

Policy Optimization Using Semiparametric Models for Dynamic Pricing

Jianqing Fan Yongyi Guo Mengxin Yu

Abstract

In this paper, we study the contextual dynamic pricing problem where the market value of a product is linear in its observed features plus some market noise. Products are sold one at a time, and only a binary response indicating success or failure of a sale is observed. Our model setting is similar to [Javanmard and Nazerzadeh \[2019\]](#) except that we expand the demand curve to a semiparametric model and learn dynamically both parametric and nonparametric components. We propose a dynamic statistical learning and decision making policy that minimizes regret (maximizes revenue) by combining semiparametric estimation for a generalized linear model with unknown link and online decision making. Under mild conditions, for a market noise c.d.f. $F(\cdot)$ with m -th order derivative ($m \geq 2$), our policy achieves a regret upper bound of $\tilde{O}_d(T^{\frac{2m+1}{4m-1}})$, where T is the time horizon and $\tilde{O}_d$ is the order hiding logarithmic terms and the feature dimension d . The upper bound is further reduced to $\tilde{O}_d(\sqrt{T})$ if F is super smooth. These upper bounds are close to $\Omega(\sqrt{T})$, the lower bound where F belongs to a parametric class. We further generalize these results to the case with dynamic dependent product features under the strong mixing condition.

Keyword: Contextual dynamic pricing, generalized linear model with unknown link, policy optimization, non-parametric statistics.

¹Jianqing Fan is Frederick L. Moore '18 Professor of Finance, Professor of Statistics, and Professor of Operations Research and Financial Engineering at the Princeton University. Yongyi Guo and Mengxin Yu are Ph.D. students at Department of Operations Research and Financial Engineering, Princeton University, Princeton, NJ 08544, USA. Emails: {jqfan, yongyig, mengxiny}@princeton.edu. Research supported by the NSF grant DMS-2052926, DMS-2053832, and the ONR grant N00014-19-1-2120

1 Introduction

Dynamic pricing is the study of determining and adjusting the selling prices of products over time based on statistical learning and policy optimization. As an integral part of revenue management, it has wide applications to various industries. Research on dynamic pricing has spanned across the fields of statistics, machine learning, economics, and operations research [den Boer, 2015, Wei and Zhang, 2018, Misic and Perakis, 2020]. In general, a good pricing strategy often involves good statistical learning of the demand function as well as revenue optimization over time.

Recent works particularly focus on feature-based (or contextual) pricing models, where the market value of a product as well as the pricing strategy depend on some observable features of the product [Javanmard and Nazerzadeh, 2019, Ban and Keskin, 2020]. Given the product features (covariates) available through the massive real-time data in online platforms today, feature-based pricing models take product heterogeneity into account, which enable customized pricing for products.

In this work, we consider the following dynamic pricing problem: We assume that a seller sells one product at each time $t = 1, \dots, T$. Each product is attached with a known feature vector $\mathbf{x}_t \in \mathbb{R}^d$. In addition, the product's market value v_t is linear in the features plus some i.i.d. market noise z_t with an *unknown* cumulative distribution $F(\cdot)$:

$$v_t = \boldsymbol{\theta}_0^\top \tilde{\mathbf{x}}_t + z_t, \quad z_t \sim F.$$

Here $\tilde{\mathbf{x}}_t = (\mathbf{x}_t^\top, 1)^\top$ and $\boldsymbol{\theta}_0$ is some unknown parameter. The customer makes an independent purchase decision for each product depending on whether the seller's posted price p_t is higher than the market value v_t , after which the revenue is collected. In this case, the demand curve $P(v_t \geq p_t)$ actually depends on both the parameter $\boldsymbol{\theta}_0$ as well as the distribution of z_t , which admits a semi-parametric form. They need to be learned or estimated dynamically from the observed binary data that indicates whether a sale is successful. Under this setting, we propose a policy which utilizes semi-parametric estimation techniques to achieve a low regret. In particular, under mild regularity conditions, if the c.d.f. $F(\cdot)$ of z_t has m^{th} derivative, the regret over a time horizon T is upper bounded by $\mathcal{O}((Td)^{\frac{2m+1}{4m-1}} \log T(1 + \log T/d))$, where d is the number of features. This result is further generalized to a setting where the product features $\mathbf{x}_t$ are not independent, as long as $\{\mathbf{x}_t\}_{t \geq 1}$ is a stationary

series that satisfies certain β -mixing conditions. Moreover, when F is infinitely differentiable, the total regret can be upper bounded by $\tilde{\mathcal{O}}((Td)^{\frac{1}{2}}(\log T)^{\frac{3}{2}+\frac{3}{2\alpha}}(\log(d+1) + \log T/d))$. This rate is the same as the parametric lower bound up to some logarithmic factors, i.e. where the distribution of z_t is generated from a parametric class.

1.1 Related Literatures

Our work contributes to the recent line of dynamic pricing literature as well as the growing literature on decision making with covariate information and contributes to kernel regression. Our work is also closely related to the non-parametric statistics literature. We'll briefly review the related works in the below.

Dynamic pricing. In the classical pricing models, one aims at maximizing the revenue over time by posting price sequentially while learning the underlying demand curve or market evaluation of a product. The demand curve is typically fixed over time, and falls into a known function class. Related literature includes [Kleinberg and Leighton \[2003\]](#), [Rusmevichientong et al. \[2006\]](#), [Besbes and Zeevi \[2009\]](#), [Broder and Rusmevichientong \[2012\]](#), [Keskin and Zeevi \[2014\]](#), [den Boer and Zwart \[2014\]](#), [Wang et al. \[2014\]](#), [den Boer and Zwart \[2015\]](#), [Babaioff et al. \[2015\]](#), [Cesa-Bianchi et al. \[2019\]](#), [Chen et al. \[2019\]](#). As an example, [Cesa-Bianchi et al. \[2019\]](#) study the dynamic pricing problem where the buyer's valuation of a product is supported on a finite K unknown points, and the success of a sale is determined by comparing the valuation to the proposed price. Using a generalization of UCB algorithm, the authors achieve the regret with order $\mathcal{O}(K \log T)$. For a comprehensive survey on this topic, see [den Boer \[2015\]](#).

Recently, many papers have been focusing on contextual dynamic pricing, where product heterogeneity is taken into account when modeling the demand curve or market evaluation. A common and natural choice is to model the market value of the product at time t as a linear function of its features $\mathbf{x}_t$ plus some market noise z_t , i.e. $v_t = \boldsymbol{\theta}^\top \mathbf{x}_t + z_t$ where $\boldsymbol{\theta}$ is some unknown parameter [[Qiang and Bayati, 2016](#), [Javanmard, 2017](#), [Miao et al., 2019](#), [Javanmard and Nazerzadeh, 2019](#), [Ban and Keskin, 2020](#), [Wang et al., 2020a](#), [Chen et al., 2021](#), [Tang et al., 2020](#), [Golrezaei et al., 2020](#)]. Under this setting, for 'truthful' buyers whose decision is based on comparing v_t and offered price p_t , the demand curve

can be expressed as a generalized linear model given feature covariates $\mathbf{x}_t$, where the link function is closely related to the distribution of the market noise z_t (see (3) for a detailed reasoning). Qiang and Bayati [2016] assume a linear model between the demand curve and the product features. They prove that the greedy iterative least squares (GILS) algorithm achieves a regret upper bound of $\mathcal{O}_d(\log T)$, where $\tilde{\mathcal{O}}_d$ is the order that hides logarithmic terms and the dimensionality of feature d , and provide a matching lower bound under their setting. Miao et al. [2019] and Ban and Keskin [2020] consider a generalized linear model with known link, while Javanmard and Nazerzadeh [2019] and Wang et al. [2020a] study the same problem with high dimensional sparse parameters. The algorithms are usually a combination of statistical estimation procedures and online learning techniques. Depending on the setting, the optimal regret ranges from $\mathcal{O}_d(\log T)$ to $\tilde{\mathcal{O}}_d(\sqrt{T})$. Other related works include Chen et al. [2021], Tang et al. [2020] where the authors explore certain differentially private policies under similar model setting; Golrezaei et al. [2020] where the authors consider the second price auction problem with multiple customers, each of which has his/her own product evaluation; and Javanmard [2017] where the parameter θ in the generalized linear model changes through time.

In practice, however, the distribution of the market noise z_t is usually unknown to the seller. Thus, it might be desirable to only assume that the noise density falls into some general class. As will be discussed in §2, this leads to modeling the demand curve as a generalized linear model with unknown link, and will be our main focus in this paper. Compared to the previous setting, this setting is more challenging, and the related literature is sparse. Javanmard and Nazerzadeh [2019] propose a preliminary algorithm that achieves a regret upper bound of $\mathcal{O}_d(T)$. Golrezaei et al. [2019] consider a second price auction with reserve where there are more than one customers, each of whom has his/her individual parameters in their demand curve model, and the customer bids are available as additional information. The authors propose the NPAC-T/NPAC-S policy that achieves a regret $\tilde{\mathcal{O}}_d(\sqrt{T})$. Golrezaei et al. [2020] also explore the second price auction and derive a regret upper bound of $\tilde{\mathcal{O}}_d(T^{2/3})$ compared to a ‘robust benchmark’ where the price maximizes the revenue of the worst link function in the class. Shah et al. [2019] explore an alternative setting where the market value $v_t = \exp(\theta^\top \mathbf{x}_t + z_t)$ and z_t has unknown distribution. By utilizing this specific structure, the authors propose the DEEP-C algorithm based on multi-arm bandit that has a regret upper bound of $\tilde{\mathcal{O}}_d(\sqrt{T})$. The authors also pro-

pose some variants of the algorithm and study them via simulations. Recently, Luo et al. [2021] study a similar problem to ours, assuming a linear market valuation with unknown noise distribution. They provide a DIP policy that achieves regret $\mathcal{O}_d(T^{2/3} + \|\hat{\theta} - \theta\|_1 T)$, where $\|\hat{\theta} - \theta\|_1$ is the estimation accuracy of the parameter θ .

There are some literature studying other dynamic pricing algorithms [Amin et al., 2014, Cohen et al., 2016, Mao et al., 2018, Leme and Schneider, 2018, Nambiar et al., 2019, Anton and Alexey, 2020, Alexey, 2020, Ban and Keskin, 2020, Li and Zheng, 2020, Javanmard et al., 2020, Chen and Gallego, 2020, Liu et al., 2021]. For example, Mao et al. [2018] study a non-parametric dynamic pricing where the market value is modeled as a general non-parametric function $f(\mathbf{x}_t)$, where $\mathbf{x}_t$ are the features. A binary feedback is similarly observed based on the comparison between $f(\mathbf{x}_t)$ and the proposed price. The authors apply a variation of midpoint algorithm and achieve a regret upper bound of $\mathcal{O}(T^{d/(d+1)})$ with d being the dimension of $\mathbf{x}_t$.

Semi-parametric and non-parametric statistical estimation. Our work is also closely related to estimation of the single index model, or the generalized linear model with an unknown link. Such model has been studied in the statistics and econometrics literature for decades, and has wide applications in fields like econometrics and finance [Powell et al., 1989, Ichimura, 1993, Hardle et al., 1993, Klein and Spady, 1993, Weisberg and Welsh, 1994, Mallick and Gelfand, 1994, Horowitz and Härdle, 1996, Carroll et al., 1997, Xia and Li, 1999, Delecroix et al., 2003, Fan and Li, 2004]. For a comprehensive summary of these works, please refer to McCulloch [2000], Györfi et al. [2002], Fan and Yao [2003], Ruppert et al. [2003], Tsybakov [2008], Horowitz [2012]. Various methods have been proposed to estimate the parametric part that achieves root- n consistency under certain conditions [Powell et al., 1989, Ichimura, 1993, Klein and Spady, 1993]. Carroll et al. [1997] study the generalized partial linear single index models, where the authors leverage local linear kernel regression with quasi-likelihood method to estimate both the parametric and non-parametric parts of the model. Xia and Li [1999] investigate in the single index coefficient model with strong-mixing features. Estimators with uniform convergence rate to the ground truth based on kernel regression are proposed.

Given a root- n consistent estimation of the coefficients, standard univariate non-parametric regression techniques can be used to estimate the non-parametric part of the single index model that achieves

ℓ_∞ consistency, which is necessary in deriving regret upper bounds. One common estimator is the Nadaraya-Watson estimator [Nadaraya, 1964, Watson, 1964]. Silverman [1978] and Mack and Silverman [1982] establish uniform convergence results for kernel density estimator and Nadaraya-Watson estimator for regression functions. In addition, Stone [1980, 1982] derive uniform convergence results for the more general local polynomial regression estimators. Masry [1996] prove similar results when the covariates satisfy strong-mixing conditions.

In this paper, we'll provide non-asymptotic error bounds for both coefficient estimation as well as the plug-in Nadaraya-Watson estimator in a uniform sense. These non-asymptotic bounds are useful for constructing regret bounds within a finite horizon.

1.2 Our Contributions

Our contributions are the following: First, compared to related works, our policy achieves a low regret with few assumptions on the market noise distribution and little additional information. Given $F \in \mathbb{C}^{(m)}$ where F is the c.d.f. of z_t , the regret over a time horizon T is upper bounded by $\tilde{O}((Td)^{\frac{2m+1}{4m-1}})$; If F is 'super smooth', the bound is further reduced to $\tilde{O}(\sqrt{Td})$, which is nearly the same regret order by assuming a parametric distribution for z_t as in Javanmard and Nazerzadeh [2019] where the s -sparsity on β_0 is imposed. Table 1 illustrates the settings of our work as well as several related literatures. Golrezaei et al. [2020] choose a more 'conservative' regret by comparing to a benchmark policy which minimizes revenue with the worst demand function over the whole ambiguity function class. In contrast, our notation of regret is more standard and 'accurate' in that our benchmark policy knows the exact demand function given any product features. Shah et al. [2019] consider a log-linear relation between the market value and the covariates instead of a linear relation and derive a regret upper bound of $\tilde{O}(\sqrt{T}d^{11/4})$. Their algorithm based on multi-arm bandit has sub-optimal dependence on the dimension d in terms of both regret and complexity, and is quite difficult to implement under general conditions. Interestingly, the authors conjecture that under the linear settings, there is no policy that achieves an $\tilde{O}_d(\sqrt{T})$ regret. Our work partly answers their guess by providing a policy with a $\tilde{O}(\sqrt{Td})$ regret when the demand function is sufficiently smooth. Compared to the DIP policy in Luo et al. [2021] and its regret $\mathcal{O}_d(T^{2/3} + \|\hat{\theta} - \theta\|_1 T)$, we are more clear on how $\hat{\theta}$ are estimated

	Feature-based	Non-parametric noise	Regret
Kleinberg and Leighton [2003]		✓	$\tilde{O}(\sqrt{T})$
Javanmard and Nazerzadeh [2019]	✓		$\tilde{O}(s\sqrt{T})$
Shah et al. [2019]	✓ (log-linear model)	✓	$\tilde{O}(\sqrt{T}d^{11/4})$
Golrezaei et al. [2020]	✓	✓	$\tilde{O}(dT^{2/3})$ (changed benchmark)
Luo et al. [2021]	✓	✓	$\mathcal{O}_d(T^{2/3} + \ \hat{\theta} - \theta_0\ _1 T)$
Our work	✓ (linear model)	✓	$\tilde{O}((Td)^{\frac{2m+1}{4m-1}})$

Table 1: Comparison with related works.

within the pricing algorithm, and we provide explicit rate on both the estimation error and the regret. Moreover, compared to several fully non-parametric dynamic pricing literatures, such as Mao et al. [2018] and Chen and Gallego [2020], our algorithm scales more nicely with dimension d , and can easily be generalized to a high-dimensional setting. Our algorithm is also easy to implement compared to some bandit-based algorithms that need dividing the feature space into bins.

Second, we generalize our results to the regime where the product features $\{\mathbf{x}_t\}_{t \geq 1}$ are weakly dependent instead of independent, which is more likely in practice. For example, for many products (such as softwares, electric products, etc.), the features of the products evolve over time and definitely inherit some past information. In other situations, the products for sale might have some common time-dependent factors shared by all products in the same industry (such as weather condition, population composition, etc.). This setting with weakly-dependent features can also be found in literatures such as Chen et al. [2022], where the authors study an offline pricing problem with parametric models and

dependent covariates.

Last but not least, we establish non-asymptotic results on the ℓ_∞ error bound of the nonparametric kernel density and regression estimation, which are potentially useful in other related study as well. As mentioned in the related literatures, most results on non-parametric kernel regression estimation are established under the asymptotic settings. Meanwhile, we believe that non-asymptotic results are necessary to achieve a finite-sample regret upper bound in the pricing problem. Please refer to Appendix C.2 for related lemmas.

1.3 Notation

Throughout this work, we use $[n]$ to denote $\{1, 2, \dots, n\}$. For any vector $\mathbf{x} \in \mathbb{R}^n$ and $q \geq 0$, we use $\|\mathbf{x}\|_q$ to represent the vector ℓ_q norm, i.e. $\|\mathbf{x}\|_q = (\sum_{i=1}^n |x_i|^q)^{1/q}$. In addition, we let $\nabla_{\mathbf{x}} L(\cdot)$, $\nabla_{\mathbf{x}}^2 L(\cdot)$ be the gradient vector and Hessian matrix of loss function $L(\cdot)$ with respect to $\mathbf{x}$. For any given matrix $\mathbf{X} \in \mathbb{R}^{d_1 \times d_2}$, we use $\|\cdot\|$ to denote the spectral norm of $\mathbf{X}$ and we write $\mathbf{X} \succeq 0$ or $\mathbf{X} \preceq 0$ if $\mathbf{X}$ or $-\mathbf{X}$ is semidefinite. For any event A , we let $\mathbb{I}_A$ be a indicator random variable which is equal to 1 if A is true and 0 otherwise. In addition, we use $\mathbb{C}^{(m)}$ with $m \in \mathbb{N}$ to denote the function class which contains all functions with m -th order continuous derivatives. For two positive sequences $\{a_n\}_{n \geq 1}$, $\{b_n\}_{n \geq 1}$, we write $a_n = \mathcal{O}(b_n)$ or $a_n \lesssim b_n$ if there exists a positive constant C such that $a_n \leq C \cdot b_n$ and we write $a_n = o(b_n)$ if $a_n/b_n \rightarrow 0$. In addition, we write $a_n = \Omega(b_n)$ or $a_n \gtrsim b_n$ if $a_n/b_n \geq c$ with some constant $c > 0$. We use $a_n = \Theta(b_n)$ if $a_n = \mathcal{O}(b_n)$ and $a_n = \Omega(b_n)$. We use notations $\mathcal{O}_d(\cdot)$, $\Omega_d(\cdot)$ and $\Theta_d(\cdot)$ to denote similar meanings as above while treating the variable d as fixed. Moreover, we let $\tilde{\mathcal{O}}(\cdot)$, $\tilde{\Omega}(\cdot)$, $\tilde{\Theta}(\cdot)$ represent the same meaning with $\mathcal{O}(\cdot)$, $\Omega(\cdot)$ and $\Theta(\cdot)$ except for ignoring log factors.

