
Smoothness-Adaptive Dynamic Pricing with Nonparametric Demand Learning

Zeqi Ye
Nankai University

Hansheng Jiang
University of Toronto

Abstract

We study the dynamic pricing problem where the demand function is nonparametric and Hölder smooth, and we focus on adaptivity to the unknown Hölder smoothness parameter β of the demand function. Traditionally the optimal dynamic pricing algorithm heavily relies on the knowledge of β to achieve a minimax optimal regret of $\tilde{O}(T^{\frac{\beta+1}{2\beta+1}})$. However, we highlight the challenge of adaptivity in this dynamic pricing problem by proving that no pricing policy can adaptively achieve this minimax optimal regret without knowledge of β . Motivated by the impossibility result, we propose a self-similarity condition to enable adaptivity. Importantly, we show that the self-similarity condition does not compromise the problem’s inherent complexity since it preserves the regret lower bound $\Omega(T^{\frac{\beta+1}{2\beta+1}})$. Furthermore, we develop a smoothness-adaptive dynamic pricing algorithm and theoretically prove that the algorithm achieves this minimax optimal regret bound without the prior knowledge β .

1 INTRODUCTION

Dynamic pricing, the practice of adjusting prices in real-time based on varying market demand, has become an integral strategy in domains like e-commerce and transportation. An effective dynamic pricing model needs to adequately balance the exploration by learning demand at various prices and the exploitation by optimizing prices based on observed price and demand data. We consider a canonical dynamic pricing problem with nonparametric demand learning. At the

period t , the decision maker chooses a price p_t and observes a noisy demand d_t , where $\mathbb{E}[d_t|p_t = p] = f(p)$ for some unknown function $f : \mathcal{P} \rightarrow \mathbb{R}_{\geq 0}$ mapping from price set \mathcal{P} to demand. The goal of dynamic pricing is to maximize the total revenue collected over a finite time horizon. The performance of a dynamic pricing policy or algorithm is measured by the cumulative regret when compared with the maximal revenue in hindsight. More broadly framed as an online optimization problem, the dynamic pricing problem features nonparametric demand learning in that f can be of any functional form and continuous action space where price can be chosen at any value in a given price interval. Dynamic pricing problem has been an active topic for decades (Kleinberg and Leighton, 2003) and has found numerous applications in retailing, auctions, and advertising (Den Boer, 2015).

Without much regularity assumption on the demand function f , the optimal regret is shown to be $\tilde{O}(T^{2/3})$. This regret rate can be improved to $\tilde{O}(T^{1/2})$ if the uniqueness of the maximum and certain local concavity property of the revenue function $r(p) = p \cdot f(p)$ is imposed. However, such uniqueness assumption can be restrictive in practice and therefore other regularity assumptions, notably the smoothness condition, of the demand functions are considered. Nonetheless, a prevalent limitation in these methodologies is the presupposed exact knowledge of the Hölder smoothness level β . In reality, such assumptions are frequently misaligned with the complexities of real-world applications, thus constraining the practical applicability of these algorithms. Against this backdrop, our work distinguishes itself by delving into the uncharted territories of adaptability in dynamic pricing. Specifically, we address the pressing challenge of how to adapt when the Hölder smoothness level β is not known.

Facing the challenge of unknown smoothness parameter, it is natural to ask the following question:

Can we design a dynamic pricing strategy that does not require the prior knowledge of β while maintaining the optimal regret of $\tilde{O}(T^{\frac{\beta+1}{2\beta+1}})$?

Proceedings of the 27th International Conference on Artificial Intelligence and Statistics (AISTATS) 2024, Valencia, Spain. PMLR: Volume 238. Copyright 2024 by the author(s).

Our answer to this question is two-fold: on the one hand, it is impossible to achieve adaptivity without imposing additional assumptions; on the other hand, we identify a novel condition that achieves adaptivity without reducing the original pricing problem’s complexity. Our contributions in this paper can be summarized as follows:

- **Characterizing Adaptivity Challenge:** We formally characterize the challenge of adaptivity. In particular, we prove that without additional conditions, achieving the optimal regret for functions without knowing the Hölder smoothness parameter is impossible. We show that one algorithm with optimal regret for a certain Hölder smoothness parameter can have sub-optimal regret when directly applied to function class with lower Hölder smoothness levels.
- **Proposing a Self-Similarity Condition:** To make adaptivity possible, we propose a self-similarity condition, which serves as a dual to the Hölder smoothness assumption. Furthermore, our analysis reveals notable properties of the self-similarity condition, in particular regarding its practical applicability and sustenance of the dynamic pricing problem’s complexity. We find that the self-similarity condition not only enables adaptivity but also does not decrease the intrinsic complexity of the original pricing problem in that the lower bound $\Omega(T^{\frac{\beta+1}{2\beta+1}})$ does not change.
- **Optimal Minimax Regret Rate:** We design a Smoothness-Adaptive Dynamic Pricing (SADP) algorithm by incorporating a dedicated phase for the estimation of the smoothness parameter. Under the self-similarity condition, we establish a tight confidence interval for the estimated Hölder smoothness parameter. We derive an optimal regret bound $\tilde{O}(T^{\frac{\beta+1}{2\beta+1}})$ that matches the same optimal bound obtained by previous algorithms that require the knowledge of β .

Organization and Notation In Section 2, we introduce related literature on dynamic pricing, bandits, and statistics. In Section 3, we explicitly formulate the dynamic pricing problem under Hölder smooth demand functions and introduce the adaptivity problem by first presenting the non-adaptive dynamic pricing algorithm. We discuss in-depth the adaptivity challenge in Section 4 and present two key favorable properties of the self-similarity condition. In Section 5, we present our smoothness adaptive dynamic pricing algorithm and give a detailed regret analysis. In Section 6, we conduct numerical experiment and present the

corresponding results. Lastly, we conclude the paper with discussions and future directions in Section 7.

Throughout the paper, the vectors are column vectors unless specified otherwise. The notation $\|x\|$ denotes the L_2 norm of vector x , and given matrix A , the notation $\|x\|_A = (x^T A x)^{1/2}$ denotes the A -norm of vector x . For matrix A , $\|A\| = \sup_{x \neq 0} \|x^T A x\| / \|x\|$ denotes the L_2 operator norm of matrix A . We employ the notation $O(\cdot)$, $\Omega(\cdot)$, $\Theta(\cdot)$ to conceal constant factors, and $\tilde{O}(\cdot)$, $\tilde{\Omega}(\cdot)$, $\tilde{\Theta}(\cdot)$ are used to mask both constant and logarithmic factors.

2 RELATED LITERATURE

Dynamic Pricing with Demand Learning Motivated by the applications in e-commerce and transportation, numerous works have studied dynamic pricing with continuous price space and demand learning (Kleinberg and Leighton, 2003; Besbes and Zeevi, 2009; Broder and Rusmevichientong, 2012; Besbes and Zeevi, 2012; Keskin and Zeevi, 2014; Chen and Gallego, 2022). The crux of non-contextual dynamic pricing lies in modeling and learning the unknown price and demand relationship. Earlier works mainly focus on continuous parametric demand models with additional concavity property of the revenue function where a regret $\tilde{O}(\sqrt{T})$ is typically shown to be optimal. For discontinuous demand functions with finite many valuations, a $\tilde{O}(\sqrt{T})$ regret bound that also depends on the number of different valuations is shown to be optimal (Cesa-Bianchi et al., 2019). For non-parametric demand models, $\tilde{O}(T^{\frac{k+1}{2k+1}})$ regret can be achieved if the demand function is k times differentiable reward function for some integer $k > 0$, and moreover a matching lower bound of $\Theta(T^{\frac{k+1}{2k+1}})$ can be established (Wang et al., 2021). However, the smoothness level k needs to be known prior to the algorithmic design, and it is thus unclear if existing algorithms are able to adapt to different smoothness levels. Our work improves upon Wang et al. (2021) by proposing a smoothness-adaptive dynamic pricing algorithm with the same minimax optimal regret rate and additionally, we extend the integer k to more generally β -smooth for any $\beta \in \mathbb{R}_+$.

In certain applications, consumer or product features, also known as contexts, are available and can be parametrized into the demand valuation (Qiang and Bayati, 2016; Javanmard, 2017; Cohen et al., 2020; Ban and Keskin, 2021; Xu and Wang, 2021). The landscape of regret analysis in contextual cases typically ranges from $\log(T)$ to $\tilde{O}(\sqrt{T})$ depending on different parametric or semiparametric assumptions on demand valuation and market noise. The smoothness level of both the demand function and the noise func-

tion may affect the regret bound, and theoretical results for adaptively learning the smoothness level are not known (Fan et al., 2022; Bu et al., 2022).

Continuum-Armed Bandit Problems Dynamic pricing is closely related to the continuum-armed bandit problem, where the actions are not discrete but rather lie in a continuous space as in the case of the continuous price space. Adaption to Hölder smoothness level β while achieving the minimax regret rate has been considered in continuum-armed bandits as well. It is shown in Locatelli and Carpentier (2018) and Hadiji (2019) that adaptivity for free is generally impossible. Our non-adaptivity result for dynamic pricing shares the same spirit as in the continuum-armed bandit problem but requires different construction of function classes in the arguments. Liu et al. (2021) propose to use a general model approach for bandit problems, but the analysis only applies to the subcase of $0 < \beta \leq 1$. Due to non-adaptivity, additional assumptions are therefore necessary for establishing adaptivity. Specifically, the assumption of self-similarity emerges as a promising candidate because it has been demonstrated to maintain the minimax regret rates in both continuum-armed bandits (Cai and Pu, 2022) and contextual bandits (Gur et al., 2022) scenarios, where only $0 < \beta \leq 1$ is considered in the latter.

Adaptivity in Statistics More broadly, adaptive inference and adaptive estimators have been widely considered in statistics, but less is known if these techniques are suited for regret minimization. While several structural conditions have profound implications in nonparametric regression, such as monotonicity, concavity, as discussed in Cai et al. (2013), introducing any of these assumptions may either significantly diminish the problem’s complexity or do not directly contribute to the learning of the smoothness parameter (Slivkins et al., 2019; Cai and Pu, 2022). Consequently, with any of these structural assumptions at play, the minimax regret operates at the parametric rate, making it agnostic to smoothness variations.

3 PRELIMINARIES

Problem Description We consider the dynamic pricing problem with demand learning over a finite time horizon of length T . At every time period $t = 1, \dots, T$, the seller selects a price $p_t \in [p_{\min}, 1]$, where $0 < p_{\min} < 1$ is a predetermined price lower bound and the price upper bound is normalized to 1 without loss of generality. After the seller sets the price, the customers then arrive and a randomized demand $d_t \in [0, d_{\max}]$ is incurred. The randomized

demand d_t given price is determined by a demand function $f : [p_{\min}, 1] \rightarrow [0, d_{\max}]$ and some random market noise, and the expectation of the randomized demand $\mathbb{E}[d_t | p_t = p] = f(p)$. The noise in demand $d_t - f(p)$ follows a sub-gaussian distribution with respect to some parameters. The revenue collected at time t is $r_t = p_t \cdot d_t$, and the expected revenue given p_t is $p_t \times f(p_t)$.

As is common in previous literature on pricing (Besbes and Zeevi, 2012; Wang et al., 2021; Bu et al., 2022), the Hölder smoothness assumption is used to constrain the volatility of the demand function f in any given region. Throughout the paper, the demand function f is assumed to belong to the Hölder smooth function class $\mathcal{H}(\beta, L)$ for certain $\beta, L > 0$ that are defined as follows.

Definition 1 (Hölder Smooth Function Class). *The Hölder class of functions $\mathcal{H}(\beta, L)$ is defined to be the set of $w(\beta)$ times continuously differentiable functions $g : [p_{\min}, 1] \rightarrow \mathbb{R}$ such that for any $p, p' \in [p_{\min}, 1]$,*

$$\sup_{p \in [p_{\min}, 1]} \left| g^{(k)}(p) \right| \leq L, \forall 0 \leq k \leq w(\beta),$$

$$\left| g^{(w(\beta))}(p) - g^{(w(\beta))}(p') \right| \leq L \cdot |p - p'|^{\beta - w(\beta)},$$

where $w(\beta)$ is the largest integer that is strictly smaller than β .

Policy and Regret An admissible dynamic pricing policy π over T selling periods is a sequence of T random functions $\pi_1, \pi_2, \dots, \pi_T$ such that $\pi_t : (p_1, d_1, \dots, p_{t-1}, d_{t-1}) \mapsto p_t$ is a mapping function that maps the history prior to time t to a price p_t . Since the demand function belongs to $\mathcal{H}(\beta, L)$ and thus continuous over $[p_{\min}, 1]$, there exists some optimal price $p^* \in \arg \max_{p \in [p_{\min}, 1]} \mathbb{E}[r_t | p_t = p]$. Note that here we do not require the optimal price to be unique.

The performance of dynamic pricing policies is evaluated by the cumulative regret defined as follows. For an admissible dynamic pricing policy π over T selling periods, the regret $R^\pi(T)$ over a time horizon T is

$$R^\pi(T) = \mathbb{E}^\pi \left[\sum_{t=1}^T \{p^* f(p^*) - p_t f(p_t)\} \right],$$

where the price sequence $\{p_t\}_{t=1}^T$ is determined by the policy.

Non-Adaptive Pricing If the smoothness parameter β is known, non-adaptive dynamic pricing algorithms can achieve the optimal regret rate, which is called Hölder-Smooth Dynamic Pricing (HSDP) algorithm and presented in Algorithm 1. The algorithm

is designed based on the following idea. We first segment the price interval into many small intervals, and the length of each small bin depends on β , and then we can run local polynomial regression to approximate the true demand function in each small price interval separately. As we formally show later in Lemma 3, this non-adaptive algorithm achieves the optimal regret if it is run with the correct smoothness parameter.

