k \ge 3$ since $L_0 \ge \log^3 k$), there exists a constant $K_1 > 0$ such that, for all $k \ge K_1$,

$$\omega\psi \le \phi/4. \tag{30}$$

Third, note that R4 implies $\lambda^D \succ \lambda_L \mathbf{e} \succ 0$ for some $\lambda_L \in \mathbb{R}$. Hence, there exits a constant $K_2 > 0$ such that, for all $k \ge K_2$, $\psi \le \lambda_L$ (because $\psi := \sqrt{\epsilon(L_0)} \to 0$ as $k \to \infty$) and

$$\lambda_{\delta}^{D}(\hat{\theta}) = \lambda^{D} - \lambda_{\mathbf{0}}^{D}(\theta^{*}) + \lambda_{\delta}^{D}(\hat{\theta}) \succeq \lambda^{D} - ||\lambda_{\delta}^{D}(\hat{\theta}) - \lambda_{\mathbf{0}}^{D}(\theta^{*})||_{2} \mathbf{e}$$
$$\succeq \lambda^{D} - \kappa(||\theta^{*} - \hat{\theta}||_{2} + ||\delta||_{2}) \mathbf{e} \succeq \lambda^{D} - \kappa M_{2}\epsilon(L_{0})\mathbf{e} - \kappa M_{4}\epsilon(L_{0})^{2}\mathbf{e} \succeq \lambda_{L}\mathbf{e}, \qquad (31)$$

where the first equality and the second inequality follow by part (e) and (f) of Lemma 3 respectively, and the last inequality follows by (9) and the definition of \mathcal{E} .

Finally, $L_0/k \to 0$ as $k \to \infty$. So, there exists a constant $K_3 > 0$ such that, for all $k \ge K_3$,

$$C_{\tilde{L}_0} \succeq kC - \tilde{L}_0 A \mathbf{e} \succeq A \mathbf{e}.$$
(32)

The rest of the proof follows by induction. Fix some $k \ge \Omega_2$. If $\tau \le \tilde{L}_0 + 1$, there is nothing to prove. Suppose that $\tau > \tilde{L}_0 + 1$. Note that $p_{\tilde{L}_0+1} = \hat{p}_{\tilde{L}_0+1}$ because $C_{\tilde{L}_0} \succeq A\mathbf{e}$ (by (32)) and $\hat{p}_{\tilde{L}_0+1} = p_{\delta}^D(\hat{\theta}) \in \mathsf{Ball}(p^D, \phi/2) \subseteq \mathcal{P}$ (by (29)). Hence,

$$C_{\tilde{L}_{0}+1} = C_{\tilde{L}_{0}} - AD_{\tilde{L}_{0}+1} = C_{\tilde{L}_{0}} - A\lambda_{\delta}^{D}(\hat{\theta}) - A\tilde{\Delta}_{\tilde{L}_{0}+1} \succeq (k - \tilde{L}_{0} - 1)A\left(\lambda_{\delta}^{D}(\hat{\theta}) - \frac{\Delta_{\tilde{L}_{0}+1}}{k - \tilde{L}_{0} - 1}\right)$$
$$\succeq (k - \tilde{L}_{0} - 1)A\mathbf{e}\left(\lambda_{L} - \left|\left|\frac{\tilde{\Delta}_{\tilde{L}_{0}+1}}{k - \tilde{L}_{0} - 1}\right|\right|_{2}\right) \succeq (k - \tilde{L}_{0} - 1)A\mathbf{e}\left(\psi - \left|\left|\frac{\tilde{\Delta}_{\tilde{L}_{0}+1}}{k - \tilde{L}_{0} - 1}\right|\right|_{2}\right) \succeq A\mathbf{e}, \quad (33)$$

where the first inequality follows since $A\lambda_{\delta}^{D}(\hat{\theta}) \leq kC/k - \delta = C_{\tilde{L}_{0}}/(k - \tilde{L}_{0})$ (recall that $\lambda_{\delta}^{D}(\hat{\theta})$ is feasible to $\mathbf{QP}(\hat{\theta}; \delta)$), the second inequality follows by (31), the third inequality follows from the fact that $\psi \leq \lambda_{L}$ for $k \geq \Omega_{2} \geq K_{2}$, and the fourth inequality follows by the definition of τ in (†). Since A is non-negative, $C_{\tilde{L}_{0}+1} \geq A\mathbf{e} \geq A_{j}$ for all $j \in [1, n]$. This is our induction base case. Now, suppose that $C_{s} \geq A_{j}$ and $\hat{p}_{s} \in \mathcal{P}$ for all $j = 1, \ldots, n$ for all $s \in [\tilde{L}_{0}+1, t-1]$ and $t-1 < \tau$. If $t \geq \tau$, we have finished the induction. Otherwise,

$$\left| \left| \hat{p}_t - p_{\delta}^D(\hat{\theta}) \right| \right|_2 = \left| \left| (\hat{B}')^{-1} \sum_{s=\tilde{L}_0+1}^{t-1} \frac{\tilde{\Delta}_s}{k-s} \right| \right|_2 \le \omega \left| \left| \sum_{s=\tilde{L}_0+1}^{t-1} \frac{\tilde{\Delta}_s}{k-s} \right| \right|_2 \le \omega \psi \le \frac{\phi}{4}, \tag{34}$$

where the first equality follows because by the definition of $p(., \hat{\theta}_{\iota})$ we have $\nabla_{\lambda} p(\lambda; \hat{\theta}_{\iota})' = (\hat{B}')^{-1}$ for all λ , the first inequality follows by Lemma 3 part (a), the second inequality follows by (†), and the last inequality follows by (30). Combining (29) and (34), we have $\hat{p}_t \in \mathsf{Ball}(p^D, 3\phi/4) \subseteq \mathcal{P}$. By the same arguments as in (33),

$$C_{t} = C_{\tilde{L}_{0}} - \sum_{s=\tilde{L}_{0}+1}^{t} AD_{s} = C_{\tilde{L}_{0}} - \sum_{s=\tilde{L}_{0}+1}^{t} A(\hat{\lambda}_{s} + \tilde{\Delta}_{s})$$

$$= C_{\tilde{L}_{0}} - \sum_{s=\tilde{L}_{0}+1}^{t} A\left(\lambda_{\delta}^{D}(\hat{\theta}) - \sum_{v=\tilde{L}_{0}+1}^{s-1} \frac{\tilde{\Delta}_{v}}{k-v} + \tilde{\Delta}_{s}\right) \succeq \sum_{s=t+1}^{k} A\lambda_{\delta}^{D}(\hat{\theta}) - \sum_{s=\tilde{L}_{0}+1}^{t} \left(A\tilde{\Delta}_{s} - \sum_{v=\tilde{L}_{0}+1}^{s-1} \frac{A\tilde{\Delta}_{v}}{k-v}\right)$$

$$= \sum_{s=t+1}^{k} A\lambda_{\delta}^{D}(\hat{\theta}) - \sum_{s=\tilde{L}_{0}+1}^{t} \frac{(k-t)A\tilde{\Delta}_{s}}{k-s} \succeq (k-t)A\mathbf{e}\left(\lambda_{L} - \left|\left|\sum_{s=\tilde{L}_{0}+1}^{t} \frac{\tilde{\Delta}_{s}}{k-s}\right|\right|_{2}\right) \succeq A\mathbf{e} \succeq A_{j},$$

for all $j \in \overline{[1,n]}$. This completes the induction. \Box

Proof of Lemma 6: Fix $k \ge \Omega_2$. Recall that we have defined at the end of Part 2 in §5.1 an auxiliary stopping time $\tilde{\tau}$ as the minimum of k and the first time $t \ge \tilde{L}_0 + 1$ such that the following condition (††) is violated:

$$\begin{array}{l} (\dagger\dagger) \qquad \psi > \tilde{S}(t), \quad \text{where } \tilde{S}(k) := \infty \ \text{and } \forall t \in \overline{[1, k-1]}, \\ \tilde{S}(t) := \left| \left| \sum_{s = \tilde{L}_0 + 1}^t \frac{\Delta_s}{k - s} + \sum_{s = \tilde{L}_0 + 1}^t \frac{(\lambda_s - \hat{\lambda}_s)}{k - s} \mathbf{1}_{\{s \le \tau\}} \right| \right|_2 + \frac{1}{k - t}. \end{array}$$

where ψ is as defined in the definition of τ in (†). We first show that $\tau = \tilde{\tau}$ almost surely. If $\tau = t' < k$, by definition of τ , $\psi \leq S(t')$ and $\psi > S(t)$ for all t < t'. So,

$$\tilde{S}(t) = \left\| \sum_{s=\tilde{L}_{0}+1}^{t} \frac{\Delta_{s}}{k-s} + \sum_{s=\tilde{L}_{0}+1}^{t} \frac{(\lambda_{s}-\hat{\lambda}_{s})}{k-s} \mathbf{1}_{\{s\leq\tau\}} \right\|_{2}^{2} + \frac{1}{k-t}$$
$$= \left\| \sum_{s=\tilde{L}_{0}+1}^{t} \frac{\Delta_{s}}{k-s} + \sum_{s=\tilde{L}_{0}+1}^{t} \frac{(\lambda_{s}-\hat{\lambda}_{s})}{k-s} \right\|_{2}^{2} + \frac{1}{k-t} = S(t).$$

for all $t \leq t'$. But, this implies $\psi \leq \tilde{S}(t')$ and $\psi > \tilde{S}(t)$ for all t < t', which means that $\tilde{\tau} = t'$. If $\tau = k$, immediately we have $\tilde{S}(t) = S(t)$ for all t < k. Moreover, since, $\psi > S(t) = \tilde{S}(t)$ for t < k, we must have $\tilde{\tau} = k = \tau$. Define the following two terms:

$$\tilde{S}_r(t) = \left\| \sum_{s=\tilde{L}_0+1}^t \frac{\Delta_s}{k-s} \right\|_2 \quad \text{and} \quad \tilde{S}_s(t) = \sum_{s=\tilde{L}_0+1}^t \frac{||\lambda_s - \hat{\lambda}_s||_2}{k-s} \mathbf{1}_{\{s \le \tau\}}.$$

Since $\tilde{S}(t) \leq \tilde{S}_r(t) + \tilde{S}_s(t) + (k-t)^{-1}$ and $\tilde{\tau}$ is non-negative, we have:

. .

. .

$$\mathbb{E}_{\lambda^*}^{\pi}[k-\tau|\mathcal{E}] = \mathbb{E}_{\lambda^*}^{\pi}[k-\tilde{\tau}|\mathcal{E}] = \sum_{t=1}^{k-1} \mathbb{P}_{\lambda^*}^{\pi}(\tilde{\tau} \le t|\mathcal{E}) \le \sum_{t=1}^{k-1} \mathbb{P}_{\lambda^*}^{\pi}\left(\max_{\tilde{L}_0+1\le s\le t} \{\tilde{S}_r(s)\} \ge \frac{\psi}{4} \middle| \mathcal{E}\right) + \sum_{t=1}^{k-1} \mathbb{P}_{\lambda^*}^{\pi}\left(\max_{\tilde{L}_0+1\le s\le t} \{\tilde{S}_s(s)\} \ge \frac{\psi}{2} \middle| \mathcal{E}\right) + \sum_{t=1}^{k-1} \mathbb{P}_{\lambda^*}^{\pi}\left(\frac{1}{k-t} \ge \frac{\psi}{4} \middle| \mathcal{E}\right).$$
(35)

Note that $\{\tilde{S}_r(t)^2\}_{t=\tilde{L}_0+1}^{k-1}$ is a submartingale with respect to $\{\mathcal{H}_t\}_{t=\tilde{L}_0+1}^{k-1}$. So, we can bound the first term after inequality in (35) as follows. For $t \leq k-1$,

$$\mathbb{P}_{\lambda^{*}}^{\pi} \left(\max_{\tilde{L}_{0}+1 \leq s \leq t} \{\tilde{S}_{r}(s)\} \geq \frac{\psi}{4} \middle| \mathcal{E} \right) = \mathbb{P}_{\lambda^{*}}^{\pi} \left(\max_{\tilde{L}_{0}+1 \leq s \leq t} \{\tilde{S}_{r}(s)^{2}\} \geq \frac{\psi^{2}}{16} \middle| \mathcal{E} \right) \leq \frac{1}{\mathbb{P}_{\lambda^{*}}^{\pi}(\mathcal{E})} \mathbb{P}_{\lambda^{*}}^{\pi} \left(\max_{\tilde{L}_{0}+1 \leq s \leq t} \{\tilde{S}_{r}(s)^{2}\} \geq \frac{\psi^{2}}{16} \right) \\
\leq \frac{32}{\psi^{2}} \mathbb{E}_{\lambda^{*}}^{\pi} \left[\tilde{S}_{r}(t)^{2} \right] \leq \frac{32}{\psi^{2}} \mathbb{E}_{\lambda^{*}}^{\pi} \left[\sum_{s=1}^{t} \frac{||\Delta_{s}||_{2}^{2}}{(k-s)^{2}} \right] \leq \frac{32}{\psi^{2}} \left[\frac{2}{(k-t)^{2}} + \frac{2}{k-t} \right] \leq \frac{128}{\psi^{2}(k-t)},$$
(36)

where the second inequality follows by Doob's submartingale inequality and (7), the third inequality follows because $\mathbb{E}_{\lambda^*}^{\pi}[\Delta'_s \Delta_t] = 0$ if $s \neq t$, the fourth inequality follows by integral comparison and the fact that $||\Delta_t||_2^2 \leq 2$, and the last inequality follows because $k - t \geq 1$. Thus, there exists $K_5 > 0$ independent of $k \geq \Omega_2$ such that

$$\sum_{t=1}^{k-1} \mathbb{P}_{\lambda^*}^{\pi} \left(\max_{\tilde{L}_0 + 1 \le s \le t} \{ \tilde{S}_r(s) \} \ge \frac{\psi}{4} \, \middle| \, \mathcal{E} \right) \le \frac{128}{\psi^2} \sum_{t=1}^{k-1} \frac{1}{k-t} \le K_5 \psi^{-2} \log k = K_5 \epsilon(L_0)^{-1} \log k$$

where the equality follows by the definition of ψ in (†). We now bound the second term in (35). By Lemma 3 part (b) and (8), there exists a constant $K_4 > 0$ independent of k such that, for all $k \ge \Omega_2 \ge \Omega_1$, we have:

$$\begin{aligned} ||\lambda_s - \hat{\lambda}_s||_2 &= ||\lambda^*(p_s) - \lambda(p_s; \hat{\theta}_\iota)||_2 \\ &\leq ||\lambda^*(p_s)||_2 + ||\lambda(p_s; \theta^*_\iota)||_2 + ||\lambda(p_s; \theta^*_\iota) - \lambda(p_s; \hat{\theta}_\iota)||_2 \\ &\leq \sqrt{n} ||\lambda^*(.)||_{\infty} + \sqrt{n} ||\lambda(.; \theta^*_\iota)||_{\infty} + \omega \bar{\delta} \leq K_4, \end{aligned}$$
(37)

Conditioning on \mathcal{E} , for $s < \tau$, we can derive a sharper bound, i.e., $||\lambda_s - \hat{\lambda}_s||_2 \leq \omega_0 \epsilon(L_0)$ for some constant ω_0 independent of k (see (44) for the derivation). Now, observe that the following holds:

$$\mathbb{P}_{\lambda^{*}}^{\pi} \left(\max_{\tilde{L}_{0}+1 \leq s \leq t} \{\tilde{S}_{s}(s)\} \geq \frac{\psi}{2} \middle| \mathcal{E} \right) = \mathbb{P}_{\lambda^{*}}^{\pi} \left(\tilde{S}_{s}(t) \geq \frac{\psi}{2} \middle| \mathcal{E} \right) \leq \frac{16}{\psi^{4}} \mathbb{E}_{\lambda^{*}}^{\pi} \left[\left(\sum_{s=\tilde{L}_{0}+1}^{t} \frac{||\lambda_{s} - \hat{\lambda}_{s}||_{2} \mathbf{1}_{\{s < \tau\}}}{k-s} + \sum_{s=\tilde{L}_{0}+1}^{t} \frac{||\lambda_{s} - \hat{\lambda}_{s}||_{2} \mathbf{1}_{\{s = \tau\}}}{k-s} \right)^{4} \middle| \mathcal{E} \right] \\
\leq \frac{128}{\psi^{4}} \mathbb{E}_{\lambda^{*}}^{\pi} \left[\left(\sum_{s=\tilde{L}_{0}+1}^{t} \frac{||\lambda_{s} - \hat{\lambda}_{s}||_{2} \mathbf{1}_{\{s < \tau\}}}{k-s} \right)^{4} \middle| \mathcal{E} \right] + \frac{128}{\psi^{4}} \mathbb{E}_{\lambda^{*}}^{\pi} \left[\left(\sum_{s=\tilde{L}_{0}+1}^{t} \frac{||\lambda_{s} - \hat{\lambda}_{s}||_{2} \mathbf{1}_{\{s < \tau\}}}{k-s} \right)^{4} \middle| \mathcal{E} \right] \\
\leq \frac{128\omega_{0}^{4}\epsilon(L_{0})^{4}}{\psi^{4}} \log^{4} \left(\frac{k}{k-t} \right) + \frac{128}{\psi^{4}} \left(\frac{K_{4}}{k-t} \right)^{4} \\
= 128\omega_{0}^{4}\epsilon(L_{0})^{2} \log^{4} \left(\frac{k}{k-t} \right) + 128K_{4}^{4}\epsilon(L_{0})^{-2} \left(\frac{1}{k-t} \right)^{4}$$
(38)

where the first equality follows by the monotonicity of $\tilde{S}_s(t)$, the first inequality follows by Markov's inequality, the second inequality follows since $(a+b)^4 \leq 8a^4 + 8b^4$, the third inequality follows by (37) and (44), and the last equality follows since $\psi = \sqrt{\epsilon(L_0)}$. So, there exists a constant $K_6 > 0$ independent of $k \geq \Omega_2$ such that:

$$\begin{split} \sum_{t=1}^{k-1} \mathbb{P}_{\lambda^*}^{\pi} \left(\max_{\tilde{L}_0 + 1 \le s \le t} \{ \tilde{S}_s(t) \} \ge \frac{\psi}{2} \left| \mathcal{E} \right) &\le \sum_{t=1}^{k-1} 128\omega_0^4 \epsilon(L_0)^2 \log^4 \left(\frac{k}{k-t} \right) + 128K_4^4 \epsilon(L_0)^{-2} \sum_{t=1}^{k-1} \left(\frac{1}{k-t} \right)^4 \\ &\le K_6 \left(\epsilon(L_0)^2 k + \epsilon(L_0)^{-2} \right), \end{split}$$

where the last inequality follows from the following claim:

CLAIM 1. For any
$$s \in \mathbb{Z}_{++}$$
, $\sum_{t=1}^{k-1} \log^s(\frac{k}{k-t}) \leq s! h$