1.4 Roadmap

The rest of this paper is organized as follows. We describe the problem in §2 and propose a solution in §3 where some heuristic arguments are offered for bounding the regret. In §4, we provide our theoretical results on the upper bounds of the regret and in §B, we discuss a lower bound result. Our algorithm is illustrated in §5 by intensive simulation experiments.

2 Problem Setting

We consider the pricing problem where a seller has a single product for sale at each time period $t = 1, 2, \dots, T$. Here T is the total number of periods (i.e. length of horizon) and may be unknown to the seller. The market value of the product at time t is v_t and is unknown. We assume that the range of v_t is contained in a closed interval in $(0, B)$. In particular, we assume that $v_t \in [\delta_v, B - \delta_v]$ for some constant $\delta_v > 0$. At each period t , the seller posts a price p_t . If $p_t \leq v_t$, a sale occurs, and the seller collects a revenue of p_t ; otherwise, no sale occurs and no revenue is obtained. Let y_t be the response variable that indicates whether a sale has occurred at period t . Then

$$y_t = \begin{cases} +1 & \text{if } v_t \geq p_t, \\ 0 & \text{if } v_t < p_t. \end{cases} \quad (1)$$

The goal of the seller is to design a pricing policy that maximizes the collected revenue.

In this paper, we further model the market value v_t as a linear function of the product's observable feature covariate $\mathbf{x}_t \in \mathbb{R}^d$. In particular, define $\tilde{\mathbf{x}}_t = (\mathbf{x}_t^\top, 1)^\top$, where we assume $\{\mathbf{x}_t\}_{t \geq 1}$ are i.i.d. samples from an unknown distribution $\mathbb{P}_X$ supported on a bounded subset $\mathcal{X} \subseteq \mathbb{R}^d$. Assume that

$$v_t = \boldsymbol{\theta}_0^\top \tilde{\mathbf{x}}_t + z_t, \quad (2)$$

where $\boldsymbol{\theta}_0 = (\boldsymbol{\beta}_0^\top, \alpha_0)^\top \in \mathbb{R}^{d+1}$ is an unknown parameter, and $\{z_t\}_{t \geq 1}$ is an i.i.d. sequence of idiosyncratic noise drawn from an **unknown** distribution with zero mean and bounded support $(-\delta_z, \delta_z)$. The cumulative distribution function of z_t is denoted by $F(\cdot)$. The above model implies that

$$y_t = \begin{cases} +1 & \text{with probability } 1 - F(p_t - \boldsymbol{\theta}_0^\top \tilde{\mathbf{x}}_t), \\ 0 & \text{with probability } F(p_t - \boldsymbol{\theta}_0^\top \tilde{\mathbf{x}}_t). \end{cases} \quad (3)$$

Remark 1. In fact, each $\mathbf{x}_t$ here can contain both product information and the buyer information, as long as this information is revealed to the seller.

Remark 2. The reason that we assume z_t has bounded support $[-\delta_z, \delta_z]$ is to ensure the market valuation $v_t \geq 0$, which is more reasonable in practice (Otherwise v_t has positive probability to be negative, since z_t is independent with the covariates $\mathbf{x}_t$). The truncated Gaussian distribution falls in such category. If the market allows v_t, p_t to be negative, then we can replace the boundness of z_t by any sub-Gaussian distributions.

In a non-dynamic setting, the model (3) is closely related to the single index model, or generalized linear (logistic regression) model with unknown link function [Ichimura, 1993, Fan et al., 1995, Carroll et al., 1997]. In their works, it's usually assumed that $p_t = 0$ and $\{\tilde{\mathbf{x}}_t\}_{t \geq 1}$ are independent observations, and the goal is to estimate θ_0 and F . Meanwhile, we work on the dynamic setting where we need to optimize some revenue function by iteratively deciding p_t given previous observations based on dynamically learned θ_0 and F . These two problems are closely related but also decisively different.

We now state our objective in more details. Given observed features $\mathbf{x}_t$, the expected revenue at time t with a posted price p is

$$\text{rev}_t(p, \theta_0, F) := \mathbb{E} p \cdot \mathbb{1}(v_t \geq p) = p(1 - F(p - \theta_0^\top \tilde{\mathbf{x}}_t)). \quad (4)$$

The optimal posted price p_t^* for a product with attribute $\mathbf{x}_t$ is given by

$$p_t^* = \underset{p \geq 0}{\operatorname{argmax}} p(1 - F(p - \theta_0^\top \tilde{\mathbf{x}}_t)), \quad (5)$$

which depends on unknown parameters and needs to be learned dynamically from the data. As in common practice, we evaluate the performance of any policy π that governs the rule of posted prices $\{p_t\}_{t \geq 1}$ by investigating the regret compared to the ‘oracle pricing policy’ that uses the knowledge of both θ_0 and $F(\cdot)$ and offers p_t^* according to (5) for any given t . In other words, we consider the problem of maximizing revenue as minimizing the following maximum regret

$$\text{Regret}_\pi(T) \equiv \max_{\substack{\theta_0 \in \Omega \\ \mathbb{P}_X \in \mathcal{Q}(\mathcal{X})}} \mathbb{E} \left[\sum_{t=1}^T \left(p_t^* \mathbb{1}(v_t \geq p_t^*) - p_t(\pi) \mathbb{1}(v_t \geq p_t(\pi)) \right) \right], \quad (6)$$

where the expectation is taken with respect to the the idiosyncratic noise z_t and $\mathbf{x}_t$, and $p_t(\pi)$ denotes the price offered at time t by following policy π . Here $\mathcal{Q}(\mathcal{X})$ represents the set of probability distributions supported on a bounded set $\mathcal{X}$. Our goal is to choose a good strategy π such that the above total regret is small.

Apparently, learning θ_0 and $F(\cdot)$ over time gives the seller much more information to estimate the market value of a new product given its feature covariates. On the other hand, the seller also wants to always give optimized price so as to maximize the expected revenue by (5). Therefore, it's necessary to have a good policy that strikes a balance between exploration (collecting data information for learning parameters) and exploitation (offering optimal pricing based on learned parameters).

Before proposing our algorithm, we first impose some regularity condition on F so that the optimization problem (5) is 'well-behaved'.

Assumption 1. *There exists a positive constant c_ϕ such that $\phi'(u) \geq c_\phi$ for all $u \in (-\delta_z, \delta_z)$, where $\phi(u) := u - \frac{1-F(u)}{F'(u)}$.*

Assumption 1 ensures that $\phi(\cdot)$ is strictly increasing, which implies a unique solution to (5). In fact, the first order condition of (5) yields

$$p_t^* = g(\theta_0^\top \tilde{\mathbf{x}}_t),$$

where $g(u) \triangleq u + \phi^{-1}(-u)$.

Remark 3. We only put some necessary assumptions on F in order to guarantee the existence of the unique optimal price p_t^* in (5), given observed $\tilde{\mathbf{x}}_t$ and unknown but fixed θ_0 . Comparing to the Assumption 2.1 in Javanmard and Nazerzadeh [2019], our Assumption 1 is weaker, since assumption that $1 - F(u)$ is log-concave is a special case of our assumption with $c_\phi \geq 1$.

3 Algorithm and Basic Regret Analysis

We first propose Algorithm 1 in §3.1 which describes our policy for minimizing the regret given in (6), and then provide the main idea for the regret analysis achieved by our Algorithm 1 in §3.2.

3.1 A Proposed Algorithm

In the following algorithm, we divide the time horizon into 'episodes' with increasing lengths. The first part of each episode is a short exploration phase where the offered prices are i.i.d. to collect the

data and model parameters (i.e. $\hat{\theta}$, $\hat{F}$) are then updated based on the collect data. The second part is an exploitation phase, where the optimal p_t is offered according to the current estimate of parameters and the new $\tilde{\mathbf{x}}_t$. The details are stated in Algorithm 1.

Algorithm 1: Feature-based dynamic pricing with unknown noise distribution

- 1: **Input:** Upper bound of market value ($\{v_t\}_{t \geq 1}$): $B > 0$, minimum episode length: ℓ_0 , degree of smoothness: m .
- 2: **Initialization:** $p_1 = 0$, $\hat{\theta}_1 = 0$.
- 3: **for** each episode $k = 1, 2, \dots$, **do**
- 4: Set length of the k -th episode $\ell_k = 2^{k-1}\ell_0$; Length of the exploration phase $a_k = \lceil (\ell_k d)^{\frac{2m+1}{4m-1}} \rceil$.
- 5: **Exploration Phase** ($t \in I_k := \{\ell_k, \dots, \ell_k + a_k - 1\}$):
- 6: Offer price $p_t \sim \text{Unif}(0, B)$.
- 7: **Updating Estimates (at the end of the exploration phase with data $\{(\tilde{\mathbf{x}}_t, y_t)\}_{t \in I_k}$):**
- 8: Update estimate of θ_0 by $\hat{\theta}_k = \hat{\theta}_k(\{(\tilde{\mathbf{x}}_t, y_t)\}_{t \in I_k})$;

$$\hat{\theta}_k = \underset{\theta}{\operatorname{argmin}} L_k(\theta) := \frac{1}{|I_k|} \sum_{t \in I_k} (By_t - \theta^\top \tilde{\mathbf{x}}_t)^2 \quad (7)$$

- 9: Update estimates of F , F' by $F_k(u, \hat{\theta}_k) = F_k(u; \hat{\theta}_k, \{(\tilde{\mathbf{x}}_t, y_t, p_t)\}_{t \in I_k})$,
 $F_k^{(1)}(u, \hat{\theta}_k) = F_k^{(1)}(u, \hat{\theta}_k, \{(\tilde{\mathbf{x}}_t, y_t, p_t)\}_{t \in I_k})$. The detailed formulas are given by (14) and (16).
- 10: Update estimate of ϕ by $\hat{\phi}_k(u) = u - \frac{1 - \hat{F}_k(u)}{\hat{F}_k^{(1)}(u)}$ and estimate of g by $\hat{g}_k(u) = u + \hat{\phi}_k^{-1}(-u)$.
- 11: **Exploitation Phase** ($t \in I'_k := \{\ell_k + a_k, \dots, \ell_{k+1} - 1\}$):
- 12: Offer p_t as

$$p_t = \min\{\max\{\hat{g}_k(\tilde{\mathbf{x}}_t^\top \hat{\theta}_k), 0\}, B\}. \quad (8)$$

13: **end for**

Despite semiparametric model (3) with unknown link, by offering $p_t \sim \text{Unif}(0, B)$, By_t follows the linear model with regression $\tilde{\mathbf{x}}_t^\top \theta_0$ and this leads to the least-squares estimate (7). To see this, it follows that

$$\mathbb{E}[By_t | \tilde{\mathbf{x}}_t] = B\mathbb{E}_{z_t}\mathbb{E}[y_t | \tilde{\mathbf{x}}_t, z_t] = B\mathbb{E}_{z_t}\mathbb{E}[\mathbb{1}(p_t \leq \theta_0^\top \tilde{\mathbf{x}}_t + z_t) | \tilde{\mathbf{x}}_t, z_t] = B\mathbb{E}\frac{\theta_0^\top \tilde{\mathbf{x}}_t + z_t}{B} = \tilde{\mathbf{x}}_t^\top \theta_0.$$

On the other hand, a uniform distribution for p_t is actually critical for the above property. Suppose that p_t is drawn from a c.d.f. $F_p(\cdot)$ and there is a transform f_1 of y_t that satisfies

$$\mathbb{E}f_1(y_t) = \mathbb{E}\tilde{\mathbf{x}}_t^\top \boldsymbol{\theta}_0 = \mathbb{E}v_t$$

for all $\mathbb{P}_X$, then according to (3), we have

$$\begin{aligned} \mathbb{E}v_t &= \mathbb{E}\mathbb{E}[f_1(y_t) \mid \tilde{\mathbf{x}}_t, z_t] = \mathbb{E}\mathbb{E}[f_1(\mathbb{1}(p_t \leq \tilde{\mathbf{x}}^\top \boldsymbol{\theta}_0 + z_t)) \mid \tilde{\mathbf{x}}_t, z_t] \\ &= \mathbb{E}F_p(\tilde{\mathbf{x}}^\top \boldsymbol{\theta}_0 + z_t)f_1(1) + \mathbb{E}(1 - F_p(\tilde{\mathbf{x}}^\top \boldsymbol{\theta}_0 + z_t))f_1(0) \\ &= f_1(0) + (f_1(1) - f_1(0))\mathbb{E}F_p(v_t). \end{aligned}$$

Since the above equation holds for all $\mathbb{P}_X \in \mathcal{Q}(X)$, it can only be the case that F_p is linear within the region $[0, B]$, which implies that p_t should follow a uniform distribution.

Remark 4. In Algorithm 1, the interval $[0, B]$ can be replaced with any interval that covers the range of the market value v_t . In practice, we can shrink the sampling interval at each exploration phase according to the feedback information observed in the past.

Remark 5. If z_t follows distributions with unbounded support and sub-Gaussian tails, in Algorithm 1, we only need to replace B by $B_k = C\sqrt{\log |I_k|}$ such that v_t falls in $(-B_k, B_k)$ with high probability. We then offer $p_t \sim \text{Unif}(-B_k, B_k)$. Conditional on $v_t \in (-B_k, B_k)$, $B_k(2y_t - 1)$ serves as an unbiased estimator for $\tilde{\mathbf{x}}_t^\top \boldsymbol{\theta}_0$. Thus, all the following theoretical results work.

3.2 Main Idea for Regret Analysis

The main idea behind our regret analysis is a balance between exploration and exploitation. This idea is shown in the following heuristic arguments. For simplicity, we assume for now that there is only one episode, and that the total length of time (horizon) ℓ is known and d is bounded.

First, denote ℓ_1 as the length of the exploration phase. During this phase, the regret R_1 at each time is bounded by a constant due to bounded distribution $F(\cdot)$ that entails bounded p_t^* in (5). Therefore, the total regret in this phase is

$$R_1 = \mathcal{O}(\ell_1). \tag{9}$$

For the second phase, the expected regret can be controlled by the estimation error of both θ and g (which is a functional of F as mentioned in (8)). In fact, let the regret at each time point t be

$$R_t := p_t^* \mathbb{I}_{(v_t \geq p_t^*)} - p_t \mathbb{I}_{(v_t \geq p_t)}.$$

Then the conditional expectation of regret at time t given previous information and $\tilde{\mathbf{x}}_t$ is

$$\begin{aligned} \mathbb{E}[R_t | \bar{\mathcal{H}}_{t-1}] &= \mathbb{E}[p_t^* \mathbb{I}_{(v_t \geq p_t^*)} - p_t \mathbb{I}_{(v_t \geq p_t)} | \bar{\mathcal{H}}_{t-1}] \\ &= p_t^* (1 - F(p_t^* - \tilde{\mathbf{x}}_t^\top \theta_0)) - p_t (1 - F(p_t - \tilde{\mathbf{x}}_t^\top \theta_0)) \\ &= \text{rev}_t(p_t^*, \theta_0, F) - \text{rev}_t(p_t, \theta_0, F) \end{aligned} \quad (10)$$

Here $\bar{\mathcal{H}}_t = \sigma(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{t+1}; z_1, \dots, z_t)$. On the other hand, under mild conditions, the above difference in revenue can further be upper bounded by an order of $(p_t - p_t^*)^2$ using Taylor expansion.

Therefore, we have

$$\begin{aligned} \mathbb{E}[R_t | \bar{\mathcal{H}}_{t-1}] &\lesssim (p_t - p_t^*)^2 = (\hat{g}(\hat{\theta}^\top \tilde{\mathbf{x}}_t) - g(\theta_0^\top \tilde{\mathbf{x}}_t))^2 \\ &\leq 2(\hat{g}(\hat{\theta}^\top \tilde{\mathbf{x}}_t) - g(\hat{\theta}^\top \tilde{\mathbf{x}}_t))^2 + 2(g(\hat{\theta}^\top \tilde{\mathbf{x}}_t) - g(\theta_0^\top \tilde{\mathbf{x}}_t))^2 \\ &:= \mathbf{J}_1 + \mathbf{J}_2. \end{aligned} \quad (11)$$

In fact, $\mathbf{J}_2$ is upper bounded by $\|\hat{\theta} - \theta_0\|_2^2$ (given the Lipschitz property of g according to Assumption 1 and suitable conditions over $\mathbb{P}_X$). By solving (7), we prove that the squared ℓ_2 error is of order $\mathcal{O}(\ell_1^{-1})$, which is the order of $\mathbf{J}_2$. The term $\mathbf{J}_1$ is upper bounded by $\|\hat{g} - g\|_\infty^2$, and is further bounded by $\max\{\|\hat{F} - F\|_\infty^2, \|\hat{F}' - F'\|_\infty^2\}$. Note that by (1), $F(\cdot)$ is the non-parametric function of $1 - Y_t$ given $w_t = p_t - \tilde{\mathbf{x}}_t^\top \theta_0$, in which p_t is the observed price given in the exploration phase. Since θ_0 is estimated at a faster rate, we can assume that w_t is observable given a proper estimator of θ_0 . Therefore, the error rate is dominated by estimating $F'(\cdot)$. Assuming F has an m -th continuous derivative, we construct $\hat{g}$ using the kernel estimator with a m -th order kernel, and prove that $\max\{\|\hat{F} - F\|_\infty, \|\hat{F}' - F'\|_\infty\} \lesssim \mathcal{O}(\ell_1^{-(m-1)/(2m+1)})$ in which a logarithmic order is ignored for simplicity of presentation. Therefore, the total regret during the exploitation phase can be upper bounded by

$$R_2 \lesssim \ell \cdot \ell_1^{-2(m-1)/(2m+1)}. \quad (12)$$

Combining (9) and (12), we know that by choosing ℓ_1 of the order of $\ell^{(2m+1)/(4m-1)}$, we balance the regret of both exploration and exploitation phase, and the total regret during the episode is given by

$$R_1 + R_2 = \mathcal{O}(\ell^{(2m+1)/(4m-1)}).$$

For a second order kernel, the above regret is of order $\mathcal{O}(\ell^{5/7})$. For a relatively large m , the regret is close to $\mathcal{O}(\ell^{1/2})$, which is actually proven to be the lower bound for a wider class of problems.

4 Regret Results on Proposed Policy

In this section, we divide our results into three parts. In §4.1, we consider the setting with independent covariates and finite differentiable noise distributions. In §4.2, we further extend our results in §4.1 to the setting with correlated features. Finally we extend the aforementioned results to the regime with infinitely differentiable noise distributions i.e. $m = \infty$ in §4.3.

4.1 Result under Independence Settings

The main result of this section is Theorem 1. To obtain this results, we first state some technical conditions and technical lemmas, which demonstrate the accuracy of statistical learning in each episode. These lemmas provide insights how statistical accuracy influences on the regret of our policy and have interests of their own rights.

Assume that $\|\theta_0\| \leq R_\Theta$ for some constant $R_\Theta > 0$. We also define $R_{\mathcal{X}} := \sup_{\mathbf{x} \in \mathcal{X}} \|\mathbf{x}\|_2$. Before stating our main results, we first make the following assumptions on $\mathbf{x}_t$.