Algorithm 1 Hölder-Smooth Dynamic Pricing (HSDP)

Input: Time horizon T , Hölder smoothness β , minimum price p_{\min} , maximum demand d_{\max} , number of bins N , parameter $L > 0$, optional initial history $\mathcal{D}^{(0)}$;

- 1: Set polynomial degree $k = w(\beta)$;
- 2: Partition $[p_{\min}, 1]$ into N segments of equal lengths, denoted as $\mathbf{I}_j = [a_j, b_j]$ where $a_j = p_{\min} + \frac{(j-1)(1-p_{\min})}{N}$, $b_j = p_{\min} + \frac{j(1-p_{\min})}{N}$ for $j = 1, 2, \dots, N$, and let $\Delta = L \left(\frac{1-p_{\min}}{N} \right)^\beta$;
- 3: Initialize segment history, realized demands and trial numbers $\mathcal{D}_j := \{(p_t, d_t) : p_t \in \mathbf{I}_j\}$, $\tau_j := \sum_{p_t \in \mathcal{D}_j} p_t d_t$, $n_j := |\mathcal{D}_j|$ where $\mathcal{D}_j \subset \mathcal{D}^{(0)}$ for all $1 \leq j \leq N$;
- 4: **for** $t = 1, 2, \dots, T$ **do**
- 5: Compute $CI_j := \left[\Delta + \frac{(3d_{\max}+L)\sqrt{2}}{\sqrt{n_j}} \right] (k+1) \ln(2(k+1)T)$;
- 6: Select $j_t := \arg \max_{1 \leq j \leq N} \frac{\tau_j}{n_j} + CI_j$;
- 7: Let $\delta = \frac{1}{T^2}$, compute $\gamma = L\sqrt{k+1} + \Delta\sqrt{|\mathcal{D}_{j_t}|} + d_{\max}\sqrt{2(k+1)\ln\left(\frac{4(k+1)t}{\delta}\right)} + 2$ and $\Lambda = I_{(k+1) \times (k+1)} + \sum_{(p,d) \in \mathcal{D}_{j_t}} \phi^{(k)}(p)\phi^{(k)}(p)^T$;
- 8: Do local polynomial regression on \mathbf{I}_{j_t} with ridge type penalty and the estimator $\hat{\theta} = \arg \min_{\theta \in \mathbb{R}^{k+1}} \sum_{(p,d) \in \mathcal{D}_{j_t}} |d - \langle \hat{\theta}, \phi^{(k)}(p) \rangle|^2 + \|\theta\|_2^2$;
- 9: Set price $p_t = \arg \max_{p \in \mathbf{I}_{j_t}} p \times \min\{d_{\max}, \langle \hat{\theta}, \phi^{(k)}(p) \rangle + \gamma\sqrt{\phi^{(k)}(p)^T \Lambda^{-1} \phi^{(k)}(p)} + \Delta\}$;
- 10: Observe realized demand $d_t \in [0, d_{\max}]$;
- 11: Update $\tau_j \leftarrow \tau_j + d_t p_t$, $n_j \leftarrow n_j + 1$, $\mathcal{D}_j \leftarrow \mathcal{D}_j \cup \{(p_t, d_t)\}$ for $j = j_t$;
- 12: **end for**

To help illustrate Algorithm 1, we introduce the concept of local polynomial regression, a crucial component of both Algorithm 1 and our smoothness-adaptive dynamic pricing algorithm that will be introduced later. A systematic treatment of local polynomial regression in the statistics literature can be found in the monograph by Fan (2018). Compared to conventional regression methods, the local polynomial regression approach incorporates a scaling process that offers several advantages in terms of flexibility, adaptability, and efficiency. By applying the local polynomial regres-

sion with respect to a carefully chosen support set and focusing on specific small intervals, our method can effectively estimate the mean demand function with greater accuracy, making it suitable for a wide range of applications in dynamic pricing. The scaling process also allows for more efficient computation and model fitting, particularly in situations where data is limited or sparse.

Definition 2 (Local polynomial regression). *Let $\mathbb{O} = \{(p_{(1)}, d_{(1)}), \dots, (p_{(m_0)}, d_{(m_0)})\}$ be a sequence of observations, where $p_{(1)}$ has support $\subset [p_{\min}, 1]$. Our goal is to estimate $\mathbb{E}[d_{(1)}|p_{(1)}]$ with these samples nonparametrically. Let $\mathbf{I} = [a, b] \subset [p_{\min}, 1]$, and let those observations such that $p_{(i)} \in \mathbf{I}$ be $\{\mathbb{O}_{\mathbf{I}} = (p_{(1)}, d_{(1)}), \dots, (p_{(m_0)}, d_{(m_0)})\}$. Then we can estimate $\mathbb{E}[d_{(1)}|p_{(1)}]$ by fitting a polynomial regression on $[a, b]$ with samples in $\mathbb{O}_{\mathbf{I}}$.*

Let $t_m(p) = \left(\frac{1}{2} + \frac{p - \frac{a+b}{2}}{b-a}\right)^m$ and vector $\phi^{(l)}(p) = (t_0(p), t_1(p), \dots, t_l(p))^T$ for some integer l . Define

$$\hat{\theta} = \arg \min_{\theta \in \mathbb{R}^{l+1}} \sum_{j=1}^{m_0} \left(d_{(j)} - \langle \phi^{(l)}(p_{(j)}), \theta \rangle \right)^2.$$

For concreteness, if the minimizer is not unique we define $\hat{\theta} = 0$. The local polynomial regression estimate on \mathbf{I} is given by

$$\hat{f}(p; \mathbb{O}, l, \mathbf{I}) := \langle \phi^{(l)}(p), \hat{\theta} \rangle.$$

By using the local polynomial regression, we can leverage the Hölder smoothness condition to improve the approximation error at each small price interval. Suppose the length of a price interval is ϵ , a constant approximation would lead to approximation error $O(\epsilon)$ if the demand function f is Lipschitz and generally no approximation guarantee without Lipschitz condition. On the other hand, if $f \in \mathcal{H}(\beta, L)$, we can bound the approximation error by $O(\epsilon^\beta)$, which improves $O(\epsilon)$ error for $\beta > 1$, and is also strictly better when $\beta < 1$ and the Lipschitz condition fails to hold. Details of this approximation guarantee are presented in Lemma 12 in Appendix C.1.

Importantly, Algorithm 1 relies on the input of β to construct the number of small intervals N , which is decided as $\lceil T^{\frac{1}{2\beta+1}} \rceil$. The parameter β affects the number of small intervals and therefore the length of each small interval. The underlying reason is that the approximation error of the local polynomial regression step crucially depends on the interval length, which plays an important role when establishing the optimal regret bound $\tilde{O}(T^{\frac{\beta+1}{2\beta+1}})$. It is therefore highly nontrivial, if not impossible, to remove the dependence on β from the design of Algorithm 1.

4 ADAPTIVITY TO UNKNOWN SMOOTHNESS

4.1 Difficulty of Adaption

We show that it is impossible for any policy to achieve adaptivity without additional assumptions. This statement is formalized by establishing Theorem 1 where we consider two different smoothness levels. In Theorem 1, we show that a policy that achieves the optimal regret rate on a smoothness level α could not simultaneously do so on a smoothness level $\beta < \alpha$.

Theorem 1. *It is impossible to achieve adaption without additional assumptions. Fix any two positive Hölder smoothness parameters $\alpha > \beta > 0$, and parameters $L(\alpha), L(\beta) > 0$. Suppose that there is a policy π achieves the optimal regret $\tilde{O}\left(T^{\frac{\alpha+1}{2\alpha+1}}\right)$ over $\mathbb{E}[d|p] = f(p) \in \mathcal{H}(\alpha, L(\alpha))$, then there exists a constant $C > 0$ that is independent of π such that*

$$\sup_{f \in \mathcal{H}(\beta, L(\beta))} R^\pi(T) \geq \Omega\left(T^{\frac{\beta+1}{2\beta+1} + \frac{\beta(\alpha-\beta)}{2(2\beta+1)^2(2\alpha+1)}}\right),$$

which means that it cannot achieve the optimal regret over $\mathbb{E}[d|p] = f(p) \in \mathcal{H}(\beta, L(\beta))$.

The proof of Theorem 1 is accomplished by constructing a single basis function. From this basis function, we then generate demand functions with distinct Hölder smoothness levels of α and β . By comparing the Kullback-Leibler divergence between the resulting probability measures under these different conditions, we establish the existence of a regret gap. The comprehensive proof of Theorem 1 can be found in Appendix A.1.

The negative result in Theorem 1 highlights the difficulty of adaption and therefore necessitates the need for introducing additional conditions. A potential condition should ideally *not only make the adaptivity possible* for a wide range of functions as large as possible *but also not trivialize the pricing problem's complexity*.

4.2 Self-Similarity Condition

We identify the self-similarity condition to enable adaptivity with desirable properties. Before introducing the definition, we need some notation regarding function projections onto the space of polynomial functions. For any positive integer l , let $\text{Poly}(l)$ denote the set of all polynomials of degree less than or equal to l . For any function $g(\cdot)$, we use $\Gamma_l^U g(\cdot)$ to denote the L_2 -projection of the function $g(\cdot)$ onto $\text{Poly}(l)$ over some interval U , which can be computed by the following

minimization

$$\Gamma_l^U g(p) := \min_q \int_U |g(u) - q(u)|^2 du, \\ \text{s.t. } q \in \text{Poly}(l).$$

Definition 3 (Self-Similarity Condition). *A function $g : [a, b] \rightarrow \mathbb{R}$, $[a, b] \subseteq [0, 1]$ is self-similar on $[a, b]$ with parameters $\beta, l \in \mathbb{Z}^+$, $M_1 \in \mathbb{R}_{\geq 0}$, $M_2 \in \mathbb{R}_+$ if for some positive integer $c > M_1$ it holds that*

$$\max_{V \in \mathcal{V}_c} \sup_{p \in V} |\Gamma_l^V g(p) - g(p)| \geq M_2 \cdot 2^{-c\beta},$$

where we define

$$\mathcal{V}_c = \left\{ \left[a + \frac{i}{2^c}, a + \frac{i+1}{2^c} \right] \cap [0, 1], i = 0, 1, \dots, 2^c - 1 \right\}$$

for any positive integer c . We denote the class of self-similar functions by $\mathcal{S}(\beta, l, M_1, M_2)$.

In contrast to Hölder smoothness, the self-similarity condition provides a global lower bound on the approximation error using polynomial regression. This dual nature facilitates the estimation of the smoothness of payoff functions by comparing the approximation on different scales. The self-similarity condition has previously appeared in the literature of nonparametric regression for constructing adaptive confidence intervals (Picard and Tribouley, 2000; Giné and Nickl, 2010).

Example 1 (Example of Self-Similar Functions). *Let f be any function with continuous first-order derivative uniformly bounded by C_1 . We define the function class \mathcal{F} as $\mathcal{F}(f) = \{f : x \mapsto c_0 \cdot x^\beta + f : c_0 \in \mathbb{R}, |c_0| \geq C_1\}$, then all function in $\mathcal{F}(f)$ is self-similar with parameters $\beta, l = 0$ for some constants M_1, M_2 depending on C_1 and C_2 .*

To enable adaptivity, we assume that the self-similar condition holds for the demand function.

Assumption 1. *The demand function f is self-similar with some parameters β, l, M_1, M_2 .*

4.3 Lower Bound

We emphasize that adding this self-similarity condition does not diminish the hardness of the dynamic pricing problem. We validate this statement by showing a worst-case lower bound of *any* policy $\Omega(T^{\frac{\beta+1}{2\beta+1}})$ even in the presence of the self-similarity condition. Notably, this rate matches exactly with the lower bound given in the original problem without self-similarity condition (Wang et al., 2021).

Theorem 2 (Lower Bound). *Self-similarity does not change the minimax regret rate and therefore does not*

lower the problem difficulty for any admissible dynamic pricing policy π . Formally, for any positive parameters $\beta, M_1, L > 0$, there exists a constant $M_2 > 0$ satisfying that

$$\inf_{\pi} \sup_{f \in \mathcal{H}(\beta, L) \cap \mathcal{S}(\beta, w(\beta), M_1, M_2)} R^\pi(T) \geq \Omega(T^{\frac{\beta+1}{2\beta+1}}).$$

Theorem 2 says that there exists a class of non-trivial instances belonging to the self-similar function class such that the worst-case regret of any pricing policy is lower bounded by $\Omega(T^{\frac{\beta+1}{2\beta+1}})$. A main challenge of proving Theorem 2 is therefore constructing such a class of demand functions to establish the lower bound. We now explain our constructions.

Let $u(\cdot) : [0, 1] \rightarrow \mathbb{R}$ be a \mathcal{C}^∞ function with $u(0) = 1, u(1) = 0, u^{(k)}(0) = 0, u^{(k)}(1) = 0, \forall k \in \mathbb{Z}^+$. Let $u_1(x) = \sin^2(\frac{\pi}{2} \cdot u(4x-1))$. $\forall x \in [0, 1]$, let $\sigma(x) = |x - \frac{1}{2}|$ and define

$$g(x) = \begin{cases} c_1 \cdot u_1(\sigma(x)) \cdot \frac{(1-\sigma(x)^\beta)}{2} & \text{if } \frac{1}{4} \leq \sigma(x) \leq \frac{1}{2}, \\ c_1 \cdot \frac{(1-\sigma(x)^\beta)}{2} & \text{if } \sigma(x) < \frac{1}{4}, \end{cases}$$

for some sufficiently small constant $c_1 > 0$, then we have $g \in \mathcal{H}(\beta, L)$. On the other hand, we know that for some constant $M_2 > 0$, $g \in \mathcal{S}(\beta, w(\beta), 0, M_2)$ (Gur et al., 2022, Lemma 1.7). Also, g has a unique maximum point at $x = \frac{1}{2}$ and for any $x \in \{0, 1\}, l \in \{0, 1, \dots, w(\beta)\}$ it holds that $g^{(l)}(x) = 0$.

As a result, we have shown that for each $L, \beta > 0$, there exists a function $g : [0, 1] \rightarrow [0, 1]$ satisfying that $g \in \mathcal{H}(\beta, L) \cap \mathcal{S}(\beta, w(\beta), 0, M_2)$ for some constant $M_2 > 0$; and g has a unique maximizer at $\frac{1}{2}$; and for any $x \in \{0, 1\}, k \in \{0, 1, \dots, w(\beta)\}$ it holds that $g^{(k)}(x) = 0$. The constructed function class, combined with a classical argument with the Kullback–Leibler divergence, leads to the lower bound stated in Theorem 2, and a complete proof of Theorem 2 can be found in Appendix A.4.

5 ALGORITHM AND REGRET ANALYSIS

In this section, we introduce our Smoothness-Adaptive Dynamic Pricing (SADP) algorithm and provide a detailed regret analysis.