Proof: Note that by integral inequality $\sum_{t=1}^{k-1} \log^s(\frac{k}{k-t}) \leq \int_0^{k-1} \log^s(\frac{k}{k-t}) dt = \int_1^k \log^s(\frac{k}{t}) dt$. Then the proof follows by an induction argument. When s = 1, $\int_1^k \log(\frac{k}{t}) dt = \int_1^k \log k dt - \int_1^k \log t dt = (k-1)\log k - (k\log k - (k-1)) \leq k$. Suppose for all $s \leq n-1$, $\int_1^k \log^s(\frac{k}{t}) dt \leq s!k$, then

$$\begin{split} \int_{1}^{k} \log^{n}\left(\frac{k}{t}\right) dt &= \log k \int_{1}^{k} \log^{n-1}\left(\frac{k}{t}\right) dt - \int_{1}^{k} \log^{n-1}\left(\frac{k}{t}\right) \log t dt \\ &\leq (n-1)! k \log k - \left[t \log t \log^{n-1}\left(\frac{k}{t}\right)\Big|_{t=1}^{k} - \int_{1}^{k} \left(\log^{n-1}\left(\frac{k}{t}\right) - (n-1) \log^{n-2}\left(\frac{k}{t}\right) \log t\right) dt \\ &= (n-1)! k \log k + \int_{1}^{k} \log^{n-1}\left(\frac{k}{t}\right) dt - \int_{1}^{k} (n-1) \log^{n-2}\left(\frac{k}{t}\right) (\log t - \log k + \log k) dt \\ &= (n-1)! (k \log k + k) + (n-1) \int_{1}^{k} \log^{n-1}\left(\frac{k}{t}\right) dt - \log k(n-1) \int_{1}^{k} \log^{n-2}\left(\frac{k}{t}\right) dt \\ &= (n-1)! (k \log k + k) + (n-1)(n-1)! k - (n-1)! k \log k = n! k. \end{split}$$

This completes the proof of Claim 1. \Box

The third term of (35) can be bounded as follows:

$$\sum_{t=1}^{k-1} \mathbb{P}_{\lambda^*}^{\pi} \left(\left. \frac{1}{k-t} \ge \frac{\psi}{4} \right| \mathcal{E} \right) = \sum_{t=1}^{k-1} \mathbb{P}_{\lambda^*}^{\pi} \left(\frac{1}{k-t} \ge \frac{\psi}{4} \right) = 4/\psi = 4\epsilon (L_0)^{-1/2}.$$

Putting the bounds for the three terms in (35) together, we conclude that there exists a constant $M_5 > 0$ such that, for all $k \ge \Omega_2$, we have $\mathbb{E}_{\lambda^*}^{\pi}[k-\tau|\mathcal{E}] = \mathbb{E}_{\lambda^*}^{\pi}[k-\tilde{\tau}|\mathcal{E}] \le K_5 \epsilon(L_0)^{-1} \log k + K_6 \epsilon(L_0)^2 k + K_6 \epsilon(L_0)^{-2} + 4\epsilon(L_0)^{-1/2} \le M_5(\epsilon(L_0)^{-1} \log k + \epsilon(L_0)^{-1}) \log k + \epsilon(L_0)^{-2}).$

5.4. Part 3: Derivation of (11) and (12). For simplicity, we suppress the dependency of $\epsilon(L_0)$ on L_0 and simply write $\epsilon(L_0)$ as ϵ throughout this section. Recall that we define $\lambda_t := \lambda^*(p_t)$ and $\hat{\lambda}_t := \lambda(p_t; \hat{\theta}_t)$ in §5.3. Also, by Lemma 5, for all $k \ge \Omega_2$ and all sample paths on \mathcal{E} ,

$$\hat{\lambda}_t = \lambda(\hat{p}_t; \hat{\theta}_\iota) = \lambda_\delta^D(\hat{\theta}) - \sum_{s=\tilde{L}_0+1}^{t-1} \frac{\tilde{\Delta}_s}{k-s} \quad \text{for all } t < \tau.$$
(39)

These will be used multiple times in the derivation of (11) and (12).

Derivation of inequality (11). Note that $\nabla r_{\lambda}^{*}(\lambda^{D}) = \nabla_{\lambda}q(p(\lambda^{D};\theta_{\iota}^{*});\theta_{o}^{*})$ by (3) and $\lambda_{\mathbf{0}}^{D}(\theta^{*}) = \lambda^{D}$ by Lemma 3 part (e). So, we can write: $\nabla r_{\lambda}^{*}(\lambda^{D}) \cdot (\lambda^{D} - \lambda_{t}) = \nabla_{\lambda}q(p(\lambda_{\mathbf{0}}^{D}(\theta^{*});\theta_{\iota}^{*});\theta_{o}^{*}) \cdot (\lambda_{\mathbf{0}}^{D}(\theta^{*}) - \lambda_{t}) = \mu_{\mathbf{0}}^{D}(\theta^{*})'A(\lambda_{\mathbf{0}}^{D}(\theta^{*}) - \lambda_{\delta}^{D}(\hat{\theta}) + \lambda_{\delta}^{D}(\hat{\theta}) - \hat{\lambda}_{t} + \hat{\lambda}_{t} - \lambda_{t})$, where the last equality follows by the Karush-Kuhn-Tucker(KKT) optimality condition. Therefore, for all $k \geq \Omega$, the first term of (11) can be broken into two parts:

$$\mathbb{E}_{\lambda^{*}}^{\pi} \left[\sum_{t=\tilde{L}_{0}+1}^{\tau-1} \nabla r_{\lambda}^{*}(\lambda^{D}) \cdot (\lambda^{D} - \lambda_{t}) \middle| \mathcal{E} \right] \\
= \mathbb{E}_{\lambda^{*}}^{\pi} \left[\sum_{t=\tilde{L}_{0}+1}^{\tau-1} \mu_{\mathbf{0}}^{D}(\theta^{*})' (A\lambda_{\mathbf{0}}^{D}(\theta^{*}) - A\lambda_{\delta}^{D}(\hat{\theta})) \middle| \mathcal{E} \right] \\
+ \mathbb{E}_{\lambda^{*}}^{\pi} \left[\sum_{t=\tilde{L}_{0}+1}^{\tau-1} \mu_{\mathbf{0}}^{D}(\theta^{*})' A(\lambda_{\delta}^{D}(\hat{\theta}) - \hat{\lambda}_{t} + \hat{\lambda}_{t} - \lambda_{t}) \middle| \mathcal{E} \right] \tag{40}$$

By Lemma 3 part (f), for all sample paths on \mathcal{E} , the set of constraints of $\mathbf{QP}(\theta^*; \mathbf{0})$ that have nonzero optimal dual variables are binding at the optimal solution $\lambda_{\delta}^{D}(\hat{\theta})$ in $\mathbf{QP}(\hat{\theta}; \delta)$. This implies that the first expectation after the equality in (40) is zero because, for all *i*, either we have $\mu_{\mathbf{0},i}^{D}(\theta^*) = 0$ or $(A\lambda_{\mathbf{0}}^{D}(\theta^*))_{i} - (A\lambda_{\delta}^{D}(\hat{\theta}))_{i} = 0$. As for the second expectation, by (39) and the definition of $\tilde{\Delta}_{t}$ (i.e., $\tilde{\Delta}_{t} = \Delta_{t} + \lambda_{t} - \hat{\lambda}_{t}$), we can write:

$$\mathbb{E}_{\lambda^{*}}^{\pi} \left[\sum_{t=\tilde{L}_{0}+1}^{\tau-1} \mu_{\mathbf{0}}^{D}(\theta^{*})' A(\lambda_{\delta}^{D}(\hat{\theta}) - \hat{\lambda}_{t} + \hat{\lambda}_{t} - \lambda_{t}) \middle| \mathcal{E} \right] \\
= \mathbb{E}_{\lambda^{*}}^{\pi} \left[\sum_{t=\tilde{L}_{0}+1}^{\tau-1} \mu_{\mathbf{0}}^{D}(\theta^{*})' \left(\sum_{s=\tilde{L}_{0}+1}^{t-1} \frac{A\tilde{\Delta}_{s}}{k-s} + A\Delta_{t} - A\tilde{\Delta}_{t} \right) \middle| \mathcal{E} \right] \\
= \mathbb{E}_{\lambda^{*}}^{\pi} \left[\sum_{t=\tilde{L}_{0}+1}^{\tau-1} \mu_{\mathbf{0}}^{D}(\theta^{*})' A\Delta_{t} \middle| \mathcal{E} \right] + \mathbb{E}_{\lambda^{*}}^{\pi} \left[\sum_{t=\tilde{L}_{0}+1}^{\tau-1} \left(\frac{\tau-t-1}{k-t} - 1 \right) \mu_{\mathbf{0}}^{D}(\theta^{*})' A\tilde{\Delta}_{t} \middle| \mathcal{E} \right]. \quad (41)$$

Since $\{\Delta_t\}_{t=\tilde{L}_0+1}^{k-1}$ is a martingale difference sequence with respect to $\{\mathcal{H}_t\}_{t=\tilde{L}_0+1}^{k-1}$, we can bound:

$$\mathbb{E}_{\lambda^{*}}^{\pi} \left[\sum_{t=\tilde{L}_{0}+1}^{\tau-1} \mu_{\mathbf{0}}^{D}(\theta^{*})' A \Delta_{t} \middle| \mathcal{E} \right] = \frac{\mu_{\mathbf{0}}^{D}(\theta^{*})' A}{\mathbb{P}_{\lambda^{*}}^{\pi}(\mathcal{E})} \left\{ \mathbb{E}_{\lambda^{*}}^{\pi} \left[\sum_{t=\tilde{L}_{0}+1}^{\tau-1} \Delta_{t} \right] - \mathbb{E}_{\lambda^{*}}^{\pi} \left[\sum_{t=\tilde{L}_{0}+1}^{\tau-1} \Delta_{t} \middle| \mathcal{E}^{c} \right] \mathbb{P}_{\lambda^{*}}^{\pi}(\mathcal{E}^{c}) \right\} \\
\leq \mu_{\mathbf{0}}^{D}(\theta^{*})' A \mathbf{e} \frac{1 + k \mathbb{P}_{\lambda^{*}}^{\pi}(\mathcal{E}^{c})}{\mathbb{P}_{\lambda^{*}}^{\pi}(\mathcal{E})} \leq 2(1 + M_{3}) \mu_{\mathbf{0}}^{D}(\theta^{*})' A \mathbf{e},$$
(42)

where the first inequality follows from $\mathbb{E}_{\lambda^*}^{\pi}[\sum_{t=L+1}^{\tau-1} \Delta_t] = \mathbb{E}_{\lambda^*}^{\pi}[\sum_{t=L+1}^{\tau} \Delta_t] - \mathbb{E}_{\lambda^*}^{\pi}[\Delta_{\tau}]$ (by Optional Stopping Time Theorem) and the fact that $|\Delta_t| \prec \mathbf{e}$, and the second inequality follows by Lemma 4 and (7). The second term of (41) can be bounded as follows:

$$\mathbb{E}_{\lambda^{*}}^{\pi} \left[\sum_{t=\tilde{L}_{0}+1}^{\tau-1} \left(\frac{\tau-t-1}{k-t} - 1 \right) \mu_{\mathbf{0}}^{D}(\theta^{*})' A \tilde{\Delta}_{t} \middle| \mathcal{E} \right] \leq \mathbb{E}_{\lambda^{*}}^{\pi} \left[(k-\tau+1) \left| \mu_{\mathbf{0}}^{D}(\theta^{*})' \sum_{t=\tilde{L}_{0}+1}^{\tau-1} \frac{A \tilde{\Delta}_{t}}{k-t} \right| \middle| \mathcal{E} \right] \\ \leq \mathbb{E}_{\lambda^{*}}^{\pi} \left[(k-\tau+1) \left| \left| \mu_{\mathbf{0}}^{D}(\theta^{*}) \right| \right|_{2} \left| \left| A \right| \right|_{2} \left\| \sum_{t=\tilde{L}_{0}+1}^{\tau-1} \frac{\tilde{\Delta}_{t}}{k-t} \right\|_{2} \middle| \mathcal{E} \right] \\ \leq \psi \left| \left| \mu_{\mathbf{0}}^{D}(\theta^{*}) \right| \right|_{2} \left| \left| A \right| \right|_{2} \left(\mathbb{E}_{\lambda^{*}}^{\pi} \left[k-\tau | \mathcal{E} \right] + 1 \right) \\ \leq \left| \left| \mu_{\mathbf{0}}^{D}(\theta^{*}) \right| \right|_{2} \left| \left| A \right| \right|_{2} \left(\mathbb{E}_{\lambda^{*}}^{\pi} \left[k-\tau | \mathcal{E} \right] + 1 \right) \\ \leq \left| \left| \mu_{\mathbf{0}}^{D}(\theta^{*}) \right| \right|_{2} \left| \left| A \right| \right|_{2} \left[M_{5}(\epsilon^{2}k + \epsilon^{-1}\log k + \epsilon^{-2}) + 1 \right]. \quad (43)$$

where the third inequality follows by (†), and the fourth inequality follows by Lemma 6 and the fact that $\psi = \sqrt{\epsilon(L_0)} < 1$ for $k \ge \Omega \ge \Omega_3$. Putting together (40) - (43) yields

$$\mathbb{E}_{\lambda^*}^{\pi} \left[\sum_{t=\tilde{L}_0+1}^{\tau-1} \nabla r_{\lambda}^*(\lambda^D) \cdot (\lambda^D - \lambda_t) \right| \mathcal{E} \right] \le M_6 (1 + \epsilon^2 k + \epsilon^{-1} \log k + \epsilon^{-2})$$

where $M_6 = 2(1+M_3)\mu_0^D(\theta^*)'A\mathbf{e} + ||\mu_0^D(\theta^*)||_2||A||_2(1+M_5).$

Derivation of inequality (12). By definition, $\lambda(p_t; \theta_{\iota}^*) = \lambda^*(p^D) + (\nabla \lambda^*(p^D))'(p_t - p^D)$. Since $\lambda^*(p_t) = \lambda^*(p^D) + \nabla \lambda^*(p^D)'(p_t - p^D) + (p_t - p^D)'\nabla^2 \lambda^*(\xi)(p_t - p^D)$ for some $\xi \in \mathcal{P}$ and $\sup_{\hat{\xi} \in \mathcal{P}} ||(p_t - p^D)'\nabla^2 \lambda^*(\hat{\xi})(p_t - p^D)||_2 \le \kappa_0 ||p_t - p^D||_2^2$ for some $\kappa_0 > 0$ (by R1 and the compactness of \mathcal{P}),

$$||\lambda^{*}(p_{t}) - \lambda(p_{t};\theta_{\iota}^{*})||_{2} \leq \kappa_{0}||p_{t} - p^{D}||_{2}^{2}$$

So, conditioning on \mathcal{E} , for all $t < \tau$, we have:

$$\begin{aligned} |\lambda_{t} - \hat{\lambda}_{t}||_{2} &= ||\lambda^{*}(p_{t}) - \lambda(p_{t};\hat{\theta}_{\iota})||_{2} \\ &\leq ||\lambda^{*}(p_{t}) - \lambda(p_{t};\theta_{\iota}^{*})||_{2} + ||\lambda(p_{t};\theta_{\iota}^{*}) - \lambda(p_{t};\hat{\theta}_{\iota})||_{2} \\ &\leq \kappa_{0}||p_{t} - p^{D}||_{2}^{2} + \omega M_{2}\epsilon \\ &= \kappa_{0}||p_{t} - p^{D}_{\delta}(\hat{\theta}) + p^{D}_{\delta}(\hat{\theta}) - p^{D}_{0}(\theta^{*})||_{2}^{2} + \omega M_{2}\epsilon \\ &\leq 2\kappa_{0}||p_{t} - p^{D}_{\delta}(\hat{\theta})||_{2}^{2} + 2\kappa_{0}||p^{D}_{\delta}(\hat{\theta}) - p^{D}_{0}(\theta^{*})||_{2}^{2} + \omega M_{2}\epsilon \\ &\leq 2\kappa_{0}||\hat{p}_{t} - p^{D}_{\delta}(\hat{\theta})||_{2}^{2} + 2\kappa_{0}\kappa^{2}(||\delta||_{2} + ||\hat{\theta} - \theta^{*}||_{2})^{2} + \omega M_{2}\epsilon \\ &\leq 2\kappa_{0}\omega^{2}\psi^{2} + 2\kappa_{0}\kappa^{2}(M_{4}\epsilon^{2} + M_{2}\epsilon)^{2} + \omega M_{2}\epsilon \leq \omega_{0}\epsilon \end{aligned}$$
(44)

where the second inequality follows by Lemma 3 part (b) and the definition of \mathcal{E} , the fourth inequality follows by Lemma 5 (i.e., $p_t = \hat{p}_t$ for $t < \tau$) and Lemma 3 part (f), and the fifth inequality follows by (9), (34) and the definition of \mathcal{E} . Now,

$$\begin{split} &\frac{1}{2} \mathbb{E}_{\lambda^{*}}^{\pi} \left[\sum_{t=\tilde{L}_{0}+1}^{\tau-1} ||\lambda_{\delta}^{D}(\hat{\theta}) - \lambda_{t}||_{2}^{2} \middle| \mathcal{E} \right] \leq \mathbb{E}_{\lambda^{*}}^{\pi} \left[\sum_{t=\tilde{L}_{0}+1}^{\tau-1} ||\lambda_{\delta}^{D}(\hat{\theta}) - \hat{\lambda}_{t}||_{2}^{2} \middle| \mathcal{E} \right] + \mathbb{E}_{\lambda^{*}}^{\pi} \left[\sum_{t=\tilde{L}_{0}+1}^{\tau-1} ||\hat{\lambda}_{t} - \lambda_{t}||_{2}^{2} \middle| \mathcal{E} \right] \\ &= \mathbb{E}_{\lambda^{*}}^{\pi} \left[\sum_{t=\tilde{L}_{0}+1}^{\tau-1} \left| \left| \sum_{s=\tilde{L}_{0}+1}^{t-1} \frac{\tilde{\Delta}_{s}}{k-s} \right| \right|_{2}^{2} \middle| \mathcal{E} \right] + \mathbb{E}_{\lambda^{*}}^{\pi} \left[\sum_{t=\tilde{L}_{0}+1}^{\tau-1} ||\hat{\lambda}_{t} - \lambda_{t}||_{2}^{2} \middle| \mathcal{E} \right] \\ &\leq 2 \mathbb{E}_{\lambda^{*}}^{\pi} \left[\sum_{t=\tilde{L}_{0}+1}^{\tau-1} \left(\left| \left| \sum_{s=\tilde{L}_{0}+1}^{t-1} \frac{\Delta_{s}}{k-s} \right| \right|_{2}^{2} + \left(\sum_{s=\tilde{L}_{0}+1}^{t-1} \frac{||\lambda_{s} - \hat{\lambda}_{s}||_{2}}{k-s} \right)^{2} \right) \middle| \mathcal{E} \right] + \mathbb{E}_{\lambda^{*}}^{\pi} \left[\sum_{t=\tilde{L}_{0}+1}^{\tau-1} ||\hat{\lambda}_{t} - \lambda_{t}||_{2}^{2} \middle| \mathcal{E} \right] \\ &\leq 2 \mathbb{E}_{\lambda^{*}}^{\pi} \left(\mathcal{E}_{\lambda^{*}}^{\tau-1} \left[\sum_{s=1}^{t-1} \frac{||\sum_{s=1}^{t-1} \Delta_{s}}{k-s} \right] \right|_{2}^{2} + 2 \mathbb{E}_{\lambda^{*}}^{\pi} \left[\sum_{s=\tilde{L}_{0}+1}^{t-1} \frac{\omega_{0}\epsilon}{k-s} \right)^{2} \middle| \mathcal{E} \right] + \mathbb{E}_{\lambda^{*}}^{\pi} \left[\sum_{t=\tilde{L}_{0}+1}^{\tau-1} ||\hat{\lambda}_{t} - \lambda_{t}||_{2}^{2} \middle| \mathcal{E} \right] \end{aligned}$$