Assumption 2. *There exist positive constants $c_{\min}$ and $c_{\max}$, such that the covariance matrix Σ given by $\Sigma = \mathbb{E}[\tilde{\mathbf{x}}_t \tilde{\mathbf{x}}_t^\top]$ satisfies $c_{\min} \mathbb{I} \preceq \Sigma \preceq c_{\max} \mathbb{I}$, where $\tilde{\mathbf{x}}_t = (\mathbf{x}_t^\top, 1)^\top$*

As we observe from $\mathbf{J}_1, \mathbf{J}_2$ given in (11), bounding the regret in the exploitation phase needs to estimate both parameter θ_0 and function $g(\cdot)$. In the following, we first present an upper bound of estimating θ_0 at the end of the exploration phase within each episode in the following Lemma 1. Recall $|I_k|$ is the length of the k -th exploration phase.

Lemma 1. *Under Assumption 2, there exist positive constants c_0 and c_1 depending only on absolute constants given in assumptions such that for any episode k , as long as $|I_k| \geq c_0(d+1)$, with probability at least $1 - 2e^{-c_1 c_{\min}^2 |I_k|/16} - 2/|I_k|$,*

$$\|\hat{\boldsymbol{\theta}}_k - \boldsymbol{\theta}_0\|_2 \leq \frac{8 \max\{R_{\mathcal{X}}, 1\}(R_{\mathcal{X}}R_{\Theta} + B)}{c_{\min}} \sqrt{\frac{(d+1) \log |I_k|}{|I_k|}}. \quad (13)$$

Let $\Theta_k := B(\boldsymbol{\theta}_0, R_k)$, where R_k is the right hand side of (13). We conclude from Lemma 1 that with high probability, R_k is of order at most $\sqrt{d \log |I_k|/|I_k|}$, and we can achieve similar upper bounds for $\mathbf{J}_2$ for any episode k .

Next, we proceed to construct the estimator $\hat{g}_k$ in each episode and bound its distance to g . Notice that $g(u) = u + \phi^{-1}(-u)$, and $\phi(u) = u - \frac{1-F(u)}{F'(u)}$. Thus, a natural way to construct $\hat{g}_k$ is from an estimate of F and F' , as mentioned in our algorithm. Moreover, the uniform error bounds of our estimators $\hat{F}_k$ and $\hat{F}_k^{(1)}$ guarantee a uniform error bound of $\hat{g}_k$.

We use the kernel regression method and $\hat{\boldsymbol{\theta}}_k$ obtained above to construct $\hat{F}_k$ and $\hat{F}_k^{(1)}$. Recall that by (3), we have $E(y_t|w_t(\boldsymbol{\theta}_0)) = 1 - F(w_t(\boldsymbol{\theta}_0))$ where $w_t(\boldsymbol{\theta}) := p_t - \tilde{\mathbf{x}}_t^\top \boldsymbol{\theta}$. Recall p_t is the observed price offered in the k -th exploration phase. Thus, given $\hat{\boldsymbol{\theta}}_k$, $F(\cdot)$ can be estimated by using the Nadaraya-Watson kernel regression estimator and $F'(\cdot)$ can be estimated by the derivative of the estimator. Specifically, we define

$$\hat{F}_k(u, \boldsymbol{\theta}) = 1 - \hat{r}_k(u, \boldsymbol{\theta}) = 1 - \frac{h_k(u, \boldsymbol{\theta})}{f_k(u, \boldsymbol{\theta})}, \quad (14)$$

and $\hat{F}_k^{(1)}(u) = \hat{F}_k^{(1)}(u, \hat{\boldsymbol{\theta}}_k)$, where

$$h_k(u, \boldsymbol{\theta}) = \frac{1}{|I_k|b_k} \sum_{t \in I_k} K\left(\frac{w_t(\boldsymbol{\theta}) - u}{b_k}\right) Y_t, \quad f_k(u, \boldsymbol{\theta}) = \frac{1}{|I_k|b_k} \sum_{t \in I_k} K\left(\frac{w_t(\boldsymbol{\theta}) - u}{b_k}\right), \quad (15)$$

for a chosen m -th order kernel K and a suitable bandwidth b_k . Now, we estimate the derivative $F'(\cdot)$ by taking the derivative of the estimator. That is, $\hat{F}_k^{(1)}(u) = \hat{F}_k^{(1)}(u, \hat{\boldsymbol{\theta}}_k)$ where

$$\hat{F}_k^{(1)}(u, \boldsymbol{\theta}) = -\hat{r}_k^{(1)}(u, \boldsymbol{\theta}) = -\frac{h_k^{(1)}(u, \boldsymbol{\theta})f_k(u, \boldsymbol{\theta}) - h_k(u, \boldsymbol{\theta})f_k^{(1)}(u, \boldsymbol{\theta})}{f_k^2(u, \boldsymbol{\theta})}, \quad (16)$$

$$h_k^{(1)}(u, \boldsymbol{\theta}) = \frac{-1}{|I_k|b_k^2} \sum_{t \in I_k} K'\left(\frac{w_t(\boldsymbol{\theta}) - u}{b_k}\right) Y_t, \quad f_k^{(1)}(u, \boldsymbol{\theta}) = \frac{-1}{|I_k|b_k^2} \sum_{t \in I_k} K'\left(\frac{w_t(\boldsymbol{\theta}) - u}{b_k}\right). \quad (17)$$

Recall we mention in §2 that $(-\delta_z, \delta_z)$ is the support of noise z_t . In addition, we also mention that T denotes the length of time horizon which is unknown. In the following, we will state other necessary assumptions to derive the regret upper bound:

Assumption 3. *The density of $w_t(\boldsymbol{\theta})$ (denoted as $f_{\boldsymbol{\theta}}$) satisfies the following:*

- (Smoothness) *There exists an integer $m \geq 2$ and a constant l_f such that for all $\boldsymbol{\theta} \in \Theta_0 := \{\boldsymbol{\theta} \mid \|\boldsymbol{\theta} - \boldsymbol{\theta}_0\|_2 \leq C_{\boldsymbol{\theta}} T^{-\frac{2m+1}{4(4m-1)}} d^{\frac{m-1}{4m-1}} \sqrt{\log T + 2 \log d}\}$, $f_{\boldsymbol{\theta}}(u) \in \mathbb{C}^{(m)}$, and $f_{\boldsymbol{\theta}}^{(m)}$ is l_f -Lipschitz on $I := [-\delta_z, \delta_z]$.*
- (Boundedness) *There exists a constant $\bar{f} > 0$ such that $\forall u \in \mathbb{R}$ and $\boldsymbol{\theta} \in \Theta_0$, $\max\{|f_{\boldsymbol{\theta}}(u)|, |f'_{\boldsymbol{\theta}}(u)|\} \leq \bar{f}$. In addition, there exists a universal constant $c > 0$ such that $f_{\boldsymbol{\theta}}(u) \geq c$ for all $u \in I$, $\boldsymbol{\theta} \in \Theta_0$.*

Remark 6. We provide some examples for Assumption 3. For any covariate $\mathbf{x} \in \mathbb{R}^d$, as long as there exists an entry of it that follows a continuous distribution in $\mathbb{C}^{(m)}$, $m \geq 1$, such as Beta-distribution or truncated Gaussian distribution, we can ensure the density of $w(\boldsymbol{\theta}) = p_t - \tilde{\mathbf{x}}_t^\top \boldsymbol{\theta}$ satisfies both the smoothness and boundedness conditions in Assumption 3.

Assumption 4. $r_{\boldsymbol{\theta}}(u) := \mathbb{E}[y_t \mid w_t(\boldsymbol{\theta}) = u]$ satisfies the following:

- (Smoothness) $h_{\boldsymbol{\theta}}(u) = f_{\boldsymbol{\theta}}(u)r_{\boldsymbol{\theta}}(u) \in \mathbb{C}^{(m)}$; $h_{\boldsymbol{\theta}}^{(m)}$ is l_f -Lipschitz on I for all $\boldsymbol{\theta} \in \Theta_0$. Here m and l_f are defined in Assumption 3.
- (Lipschitz) *There exists a constant l_r such that $r_{\boldsymbol{\theta}_0} = 1 - F$ is l_r -Lipschitz, and for any $\epsilon > 0$,*

$$\sup_{\|\boldsymbol{\theta} - \boldsymbol{\theta}_0\|_2 \leq \epsilon, u \in I} |r'_{\boldsymbol{\theta}}(u) - r'_{\boldsymbol{\theta}_0}(u)| \leq l_r \epsilon.$$

Assumption 5. *The kernel K satisfies the following:*

- (Order- m kernel) $\int K(s)ds = 1$, $\int s^j K(s)ds = 0$ for $j \in \{1, \dots, m-1\}$, and that $\int |s^m K(s)|ds < +\infty$. Here m is the same as in Assumption 3.
- (Lipschitz) *Both $K(s)$ and $K'(s)$ are l_K -Lipschitz continuous with bounded support.*

The Assumptions 3-5 are quite standard assumptions in non-parametric statistics; see Fan and Gijbels [1996], Tsybakov [2008] for more details. Given these assumptions, we will prove that with high probability, the estimators $\hat{F}_k(u, \theta)$ and $\hat{F}_k^{(1)}(u, \theta)$ are sufficiently close to $F(u)$ and $F'(u)$ respectively given any $\theta \in \Theta_0$ for every sufficiently large k . Specifically, we obtain the desired error bound for $\hat{F}_k(u) = \hat{F}_k(u, \hat{\theta}_k)$ and $\hat{F}_k^{(1)}(u) = \hat{F}_k^{(1)}(u, \hat{\theta}_k)$.

Remark 7. Assumptions 3 and 4 can be relaxed in terms of the smoothness requirements: For all $m \geq 3$, we only need $f_\theta(u), h_\theta(u) \in \mathbb{C}^{(m-1)}$, and that $f_\theta^{(m-1)}(u), h_\theta^{(m-1)}(u)$ are ℓ -Lipschitz for some constant ℓ . For $m = 2$, we only need $f_\theta(u), h_\theta(u) \in \mathbb{C}^{(1)}$, and that the second order derivatives of $f_\theta(u), h_\theta(u)$ exist and are bounded. One is able to see assuming functions in $\mathbb{C}^{(m)}$ is a sufficient condition for the aforementioned conditions to hold, for the simplicity of our notations here, we keep the original assumptions.

Remark 8. If we only assume $F(\cdot)$ is ℓ -Lipschitz continuous (i.e. it may not be differentiable), we also provide an alternative algorithm in §H which achieves a regret upper bound $\tilde{O}(T^{3/4})$.

Remark 9. One is also able to estimate $F(u), F'(u)$ with the local polynomial estimator (see e.g. Fan and Gijbels [1996]). In this case, the assumptions can be weakened further. Specifically, the local polynomial estimators for F and F' enjoy all the theoretical guarantees given only the second part of Assumptions 3 and 5 instead of both Assumptions 3 and 5. For example, Lipschitz continuous density functions on $[-\delta_z, \delta_z]$ satisfy Assumption 4.2. The proof is very similar. For simplicity, we only focus on studying kernel regression in this paper.

Lemma 2. Under Assumptions 3, 4 and 5, there exist constants $B_{x,K}$, $B'_{x,K}$ and $C_{x,K}$ (depending only on the absolute constants within the assumptions) such that as long as

$$T \geq B_{x,K}(\log T + 2 \log d)^{\frac{4m-1}{m}} d^{\frac{2m-1}{m}},$$

we have for any $k \geq \lceil (\log(\sqrt{T} + \ell_0) - \log \ell_0) / \log 2 \rceil + 2$ and $\delta \in [4 \exp(-B'_{x,K} |I_k|^{\frac{2m}{2m+1}} / \log |I_k|), \frac{1}{2}]$, with probability at least $1 - 2\delta$,

$$\sup_{u \in I, \theta \in \Theta_k} |\hat{F}_k(u, \theta) - F(u)| \leq C_{x,K} |I_K|^{-\frac{m}{2m+1}} \sqrt{\log |I_K|} (\sqrt{d} + \sqrt{\log \frac{1}{\delta}}). \quad (18)$$

Here $I = [-\delta_z, \delta_z]$ and we choose the bandwidth $b_k = |I_k|^{-\frac{1}{2m+1}}$.

Lemma 3. *Under the same conditions as Lemma 2, with probability at least $1 - 4\delta$, we have*

$$\sup_{u \in I, \theta \in \Theta_k} |\hat{F}_k^{(1)}(u, \theta) - F'(u)| \leq \tilde{C}_{x,K} |I_K|^{-\frac{m-1}{2m+1}} \sqrt{\log |I_K|} (\sqrt{d} + \sqrt{\log \frac{1}{\delta}}). \quad (19)$$

We next develop a uniform upper bound for term $\mathbf{J}_1$ given in (11) for the k -th episode in Lemma 4 below.

Lemma 4. *Reinstating the notations and conditions in Lemma 2, with probability at least $1 - 6\delta$, we have*

$$\sup_{u \in [\delta_z, B - \delta_z]} |\hat{g}_k(u) - g(u)| \leq \bar{C}_{x,K} |I_K|^{-\frac{m-1}{2m+1}} \sqrt{\log |I_K|} (\sqrt{d} + \sqrt{\log \frac{1}{\delta}}).$$

Remark 10. In Algorithm 1 we define $\hat{g}_k(u) = u + \hat{\phi}_k^{-1}(-u)$ with $u \in [\delta_z, B - \delta_z]$. Thus, computing $\hat{g}_k(u)$ involves obtaining the inverse of $\hat{\phi}_k$, which is not necessarily monotonic. Nevertheless, it's not difficult to define or compute $\hat{\phi}_k^{-1}$. In fact, we'll show in the proof of Lemma 4 that $\hat{\phi}_k$ is very 'close' to ϕ in some main interval of interest, which contains $[\phi^{-1}(\delta_z - B), \phi^{-1}(-\delta_z)]$ and depends only on F . (Recall in Assumption 1 that ϕ' is bounded below from 0, so ϕ is strictly increasing). Thus, for any $u \in [\delta_z, B - \delta_z]$, the above fact will guarantee the existence of $\hat{\phi}_k^{-1}(-u)$ as some x within the interval such that $\hat{\phi}_k(x) = -u$.

Combining the above lemmas, which give us upper bounds for terms $\mathbf{J}_1, \mathbf{J}_2$ in every episode, we have the following Theorem 1, which provides an upper bound for the regret.

Theorem 1. *Under Assumptions 1, 3, 4 and 5, there exist constants $\bar{B}_{x,K}$, $\bar{B}'_{x,K}$ and $C_{x,K}^*$ (depending only on the absolute constants within the assumptions) such that for all T satisfying*

$$T \geq \max\{\bar{B}_{x,K}(\log T + 2 \log d)^{\frac{4m-1}{m-1}} d^{\frac{2m+1}{m-1}}, 4d^{\frac{2m+1}{m-1}}\},$$

the regret of Algorithm 1 over time T is no more than $C_{x,K}^(Td)^{\frac{2m+1}{4m-1}} \log T(1 + \log T/d)$.*

Remark 11. We note that Golrezaei et al. [2020] shares a similar framework with ours, although with a different regret measure. Specifically, we use a more traditional notion of regret by setting the

benchmark p_t^* from (5) with true θ_0 and $F(\cdot)$. In Golrezaei et al. [2020], the authors instead set the benchmark p_t^* so as to maximize the worst function in their function class $\mathcal{F}$, i.e.

$$p_t^* = \arg\max_{p \geq 0} \min_{F \in \mathcal{F}} p(1 - F(p - \theta_0^\top \tilde{\mathbf{x}}_t)).$$

Their optimal regret is of order $\tilde{\mathcal{O}}_d(T^{2/3})$, while ours is $\tilde{\mathcal{O}}_d(T^{\frac{2m+1}{4m-1}})$, which is closer to $\mathcal{O}_d(T^{1/2})$ when m is sufficiently large. Intuitively, a benchmark being the price maximizing the worst function is too conservative when their ambiguity function class is very large and the market noises are only sampled from a fixed distribution function in that function class, which is true in our semi-parametric setting.

On the other hand, Golrezaei et al. [2019] also work on similar but simpler settings, where they assume having unknown demanding curves but observable valuations instead of censored responses. By contrast, we work on a more common setting where the actual market values of products are unknown.

Remark 12. Both Algorithm 1 and Theorem 1 depend on the smoothness class of the function $F(\cdot)$. A popular choice in nonparametric curve estimation literature is $m = 2$, as other choices do not improve much for practical sample sizes. Nevertheless, we provide two ways to choose m that addresses a referee’s query.

- **Estimate m using cross-validation.** Specifically, we pick some relatively small m during the first episode. At each episode $k \geq 2$, before entering the exploration phase, we update the estimate of m using cross-validation [Hall and Racine, 2015] with the data gathered from the previous exploration phase. Then, we proceed with the main algorithm with this updated estimate until the next episode. For more details of the cross-validation procedure and the combined algorithm, see Section I.
- **Pick a constant pessimistic estimation of m .** In fact, we can directly fix a relatively small m (e.g. $m = 2$ or $m = 4$). In many cases, the performance of the algorithm ($\tilde{\mathcal{O}}((Td)^{5/7})$ and $\tilde{\mathcal{O}}((Td)^{3/5})$) will not be significantly different from where m is known (at least $\Omega((Td)^{1/2})$).

The above two ways can be applied to all settings in this paper as long as F is only required to be smooth to a finite degree.

4.2 Results under the setting with strong-mixing features

As mentioned in the introduction, we believe that in many situations, the dependence of features over time is inevitable. Thus, in this section, we generalize our results to the case where $\mathbf{x}_t$ can be dependent. For this purpose, we first impose the strong-mixing condition which measure the dependence between covariates over time.

Definition 1. [β -mixing] For a sequence of random vectors $\mathbf{x}_t \in \mathbb{R}^{d \times 1}$ on a probability space $(\Omega, \mathcal{X}, \mathbb{P})$, define β -mixing coefficient

$$\beta_k = \sup_{l \geq 0} \beta(\sigma(\mathbf{x}_t, t \leq l), \sigma(\mathbf{x}_t, t \geq l + k))$$

in which

$$\beta(\mathcal{A}, \mathcal{B}) = \frac{1}{2} \sup \left\{ \sum_{i \in I} \sum_{j \in J} |\mathbb{P}(A_i \cap B_j) - \mathbb{P}(A_i)\mathbb{P}(B_j)| \right\},$$

the maximum being taken over all finite partitions $(A_i)_{i \in I}$ and $(B_j)_{j \in J}$ of Ω with elements in $\mathcal{A}$ and $\mathcal{B}$.

The following assumption ensures that $\{\mathbf{x}_t\}_{t \geq 1}$ are not too strongly dependent. Combining with other assumptions, we ensure that the empirical covariance matrix $\frac{1}{n} \sum_{i=1}^n \tilde{\mathbf{x}}_i \tilde{\mathbf{x}}_i^\top$ concentrate around the population version, which is necessary in deriving the regret in every episode.