5.1 Algorithm Description

We now present our SADP algorithm described in Algorithm 2, which incorporates an efficient smoothness parameter selection phase and is designed to adapt to the unknown Hölder smoothness level.

Algorithm 2 Smoothness-Adaptive Dynamic Pricing (SADP)

Input: Time horizon T , Hölder smoothness range $[\beta_{\min}, \beta_{\max}]$, minimum price p_{\min} , maximum demand d_{\max} , parameter $L > 0$;

- 1: Set local polynomial regression degree $l = w(\beta_{\max})$;
- 2: Set $k_1 = \frac{1}{2\beta_{\max}+2}, k_2 = \frac{1}{4\beta_{\max}+2}, K_1 = 2^{\lfloor k_1 \log_2(T) \rfloor}, K_2 = 2^{\lfloor k_2 \log_2(T) \rfloor}$;
- 3: **for** $i = 1, 2$ **do**
- 4: Set trial time $T_i = T^{\lfloor \frac{1}{2} + k_i \rfloor}$;
- 5: Pull arms T_i times from $U(p_{\min}, 1)$ independently;
- 6: **for** $m = 1, 2, \dots, K_i$ **do**
- 7: Let the samples which fall in $[p_{\min} + \frac{(m-1)(1-p_{\min})}{K_i}, p_{\min} + \frac{m(1-p_{\min})}{K_i}]$ be $\mathcal{O}_{i,m} = \{(p_t, d_t) : p_t \in [p_{\min} + \frac{(m-1)(1-p_{\min})}{K_i}, p_{\min} + \frac{m(1-p_{\min})}{K_i}]\}$;
- 8: Fit local polynomial regression on $[p_{\min} + \frac{(m-1)(1-p_{\min})}{K_i}, p_{\min} + \frac{m(1-p_{\min})}{K_i}]$ with $\mathcal{O}_{i,m}$, construct estimate $\hat{f}_i(p)$ on the interval;
- 9: **end for**
- 10: **end for**
- 11: Let $\hat{\beta} = -\frac{\ln(\max\|\hat{f}_2 - \hat{f}_1\|_\infty)}{\ln(T)} - \frac{\ln(\ln(T))}{\ln(T)}$;
- 12: Set $N = \lceil T^{\frac{1}{2\hat{\beta}+1}} \rceil, \mathcal{D} = \cup_i \cup_m \mathcal{O}_{i,m}$;
- 13: Call HSDP($T - T_1 - T_2, \hat{\beta}, p_{\min}, d_{\max}, N, \Delta, L, \mathcal{D}$)

We provide some intuition behind the design of Algorithm 2. Harnessing the Hölder smoothness assumption, we can employ local polynomial regression to estimate the demand function reasonably well. However, the demand function is not easily approximated by polynomials, due to the inherent self-similarity condition presented by the dual nature of Hölder smoothness. The estimation granularity refers to the number of intervals into which the domain of price is partitioned for better piecewise polynomial approximation.

In Algorithm 2, we employ two distinct levels of granularity to estimate the demand function, indexed by 1 and 2 respectively. For estimation $i \in \{1, 2\}$, the price range is segmented into small intervals of size $(1 - p_{\min})/K_i$, where K_i is a quantity depending on T . The algorithm then allocates T_i time periods to collect price and demand data points and fit with local polynomial regression. The constructed estimates of demand function f are \hat{f}_1 and \hat{f}_2 , which helps us to establish an estimate of the Hölder smoothness parameter as defined in (1).

$$\hat{\beta} = -\frac{\ln(\max\|\hat{f}_2 - \hat{f}_1\|_\infty)}{\ln(T)} - \frac{\ln(\ln(T))}{\ln(T)}. \quad (1)$$

The estimator $\hat{\beta}$ is then fed into the Hölder-smooth dynamic pricing algorithm (Algorithm 1) for the remaining time horizon $T - T_1 - T_2$. As evident from the algorithm design, the accuracy of $\hat{\beta}$ estimation is critical to the regret bound of our smoothness-adaptive dynamic pricing algorithm in Algorithm 2. In the next subsection, we provide a tight confidence interval for the estimator $\hat{\beta}$, which plays a central role in our final regret analysis of the SADP algorithm.

5.2 Accuracy of Estimation

By employing two distinct levels of granularity to estimate the demand function and, in conjunction with the previously established upper and lower bounds of the approximation error, we can prove a confidence interval for the distance between the two estimations $\|\hat{f}_2 - \hat{f}_1\|_\infty$. This distance is directly related to the Hölder smoothness parameter β . Ultimately, we arrive at a reasonably narrow confidence interval for β , which converges rapidly as T increases. Formally, we have the following theorem, which plays a key role in establishing the regret bound.

Theorem 3. *With an upper bound β_{\max} of the smoothness parameter, under the assumptions and settings in the Algorithm 2, for some constant $C > 0$, with probability at least $1 - O\left(e^{-C \ln^2(T)}\right)$,*

$$\hat{\beta} \in \left[\beta - \frac{4(\beta_{\max} + 1) \ln(\ln(T))}{\ln(T)}, \beta \right].$$

Theorem 3 demonstrates the effectiveness of our proposed SADP algorithm in estimating the Hölder smoothness parameter β without prior knowledge. This adaptability, along with the effective smoothness parameter selection phase, enables our algorithm to construct a tight confidence interval for the Hölder smoothness parameter β and achieve a high convergence rate. These characteristics contribute to the desired regret bound in dynamic pricing scenarios, opening up possibilities for the development of more robust and adaptive dynamic pricing algorithms.

In order to prove Theorem 3, we firstly introduce a lemma to characterize the convergence on \hat{f} .

Lemma 1. *Let $\{p_{(i)}, i = 1, 2, \dots, n\}$ be an i.i.d. uniform sample in an interval $\mathbf{I} = [a, b] \subset [p_{\min}, 1]$, and $\mathbb{O}_{\mathbf{I}} = \{(p_{(1)}, d_{(1)}), \dots, (p_{(n)}, d_{(n)})\}$. With the assumptions, suppose sub-gaussian parameter $u_1 \leq \exp(u'_1 \cdot n^v)$ for some positive constants v, u'_1 , polynomial degree $l \geq w(\beta)$. Let*

$$\delta_1 = \left| \mathbb{E} [d_{(1)} | p_{(1)} = p] - \hat{f}(p; \mathbb{O}, l, [a, b]) \right|,$$

$$\delta_2 = \left| \Gamma_l^{\mathbf{I}} \{ \mathbb{E} [d_{(1)} | p_{(1)} = p] \} - \hat{f}(p; \mathbb{O}, l, \mathbf{I}) \right|,$$

Then, there exist positive constants C_1, C_2 such that with probability at least $1 - O\left(e^{-C_2 \ln^2(n)}\right)$, for any $p \in \mathbf{I}$ and $n > C_1$, the following inequality holds:

$$\delta_1 < (b - a)^\beta \ln(n) + \ln^3(n) \cdot n^{-\frac{1}{2}(1-v)}.$$

Also,

$$\delta_2 < \ln^3(n) \cdot n^{-\frac{1}{2}(1-v)}.$$

With Lemma 1 in place, we can now proceed to prove Theorem 3. The proof for this theorem relies on the concentration result stated in Lemma 1 and the construction of the confidence interval based on the two-level granularity approach. By analyzing the relationship between the distance of the estimated demand functions and the Hölder smoothness parameter β , we can show that the estimated $\hat{\beta}$ falls within the stated confidence interval with high probability.

We also present a lemma to mitigate the issue of insufficient sample points falling within certain intervals.

Lemma 2. *Let $\{B_i : i = 1, 2, \dots, n\}$ be i.i.d random variables. Suppose $B_1 \sim \text{Bernoulli}(\frac{1}{m})$. Let $\bar{B} = \frac{\sum_{i=1}^n B_i}{n}$. Then*

$$\mathbb{P}(\bar{B} < \frac{1}{2m}) \leq \exp\left(-\frac{n}{50m}\right).$$

Proof Sketch of Theorem 3. Our first objective is to ascertain an upper bound for the distance between \hat{f}_1 and \hat{f}_2 . Define the interval $\mathbf{I}_{i,m} = \left[p_{\min} + \frac{(m-1)(1-p_{\min})}{K_i}, p_{\min} + \frac{m(1-p_{\min})}{K_i} \right)$. Invoking the first part of Lemma 1 and corroborated by Lemma 2, with a probability of at least $1 - O\left(e^{-C \ln^2(n)}\right)$, we have $\forall p \in \mathbf{I}_{i,m}$,

$$\begin{aligned} \left| f(p) - \hat{f}_i(p) \right| &< K_i^{-\beta} \ln(T) \\ &+ \ln^3(T) \cdot \left(\frac{T_i}{2K_i} \right)^{-\frac{1}{2}(1-v_i)}, \end{aligned} \quad (2)$$

for some sufficiently small constants v_1, v_2 .

Subsequently, utilizing inequality (2), we deduce the following upper bound

$$\begin{aligned} \left\| \hat{f}_2 - \hat{f}_1 \right\|_\infty &\leq (K_1 + K_2)^{-\beta} \ln(T) \\ &+ \ln^3(T) \cdot \left[\left(\frac{T_1}{2K_1} \right)^{-\frac{1}{2}(1-v_1)} \right. \\ &\quad \left. + \left(\frac{T_2}{2K_2} \right)^{-\frac{1}{2}(1-v_2)} \right] \\ &\leq 2^c T^{-\beta k_2} \ln(T), \end{aligned}$$

for a small constant c .

However, to prove the theorem, it's necessary to establish a lower bound for the distance between f_1 and \hat{f}_2 .

Firstly, using inequality (2), we can establish an upper bound for the distance between f and \hat{f}_1 . Furthermore, by invoking the second part of Lemma 1 as well as Lemma 2, with probability at least $1 - O(e^{-C \ln^2(n)})$, $\forall p \in \mathbf{I}_{i,m}$,

$$\left| \Gamma_l^{\mathbf{I}_{i,m}} f(p) - \hat{f}_2(p) \right| \leq \ln^3(T) \cdot \left(\frac{T_2}{2K_2} \right)^{-\frac{1}{2}(1-v_2)}, \quad (3)$$

where v_1, v_2 are sufficiently small.

Given the self-similar properties of f , a lower bound for the distance between f and $\Gamma_l f$ is established:

$$\|f - \Gamma_l f\|_\infty \geq M_2 \cdot K_2^{-\beta}. \quad (4)$$

Subsequently, combining inequalities (2), (3), and (4), a lower bound can be derived as

$$\begin{aligned} \|\hat{f}_2 - \hat{f}_1\|_\infty &\geq M_2 \cdot K_2^{-\beta} - K_1^{-\beta} \ln(T) \\ &\quad - \ln^3(T) \cdot \left[\left(\frac{T_1}{2K_1} \right)^{-\frac{1}{2}(1-v_1)} \right. \\ &\quad \left. + \left(\frac{T_2}{2K_2} \right)^{-\frac{1}{2}(1-v_2)} \right] \\ &\geq \frac{M_2}{2} T^{-\beta k_2}, \end{aligned}$$

To advance the proof, we employ the probability union bound. For some constant $C > 0$, with probability at least $1 - O(e^{-C \ln^2(T)})$, the following holds:

$$\begin{aligned} \hat{\beta} &= -\frac{\ln\left(\max\|\hat{f}_2 - \hat{f}_1\|_\infty\right)}{\ln(T)} - \frac{\ln(\ln(T))}{\ln(T)} \\ &\in \left[\beta - \frac{c \ln(2) + \ln(\ln(T))}{k_2 \ln(T)} \right. \\ &\quad \left. - \frac{\ln(\ln(T))}{\ln(T)}, \beta - \frac{\ln\left(\frac{M_2}{2}\right)}{k_2 \ln(T)} - \frac{\ln(\ln(T))}{\ln(T)} \right] \\ &\subset \left[\beta - \frac{4(\beta_{\max} + 1) \ln(\ln(T))}{\ln(T)}, \beta \right], \end{aligned}$$

which completes the proof. \square

5.3 Regret Analysis

After obtaining an estimation of β , we can provide a more precise bound for the distance between the local polynomial projection and the demand function.

Theorem 4. *Suppose $f \in \mathcal{H}(\beta, L)$, SADP has an estimation of β with $\hat{\beta}$, and is run with $N \geq \lceil T^{\frac{1}{2\hat{\beta}+1}} \rceil, \Delta \leq L \left(\frac{1-p_{\min}}{N} \right)^{\hat{\beta}}$, then with probability $1 - O(T^{-1})$ the cumulative regret of SADP is upper bounded by $\tilde{O}\left(T^{\frac{\beta+1}{2\beta+1}}\right)$.*

We introduce Lemma 3 to bound the regret rate of our non-adaptive algorithm HSDP, whose proof is in Appendix C.2.

Lemma 3. *If HSDP is run with $\hat{\beta} \leq \beta$ and other conditions stays the same as Theorem 4, then with probability $1 - O(T^{-1})$, the cumulative regret is upper bounded by $O\left(T^{\frac{\beta+1}{2\beta+1}}\right)$.*

Proof Sketch of Theorem 4. Considering the event $A^* : \{\hat{\beta} \in [\beta - \frac{4(\beta_{\max}+1)\ln(\ln(T))}{\ln(T)}, \beta]\}$, and by theorem 3, we know that $\mathbb{P}(A^*) \geq 1 - O(e^{-C \ln^2(T)})$. Under event A^* , $\hat{\beta}$ converges to β with rate $O\left(\frac{\ln(\ln(T))}{\ln(T)}\right)$, we have $O\left(T^{\frac{\hat{\beta}+1}{2\hat{\beta}+1}} - T^{\frac{\beta+1}{2\beta+1}}\right) \leq O\left(T^{\frac{\beta+1}{2\beta+1}} \cdot T^{\frac{\beta-\hat{\beta}}{(2\beta+1)(\beta+1)}}\right) \leq \tilde{O}\left(T^{\frac{\beta+1}{2\beta+1}}\right)$. Considering lemma 3, we can derive the regret bound for HSDP under event A^* :

$$\sum_{t=1}^{T-T_1-T_2} [p^* f(p^*) - \hat{p}_t f(\hat{p}_t)] \leq \tilde{O}\left(T^{\frac{\beta+1}{2\beta+1}}\right).$$

For the adaptive part of SADP, note that $T_1, T_2 \leq T^{\frac{\beta+1}{2\beta+1}}$, which means that the regret is bounded by $O\left(T^{\frac{\beta+1}{2\beta+1}}\right)$. Applying the union bound with event A^* , we can derive that with probability $1 - O(T^{-1})$,

$$\begin{aligned} R^\pi(T) &= \mathbb{E} \left[\sum_{t'=1}^{T_1+T_2} \{p^* f(p^*) - p_{t'} f(p_{t'})\} \right] \\ &\quad + \mathbb{E} \left[\sum_{t=1}^{T-T_1-T_2} \{p^* f(p^*) - p_t f(p_t)\} \right] \\ &\leq \tilde{O}\left(T^{\frac{\beta+1}{2\beta+1}}\right). \end{aligned} \quad \square$$

Theorem 4 highlights the effectiveness of the SADP algorithm in achieving the desired regret bound under the specified conditions. By estimating the Hölder smoothness parameter β and generalizing the non-adaptive dynamic-pricing algorithm with non-integer Hölder smoothness parameter, our algorithm is capable of maintaining a high level of performance in dynamic pricing scenarios. We note that by employing the doubling trick the same regret can be achieved without the knowledge of T .