$$\leq 4 \sum_{t=1}^{k-1} \sum_{s=1}^{t-1} \frac{\mathbb{E}_{\lambda^*}^{\pi}[||\Delta_s||_2^2]}{(k-s)^2} + 2\omega_0^2 \epsilon^2 \sum_{t=1}^{k-1} \left(\sum_{s=1}^{t-1} \frac{1}{k-s}\right)^2 + \omega_0^2 k \epsilon^2$$

$$\leq 8 \sum_{t=1}^{k-1} \sum_{s=1}^{t-1} \frac{1}{(k-s)^2} + 2\omega_0^2 \epsilon^2 \sum_{t=1}^{k-1} \log^2 \left(\frac{k}{k-t}\right) + \omega_0^2 \epsilon^2 k \leq 8 \log k + 5\omega_0^2 \epsilon^2 k = M_{10}(\log k + \epsilon^2 k), (45)$$

where $M_{10} := 8 + 5\omega_0^2$, the third inequality follows from (44), the fourth inequality follows by (7) and the fact that $\mathbb{E}_{\lambda^*}^{\pi}[\Delta'_s \Delta_t] = 0$ if $s \neq t$, and the last two inequalities follows by integral comparison and the fact that $\sum_{t=1}^{k-1} \log^s(\frac{k}{k-t}) \leq \int_1^k \log^s(\frac{k}{k-t}) dt \leq s!k$. The derivation of inequality (12) is completed by noting that

$$\begin{split} & \frac{\bar{v}}{2} \mathbb{E}_{\lambda^*}^{\pi} \left[\sum_{t=\tilde{L}_0+1}^{\tau-1} \left| \left| \lambda_0^D(\theta^*) - \lambda_t \right) \right| \right|_2^2 \middle| \mathcal{E} \right] \leq \bar{v} \mathbb{E}_{\lambda^*}^{\pi} \left[\sum_{t=\tilde{L}_0+1}^{\tau-1} \left(\left| \left| \lambda_0^D(\theta^*) - \lambda_\delta^D(\hat{\theta}) \right| \right|_2^2 + \left| \left| \lambda_\delta^D(\hat{\theta}) - \lambda_t \right| \right|_2^2 \right) \middle| \mathcal{E} \right] \\ & \leq \bar{v} k \mathbb{E}_{\lambda^*}^{\pi} \left[\left(\kappa ||\theta^* - \hat{\theta}||_2 + \kappa ||\delta||_2 \right)^2 \middle| \mathcal{E} \right] + \bar{v} \mathbb{E}_{\lambda^*}^{\pi} \left[\sum_{t=\tilde{L}_0+1}^{\tau-1} \left| \left| \lambda_\delta^D(\hat{\theta}) - \lambda_t \right| \right|_2^2 \middle| \mathcal{E} \right] \\ & \leq 2\bar{v} k \left(\kappa^2 M_2^2 \epsilon^2 + \kappa^2 M_4^2 \epsilon^4 \right) + 2\bar{v} M_{10} (\log k + \epsilon^2 k) \\ & \leq 2\bar{v} M_{10} \log k + (2\bar{v} M_{10} + 2\bar{v} \kappa^2 M_4^2 + 2\bar{v} \kappa^2 M_2^2) \epsilon^2 k = M_7 (\log k + \epsilon^2 k) \end{split}$$

where $M_7 := 4\bar{v}M_{10} + 2\bar{v}\kappa^2 M_4^2 + 2\bar{v}\kappa^2 M_2^2$. The second inequality follows by Lemma 3 part (f), the third inequality follows by the definition of \mathcal{E} , (9) and (45), and the fourth inequality follows by the fact that $\epsilon < 1$ for $k \ge \Omega \ge \Omega_3$.

To summarize, the proofs of Lemma 4 in §5.2, Lemma 5 and Lemma 6 in §5.3, and the derivation of (11) and (12) in §5.4 fill in the gaps in the outline in §5.1. This completes the proof of Theorem 1.

6. Closing remarks. We study the problem of joint learning and pricing in a general capacitated network RM problem with multiple products and multiple limited resources. We develop a heuristic called NSC that combines Spline Estimation, linear approximation of the estimated demand function, quadratic approximation of the estimated revenue function, and self-adjusting price updates. We show analytically that if the underlying demand function is sufficiently smooth, the revenue loss under NSC is $\mathcal{O}(k^{1/2+\epsilon}\log k)$ for any fixed $\epsilon > 0$. This is the tightest bound of its kind and is very close to the known theoretical lower bound of $\Omega(\sqrt{k})$. Our result suggests the applicability of self-adjusting controls in dynamic pricing problems. Moreover, in proving our main result, we derive large deviation bounds for spline estimation and prove a nonparametric stability result of the optimal solution of a constrained optimization problem. These results are of independent interest and are potentially useful for other application areas.

Acknowledgments. The authors thank the Area Editor Adam Shwartz, the Associate Editor and four anonymous referees whose insightful comments greatly improved the paper.

References

- [1] Atar R, Reiman M (2012) Aymptotically optimal dynamic pricing for network revenue management. Stochastic Systems 2:232 – 276.
- [2] Auer P, Cesa-Bianchi N, Fischer P (2002a) Finite-time analysis of the multiarmed bandit problem. Machine Learning 47:235–256.
- [3] Auer P, Cesa-Bianchi N, Freund Y, Schapire R (2002b) The nonstochastic multiarmed bandit problem. SIAM J. COMPUT. 32:48-77.
- [4] Badanidiyuru A, Kleinberg R, Slivkins A (2013) Bandits with knapsacks. FOCS '13 Proceedings of the 2013 IEEE 54th Annual Symposium on Foundations of Computer Science, 207–216.

- [5] Badanidiyuru A, Kleinberg R, Slivkins A (2015) Bandits with knapsacks. Working paper .
- [6] Besbes O, Zeevi A (2009) Dynamic pricing without knowing the demand function: Risk bound and near-optimal algorithms. Oper. Res. 57:1407–1420.
- [7] Besbes O, Zeevi A (2012) Blind network revenue management. Oper. Res. 60:1537–1550.
- [8] Besbes O, Zeevi A (2015) On the (surprising) sufficiency of linear models for dynamic pricing with demand learning. *Management Sci.* 61:723–739.
- [9] Bitran G, Caldentey R (2003) An overview of pricing models for revenue management. Manufacturing Service Oper. Management 5:203-229.
- [10] Bonnans J, Shapiro A (2000) Perturbation Analysis of Optimization Problems (Springer).
- Broder J, Rusmevichientong P (2012) Dynamic pricing under a general parametric choice model. Oper. Res. 60:965–980.
- [12] Chen Q, Jasin S, Duenyas I (2016) Realtime pricing with minimal and flexible price adjustment. Management Sci. 62:2437–2455.
- [13] Chen Q, Jasin S, Duenyas I (2017) Nonparametric self-adjusting control: An extension to compound poisson process. https://sites.google.com/a/umich.edu/georgeqc/research.
- [14] Combes R, Jiang C, Srikant R (2015) Bandits with budgets: Regret lower bounds and optimal algorithms. Proceedings of the 2015 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems, 245–257.
- [15] Elmaghraby W, Keskinocak P (2003) Dynamic pricing in the presence of inventory considerations: Research overview, current practices, and future directions. *Management Sci.* 49:1287–1309.
- [16] Ferreira KJ, Simchi-Levi D, Wang H (2015) Online network revenue management using thompson sampling. Under 2nd round review in Operations Research .
- [17] Flajolet A, Jaillet P (2017) Logarithmic regret bounds for bandits with knapsacks. Working paper.
- [18] Gallego G, van Ryzin G (1994) Optimal dynamic pricing of inventory with stochastic demand over finite horizons. *Management Sci.* 40:999–1020.
- [19] Gallego G, van Ryzin G (1997) A multiproduct dynamic pricing problem and its applications to network yield management. Oper. Res. 45:24–41.
- [20] Gyorfi L, Kohler M, Krzyzak A, Walk H (2002) A Distribution-Free Theory of Nonparametric Regression (Springer).
- [21] Jasin S (2014) Reoptimization and self-adjusting price control for network revenue management. Oper. Res. 62:1168–1178.
- [22] Jasin S (2015) Performance of an lp-based control for revenue management with unknown demand parameters. Oper. Res. 63:909–915.
- [23] Keskin N, Zeevi A (2014) Dynamic pricing with an unknown demand model: Asymptotically optimal semi-myopic policies. Oper. Res. 62:1142–1167.
- [24] Lai T, Robbins H (1985) Asymptotic efficient adaptive allocation rules. Advances in Applied Mathematics 6:4–22.
- [25] Lei Y, Jasin S, Sinha A (2014) Near-optimal bisection search for nonparametric dynamic pricing with inventory constraint. Working paper.
- [26] Maglaras C, Meissner J (2006) Dynamic pricing strategies for multiproduct revenue management problems. Manufacturing Service Oper. Management 8:136–148.
- [27] Ozer O, Phillips R (2012) The Oxford Handbook of Pricing Management (Oxford University Press).
- [28] Schumaker L (2007) Spline Functions: Basic Theory (Third Edition) (Cambridge University Press).
- [29] Talluri K, van Ryzin G (2005) The theory and Practice of Revenue Management (Springer).
- [30] Wang Z, Deng S, Ye Y (2014) Closing the gaps: A learning-while-doing algorithm for single-product revenue management problems. *Oper. Res.* 62:318–331.

Online Appendices: A Nonparametric Self-Adjusting Control for Joint Learning and Optimization of Multi-Product Pricing with Finite Resource Capacity

Qi (George) Chen

Naveen Jindal School of Management, University of Texas at Dallas, Richardson, TX 75080 USA georgeqc@utdallas.edu

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

Appendix A: Proof of Lemma 1

Let $\bar{\delta}_l := (\bar{p}_l - p_l)/(d+1)$. The proof of Lemma 1 depends on two important lemmas, which we first state and prove later:

LEMMA A1. Define $\mathcal{X} := \bigotimes_{l=1}^{n} [0, \bar{x}_l]$ where $0 < \bar{x}_l \leq 1$ for all $l \in \overline{[1, n]}$. Let $f : \mathcal{X} \to \mathbb{R}$ be a function that satisfies N1-N2. Let s be a positive integer and \bar{s} be as defined in N1. There exists $g \in \bigotimes_{l=1}^{n} \mathsf{P}^{(s \wedge \bar{s})-1}([0, \bar{x}_{l}])$ such that for any $r \in [0, s \wedge \bar{s}]$, and any $r_{l} \in \mathbb{Z}_{+}, l \in [1, n]$ satisfying $\sum_{l=1}^{n} r_{l} = r$, the following holds

 $\left\| \left\| \frac{\partial^r (f-g)(.)}{\partial x_1^{r_1} \dots \partial x_n^{r_n}} \right\|_{\infty} \le C_{n,r} W \left[\max_{l=1,\dots,n} \{ \bar{x}_l \} \right]^{s \wedge \bar{s} - r},$

where $C_{n,r} > 0$ only depends on n, r and W is as defined in N2.

LEMMA A2. Suppose $s \ge 2$. Let \mathcal{L} , $\{y_{l,i}\}_{l=1,i=1}^{n,2s+d}$, $\{\beta_{l,i,j}\}_{l=1,i=1,j=1}^{n,s+d,s}$ and $\{N_{i_1,\ldots,i_n}(.)\}_{i_1=1,\ldots,i_n=1}^{s+d,\ldots,s+d}$ be as defined in the Technical Details for Spline Approximations in §3.1. The following properties hold:

a. $\mathcal{L}f = f, \forall f \in \bigotimes_{l=1}^{n} \mathsf{P}^{s-1}([p_l, \bar{p}_l]).$

 $\begin{array}{l} b. \ For \ all \ l \in \overline{[1,n]}, i \in \overline{[1,s+d]}, j \in \overline{[1,s]}, \ we \ have \ |\beta_{l,i,j}| \leq (y_{l,i+s} - y_{l,i})^{j-1} \leq (s\bar{\delta}_l)^{j-1}.\\ c. \ For \ all \ i_l \in \overline{[1,s+d]}, l \in \overline{[1,n]}, \ any \ r \in \overline{[0,s-2]} \ and \ any \ r_l \in \mathbb{Z}_+, l \in \overline{[1,n]} \ satisfying \ \sum_{l=1}^n r_l = r,\\ N_{i_1,\ldots,i_n}(.) \ is \ nonnegative \ and \ \partial^r N_{i_1,\ldots,i_n}(p)/(\partial p_1^{r_1} \ldots \partial p_n^{r_n}) = 0 \ for \ all \ p \notin \otimes_{l=1}^n (y_{l,i_l}, y_{l,i_l+s}).\\ d. \ \sum_{i_1=1}^{s+d} \cdots \sum_{i_n=1}^{s+d} |N_{i_1,\ldots,i_n}(p)| = 1 \ for \ all \ p \in \mathcal{P}.\\ e. \ Fix \ any \ r \in \overline{[0,s-2]} \ and \ any \ r_l \in \mathbb{Z}_+, l \in \overline{[1,n]} \ satisfying \ \sum_{l=1}^n r_l = r, \end{array}$

$$\left| \left| \frac{\partial^r N_{i_1,\dots,i_n}(.)}{\partial p_1^{r_1}\dots\partial p_n^{r_n}} \right| \right|_{\infty} \le C_{r,s} \left[\min_{l=1,\dots,n} \{ \bar{\delta}_l \} \right]^{-1}$$

where $C_{r,s} > 0$ is a constant that only depends on r and s.

We first discuss the meaning of the two lemmas above. Since a spline function is essentially a sequence of local polynomial functions attached together, to understand its approximation accuracy, we need to first answer the following question: Suppose that we use a polynomial function g to approximate a deterministic function f on a small region, how does the approximation error depend on the smoothness index of f, the degree of g, and the size of the region? Lemma A1 derives a bound for approximation error as a function of these factors. Lemma A2 summarizes some useful properties of the spline function constructed using the B-Spline approach (see §3.1 for more details). We now proceed to prove Lemma 1.

Let $K' = \max\{K_1, K_2, K_3\}$ where the constants K_1 is defined below and K_2, K_3 are defined later in Step 3 (below (A12)). Let $K = \exp(\log^2 K')$. Since $L_0 \ge \log^3 k$, $d \to \infty$ as $k \to \infty$. So there exists a constant $K_1 \ge 3$ such that for all $k \ge K_1$ and for all $l \in [1, n]$, $2s\bar{\delta}_l \le 1$. This observation allows us to invoke Lemma A1 later. Note that for $k \le K'$, the desired result holds because for any x > 0,

$$\mathbb{P}\left(\left|\left|\frac{\partial^r(\lambda_j^*(p) - \tilde{\lambda}_j(p))}{\partial p_1^{r_1} \dots \partial p_n^{r_n}}\right|\right|_{\infty} \ge x\right) \le 1 = K \exp(-\log^2 K').$$

Hence, in the remaining of the proof, we will focus only on the case when k > K'. We proceed in several steps:

Step 1

Our objective in this step is to compute an upper bound for

$$\left\| \frac{\partial^r (\lambda_j^*(.) - \tilde{\lambda}_j(.))}{\partial p_1^{r_1} \dots \partial p_n^{r_n}} \right\|_{\infty}$$

Fix some $\tilde{i}_l \in \overline{[s,s+d]}$ for all $l \in \overline{[1,n]}$, $j \in \overline{[1,n]}$ and $r \in \overline{[0,(s-2)\wedge \bar{s}]}$. Define two hypercubes $H_{\tilde{i}_1,\ldots,\tilde{i}_n} := \bigotimes_{l=1}^n [y_{l,\tilde{i}_l}, y_{l,\tilde{i}_l+1}]$ and $\tilde{H}_{\tilde{i}_1,\ldots,\tilde{i}_n} := \bigotimes_{l=1}^n [y_{l,\tilde{i}_l-s+1}, y_{l,\tilde{i}_l+s}]$. For any $p \in H_{\tilde{i}_1,\ldots,\tilde{i}_n}$, and any $r_l \in \mathbb{Z}_+, l \in \overline{[1,n]}$ satisfying $\sum_{l=1}^n r_l = r$, we have:

$$\left|\frac{\partial^r (\lambda_j^*(p) - \tilde{\lambda}_j(p))}{\partial p_1^{r_1} \dots \partial p_n^{r_n}}\right| \le \left|\frac{\partial^r (\lambda_j^*(p) - \mathcal{L}\lambda_j^*(p))}{\partial p_1^{r_1} \dots \partial p_n^{r_n}}\right| + \left|\frac{\partial^r (\mathcal{L}\lambda_j^*(p) - \tilde{\lambda}_j(p))}{\partial p_1^{r_1} \dots \partial p_n^{r_n}}\right|.$$
 (A1)

We now bound the terms after the inequality separately.