Assumption 6. The sequence $\mathbf{x}_t, t \geq 0$ are strictly stationary time series and follow β -mixing condition, in a sense we assume that $\beta_k \leq e^{-ck}$ holds with some constant c .

In order to derive the final regret upper bound under the strong-mixing setting, we also need an additional technical assumption stated below:

Assumption 7. Let $r_\theta(u_i, u_j) := \mathbb{E}[y_i y_j \mid w_j(\theta) = u_j, w_i(\theta) = u_i]$, $j > i \geq 0$, $r_\theta(u_j) := \mathbb{E}[y_j \mid w_j(\theta) = u_j]$, $j \geq 0$ be the joint regression function and marginal regression function. In addition, we also set $f_\theta(u_i, u_j)$, $j > i \geq 0$, $f_\theta(u_i)$, $i \geq 0$ as the joint density of $w_i(\theta)$ and $w_j(\theta)$ and marginal density of $w_i(\theta)$ respectively. Then we define $g_{1,\theta}(u_i, u_j) := r_\theta(u_i, u_j)f_\theta(u_i, u_j) - r_\theta(u_i)f_\theta(u_i)r_\theta(u_j)f_\theta(u_j)$

and $g_{2,\theta}(u_i, u_j) = f_\theta(u_i, u_j) - f_\theta(u_i)f_\theta(u_j)$. We assume $g_{1,\theta}(u_i, u_j)$ and $g_{2,\theta}(u_i, u_j)$ follow l -Lipschitz continuous condition, in a sense that

$$|g_{q,\theta}(u_i, u_j) - g_{q,\theta}(u'_i, u'_j)| \leq l\sqrt{(u_i - u'_i)^2 + (u_j - u'_j)^2}, \quad q \in \{1, 2\}$$

holds for all (u_i, u_j) , with $i, j \in [n]$ and $\theta \in \Theta_0$.

When the covariates $\mathbf{x}_i, \mathbf{x}_j$ are independent, we have $g_{q,\theta}(u_i, u_j) = 0, q \in \{1, 2\}$, for all (u_i, u_j) . Under such a mild assumption, we obtain a uniform upper bound of $|g_{q,\theta}(u_i, u_j)|$, which is dominated by the β -mixing constant $\beta_{j-i}^{1/3}$, for all $\theta \in \Theta_0$ and (u_i, u_j) (see Appendix F.7). Thus, this assumption essentially guarantees that the joint regression and density functions of the features still stay close to the products of their marginal ones even if they are correlated.

Following similar analysis with §4.1, we reach the following theorem which gives a regret upper bound at similar rate with Theorem 1 under the strong-mixing feature setting.

Theorem 2. *Let Assumptions 1, 3, 4, 5, 6 and 7 hold. Then there exist constants $B_{m,x,K}^*$ and $C_{m,x,K}^*$ (depending only on the absolute constants within the assumptions) such that for all T satisfying*

$$T \geq \max\{B_{m,x,K}^*(\log T + 2 \log d)^{\frac{12m-3}{m-1}}[(d+1) \log(d+1)]^{\frac{4m-1}{m-1}}/d^2, d^{\frac{2m+1}{m-1}}\}$$

the regret of Algorithm 1 over time T is no more than $C_{m,x,K}^(Td)^{\frac{2m+1}{4m-1}} \log^4 T$.*

4.3 Result on infinitely differentiable market noise distribution

In §4.1 and §4.2, we analyze the regret upper bounds when the noise distribution F has an m -th order continuous derivative, with any finite $m \geq 2$. The regret of our algorithm is of order $\tilde{\mathcal{O}}((Td)^{\frac{2m+1}{4m-1}})$, which gets closer to $\tilde{\mathcal{O}}(\sqrt{Td})$ as the degree of smoothness m goes to infinity. In fact, this is mainly due to inaccurate estimation of F and F' resulting from the bias of the kernel estimator. In this section, we deal with super smooth noise distributions [Fan, 1991], where F is infinitely differentiable. Under mild conditions, we're able to control the bias within $\mathcal{O}(1/|I_k|^{\frac{1}{2}})$ for each episode k by using extremely smooth kernels. As a reminder, here $|I_k|$ is the length of the k -th exploration phase. This leads to a $\tilde{\mathcal{O}}_d(\sqrt{T})$ regret bound in our algorithm. In particular, we assume the following:

Assumption 8. Define ϕ_θ , ξ_θ , $\phi_\theta^{(1)}$ and $\xi_\theta^{(1)}$ as the Fourier transform of the function f_θ , h_θ , f'_θ and h'_θ respectively:

$$\begin{aligned}\phi_\theta(s) &= \int_{-\infty}^{\infty} f_\theta(x) e^{isx} dx, \quad \xi_\theta(s) = \int_{-\infty}^{\infty} h_\theta(x) e^{isx} dx, \\ \phi_\theta^{(1)}(s) &= \int_{-\infty}^{\infty} f'_\theta(x) e^{isx} dx, \quad \xi_\theta^{(1)}(s) = \int_{-\infty}^{\infty} h'_\theta(x) e^{isx} dx,\end{aligned}$$

and $h_\theta(x) = f_\theta(x)r_\theta(x)$. There exist positive constant D_ϕ and d_ϕ and $\alpha > 0$ such that

$$\max\{|\phi_\theta(s)|, |\xi_\theta(s)|, |\phi_\theta^{(1)}(s)|, |\xi_\theta^{(1)}(s)|\} \leq D_\phi e^{-d_\phi |s|^\alpha}$$

for all $s \in \mathbb{R}$.

Remark 13. -This assumption is quite standard, and ensures that $f_\theta(u)$, $F_\theta(u) \in \mathbb{C}^\infty$. The class of functions are still infinite dimensional nonparametric functions. The class of supersmooth functions has been used in non-parametric density literature. In particular, it has been used in [Fan \[1991\]](#) for characterizing the difficulty of non-parametric deconvolution.

Under the Assumption of [8](#), for each episode k , we can successfully control the bias within $\mathcal{O}(1/\sqrt{|I_k|})$ via an infinite order kernel [[McMurphy and Politis, 2004](#), [Berg and Politis, 2009](#)]. In order to construct an infinite order kernel K , we simply let K be the Fourier inverse transform of some ‘well-behaved’ function. In particular, let

$$K(x) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \kappa(s) e^{-isx} ds, \tag{20}$$

be the Fourier inversion of κ satisfying

$$\kappa(s) = \begin{cases} 1, & |s| \leq c_\kappa \\ g_\infty(|s|), & \text{otherwise.} \end{cases}$$

Here g_∞ is any continuous, square-integrable function that is bounded in absolute value by 1 and satisfies $g_\infty(|c_\kappa|) = 1$. This defines an infinity order kernel function [[Fan and Gijbels, 1996](#)].

By plugging the infinite order kernel K into our algorithm, we’re able to obtain the following lemma:

Lemma 5. Under Assumption 8, there exists a positive constant C_{inf} depending only on α , D_ϕ and d_ϕ such that for all kernel K satisfying (20), for each episode k , by choosing the bandwidth $b_k = c_\kappa(d_\phi/\log |I_k|)^{1/\alpha}$ in (15) and (17), we have

$$\begin{aligned} \sup_{u \in I, \theta \in \Theta_k} |\mathbb{E}[f_k(u, \theta)] - f_\theta(u)| &\leq \frac{C_{\text{inf}}}{\sqrt{|I_k|}}, & \sup_{u \in I, \theta \in \Theta_k} |\mathbb{E}[h_k(u, \theta)] - h_\theta(u)| &\leq \frac{C_{\text{inf}}}{\sqrt{|I_k|}}, \\ \sup_{u \in I, \theta \in \Theta_k} |\mathbb{E}[f_k^{(1)}(u, \theta)] - f'_\theta(u)| &\leq \frac{C_{\text{inf}}}{\sqrt{|I_k|}}, & \sup_{u \in I, \theta \in \Theta_k} |\mathbb{E}[h_k^{(1)}(u, \theta)] - h'_\theta(u)| &\leq \frac{C_{\text{inf}}}{\sqrt{|I_k|}}. \end{aligned}$$

Following similar proof procedures of Theorems 1 and 2, Lemma 5 leads to the following theorem, which gives a regret upper bound of $\tilde{\mathcal{O}}_d(\sqrt{T})$, achieving the same convergence rate with the parametric case up to logarithmic terms [Javanmard and Nazerzadeh, 2019].

Theorem 3. Let Assumptions 1, 3, 4, 5, 6, 7 and 8 hold. Then there exist constants B_{inf}^* and C_{inf}^* (depending only on the absolute constants within the assumptions) such that by choosing $|I_k| = \lceil \sqrt{l_k d} \rceil$ instead in Algorithm 1, for all T satisfying

$$T \geq B_{\text{inf}}^* d^2 (\log T + 2 \log d)^{12+12/\alpha} \log^4(d+1),$$

the regret of the algorithm over time T is no more than $C_{\text{inf}}^* (Td)^{\frac{1}{2}} (\log T)^{\frac{3}{2} + \frac{3}{2\alpha}} [\log(d+1) + \log T/d]$.

Remark 14. Theorem 3 partly overturns the conjecture in Shah et al. [2019] that there is no policy can achieve an $\tilde{\mathcal{O}}_d(\sqrt{T})$ regret under the setting where the market value is linear in the features as in (2). We provide a regime with super smooth market noise in which $\tilde{\mathcal{O}}_d(\sqrt{T})$ regret upper bound is attainable by our policy.

Remark 15. In §B, we provide discussions on minimax lower bound, adversarial agents, inference for the demand, practical implementation, and extensions to the high dimensional setting.

5 Simulations

5.1 Justification of theoretical results

In this section, we illustrate the performance of our policy through large-scale simulations under various settings. Recall our model (2), where $\mathbf{x}_t \in \mathbb{R}^d$ and z_t follows distributions with bounded support

and smooth c.d.f. Throughout this section, we let the dimension $d = 3$ and the coefficients $\alpha_0 = 3$, $\beta_0 = \sqrt{2/3} \cdot \mathbf{1}_{3 \times 1}$. For each value of smoothness degree $m \in \{2, 4, 6\}$, we fix a density function from $\mathbb{C}^{(m-1)}$ for all z_t (thus the c.d.f. F belongs to $\mathbb{C}^{(m)}$). Specifically, we set the p.d.f. of z_t as $f_m(x) \propto (1/4 - x^2)^{m/2} \cdot \mathbb{I}_{\{|x| \leq 1/2\}}$ for $m \in \{2, 4, 6\}$. Moreover, for each m , the covariates $\mathbf{x}_t \in \mathbb{R}^3$ are generated from a p.d.f. in $\mathbb{C}^{(m)}$ in the following ways:

- **i.i.d. $\mathbf{x}_t$ with independent entries:** Each coordinate of $\mathbf{x}_t$ is generated from density $f_m(x) \propto (2/3 - x^2)^{m+1} \cdot \mathbb{I}_{\{|x| \leq \sqrt{2/3}\}}$.
- **i.i.d. $\mathbf{x}_t$ with dependent entries:** $\mathbf{x}_t$ is generated from the density function $f_m(\mathbf{x}) \propto (1 - \mathbf{x}^\top \Sigma^{-1} \mathbf{x})^{m+1}$. Here Σ is a positive definite matrix with (i, j) -th entry being equal to $0.2^{|i-j|}$, $1 \leq i, j \leq 3$.
- **Strong mixing $\mathbf{x}_t$ with dependent entries:** We generate $\mathbf{x}_t$ from the VAR (vector autoregression) model, where $\mathbf{x}_t = \mathbf{A}\mathbf{x}_{t-1} + \mathbf{B}\mathbf{x}_{t-2} + \boldsymbol{\xi}_t$. Here $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{3 \times 3}$ with $\mathbf{A}_{i,j} = 0.4^{|i-j|+1}$, $\mathbf{B}_{i,j} = 0.1^{|i-j|+1}$, $i, j \in \{1, 2, 3\}$. In addition, $\{\boldsymbol{\xi}_t\}_{t \geq 1}$ are i.i.d. with density $f_m(\boldsymbol{\xi}) \propto (1 - \boldsymbol{\xi}^\top \Sigma^{-1} \boldsymbol{\xi})^{m+1}$ where the Σ is the same as the one given in (ii).

When implementing our algorithm, we divide the time horizon into consecutive episodes by setting the length of the k -th episode as $\ell_k = 2^{k-1} \ell_0$ with $k \in \mathbb{N}^+$ and $\ell_0 = 200$. We further separate every episode into an exploration phase with length $|I_k| = \min\{(d\ell_k)^{(2m+1)/(4m-1)}, \ell_k\}$ depending on the values of m and d . The exploitation phase contains the rest of the time in that episode. In the exploration phase, we sample p_t from $\text{Unif}(0, B = 6)$, since $B = 6$ is a valid upper bound of v_t . In the exploitation phase, we set the kernels as follows: For any given $m \in \{2, 4, 6\}$ prefixed at the beginning of the algorithm, we choose the kernel function with m -th order. Here we choose the second, fourth, sixth-order kernel functions as $K_2(u) = 35/12(1 - u^2)^3 \cdot \mathbb{I}_{\{|u| \leq 1\}}$, $K_4(u) = 27/16(1 - 11/3u^2) \cdot K_2(u)$ and $K_6(u) = 297/128(1 - 26/3u^2 + 13u^4) \cdot K_2(u)$ respectively. In episode k , we set the bandwidth b_k as $3 \cdot |I_k|^{-\frac{1}{2m+1}}$ in (14) and (16) according to the settings in the theoretical analysis. In reality, one can also tune the bandwidth by using cross validation at the end of every exploration phase. Moreover, when calculating $p_t = \hat{g}(\tilde{\mathbf{x}}_t^\top \hat{\boldsymbol{\theta}}_k) = \tilde{\mathbf{x}}_t^\top \hat{\boldsymbol{\theta}}_k + \hat{\phi}_k^{-1}(-\tilde{\mathbf{x}}_t^\top \hat{\boldsymbol{\theta}}_k)$, we find $\hat{\phi}_k^{-1}(-\tilde{\mathbf{x}}_t^\top \hat{\boldsymbol{\theta}}_k)$ as follows: First, we look for




Figure 1: Regret log-log plot in the setting with i.i.d. covariates with independent entries. The three subplots show the case $m \in [2, 4, 6]$ respectively. The x-axis is $\log(T) - \log(1500)$ for $T \in [1500, 2000, 3100, 4000, 5000, 6300]$, while the y-axis is $\widetilde{\text{reg}}(T) := \log(\text{reg}(T)) - 2 \log \log T - (\log(\text{reg}(1500)) - 2 \log \log 1500)$. The solid blue, green and red lines represent the mean $\widetilde{\text{reg}}(T)$ of the Algorithm 1 with unknown $g(\cdot)$ and θ_0 , unknown $g(\cdot)$ but known θ_0 , and known $g(\cdot)$ but unknown θ_0 respectively over 30 independent runs. The light color areas around those solid lines depict the standard error of our estimation of $\log(\text{reg}(T)) - 2 \log \log T$. The dashed black lines in (a) – (c) represents the benchmark whose slopes are equal to $\frac{2m+1}{4m-1}$ with $m \in \{2, 4, 6\}$.

$x \in [-1, 1]$ such that $\hat{\phi}_k(x) = -\tilde{\mathbf{x}}_t^\top \hat{\theta}_k$ (The interval $[-1, 1]$ contains the true support of $\phi(x)$ $[-0.5, 0.5]$, since in reality, we might only know a range of the true support). Then, we do a transformation of variable x to $x(y) = -2 \cdot \exp(y)/(1 + \exp(y)) + 1$ and solve y as the root of $\hat{\phi}_k(x(y)) + \tilde{\mathbf{x}}_t^\top \hat{\theta}_k = 0$ by using Newton's method starting at $y = 0$. Finally, we set $x = -2 \cdot \exp(y)/(1 + \exp(y)) + 1$ as $\hat{\phi}_k^{-1}(-\tilde{\mathbf{x}}_t^\top \hat{\theta}_k)$ and offer p_t according to the algorithm.

For any given $m \in \{2, 4, 6\}$, under the three covariate settings discussed above, we input m into the algorithm, select the corresponding kernel and repeat Algorithm 1 for 30 times until $T = 6300$. For each $T \in [1500, 2000, 3100, 4000, 5000, 6300]$, we record the cumulative regret $\text{reg}(T)$. For the first two covariate settings, recall from Theorem 1 that the regret $\text{reg}(T) \lesssim T^{\frac{2m+1}{4m-1}} \log^2 T$. Thus, we plot $\widetilde{\text{reg}}(T)$ against $\log(T) - \log(1500)$ in Figure 1, 7, 8, where $\widetilde{\text{reg}}(T) := \log(\text{reg}(T)) - 2 \log \log T - (\log(\text{reg}(1500)) - 2 \log \log 1500)$;

From Figures 1, 7, 8, we conclude that under all settings, the rates of the empirical regrets' in-

crements produced by Algorithm 1 (as shown by the solid blue lines) do not exceed their theoretical counterparts given in Theorems 1 and 2 (as shown by the dashed black lines). In many cases, the growth rates of the empirical regrets are very close to those of the theoretical lines. This demonstrates the tightness of our theoretical results. Moreover, as all the solid lines have similar growth rates, we show that Algorithm 1 is robust to the estimation of θ_0 and $g(\cdot)$. This is further proved in Appendix G, where we directly plot $\text{reg}(T)$ for all the settings discussed here. See Appendix G for more plots and discussions.

5.2 Comparison with other methods

In this subsection, we provide numerical studies which illustrate differences between our methods and two highly related prior arts ('RMLP-2' and 'Bandit') using both synthetic and real data. Here, 'RMLP-2' is the policy proposed in Javanmard and Nazerzadeh [2019] that solves the same problem as ours except that the noise distribution falls in a *parametric* function class. In addition, we denote the policy proposed in Kleinberg and Leighton [2003] as 'Bandit', which leverages a variant of UCB algorithm under non-parametric noise distribution that achieves $\mathcal{O}(\sqrt{T})$ regret *without* modeling covariate information.

We first use synthetic data to illustrate the efficiency of our method over 'RMLP-2' and 'Bandit'. For each smoothness degree $m = \{2, 4, 6\}$, we generate our data following the same way given in §5.1, except that we only generate the distribution of $\mathbf{x}_t$ according to the first option discussed in §5.1. We illustrate the performance of our method against those two prior arts in the following figures. Here we follow Algorithm 4 which uses a data-driven way to determine m before every episode. For RMLP-2, since there is no way the algorithm knows the true noise distribution, we instead assume the noise falls into a Gaussian distribution when executing the algorithm.

We see from the simulation results that the regret we achieved is much smaller than those two benchmarks. As for the comparison with RMLP-2, our method is robust to the mis-specification of the parametric function class since our algorithm can adapt to all functions in the non-parametric class. For the comparison with 'Bandit', we see that only using the non-parametric bandit algorithm without considering the contextual information (heterogeneity of product) will lose much efficiency in gaining




Figure 2: Regret Comparison between our methods and two benchmarks (RMLP-2 and Bandit). From the left to the right, the true underlying degree of smoothness is $m = \{2, 4, 6\}$ respectively. The x-axis denotes the time stamp T ranges from 1 $\sim$ 12000, and the y-axis denotes the regret at the time T defined in (6). We repeat the experiment 30 times and record the averaged regrets (solid lines) and standard errors (light areas) of every policy. The blue line denotes the regret of our policy (in Algorithm 1) with knowing degree of smoothness m and the orange line represents the regret of our policy (given in Algorithm 4) without knowing degree of smoothness m . The green and red lines are the regrets of implementing 'RMLP-2' and 'Bandit' policy respectively.

revenue.