6 NUMERICAL EXPERIMENT

We simulate the performance of our policies SADP with β_{\max} is set to 3.0, and HSDP($\tilde{\beta}$) that is initiated by some misspecified smoothness parameter $\tilde{\beta} \in \{0.5, 1, \dots, 2.5, 3.0\}$. Our simulation code can be found at <https://github.com/liamyzq/Smoothness-Adaptive-Dynamic-Pricing>.

In our numerical experiment, we set $p_{\min} = 1, d_{\max} = 1$, and rewards that are Gaussian with standard deviation $\sigma = 0.05$. We then consider a demand function based on the setting in Gur et al. (2022) with a Hölder smoothness parameter of $\beta_1 = 0.8$. We set $T = 2 \times 10^4$, and this choice of time periods is at a similar scale as in Gur et al. (2022). Specifically, we let $\phi(p) = \mathbf{1}\{|p| \leq 1\}(1 - p)$, and the demand function is defined as

$$f_1(p) = \begin{cases} \frac{1+L_1(\frac{1}{2})^{\beta_1}-L_1p^{\beta_1}}{4}, & \text{if } p_{\min} \leq p \leq \frac{1}{2}, \\ \frac{1}{4} + C \sum_{j=1}^m (-1)^j (2M)^{-\beta_1}, & \text{if } \frac{1}{2} < p \leq 1, \end{cases}$$

where $L_1 = C = 1, \alpha = 0.01, \tau = 0.8, M = \frac{1}{16} \left[\frac{1}{4} \left(\frac{2 \ln 2}{T} \right)^{\frac{-\tau}{2\tau+1}} \right]^{\frac{1}{\beta_1}}$, $m = \lfloor M^{1-\alpha\beta_1} \rfloor$, and $\xi_j = \phi(2M[(2-2p) - \frac{j+\frac{1}{2}}{M}])$ for each $j = 1, 2, \dots, m$.

We utilize the concept of relative regret to present the results of our numerical experiments, defined as $R_{relative}^{\pi}(T) = \frac{R^{\pi}(T)}{T \cdot p^* f(p^*)}$. The experiments are conducted repeatedly, and Figure 1 displays the mean relative regrets obtained from 30 repeated experiments under our setting. Our numerical findings indicate that under the setting, as $\tilde{\beta}$ gets larger compared to β , the cost of smoothness misspecification incurred by the misspecified HSDP($\tilde{\beta}$) policy increases and dominates the cost of adaptation incurred by the SADP policy. Therefore, the value of adaptivity is even larger in the setting. In particular, we note that HSDP(β) with true smoothness β (or the nearest β) is not guaranteed to minimize regret under every function in the class, and this setting is an example.

7 CONCLUSION

Motivated by the challenge of unknown smoothness levels in applications, we develop a smoothness-adaptive dynamic pricing algorithm under self-similarity conditions. To make dynamic pricing algorithms, it is very desirable to remove the parameter dependence of algorithms. Moving forward, it is promising to explore whether our approach and the notion of self-similarity can be generalized to other dynamic pricing problems such as feature-based dynamic pricing or more broadly other online learning problems.




Figure 1: The comparison of mean relative regret over 30 experiments.

To further improve adaptivity, it is of interest to consider other parameters that are implicitly used in the algorithm design.

Acknowledgements

The authors are grateful to the area chair and three anonymous reviewers for helpful and constructive feedback.

References

- Abbasi-Yadkori, Y., Pal, D., and Szepesvari, C. (2012). Online-to-confidence-set conversions and application to sparse stochastic bandits. In *Artificial Intelligence and Statistics*, pages 1–9. PMLR.
- Ban, G.-Y. and Keskin, N. B. (2021). Personalized dynamic pricing with machine learning: High-dimensional features and heterogeneous elasticity. *Management Science*, 67(9):5549–5568.
- Besbes, O. and Zeevi, A. (2009). Dynamic pricing without knowing the demand function: Risk bounds and near-optimal algorithms. *Operations research*, 57(6):1407–1420.
- Besbes, O. and Zeevi, A. (2012). Blind network revenue management. *Operations research*, 60(6):1537–1550.
- Broder, J. and Rusmevichientong, P. (2012). Dynamic pricing under a general parametric choice model. *Operations Research*, 60(4):965–980.
- Bu, J., Simchi-Levi, D., and Wang, C. (2022). Context-based dynamic pricing with partially linear demand model. *Advances in Neural Information Processing Systems*, 35:23780–23791.
- Cai, T. T., Low, M. G., and Xia, Y. (2013). Adaptive confidence intervals for regression functions un-

- der shape constraints. *The Annals of Statistics*, 41(2):722–750.
- Cai, T. T. and Pu, H. (2022). Stochastic continuum-armed bandits with additive models: Minimax regrets and adaptive algorithm. *The Annals of Statistics*, 50(4):2179–2204.
- Cesa-Bianchi, N., Cesari, T., and Perchet, V. (2019). Dynamic pricing with finitely many unknown valuations. In *Algorithmic Learning Theory*, pages 247–273. PMLR.
- Chen, N. and Gallego, G. (2022). A primal–dual learning algorithm for personalized dynamic pricing with an inventory constraint. *Mathematics of Operations Research*, 47(4):2585–2613.
- Cohen, M. C., Lobel, I., and Paes Leme, R. (2020). Feature-based dynamic pricing. *Management Science*, 66(11):4921–4943.
- Den Boer, A. V. (2015). Dynamic pricing and learning: historical origins, current research, and new directions. *Surveys in operations research and management science*, 20(1):1–18.
- Fan, J. (2018). *Local polynomial modelling and its applications: monographs on statistics and applied probability 66*. Routledge.
- Fan, J., Guo, Y., and Yu, M. (2022). Policy optimization using semiparametric models for dynamic pricing. *Journal of the American Statistical Association*, pages 1–29.
- Giné, E. and Nickl, R. (2010). Confidence bands in density estimation. *Annals of statistics*, 38(2):1122–1170.
- Gur, Y., Momeni, A., and Wager, S. (2022). Smoothness-adaptive contextual bandits. *Operations Research*, 70(6):3198–3216.
- Hadji, H. (2019). Polynomial cost of adaptation for x-armed bandits. *Advances in Neural Information Processing Systems*, 32.
- Javanmard, A. (2017). Perishability of data: dynamic pricing under varying-coefficient models. *The Journal of Machine Learning Research*, 18(1):1714–1744.
- Keskin, N. B. and Zeevi, A. (2014). Dynamic pricing with an unknown demand model: Asymptotically optimal semi-myopic policies. *Operations research*, 62(5):1142–1167.
- Kleinberg, R. and Leighton, T. (2003). The value of knowing a demand curve: Bounds on regret for on-line posted-price auctions. In *44th Annual IEEE Symposium on Foundations of Computer Science, 2003. Proceedings.*, pages 594–605. IEEE.
- Liu, Y., Wang, Y., and Singh, A. (2021). Smooth bandit optimization: generalization to holder space. In *International Conference on Artificial Intelligence and Statistics*, pages 2206–2214. PMLR.
- Locatelli, A. and Carpentier, A. (2018). Adaptivity to smoothness in x-armed bandits. In *Conference on Learning Theory*, pages 1463–1492. PMLR.
- Picard, D. and Tribouley, K. (2000). Adaptive confidence interval for pointwise curve estimation. *The Annals of Statistics*, 28(1):298–335.
- Qiang, S. and Bayati, M. (2016). Dynamic pricing with demand covariates. *arXiv preprint arXiv:1604.07463*.
- Slivkins, A. et al. (2019). Introduction to multi-armed bandits. *Foundations and Trends® in Machine Learning*, 12(1-2):1–286.
- Stewart, G. W. (1977). On the perturbation of pseudo-inverses, projections and linear least squares problems. *SIAM review*, 19(4):634–662.
- Tropp, J. A. (2012). User-friendly tail bounds for sums of random matrices. *Foundations of computational mathematics*, 12:389–434.
- Wang, Y., Chen, B., and Simchi-Levi, D. (2021). Multimodal dynamic pricing. *Management Science*, 67(10):6136–6152.
- Xu, J. and Wang, Y.-X. (2021). Logarithmic regret in feature-based dynamic pricing. *Advances in Neural Information Processing Systems*, 34:13898–13910.

Checklist

1. For all models and algorithms presented, check if you include:
 - (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. [Yes]
 - (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. [Yes]
 - (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [Not Applicable]
2. For any theoretical claim, check if you include:
 - (a) Statements of the full set of assumptions of all theoretical results. [Yes]
 - (b) Complete proofs of all theoretical results. [Yes]
 - (c) Clear explanations of any assumptions. [Yes]
3. For all figures and tables that present empirical results, check if you include:

- (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). [Yes]
 - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [Yes]
 - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Yes]
 - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
- (a) Citations of the creator If your work uses existing assets. [Not Applicable]
 - (b) The license information of the assets, if applicable. [Not Applicable]
 - (c) New assets either in the supplemental material or as a URL, if applicable. [Not Applicable]
 - (d) Information about consent from data providers/curators. [Not Applicable]
 - (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [Not Applicable]
5. If you used crowdsourcing or conducted research with human subjects, check if you include:
- (a) The full text of instructions given to participants and screenshots. [Not Applicable]
 - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [Not Applicable]
 - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [Not Applicable]

Smoothness-Adaptive Dynamic Pricing with Nonparametric Demand Learning

Supplementary Materials

A PROOFS ON SELF-SIMILARITY

A.1 Proof of Theorem 1

Proof of Theorem 1. Let $d \sim \mathcal{N}(f(p), \sigma^2)$ be a normal random variable, σ^2 here is small enough such that $\mathcal{N}(0, \sigma^2)$ is sub-Gaussian under different demand settings. Also let $\epsilon_1 = \frac{\alpha - \beta}{2(2\alpha + 1)(2\beta + 1)}$ and $\epsilon_2 = 2\epsilon_1, \epsilon_3 = \frac{\epsilon_1}{2}$. Denote

$$\psi(p) = \begin{cases} ce^{-\frac{1}{(p-p_{\min})^{(1-p)}}} & \text{if } p_{\min} < p < 1 \\ 0, & \text{otherwise.} \end{cases}$$

when $c < \frac{1}{2}$ is small enough, $\psi \in \mathcal{H}(\alpha, L'(\alpha)) \cap \mathcal{H}(\beta, L'(\beta))$. Define a counting random variable $Z_{k,m} = \sum_{t=1}^T \mathbb{1} \left\{ p_t \in \left[p_{\min} + \frac{m(1-p_{\min})}{k}, p_{\min} + \frac{(m+1)(1-p_{\min})}{k} \right) \right\}$ for any positive integer k . Let $a \propto T^{\frac{1}{2\alpha+1} - \frac{\epsilon_2}{\alpha}}, b \propto T^{\frac{1-\epsilon_1}{2\beta+1}}$. Define index set $S_{a,b} = \{0, 1, \dots, b-1\} \cap (\frac{b}{a}, +\infty)$ and let m_0 be the index in the index set such that $\mathbb{E}[Z_{k,m_0}]$ is the smallest. Define functions $\psi_a(p) = a^{-\alpha} \psi(a(p-p_{\min}))$ and $\psi_b(p) = b^{-\beta} \psi(b(p-p_{\min}) - m_0)$. And let $g_1(p) = \frac{1}{p} \left[\frac{1}{2} + \psi_a(p) \right], g_2(p) = \frac{1}{p} \left[\frac{1}{2} + \psi_a(p) + \psi_b(p) \right]$, by Lemma 4, $g_1 \in \mathcal{H}(\alpha, L(\alpha)), g_2 \in \mathcal{H}(\beta, L(\beta))$.

Denote the probability measure determined by \mathcal{A} and $f = g_i$ by \mathbb{P}_i for $i = 1, 2$. Let $\mathbb{E}_i[Z]$ be the expectation of random variable $Z(p_1, \dots, p_T, d_1, \dots, d_T)$ if the probability measure is \mathbb{P}_i .