Bounding the first term in (A1). Let $\mathcal{X} = \tilde{H}_{\tilde{i}_1,...\tilde{i}_n}$. Since $2s\bar{\delta}_l \leq 1$ for $k \geq K' \geq K_1$, by Lemma A1, there exists $g \in \bigotimes_{l=1}^n \mathsf{P}^{(s \wedge \bar{s})-1}$ $([\underline{p}_l, \bar{p}_l])$ such that for all $p \in \tilde{H}_{\tilde{i}_1,...,\tilde{i}_n}$ and $r \in [0, (s-2) \wedge \bar{s}]$,

$$\left|\frac{\partial^r (\lambda_j^*(p) - g(p))}{\partial p_1^{r_1} \dots \partial p_n^{r_n}}\right| \le C_{n,r} W \left[\max_{l=1,\dots,n} \{2s\bar{\delta}_l\}\right]^{s\wedge\bar{s}-r} \le C_{n,r} W (2s)^{s\wedge\bar{s}-r} \left[\max_{l=1,\dots,n} \left\{\frac{\bar{p}_l - \underline{p}_l}{d}\right\}\right]^{s\wedge\bar{s}-r} (A2)$$

where $C_{n,r}$ is a positive constant that only depends on n and r. Note that for all $i_l \in [\tilde{i}_l - s + 1, \tilde{i}_l]$ and for all $r_l \in [\overline{1,s}], l \in [\overline{1,n}]$, we have $(\tau_{1,i_1,r_1}, \ldots, \tau_{n,i_n,r_n}) \in \tilde{H}_{\tilde{i}_1,\ldots,\tilde{i}_n}$. Thus, there exists a constant C_0 independent of k such that for any $i_l \in [\tilde{i}_l - s + 1, \tilde{i}_l], l \in [\overline{1,n}]$, the following holds:

$$\begin{split} &\gamma_{i_{1},\dots,i_{n}}\lambda_{j}^{*}-\gamma_{i_{1},\dots,i_{n}}g|\\ &\leq \sum_{j_{1}=1}^{s}\sum_{r_{1}=1}^{j_{1}}\cdots\sum_{j_{n}=1}^{s}\sum_{r_{n}=1}^{j_{n}}\frac{\left(\prod_{l=1}^{n}\beta_{l,i_{l},j_{l}}\right)|\lambda_{j}^{*}(\tau_{1,i_{1},r_{1}},\dots,\tau_{n,i_{n},r_{n}})-g(\tau_{1,i_{1},r_{1}},\dots,\tau_{n,i_{n},r_{n}})|}{\left|\prod_{l=1}^{n}\prod_{l=1}^{j_{l}}\sum_{s_{l}=1,s_{l}\neq r_{l}}^{j_{l}}(\tau_{l,i_{l},r_{l}}-\tau_{l,i_{l},s_{l}})\right|\\ &\leq \sum_{j_{1}=1}^{s}\sum_{r_{1}=1}^{j_{1}}\cdots\sum_{j_{n}=1}^{s}\sum_{r_{n}=1}^{j_{n}}\frac{\prod_{l=1}^{n}(\bar{\delta}_{l}s)^{j_{l}-1}}{\prod_{l=1}^{n}(\bar{\delta}_{l}/s)^{j_{l}-1}}|\lambda_{j}^{*}(\tau_{1,i_{1},r_{1}},\dots,\tau_{n,i_{n},r_{n}})-g(\tau_{1,i_{1},r_{1}},\dots,\tau_{n,i_{n},r_{n}})|\\ &=\sum_{j_{1}=1}^{s}\sum_{r_{1}=1}^{j_{1}}\cdots\sum_{j_{n}=1}^{s}\sum_{r_{n}=1}^{j_{n}}s^{2(\sum_{l=1}^{n}j_{l}-n)}|\lambda_{j}^{*}(\tau_{1,i_{1},r_{1}},\dots,\tau_{n,i_{n},r_{n}})-g(\tau_{1,i_{1},r_{1}},\dots,\tau_{n,i_{n},r_{n}})|\\ &\leq \left(\frac{s+s^{2}}{2}\right)^{n}s^{2(ns-n)}C_{n,0}W(2s)^{s\wedge\bar{s}}\left[\max_{l=1,\dots,n}\left\{\frac{\bar{p}_{l}-\underline{p}_{l}}{d}\right\}\right]^{s\wedge\bar{s}}\leq\frac{C_{0}}{d^{s\wedge\bar{s}}}, \end{split}$$
(A3)

where the first inequality follows by the definition of γ_{i_1,\ldots,i_n} , the second inequality follows by Lemma A2 part (b), and the third inequality follows by (A2). Then, for any $p \in H_{\tilde{i}_1,\ldots,\tilde{i}_n}$, we have:

$$\left|\frac{\partial^{r}(\mathcal{L}\lambda_{j}^{*}(p)-\mathcal{L}g(p))}{\partial p_{1}^{r_{1}}\dots\partial p_{n}^{r_{n}}}\right| \leq \sum_{i_{1}=1}^{s+d}\cdots\sum_{i_{n}=1}^{s+d}\left|\gamma_{i_{1},\dots,i_{n}}\lambda_{j}^{*}-\gamma_{i_{1},\dots,i_{n}}g\right| \left|\frac{\partial^{r}N_{i_{1},\dots,i_{n}}(p)}{\partial p_{1}^{r_{1}}\dots\partial p_{n}^{r_{n}}}\right|$$
$$=\sum_{i_{1}=\tilde{i}_{1}-s+1}^{\tilde{i}_{1}}\cdots\sum_{i_{n}=\tilde{i}_{n}-s+1}^{\tilde{i}_{n}}\left|\gamma_{i_{1},\dots,i_{n}}\lambda_{j}^{*}-\gamma_{i_{1},\dots,i_{n}}g\right| \left|\frac{\partial^{r}N_{i_{1},\dots,i_{n}}(p)}{\partial p_{1}^{r_{1}}\dots\partial p_{n}^{r_{n}}}\right|$$
$$\leq s^{n}\frac{C_{0}}{d^{s\wedge\bar{s}}}C_{r,s}\left[\min_{l=1,\dots,n}\left\{\bar{\delta}_{l}\right\}\right]^{-r}\leq \frac{s^{n}C_{0}C_{r,s}}{d^{s\wedge\bar{s}}}\left[\min_{l=1,\dots,n}\left\{\frac{\bar{p}_{l}-\underline{p}_{l}}{2d}\right\}\right]^{-r}$$
$$\leq 2^{r}s^{n}C_{0}C_{r,s}\left[\min_{l=1,\dots,n}\left\{\bar{p}_{l}-\underline{p}_{l}\right\}\right]^{-r}d^{r-s\wedge\bar{s}},$$
(A4)

where the equality follows because, by Lemma A2 part (c), $\partial^r N_{i_1,\ldots,i_n}(p)/(\partial p_1^{r_1}\ldots\partial p_n^{r_n}) = 0$ for $p \notin \bigotimes_{l=1}^n (y_{l,i_l}, y_{l,i_l+s})$, the second inequality follows by (A3) and Lemma A2 part (e), and the third inequality follows since $d+1 \leq 2d$.

Putting things together, by Lemma A2 part (a) (note that $s \wedge \bar{s} \leq s$), (A2) and (A4), we have the following inequality for all $p \in H_{\tilde{i}_1,\ldots,\tilde{i}_n}$:

$$\begin{aligned} \left| \frac{\partial^{r} (\lambda_{j}^{*}(p) - \mathcal{L}\lambda_{j}^{*}(p))}{\partial p_{1}^{r_{1}} \dots \partial p_{n}^{r_{n}}} \right| &\leq \left| \frac{\partial^{r} (\lambda_{j}^{*}(p) - g(p))}{\partial p_{1}^{r_{1}} \dots \partial p_{n}^{r_{n}}} \right| + \left| \frac{\partial^{r} (g(p) - \mathcal{L}g(p))}{\partial p_{1}^{r_{1}} \dots \partial p_{n}^{r_{n}}} \right| + \left| \frac{\partial^{r} (\mathcal{L}g(p) - \mathcal{L}\lambda_{j}^{*}(p))}{\partial p_{1}^{r_{1}} \dots \partial p_{n}^{r_{n}}} \right| \\ &\leq \left[C_{n,r} W \left[2s \max_{1 \leq l \leq n} \{\bar{p}_{l} - \underline{p}_{l}\} \right]^{s \wedge \bar{s} - r} + 2^{r} s^{n} C_{0} C_{r,s} \left[\min_{l=1,\dots,n} \left\{ \bar{p}_{l} - \underline{p}_{l} \right\} \right]^{-r} \right] \frac{1}{d^{s \wedge \bar{s} - r}} \\ &\leq \frac{C_{1}}{d^{s \wedge \bar{s} - r}}, \end{aligned}$$
(A5)

for some C_1 independent of k. Since the right hand side of (A5) does not depend on $\tilde{i}_1, \ldots, \tilde{i}_n$, the inequality holds uniformly for all $p \in \mathcal{P}$.

Bounding the second term in (A1). Define $\xi_{i_1,\ldots,i_n}^j := \max_{1 \le r_1,\ldots,r_n \le s} \{ |\lambda_j^*(\tau_{1,i_1,r_1},\ldots,\tau_{n,i_n,r_n}) - \tilde{\lambda}_j(\tau_{1,i_1,r_1},\ldots,\tau_{n,i_n,r_n})| \}$ and $\xi^j := \max_{1 \le i_1,\ldots,i_n \le s+d} \{\xi_{i_1,\ldots,i_n}^j\}$. For any $i_l \in [\tilde{i}_l - s + 1, \tilde{i}_l], l \in [1,n]$, by similar argument as in (A3), we have:

$$\begin{aligned} \left| \gamma_{i_1,\dots,i_n} \lambda_j^* - c_{i_1,\dots,i_n}^j \right| &\leq \sum_{j_1=1}^s \sum_{r_1=1}^{j_1} \cdots \sum_{j_n=1}^s \sum_{r_n=1}^{j_n} \frac{\left(\prod_{l=1}^n \beta_{l,i_l,j_l} \right) |\lambda_j^*(\tau_{1,i_1,r_1},\dots,\tau_{n,i_n,r_n}) - \tilde{\lambda}_j(\tau_{1,i_1,r_1},\dots,\tau_{n,i_n,r_n}) |}{\left| \prod_{l=1}^n \prod_{s_l=1,s_l \neq r_l}^{j_l} (\tau_{l,i_l,r_l} - \tau_{l,i_l,s_l}) \right|} \\ &\leq \sum_{j_1=1}^s \sum_{r_1=1}^{j_1} \cdots \sum_{j_n=1}^s \sum_{r_n=1}^{j_n} s^{2(\sum_{l=1}^n j_l - n)} |\lambda_j^*(\tau_{1,i_1,r_1},\dots,\tau_{n,i_n,r_n}) - \tilde{\lambda}_j(\tau_{1,i_1,r_1},\dots,\tau_{n,i_n,r_n})| \\ &\leq \left(\frac{s+s^2}{2} \right)^n s^{2(ns-n)} \xi_{i_1,\dots,i_n}^j. \end{aligned}$$

So, for all $p \in H_{\tilde{i}_1,\ldots,\tilde{i}_n}$, there exists constant $C_2 > 0$ independent of k and $\tilde{i}_1,\ldots,\tilde{i}_n$ such that

$$\begin{aligned} \left| \frac{\partial^r (\mathcal{L}\lambda_j^*(p) - \tilde{\lambda}_j(p))}{\partial p_1^{r_1} \dots \partial p_n^{r_n}} \right| &\leq \sum_{i_1=1}^{s+d} \dots \sum_{i_n=1}^{s+d} \left| \gamma_{i_1,\dots,i_n} \lambda_j^* - c_{i_1,\dots,i_n}^j \right| \ \left| \frac{\partial^r N_{i_1,\dots,i_n}(p)}{\partial p_1^{r_1} \dots \partial p_n^{r_n}} \right| \\ &= \sum_{i_1=\tilde{i}_1-s+1}^{\tilde{i}_1} \dots \sum_{i_n=\tilde{i}_n-s+1}^{\tilde{i}_n} \left| \gamma_{i_1,\dots,i_n} \lambda_j^* - c_{i_1,\dots,i_n}^j \right| \ \left| \frac{\partial^r N_{i_1,\dots,i_n}(p)}{\partial p_1^{r_1} \dots \partial p_n^{r_n}} \right| \end{aligned}$$

$$\leq s^{n} \left(\frac{s+s^{2}}{2}\right)^{n} s^{2(ns-n)} \xi^{j} C_{r,s} \left[\min_{l=1,...,n} \{\bar{\delta}_{l}\}\right]^{-r}$$

$$\leq \left(\frac{1+s}{2}\right)^{n} s^{2ns} \xi^{j} C_{r,s} \left[\min_{l=1,...,n} \left\{\frac{\bar{p}_{l}-\underline{p}_{l}}{2d}\right\}\right]^{-r}$$

$$= 2^{r-n} C_{r,s} \left(1+s\right)^{n} s^{2ns} \left[\min_{l=1,...,n} \left\{\bar{p}_{l}-\underline{p}_{l}\right\}\right]^{-r} \xi^{j} d^{r} = C_{2} \xi^{j} d^{r}, \qquad (A6)$$

where the second inequality follows by Lemma A2 part (e) and $C_{r,s}$ only depends on r and s, the third inequality follows since $d+1 \leq 2d$.

Note that the right hand side of (A6) does not depend on $\tilde{i}_1, \ldots, \tilde{i}_n$, (A6) holds uniformly for all $p \in \mathcal{P}$. So, by (A1), (A5) and (A6), we conclude that there exists C_3 independent of k such that

$$\left\| \left\| \frac{\partial^{r}(\lambda_{j}^{*}(.) - \tilde{\lambda}_{j}(.))}{\partial p_{1}^{r_{1}} \dots \partial p_{n}^{r_{n}}} \right\|_{\infty} = \sup_{p \in \mathcal{P}} \left| \frac{\partial^{r}(\lambda_{j}^{*}(p) - \tilde{\lambda}_{j}(p))}{\partial p_{1}^{r_{1}} \dots \partial p_{n}^{r_{n}}} \right| \\
\leq \max_{s \leq \tilde{i}_{1}, \dots, \tilde{i}_{n} \leq s+d} \left\{ \sup_{p \in H_{\tilde{i}_{1}, \dots, \tilde{i}_{n}}} \left\{ \left| \frac{\partial^{r}(\lambda_{j}^{*}(p) - \mathcal{L}\lambda_{j}^{*}(p))}{\partial p_{1}^{r_{1}} \dots \partial p_{n}^{r_{n}}} \right| + \left| \frac{\partial^{r}(\mathcal{L}\lambda_{j}^{*}(p) - \tilde{\lambda}_{j}(p))}{\partial p_{1}^{r_{1}} \dots \partial p_{n}^{r_{n}}} \right| \right\} \right\} \\
\leq \left(\frac{C_{1}}{d^{s \wedge \bar{s} - r}} + C_{2}\xi^{j}d^{r} \right) \leq C_{3} \left(\frac{1}{d^{s \wedge \bar{s} - r}} + \xi^{j}d^{r} \right).$$
(A7)

Step 2

We now analyze the term ξ^j . Note that $\xi^j = \max_{p \in \tilde{\mathcal{G}}} |\lambda_j^*(p) - \tilde{\lambda}_j(p)|$ where $\tilde{\mathcal{G}} := \{(\tau_{1,i_1,j_1}; \ldots; \tau_{n,i_n,j_n}): i_l \in \overline{[1,s]}, j_l \in \overline{[1,s]}, \forall l \in \overline{[1,n]}\}$ is as defined in §3.1. So, for all $x \ge 0$, we can bound

$$\mathbb{P}\left(\max_{p\in\tilde{\mathcal{G}}}|\lambda_{j}^{*}(p)-\tilde{\lambda}_{j}(p)|\geq x\right)\leq\mathbb{P}\left(\max_{p\in\tilde{\mathcal{G}}}\{\tilde{\lambda}_{j}(p)-\lambda_{j}^{*}(p)\}\geq x\right)+\mathbb{P}\left(\max_{p\in\tilde{\mathcal{G}}}\{\lambda_{j}^{*}(p)-\tilde{\lambda}_{j}(p)\}\geq x\right)(A8)$$

We now bound the two terms after the inequality separately. For $x \ge 0$ and t > 0, since $|\tilde{\mathcal{G}}| = s^n (s+d)^n$, the following holds:

$$\mathbb{P}\left(\max_{p\in\tilde{\mathcal{G}}}\{\tilde{\lambda}_{j}(p)-\lambda_{j}^{*}(p)\}\geq x\right) = \mathbb{P}\left(t\max_{p\in\tilde{\mathcal{G}}}\{\tilde{\lambda}_{j}(p)-\lambda_{j}^{*}(p)\}\geq tx\right) \\
\leq \exp(-tx)\mathbb{E}\left[\exp\left(t\max_{p\in\tilde{\mathcal{G}}}\{\tilde{\lambda}_{j}(p)-\lambda_{j}^{*}(p)\}\right)\right] \\
\leq \exp(-tx)\sum_{p\in\tilde{\mathcal{G}}}\mathbb{E}\left[\exp\left(t(\tilde{\lambda}_{j}(p)-\lambda_{j}^{*}(p))\right)\right] \\
\leq \exp(-tx)s^{n}(s+d)^{n}\max_{p\in\tilde{\mathcal{G}}}\left\{\mathbb{E}\left[\exp\left(t(\tilde{\lambda}_{j}(p)-\lambda_{j}^{*}(p))\right)\right]\right\}.$$
(A9)

Note that there exists a $p^* \in \tilde{\mathcal{G}}$ such that the expectation $\mathbb{E}\left[\exp\left(t(\tilde{\lambda}_j(p) - \lambda_j^*(p))\right)\right]$ in (A9) attains its maximum. So, for all $0 < t \leq L_0$,

$$\begin{split} \max_{p \in \tilde{\mathcal{G}}} \left\{ \mathbb{E} \left[\exp \left(t(\tilde{\lambda}_j(p) - \lambda_j^*(p)) \right) \right] \right\} &= \mathbb{E} \left[\exp \left(t(\tilde{\lambda}_j(p^*) - \lambda_j^*(p^*)) \right) \right] \\ &= \exp(-t\lambda_j^*(p^*)) \left\{ \mathbb{E} \left[\exp \left(\frac{t}{L_0} \mathsf{Bernoulli}(\lambda_j^*(p^*)) \right) \right] \right\}_{L_0}^{L_0} \\ &= \exp(-t\lambda_j^*(p^*)) \left\{ 1 - \lambda_j^*(p^*) + \lambda_j^*(p^*) \exp \left(\frac{t}{L_0} \right) \right\} \end{split}$$

$$\leq \exp(-t\lambda_{j}^{*}(p^{*}))\left\{\exp\left(\lambda_{j}^{*}(p^{*})\left[\exp\left(\frac{t}{L_{0}}\right)-1\right]\right)\right\}^{L_{0}}$$

$$= \exp(-t\lambda_{j}^{*}(p^{*}))\exp\left(\lambda_{j}^{*}(p^{*})L_{0}\sum_{j=1}^{\infty}\frac{1}{j!}\left(\frac{t}{L_{0}}\right)^{j}\right)$$

$$= \exp\left(\lambda_{j}^{*}(p^{*})L_{0}\sum_{j=2}^{\infty}\frac{1}{j!}\left(\frac{t}{L_{0}}\right)^{j}\right)$$

$$\leq \exp\left(\lambda_{j}^{*}(p^{*})t^{2}/L_{0}\right) \leq \exp(t^{2}/L_{0}), \qquad (A10)$$

where the second equality follows because $\tilde{\lambda}_j(p^*)$ is the average of L_0 independent Bernoulli random variables with success probability $\lambda_j^*(p^*)$ and the last inequality follows from the fact that $\sum_{j=2}^{\infty} (j!)^{-1} (t/L_0)^j \leq (t/L_0)^2 \sum_{j=2}^{\infty} [j(j-1)]^{-1} \leq (t/L_0)^2$. Hence, by (A9) and (A10), for all $0 < t \leq L_0$,

$$\mathbb{P}\left(\max_{p\in\tilde{\mathcal{G}}}\{\tilde{\lambda}_j(p)-\lambda_j^*(p)\}\geq x\right)\leq \exp(t^2/L_0-tx+\log(s^n(s+d)^n)).$$
(A11)