5.2.1 Real Application

Next, we leverage a simulation based on the real data to further illustrate the merits of our Algorithm over 'RMLP-2' and 'Bandit'.

We use the real-life auto loan dataset provided by the Center for Pricing and Revenue Management at Columbia University. This dataset is used by several related works [Phillips et al., 2015, Ban and Keskin, 2020, Luo et al., 2021, Wang et al., 2020a] and many others. The dataset contains 208,085 auto loan applications received from July 2002 to November 2004. Some features such as the amount of loan, the borrower's information is contained in that dataset. We adopt the feature selection in the same way with Ban and Keskin [2020], Luo et al. [2021], Wang et al. [2020a] and consider the




Figure 3: Comparison between our policy and ‘RMLP-2’ and ‘Bandit’ based on real data application.

following four features: the loan amount approved, FICO score, prime rate and competitor’s rate. As for the price variable, we also computed it in the same way with the aforementioned literature, where $p_t = \text{Monthly Payment} \cdot \sum_{t=1}^{\text{Term}} (1 + \text{Rate})^{-t} - \text{Loan Amount}$. The rate is set as 0.12%, which is an approximate average of the monthly London interbank rate for the studied time period. Moreover, this dataset also records purchasing decision of the borrowers given the price set by the lender. For more details on this dataset, please refer to Phillips et al. [2015], Ban and Keskin [2020].

Note that one is not able to obtain online responses to any algorithms, thus, we follow the calibration idea proposed in Ban and Keskin [2020], Luo et al. [2021], Wang et al. [2020b] to first estimate the binary choice model and leverage it as the ground truth to conduct online numerical experiments. To be more specific, we first scale all variables into the scale of $[0, 1]$ (since the prediction results of single index model won’t be affected by scale of the covariates). We randomly sample 5000 data points, estimate θ_0 and F using semi-parametric estimation tools from these data. We next treat them as the underlying true parameters for our binary choice model stated in (3). Given these key components, the remaining experiments remain almost the same as discussed in §5.1 and §5.2, except that here we set θ_0 , distribution $F(\cdot)$ as the estimated one given above and sample $\mathbf{x}_t$ from those four features above. We set $B_0 = 4$, $\ell_0 = 200$ and conduct Algorithm 4 (in this algorithm, we use cross-validation to select m at the beginning of every episode, details are given in Algorithm 5).

We next compare Algorithm 4 with ‘RMLP-2’ and ‘Bandit’ policies. The details are given in Figure 3. To summarize, our policy outperforms the RMLP-2 [Javanmard and Nazerzadeh, 2019] and

non-parametric bandit policy [Kleinberg and Leighton, 2003] in terms of both the regret performance and the ability to adapt to different noise distributions.

6 Conclusion

In this paper, we study the contextual dynamic pricing problem where the market value is linear in features, and the market noise has unknown distribution. We propose a policy that combines semi-parametric statistical estimation and online decision making. Our policy achieves near optimal regret, and is close to the regret lower bound where the market noise distribution belongs to a parametric class. We further generalize these results to the case when the product features satisfy the strong mixing condition. The practical performance of the algorithm is proved by extensive simulations.

There are several directions worth exploring in the future. First, we conjecture that the estimation accuracy of the market noise distribution F is crucial in the regret. Thus, within the function class $F \in \mathbb{C}^{(m)}$, we conjecture that a tighter regret lower bound $\Omega_d(T^{\frac{2m+1}{4m-1}})$ can be achieved instead of $\Omega_d(\sqrt{T})$, namely, our procedure is optimal. Second, in this work, we consider a linear model for the market value. In case a more complex model is appropriate, it's possible to extend our methodology to where the market value is nonlinear in product features, e.g. $v_t = \phi(\boldsymbol{\theta}_0^\top \mathbf{x}_t) + z_t$ or other structured statistical machine learning model such as the additive model $v_t = f_1(x_{t1}) + \dots + f_d(x_{td}) + z_t$. Finally, it's worth studying similar pricing problems with adversarial or strategic buyers, which is potentially more suitable in some specific applications.

References

- D. Alexey. Optimal non-parametric learning in repeated contextual auctions with strategic buyer. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 2668–2677. PMLR, 13–18 Jul 2020.
- K. Amin, A. Rostamizadeh, and U. Syed. Repeated contextual auctions with strategic buyers. In *Advances in Neural Information Processing Systems*, volume 27, 2014.
- Z. Anton and D. Alexey. Bisection-based pricing for repeated contextual auctions against strategic buyer. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 11469–11480. PMLR, 2020.

- M. Babaioff, S. Dughmi, R. Kleinberg, and A. Slivkins. Dynamic pricing with limited supply. *The ACM Transactions on Economics and Computation*, 2015. doi: 10.1145/2559152.
- G. Ban and N. Keskin. Personalized dynamic pricing with machine learning: High dimensional features and heterogeneous elasticity. *Management Science*, 67(9):5549–5568, 2020.
- A. Berg and D. Politis. Cdf and survival function estimation with infinite-order kernels. *Electronic Journal of Statistics*, 3:1436–1454, 2009.
- O. Besbes and A. Zeevi. Dynamic pricing without knowing the demand function: Risk bounds and near-optimal algorithms. *Operations Research*, 57(6):1407–1420, 2009.
- J. Broder and P. Rusmevichientong. Dynamic pricing under a general parametric choice model. *Operations Research*, 60(4):965–980, 2012.
- R. J. Carroll, J. Fan, I. Gijbels, and M. P. Wand. Generalized partially linear single-index models. *Journal of the American Statistical Association*, 92(438):477–489, 1997.
- N. Cesa-Bianchi, T. Cesari, and V. Perchet. Dynamic pricing with finitely many unknown valuations. volume 98 of *Proceedings of Machine Learning Research*, pages 247–273. PMLR, 2019.
- N. Chen and G. Gallego. Nonparametric pricing analytics with customer covariates. *Operations Research*, 69(3):974–984, 2020.
- Q. Chen, S. Jasin, and I. Duenyas. Nonparametric self-adjusting control for joint learning and optimization of multiproduct pricing with finite resource capacity. *Mathematics of Operations Research*, 44(2):601–631, 2019.
- X. Chen, D. Simchi-Levi, and Y. Wang. Privacy-preserving dynamic personalized pricing with demand learning. *Management Science*, 2021.
- X. Chen, Z. Owen, C. Pixton, and D. Simchi-Levi. A statistical learning approach to personalization in revenue management. *Management Science*, 68(3):1923–1937, 2022.
- M. C. Cohen, I. Lobel, and R. Paes Leme. Feature-based dynamic pricing. *Management Science*, 66(11):4921–5484, 2016.
- M. Delecroix, W. Härdle, and M. Hristache. Efficient estimation in conditional single-index regression. *Journal of Multivariate Analysis*, 86(2):213–226, 2003.
- A. V. den Boer. Dynamic pricing and learning: historical origins, current research, and new directions. *Surveys in Operations Research and Management Science*, 20(1):1–18, 2015.
- A. V. den Boer and B. Zwart. Simultaneously learning and optimizing using controlled variance pricing. *Management Science*, 60(3):770–783, 2014.
- A. V. den Boer and B. Zwart. Mean square convergence rates for maximum quasi-likelihood estimators. *Stochastic systems*, 4(2):375–403, 2015.

- J. Fan. On the optimal rates of convergence for nonparametric deconvolution problems. *The Annals of Statistics*, pages 1257–1272, 1991.
- J. Fan and I. Gijbels. *Local polynomial modelling and its applications*. Chapman and Hall, 1996.
- J. Fan and R. Li. New estimation and model selection procedures for semiparametric modeling in longitudinal data analysis. *Journal of the American Statistical Association*, 99(467):710–723, 2004.
- J. Fan and Q. Yao. *Nonlinear Time Series: Nonparametric and Parametric Methods*. Springer, 2003.
- J. Fan, N. E. Heckman, and M. P. Wand. Local polynomial kernel regression for generalized linear models and quasi-likelihood functions. *Journal of the American Statistical Association*, 90(429):141–150, 1995.
- N. Golrezaei, P. Jaillet, and J. C. N. Liang. Incentive-aware contextual pricing with non-parametric market noise. *arXiv:1911.03508*, 2019.
- N. Golrezaei, A. Javanmard, and V. Mirrokni. Dynamic incentive-aware learning: Robust pricing in contextual auctions. *Operations Research*, 69(1):297–314, 2020.
- L. Györfi, A. Krzyżak, M. Kohler, and H. Walk. *A distribution-free theory of nonparametric regression*. Springer, 2002.
- P. G. Hall and J. S. Racine. Infinite order cross-validated local polynomial regression. *Journal of Econometrics*, 185(2):510–525, 2015.
- W. Härdle, P. Hall, and H. Ichimura. Optimal Smoothing in Single-Index Models. *The Annals of Statistics*, 21(1):157 – 178, 1993.
- J. L. Horowitz. *Semiparametric methods in econometrics*, volume 131. Springer Science & Business Media, 2012.
- J. L. Horowitz and W. Härdle. Direct semiparametric estimation of single-index models with discrete covariates. *Journal of the American Statistical Association*, 91(436):1632–1640, 1996.
- H. Ichimura. Semiparametric least squares (sls) and weighted sls estimation of single-index models. *Journal of Econometrics*, 58(1):71–120, 1993. ISSN 0304-4076.
- A. Javanmard. Perishability of data: Dynamic pricing under varying-coefficient models. *The Journal of Machine Learning Research*, 18(53):1–31, 2017.
- A. Javanmard and H. Nazerzadeh. Dynamic pricing in high-dimensions. *The Journal of Machine Learning Research*, 20(1):315–363, 2019.
- A. Javanmard, H. Nazerzadeh, and S. Shao. Multi-product dynamic pricing in high-dimensions with heterogeneous price sensitivity. In *2020 IEEE International Symposium on Information Theory (ISIT)*, pages 2652–2657, 2020.

- N. B. Keskin and A. Zeevi. Dynamic pricing with an unknown demand model: Asymptotically optimal semi-myopic policies. *Operations Research*, 26(5):1142–1167, 2014.
- R. W. Klein and R. H. Spady. An efficient semiparametric estimator for binary response models. *Econometrica: Journal of the Econometric Society*, 61:387–421, 1993.
- R. Kleinberg and T. Leighton. The value of knowing a demand curve: bounds on regret for online posted-price auctions. *44th Annual IEEE Symposium on Foundations of Computer Science*, pages 594–605, 2003.
- P. R. Leme and J. Schneider. Contextual search via intrinsic volumes. In *2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 268–282, 2018.
- X. Li and Z. Zheng. Dynamic pricing with external information and inventory constraint. *Available at SSRN*, 2020.
- A. Liu, R. P. Leme, and J. Schneider. Optimal contextual pricing and extensions. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1059–1078, 2021.
- Y. Luo, W. W. Sun, and Y. Liu. Distribution-free contextual dynamic pricing. *arXiv:2109.07340*, 2021.
- Y. P. Mack and B. W. Silverman. Weak and strong uniform consistency of kernel regression estimates. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 61:405–415, 1982.
- B. K. Mallick and A. E. Gelfand. Generalized linear models with unknown link functions. *Biometrika*, 81(2): 237–245, 1994.
- J. Mao, R. Leme, and J. Schneider. Contextual pricing for lipschitz buyers. In *Advances in Neural Information Processing Systems*, volume 31, pages 5643–5651, 2018.
- E. Masry. Multivariate local polynomial regression for time series: uniform strong consistency and rates. *Journal of Time Series Analysis*, 17:571–599, 1996.
- C. E. McCulloch. Generalized linear models. *Journal of the American Statistical Association*, 95(452):1320–1324, 2000.
- L. T. McMurry and N. D. Politis. Nonparametric regression with infinite order flat-top kernels. *Journal of Nonparametric Statistics*, 16(3-4):549–562, 2004.
- S. Miao, X. Chen, X. Chao, J. Liu, and Y. Zhang. Context-based dynamic pricing with online clustering. *ArXiv:1902.06199*, 2019.
- V. V. Misic and G. Perakis. Data analytics in operations management: A review. *Manufacturing & Service Operations Management*, 22(1):158–169, 2020.
- E. A. Nadaraya. On estimating regression. *Theory of Probability and Its Applications*, 9(1):141–142, 1964.
- M. Nambiar, D. Simchi-Levi, and H. Wang. Dynamic learning and pricing with model misspecification. *Management Science*, 65(11):4980–5000, 2019.

- R. L. Phillips, A. S. Simsek, and G. J. v. Ryzin. The effectiveness of field price discretion: Empirical evidence from auto lending. *Management Science*, 61:1741–1759, 2015.
- J. L. Powell, J. H. Stock, and T. M. Stoker. Semiparametric estimation of index coefficients. *Econometrica: Journal of the Econometric Society*, 57(6):1403–1430, 1989.
- S. Qiang and M. Bayati. Dynamic pricing with demand covariates. *Stochastic Models eJournal*, 2016.
- D. Ruppert, M. P. Wand, and R. J. Carroll. *Semiparametric regression*. Cambridge university press, 2003.
- P. Rusmevichientong, B. Van Roy, and P. W. Glynn. A nonparametric approach to multiproduct pricing. *Operations Research*, pages 82–98, 2006.
- V. Shah, R. Johari, and J. Blanchet. Semi-parametric dynamic contextual pricing. In *Advances in Neural Information Processing Systems*, volume 32, pages 2363–2373, 2019.
- B. W. Silverman. Weak and Strong Uniform Consistency of the Kernel Estimate of a Density and its Derivatives. *The Annals of Statistics*, 6(1):177 – 184, 1978.
- C. Stone. Optimal rates of convergence for nonparametric estimators. *The Annals of Statistics*, 8(6):1348–1360, 1980.
- C. J. Stone. Optimal Global Rates of Convergence for Nonparametric Regression. *The Annals of Statistics*, 10(4):1040 – 1053, 1982.
- W. Tang, C.-J. Ho, and Y. Liu. *Differentially Private Contextual Dynamic Pricing*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 2020.
- A. B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer Publishing Company, Incorporated, 2008. ISBN 0387790519.
- C.-H. Wang, Z. Wang, W. W. Sun, and G. Cheng. Online regularization for high-dimensional dynamic pricing algorithms. *arXiv preprint arXiv:2007.02470*, 2020a.
- Y. Wang, X. Chen, X. Chang, and D. Ge. Uncertainty quantification for demand prediction in contextual dynamic pricing. *Production and Operations Management*, 30, 12 2020b.
- Z. Wang, S. Deng, and Y. Ye. Close the gaps: A learning-while-doing algorithm for single-product revenue management problems. *Operations Research*, 62(2):318–331, 2014.
- G. S. Watson. Smooth regression analysis. *Sankhyā: The Indian Journal of Statistics, Series A.*, 26(4):359–372, 1964.
- M. M. Wei and F. Zhang. Recent research developments of strategic consumer behavior in operations management. *Computers and Operations Research*, 93:166–176, 2018. ISSN 0305-0548.
- S. Weisberg and A. H. Welsh. Adapting for the missing link. *The Annals of Statistics*, 22(4):1674–1700, 1994.
- Y. Xia and W. K. Li. On single-index coefficient regression models. *Journal of the American Statistical Association*, 94(448):1275–1285, 1999.

Supplement material for “Policy Optimization Using Semi-parametric Models for Dynamic Pricing”

A Addition Figures




Figure 7: Regret log-log plot in the setting with i.i.d. covariates with dependent entries. The remaining caption is the same as Figure 1.




Figure 8: Regret log-log plot in the setting with strong mixing covariates. The remaining caption is the same as Figure 1.

B Discussion

1. [Minimax Lower Bound] Our work shares a similar setting with [Broder and Rusmevichientong \[2012\]](#), in which they study a general choice model with parametric structure and binary

response, but without any covariates. A lower bound of order $\Omega(\sqrt{T})$ is established by constructing an ‘uninformative price’ in their work. To be more precise, an uninformative price is a price that all demand curves (probability of successful sales) as offered price indexed by unknown parameters intersect. Namely, the demands at this uninformative price are the same for all unknown parameters. In addition, such price is also the optimal price with some parameters. In this case, the price is uninformative because it doesn’t reveal any information on the true parameter. Intuitively, if one tries to learn model parameters, the only way is to offer prices that are sufficiently far from the uninformative price (optimal price) which leads to a larger regret.

Borrowing the idea from [Broder and Rusmevichientong \[2012\]](#) and [Javanmard and Nazerzadeh \[2019\]](#), we deduce that there exists an ‘uninformative price’ in the following class of models: Consider a class of distributions $\mathcal{F}$ which satisfies Assumption 1:

$$\mathcal{F} := \{F_\sigma : \sigma > 0, F_\sigma = F(x/\sigma)\}.$$

Here, F is the c.d.f. of a known distribution with mean zero. Moreover, we assume the support of F'_σ is contained in $[-a, a]$ (For instance, the class of distributions with density $f_\sigma(x) = 4/(3\sigma^3)(\sigma - x)^k(\sigma + x)^k \cdot \mathbb{I}_{\{|x| \leq \sigma\}}$, $k \geq 1$ or $f_\sigma(x) = C_\sigma \exp\left(-\frac{\sigma^2}{\sigma^2 - x^2}\right) \cdot \mathbb{I}_{\{|x| \leq \sigma\}}$ with $\sigma \leq a$ etc.)

Let $\beta = 1/\sigma$ and multiply β on both sides of (2), which leads to

$$\tilde{v}(\mathbf{x}_t) = \tilde{\beta}_0^\top \mathbf{x}_t + \tilde{\alpha}_0 + \tilde{z}_t.$$

Here, $\tilde{v}_t = \beta v_t$, $\tilde{\beta}_0 = \beta \beta_0$, $\tilde{\alpha}_0 = \beta \alpha_0$ and $\tilde{z}_t = \beta z_t$. The distribution of $\tilde{z}_t$ is F_1 , which is denoted as F here for convenience. Next, in our sub-parameter class, we first let $\beta_0 = 0$ and fix a number ξ with $F'(\xi) \neq 0$. Then we choose a collection of $\{(\sigma, \alpha_0)\}$ which satisfies $\beta = 1/\sigma = (\xi + \tilde{\alpha}_0)$. Following the same arguments as in [Javanmard and Nazerzadeh \[2019\]](#), one can prove that $p = 1$ is indeed an uninformative price. Since in the sub-parametric class given above, all demand curves intersect at a point $1 - F(\xi)$ when $p = 1$, and for a special $(\sigma, \alpha_0) = (1/(\xi - \phi(\xi)), -\phi(\xi)/(\xi - \phi(\xi)))$, $p = 1$ is the optimal price. Thus the $\Omega(\sqrt{T})$ lower bound applies.