If $f = g_1$, we have

$$\begin{aligned} R^\pi(T) &= \mathbb{E}_1 \left[\sum_{t=1}^T \{p^* \cdot g_1(p^*) - p_t \cdot d_t\} \right] \\ &\geq \mathbb{E} \left[\sum_{t=1}^T \{ \psi_a(p^*) - \mathcal{N}(\psi_a(p_t), \sigma^2) \} \mathbb{1} \left\{ p_t \notin \left[p_{\min}, p_{\min} + \frac{1-p_{\min}}{a} \right) \right\} \right] \\ &\geq \mathbb{E} \left[\sum_{t=1}^T \psi_a(p^*) \mathbb{1} \left\{ p_t \notin \left[p_{\min}, p_{\min} + \frac{1-p_{\min}}{a} \right) \right\} \right] \\ &= \Omega(a^{-\alpha}) \mathbb{E}_1[T - Z_{a,0}]. \end{aligned} \tag{A.1}$$

By the conditions from the theorem, we have $R_T(\mathcal{A}; g_1) \leq O\left(T^{\frac{\alpha+1}{2\alpha+1} + \epsilon_3}\right)$. So we have

$$\mathbb{E}_1[T - Z_{a,0}] \leq O\left(a^\alpha \cdot T^{\frac{\alpha+1}{2\alpha+1} + \epsilon_3}\right).$$

And by the definition of m_0 and notice that $b > a$, we can derive that

$$\mathbb{E}_1[Z_{b,m_0}] \leq \frac{\mathbb{E}_1[T - Z_{a,0}]}{|S_{a,b}|} \leq O\left(a^\alpha \cdot b^{-1} \cdot T^{\frac{\alpha+1}{2\alpha+1} + \epsilon_3}\right). \tag{A.2}$$

Then we can decompose the KL-divergence between $\mathbb{P}_1, \mathbb{P}_2$ as

$$KL(\mathbb{P}_1 || \mathbb{P}_2) = \mathbb{E}_1 \left[\prod_{t=1}^T \frac{\mathbb{P}_1(d_t | p_t)}{\mathbb{P}_2(d_t | p_t)} \right] = \sum_{t=1}^T \mathbb{E}_1 [KL(\mathcal{N}(g_1(p_t), \sigma^2) || \mathcal{N}(g_2(p_t), \sigma^2))].$$

By inequality (A.2) and Lemma 5 we can obtain

$$\begin{aligned} KL(\mathbb{P}_1 || \mathbb{P}_2) &= \frac{1}{2\sigma^2} \mathbb{E}_1 \left[\sum_{t=1}^T \{g_1(p_t) - g_2(p_t)\}^2 \right] \\ &\leq \frac{b^{-2\beta}}{2\sigma^2 p_{\min}} \sum_{t=1}^T \mathbb{E}_1 \left[\mathbb{1} \left\{ p_t \in \left[p_{\min} + \frac{m_0(1-p_{\min})}{b}, p_{\min} + \frac{(m_0+1)(1-p_{\min})}{b} \right] \right\}^2 \right] \\ &\leq \frac{b^{-2\beta}}{2\sigma^2 p_{\min}} \mathbb{E}_1 [Z_{b, m_0}] \\ &\leq O \left(b^{-2\beta-1} \cdot a^\alpha \cdot b^{-1} \cdot T^{\frac{\alpha+1}{2\alpha+1} + \epsilon_3} \right) \\ &= O \left(T^{\epsilon_1-1} \cdot T^{\frac{\alpha+1}{2\alpha+1} + \epsilon_3} \cdot T^{\frac{\alpha}{2\alpha+1} - \epsilon_2} \right) \\ &= O \left(T^{\epsilon_1 - \epsilon_2 + \epsilon_3} \right) \\ &= O \left(T^{-\epsilon_3} \right) \\ &= o(1). \end{aligned}$$

Let $A = \{Z_{b, m_0} > b^{-1} \cdot a^\alpha \cdot T^{\frac{\alpha+1}{2\alpha+1} + \epsilon_3} \cdot \ln(T)\}$. Inequality A.2 implies that $\mathbb{P}_1(A) = o(1)$. By Lemma 6, we have $|\mathbb{P}_1(A) - \mathbb{P}_2(A)| = o(1)$, so we can derive that $\mathbb{P}_2(A) = o(1)$.

Since

$$b^{-1} \cdot a^\alpha \cdot T^{\frac{\alpha+1}{2\alpha+1} + \epsilon_3} \cdot \ln(T) = T^{\frac{\epsilon_1-1}{2\beta+1}} \cdot T^{\frac{\alpha+1}{2\alpha+1} + \epsilon_3} \cdot T^{\frac{\alpha}{2\alpha+1} - \epsilon_2} = o(T),$$

On A^c , we have $T - Z_{b, m_0} > \frac{T}{2}$. Note that $\max_{p \in [p_{\min}, p_{\min} + \frac{1-p_{\min}}{a}]} \psi_a(p) \ll \max_{p \in [p_{\min} + \frac{m_0(1-p_{\min})}{b}, p_{\min} + \frac{(m_0+1)(1-p_{\min})}{b}]} \psi_b(p)$, then with the similar procedure in inequality A.1, if $f = g_2$, we have

$$\begin{aligned} R_T(\mathcal{A}; g_2) &\geq \Omega \left(b^{-\beta} \cdot \mathbb{E}_2 [T - Z_{b, m_0}] \right) \\ &\geq \Omega \left(b^{-\beta} \cdot \mathbb{E}_2 [(T - Z_{b, m_0}) \mathbb{1}_{A^c}] \right) \\ &\geq \Omega \left(b^{-\beta} \cdot \frac{T}{2} \cdot \mathbb{P}_2(A^c) \right) \\ &\geq \Omega \left(T^{1 + \frac{\beta(\epsilon_1-1)}{2\beta+1}} \right) \\ &\geq \Omega \left(T^{\frac{\beta+1}{2\beta+1} + \frac{\beta(\alpha-\beta)}{2(2\beta+1)^2(2\alpha+1)}} \right), \end{aligned}$$

which completes the proof. □

A.2 Technical Lemmas for Theorem 1

Lemma 4. Suppose $r(p) \in \mathcal{H}(\beta, L')$, $p \in [p_{\min}, 1]$, $0 < p_{\min} < 1$, then $f(p) = \frac{r(p)}{p} \in \mathcal{H}(\beta, L)$ for some constant L .

Proof of Lemma 4. This is a basic property of Hölder class of functions, whose proof follows directly from Lemma 1.10 of Gur et al. (2022) and the fact that the function $p \mapsto 1/p$ is Hölder smooth of any levels when restricted to the interval $[p_{\min}, 1]$. □

Then we introduce three lemmas about KL-divergence which are standard results in the literature and therefore we omit the proofs.

Lemma 5. Let $d_1 \sim \mathcal{N}(\mu_1, \sigma^2)$, $d_2 \sim \mathcal{N}(\mu_2, \sigma^2)$, then the KL-divergence between d_1 and d_2 is $\frac{(\mu_1 - \mu_2)^2}{2\sigma^2}$.

Lemma 6. Let $\mathbb{P}_1, \mathbb{P}_2$ be two probability measures on the same σ -algebra, then for any event A on this σ -algebra, we have

$$KL(\mathbb{P}_1 || \mathbb{P}_2) \geq 2(\mathbb{P}_1(A) - \mathbb{P}_2(A))^2.$$

Lemma 7. For $\mathbb{P}_1, \mathbb{P}_2$ defined as above, in terms of the total variation norm $\|\cdot\|_{TV}$, we have

$$\|\mathbb{P}_1 - \mathbb{P}_2\|_{TV} \leq \sqrt{2KL(\mathbb{P}_1 || \mathbb{P}_2)}.$$

A.3 Proof of Proposition 1 and Lemma 8 - 9

The results in this section will be used to prove Theorem 2.

A.3.1 Proof of Proposition 1

Proposition 1. For each $L, \beta > 0$, there exists a function $g : [0, 1] \rightarrow [0, 1]$ satisfying that $g \in \mathcal{H}(\beta, L) \cap \mathcal{S}(\beta, w(\beta), 0, M_2)$ for some constant $M_2 > 0$; and g has a unique maximizer at $\frac{1}{2}$; and for any $x \in \{0, 1\}, k \in \{0, 1, \dots, w(\beta)\}$ it holds that $g^{(k)}(x) = 0$.

Proof. Let $u(\cdot) : [0, 1] \rightarrow \mathbb{R}$ be a \mathcal{C}^∞ function with $u(0) = 1, u(1) = 0, u^{(k)}(0) = 0, u^{(k)}(1) = 0, \forall k \in \mathbb{Z}^+$. Let $u_1(x) = \sin^2\left(\frac{\pi}{2} \cdot u(4x - 1)\right)$. $\forall x \in [0, 1]$, let $\sigma(x) = |x - \frac{1}{2}|$ and define

$$g(x) = \begin{cases} c_1 \cdot u_1(\sigma(x)) \cdot \frac{(1 - \sigma(x)^\beta)}{2} & \text{if } \frac{1}{4} \leq \sigma(x) \leq \frac{1}{2}, \\ c_1 \cdot \frac{(1 - \sigma(x)^\beta)}{2} & \text{if } \sigma(x) < \frac{1}{4} \end{cases},$$

for some sufficiently small constant $c_1 > 0$, then we have $g \in \mathcal{H}(\beta, L)$, and following the proof procedure as Lemma 1.7 of Gur et al. (2022), for some constant $M_2 > 0$, $g \in \mathcal{S}(\beta, w(\beta), 0, M_2)$. Also, g has a unique maximum point at $x = \frac{1}{2}$ and for any $x \in \{0, 1\}, l \in \{0, 1, \dots, w(\beta)\}$ it holds that $g^{(l)}(x) = 0$. \square

A.3.2 Proof of Lemma 8

Lemma 8. The worst-case regret of algorithm \mathcal{A} over time period T can be lower bounded as

$$\sup_{f \in \mathcal{H}(\beta, L)} R_T(\mathcal{A}; f, \mathcal{N}(0, \sigma^2)) \geq \Omega\left(\epsilon_T^\beta\right) \cdot \max_{1 \leq j \leq J} (T - \mathbb{E}_j[T_j]),$$

where $f \in \mathcal{S}(\beta, w(\beta), 0, M_2)$ on $[p_{\min}, 1]$ for some constant M_2 .

Proof. For $\forall j \in \{1, 2, \dots, J\}$,

$$\begin{aligned} \sup_{f \in \mathcal{H}(\beta, L)} R_T(\mathcal{A}; f, \mathcal{N}(0, \sigma^2)) &\geq \mathbb{E}_j \left[\sum_{i=1}^T \{p^* f_j(p^*) - p_i f_j(p_i)\} \right] \\ &\geq \Omega\left(\epsilon_T^\beta\right) \cdot \mathbb{E}_j \left[\sum_{i=1}^T \mathbf{1}\{p_i \notin \mathbf{I}_j\} \right] \\ &\geq \Omega\left(\epsilon_T^\beta\right) \cdot (T - \mathbb{E}_j[T_j]), \end{aligned}$$

which completes the proof. \square

A.3.3 Proof of Lemma 9

Lemma 9. For fixed $j \in \{1, 2, \dots, J\}$, we have

$$|\mathbb{E}_0 [T_j] - \mathbb{E}_j [T_j]| \leq O\left(T\epsilon_T^\beta\right) \cdot \sqrt{\mathbb{E}_0 [T_j]}.$$

Proof. Because f_0 and f_j only differs on \mathbf{I}_j , we have that

$$KL(\mathbb{P}_0 || \mathbb{P}_j) \leq \mathbb{E}_0 [T_j] \cdot \sup_{p \in \mathbf{I}_j} KL(\mathcal{N}(f_0(p), \sigma^2) || \mathcal{N}(f_j(p), \sigma^2)). \quad (\text{A.3})$$

Then by Lemma 6, we have

$$\sup_{p \in \mathbf{I}_j} KL(\mathcal{N}(f_0(p), \sigma^2) || \mathcal{N}(f_j(p), \sigma^2)) \leq O\left(\epsilon_T^{2\beta}\right), \quad (\text{A.4})$$

and by Lemma 7, inequalities A.3,A.4 we have

$$\|\mathbb{P}_0 - \mathbb{P}_j\|_{TV} \leq O\left(\epsilon_T^\beta\right) \cdot \sqrt{\mathbb{E}_0 [T_j]}.$$

Subsequently,

$$|\mathbb{E}_0 [T_j] - \mathbb{E}_j [T_j]| \leq \sum_{t=1}^T \{t \cdot |\mathbb{P}_0 (T_j = t) - \mathbb{P}_j (T_j = t)|\} \leq T \cdot \|\mathbb{P}_0 - \mathbb{P}_j\|_{TV} \leq O\left(T\epsilon_T^\beta\right) \cdot \sqrt{\mathbb{E}_0 [T_j]},$$

which completes the proof. □

A.4 Proof of Theorem 2

Proof of Theorem 2. Let β denote the true Hölder smoothness parameter here.

Firstly, we introduce a proposition constructing the reward function we need.

Define number of intervals $J = \lceil \epsilon_T^{-1} \rceil$ for ϵ_T defined as $cT^{-\frac{1}{2\beta+1}}$ where c is some sufficiently small constant depending only on β . Let $\mathbf{I}_j = [a_j, b_j]$, $a_j = p_{\min} + \frac{(j-1)(1-p_{\min})}{J}$, $b_j = p_{\min} + \frac{j(1-p_{\min})}{J}$, for $j = 1, 2, \dots, J$. Then define J different demand functions f_1, f_2, \dots, f_J , let

$$f_j(p) = \begin{cases} \frac{1}{2p} & \text{if } p \notin \mathbf{I}_j \\ \frac{1}{2p} + \frac{1}{p}\epsilon_T^\beta g\left(\frac{p-a_j}{\epsilon_T}\right) & \text{if } p \in \mathbf{I}_j \end{cases},$$

where $f_j \in \mathcal{H}(\beta, L) \cap \mathcal{S}(\beta, w(\beta), 0, M_2)$. Define also $f_0(p) \equiv \frac{1}{2p}$, $p \in [p_{\min}, 1]$.

Denote the probability measure determined by algorithm \mathcal{A} and $f = f_j$ by \mathbb{P}_j for $j = 0, 1, \dots, J$. Let $\mathbb{E}_i [Z]$ be the expectation of random variable $Z(p_1, \dots, p_T, d_1, \dots, d_T)$ if the probability measure is \mathbb{P}_j . Let demand $d_i \sim \mathcal{N}(f(p_i), \sigma^2)$, where σ^2 is small enough that $\mathcal{N}(f(p_i), \sigma^2)$ is sub-Gaussian with parameters u_1, u_2 .

Then we can upper bound the difference between $\mathbb{E}_0 [T_j]$ and $\mathbb{E}_j [T_j]$ by the properties of KL-divergence.