Following similar arguments, for all $0 < t \leq L_0$, there exists some $q^* \in \tilde{\mathcal{G}}$ such that

$$\mathbb{P}\left(\max_{p\in\tilde{\mathcal{G}}}\{\lambda_{j}^{*}(p)-\tilde{\lambda}_{j}(p)\}\geq x\right)\leq\exp(-tx)\left[\max_{p\in\tilde{\mathcal{G}}}\left\{\mathbb{E}\left[\exp\left(t(\lambda_{j}^{*}(p)-\tilde{\lambda}_{j}(p))\right)\right]\right\}\right]s^{n}(s+d)^{n}\\\leq\exp(-tx)\left[\exp(t\lambda_{j}^{*}(q^{*}))\exp\left(\lambda_{j}^{*}(q^{*})L_{0}\sum_{j=1}^{\infty}\frac{(-1)^{j}}{j!}\left(\frac{t}{L_{0}}\right)^{j}\right)\right]s^{n}(s+d)^{n}\\\leq\exp(\lambda_{j}^{*}(q^{*})t^{2}/L_{0}-tx)s^{n}(s+d)^{n}\\\leq\exp(t^{2}/L_{0}-tx+\log(s^{n}(s+d)^{n})).$$
(A12)

Pick $x = 4L_0^{-1/2}(s+d)^{n/2}s^{n/2}\log k$ and $t = L_0x/2$. We now show that under this choice of x and t, the inequalities (A11) and (A12) hold for large k, i.e., $t \le L_0$ when k is large. Recall that we have set $d = \lceil (L_0^{1/2}/\log k)^{1/(s+n/2)} \rceil$. Since $L_0 \ge \log^3 k$, for $k \ge 3$, we have $L_0^{1/2}/\log k \ge 1$. This implies that

$$(L_0^{1/2}/\log k)^{1/(s+n/2)} \le d \le 2(L_0^{1/2}/\log k)^{1/(s+n/2)}.$$
(A13)

We then have that for all $k \ge 3$, the following holds

$$x = \frac{4\log k}{\sqrt{L_0}} (s+d)^{\frac{n}{2}} s^{\frac{n}{2}} \le \frac{4\log k}{\sqrt{L_0}} (s+1)^{\frac{n}{2}} s^{\frac{n}{2}} d^{\frac{n}{2}} \le 4(s+1)^{\frac{n}{2}} s^{\frac{n}{2}} 2^{\frac{n}{2}} \left(\frac{\log k}{\sqrt{L_0}}\right)^{\frac{s}{s+n/2}} \to 0, \text{ as } k \to \infty.$$

Hence, there exists a constant $K_2 \ge 3$ such that for all $k \ge K_2 \ge 3$, we have $x \le 2$ and hence $t = L_0 x/2 \le L_0$. The following inequality holds for $k \ge K_2$:

$$\mathbb{P}\left(\xi^{j} \ge 4L_{0}^{-1/2}d^{\frac{n}{2}}(s+s^{2})^{\frac{n}{2}}\log k\right) \le \mathbb{P}\left(\xi^{j} \ge 4L_{0}^{-1/2}s^{\frac{n}{2}}(s+d)^{\frac{n}{2}}\log k\right) \\
= \mathbb{P}\left(\max_{p\in\tilde{\mathcal{G}}}|\tilde{\lambda}_{j}(p)-\lambda_{j}^{*}(p)|\ge 4L_{0}^{-1/2}s^{\frac{n}{2}}(s+d)^{\frac{n}{2}}\log k\right) \\
\le 2\exp\left(-\frac{x^{2}L_{0}}{4}\right)s^{n}(s+d)^{n} = 2s^{n}(s+d)^{n}\exp\left(-4s^{n}(s+d)^{n}\log^{2}k\right) \\
\le 2s^{n}(s+d)^{n}\exp\left(-2s^{n}(s+d)^{n}\right)\exp\left(-\log^{2}k\right) \le K_{3}\exp\left(-\log^{2}k\right), \quad (A14)$$

where $K_3 = \sup_{x\geq 0} \{x \exp(-x)\}$ is a positive constant, the first inequality follows since $s + d \leq (s+1)d$ for $d \geq 1$, the second inequality follows by (A8), (A11) and (A12). Let $\Psi_r = C_3(1 + 2^{n/2+r+2}(s+s^2)^{n/2})$ which is independent of k (C_3 is defined in (A7)), then

$$\Psi_r \left(\frac{\log k}{\sqrt{L_0}}\right)^{\frac{s\wedge\bar{s}-r}{s+n/2}} \ge C_3 \left[\left(\frac{\log k}{\sqrt{L_0}}\right)^{\frac{s\wedge\bar{s}-r}{s+n/2}} + 2^{\frac{n}{2}+r+2}(s+s^2)^{\frac{n}{2}} \left(\frac{\log k}{\sqrt{L_0}}\right)^{\frac{s-r}{s+n/2}} \right] \\ \ge C_3 \left(\frac{1}{d^{s\wedge\bar{s}-r}} + \frac{4\log k}{\sqrt{L_0}}(s+s^2)^{\frac{n}{2}}d^{\frac{n}{2}+r}\right),$$
(A15)

where the first inequality follows since $\log k/L_0^{1/2} \le 1$ and the second inequality follows by (A13). So,

$$\begin{aligned} & \mathbb{P}\left(\left|\left|\frac{\partial^{r}(\lambda_{j}^{*}(p)-\tilde{\lambda}_{j}(p))}{\partial p_{1}^{r_{1}}\dots\partial p_{n}^{r_{n}}}\right|\right|_{\infty} \geq \Psi_{r}\left(\frac{\log k}{\sqrt{L_{0}}}\right)^{\frac{s\wedge s-r}{s+n/2}}\right) \\ & \leq \mathbb{P}\left(\left|\left|\frac{\partial^{r}(\lambda_{j}^{*}(p)-\tilde{\lambda}_{j}(p))}{\partial p_{1}^{r_{1}}\dots\partial p_{n}^{r_{n}}}\right|\right|_{\infty} \geq C_{3}\left(\frac{1}{d^{s\wedge\bar{s}-r}}+\frac{4\log k}{\sqrt{L_{0}}}(s+s^{2})^{\frac{n}{2}}d^{\frac{n}{2}+r}\right)\right) \\ & \leq \mathbb{P}\left(C_{3}\left(\frac{1}{d^{s\wedge\bar{s}-r}}+\xi^{j}d^{r}\right) \geq C_{3}\left(\frac{1}{d^{s\wedge\bar{s}-r}}+\frac{4\log k}{\sqrt{L_{0}}}(s+s^{2})^{\frac{n}{2}}d^{\frac{n}{2}+r}\right)\right) \\ & \leq \mathbb{P}\left(\xi^{j}\geq\frac{4\log k}{\sqrt{L_{0}}}d^{\frac{n}{2}}(s+s^{2})^{\frac{n}{2}}\right) \leq K_{3}\exp(-\log^{2}k) \leq K\exp(-\log^{2}k), \end{aligned}$$

where the first inequality follows by (A15), the second inequality follows by (A7) and the fourth inequality follows by (A14). This completes the proof. \Box

Proof of Lemma A1: For $k \in \mathbb{Z}_{++}$, define $\mathcal{I}^k := \{a = (a_1; \ldots; a_n) : a_l \in \overline{[0,k]}, \text{ for all } l \in \overline{[1,n]}, and \sum_{l=1}^n a_l = k\}$. Define $g(x) = \sum_{k=0}^{(s \wedge \bar{s})-1} \sum_{a \in \mathcal{I}^k} h_f(x, \mathbf{0}, a)$ for all $x \in \mathcal{X}$ where

$$h_f(x,y,a) := \frac{\partial^{a_1 + \dots + a_n} f(y)}{\partial x_1^{a_1} \dots \partial x_n^{a_n}} \prod_{l=1}^n \frac{(x_l - y_l)^{a_l}}{a_l!}, \ \forall y \in \mathcal{X}.$$

It is straightforward to verify that $g(.) \in \bigotimes_{l=1}^{n} \mathsf{P}^{s \wedge \overline{s}-1}([0, \overline{x}_{l}])$. For some $k \in \overline{[0, s \wedge \overline{s}]}$, consider any $a = (a_{1}; \ldots; a_{n}) \in \mathcal{I}^{k}$, any $x, y \in \mathcal{X}$, any $r \in \overline{[0, k]}$, and any $r_{l} \in \mathbb{Z}_{+}, l \in \overline{[1, n]}$ satisfying $\sum_{l=1}^{n} r_{l} = r$. If $a_{l} < r_{l}$ for some $l \in \overline{[1, n]}, \partial^{r} h_{f}(x, y, a) / (\partial x_{1}^{r_{1}} \dots \partial x_{n}^{r_{n}}) = 0$. Hence, the following hold:

$$\left|\frac{\partial^{r} h_{f}(x,y,a)}{\partial x_{1}^{r_{1}} \dots \partial x_{n}^{r_{n}}}\right| = \begin{cases} 0, & \text{if } a_{l} < r_{l}, \text{ for some } l \in [1,n]; \\ \left|\frac{\partial^{a_{1}+\dots+a_{n}} f(y)}{\partial x_{1}^{a_{1}} \dots \partial x_{n}^{a_{n}}} \prod_{l=1}^{n} \frac{(x_{l}-y_{l})^{a_{l}-r_{l}}}{(a_{l}-r_{l})!}\right| \leq \frac{W}{\prod_{l=1}^{n} (a_{l}-r_{l})!} \left[\max_{l=1,\dots,n} \{\bar{x}_{l}\}\right]^{k-r}, \text{ otherwise,} \end{cases}$$
(A16)

where the inequality follows by N2, $\bar{x}_l \leq 1$ and $a_l - r_l \geq 0$ for all $l \in \overline{[1,n]}$, and $\sum_{l=1}^n (a_l - r_l) = k - r$. So, for any $r \in \overline{[0, s \land \bar{s}]}$ and for any $r_l \in \mathbb{Z}_+, l \in \overline{[1,n]}$ satisfying $\sum_{l=1}^n r_l = r$,

$$\begin{split} \left| \left| \frac{\partial^r (f-g)(.)}{\partial x_1^{r_1} \dots \partial x_n^{r_n}} \right| \right|_{\infty} &= \sup_{x \in \mathcal{X}} \left| \frac{\partial^r (f-g)(x)}{\partial x_1^{r_1} \dots \partial x_n^{r_n}} \right| \leq \sum_{a \in \mathcal{I}^{s \wedge \bar{s}}} \sup_{x,y \in \mathcal{X}} \left| \frac{\partial^r h_f(x,y,a)}{\partial x_1^{r_1} \dots \partial x_n^{r_n}} \right| \\ &\leq W \left[\max_{l=1,\dots,n} \{\bar{x}_l\} \right]^{s \wedge \bar{s} - r} \sum_{\substack{\sum_{l=1}^n a_l = s \wedge \bar{s} \\ a_l \geq r_l, \forall l \in [\bar{1}, n]}} \frac{1}{\prod_{l=1}^n (a_l - r_l)!} \\ &= W \left[\max_{l=1,\dots,n} \{\bar{x}_l\} \right]^{s \wedge \bar{s} - r} \frac{1}{(s \wedge \bar{s} - r)!} \sum_{\substack{\sum_{l=1}^n w_l = s \wedge \bar{s} - r \\ w_l \in \mathbb{Z}_+, \forall l \in [\bar{1}, n]}} \frac{(s \wedge \bar{s} - r)!}{\prod_{l=1}^n w_l!} \\ &= W \left[\max_{l=1,\dots,n} \{\bar{x}_l\} \right]^{s \wedge \bar{s} - r} \frac{n^{s \wedge \bar{s} - r}}{(s \wedge \bar{s} - r)!} \end{split}$$

where the first inequality follows by the Lagrangian remainder formula, the second inequality follows by (A16), and the last equality follows by the multinomial theorem. The result follows by letting $C_{n,r} = \sup_{k \ge r} n^{k-r}/(k-r)! \le n^n/n! < \infty$. \Box

Proof of Lemma A2: Recall that $N_{l,i_l}^s(.)$ are defined in §3.1 as the building blocks of the tensorproduct B-Spline basis functions. Let $D^{\sigma}, D^{\sigma}_{+}, D^{\sigma}_{-}$ respectively denote the σ^{th} order derivative, right derivative, left derivative of a single variate real function. We first state some known results of spline functions that will be used to prove Lemma A2.

THEOREM A1. (THEOREM 6.18 IN SCHUMAKER [1]) For any $l \in \overline{[1,n]}$, we have $\mathcal{L}_l f = f$ for all $f \in \mathsf{P}^{s-1}([p_l, \bar{p}_l])$.

THEOREM A2. (LEMMA 6.19 IN SCHUMAKER [1]) For all $l \in [\overline{[1,n]}, i \in [\overline{[1,s+d]}, j \in [\overline{[1,s]}, we have |\beta_{l,i,j}| \leq (y_{l,i+s} - y_{l,i})^{j-1} \leq (s \bar{\delta}_l)^{j-1}$.

THEOREM A3. (THEOREM 4.17 IN SCHUMAKER [1]) Let s > 1. Fix $l \in \overline{[1,n]}$ and $i_l \in \overline{[1,s+d]}$. Suppose $y_{l,i_l} < y_{l,i_{l+s}}$. Then $N_{l,i_l}^s(p_l) > 0$ when $p_l \in (y_{l,i_l}, y_{l,i_{l+s}})$, and $N_{l,i_l}^s(p_l) = 0$ when $p_l \notin [y_{l,i_l}, y_{l,i_{l+s}}]$. At the end points of $(y_{l,i_l}, y_{l,i_{l+s}})$, we have

$$(-1)^{k+s-\mu_{l,i_l}} D^k_+ N^s_{l,i_l}(y_{l,i_l}) = 0 \qquad k = 0, 1, \dots, s-1-\mu_{l,i_l} (-1)^{s-\nu_{l,i_l+s}} D^k_- N^s_{l,i_l}(y_{l,i_l+s}) = 0 \qquad k = 0, 1, \dots, s-1-\nu_{l,i_l+s}$$

where $\mu_{l,i_l} = \max\{j: y_{l,i_l} =, \dots, = y_{l,i_l+j-1}\}$ and $\nu_{l,i_l+s} = \max\{j: y_{l,i_l+s} =, \dots, = y_{l,i_l+s-j+1}\}$.

THEOREM A4. (THEOREM 4.20 IN SCHUMAKER [1]) Fix $l \in \overline{[1,n]}$ and $i_l \in \overline{[s,s+d]}$. For all $p_l \in [y_{l,i_l}, y_{l,i_l+1}), \sum_{v_l=i_l+1-s}^{i_l} N_{l,v_l}^s(p_l) = 1.$

THEOREM A5. (THEOREM 4.22 IN SCHUMAKER [1]) Fix $l \in \overline{[1,n]}$. Suppose that k and p_l are such that $y_{l,k} \leq p_l < y_{l,k+1}$, and define $\delta_{l,i_l,k,j} = \min\{(y_{l,v+j} - y_{l,v}) : y_{l,i_l} \leq y_{l,v} \leq y_{l,k} < y_{l,k+1} \leq y_{l,v+j} \leq y_{l,i_l+s}\}$, for $j \in \overline{[1,s]}$. Suppose $\sigma > 0$ and $\delta_{l,i_l,k,s-\sigma+1} > 0$. Then $|D^{\sigma}_+ N^s_{l,i_l}(p_l)| \leq \Gamma_{s,\sigma}/(\prod_{q=1}^{\sigma} \delta_{l,i_l,k,s-q})$ where $\Gamma_{s,\sigma} = \frac{(s-1)!}{(s-\sigma-1)!} \begin{pmatrix} \sigma \\ \lfloor \sigma/2 \rfloor \end{pmatrix} \leq 2^{\sigma} \frac{(s-1)!}{(s-\sigma-1)!}$.

We now proceed to prove each part in Lemma A2 one by one.

Proof of part (a)

Note that $\mathcal{L} = \mathcal{L}_1 \circ \mathcal{L}_2 \circ \cdots \circ \mathcal{L}_n$. For any $f \in \bigotimes_{l=1}^n \mathsf{P}^{s-1}[\underline{p}_l, \overline{p}_l]$, we can apply Theorem A1 iteratively n times to obtain $\mathcal{L}f = \mathcal{L}_1 \circ \cdots \circ \mathcal{L}_n f = \mathcal{L}_1 \circ \cdots \circ \mathcal{L}_{n-1} f = \cdots = f$, where $\mathcal{L}_l f$ is understood as applying \mathcal{L}_l to f which is viewed a single variate function of p_l .

Proof of part (b) This follows directly from Theorem A2.

Proof of part (c)

By our definition of $\{y_{l,i}\}_{l=1,i=1}^{n,2s+d}$, $\mu_{l,i_l} = \nu_{l,i_l+s} = 1$ for all $l \in \overline{[1,n]}$ and $i_l \in \overline{[1,s+d]}$. Hence, by Theorem A3, for any $l \in \overline{[1,n]}$, $i_l \in \overline{[1,s+d]}$ and $r_l \in \overline{[0,s-2]}$, $N_{l,i_l}^s(.)$ is nonnegative and $D^{r_l}N_{l,i_l}^s(p_l) = 0$

for all $p_l \notin (y_{l,i_l}, y_{l,i_l+s})$. Hence N_{i_1,\dots,i_n} is nonnegative and, for any $r \in \overline{[0, s-2]}$ and any $r_l \in \mathbb{Z}_+, l \in \overline{[1,n]}$ satisfying $\sum_{l=1}^n r_l = r$,

$$\frac{\partial^r N_{i_1,\dots,i_n}(p)}{\partial p_1^{r_1}\dots \partial p_n^{r_n}} = \prod_{l=1}^n D^{r_l} N_{l,i_l}^s(p_l) = 0$$

for all $p \notin \bigotimes_{l=1}^{n} (y_{l,i_l}, y_{l,i_l+s})$, where the second equality follows since $r_l \leq r \leq s-2$.