Remark 1. When we only consider explore-then-commit algorithms and offer price as $p_t = \hat{\phi}_k^{-1}(-\mathbf{x}_t^\top \hat{\theta}) + \mathbf{x}_t^\top \hat{\theta}$, with $\hat{\phi}_k(u) = u - \frac{1 - \hat{F}_k(u)}{\hat{F}^{(1)}(u)}$, the optimality of p_t reduces to the optimality of estimating $F(\cdot)$, $f(\cdot)$ and θ . According to Stone [1980, 1982], Tsybakov [2008], the statistical rates of our estimators on $\hat{F}$, $\hat{F}^{(1)}$ and $\hat{\theta}$ are minimax optimal in every episode. Thus, our posted price is optimal constrained on this type of policies. However, if we consider a general policy class, there is currently no lower bound for feature-based pricing given unknown noise distribution with finite smoothness degree besides the general $\sqrt{T}$ lower bound mentioned above. It remains an open problem whether our upper bound is tight for finite m .

2. [The adversarial setting] We note that in some real applications with potentially adversarial contexts, the covariance of the feature vectors might be singular or ill-conditioned (e.g. due to repeated buyers recorded in $\mathbf{x}_t$). However, our algorithm can be adjusted to cope with such situations. The key observation here is that this assumption is *only* required in our exploration phase: For any k , we allow arbitrary $\mathbf{x}_t$ in the k -th exploitation phase, since we have already obtained accurate estimators $\hat{\boldsymbol{\theta}}_k$ and $\hat{g}_k(\cdot)$ for $\boldsymbol{\theta}_0$ and $g(\cdot)$. Therefore, whenever there is a sign of a repeated buyer, we can modify our algorithm slightly by using the $\hat{g}_{k-1}(\cdot)$ in the last episode to offer a price, and then move this buyer to the corresponding exploitation phase. If the number of similar buyers in the k -th episode is ℓ_k^r with any $r < 1$ and we assume the remaining buyers are sampled i.i.d. from a distribution, we are still able to proceed by only arranging some contexts with similar buyers into the exploitation phase directly. This matches with some real situation in online shopping where personal preference features will be recorded by the seller in order to make recommendation in the future.
3. [Online inference of the demand] Recently, Wang et al. [2020] use a de-biased approach to quantify the uncertainty of the demand function in a parametric class which offers new insight to the field of statistical decision making.

In our work, we combine the non-parametric statistical estimation and online decision making to derive a policy that maximize the seller's revenue. We next also briefly discuss our intuition on depicting the uncertainty of the demand curve in a non-parametric class. Recall the demand curve given in (4). For given $p, \mathbf{x}$, and estimators $\hat{F}_k, \hat{\boldsymbol{\theta}}_k$, in the k -th exploitation phase, deriving asymptotic behavior of the demand curve reduces to deriving the asymptotic behavior of our estimator on $\hat{F}_k(\cdot)$. This is due to the statistical rate of $\hat{F}_k(\cdot)$ dominates that of $\hat{\boldsymbol{\theta}}_k$. According to asymptotic behavior of the kernel regression [Fan and Gijbels, 1996, Carroll et al., 1997, Fan et al., 1998], we have the following pointwise confidence interval for $\hat{F}$:

$$\sqrt{|I_k| h_k} (\hat{F}_k(u) - F(u) - h_k^m \kappa_m B(u)) \rightarrow N\left(0, \int K^2(x) dx \sigma^2(u) / f(u)\right),$$

where $f(\cdot)$ is the density of $p_t - \mathbf{x}_t^\top \boldsymbol{\theta}_0$ with $p_t \sim \text{Unif}(0, B)$ and we recall that $|I_k|$ is the length of our k -th exploration phase. In addition, $\kappa_m = \int K(x) x^m dx$, $B(u) = F^{(m)}(u) f(u) / m! + F^{(m-1)}(u) f^{(1)}(u) / (m-1)! + \dots + F^{(1)}(u) f^{(m-1)}(u) / (m-1)!$, and $\sigma^2(u) = \text{Var}(y_t | p_t - \mathbf{x}_t^\top \boldsymbol{\theta}_0 = u)$. Thus, for any given $p, \mathbf{x}$, and an $\hat{\boldsymbol{\theta}}_k$, we are able to derive the pointwise asymptotic behavior of our demand curve as follows:

$$\begin{aligned} & \sqrt{|I_k| h_k} (p \hat{F}_k(p - \mathbf{x}^\top \hat{\boldsymbol{\theta}}_k) - p F(p - \mathbf{x}^\top \boldsymbol{\theta}_0) - p h_k^m \kappa_m B(p - \mathbf{x}^\top \boldsymbol{\theta}_0)) \\ & \rightarrow N\left(0, p^2 \int K^2(s) ds \sigma^2(p - \mathbf{x}^\top \boldsymbol{\theta}_0) / f(p - \mathbf{x}^\top \boldsymbol{\theta}_0)\right). \end{aligned}$$

The data-driven confidence interval for our demand curve given in (4) can be established via bootstrap and the undersmoothing technique (to remove the bias), see e.g. Hall [1992], Horowitz [2001] for more details. Similarly, uniform statistical inference results can also be established

by using similar non-parametric tools, see e.g. Eubank and Speckman [1993], Neumann and Polzehl [1998], Hall and Horowitz [2013] for more details. We will leave the detailed proof for future work.

4. In some situations, it might be difficult for retailers to adopt a uniform pricing strategy even during a short period of time. An alternative strategy might be the following: As in Algorithm 1, we divide the time horizon into episodes according to the doubling strategy. However, now we no longer divide an episode into explore-then-exploitation phases. Instead, at the beginning of each episode $k > 1$, we leverage all the data $\{p_t, \mathbf{x}_t, y_t\}$ collected from the previous episode to estimate θ_0 and F . Then, we compute $\hat{g}_k$ from the estimates $\hat{F}_k$ and $\hat{F}_k^{(1)}$, and perform exploitation directly throughout this episode. This procedure can help us to get rid of uniform exploration in practice. We leave the theoretical guarantees for this refined algorithm as our future work.

B.1 Extension: High-dimensional Feature-based Dynamic Pricing

Algorithm 1 can be naturally extended to the high-dimensional setting, where $\theta_0 \in \mathbb{R}^d$, d can be large compared to T , while $\|\theta_0\|_0 \leq s$ for a relatively small sparsity s . This happens in applications when a large amount of covariate information is available, and the actual market value only depends on some essential factors. One way of extension is the following: at each episode, we can replace estimation of $\hat{\theta}_k$ in (7) with the two steps below.

Step 1. Let

$$\tilde{\theta}_k = \underset{\theta}{\operatorname{argmin}} L_k(\theta) + \lambda p(\theta), \quad (1)$$

where

$$L_k(\theta) := \frac{1}{|I_k|} \sum_{t \in I_k} (By_t - \theta^\top \tilde{\mathbf{x}}_t)^2, \quad p(\theta) = \sum_{j=1}^p p(|\theta^{(j)}|)$$

for some penalty function $p(\cdot)$. As in Zhao and Yu [2006], Fan and Li [2001], Zhang [2010], by choosing different $p(\cdot)$ such as in the ℓ_1 , SCAD or MCP penalty, under suitable conditions such as irrepresentable condition, variable selection consistency is achieved with high probability.

Step 2. Let $\hat{S}_k = \operatorname{supp}(\tilde{\theta}_k)$, we then refit the least squares (7) on $\hat{S}_k$:

$$\hat{\theta}_k = \underset{\operatorname{supp}(\theta) \subseteq \hat{S}_k}{\operatorname{argmin}} L_k(\theta). \quad (2)$$

Then the conclusions of Lemma 2 hold with high probability.

After Step 2, we continue the remaining steps of Algorithm 1 in the episode. In fact, if we can learn the support of θ_0 , we essentially translate the problem into a low-dimensional one, and we can prove that Algorithm 1 achieves a regret upper bound of $\tilde{\mathcal{O}}((Ts)^{\frac{4m+1}{2m-1}})$ if $F \in \mathcal{C}^{(m)}$ (or $\mathcal{O}((Ts)^{1/2})$ if F is super smooth).

C Proof under the time-independent feature setting

C.1 Proof of Lemma 1

First, recall that $R_{\mathcal{X}} := \sup_{\mathbf{x} \in \mathcal{X}} \|\mathbf{x}\|_2$, we deduce that $\mathbf{x}_t$ is also subgaussian with norm upper bounded by $\psi_x = R_{\mathcal{X}}$. This fact is useful in later proofs as well. Now according to (7), for the k -th episode, our loss function $L_k(\boldsymbol{\theta})$ is defined as

$$L_k(\boldsymbol{\theta}) = \frac{1}{|I_k|} \sum_{t \in I_k} (By_t - \boldsymbol{\theta}^\top \tilde{\mathbf{x}}_t)^2. \quad (3)$$

For notational convenience, denote $n = |I_k|$. Then the gradient and Hessian of $L_k(\boldsymbol{\theta})$ is given by

$$\nabla_{\boldsymbol{\theta}} L_k(\boldsymbol{\theta}) = \frac{1}{n} \sum_{t \in I_k} 2(\boldsymbol{\theta}^\top \tilde{\mathbf{x}}_t - By_t) \tilde{\mathbf{x}}_t, \quad (4)$$

$$\nabla_{\boldsymbol{\theta}}^2 L_k(\boldsymbol{\theta}) = \frac{1}{n} \sum_{t \in I_k} 2\tilde{\mathbf{x}}_t \tilde{\mathbf{x}}_t^\top. \quad (5)$$

Let $\hat{\boldsymbol{\theta}}_k$ be the global minimizer of $L_k(\boldsymbol{\theta})$. We do a Taylor expansion of $L_k(\hat{\boldsymbol{\theta}}_k)$ at $\boldsymbol{\theta}_0$:

$$L_k(\hat{\boldsymbol{\theta}}_k) - L_k(\boldsymbol{\theta}_0) = \langle \nabla L_k(\boldsymbol{\theta}_0), \hat{\boldsymbol{\theta}}_k - \boldsymbol{\theta}_0 \rangle + \frac{1}{2} \langle \hat{\boldsymbol{\theta}}_k - \boldsymbol{\theta}_0, \nabla_{\boldsymbol{\theta}}^2 L_k(\tilde{\boldsymbol{\theta}})(\hat{\boldsymbol{\theta}}_k - \boldsymbol{\theta}_0) \rangle. \quad (6)$$

Here $\tilde{\boldsymbol{\theta}}$ is a point lying between $\hat{\boldsymbol{\theta}}_k$ and $\boldsymbol{\theta}_0$. As $\hat{\boldsymbol{\theta}}_k$ is the global minimizer of loss (3), we have

$$\langle \nabla L_k(\boldsymbol{\theta}_0), \hat{\boldsymbol{\theta}}_k - \boldsymbol{\theta}_0 \rangle + \frac{1}{2} \langle \hat{\boldsymbol{\theta}}_k - \boldsymbol{\theta}_0, \nabla_{\boldsymbol{\theta}}^2 L_k(\tilde{\boldsymbol{\theta}})(\hat{\boldsymbol{\theta}}_k - \boldsymbol{\theta}_0) \rangle \leq 0$$

which implies

$$\langle \hat{\boldsymbol{\theta}}_k - \boldsymbol{\theta}_0, \frac{1}{n} \sum_{t \in I_k} \tilde{\mathbf{x}}_t \tilde{\mathbf{x}}_t^\top (\hat{\boldsymbol{\theta}}_k - \boldsymbol{\theta}_0) \rangle \leq \langle \nabla L_k(\boldsymbol{\theta}_0), \boldsymbol{\theta}_0 - \hat{\boldsymbol{\theta}}_k \rangle \leq \sqrt{d} \|\nabla L_k(\boldsymbol{\theta}_0)\|_\infty \cdot \|\boldsymbol{\theta}_0 - \hat{\boldsymbol{\theta}}_k\|_2. \quad (7)$$

In order to achieve ℓ_2 -convergence rate of $\hat{\boldsymbol{\theta}}_k$, we separate our following analysis into two steps.

Step I: In this step, we lower bound the minimum eigenvalue of

$$\boldsymbol{\Sigma}_k := \frac{1}{n} \sum_{t \in I_k} \tilde{\mathbf{x}}_t \tilde{\mathbf{x}}_t^\top. \quad (8)$$

using concentration inequalities.

Since $\boldsymbol{\Sigma}_k$ is an average of n i.i.d. random matrices with mean $\boldsymbol{\Sigma} = \mathbb{E}[\tilde{\mathbf{x}}_t \tilde{\mathbf{x}}_t^\top]$ and that $\{\tilde{\mathbf{x}}_t\}$ are sub-Gaussian random vectors, according to Remark 5.40 in Vershynin [2012], there exist c_1 and $C > c_{\min}$ such that with probability at least $1 - 2e^{-c_1 t^2}$,

$$\|\boldsymbol{\Sigma}_k - \boldsymbol{\Sigma}\| \leq \max\{\delta, \delta^2\}, \quad \text{where } \delta := C \sqrt{\frac{d+1}{n}} + \frac{t}{\sqrt{n}}. \quad (9)$$

Here c_1, C are both constants that are only related to sub-Gaussian norm of $\tilde{\mathbf{x}}_t$. Now we plug in $t = c_{\min}\sqrt{n}/4$ and $c_0 = 16C^2/c_{\min}^2$, then as long as $n \geq c_0(d+1)$, with probability at least $1 - 2e^{-c_1 c_{\min}^2 n/16}$,

$$(c_{\min}/2) \cdot \mathbb{I} \leq \Sigma_k. \quad (10)$$

Step II: In this step, we provide an upper bound of $\|\nabla_{\theta} L_k(\theta_0)\|_{\infty}$.

First, we prove $\mathbb{E}[\nabla_{\theta} L_k(\theta_0)] = 0$. By definition we have

$$\nabla_{\theta} L_k(\theta_0) = \frac{1}{n} \sum_{t \in I_k} 2(\theta_0^{\top} \tilde{\mathbf{x}}_t - By_t) \tilde{\mathbf{x}}_t$$

We take the conditional expectation of $\nabla_{\theta} L_k(\theta_0)$ and obtain

$$\mathbb{E}[\nabla_{\theta} L_k(\theta_0) \mid \tilde{\mathbf{x}}_t] = \frac{1}{n} \sum_{t \in I_k} 2\mathbb{E}[(\theta_0^{\top} \tilde{\mathbf{x}}_t - By_t) \mid \tilde{\mathbf{x}}_t] \tilde{\mathbf{x}}_t.$$

By our definition on y_t ,

$$\begin{aligned} \mathbb{E}[\theta_0^{\top} \tilde{\mathbf{x}}_t - By_t \mid \tilde{\mathbf{x}}_t] &= \theta_0^{\top} \tilde{\mathbf{x}}_t - \mathbb{E}[B\mathbb{I}_{\{p_t \leq v_t\}} \mid \tilde{\mathbf{x}}_t] \\ &= \theta_0^{\top} \tilde{\mathbf{x}}_t - \mathbb{E}[\mathbb{E}[B\mathbb{I}_{\{p_t \leq v_t\}} \mid v_t] \mid \tilde{\mathbf{x}}_t] \\ &= \theta_0^{\top} \tilde{\mathbf{x}}_t - B \cdot \mathbb{E}[v_t/B \mid \tilde{\mathbf{x}}_t] = 0, \end{aligned}$$

where the third equality follows from $p_t \sim \text{Uniform}(0, B)$. After finally taking expectation with respect to $\tilde{\mathbf{x}}_t$ we deduce that $\mathbb{E}[\nabla_{\theta} L_k(\theta_0)] = 0$.

Next, we get an upper bound of $\|\nabla_{\theta} L_k(\theta_0)\|_{\infty}$. By (4), we have every entry of $\nabla_{\theta} L_k(\theta_0)$ is mean zero. In addition, according to our Assumption 2, we have $\mathbf{x}_t$ are i.i.d. sub-Gaussian random vectors with sub-Gaussian norm ψ_x . Thus, we have $\max_{i \in [d]} \|\mathbf{x}_{t,i}\|_{\psi_2} \leq \psi_x$. On the other hand, $\tilde{\mathbf{x}}_t^{\top} \theta_0 - By_t$ is bounded by the constant $R_{\mathcal{X}} R_{\Theta} + B$. Therefore,

$$\mathbb{P}(|2(\theta_0^{\top} \tilde{\mathbf{x}}_t - By_t) \tilde{\mathbf{x}}_{t,i}| \geq u) \leq \mathbb{P}(2(R_{\mathcal{X}} R_{\Theta} + B) |\tilde{\mathbf{x}}_{t,i}| \geq u) \leq 2 \exp\left(\frac{-u^2}{8\psi_x^2 (R_{\mathcal{X}} R_{\Theta} + B)^2}\right)$$

for $i \in [2 : (d+1)]$, which implies that $2(\theta_0^{\top} \tilde{\mathbf{x}}_t - By_t) \tilde{\mathbf{x}}_{t,i}$, $i \in [2 : (d+1)]$ are sub-Gaussian random variables with variance proxy $2\psi_x(R_{\mathcal{X}} R_{\Theta} + B)$. Moreover, We can also obtain $\|2(\theta_0^{\top} \tilde{\mathbf{x}}_t - By_t) \tilde{\mathbf{x}}_{t,1}\|_{\psi_2} \leq 2(R_{\mathcal{X}} R_{\Theta} + B)$ by Hoeffding's inequality.

We now take the union bound of all entries of $\nabla_{\theta} L_k(\theta_0)$:

$$\mathbb{P}(\|\nabla_{\theta} L_k(\theta_0)\|_{\infty} \geq t) \leq 2(d+1) \exp\left(\frac{-t^2}{8 \max\{\psi_x^2, 1\} (R_{\mathcal{X}} R_{\Theta} + B)^2}\right) \quad (11)$$

$$= 2 \exp\left(\frac{-nt^2}{8 \max\{\psi_x^2, 1\} (R_{\mathcal{X}} R_{\Theta} + B)^2} + \log(d+1)\right). \quad (12)$$

As we assume $n \geq d + 1$, by taking $t = 4 \max\{\psi_x, 1\}(R_{\mathcal{X}}R_{\Theta} + B)\sqrt{\log n/n}$ in (12), then with probability $1 - 2/n$, we have

$$\|\nabla_{\theta} L_k(\theta_0)\|_{\infty} \leq 4 \max\{\psi_x, 1\}(R_{\mathcal{X}}R_{\Theta} + B)\sqrt{\frac{\log n}{n}}. \quad (13)$$

Finally, combining (7), (10) and (13), we obtain that with probability at least $1 - 2e^{-c_1 c_{\min}^2 |I_k|/16} - 2/|I_k|$,

$$\|\hat{\theta}_k - \theta_0\|_2 \leq \frac{8 \max\{\psi_x, 1\}(R_{\mathcal{X}}R_{\Theta} + B)}{c_{\min}} \sqrt{\frac{(d+1) \log |I_k|}{|I_k|}}.$$

C.2 Proof of Lemma 2

For the following analysis, we fix any episode index k satisfying the conditions of Lemma 2. It's easy to verify that for any $k \geq (\log(\sqrt{T} - \log \ell_0))/\log 2$, $\Theta_k \subset \Theta_0$. Therefore, all the assumptions hold for $\theta \in \Theta_k$. Our goal is to prove (18) holds with high probability on the k -th episode.