Let $j^* = \arg \min_{j \in \{1, 2, \dots, J\}} \mathbb{E}_0 [T_j]$, it is obvious that $\mathbb{E}_0 [T_j] \leq \frac{T}{j} \leq T\epsilon_T$. Then by Lemma 9, for some sufficiently small c , we can obtain

$$\mathbb{E}_j [T_j] \leq O\left(T\epsilon_T^\beta\right) \cdot \sqrt{\mathbb{E}_0 [T_j]} \leq O\left(T \cdot c^\beta \cdot T^{-\frac{\beta}{2\beta+1}}\right) \cdot \sqrt{T \cdot c \cdot T^{-\frac{1}{2\beta+1}}} \leq \frac{T}{2}.$$

Consequently, by Lemma 8

$$\begin{aligned} \sup_{f \in \mathcal{H}(\beta, L)} R_T(\mathcal{A}; f, \mathcal{N}(0, \sigma^2)) &\geq \Omega\left(\epsilon_T^\beta\right) \cdot \max_{1 \leq j \leq J} (T - \mathbb{E}_j [T_j]) \\ &\geq \Omega\left(c^\beta \cdot T^{-\frac{\beta}{2\beta+1}}\right) \cdot \frac{T}{2} \\ &\geq \Omega\left(T^{\frac{\beta+1}{2\beta+1}}\right), \end{aligned}$$

which completes the proof. \square

B PROOFS OF CONSTRUCTED CONFIDENCE INTERVAL

B.1 Proof of Lemma 1

To prove Lemma 1, we first introduce the following two lemmas.

Lemma 10. *Suppose A, B are two $n_0 \times n_0$ symmetric matrices. If $\|A - B\| \leq \frac{\lambda_{\min}(B)}{2}$, then*

$$\|A^{-1} - B^{-1}\| \leq 2(1 + \sqrt{5}) \|A - B\| \lambda_{\min}^{-2}(B)$$

The proof of Lemma 1 directly follows Theorem 3.3 of Stewart (1977).

Lemma 11. *Suppose $d_{(1)}$ is a sub-Gaussian random variable with parameter u_1, u_2 , then we can upper bound $\mathbb{E}[|d_{(1)}|]$ with $\frac{(2\sqrt{\ln(2u_1)}+1)}{\sqrt{u_2}}$.*

Proof of Lemma 11. Let $K = 2\frac{\sqrt{\ln(2u_1)}}{u_2}$, and we have

$$\begin{aligned} \mathbb{E}[|d_{(1)}|] &= \int_0^\infty \mathbb{P}(|d_{(1)}| \geq x) dx = \int_0^K \mathbb{P}(|d_{(1)}| \geq x) dx + \int_K^\infty \mathbb{P}(|d_{(1)}| \geq x) dx \\ &\leq K + \int_K^\infty u_1 \cdot e^{-u_2 x^2} dx \\ &\leq K + \frac{u_1}{2u_2 K} e^{-u_2 K^2} \\ &\leq \frac{(2\sqrt{\ln(2u_1)} + 1)}{\sqrt{u_2}}. \end{aligned}$$

\square

Proof of Lemma 1. Without loss of generality, in this part, we do a translation for $d_{(1)}|p_{(1)}$ to make its expectation 0. Let \mathbf{P}_n be a $n \times (l+1)$ matrix with its m th row $\phi^{(l)}(p_{(m)})^T$ for every m and $\mathbf{d}_n = (d_{(1)}, \dots, d_{(n)})^T$. By least square regression, we obtain

$$\hat{\theta} = (\mathbf{P}_n^T \mathbf{P}_n)^{-1} \mathbf{P}_n^T \mathbf{d}_n = \left(\frac{\mathbf{P}_n^T \mathbf{P}_n}{n} \right)^{-1} \frac{\mathbf{P}_n^T \mathbf{d}_n}{n}$$

Define

$$\theta_0 = \left(\mathbb{E} \left[\phi^{(l)}(p_{(1)}) \phi^{(l)}(p_{(1)})^T \right] \right)^{-1} \mathbb{E} \left[\phi^{(l)}(p_{(1)})^T d_{(1)} \right].$$

The goal of this lemma is to obtain the convergence properties of θ , here we firstly prove the convergence of $\left(\frac{\mathbf{P}_n^T \mathbf{P}_n}{n} \right)^{-1}$. Let $U_1 = \phi^{(l)}(p_{(1)}) \phi^{(l)}(p_{(1)})^T - \mathbb{E} \left[\phi^{(l)}(p_{(1)}) \phi^{(l)}(p_{(1)})^T \right]$, by Bernstein inequality (Tropp (2012), Theorem 1.6),

$$\mathbb{P} \left(\left\| \frac{\mathbf{P}_n^T \mathbf{P}_n}{n} - \mathbb{E} \left[\phi^{(l)}(p_{(1)}) \cdot \phi^{(l)}(p_{(1)})^T \right] \right\| \geq w \right) \leq (2l+2) \cdot \exp \left(\frac{-nw^2}{R_1 + \frac{wR_2}{3}} \right)$$

where $R_1 = \max\{\|\mathbb{E}[U_1 \cdot U_1^T]\|, \|\mathbb{E}[U_1^T \cdot U_1]\|\}$, $R_2 = \sup \|U_1\|$. For each m , we have $t_m(p_{(1)}) \in [0, 1]$ with probability 1, so $R_1, R_2 \leq O(1)$. Let $w = \frac{\ln(n)}{\sqrt{n}}$, then we have

$$\mathbb{P} \left(\left\| \frac{\mathbf{P}_n^T \mathbf{P}_n}{n} - \mathbb{E} \left[\phi^{(l)}(p_{(1)}) \cdot \phi^{(l)}(p_{(1)})^T \right] \right\| \geq \frac{\ln(n)}{\sqrt{n}} \right) \leq O \left(e^{-C \ln^2(n)} \right)$$

for some constant $C > 0$.

Let $V(p_{(1)}) = \mathbb{E} \left[\phi^{(l)}(p_{(1)}) \phi^{(l)}(p_{(1)})^T \right]$, we can prove there exists a constant M_0 which satisfies $\lambda_{\min}(V(p_{(1)})) \geq M_0 > 0$ where λ_{\min} denotes the least eigenvalue of the matrix. We have

$$\lambda_{\min}(V(p_{(1)})) = \inf_{u \in \mathbb{R}^{l+1}, \|u\|_2=1} u^T V(p_{(1)}) u = \inf_{u \in \mathbb{R}^{l+1}, \|u\|_2=1} \mathbb{E} \left\| \phi^{(l)}(p_{(1)})^T u \right\|^2 \geq M_0 > 0$$

Based on Lemma 10, we have

$$\mathbb{P} \left(\left\| \left(\frac{\mathbf{P}_n^T \mathbf{P}_n}{n} \right)^{-1} - \left(\mathbb{E} \left[\phi^{(l)}(p_{(1)}) \cdot \phi^{(l)}(p_{(1)})^T \right] \right)^{-1} \right\| \geq C_1 \frac{\ln(n)}{\sqrt{n}} \right) \leq O \left(e^{-C_2 \ln^2(n)} \right) \quad (\text{B.1})$$

Then we prove the convergence of $\frac{\mathbf{P}_n^T \mathbf{d}_n}{n}$. Let $\mathbf{1}(|d_{(m)}| > M)$ be a n -dimensional vector whose m th element is $\mathbf{1}(|d_{(m)}| > M)$, $F_M(\mathbf{d}_n)$ be a n -dimensional vector whose m th element is $d_{(m)} \mathbf{1}(|d_{(m)}| \leq M)$. Let $U_2 = \phi^{(l)}(p_{(1)}) d_{(1)} \cdot \mathbf{1}(|d_{(1)}| \leq M) - \mathbb{E} \left[\phi^{(l)}(p_{(1)}) d_{(1)} \cdot \mathbf{1}(|d_{(1)}| \leq M) \right]$, and by Bernstein Inequality (Tropp (2012), Theorem 1.6),

$$\mathbb{P} \left(\left\| \frac{\mathbf{P}_n^T F_M(\mathbf{d}_n)}{n} - \mathbb{E} \left[\phi^{(l)}(p_{(1)}) d_{(1)} \cdot \mathbf{1}(|d_{(1)}| \leq M) \right] \right\| \geq w \right) \leq (2l+2) \cdot \exp \left(\frac{-nw^2}{R_3 + \frac{wR_4}{3}} \right)$$

where $R_3 = \max\{\|\mathbb{E}[U_2 \cdot U_2^T]\|, \|\mathbb{E}[U_2^T \cdot U_2]\|\}$, $R_4 = \sup \|U_2\|$. Each element of U_2 is upper bounded by $O(M)$ so we have $R_3 \leq O(M^2)$, $R_4 \leq O(M)$. Let $w = \frac{M \ln(n)}{\sqrt{n}}$, then we have

$$\mathbb{P} \left(\left\| \frac{\mathbf{P}_n^T F_M(\mathbf{d}_n)}{n} - \mathbb{E} \left[\phi^{(l)}(p_{(1)}) d_{(1)} \cdot \mathbf{1}(|d_{(1)}| \leq M) \right] \right\| \geq w \right) \leq O \left(e^{-C \ln^2(n)} \right) \quad (\text{B.2})$$

for some constant $C > 0$. And by the sub-gaussian assumption, we have

$$\mathbb{P}(\exists i \in \{1, \dots, n\}, |d_{(i)}| > M) \leq u_1 n \cdot e^{-u_2 M^2}$$

Taking $M = \sqrt{\frac{\ln(u_1 n)}{u_2}} \ln(n)$, we have

$$\mathbb{P}(\exists i \in \{1, \dots, n\}, |d_{(i)}| > M) \leq e^{-C \ln^2(n)} \quad (\text{B.3})$$

and C is a constant that independent of u_1, u_2 . And also by the we can deduce that

$$\begin{aligned} & \left| \mathbb{E} \left[\phi^{(l)}(p_{(1)}) d_{(1)} \cdot \mathbf{1}(|d_{(1)}| \leq M) \right] - \mathbb{E} \left[\phi^{(l)}(p_{(1)}) d_{(1)} \right] \right| \\ &= O \left(\int_M^{+\infty} p dF_d \right) = O \left(\int_M^{+\infty} (p - M) dF_d + M \mathbb{P}(|d_{(1)}| \geq M) \right) \\ &= O \left(\int_M^{+\infty} \mathbb{P}(|d| \geq p) dp + M \mathbb{P}(|d_{(1)}| \geq M) \right) \\ &= O \left(\int_M^{+\infty} u_1 \cdot e^{-u_2 p^2} dp + u_1 \cdot M \cdot e^{-u_2 M^2} \right) \leq O \left(\frac{u_1}{2Mu_2} e^{-u_2 M^2} + u_1 M \cdot e^{-u_2 M^2} \right) \\ &= O \left(\left(\frac{u_1}{2\sqrt{\ln(u_1 n)} u_2 \ln(n)} + u_1 \sqrt{\frac{\ln(u_1 n)}{u_2}} \ln(n) \right) \cdot \exp(-\ln(u_1 n) \ln^2(n)) \right) \end{aligned}$$

which leads to

$$\left| \mathbb{E} \left[\phi^{(l)}(p_{(1)}) d_{(1)} \cdot \mathbf{1}(|d_{(1)}| \leq M) \right] - \mathbb{E} \left[\phi^{(l)}(p_{(1)}) d_{(1)} \right] \right| \leq O \left(\frac{1}{n} \right) \quad (\text{B.4})$$

Combining the three inequalities B.2, B.3 and B.4, we have the following

$$\mathbb{P} \left(\left\| \frac{\mathbf{P}_n^T \mathbf{d}_n}{n} - \mathbb{E} \left[\phi^{(l)}(p_{(1)}) d_{(1)} \right] \right\| \geq C \cdot \frac{\sqrt{\ln(u_1 \cdot n) \ln^2(n)}}{\sqrt{n}} \right) \leq C_0 e^{-C_2 \ln^2(n)} \quad (\text{B.5})$$

where C, C_0, C_2 are constants depending on u_2 and l .

And by Lemma 11, we know that

$$\left\| \mathbb{E} \left[\phi^{(l)}(p_{(1)}) d_{(1)} \right] \right\| \leq O \left(\frac{\sqrt{\ln(u_1)} + 1}{u_2} \right). \quad (\text{B.6})$$

Then combining inequalities B.1, B.5 and B.6, recall the definition of θ_0 and $\hat{\theta}$, with probability at least $1 - O(e^{-C_2 \ln^2(n)})$ for some constants C_2 depending on u_2 and l , we have

$$\begin{aligned} \|\hat{\theta} - \theta_0\| &= \left\| \left(\frac{\mathbf{P}_n^T \mathbf{P}_n}{n} \right)^{-1} \frac{\mathbf{P}_n^T \mathbf{d}_n}{n} - \left(\mathbb{E} \left[\phi^{(l)}(p_{(1)}) \phi^{(l)}(p_{(1)})^T \right] \right)^{-1} \mathbb{E} \left[\phi^{(l)}(p_{(1)})^T d_{(1)} \right] \right\| \\ &= \left\| \left(\frac{\mathbf{P}_n^T \mathbf{P}_n}{n} \right)^{-1} \left(\frac{\mathbf{P}_n^T \mathbf{d}_n}{n} - \mathbb{E} \left[\phi^{(l)}(p_{(1)}) d_{(1)} \right] \right) \right. \\ &\quad \left. + \left(\left(\frac{\mathbf{P}_n^T \mathbf{P}_n}{n} \right)^{-1} - \mathbb{E} \left[\phi^{(l)}(p_{(1)}) \phi^{(l)}(p_{(1)})^T \right] \right)^{-1} \mathbb{E} \left[\phi^{(l)}(p_{(1)}) d_{(1)} \right] \right\| \\ &\leq \left\| \left(\frac{\mathbf{P}_n^T \mathbf{P}_n}{n} \right)^{-1} \left(\frac{\mathbf{P}_n^T \mathbf{d}_n}{n} - \mathbb{E} \left[\phi^{(l)}(p_{(1)}) d_{(1)} \right] \right) \right\| \\ &\quad + \left\| \left(\left(\frac{\mathbf{P}_n^T \mathbf{P}_n}{n} \right)^{-1} - \mathbb{E} \left[\phi^{(l)}(p_{(1)}) \phi^{(l)}(p_{(1)})^T \right] \right)^{-1} \mathbb{E} \left[\phi^{(l)}(p_{(1)}) d_{(1)} \right] \right\| \\ &\leq O \left(\frac{\sqrt{\ln(u_1 \cdot n) \ln^2(n)}}{\sqrt{n}} \right) + O \left(\frac{(\sqrt{\ln(u_1)} + 1) \ln(n)}{\sqrt{nu_2}} \right) \\ &< \ln^3(n) \cdot n^{-\frac{1}{2}(1-v)}, \end{aligned} \quad (\text{B.7})$$

for n larger than some constant C_1 depending on u'_1, u_2, l .