Proof of part (d)

Let $H_{i_1,\dots,i_n} = \bigotimes_{l=1}^n [y_{l,i_l}, y_{l,i_l+1}]$ for any $i_l \in \overline{[s,s+d]}, l \in \overline{[1,n]}$. By Theorem A4 and the fact that $\{N_{i_1,\dots,i_n}(.)\}_{i_1=1,\dots,i_n=1}^{s+d}$ are all continuous functions (because $s \ge 2$), we have $\sum_{v_1=i_1+1-s}^{i_1} \cdots \sum_{v_n=i_n+1-s}^{i_n} N_{v_1,\dots,v_n}(p) = 1$ for $p \in H_{i_1,\dots,i_n}$. Moreover, by Lemma A2 part (c), $N_{i_1,\dots,i_n}(.)$ is nonnegative and $N_{i_1,\dots,i_n}(p) = 0$ for $p = (p_1;\dots;p_n) \notin \bigotimes_{l=1}^n (y_{l,i_l}, y_{l,i_l+s})$. Fix some $\tilde{i}_l \in \overline{[s,s+d]}, l \in \overline{[1,n]}$. For all $p \in H_{\tilde{i}_1,\dots,\tilde{i}_n}, \sum_{i_1=1}^{s+d} \cdots \sum_{i_n=1}^{s+d} |N_{i_1,\dots,i_n}(p)| = \sum_{i_1=1}^{s+d} \cdots \sum_{i_n=1}^{s+d} N_{i_1,\dots,i_n}(p) = \sum_{i_1=1}^{\tilde{i}_1+1-s} \cdots \sum_{i_n=\tilde{i}_n+1-s} N_{v_1,\dots,v_n}(p) = 1$. The result follows since the equality holds for all $\tilde{i}_l \in \overline{[s,s+d]}, l \in \overline{[1,n]}, and \mathcal{P} = \bigcup_{\tilde{i}_1=s,\dots,\tilde{i}_n=s}^{s+d} H_{\tilde{i}_1,\dots,\tilde{i}_n}$.

Proof of part (e)

Fix $r \in \overline{[0, s-2]}$, and consider any $r_l \in \mathbb{Z}_+, l \in \overline{[1, n]}$ satisfying $\sum_{l=1}^n r_l = r$. Since $N_{l, i_l}^s(.) \in \mathbb{C}^{s-2}([\underline{p}_l, \bar{p}_l])$ and $r_l \leq r \leq s-2$, $D_+^{r_l}N_{l, i_l}^s(p_l) = D_-^{r_l}N_{l, i_l}^s(p_l) = D^{r_l}N_{l, i_l}^s(p_l)$ for all $p_l \in [\underline{p}_l, \bar{p}_l]$. Fix some $i_l \in \overline{[s, s+d]}$ for all $l \in \overline{[1, n]}$. Suppose that $p_l \in [y_{l, i_l}, y_{l, i_l+1})$. Then, if $r_l = 0$, $|D^{r_l}N_{l, i_l}^s(p_l)| = |N_{l, i_l}^s(p_l)| \leq 1 = 2^0 \frac{(s-1)!}{(s-0-1)!} \overline{\delta}_l^0$ where the inequality follows by Lemma A2 part (d). Otherwise, $r_l \geq 1$, and $s - r_l \geq s - r \geq 2 > 1$. So $\delta_{l, i_l, k, s-q} \geq \overline{\delta}_l > 0$ for all $q = 1, \ldots, r_l$ (recall that $\delta_{l, i_l, k, j}$ is as defined in Theorem A5). Then, by Theorem A5, $|D^{r_l}N_{l, i_l}^s(p_l)| \leq 2^{r_l} \frac{(s-1)!}{(s-r_l-1)!} \overline{\delta}_l^{-r_l}$.

Now, for any $p = (p_1; \ldots; p_n) \in \mathcal{P}$, if $p_l \in (y_{l,i_l}, y_{l,i_l+s})$ for all $l \in \overline{[1,n]}$, the following holds,

$$\left|\frac{\partial^r N_{i_1,\dots,i_n}(p)}{\partial p_1^{r_1}\dots\partial p_n^{r_n}}\right| = \prod_{l=1}^n \left|D^{r_l} N_{l,i_l}^s(p_l)\right| \le \prod_{l=1}^n 2^{r_l} \frac{(s-1)!}{(s-r_l-1)!} \bar{\delta}_l^{-r_l} \le 2^r \left[\frac{(s-1)!}{(s-r-1)!}\right]^n \left[\min_{l=1,\dots,n} \{\bar{\delta}_l\}\right]^{-r}.$$

Otherwise, there exists some l_0 such that $p_{l_0} \notin (y_{l_0,i_{l_0}}, y_{l_0,i_{l_0}+s})$. By Lemma A2 part (c),

$$\frac{\partial^r N_{i_1,\dots,i_n}(p)}{\partial p_1^{r_1}\dots\partial p_n^{r_n}}\bigg| = 0 \le 2^r \left[\frac{(s-1)!}{(s-r-1)!}\right]^n \left[\min_{l=1,\dots,n} \{\bar{\delta}_l\}\right]^{-r}$$

So the result follows by letting $C_{r,s} = 2^r [(s-1)!/(s-r-1)!]^n$. \Box

Appendix B: Proof of Proposition 1

We first show the feasibility of $\tilde{\mathbf{NP}}(\delta)$. Define $\tilde{h}(.) = \tilde{f}(g^{-1}(.)) : \mathcal{Y} \to \mathbb{R}^n$ and $h(.) = f(g^{-1}(.)) : \mathcal{Y} \to \mathbb{R}^n$. By condition (ii), h(.) is strongly concave. Also, define $\tilde{\delta}(y) := \tilde{g}(g^{-1}(y)) - y$. Consider the following two optimization problems:

$$(\mathbf{NP}_y) \quad \max_{y \in \mathcal{Y}} \{h(y) \colon Uy \preceq V\} \quad \text{and} \quad (\tilde{\mathbf{NP}}_y(\delta)) \quad \max_{y \in \mathcal{Y}} \left\{\tilde{h}(y) \colon Uy + U\tilde{\delta}(y) \preceq V - \delta\right\}.$$

Note that \mathbf{NP}_y is equivalent to \mathbf{NP} and $\tilde{\mathbf{NP}}_y(\delta)$ is equivalent to $\tilde{\mathbf{NP}}(\delta)$. Thus, $y^* := g(x^*)$ is the optimal solution to \mathbf{NP}_y and $Uy^* \leq V$. Since $g^{-1}(.)$ is continuous by condition (i) and x^* is in the interior of \mathcal{X} by condition (iv), y^* is in the interior of \mathcal{Y} and there exists a constant $\bar{\phi} > 0$ such that $y^* - \bar{\phi} \mathbf{e} \subseteq \mathcal{Y}$. Let $\bar{\delta} = \min_i \{\bar{\phi}(U\mathbf{e})_i\}$ (note that since U does not have zero rows and all its components are non-negative, $\bar{\delta} > 0$). We claim that if $||Ug(.) - U\tilde{g}(.) + \delta||_{\infty} \leq \bar{\delta}$, then $y^* - \bar{\phi}\mathbf{e}$ is a feasible solution of $\tilde{\mathbf{NP}}_{y}(\delta)$. To see this, simply note that

$$\begin{aligned} ||U\tilde{\delta}(.) + \delta||_{\infty} &= \sup_{y \in \mathcal{Y}} ||U\tilde{g}(g^{-1}(y)) - Uy + \delta||_{\infty} = \sup_{y \in \mathcal{Y}} ||U\tilde{g}(g^{-1}(y)) - Ug(g^{-1}(y)) + \delta||_{\infty} \\ &= \sup_{x \in \mathcal{X}} ||U\tilde{g}(x) - Ug(x) + \delta||_{\infty} = ||Ug(.) - U\tilde{g}(.) + \delta||_{\infty}. \end{aligned}$$

So, $U(y^* - \bar{\phi}\mathbf{e}) + U\tilde{\delta}(y^* - \bar{\phi}\mathbf{e}) + \delta \preceq Uy^* - \bar{\phi}U\mathbf{e} + ||U\tilde{\delta}(.) + \delta||_{\infty}\mathbf{e} \preceq Uy^* + \bar{\delta}\mathbf{e} - \bar{\phi}U\mathbf{e} \preceq V$, where the last inequality follows by the definition of $\bar{\delta}$ and the fact that y^* is feasible to \mathbf{NP}_y . This proves that $\tilde{\mathbf{NP}}_y(\delta)$ is feasible. Thus, $\tilde{\mathbf{NP}}(\delta)$ is feasible. Since the feasible region of $\tilde{\mathbf{NP}}_y(\delta)$ is compact and $\tilde{h}(.)$ is continuous, $\tilde{\mathbf{NP}}_y(\delta)$ has an optimal solution. Let \tilde{y}_{δ} denote an optimal solution of $\tilde{\mathbf{NP}}_y(\delta)$ (note that \tilde{y}_{δ} may not be unique).

We now proceed to derive a bound of $||y^* - \tilde{y}_{\delta}||_2$, which will be used later to obtain the desired bound for $||x^* - \tilde{x}_{\delta}||_2$. To bound $||y^* - \tilde{y}_{\delta}||_2$, we will use the optimal solution of an auxiliary optimization problem below:

$$(\tilde{\mathbf{NP}}_{y}^{ax}(\delta)) = \max_{y \in \mathcal{Y}} \left\{ h(y) : Uy + U\tilde{\delta}(y) \leq V - \delta \right\}.$$

The above problem has the same feasible region as $\tilde{\mathbf{NP}}_{y}(\delta)$, so it is feasible. Let y_{δ}^{ax} denote an optimal solution of $\tilde{\mathbf{NP}}_{y}^{ax}(\delta)$. Since $||y^* - \tilde{y}_{\delta}||_2 \leq ||y^* - y_{\delta}^{ax}||_2 + ||y_{\delta}^{ax} - \tilde{y}_{\delta}||_2$, to bound $||y^* - \tilde{y}_{\delta}||_2$, we only need to bound $||y^* - y_{\delta}^{ax}||_2$ and $||y_{\delta}^{ax} - \tilde{y}_{\delta}||_2$. To derive an upper bound of $||y^* - y_{\delta}^{ax}||_2$, we need to use the following lemma (the proof is given later).

LEMMA A3. Consider the family of perturbed optimization problems below:

$$(\mathbf{NP}_y(\epsilon)) \quad \max_{y \in \mathcal{Y}} \{h(y): Uy \leq V + \epsilon\}.$$

Suppose that h(.) is strongly concave and twice continuously differentiable, \mathcal{Y} is a convex compact set, U is a non-negative matrix and has full row rank, and the optimal solution of $\mathbf{NP}_y(\mathbf{0})$ lies in the interior of \mathcal{Y} . If $y^*(\epsilon)$ is an optimal solution for $\mathbf{NP}_y(\epsilon)$, then $||y^*(\mathbf{0}) - y^*(\epsilon)||_2 \leq K||\epsilon||_{\infty}$ for some K > 0 independent of ϵ .

Note that the assumptions of Lemma A3 hold (i.e., $h(.) = f(g^{-1}(.))$ is twice continuously differentiable because f(.) and $g^{-1}(.)$ are both twice continuously differentiable. Also, as shown earlier, the optimal solution of $\mathbf{NP}_y(\mathbf{0})$, $y^*(\mathbf{0}) = y^*$, is in the interior of \mathcal{Y}). By the strong concavity of h(.), y^{ax}_{δ} is the *unique* optimal solution of $\mathbf{NP}_y(-\delta - U\tilde{\delta}(y^{ax}_{\delta}))$. Thus, by Lemma A3, there exists a constant $K_1 > 0$ independent of $\tilde{f}, \tilde{g}, \delta$ such that

$$||y^* - y^{ax}_{\delta}||_2 \le K_1 ||U\tilde{\delta}(y^{ax}_{\delta}) + \delta||_{\infty} \le K_1 (||U||_{\infty} ||g(.) - \tilde{g}(.)||_{\infty} + ||\delta||_{\infty}).$$
(A17)

We now derive a bound for $||\tilde{y}_{\delta} - y_{\delta}^{ax}||_2$. Since $\tilde{\mathbf{NP}}_y^{ax}(\delta)$ and $\tilde{\mathbf{NP}}_y(\delta)$ have the same constraints, \tilde{y}_{δ} is feasible for $\tilde{\mathbf{NP}}_y^{ax}(\delta)$ and y_{δ}^{ax} is feasible for $\tilde{\mathbf{NP}}_y(\delta)$. By the strong concavity of h(.), there exists a constant v > 0 depending only on h(.) such that

$$h(\tilde{y}_{\delta}) \leq h(y_{\delta}^{ax}) + \nabla h(y_{\delta}^{ax}) \cdot (\tilde{y}_{\delta} - y_{\delta}^{ax}) - \frac{v}{2} ||\tilde{y}_{\delta} - y_{\delta}^{ax}||_{2}^{2} \leq h(y_{\delta}^{ax}) - \frac{v}{2} ||\tilde{y}_{\delta} - y_{\delta}^{ax}||_{2}^{2},$$
(A18)

where the last inequality follows because $\nabla h(y_{\delta}^{ax}) \cdot (\tilde{y}_{\delta} - y_{\delta}^{ax}) \leq 0$ (otherwise, y_{δ}^{ax} cannot be the optimal solution of $\tilde{\mathbf{NP}}_{y}^{ax}(\delta)$). Note also that $\tilde{h}(y_{\delta}^{ax}) \leq \tilde{h}(\tilde{y}_{\delta})$. Combining this with (A18), by Mean Value Theorem, we have

$$\frac{v}{2} ||\tilde{y}_{\delta} - y_{\delta}^{ax}||_{2}^{2} \leq [h(y_{\delta}^{ax}) - \tilde{h}(y_{\delta}^{ax})] - [h(\tilde{y}_{\delta}) - \tilde{h}(\tilde{y}_{\delta})] \leq (\nabla h(\xi) - \nabla \tilde{h}(\xi))'(y_{\delta}^{ax} - \tilde{y}_{\delta}) \\ \leq ||(\nabla h(.) - \nabla \tilde{h}(.))'||_{\infty} ||\tilde{y}_{\delta} - y_{\delta}^{ax}||_{\infty} \leq ||(\nabla h(.) - \nabla \tilde{h}(.))'||_{\infty} ||\tilde{y}_{\delta} - y_{\delta}^{ax}||_{2},$$

for some $\xi \in \mathcal{Y}$. This means that $||\tilde{y}_{\delta} - y_{\delta}^{ax}||_2 \leq \frac{2}{v}||(\nabla h(.) - \nabla \tilde{h}(.))'||_{\infty}$. Combining this with (A17),

$$\begin{aligned} ||y^* - \tilde{y}_{\delta}||_2 &\leq ||y^* - y_{\delta}^{ax}||_2 + ||y_{\delta}^{ax} - \tilde{y}_{\delta}||_2 \\ &\leq K_1 ||U||_{\infty} ||g(.) - \tilde{g}(.)||_{\infty} + K_1 ||\delta||_{\infty} + \frac{2}{v} ||(\nabla h(.) - \nabla \tilde{h}(.))'||_{\infty} \\ &\leq K_1 ||U||_{\infty} ||g(.) - \tilde{g}(.)||_{\infty} + K_1 ||\delta||_{\infty} + \frac{2}{v} ||(\nabla g^{-1}(.))'||_{\infty} ||(\nabla f(.) - \nabla \tilde{f}(.))'||_{\infty} (A19) \end{aligned}$$

where the last inequality holds because $\nabla h(y) - \nabla \tilde{h}(y) = \nabla g^{-1}(y) [\nabla f(g^{-1}(y)) - \nabla \tilde{f}(g^{-1}(y))]$ for all $y \in \mathcal{Y}$. This means that the following inequality also holds:

$$\begin{aligned} ||x^* - \tilde{x}_{\delta}||_{2} \\ \leq \sqrt{n} ||g^{-1}(y^*) - g^{-1}(\tilde{y}_{\delta})||_{\infty} \leq \sqrt{n} ||(\nabla g^{-1}(.))'||_{\infty} ||y^* - \tilde{y}_{\delta}||_{\infty} \leq \sqrt{n} ||(\nabla g^{-1}(.))'||_{\infty} ||y^* - \tilde{y}_{\delta}||_{2} \\ \leq \sqrt{n} ||(\nabla g^{-1}(.))'||_{\infty} \left(K_{1} ||U||_{\infty} ||g(.) - \tilde{g}(.)||_{\infty} + K_{1} ||\delta||_{\infty} + \frac{2}{v} ||(\nabla g^{-1}(.))'||_{\infty} ||(\nabla f(.) - \nabla \tilde{f}(.))'||_{\infty} \right) \\ \leq K(||(\nabla f(.) - \nabla \tilde{f}(.))'||_{\infty} + ||g(.) - \tilde{g}(.)||_{\infty} + ||\delta||_{\infty}), \end{aligned}$$

where $K = \sqrt{n} ||(\nabla g^{-1}(.))'||_{\infty} (K_1 ||U||_{\infty} + K_1 + \frac{2}{v} ||(\nabla g^{-1}(.))'||_{\infty})$. This completes the proof. \Box

Proof of Lemma A3. We claim that there exists $\bar{\epsilon} := \min\{\bar{\epsilon}_1, \bar{\epsilon}_2\} > 0$, where $\bar{\epsilon}_1, \bar{\epsilon}_2$ are strictly positive constants to be defined later, such that, for all $||\epsilon||_{\infty} \leq \bar{\epsilon}, ||y^*(\mathbf{0}) - y^*(\epsilon)||_2 \leq K_1 ||\epsilon||_{\infty}$ for some $K_1 > 0$ independent of ϵ . Note that, if this claim is true, Lemma A3 can be proven as follows. Define $l := \sup_{y_1, y_2 \in \mathcal{Y}} ||y_1 - y_2||_2$ and let $K_2 = l/\bar{\epsilon}$. Then, for all ϵ with $||\epsilon||_{\infty} > \bar{\epsilon}, ||y^*(\epsilon) - y^*(\mathbf{0})||_2 \leq l = K_2 \bar{\epsilon} \leq K_2 ||\epsilon||_{\infty}$. So, Lemma A3 follows by letting $K = \max\{K_1, K_2\}$. We now prove our claim.

We first introduce two optimization problems whose optimal solutions are closely related to $y^*(\mathbf{0})$ and $y^*(\epsilon)$. The first optimization problem is almost identical to $\mathbf{NP}_y(\mathbf{0})$ except that the domain is \mathbb{R}^n instead of \mathcal{Y} :

$$(\bar{\mathbf{NP}}_y) \quad \max_{y \in \mathbb{R}^n} \left\{ h(y) : Uy \leq V \right\}.$$

To define our second optimization problem, first note that, since U has full row rank, there exists an n by m matrix H such that UH = I. For any $\epsilon \in \mathbb{R}^m$, by a change of variables $y = z + H\epsilon$, we can transform $\mathbf{NP}_y(\epsilon)$ into an equivalent optimization problem below:

$$(\mathbf{NP}_{z}(\epsilon)) \qquad \max_{z\in\mathcal{Y}-H\epsilon} \{h_{\epsilon}(z): Uz \leq V\},\$$

where $h_{\epsilon}(z) := h(z + H\epsilon)$. Two important observations are in order. The first observation relates $y^*(\mathbf{0})$ to the first optimization problem $\bar{\mathbf{NP}}_y$ whereas the second observation relates $y^*(\epsilon)$ to the second optimization problem $\mathbf{NP}_z(\epsilon)$.

Observation 1: $y^*(\mathbf{0})$ is the unique optimal solution to \bar{NP}_y .