Now we have the i.i.d. samples $\{w_t(\theta) := p_t - \tilde{\mathbf{x}}_t^{\top} \theta, y_t\}_{t \in I_k}$ from some distribution $P_{w(\theta), y}$. According to the previous notations, the marginal distribution $P_{w(\theta)}$ has density $f_{\theta}(u)$. Moreover, $r_{\theta}(u) := \mathbb{E}[y_t | w_t(\theta) = u]$. We're interested in bounding the quantity $\sup_{u \in I, \theta \in \Theta_k} |\hat{r}_k(u, \theta) - r_{\theta}(u)|$, which leads to the desired conclusion of the lemma.

For notational simplicity, let $n = |I_k|$ be the length of the exploration phase. Recall that $\hat{r}_k(u, \theta) = h_k(u, \theta)/f_k(u, \theta)$, where

$$h_k(u, \theta) = \frac{1}{nb_k} \sum_{t \in I_k} K\left(\frac{w_t(\theta) - u}{b_k}\right) Y_t, \quad f_k(u, \theta) = \frac{1}{nb_k} \sum_{t \in I_k} K\left(\frac{w_t(\theta) - u}{b_k}\right).$$

Here, $b_k > 0$ is the bandwidth (to be chosen), and $K(\cdot)$ is some kernel function.

Note that $r_{\theta}(u) = \frac{h_{\theta}(u)}{f_{\theta}(u)}$, we can write the difference between $\hat{r}_k$ and r as

$$\hat{r}_k(u, \theta) - r_{\theta}(u) = \frac{h_k(u, \theta)}{f_k(u, \theta)} - \frac{h_{\theta}(u)}{f_{\theta}(u)} = \frac{h_k(u, \theta) - h_{\theta}(u)}{f_k(u, \theta)} + h_{\theta}(u) \cdot \left[\frac{1}{f_k(u, \theta)} - \frac{1}{f_{\theta}(u)} \right]. \quad (14)$$

The following lemmas are used as tools to control the right hand side of the above equation. The proof of the lemmas can be found in §F.1 and F.2.

Lemma 1. Under Assumptions 3–5, for any $b_k \leq 1$,

$$\sup_{u \in I, \theta \in \Theta_k} |\mathbb{E} h_k(u, \theta) - h_{\theta}(u)| \leq C_{x,K}^{(1)} b_k^m, \quad (15)$$

$$\sup_{u \in I, \theta \in \Theta_k} |\mathbb{E} f_k(u, \theta) - f_{\theta}(u)| \leq C_{x,K}^{(1)} b_k^m. \quad (16)$$

Here, $C_{x,K}^{(1)} = l_f \int \frac{|s^m K(s)| ds}{(m-1)!}$.

Lemma 2. Under Assumptions 3–5, $\forall b_k \leq 1$, $\delta \in [4e^{-nb_k/3}, \frac{1}{2})$, as long as $nb_k \geq \max\{132d(\log \frac{1}{b_k} + 1), 3 \log n\}$, either of the following inequalities holds with probability at least $1 - \delta$:

$$\sup_{u \in I, \theta \in \Theta_k} |h_k(u, \theta) - \mathbb{E}h_k(u, \theta)| \leq C_{x,K}^{(2)} \sqrt{\frac{\log n}{nb_k}} \left(\sqrt{d} + \sqrt{\log 1/\delta} \right), \quad (17)$$

$$\sup_{u \in I, \theta \in \Theta_k} |f_k(u, \theta) - \mathbb{E}f_k(u, \theta)| \leq C_{x,K}^{(2)} \sqrt{\frac{\log n}{nb_k}} \left(\sqrt{d} + \sqrt{\log 1/\delta} \right). \quad (18)$$

Here $C_{x,K}^{(2)} = l_K \left(8\sqrt{22} \max\{2\bar{f} \int K^2 \mathbf{d}s, 2\bar{f} \int K'^2 \mathbf{d}s, \frac{2}{3}\bar{K}, 1\} + \frac{60(6\sqrt{\log 2} + \sqrt{c_0})}{c_0} \sqrt{1 + R_{\mathcal{X}}^2} \max\{\delta_z, \frac{\max\{1, \psi_x\}(B + R_{\mathcal{X}}R_{\Theta})}{c_{\min}}\} \right)$ (Numerical constants are not optimized).

Now according to (14), we have

$$\begin{aligned} \sup_{u \in I, \theta \in \Theta_k} |\hat{r}_k(u, \theta) - r(u)| &\leq \sup_{u \in I, \theta \in \Theta_k} \frac{|h_k(u, \theta) - h_{\theta}(u)|}{|f_{\theta}(u) - |f_k(u, \theta) - f_{\theta}(u)||} \\ &\quad + \sup_{u \in I, \theta \in \Theta_k} \frac{h_{\theta}(u)}{f_{\theta}(u)} \cdot \frac{|f_k(u, \theta) - f_{\theta}(u)|}{|f_{\theta}(u) - |f_k(u, \theta) - f_{\theta}(u)||} \\ &\leq \frac{\sup_{u \in I, \theta \in \Theta_k} |h_k(u, \theta) - h_{\theta}(u)|}{c - \sup_{u \in I, \theta \in \Theta_k} |f_k(u, \theta) - f_{\theta}(u)|} \\ &\quad + \sup_{u \in I, \theta \in \Theta_k} r_{\theta}(u) \cdot \frac{\sup_{u \in I, \theta \in \Theta_k} |f_k(u, \theta) - f_{\theta}(u)|}{c - \sup_{u \in I, \theta \in \Theta_k} |f_k(u, \theta) - f_{\theta}(u)|} \\ &\leq \frac{\sup_{u \in I, \theta \in \Theta_k} |h_k(u, \theta) - h_{\theta}(u)|}{c - \sup_{u \in I, \theta \in \Theta_k} |f_k(u, \theta) - f_{\theta}(u)|} + \frac{\sup_{u \in I, \theta \in \Theta_k} |f_k(u, \theta) - f_{\theta}(u)|}{c - \sup_{u \in I, \theta \in \Theta_k} |f_k(u, \theta) - f_{\theta}(u)|} \end{aligned} \quad (19)$$

(20)

as long as we ensure that $\sup_{u \in I, \theta \in \Theta_k} |f_k(u, \theta) - f_{\theta}(u)| \leq \frac{c}{2}$.

Let $b_k = n^{-\frac{1}{2m+1}}$. By letting $B_{x,K} = \max\{4C_{x,K}^{(3)8}/c^8, (2c_0)^4, (2C_b)^4\}$, we can verify that for any qualifying episode k , $nb_k \geq \max\{C_b d(\log \frac{1}{b_k} + 1), 3 \log n\}$. Combining (15) and (17), we have that $\forall \delta \in [4 \exp(-n^{\frac{2m}{2m+1}}/3), \frac{1}{2})$, with probability at least $1 - \delta$,

$$\begin{aligned} \sup_{u \in I, \theta \in \Theta_k} |h_k(u, \theta) - h_{\theta}(u)| &\leq \sup_{u \in I, \theta \in \Theta_k} |h_k(u, \theta) - \mathbb{E}h_k(u, \theta)| + \sup_{u \in I, \theta \in \Theta_k} |\mathbb{E}h_k(u, \theta) - h_{\theta}(u)| \\ &\leq C_{x,K}^{(1)} n^{-\frac{m}{2m+1}} + C_{x,K}^{(2)} \sqrt{\frac{\log n}{nb_k}} \left(\sqrt{d} + \sqrt{\log 1/\delta} \right) \\ &\leq C_{x,K}^{(3)} n^{-\frac{m}{2m+1}} \sqrt{\log n} \left(\sqrt{d} + \sqrt{\log 1/\delta} \right). \end{aligned}$$

Here, $C_{x,K}^{(3)} = C_{x,K}^{(1)} + C_{x,K}^{(2)}$. Similarly, with probability at least $1 - \delta$,

$$\sup_{u \in I, \theta \in \Theta_k} |f_k(u, \theta) - f_\theta(u)| \leq C_{x,K}^{(3)} n^{-\frac{m}{2m+1}} \sqrt{\log n} \left(\sqrt{d} + \sqrt{\log 1/\delta} \right).$$

It's easily seen that as long as $n^{\frac{m}{2m+1}} / \sqrt{\log n} \geq \frac{2C_{x,K}^{(3)}}{c} (\sqrt{d} + \sqrt{\log 1/\delta})$, The right hand side of the above inequality is upper bounded by $c/2$, which guarantees that

$$\sup_{u \in I, \theta \in \Theta_k} |f_k(u, \theta) - f_\theta(u)| \leq \frac{c}{2}.$$

(Remark: From the conditions in the lemma, by letting $B_{x,K} = \max\{4C_{x,K}^{(3)8}/c^8, (2c_0)^4, (2C_b)^4\}$ and $B'_{x,K} = \min\{(\frac{c}{4C_{x,K}^{(3)}})^2, 1/3\}$, we have

$$n^{\frac{m}{2m+1}} / \sqrt{\log n} \geq \frac{4C_{x,K}^{(3)}}{c} \sqrt{d}, \quad n^{\frac{m}{2m+1}} / \sqrt{\log n} \geq \frac{4C_{x,K}^{(3)}}{c} \sqrt{\log \frac{1}{\delta}},$$

which lead to $n^{\frac{m}{2m+1}} / \sqrt{\log n} \geq \frac{2C_{x,K}^{(3)}}{c} (\sqrt{d} + \sqrt{\log 1/\delta})$.

Plugging the above results into inequality (20) gives

$$\sup_{u \in I, \theta \in \Theta_k} |\hat{r}_k(u, \theta) - r_\theta(u)| \leq \frac{4C_{x,K}^{(3)}}{c} n^{-\frac{m}{2m+1}} \sqrt{\log n} \left(\sqrt{d} + \sqrt{\log 1/\delta} \right). \quad (21)$$

Next, we proceed to upper bound the quantity $\sup_{t \in I, \theta \in \Theta_k} |r_\theta(u) - r_{\theta_0}(u)|$. We know that for any $\theta \in \Theta_k$,

$$r_\theta(u) = \mathbb{E}[Y_t | p_t - \tilde{\mathbf{x}}_t^\top \theta = u] = \mathbb{E}[\mathbb{E}[Y_t | \tilde{\mathbf{x}}_t, p_t] | p_t - \tilde{\mathbf{x}}_t^\top \theta = u] = \mathbb{E}[r_{\theta_0}(p_t - \tilde{\mathbf{x}}_t^\top \theta_0) | p_t - \tilde{\mathbf{x}}_t^\top \theta = u].$$

Moreover from the Lipchitz property of r_{θ_0} ,

$$\sup_{\mathbf{x} \in \mathcal{X}, \theta \in \Theta_k} |r_{\theta_0}(p_t - \tilde{\mathbf{x}}^\top \theta_0) - r_{\theta_0}(p_t - \tilde{\mathbf{x}}^\top \theta)| \leq l_r R_{\mathcal{X}} R_k = l_r R_{\mathcal{X}} \cdot \frac{10 \max\{\psi_x, 1\} (B + R_{\mathcal{X}} R_{\Theta})}{c_{\min}} \sqrt{\frac{(d+1) \log n}{n}}.$$

Therefore,

$$\sup_{u \in I, \theta \in \Theta_k} |r_\theta(u) - r_{\theta_0}(u)| \leq \mathbb{E} \left[\sup_{\mathbf{x} \in \mathcal{X}, \theta \in \Theta_k} |r_{\theta_0}(p_t - \tilde{\mathbf{x}}^\top \theta_0) - r_{\theta_0}(p_t - \tilde{\mathbf{x}}^\top \theta)| | p_t - \tilde{\mathbf{x}}_t^\top \theta = u \right] \leq C_{x,K}^{(4)} \sqrt{\frac{d \log n}{n}}, \quad (22)$$

where $C_{x,K}^{(4)} = l_r R_{\mathcal{X}} \cdot \frac{10 \max\{\psi_x, 1\} (B + R_{\mathcal{X}} R_{\Theta})}{c_{\min}}$.

Finally, after combing our results in (21)-(22), we claim our conclusion for Lemma 2.

C.3 Proof of Lemma 3

Following the same settings as in the proof of Lemma 2, we now aim at bounding the quantity $\sup_{u \in I, \theta \in \Theta_k} |\hat{r}_k^{(1)}(u, \theta) - r'_{\theta_0}(u)|$, where

$$\begin{aligned} r_k^{(1)}(u, \theta) &= \frac{h_k^{(1)}(u, \theta) f_k(u, \theta) - h_k(u, \theta) f_k^{(1)}(u, \theta)}{f_k^2(u, \theta)}, \\ h_k(u, \theta) &= \frac{1}{nb_k} \sum_{u \in I_k} K\left(\frac{w_t(\theta) - u}{b_k}\right) Y_t, \quad f_k(u, \theta) = \frac{1}{nb_k} \sum_{t \in I_k} K\left(\frac{w_t(\theta) - u}{b_k}\right), \\ h_k^{(1)}(u, \theta) &= \frac{-1}{nb_k^2} \sum_{t \in I_k} K'\left(\frac{w_t(\theta) - u}{b_k}\right) Y_t, \quad f_k^{(1)}(u, \theta) = \frac{-1}{nb_k^2} \sum_{t \in I_k} K'\left(\frac{w_t(\theta) - u}{b_k}\right). \end{aligned}$$

Similar to the proof of Lemma 2, we will bound $\sup_{u \in I, \theta \in \Theta_k} |\hat{r}_k^{(1)}(u, \theta) - r'_\theta(u)|$ and $\sup_{u \in I, \theta \in \Theta_k} |r'_\theta(u) - r'_{\theta_0}(u)|$ separately. First, notice that

$$r'_\theta(u) = \frac{h'_\theta(u) f_\theta(u) - f'_\theta(u) h_\theta(u)}{f_\theta^2(u)},$$

we can bound $\sup_{u \in I, \theta \in \Theta_k} |\hat{r}_k^{(1)}(u, \theta) - r'_\theta(u)|$ from the following four terms: $\sup_{u \in I, \theta \in \Theta_k} |f_k(u, \theta) - f_\theta(u)|$, $\sup_{u \in I, \theta \in \Theta_k} |h_k(u, \theta) - h_\theta(u)|$, $\sup_{u \in I, \theta \in \Theta_k} |f_k^{(1)}(u, \theta) - f'_\theta(u)|$ and $\sup_{u \in I, \theta \in \Theta_k} |h_k^{(1)}(u, \theta) - h'_\theta(u)|$. In fact, we can upper bound the first two terms from Lemma 1 and 2. The lemmas below help us bound the last two terms. The proof can be found in §F.3 and F.4.

Lemma 3. *Given Assumptions 3-5, for any $b_k \leq 1$,*

$$\sup_{u \in I, \theta \in \Theta_k} |\mathbb{E} h_k^{(1)}(u, \theta) - h'_\theta(u)| \leq C_{x,K}^{(5)} b_k^{m-1}, \quad (23)$$

$$\sup_{u \in I, \theta \in \Theta_k} |\mathbb{E} f_k^{(1)}(u, \theta) - f'_\theta(u)| \leq C_{x,K}^{(5)} b_k^{m-1}. \quad (24)$$

Here, $C_{x,K}^{(5)} = \frac{l_f}{(m-2)!} \int |K(s) s^{m-1}| ds$.

Lemma 4. *Given assumptions 3, 4 and 5, $\forall b_k \in [\frac{1}{n}, 1]$, $\delta \in [4e^{-nb_k/3}, \frac{1}{2}]$, either of the following inequalities holds with probability at least $1 - \delta$:*

$$\sup_{u \in I, \theta \in \Theta_k} |h_k^{(1)}(u, \theta) - \mathbb{E} h_k^{(1)}(u, \theta)| \leq C_{x,K}^{(2)} \sqrt{\frac{\log n}{nb_k^3}} \left(\sqrt{d} + \sqrt{\log 1/\delta} \right), \quad (25)$$

$$\sup_{u \in I, \theta \in \Theta_k} |f_k^{(1)}(u, \theta) - \mathbb{E} f_k^{(1)}(u, \theta)| \leq C_{x,K}^{(2)} \sqrt{\frac{\log n}{nb_k^3}} \left(\sqrt{d} + \sqrt{\log 1/\delta} \right). \quad (26)$$

Here $C_{x,K}^{(2)} = l_K \left(8\sqrt{22} \max\{2\bar{f} \int K^2 \mathbf{d}s, 2\bar{f} \int K'^2 \mathbf{d}s, \frac{2}{3}\bar{K}, 1\} + \frac{60(6\sqrt{\log 2} + \sqrt{c_0})}{c_0} \sqrt{1 + R_{\mathcal{X}}^2} \max\{\delta_z, \frac{\max\{1, \psi_x\}(B + R_{\mathcal{X}} R_{\Theta})}{c_{\min}}\} \right)$ (Numerical constants are not optimized).