Note that $\Gamma_l^{\mathbf{I}} \{ \mathbb{E} [d_{(1)} | p_{(1)} = p] \} = \langle \phi^{(l)}(p), \theta_0 \rangle$ and $\hat{f}(p; \mathbb{O}, l, \mathbf{I}) = \langle \phi^{(l)}(p), \hat{\theta} \rangle$, with $\phi^{(l)}(p_{(1)}) \leq O(1)$, the second part of the lemma is proved.

In order to prove the first part of the lemma, we can show that $|\mathbb{E} [d_{(1)} | p_{(1)} = p] - \langle \phi^{(l)}(p), \theta_0 \rangle| = O((b-a)^\beta)$. By the Holder assumption and Taylor expansion, there exists an $l+1$ dimensional vector θ_1 such that $|\mathbb{E} [d_{(1)} | p_{(1)} = p] - \langle \phi^{(l)}(p), \theta_1 \rangle| = O((b-a)^\beta), \forall p \in \mathbf{I}$. So we have

$$\begin{aligned} \|\theta_0 - \theta_1\| &= \left\| \left(\mathbb{E} \left[\phi^{(l)}(p_{(1)}) \phi^{(l)}(p_{(1)})^T \right] \right)^{-1} \mathbb{E} \left[\phi^{(l)}(p_{(1)})^T d_{(1)} \right] \right. \\ &\quad \left. - \left(\mathbb{E} \left[\phi^{(l)}(p_{(1)}) \phi^{(l)}(p_{(1)})^T \right] \right)^{-1} \mathbb{E} \left[\phi^{(l)}(p_{(1)}) \phi^{(l)}(p_{(1)})^T \right] \theta_1 \right\| \\ &= \left\| \left(\mathbb{E} \left[\phi^{(l)}(p_{(1)}) \phi^{(l)}(p_{(1)})^T \right] \right)^{-1} \mathbb{E} \left[\phi^{(l)}(p_{(1)}) (d_{(1)} - \langle \phi^{(l)}(p_{(1)}), \theta_1 \rangle) \right] \right\| = O((b-a)^\beta). \end{aligned}$$

Then by inequality (B.7), we can deduce that with probability at least $1 - O(e^{-C_2 \ln^2(n)})$

$$\|\hat{\theta} - \theta_1\| \leq \ln^3(n) \cdot n^{-\frac{1}{2}(1-v)} + (b-a)^\beta \ln(n).$$

Note that with $\phi^{(l)}(p_{(1)}) \leq O(1)$, $\left| \langle \phi^{(l)}(p), \hat{\theta} \rangle - \langle \phi^{(l)}(p), \theta_1 \rangle \right| \leq O\left(\left\| \hat{\theta} - \theta_1 \right\|\right)$, $\forall p \in \mathbf{I}, n > C_1$, with probability at least $1 - O\left(e^{-C_2 \ln^2(n)}\right)$, the following inequality holds

$$\left| \mathbb{E} [d_{(1)} | p_{(1)} = p] - \langle \phi^{(l)}(p_{(1)}), \hat{\theta} \rangle \right| \leq \ln^3(n) \cdot n^{-\frac{1}{2}(1-v)} + (b-a)^\beta \ln(n).$$

Therefore the first part of the lemma is proved. □

B.2 Proof of Theorem 3

Proof. We first define an event $A = \{\exists i \in \{1, 2\}, m \in \{1, 2, \dots, K_i\}, \text{s.t. } |\mathbb{O}_{i,m}| < \frac{T_i}{2K_i}\}$, by Lemma 2, we have

$$\mathbb{P}(A) \leq T \left(\exp\left(-\frac{T_1}{50K_1}\right) + \exp\left(-\frac{T_2}{50K_2}\right) \right),$$

By conditioning on A^c , we can guarantee the number of samples in each interval. Next, we aim to establish an upper bound for the distance between the distance between \hat{f}_1 and \hat{f}_2 . Let $\mathbf{I}_{i,m} = \left[p_{\min} + \frac{(m-1)(1-p_{\min})}{K_i}, p_{\min} + \frac{m(1-p_{\min})}{K_i} \right)$. Invoking the first part of Lemma 1, with probability at least $1 - O\left(e^{-C \ln^2(n)}\right)$, $\forall p \in \mathbf{I}_{i,m}$, the following inequality holds:

$$\left| f(p) - \hat{f}_i(p) \right| < K_i^{-\beta} \ln(T) + \ln^3(T) \cdot \left(\frac{T_i}{2K_i} \right)^{-\frac{1}{2}(1-v_i)}, \quad (\text{B.8})$$

for some sufficiently small constants v_1, v_2 .

Define the event

$$B = \{\exists i \in \{1, 2\}, m \in \{1, 2, \dots, K_i\}, p \in \mathbf{I}_{i,m}, \text{s.t. inequality 2 does not hold}\}, \quad (\text{B.9})$$

Applying the union bound, we find that

$$\mathbb{P}(B|A^c) \leq O\left((K_1 + K_2) e^{-C \ln^2(n)}\right),$$

for some constant $C > 0$.

Then, conditioning on $A^c \cap B^c$, we can, by inequality 2, derive an upper bound as follows:

$$\left\| \hat{f}_2 - \hat{f}_1 \right\|_\infty < (K_1 + K_2)^{-\beta} \ln(T) + \ln^3(T) \cdot \left[\left(\frac{T_1}{2K_1} \right)^{-\frac{1}{2}(1-v_1)} + \left(\frac{T_2}{2K_2} \right)^{-\frac{1}{2}(1-v_2)} \right], \quad (\text{B.10})$$

However, to prove the theorem, it's necessary to establish a lower bound for the distance between \hat{f}_1 and \hat{f}_2 . In the ensuing discussion, we aim to provide this lower bound.

Firstly, using inequality 2, we can establish an upper bound for the distance between f and \hat{f}_1 . This, in turn, aids in deducing an upper bound for the distance between $\Gamma_l f$ and \hat{f}_2 .

Furthermore, by invoking the second part of Lemma 1 as well as Lemma 2, with probability at least $1 - O\left(e^{-C \ln^2(n)}\right)$, $\forall p \in \mathbf{I}_{i,m}$,

$$\left| \Gamma_l^{\mathbf{I}_{i,m}} f(p) - \hat{f}_2(p) \right| \leq \ln^3(T) \cdot \left(\frac{T_2}{2K_2} \right)^{-\frac{1}{2}(1-v_2)}, \quad (\text{B.11})$$

where v_1, v_2 are sufficiently small.

Define the event

$$D = \{\exists p \in [p_{\min}, 1), \text{s.t. inequality 3 does not hold}\},$$

And applying the union bound we have

$$\mathbb{P}(D|A^c) \leq O\left((K_1 + K_2)e^{-C \ln^2(n)}\right),$$

Given the self-similar property of f , we can establish a lower bound for the distance between f and $\Gamma_l f$.

$$\|f - \Gamma_l f\|_\infty \geq M_2 \cdot K_2^{-\beta}, \quad (\text{B.12})$$

Subsequently, conditioning on $A^c \cap B^c \cap D^c$, and using inequalities 2, 3, and (4), we can derive a lower bound as follows:

$$\|\hat{f}_2 - \hat{f}_1\|_\infty > M_2 \cdot K_2^{-\beta} - K_1^{-\beta} \ln(T) - \ln^3(T) \cdot \left[\left(\frac{T_1}{2K_1}\right)^{-\frac{1}{2}(1-v_1)} + \left(\frac{T_2}{2K_2}\right)^{-\frac{1}{2}(1-v_2)} \right], \quad (\text{B.13})$$

Now, let's attempt to simplify the upper bound in inequality B.10:

$$\begin{aligned} \|\hat{f}_2 - \hat{f}_1\|_\infty &\leq (K_1 + K_2)^{-\beta} \ln(T) + \ln^3(T) \cdot \left[\left(\frac{T_1}{2K_1}\right)^{-\frac{1}{2}(1-v_1)} + \left(\frac{T_2}{2K_2}\right)^{-\frac{1}{2}(1-v_2)} \right] \\ &\leq 2^c T^{-\beta k_2} \ln(T), \end{aligned}$$

for some small constant c .

Similarly, we can simplify the lower bound in inequality B.13:

$$\begin{aligned} \|\hat{f}_2 - \hat{f}_1\|_\infty &\geq M_2 \cdot K_2^{-\beta} - K_1^{-\beta} \ln(T) - \ln^3(T) \cdot \left[\left(\frac{T_1}{2K_1}\right)^{-\frac{1}{2}(1-v_1)} + \left(\frac{T_2}{2K_2}\right)^{-\frac{1}{2}(1-v_2)} \right] \\ &\geq \frac{M_2}{2} T^{-\beta k_2}, \end{aligned}$$

Thus, on the event $A^c \cap B^c \cap D^c$, we have

$$\begin{aligned} \hat{\beta} &= -\frac{\ln\left(\max\|\hat{f}_2 - \hat{f}_1\|_\infty\right)}{\ln(T)} - \frac{\ln(\ln(T))}{\ln(T)} \\ &\in \left[\beta - \frac{c \ln(2) + \ln(\ln(T))}{k_2 \ln(T)} - \frac{\ln(\ln(T))}{\ln(T)}, \beta - \frac{\ln\left(\frac{M_2}{2}\right)}{k_2 \ln(T)} - \frac{\ln(\ln(T))}{\ln(T)} \right] \\ &\subset \left[\beta - \frac{4(\beta_{\max} + 1) \ln(\ln(T))}{\ln(T)}, \beta \right], \end{aligned}$$

Simultaneously we have

$$\mathbb{P}(A \cup B \cup D) = O\left(e^{-C \ln^2(n)}\right),$$

for some constant $C > 0$ which completes the proof. □

C PROOFS OF REGRET UPPER BOUNDS

C.1 Proof of Lemma 12

Lemma 12. *Suppose $f \in \mathcal{H}(\beta, L)$ and let $\mathbf{I} = [a, b] \subset [p_{\min}, 1]$ be an arbitrary interval whose length is $|\mathbf{I}|$. For estimation $\hat{\beta} \leq \beta$, there exists a polynomial with degree $k = w(\hat{\beta})$: $P_{\mathbf{I}}(p) = \Gamma_{k, \mathbf{I}} f(p) = \sum_{m=0}^k a_m \left(\frac{1}{2} + \frac{p-a+\frac{1}{2}}{b-a}\right)^m$ satisfying $|a_m| m! \leq L, \forall m \leq k$, such that*

$$\sup_{p \in \mathbf{I}} |f(p) - P_{\mathbf{I}}(p)| \leq L(b-a)^{\hat{\beta}}.$$

Proof. Firstly, let $a_m = \frac{f^{(m)}(a)}{m!} (b-a)^m$

$$\begin{aligned} P_{\mathbf{I}}(p) &= \sum_{m=0}^k a_m \left(\frac{1}{2} + \frac{p - \frac{a+b}{2}}{b-a} \right)^m \\ &= \sum_{m=0}^k \frac{f^{(m)}(a)}{m!} (p-a)^m \end{aligned}$$

By Taylor expansion with Lagrangian remainders, $\forall p \in \mathbf{I}, \exists \tilde{p} \in \mathbf{I}$ such that

$$f(p) = \sum_{m=0}^{k-1} \frac{f^{(m)}(a)}{m!} (p-a)^m + \frac{f^{(k)}(\tilde{p})}{k!} (p-a)^k$$

With $\hat{\beta} \leq \beta$, we then have that

$$\begin{aligned} |f(p) - P_{\mathbf{I}}(p)| &= \frac{|f^{(k)}(\tilde{p}) - f^{(k)}(a)|}{k!} (p-a)^k \\ &= \frac{|f^{(w(\hat{\beta}))}(\tilde{p}) - f^{(w(\hat{\beta}))}(a)|}{w(\hat{\beta})!} (p-a)^{w(\hat{\beta})} \\ &\leq \frac{L|p-a|^{\hat{\beta}-w(\hat{\beta})}}{w(\hat{\beta})!} (p-a)^{w(\hat{\beta})} \\ &= \frac{L|b-a|^{\hat{\beta}}}{k!} \\ &\leq L|b-a|^{\hat{\beta}}. \end{aligned}$$

□

C.2 Proof of Lemma 3

Proof. We first state two important lemmas, Lemma 13 and Lemma 14, whose proofs will be included later in this section.

Lemma 13. *Suppose $f \in \mathcal{H}(\beta, L)$ and let $\mathbf{I} = [a, b] \subset [p_{\min}, 1]$. If HSDP is invoked with $\Delta \geq L(b-a)^{\hat{\beta}}$ and outputs \hat{p} , then with probability $1 - \delta$ it holds that*

$$\max_{p \in \mathbf{I}} pf(p) - \hat{p}f(\hat{p}) \leq 2 \min\{d_{\max}, \gamma \sqrt{\phi^{(k)}(p)^T \Lambda^{-1} \phi^{(k)}(p) + \Delta}\},$$

where $\gamma = L\sqrt{k+1} + \Delta\sqrt{|\mathcal{D}|} + d_{\max} \sqrt{2(k+1) \ln \left(\frac{4(k+1)t}{\delta} \right) + 2}$.

Then we can use Lemma 13 to prove Lemma 14.

Lemma 14. *Keep the same setting in Lemma 13 and let $\hat{p}_1, \dots, \hat{p}_t$ be the output prices of t consecutive calls on \mathbf{I} . Then with probability $1 - O(T^{-1})$ it holds that*

$$\frac{1}{t} \sum_{i=1}^t \left[\max_{p \in \mathbf{I}} pf(p) - \hat{p}_i f(\hat{p}_i) \right] \leq \left[\Delta + \frac{(3d_{\max} + L)\sqrt{2}}{\sqrt{t}} \right] (k+1) \ln(2(k+1)T),$$

Finally, with Lemma 14 proved, we can do the UCB analysis for the Theorem.