Suppose that this is not true, i.e., there exists an optimal solution of $\overline{\mathbf{NP}}_y$, which we denote by \tilde{y} , and $\tilde{y} \neq y^*(\mathbf{0})$. Then $h(\tilde{y}) \geq h(y^*(\mathbf{0}))$ and $y_\alpha = \alpha y^*(\mathbf{0}) + (1 - \alpha)\tilde{y} \quad (\forall \alpha \in [0, 1])$ is a feasible solution. By strong concavity, $h(y_\alpha) > \alpha h(y^*(\mathbf{0})) + (1 - \alpha)h(\tilde{y}) \geq h(y^*(\mathbf{0}))$ for all $\alpha \in (0, 1)$. Since $y^*(\mathbf{0})$ is in the interior of \mathcal{Y} , so is y_α if α is sufficiently small, which contradicts with the fact that $y^*(\mathbf{0})$ maximizes $\mathbf{NP}_y(\mathbf{0})$.

Observation 2: There exists $\bar{\epsilon}_1 > 0$ such that for all ϵ with $||\epsilon||_{\infty} \leq \bar{\epsilon}_1$, $NP_y(\epsilon)$ has a unique optimal solution $y^*(\epsilon)$ and $z^*(\epsilon) := y^*(\epsilon) - H\epsilon$ is the unique optimal solution of $NP_z(\epsilon)$.

We now prove Observation 2. Since $y^*(\mathbf{0})$ lies in the interior of \mathcal{Y} , there exists a constant $\bar{\phi} > 0$ such that $\{x : ||x - y^*(\mathbf{0})||_{\infty} \leq \bar{\phi}\} \subseteq \mathcal{Y}$. Let $\bar{\epsilon}_1 := \min_i \{\bar{\phi}(U\mathbf{e})_i\}$. Note that $\bar{\epsilon}_1 > 0$ since U is nonnegative and has full row rank. Moreover, for all ϵ with $||\epsilon||_{\infty} \leq \bar{\epsilon} \leq \bar{\epsilon}_1$, $y^*(\mathbf{0}) - \bar{\phi}\mathbf{e} \in \mathcal{Y}$ is a feasible solution of $\mathbf{NP}_y(\epsilon)$ because $U(y^*(\mathbf{0}) - \bar{\phi}\mathbf{e}) \leq V - \bar{\phi}U\mathbf{e} \leq V - \bar{\epsilon}_1\mathbf{e} \leq V + \epsilon$. This means that $\mathbf{NP}_y(\epsilon)$ has a unique optimal solution $y^*(\epsilon)$ (because its feasible region is convex, compact, not empty, and its objective function h(.) is strongly concave). Hence, by definition of $\mathbf{NP}_z(\epsilon)$, $z^*(\epsilon)$ is its unique optimal solution. So, Observation 2 holds.

To bound $||y^*(\mathbf{0}) - y^*(\epsilon)||_2$, we first derive a bound for $||y^*(\mathbf{0}) - z^*(\epsilon)||_2$. Let $\bar{\epsilon}_2 := \bar{\phi}/||H||_{\infty}$. Then, for all ϵ with $||\epsilon||_{\infty} \leq \bar{\epsilon} \leq \bar{\epsilon}_2$, since $||y^*(\mathbf{0}) + H\epsilon - y^*(\mathbf{0})||_{\infty} = ||H\epsilon||_{\infty} \leq ||H||_{\infty} ||\epsilon||_{\infty} \leq \bar{\phi}$, $y^*(\mathbf{0}) + H\epsilon \in \mathcal{Y}$. This means that $y^*(\mathbf{0})$ is feasible for $\mathbf{NP}_z(\epsilon)$ and $h_\epsilon(y^*(\mathbf{0})) \leq h_\epsilon(z^*(\epsilon))$. Note that $z^*(\epsilon) \in \mathbb{R}^n$ is also feasible for \mathbf{NP}_y , so

$$h(z^{*}(\epsilon)) \leq h(y^{*}(\mathbf{0})) + \nabla h(y^{*}(\mathbf{0})) \cdot (z^{*}(\epsilon) - y^{*}(\mathbf{0})) - \frac{v}{2} ||z^{*}(\epsilon) - y^{*}(\mathbf{0})||_{2}^{2}$$

$$\leq h(y^{*}(\mathbf{0})) - \frac{v}{2} ||z^{*}(\epsilon) - y^{*}(\mathbf{0})||_{2}^{2}$$
(A20)

for some v > 0 that only depends on h(.). The first inequality follows by the strong concavity of h(.)and the second inequality follows because $\nabla h(y^*(\mathbf{0})) \cdot (z^*(\epsilon) - y^*(\mathbf{0})) \leq 0$ (otherwise, $y^*(\mathbf{0})$ cannot be the optimal solution of $\bar{\mathbf{NP}}_y$). Note also that, for all ϵ with $||\epsilon||_{\infty} \leq \bar{\epsilon}$, the following holds

$$\frac{v}{2}||z^{*}(\epsilon) - y^{*}(\mathbf{0})||_{2}^{2} \leq h(y^{*}(\mathbf{0})) - h(z^{*}(\epsilon)) - h_{\epsilon}(y^{*}(\mathbf{0})) + h_{\epsilon}(z^{*}(\epsilon))
= (\nabla h(\xi_{1}) - \nabla h_{\epsilon}(\xi_{1})) \cdot (y^{*}(\mathbf{0}) - z^{*}(\epsilon))
\leq ||\nabla h(\xi_{1}) - \nabla h_{\epsilon}(\xi_{1})||_{\infty} ||y^{*}(\mathbf{0}) - z^{*}(\epsilon)||_{\infty}
\leq ||\nabla h(\xi_{1}) - \nabla h(\xi_{1} + H\epsilon)||_{\infty} ||y^{*}(\mathbf{0}) - z^{*}(\epsilon)||_{2}
= ||\nabla^{2}h(\xi_{2})H\epsilon||_{\infty} ||y^{*}(\mathbf{0}) - z^{*}(\epsilon)||_{2}
\leq K_{0} ||H||_{\infty} ||\epsilon||_{\infty} ||y^{*}(\mathbf{0}) - z^{*}(\epsilon)||_{2}$$
(A21)

for some $\xi_1 \in \mathbb{Z}, \xi_2 \in \overline{\mathbb{Z}}$ where $\mathbb{Z} := \{z : ||z - y||_{\infty} \leq \overline{\phi} \text{ for some } y \in \mathcal{Y}\}$ and $\overline{\mathbb{Z}} := \{z : ||z - x||_{\infty} \leq \overline{\phi} \text{ for some } x \in \mathbb{Z}\}$ are both compact and convex, and $K_0 := \sup_{z \in \overline{\mathbb{Z}}} ||\nabla^2 h(z)||_{\infty}$ only depends on h(.). The first inequality of (A21) follows by (A20) and $h_{\epsilon}(y^*(\mathbf{0})) \leq h_{\epsilon}(z^*(\epsilon))$. The first equality of (A21) follows by the Mean Value Theorem and the fact that $y^*(\mathbf{0}) \in \mathbb{Z}$ and $z^*(\epsilon) \in \mathbb{Z}$ (the latter inclusion holds because, since $z^*(\epsilon) \in \mathcal{Y} - H\epsilon$, there exists a $y \in \mathcal{Y}$ such that $||z^*(\epsilon) - y||_{\infty} = ||H\epsilon||_{\infty} \leq ||H||_{\infty} \overline{\epsilon}_2 = \overline{\phi}$). Similarly, the second equality also follows by the Mean Value Theorem and the fact that $\xi_1 \in \mathbb{Z} \subseteq \overline{\mathbb{Z}}$ and $\xi_1 + H\epsilon \in \overline{\mathbb{Z}}$ (the latter inclusion holds since $\xi_1 \in \mathbb{Z}$ and $||\xi_1 + H\epsilon - \xi_1||_{\infty} \leq ||H||_{\infty} \overline{\epsilon} \leq ||H||_{\infty} \overline{\epsilon}_2 = \overline{\phi}$). Note that (A21) is equivalent to $||z^*(\epsilon) - y^*(\mathbf{0})||_2 \leq 2v^{-1}K_0||H||_{\infty}||\epsilon||_{\infty}$. Let $K_1 = 2v^{-1}K_0||H||_{\infty} + ||H||_2\sqrt{n}$. Then, for all ϵ with $||\epsilon||_{\infty} \leq \overline{\epsilon}$, we can bound:

$$\begin{aligned} ||y^{*}(\epsilon) - y^{*}(\mathbf{0})||_{2} &= ||z^{*}(\epsilon) + H\epsilon - y^{*}(\mathbf{0})||_{2} \le ||z^{*}(\epsilon) - y^{*}(\mathbf{0})||_{2} + ||H||_{2}||\epsilon||_{2} \\ &\le 2v^{-1}K_{0}||H||_{\infty}||\epsilon||_{\infty} + ||H||_{2}||\epsilon||_{2} \le K_{1}||\epsilon||_{\infty}. \end{aligned}$$

This proves the claim we stated at the beginning and completes the proof of Lemma A3. \Box

Appendix C: Proof of Lemma 2

We now prove Lemma 2 using Proposition 1. Let $g(.) = \lambda^*(.), \tilde{g}(.) = \tilde{\lambda}(.), f(.) = r^*(.), \tilde{f}(.) = \tilde{r}(.), U = A, V = C/T, \delta = \mathbf{0}, \mathcal{X} = \mathcal{P}, \mathcal{Y} = \Lambda_{\lambda^*}$. Note that $r^*(.), \lambda^*(.)$ are twice continuously differentiable by R1 and $\tilde{r}(.), \tilde{\lambda}(.)$ are continuously differentiable since $\tilde{\lambda}(.) \in \mathbb{C}^{s-2}(\mathcal{P})$ and $s - 2 \geq 1$. Also \mathcal{P} is convex and A is nonnegative with full row rank. We first verify the conditions (i) - (iv) in Proposition 1. By R1, $\lambda^*(.)$ has a twice continuously differentiable inverse function $p^*(.)$, and Λ_{λ^*} is assumed to be convex, so (i) holds. By R3, $r^*_{\lambda}(.) := r^*(p^*(.))$ is strongly concave, so (ii) holds. By

R4, **P** is feasible and its optimal solution p^D lies in the interior of \mathcal{P} , so both (iii) and (iv) hold. For any $p \in \mathcal{P}$, we have

$$\begin{aligned} ||(\nabla r^*(p) - \nabla \tilde{r}(p))'||_{\infty} &= ||(\lambda^*(p) + \nabla \lambda^*(p)p - \lambda(p) - \nabla \lambda(p)p)'||_{\infty} \\ &\leq ||(\lambda^*(p) - \tilde{\lambda}(p))'||_{\infty} + ||p'(\nabla \lambda^*(p) - \nabla \tilde{\lambda}(p))'||_{\infty} \\ &\leq n||\lambda^*(p) - \tilde{\lambda}(p)||_{\infty} + ||p'||_{\infty}||(\nabla \lambda^*(p) - \nabla \tilde{\lambda}(p))'||_{\infty} \\ &\leq n||\lambda^*(.) - \tilde{\lambda}(.)||_{\infty} + (\sum_{l=1}^n \bar{p}_l)||(\nabla \lambda^*(.) - \nabla \tilde{\lambda}(.))'||_{\infty}. \end{aligned}$$

Therefore, by Proposition 1, there exists $\bar{\delta}_1 > 0$ and $K_1 > 0$ such that for all $\tilde{\lambda}(.)$ satisfying $||A\lambda(.) - A\tilde{\lambda}(.)||_{\infty} \leq \bar{\delta}_1$, $\tilde{\mathbf{P}}$ is feasible and

$$\begin{aligned} ||p^{D} - \tilde{p}^{D}||_{2} &\leq K_{1}(||A\lambda^{*}(.) - A\lambda(.)||_{\infty} + ||(\nabla r^{*}(.) - \nabla \tilde{r}(.))'||_{\infty}) \\ &\leq K_{1}[(n + ||A||_{\infty})||\lambda^{*}(.) - \tilde{\lambda}(.)||_{\infty} + (\sum_{l=1}^{n} \bar{p}_{l})||(\nabla \lambda^{*}(.) - \nabla \tilde{\lambda}(.))'||_{\infty}] \\ &\leq K(||\lambda^{*}(.) - \tilde{\lambda}(.)||_{\infty} + ||(\nabla \lambda^{*}(.) - \nabla \tilde{\lambda}(.))'||_{\infty}) \end{aligned}$$

where $K = K_1(n+||A||_{\infty} + \sum_{l=1}^{n} \bar{p}_l)$ is independent of $\tilde{\lambda}(.)$. Let $\bar{\delta} := \bar{\delta}_1/||A||_{\infty}$. Since $||\lambda(.) - \tilde{\lambda}(.)||_{\infty} \le \bar{\delta}$ means that $||A\lambda(.) - A\tilde{\lambda}(.)||_{\infty} \le ||A||_{\infty} \bar{\delta} = \bar{\delta}_1$, the result follows. \Box

Appendix D: Derivation of the equality (2)

Recall that $u_{ij}^* := [\frac{\partial^2 \lambda_1^*(p^D)}{\partial p_i \partial p_j}; \dots; \frac{\partial^2 \lambda_n^*(p^D)}{\partial p_i \partial p_j}]$. Note that the following identity holds:

$$H_{ij}^{*} = \left[B^{*}\nabla^{2}r_{\lambda}^{*}(\lambda^{D})(B^{*})'\right]_{ij} - \left[B^{*} + (B^{*})'\right]_{ij} = \sum_{l=1}^{n}\sum_{k=1}^{n}\frac{\partial\lambda_{k}^{*}(p^{D})}{\partial p_{i}}\frac{\partial^{2}r_{\lambda}^{*}(\lambda^{D})}{\partial\lambda_{k}\partial\lambda_{l}}\frac{\partial\lambda_{l}^{*}(p^{D})}{\partial p_{j}} - \left[\frac{\partial\lambda_{i}^{*}(p^{D})}{\partial p_{j}} + \frac{\partial\lambda_{j}^{*}(p^{D})}{\partial p_{i}}\right].$$
(A22)

Note also that $r^*(p) = r^*_{\lambda}(\lambda^*(p)) = p'\lambda^*(p)$; taking the second order derivative,

$$\frac{\partial^2 r^*(p^D)}{\partial p_i \partial p_j} = \sum_{l=1}^n \sum_{k=1}^n \frac{\partial \lambda_k^*(p^D)}{\partial p_i} \frac{\partial^2 r_\lambda^*(\lambda^D)}{\partial \lambda_k \partial \lambda_l} \frac{\partial \lambda_l^*(p^D)}{\partial p_j} + \sum_{l=1}^n \frac{\partial r_\lambda^*(\lambda^D)}{\partial \lambda_l} \frac{\partial^2 \lambda_l^*(p^D)}{\partial p_i \partial p_j}
= \frac{\partial \lambda_i^*(p^D)}{\partial p_j} + \frac{\partial \lambda_j^*(p^D)}{\partial p_i} + \sum_{l=1}^n p_l^D \frac{\partial^2 \lambda_l^*(p^D)}{\partial p_i \partial p_j}.$$
(A23)

Hence, combining (A23) with (A22), we have $H_{ij}^* = \sum_{l=1}^n (p_l^D - \frac{\partial r_{\lambda}^*(\lambda^D)}{\partial \lambda_l}) \frac{\partial^2 \lambda_l^*(p^D)}{\partial p_i \partial p_j} = (u_{ij}^*)'(p^D - \nabla r_{\lambda}^*(\lambda^D))$. Note that $\lambda^*(p^D) + \nabla \lambda^*(p^D) p^D = \nabla r^*(p^D) = \nabla \lambda^*(p^D) \nabla r_{\lambda}^*(\lambda^D)$, so $p^D - \nabla r_{\lambda}^*(\lambda^D) = -\nabla \lambda^*(p^D)^{-1} \lambda^*(p^D) = -(B^*)^{-1} \lambda^D$. Hence, $H_{ij}^* = -(u_{ij}^*)'(B^*)^{-1} \lambda^D$.

Appendix E: Proof of Lemma 3

We will prove each part of the lemma in turn. Let $\bar{\delta} = \min\{\bar{\delta}_1, \bar{\delta}_2, \bar{\delta}_3, \bar{\delta}_4, \bar{\delta}_5\}$, $\kappa = \max\{\kappa_1, \kappa_2\}$, where $\bar{\delta}_1, \ldots, \bar{\delta}_5$ and κ_1, κ_2 are strictly positive constants to be defined shortly.

Proof of Parts (a)-(c)

Let $\sigma_{\max}(X)$ and $\sigma_{\min}(X)$ denote the maximum and minimum eigenvalues of a symmetric real matrix X, respectively. Since $B^* = \nabla \lambda^*(p^D)$ is invertible, $B^*(B^*)'$ is positive definite; so, $\bar{\sigma}^* := \sigma_{\max}(B^*(B^*)') > 0$ and $\underline{\sigma}^* := \sigma_{\min}(B^*(B^*)') > 0$. Define $\bar{\delta}_1 = \underline{\sigma}^*/(4\sqrt{\bar{\sigma}^*}) > 0$. Note that, for all $\theta_\iota \in \mathsf{Ball}(\theta_\iota^*, \bar{\delta})$, $||B - B^*||_2 \leq ||B - B^*||_F \leq ||\theta_\iota - \theta_\iota^*||_2 \leq \bar{\delta} \leq \bar{\delta}_1$. Therefore, for all $v \in \mathbb{R}^n$ such that $||v||_2 = 1$, we have:

$$\begin{aligned} v'B'Bv &= v'(B-B^*+B^*)'(B-B^*+B^*)v \\ &= v'(B^*)'B^*v + v'(B^*)'(B-B^*)v + v'(B-B^*)'B^*v + v'(B-B^*)'(B-B^*)v \\ &\geq \underline{\sigma}^* - 2||v||_2^2||(B^*)'||_2||B-B^*||_2 \geq \underline{\sigma}^* - 2\sqrt{\overline{\sigma}^*}\,\underline{\sigma}^*/(4\sqrt{\overline{\sigma}^*}) = \underline{\sigma}^*/2. \end{aligned}$$

This means that $\sigma_{\min}(B'B) \ge \underline{\sigma}^*/2 > 0$ and B is invertible. Since $(B'B)^{-1} = B^{-1}(B^{-1})'$,

$$|(B')^{-1}||_2 = \sqrt{\sigma_{\max}(B^{-1}(B^{-1})')} = \sqrt{\sigma_{\min}(B'B)^{-1}} \le \sqrt{2/\underline{\sigma}^*}.$$