For all $1 \leq j \leq N$, we define $r^*(\mathbf{I}_j) = \max_{p \in \mathbf{I}_j} pf(p)$. Invoking Lemma 14, by concentration inequalities, with probability $1 - O(T^{-1})$, it holds uniformly for all j that

$$r^*(\mathbf{I}_j) \leq \frac{\tau_j}{n_j} + CI_j \leq r^*(\mathbf{I}_j) + 2CI_j,$$

Let T_j be the total number of time periods that we invoke HSDP in the j th interval, and we invoke j th interval at time t . Also, denote $j^* = \arg \max_j r^*(\mathbf{I}_j)$. Note that $\hat{\beta}$ is strictly less than β , then we can still derive a bound for the regret each round using UCB analysis

$$\begin{aligned} r^*(\mathbf{I}_{j^*}) - r^*(\mathbf{I}_{j_t}) &\leq \left(\frac{\tau_{j^*}}{n_{j^*}} + CI_{j^*} \right) - \left(\frac{\tau_{j_t}}{n_{j_t}} + CI_{j_t} \right) + \left(\frac{\tau_{j_t}}{n_{j_t}} + CI_{j_t} \right) - r^*(\mathbf{I}_{j_t}) \\ &\leq \left(\frac{\tau_{j_t}}{n_{j_t}} + CI_{j_t} \right) - r^*(\mathbf{I}_{j_t}) \\ &\leq 2CI_{j_t}, \end{aligned}$$

And

$$\begin{aligned} \sum_{t=1}^{T-T_1-T_2} CI_{j_t} &\leq (k+1) \ln(2(k+1)T) \left[\Delta T + (3d_{\max} + L) \sum_{j=1}^N \sum_{i=1}^{T_j} \sqrt{\frac{2}{i}} \right] \\ &\leq (k+1) \ln(2(k+1)T) \left[\Delta T + (9d_{\max} + 3L) \sum_{j=1}^N \sqrt{T_j} \right] \\ &\leq (k+1) \ln(2(k+1)T) \left[\Delta T + (9d_{\max} + 3L) \sqrt{N \sum_{j=1}^N T_j} \right] \\ &\leq (k+1) \ln(2(k+1)T) (9d_{\max} + 4L) T^{\frac{\hat{\beta}+1}{2\hat{\beta}+1}} \\ &\leq \tilde{O}(T^{\frac{\hat{\beta}+1}{2\hat{\beta}+1}}). \end{aligned}$$

□

C.3 Proof of Lemma 13 and 14

C.3.1 Proof of Lemma 13

Proof. The (p, d) pairs in the history are labeled as $\{(p_i, d_i)\}_{i=1}^t$ in chronological order. And we can show that $d_i = f(p_i) + \xi_i = P_{\mathbf{1}}(p_i) + \xi_i + \beta_i$, where $\{\xi_i\}_{i=1}^t$ are *i.i.d* sub-gaussian random variables with zero mean and $|\beta_i| \leq \Delta$ with probability 1. Use vectors and matrices to denote them we have $\mathbf{p} = (p_i)_{i=1}^t$, $\mathbf{d} = (d_i)_{i=1}^t$, $\boldsymbol{\xi} = (\xi_i)_{i=1}^t$, $\boldsymbol{\beta} = (\beta_i)_{i=1}^t$ and $\mathbf{P} = \left(\phi^{(k)}(p_i) \right)_{i=1}^t \in \mathbb{R}^{t \times (k+1)}$. And the ridge estimator $\hat{\theta}$ can be written as $\hat{\theta} = \Lambda^{-1} \mathbf{P}^T \mathbf{d} = (\mathbf{P}^T \mathbf{P} + I)^{-1} \mathbf{P}^T \mathbf{d}$, plug in $\mathbf{d} = \mathbf{P} \theta^* + \boldsymbol{\xi} + \boldsymbol{\beta}$ with θ^* is the real coefficient of the expansion, we have

$$\hat{\theta} - \theta^* = -\Lambda^{-1} \theta^* + \Lambda^{-1} \mathbf{P}^T (\boldsymbol{\xi} + \boldsymbol{\beta}),$$

Multiplying $(\hat{\theta} - \theta^*) \Lambda$ on both sides and it leads to

$$(\hat{\theta} - \theta^*)^T \Lambda (\hat{\theta} - \theta^*) = -(\hat{\theta} - \theta^*)^T \theta^* + (\hat{\theta} - \theta^*)^T \mathbf{P}^T (\boldsymbol{\xi} + \boldsymbol{\beta}) \quad (\text{C.1})$$

$$= \sum_{i=1}^t (\xi_i + \beta_i) \langle \phi^{(k)}(p_i), \hat{\theta} - \theta^* \rangle - \langle \hat{\theta} - \theta^*, \theta^* \rangle, \quad (\text{C.2})$$

Note that

$$\left| \sum_{i=1}^t \beta_i \langle \phi^{(k)}(p_i), \hat{\theta} - \theta^* \rangle - \langle \hat{\theta} - \theta^*, \theta^* \rangle \right| \leq \sqrt{\sum_{i=1}^t \beta_i^2} \sqrt{\sum_{i=1}^t |\langle \hat{\theta} - \theta^*, \theta^* \rangle|^2} \leq \Delta \sqrt{t} \cdot \sqrt{(\hat{\theta} - \theta^*)^T \Lambda (\hat{\theta} - \theta^*)},$$

Plug the above inequality into equation (C.1), then dividing $\sqrt{(\hat{\theta} - \theta^*)^T \Lambda (\hat{\theta} - \theta^*)}$ from both sides, also noting that $\Lambda \succeq I$ which makes $\sqrt{(\hat{\theta} - \theta^*)^T \Lambda (\hat{\theta} - \theta^*)} \geq \|\hat{\theta} - \theta^*\|_2$, we obtain

$$\sqrt{(\hat{\theta} - \theta^*)^T \Lambda (\hat{\theta} - \theta^*)} \leq \|\theta^*\|_2 + \Delta\sqrt{t} + \sup_{z \in \Phi_\Lambda} |G_t(z)|, \quad (\text{C.3})$$

where

$$\Phi_\Lambda = \{z \in \mathbb{R}^{k+1} : z^T \Lambda z \leq 1\}, G_t(z) = \sum_{i=1}^t \xi_i \langle \phi^{(k)}(p_i), z \rangle.$$

Recall the definition of θ^* and the Hölder class assumption, we have $\|\theta^*\|_2 \leq L\sqrt{k+1}$. In order to bound $G_t(z)$, we introduce and prove Lemma 15:

Lemma 15. Fix k, t and a probability δ , with probability $1 - \delta$ it holds uniformly for all Λ defined above that

$$\sup_{z \in \Phi_\Lambda} |G_t(z)| \leq d_{\max} \sqrt{2(k+1) \ln \left(\frac{4(k+1)t}{\delta} \right)} + 2$$

Proof of Lemma 15. By definition we have $\|\phi^{(k)}(p_i)\|_2 \leq \sqrt{k+1}$. Let $\epsilon > 0$ be a small parameter. Denote $\|\cdot\|_\Lambda = \sqrt{(\cdot)^T \Lambda (\cdot)}$ as the Λ -norm of a vector, and $\mathbb{B}(r, \|\cdot\|) = \{z \in \mathbb{R}^{k+1} : \|z\| \leq r\}$ as a ball. Let $\mathcal{U} \subseteq \mathbb{B}(1, \|\cdot\|_2)$ be a ϵ -covering of $\mathbb{B}(1, \|\cdot\|_2)$ which means that $\sup_{z \in \mathbb{B}(1, \|\cdot\|_2)} \min_{z' \in \mathcal{U}} \|z - z'\|_2 \leq \epsilon$. Fix arbitrary $z \in \mathcal{U}$, for $|\xi_i| \leq d_{\max}$ with probability 1, by Hoffding's inequality we know that for any $\delta \in (0, 1)$,

$$\mathbb{P} \left(|G_t(z)| \leq d_{\max} \sqrt{2 \ln \left(\frac{2}{\delta} \right)} \|z\|_\Lambda \right) \geq 1 - \delta,$$

Since $\Lambda \succeq I$, we know that $\Phi_\Lambda = \mathbb{B}(1, \|\cdot\|_\Lambda) \subseteq \mathbb{B}(1, \|\cdot\|_2)$ and therefore \mathcal{U} is also a ϵ -covering of Φ_Λ , and it is easy to verify that there exists \mathcal{U} with $\ln(|\mathcal{U}|) \leq (k+1) \ln \left(\frac{2}{\epsilon} \right)$. Applying union bound we have with probability $1 - \delta$,

$$\sup_{z \in \mathcal{U} \cap \Phi_\Lambda} |G_t(z)| \leq \sqrt{2 \ln \left(\frac{2|\mathcal{U}|}{\delta} \right)} \leq d_{\max} \sqrt{2 \ln \left(\frac{2|\mathcal{U}|}{\delta} \right)} \leq d_{\max} \sqrt{2(k+1) \ln \left(\frac{4}{\epsilon} \right) + 2 \ln \left(\frac{1}{\delta} \right)},$$

Considering the covering, by $\|\phi^{(k)}(p_i)\|_2 \leq \sqrt{k+1}$, we have that $\|\Lambda\|_{\text{op}} \leq 1 + (k+1)t \leq 2(k+1)t$, then we have

$$\sup_{z \in \Phi_\Lambda} |G_t(z)| \leq d_{\max} \sqrt{2(k+1) \ln \left(\frac{4}{\epsilon} \right) + 2 \ln \left(\frac{1}{\delta} \right)} + 2(k+1)t\epsilon$$

By setting $\epsilon = \frac{1}{(k+1)t}$ we complete the proof of Lemma 15. □

Then back to the proof of Lemma 13. With inequality C.3, invoking Lemma 15, with probability $1 - \delta$ we have

$$\sqrt{(\hat{\theta} - \theta^*)^T \Lambda (\hat{\theta} - \theta^*)} \leq L\sqrt{k+1} + \Delta\sqrt{t} + d_{\max} \sqrt{2(k+1) \ln \left(\frac{4}{\epsilon} \right) + 2 \ln \left(\frac{1}{\delta} \right)} + 2,$$

And $\forall p \in \mathbf{I}$, let $\hat{f}(p) = \langle \phi^{(k)}(p), \hat{\theta} \rangle$, we can obtain

$$\begin{aligned} |\hat{f}(p) - f(p)| &\leq |\hat{f}(p) - P_{\mathbf{I}}(p)| + |P_{\mathbf{I}}(p) - f(p)| \\ &\leq \langle \phi^{(k)}(p), \hat{\theta} \rangle + \Delta \\ &\leq \sqrt{\phi^{(k)}(p)^T \Lambda^{-1} \phi^{(k)}(p)} \sqrt{(\hat{\theta} - \theta^*)^T \Lambda (\hat{\theta} - \theta^*)} + \Delta \\ &\leq \gamma \sqrt{\phi^{(k)}(p)^T \Lambda^{-1} \phi^{(k)}(p)} + \Delta. \end{aligned}$$

The upper bound $\bar{f}(p) = \min\{d_{\max}, \langle \phi^{(k)}(p), \hat{\theta} \rangle + \gamma \sqrt{\phi^{(k)}(p)^T \Lambda^{-1} \phi^{(k)}(p)} + \Delta\}$. We can infer from the above analysis that with probability $1 - \delta$, $\bar{f}(p) \geq f(p), \forall p \in \mathbf{I}$. So $\max_{p \in \mathbf{I}} pf(p) - \hat{p}f(\hat{p}) \leq \hat{p}|\bar{f}(\hat{p}) - f(p)| \leq 2 \min\{d_{\max}, \gamma \sqrt{\phi^{(k)}(p)^T \Lambda^{-1} \phi^{(k)}(p)} + \Delta\}$ which completes the proof. \square

C.3.2 Proof of Lemma 14

Proof. Invoke Lemma 13 with $\delta = \frac{1}{T^2}$ and let $\Lambda_i = I + \sum_{i' < i} \phi^{(k)}(\hat{p}_{i'}) \phi^{(k)}(\hat{p}_{i'})^T$ denote the Λ matrix at the i th call. Denote $\gamma_{\max} = \max_{i \leq t} \gamma_i$, and we can easily verify $\gamma_{\max} \leq L\sqrt{k+1} + \Delta\sqrt{t} + d_{\max}\sqrt{6(k+1)\ln((k+1)T)}$. Recalling the right side of Lemma 13, and noting that $\gamma_{\max} \geq d_{\max}$ we have

$$\begin{aligned} \sum_{i=1}^t \min\{d_{\max}, \gamma \sqrt{\phi^{(k)}(\hat{p}_i)^T \Lambda^{-1} \phi^{(k)}(\hat{p}_i)} + \Delta\} &\leq \Delta t + \sum_{i=1}^t \min\{d_{\max}, \gamma_{\max} \sqrt{\phi^{(k)}(\hat{p}_i)^T \Lambda^{-1} \phi^{(k)}(\hat{p}_i)}\} \\ &\leq \Delta t + \gamma_{\max} \sum_{i=1}^t \min\{1, \sqrt{\phi^{(k)}(\hat{p}_i)^T \Lambda^{-1} \phi^{(k)}(\hat{p}_i)}\} \\ &\leq \Delta t + \gamma_{\max} \sqrt{t} \times \sqrt{\sum_{i=1}^t \min\{1, \phi^{(k)}(\hat{p}_i)^T \Lambda^{-1} \phi^{(k)}(\hat{p}_i)\}}, \end{aligned}$$

Using the elliptical potential lemma (Abbasi-Yadkori et al. (2012), Lemma 11), we know that

$$\min\{1, \phi^{(k)}(\hat{p}_i)^T \Lambda^{-1} \phi^{(k)}(\hat{p}_i)\} \leq 2(k+1) \ln((k+1)t+1),$$

Subsequently,

$$\begin{aligned} \Delta t + \gamma_{\max} \sqrt{t} \times \sqrt{\sum_{i=1}^t \min\{1, \phi^{(k)}(\hat{p}_i)^T \Lambda^{-1} \phi^{(k)}(\hat{p}_i)\}} \\ \leq \Delta t + \gamma_{\max} \sqrt{t} \times \sqrt{2(k+1) \ln((k+1)t+1)} \\ \leq \Delta t + \left(L\sqrt{k+1} + \Delta\sqrt{t} + d_{\max}\sqrt{6(k+1)\ln((k+1)T)} \right) \times \sqrt{t} \times \sqrt{2(k+1) \ln((k+1)t+1)} \\ \leq \left[(3d_{\max} + L) \times \sqrt{2t} + \Delta t \right] \times (k+1) \ln(2(k+1)T), \end{aligned}$$

So

$$\frac{1}{t} \sum_{i=1}^t \left[\max_{p \in \mathbf{I}} pf(p) - \hat{p}_i f(\hat{p}_i) \right] \leq \left[\Delta + \frac{(3d_{\max} + L) \sqrt{2}}{\sqrt{t}} \right] (k+1) \ln(2(k+1)T).$$

\square