By definition, $\lambda(p;\theta_{\iota}) = a + B'p, \lambda(p;\theta_{\iota}^{*}) = a^{*} + (B^{*})'p, \frac{\partial\lambda_{j}}{\partial p_{i}}(p;\theta_{\iota}) = B_{ij}, \frac{\partial\lambda_{j}}{\partial p_{i}}(p;\theta_{\iota}^{*}) = B_{ij}^{*}$ for all $i, j \in \overline{[1,n]}$, and $p(\lambda;\theta_{\iota}) = -(B')^{-1}a + (B')^{-1}\lambda$ (recall that B is invertible). So, for all $p \in \mathcal{P}$ and $\lambda, \lambda' \in \lambda(\mathcal{P};\theta_{\iota})$,

$$\begin{split} ||\lambda(p;\theta_{\iota}) - \lambda(p;\theta_{\iota}^{*})||_{2} &\leq ||a - a^{*}||_{2} + ||(B - B^{*})'||_{2}||p||_{2} \\ &\leq ||a - a^{*}||_{2} + ||(B - B^{*})'||_{F} \left(\sum_{l=1}^{n} \bar{p}_{l}^{2}\right)^{\frac{1}{2}} \\ &\leq \left[1 + \left(\sum_{l=1}^{n} \bar{p}_{l}^{2}\right)^{\frac{1}{2}}\right] ||\theta_{\iota} - \theta_{\iota}^{*}||_{2}, \\ \left|\frac{\partial\lambda_{j}}{\partial p_{i}}(p;\theta_{\iota}) - \frac{\partial\lambda_{j}}{\partial p_{i}}(p;\theta_{\iota}^{*})\right| &= |B_{ij} - B_{ij}^{*}| \leq ||\theta_{\iota} - \theta_{\iota}^{*}||_{2}, \text{ and} \\ &||p(\lambda;\theta_{\iota}) - p(\lambda';\theta_{\iota})||_{2} = ||(B')^{-1}(\lambda - \lambda')||_{2} \leq ||(B')^{-1}||_{2}||\lambda - \lambda'||_{2} \leq \sqrt{2/\underline{\sigma}^{*}}||\lambda - \lambda'||_{2}. \end{split}$$

The results follow by letting $\omega = \max\{1 + (\sum_{l=1}^{n} \bar{p}_l^2)^{\frac{1}{2}}, \sqrt{2/\underline{\sigma}^*}\}.$

Proof of part (d)

To prove that $q(p(.;\theta_{\iota});\theta_{o})$ is strongly concave, we need to show that its Hessian, $B^{-1}G(B')^{-1}$, is negative definite. Let $\bar{\delta}_{2} = ||G^{*}||_{2}$. For all $\theta_{o} \in \mathsf{Ball}(\theta_{o}^{*},\bar{\delta}), ||\theta_{o} - \theta_{o}^{*}||_{2} \leq \bar{\delta} \leq \bar{\delta}_{2}$, so

$$||B^* - B||_2 \le ||B^* - B||_F \le \bar{\delta} \tag{A24}$$

$$||G^* - G||_2 \le ||G^* - G||_F \le \bar{\delta} \tag{A25}$$

$$||G||_{2} \leq ||G^{*}||_{2} + ||G - G^{*}||_{2} \leq ||G^{*}||_{2} + \bar{\delta}_{2} = 2||G^{*}||_{2}.$$
(A26)

Recall that, by Lemma 3 part (a), B is invertible and $||B^{-1}||_2 = ||(B')^{-1}||_2 \le \omega$. So,

$$\begin{split} |(B^*)^{-1}G^*((B^*)')^{-1} - B^{-1}G(B')^{-1}||_2 \\ &= ||((B^*)^{-1} - B^{-1})G^*((B^*)')^{-1} + B^{-1}(G^* - G)((B^*)')^{-1} + B^{-1}G(((B^*)')^{-1} - (B')^{-1})||_2 \\ &\leq ||(B^*)^{-1} - B^{-1}||_2||G^*||_2||((B^*)')^{-1}||_2 + ||B^{-1}||_2||G^* - G||_2||((B^*)')^{-1}||_2 \\ &+ ||B^{-1}||_2||G||_2||((B^*)')^{-1} - (B')^{-1}||_2 \\ &\leq ||(B^*)^{-1}||_2||B^* - B||_2||B^{-1}||_2||G^*||_2||((B^*)')^{-1}||_2 + ||B^{-1}||_2||G^* - G||_2||((B^*)')^{-1}||_2 \\ &+ ||B^{-1}||_2||G||_2||(B^*)^{-1}||_2||B^* - B||_2||B^{-1}||_2 \\ &\leq ||B^{-1}||_2||G^*||_2||(B^*)^{-1}||_2||G^*||_2 + 1 + 2||G^*||_2||B^{-1}||_2)\bar{\delta} \\ &\leq \omega||(B^*)^{-1}||_2(||(B^*)^{-1}||_2||G^*||_2 + 1 + 2||G^*||_2\omega)\bar{\delta} \\ &\leq C\bar{\delta}, \end{split}$$

for some C > 0 that only depends on θ^* and ω . The second inequality above holds because $(B^*)^{-1} - B^{-1} = (B^*)^{-1}(B - B^*)B^{-1}$ and the third inequality follows from (A24)-(A26).

By (3), $(B^*)^{-1}G^*((B^*)')^{-1} = \nabla^2 r_{\lambda}^*(\lambda^D)$. So, $\sigma_{\max}((B^*)^{-1}G^*((B^*)')^{-1}) \leq -\underline{v}$ by R3. Let $\bar{\delta}_3 = \underline{v}/(2C)$. Then, for all v such that $||v||_2 = 1$,

$$\begin{aligned} v'(B^{-1}G(B')^{-1})v &= v'((B^*)^{-1}G^*((B^*)')^{-1})v + v'(B^{-1}G(B')^{-1} - (B^*)^{-1}G^*((B^*)')^{-1})v \\ &\leq \sigma_{\max}((B^*)^{-1}G^*((B^*)')^{-1}) + ||B^{-1}G(B')^{-1} - (B^*)^{-1}G^*((B^*)')^{-1}||_2 \\ &\leq -\underline{v} + C\bar{\delta} \leq -\underline{v} + C\bar{\delta}_3 = -\underline{v}/2. \end{aligned}$$

This means that $B^{-1}G(B')^{-1}$ is negative definite and, thus, $q(p(:;\theta_{\iota});\theta_{o})$ is strongly concave.

Proof of part (e)

Consider the following optimization problem:

$$(\mathbf{QP}_{\lambda}(\theta;\delta)) \quad \max_{\lambda \in \lambda(\mathcal{P};\theta_{\iota})} \left\{ q(p(\lambda;\theta_{\iota});\theta_{o}): A\lambda \preceq \frac{C}{T} - \delta \right\}.$$

This problem is equivalent to $\mathbf{QP}(\theta; \delta)$, except that it optimizes over λ instead of p. Recall that, by Lemma 3 part (d), $q(p(.;\theta_{\iota});\theta_{o})$ is strongly concave; so, $\mathbf{QP}_{\lambda}(\theta;\delta)$ is a convex program and, if it is feasible, it has a unique optimal solution that is characterized by the KKT condition. Since \mathbf{P}_{λ} is a convex program and has a unique optimal solution $\lambda^{D} \in \Lambda_{\lambda^{*}}$, by the KKT conditions, we have:

$$\begin{cases} \nabla r_{\lambda}^{*}(\lambda^{D}) = A'\mu^{D} \\ (A\lambda^{D} - \frac{C}{T})'\mu^{D} = 0 \\ \mu^{D} \succeq \mathbf{0}, A\lambda^{D} \preceq \frac{C}{T} \end{cases} \text{ which, by (3), implies } \begin{cases} \nabla_{\lambda}q(p(\lambda^{D}; \theta_{\iota}^{*}); \theta_{o}^{*}) = A'\mu^{D} \\ (A\lambda^{D} - \frac{C}{T})'\mu^{D} = 0 \\ \mu^{D} \succeq \mathbf{0}, A\lambda^{D} \preceq \frac{C}{T} \end{cases}$$

Since $\lambda^D = \lambda(p^D; \theta_{\iota}^*) \in \lambda(\mathcal{P}; \theta_{\iota}^*)$ is feasible to $\mathbf{QP}_{\lambda}(\theta^*; \mathbf{0})$, by the sufficiency of KKT conditions for optimality in a strongly convex program, λ^D and μ^D are also the *unique* optimal primal and dual solution of $\mathbf{QP}_{\lambda}(\theta^*; \mathbf{0})$. Hence, $p^D = p(\lambda^D; \theta_{\iota}^*)$ is also the unique optimal solution of $\mathbf{QP}(\theta^*; \mathbf{0})$. This proves that $\lambda^D = \lambda_{\mathbf{0}}^D(\theta^*), \mu^D = \mu_{\mathbf{0}}^D(\theta^*)$ and $p^D = p_{\mathbf{0}}^D(\theta^*)$.

Proof of part (f)

The proof relies on Proposition 1. Let $g(.) = \lambda(.;\theta_{\iota}^{*}), \tilde{g}(.) = \lambda(.;\theta_{\iota}), f(.) = q(.;\theta_{o}^{*}), \tilde{f}(.) = q(.;\theta_{o}), U = A, V = C/T, \mathcal{X} = \mathcal{P}, \mathcal{Y} = \lambda(\mathcal{P};\theta_{\iota}^{*})$. We first verify conditions (i)-(iv) of Proposition 1. Since $\lambda(.;\theta_{\iota}^{*})$ is linear and $B^{*} = \nabla \lambda^{*}(p^{D})$ is invertible, it has an inverse function $p(.;\theta_{\iota}^{*})$ that is linear and, hence, twice continuously differentiable. Moreover, the set $\lambda(\mathcal{P};\theta_{\iota}^{*})$ is convex because \mathcal{P} is convex and convexity is preserved under affine transformation. So, (i) holds. By (3) and the fact that $r_{\lambda}^{*}(.)$ is strongly concave, $q(p(.;\theta_{\iota}^{*});\theta_{o}^{*})$ is strongly concave, so (ii) holds. As shown earlier, p^{D} is the optimal solution of $\mathbf{QP}(\theta^{*};\mathbf{0})$ and it is in the interior of \mathcal{P} , so (iii) and (iv) hold. Therefore, by Proposition 1, there exists some constant $\tilde{\delta} > 0$ such that if $||A\lambda(.;\theta_{\iota}^{*}) - A\lambda(.;\theta_{\iota}) + \delta||_{\infty} \leq \tilde{\delta}$, then $\mathbf{QP}(\theta;\delta)$ is feasible and there exists some constant K independent of θ such that the unique optimal solution of $\mathbf{QP}(\theta;\delta)$ (i.e., $p_{\delta}^{D}(\theta)$) satisfies the following:

$$\begin{split} ||p_{\mathbf{0}}^{D}(\theta^{*}) - p_{\delta}^{D}(\theta)||_{2} &\leq K \left(||(\nabla q(.;\theta_{o}^{*}) - \nabla q(.;\theta_{o}))'||_{\infty} + ||\lambda(.;\theta_{\iota}^{*}) - \lambda(.;\theta_{\iota})||_{\infty} + ||\delta||_{\infty} \right) \\ &= K \left(\sup_{p \in \mathcal{P}} \{ ||(F^{*} + p'G^{*} - F - p'G)'||_{\infty} \} + \sup_{p \in \mathcal{P}} \{ ||\lambda(p;\theta_{\iota}^{*}) - \lambda(p;\theta_{\iota})||_{2} \} + ||\delta||_{\infty} \right) \\ &\leq K \left(\sup_{p \in \mathcal{P}} \{ ||(F^{*} + p'G^{*} - F - p'G)'||_{2} \} + \sup_{p \in \mathcal{P}} \{ ||\lambda(p;\theta_{\iota}^{*}) - \lambda(p;\theta_{\iota})||_{2} \} + ||\delta||_{\infty} \right) \\ &\leq K \left(||F^{*} - F||_{2} + \sup_{p \in \mathcal{P}} \{ ||p'||_{2} \} ||G^{*} - G||_{2} + \omega ||\theta_{\iota}^{*} - \theta_{\iota}||_{2} + ||\delta||_{2} \right) \\ &\leq K \left(||F^{*} - F||_{2} + (\sum_{l=1}^{n} \bar{p}_{l}^{2})^{1/2} ||G^{*} - G||_{F} + \omega ||\theta_{\iota}^{*} - \theta_{\iota}||_{2} + ||\delta||_{2} \right) \\ &\leq K \left(1 + (\sum_{l=1}^{n} \bar{p}_{l}^{2})^{1/2} ||\theta_{o}^{*} - \theta_{o}||_{2} + K\omega ||\theta_{\iota}^{*} - \theta_{\iota}||_{2} + K ||\delta||_{2} \\ &\leq \kappa_{1} \left(||\theta^{*} - \theta||_{2} + ||\delta||_{2} \right) \\ &\leq \kappa (||\theta^{*} - \theta||_{2} + ||\delta||_{2}), \end{split}$$

where $\kappa_1 = K(2 + (\sum_{l=1}^n \bar{p}_l^2)^{1/2} + \omega)$. Let $\bar{\delta}_4 = \min\{\tilde{\delta}/2, \tilde{\delta}/(2\omega||A||_2), \phi/(6\kappa)\}$. Then, for $\theta_{\iota} \in \mathsf{Ball}(\theta_{\iota}^*, \bar{\delta})$ and $\delta \in \mathsf{Ball}(\mathbf{0}, \bar{\delta})$, we have

$$\begin{aligned} ||A\lambda(.;\theta_{\iota}^{*}) - A\lambda(.;\theta_{\iota}) + \delta||_{\infty} &\leq \sup_{p \in \mathcal{P}} ||A\lambda(p;\theta_{\iota}^{*}) - A\lambda(p;\theta_{\iota})||_{\infty} + ||\delta||_{\infty} \\ &\leq \sup_{p \in \mathcal{P}} ||A\lambda(p;\theta_{\iota}^{*}) - A\lambda(p;\theta_{\iota})||_{2} + ||\delta||_{2} \leq \omega ||A||_{2} ||\theta_{\iota}^{*} - \theta_{\iota}||_{2} + ||\delta||_{2} \\ &\leq \omega ||A||_{2}\bar{\delta} + \bar{\delta} \leq \omega ||A||_{2}\bar{\delta}_{4} + \bar{\delta}_{4} \leq \omega ||A||_{2} \frac{\tilde{\delta}}{2\omega ||A||_{2}} + \frac{\tilde{\delta}}{2} = \tilde{\delta}. \end{aligned}$$

Moreover, $p_{\delta}^{D}(\theta) \in \mathsf{Ball}(p_{\mathbf{0}}^{D}(\theta^{*}), \phi/2)$ because

$$\begin{aligned} ||p_{\mathbf{0}}^{D}(\theta^{*}) - p_{\delta}^{D}(\theta)||_{2} &\leq \kappa (||\theta^{*} - \theta||_{2} + ||\delta||_{2}) \leq \kappa (||\theta_{\iota}^{*} - \theta_{\iota}||_{2} + ||\theta_{o}^{*} - \theta_{o}||_{2} + ||\delta||_{2}) \\ &\leq 3\kappa\bar{\delta} \leq 3\kappa\bar{\delta}_{4} \leq 2\kappa\phi/(6\kappa) = \phi/3 < \phi/2. \end{aligned}$$

Since $\mathsf{Ball}(p_0^D(\theta^*), \phi) \subseteq \mathcal{P}$ (by R4 and $p_0^D(\theta^*) = p^D$), $\mathsf{Ball}(p_\delta^D(\theta), \phi/2) \subseteq \mathcal{P}$. Let $\kappa_2 = 2\kappa_1 ||B^*||_2 + \omega$. We have:

$$\begin{aligned} ||\lambda_{\mathbf{0}}^{D}(\theta^{*}) - \lambda_{\delta}^{D}(\theta)||_{2} &= ||\lambda(p_{\mathbf{0}}^{D}(\theta^{*});\theta_{\iota}^{*}) - \lambda(p_{\delta}^{D}(\theta);\theta_{\iota})||_{2} \\ &\leq ||\lambda(p_{\mathbf{0}}^{D}(\theta^{*});\theta_{\iota}^{*}) - \lambda(p_{\delta}^{D}(\theta);\theta_{\iota}^{*})||_{2} + ||\lambda(p_{\delta}^{D}(\theta);\theta_{\iota}^{*}) - \lambda(p_{\delta}^{D}(\theta);\theta_{\iota})||_{2} \\ &\leq ||B^{*}||_{2}||p_{\mathbf{0}}^{D}(\theta^{*}) - p_{\delta}^{D}(\theta)||_{2} + \omega||\theta^{*} - \theta||_{2} \\ &\leq (\kappa_{1}||B^{*}||_{2} + \omega)||\theta^{*} - \theta||_{2} + \kappa_{1}||B^{*}||_{2}||\delta||_{2} \\ &= \kappa_{2}(||\theta^{*} - \theta||_{2} + ||\delta||_{2}) \\ &\leq \kappa(||\theta^{*} - \theta||_{2} + ||\delta||_{2}). \end{aligned}$$

We will now show that the constraints of $\mathbf{QP}(\theta; \delta)$ that correspond to rows $\{i : \mu_{\mathbf{0},i}^{D}(\theta^{*}) > 0\}$ are binding. Since $\mathbf{QP}_{\lambda}(\theta; \delta)$ is a feasible convex program, by KKT condition, $\nabla_{\lambda}q(p(\lambda_{\delta}^{D}(\theta); \theta_{\iota}); \theta_{o}) = A'\mu_{\delta}^{D}(\theta)$. By our assumption, A has full row rank. So, there exists some m by n matrix \overline{A} such that $\mu_{\delta}^{D}(\theta) = \overline{A}\nabla_{\lambda}q(p(\lambda_{\delta}^{D}(\theta); \theta_{\iota}); \theta_{o})$. Since the right hand side is jointly continuous in $(\theta; \delta)$ at $(\theta^{*}; \mathbf{0})$, $\mu_{\delta}^{D}(\theta)$ must also be continuous in $(\theta; \delta)$ at $(\theta^{*}; \mathbf{0})$. Let $\mu := \min_{1 \leq i \leq n} \{\mu_{\mathbf{0},i}^{D}(\theta^{*}) : \mu_{\mathbf{0},i}^{D}(\theta^{*}) > 0\}$. By continuity, there exists $\overline{\delta}_{5} > 0$ such that $||\mu_{\delta}^{D}(\theta) - \mu_{\mathbf{0}}^{D}(\theta^{*})||_{2} < \mu$ for all $\theta = (\theta_{o}; \theta_{\iota})$ and δ satisfying $||\theta_{\iota} - \theta_{\iota}^{*}||_{2} \leq \overline{\delta}_{5}$, $||\theta_{o} - \theta_{o}^{*}||_{2} \leq \overline{\delta}_{5}$. Since, by definition, $\overline{\delta} \leq \overline{\delta}_{5}$, for all $\theta = (\theta_{o}; \theta_{\iota})$ and δ satisfying $||\theta_{\iota} - \theta_{\iota}^{*}||_{2} \leq \overline{\delta}$, $||\theta_{o} - \theta_{o}^{*}||_{2} \leq \overline{\delta}$ and $||\delta||_{2} \leq \overline{\delta}$, we have $\mu_{\delta,i}^{D}(\theta) > 0$ whenever $\mu_{\mathbf{0},i}^{D}(\theta^{*}) > 0$; so, the corresponding constraints in $\mathbf{QP}_{\lambda}(\theta; \delta)$ are binding due to the KKT condition that $(A\lambda_{\delta}^{N}(\theta) - \frac{C}{T})'\mu_{\delta}^{D}(\theta) = 0$. \Box

References

[1] Schumaker L (2007) Spline Functions: Basic Theory (Third Edition) (Cambridge University Press).