Now let $b_k = n^{-\frac{1}{2m+1}}$. Combining (23) and (25), we obtain that $\forall \delta \in [4 \exp(-n^{\frac{2m}{2m+1}}/3), \frac{1}{2})$, with probability at least $1 - \delta$,

$$\begin{aligned} \sup_{u \in I, \theta \in \Theta_k} |h_k^{(1)}(u, \theta) - h'_\theta(u)| &\leq \sup_{u \in I, \theta \in \Theta_k} |h_k^{(1)}(u, \theta) - \mathbb{E}h_k^{(1)}(u, \theta)| + \sup_{u \in I, \theta \in \Theta_k} |\mathbb{E}h_k^{(1)}(u, \theta) - h'_\theta(u)| \\ &\leq C_{x,K}^{(5)} n^{-\frac{m-1}{2m+1}} + C_{x,K}^{(2)} \sqrt{\frac{\log n}{nb_k^3}} \left(\sqrt{d} + \sqrt{\log 1/\delta} \right) \\ &\leq C_{x,K}^{(6)} n^{-\frac{m-1}{2m+1}} \sqrt{\log n} \left(\sqrt{d} + \sqrt{\log 1/\delta} \right) \end{aligned}$$

Here, $C_{x,K}^{(6)} = C_{x,K}^{(5)} + C_{x,K}^{(2)}$. Similarly, with probability at least $1 - \delta$,

$$\sup_{u \in I, \theta \in \Theta_k} |f_k^{(1)}(u, \theta) - f'_\theta(u)| \leq C_{x,K}^{(6)} n^{-\frac{m-1}{2m+1}} \sqrt{\log n} \left(\sqrt{d} + \sqrt{\log 1/\delta} \right).$$

Recall that when $n^{\frac{m}{2m+1}}/\sqrt{\log n} \geq \frac{2C_{x,K}^{(3)}}{c}(\sqrt{d} + \sqrt{\log 1/\delta})$, we have

$$\sup_{u \in I, \theta \in \Theta_k} |f_k(u, \theta) - f_\theta(u)| \leq \frac{c}{2}, \quad \sup_{u \in I, \theta \in \Theta_k} |h_k(u, \theta) - h_\theta(u)| \leq \frac{c}{2}.$$

Moreover, we have

$$\sup_{u \in I, \theta \in \Theta_k} \max\{|h_\theta(u)|, |f_\theta(u)|, |f'_\theta(u)|\} \leq \bar{f}, \quad \sup_{u \in I, \theta \in \Theta_k} |h'_\theta(u)| = \sup_{u \in I, \theta \in \Theta_k} |f'_\theta(u)r_\theta(u) + f_\theta(u)r'_\theta(u)| \leq l_f + l_r \bar{f}.$$

Therefore, from the definition of $r_k^{(1)}(u, \theta)$ and $r'_\theta(u)$, we have

$$\begin{aligned} &\sup_{u \in I, \theta \in \Theta_k} |r_k^{(1)}(\theta, u) - r'_\theta(u)| \\ &\leq \sup_{u \in I, \theta \in \Theta_k} \left| [h'_\theta(u)f_\theta(u) - h_\theta(u)f'_\theta(u)] \left[\frac{1}{f_k(u, \theta)^2} - \frac{1}{f_\theta(u)^2} \right] \right| \\ &\quad + \sup_{u \in I, \theta \in \Theta_k} \left| \frac{1}{f_k(u, \theta)^2} \{ [h_k^{(1)}(u, \theta)f_k(u, \theta) - h_k(u, \theta)f_k^{(1)}(u, \theta)] - [h'_\theta(u)f_\theta(u) - h_\theta(u)f'_\theta(u)] \} \right| \\ &\leq [l_f \bar{f} + (l_r + 1)\bar{f}^2] \cdot \sup_{u \in I, \theta \in \Theta_k} \left| \frac{f_k(u, \theta)^2 - f_\theta(u)^2}{f_k(u, \theta)^2 f_\theta(u)^2} \right| \\ &\quad + \sup_{u \in I, \theta \in \Theta_k} \frac{1}{f_k(u, \theta)^2} |(h_k^{(1)}(u, \theta) - h'_\theta(u))f_k(u, \theta) + h'_\theta(u)(f_k(u, \theta) - f_\theta(u))| \end{aligned}$$

$$\begin{aligned}
& - (f_k^{(1)}(u, \boldsymbol{\theta}) - f'_{\boldsymbol{\theta}}(u))h_k(u, \boldsymbol{\theta}) - f'_{\boldsymbol{\theta}}(u)(h_k(u, \boldsymbol{\theta}) - h_{\boldsymbol{\theta}}(u))| \\
\leq & [l_f \bar{f} + (l_r + 1)\bar{f}^2] \cdot \sup_{u \in I, \boldsymbol{\theta} \in \Theta_k} \frac{\frac{5}{2}f_{\boldsymbol{\theta}}(u)|f_k(u, \boldsymbol{\theta}) - f_{\boldsymbol{\theta}}(u)|}{f_k(u, \boldsymbol{\theta})^2 f_{\boldsymbol{\theta}}(u)^2} \\
& + \sup_{u \in I, \boldsymbol{\theta} \in \Theta_k} \frac{1}{f_k(u, \boldsymbol{\theta})^2} \left[\sup_{u \in I, \boldsymbol{\theta} \in \Theta_k} |f_k(u, \boldsymbol{\theta})| \cdot |h_k^{(1)}(u, \boldsymbol{\theta}) - h'_{\boldsymbol{\theta}}(u)| + \sup_{u \in I, \boldsymbol{\theta} \in \Theta_k} |h'_{\boldsymbol{\theta}}(u)| \cdot |f_k(u, \boldsymbol{\theta}) - f_{\boldsymbol{\theta}}(u)| \right. \\
& \quad \left. + \sup_{u \in I, \boldsymbol{\theta} \in \Theta_k} |h_k(u, \boldsymbol{\theta})| \cdot |f_k^{(1)}(u, \boldsymbol{\theta}) - f'_{\boldsymbol{\theta}}(u)| + \sup_{u \in I, \boldsymbol{\theta} \in \Theta_k} |f'_{\boldsymbol{\theta}}(u)| \cdot |h_k(u, \boldsymbol{\theta}) - h_{\boldsymbol{\theta}}(u)| \right]. \\
\leq & C_{x,K}^{(7)} n^{-\frac{m-1}{2m+1}} \sqrt{\log n} \left(\sqrt{d} + \sqrt{\log \frac{1}{\delta}} \right). \tag{27}
\end{aligned}$$

when $n^{\frac{m}{2m+1}}/\sqrt{\log n} \geq \frac{2C_{x,K}^{(3)}}{c}(\sqrt{d} + \sqrt{\log 1/\delta})$. Here

$$C_{x,K}^{(7)} = \left(\frac{10}{c^3} + \frac{4}{c^2}\right)[l_f(\bar{f} + 1) + (l_r + 1)\bar{f}^2]C_{x,K}^{(3)} + \left(\frac{8\bar{f}}{c^2} + \frac{4}{c}\right)C_{x,K}^{(6)}.$$

Next, we bound the term $\sup_{u \in I, \boldsymbol{\theta} \in \Theta_k} |r'_{\boldsymbol{\theta}}(u) - r'_{\boldsymbol{\theta}_0}(u)|$. In fact, according to our assumptions,

$$\sup_{u \in I, \boldsymbol{\theta} \in \Theta_k} |r'_{\boldsymbol{\theta}}(u) - r'_{\boldsymbol{\theta}_0}(u)| \leq C_{x,K}^{(4)} \sqrt{\frac{d \log n}{n}}, \tag{28}$$

where $C_{x,K}^{(4)} = l_r R_{\mathcal{X}} \cdot \frac{8 \max\{\psi_x, 1\}(B + R_{\mathcal{X}} R_{\Theta})}{c_{\min}}$. Finally, after combining our results in (27)-(28), we claim our conclusion for Lemma 3.

C.4 Proof of Lemma 4

We'll need the following auxiliary result in order to prove the lemma. The proof of Lemma 5 can be found in section F.5.

Lemma 5. *Given conditions of Lemma 4, for any $\tilde{\mathbf{x}}_t \in \mathcal{X}$ and $\boldsymbol{\theta} \in \Theta_0$, $\boldsymbol{\theta}^\top \tilde{\mathbf{x}}_t \in [\delta_z, B - \delta_z]$.*

Now we proceed to the proof. First, we seek an uniform upper bound for $|\hat{\phi}_k(u) - \phi(u)|$ from lemma 2 and 3. Recall that $\phi(u) = u - \frac{1-F(u)}{F'(u)}$ and $\hat{\phi}_k(u) = u - \frac{1-\hat{F}_k(u)}{\hat{F}_k^{(1)}(u)}$. It's easy to see that the desired uniform bound can be achieved on an interval where F' is bounded below from 0. For this reason, we choose some positive constant $c_{F'}$ and some interval $[l_{F'}, r_{F'}]$ (we'll specify how to choose them later) such that

$$\inf_{u \in [l_{F'}, r_{F'}]} F'(u) \geq c_{F'}. \tag{29}$$

From Lemma 3 we know that if in addition $|I_k|^{\frac{m-1}{2m+1}} \geq \frac{2\tilde{C}_{x,K}}{c_{F'}} \sqrt{\log |I_k|}(\sqrt{d} + \sqrt{\log 1/\delta})$, then $\sup_{u \in [l_{F'}, r_{F'}]} |\hat{F}_k^{(1)}(u) - F'(u)| \leq \frac{c_{F'}}{2}$ with probability at least $1 - 4\delta$. In fact, the above condition is

ensured by

$$T \geq \left(\frac{4\tilde{C}_{x,K}}{c_{F'}} \right)^8 (\log T + 2 \log d)^{\frac{4m-1}{m-1}} d^{\frac{2m+1}{m-1}}.$$

Combining (29), Lemma 2 and Lemma 3, we deduce that with probability at least $1 - 6\delta$,

$$\begin{aligned} \sup_{u \in [l_{F'}, r_{F'}]} |\hat{\phi}_k(u) - \phi(u)| &\leq \sup_{u \in [l_{F'}, r_{F'}]} \left| \frac{(1 - \hat{F}_k(u))(F'(u) - \hat{F}_k^{(1)}(u))}{\hat{F}_k^{(1)}(u)F'(u)} \right| \\ &\quad + \sup_{v \in [l_{F'}, r_{F'}]} \left| \frac{\hat{F}_k(u) - F(u)}{F'(u)} \right| \\ &\leq \frac{2\tilde{C}_{x,K} + C_{x,K}c_{F'}}{c_{F'}^2} |I_k|^{-\frac{m-1}{2m+1}} \sqrt{\log |I_k|} \left(\sqrt{d} + \sqrt{\log \frac{1}{\delta}} \right) \end{aligned} \quad (30)$$

Next, we proceed to bound $\sup_{u \in [\delta_z, B - \delta_z]} |\hat{g}_k(u) - g(u)|$ from $\sup_{u \in [\delta_z, B - \delta_z]} |\hat{\phi}_k^{-1}(-u) - \phi^{-1}(-u)|$ for some properly defined $\hat{\phi}_k^{-1}$. To be more specific, we will also let

$$[\delta_z - B, -\delta_z] \subseteq \phi([l_{F'}, r_{F'}]) \cap \hat{\phi}_k([l_{F'}, r_{F'}]). \quad (31)$$

The way we ensure the above is the following: First, according to the assumptions, we know $\phi'(u) \geq c_\phi > 0$, and that $\lim_{u \rightarrow \delta_z - 0} \phi(u) = \delta_z$, $\lim_{u \rightarrow l_F^{(1)} + 0} \phi(u) = -\infty$ with $l_F^{(1)} = \inf\{u : F'(u) > 0\} > -\delta_z$. We can deduce that

$$m_{F'} = \inf_{u \in [\phi^{-1}(\delta_z - B), \phi^{-1}(-\delta_z)]} F'(u) > 0.$$

Therefore, there exists some $\delta_{F'} > 0$ such that

$$\inf_{u \in [\phi^{-1}(\delta_z - B) - \delta_{F'}, \phi^{-1}(-\delta_z) + \delta_{F'}]} F'(u) > \frac{m_{F'}}{2}.$$

Now let $l_{F'} = \phi^{-1}(\delta_z - B) - \delta_{F'}$, $r_{F'} = \phi^{-1}(-\delta_z) + \delta_{F'}$, $c_{F'} = \frac{m_{F'}}{2}$. From the assumptions on ϕ , we have

$$\phi(l_{F'}) \leq \delta_z - B - c_\phi \delta_{F'}, \quad \phi(r_{F'}) \geq -\delta_z + c_\phi \delta_{F'}.$$

Combining (30), we obtain that as long as

$$\frac{2\tilde{C}_{x,K} + C_{x,K}c_{F'}}{c_{F'}^2} |I_k|^{-\frac{m-1}{2m+1}} \sqrt{\log |I_k|} \left(\sqrt{d} + \sqrt{\log \frac{1}{\delta}} \right) \leq c_\phi \delta_{F'},$$

we can ensure (31). The above condition can be obtained from the fact that

$$T \geq \left(\frac{4\tilde{C}_{x,K} + 2C_{x,K}c_{F'}}{c_{F'}^2 c_\phi \delta_{F'}} \right)^8 (\log T + 2 \log d)^{\frac{4m-1}{m-1}} d^{\frac{2m+1}{m-1}}.$$

Define

$$\hat{\phi}_k^{-1}(u) := \inf\{v \in [l_{F'}, r_{F'}] : \hat{\phi}_k(v) = u\}. \quad (32)$$

We proceed to upper bound $\sup_{u \in [\delta_z - B, -\delta_z]} |\hat{\phi}_k^{-1}(u) - \phi^{-1}(u)|$. In fact, for any u , let $v_1 = \phi^{-1}(u)$, $v_2 = \hat{\phi}_k^{-1}(u)$. Then

$$\begin{aligned} |v_1 - v_2| &\leq 1/c_\phi \cdot |\phi(v_1) - \phi(v_2)| = 1/c_\phi \cdot |\hat{\phi}_k(v_2) - \phi(v_2)| \\ &\leq 1/c_\phi \cdot \sup_{v \in [l_{F'}, r_{F'}]} |\hat{\phi}_k(v) - \phi(v)| \\ &\leq \frac{2\tilde{C}_{x,K} + C_{x,K}c_{F'}}{c_\phi c_{F'}^2} |I_k|^{-\frac{m-1}{2m+1}} \sqrt{\log |I_k|} \left(\sqrt{d} + \sqrt{\log \frac{1}{\delta}} \right) \end{aligned}$$

with probability at least $1 - 6\delta$.

Finally, since $g(u) = u + \phi^{-1}(-u)$ and $\hat{g}_k(u) = u + \hat{\phi}_k^{-1}(-u)$, we conclude Lemma 4 by choosing

$$\begin{aligned} \bar{B}_{x,K} &= \max \left\{ B_{x,K}, \left(\frac{4\tilde{C}_{x,K}}{c_{F'}} \right)^8, \left(\frac{4\tilde{C}_{x,K} + 2C_{x,K}c_{F'}}{c_{F'}^2 c_\phi \delta_{F'}} \right)^8, \left[\frac{C_\theta^2}{\delta_v^2} (1 + R_{\mathcal{X}}^2) \right]^{\frac{2(4m-1)}{2m+1}} \right\}, \\ \bar{B}'_{x,K} &= \min \left\{ B'_{x,K}, \left(\frac{c_{F'}}{4\tilde{C}_{x,K}} \right)^2, \left(\frac{c_{F'}^2 c_\phi \delta_{F'}}{4\tilde{C}_{x,K} + 2C_{x,K}c_{F'}} \right)^2 \right\}, \end{aligned}$$

and

$$\bar{C}_{x,K} = \frac{2\tilde{C}_{x,K} + C_{x,K}c_{F'}}{c_\phi c_{F'}^2}.$$

C.5 Proof of Theorem 1

In order to bound the total regret, we first try to bound the regret at each episode k . First, for all $k \leq \lfloor (\log(\sqrt{T} + \ell_0) - \log \ell_0) \log 2 \rfloor + 1$, we bound the total regret during episode k by $B\ell_k$. It can be easily verified that

$$\sum_{k \leq \lfloor (\log(\sqrt{T} + \ell_0) - \log \ell_0) \log 2 \rfloor + 1} \text{Regret}_k \leq 2B\sqrt{T}.$$

We now turn to the case where $k > \lfloor (\log(\sqrt{T} + \ell_0) - \log \ell_0) \log 2 \rfloor + 1$. Recall that the conditional expectation of regret at time t given previous information and $\tilde{\mathbf{x}}_t$ is

$$\mathbb{E}[R_t | \bar{\mathcal{H}}_{t-1}] = \mathbb{E}[p_t^* \mathbb{I}_{(v_t \geq p_t^*)} - p_t \mathbb{I}_{(v_t \geq p_t)} | \bar{\mathcal{H}}_t] = \rho_t(p_t^*) - \rho_t(p_t),$$

where $\bar{\mathcal{H}}_t = \sigma(\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{t+1}; z_1, \dots, z_t)$, and we denote $\rho_t(p) := p(1 - F(p - \boldsymbol{\theta}_0^\top \tilde{\mathbf{x}}_t))$. Using Taylor expansion and the first order condition induced by the optimality of p_t^* , we have

$$\rho_t(p_t) = \rho_t(p_t^*) + \frac{1}{2} \rho_t''(\xi_t) (p_t - p_t^*)^2,$$

where ξ_t is some value lying between p_t and p_t^* . Note that for any $p \in [0, B]$, $|\rho_t''(p)| = |2F'(p - \boldsymbol{\theta}_0^\top \tilde{\mathbf{x}}_t) - pF''(p - \boldsymbol{\theta}_0^\top \tilde{\mathbf{x}}_t)| \leq 2l_r + Bl'_r$. Thus we deduce that

$$\mathbb{E}[R_t | \bar{\mathcal{H}}_{t-1}] = \rho_t(p_t^*) - \rho_t(p_t) \leq (2l_r + Bl'_r)(p_t - p_t^*)^2,$$

which further implies that the expected regret at time t is bounded by

$$\mathbb{E}R_t \leq \frac{1}{2}(2l_r + Bl'_r)\mathbb{E}(p_t - p_t^*)^2 \quad (33)$$

On the other hand,

$$\begin{aligned} (p_t - p_t^*)^2 &\leq (\hat{g}_k(\tilde{\mathbf{x}}_t^\top \hat{\boldsymbol{\theta}}_k) - g(\tilde{\mathbf{x}}_t^\top \boldsymbol{\theta}_0))^2 \\ &\leq 2(\hat{g}_k(\tilde{\mathbf{x}}_t^\top \hat{\boldsymbol{\theta}}_k) - g(\tilde{\mathbf{x}}_t^\top \hat{\boldsymbol{\theta}}_k))^2 + 2(g(\tilde{\mathbf{x}}_t^\top \hat{\boldsymbol{\theta}}_k) - g(\tilde{\mathbf{x}}_t^\top \boldsymbol{\theta}_0))^2 \\ &:= \mathbf{J}_1 + \mathbf{J}_2. \end{aligned}$$

We first analyze $\mathbf{J}_2$. In fact, define the event

$$\mathcal{E}_k := \{\|\hat{\boldsymbol{\theta}}_k - \boldsymbol{\theta}_0\| \leq R_k\},$$

then according to Lemma 1, $\mathbb{P}(\mathcal{E}_k) \leq 1 - 2e^{-c_1 c_{\min}^2 |I_k|/16} - 2/|I_k|$. On $\mathcal{E}_k$ we have

$$\mathbf{J}_2 \leq \frac{2}{c_{\phi^2}}(\tilde{\mathbf{x}}_t^\top \hat{\boldsymbol{\theta}}_k - \tilde{\mathbf{x}}_t^\top \boldsymbol{\theta}_0)^2 \leq \frac{2}{c_{\phi^2}}R_{\mathcal{X}}^2 \|\hat{\boldsymbol{\theta}}_k - \boldsymbol{\theta}_0\|^2 \leq \frac{2}{c_{\phi^2}}R_{\mathcal{X}}^2 R_k^2.$$

Therefore,

$$\mathbb{E}\mathbf{J}_2 \leq \frac{2}{c_{\phi^2}}R_{\mathcal{X}}^2 R_k^2 + 2B^2(2e^{-c_1 c_{\min}^2 |I_k|/16} + 2/|I_k|). \quad (34)$$

As for $\mathbf{J}_1$, on the event $\mathcal{E}_k$, we deduce from Lemma 4 that for any

$$\delta \in [4 \exp(-\bar{B}_{x,K} |I_k|^{\frac{2m-2}{2m+1}} / \log |I_k|, \frac{1}{2}]$$

, with probability at least $1 - 6\delta$,

$$\mathbf{J}_1 \leq 2 \left[\sup_{u \in [\delta_z, B-\delta_z]} (\hat{g}_k(u) - g(u)) \right]^2 \leq 2\bar{C}_{x,K}^2 |I_K|^{-\frac{2(m-1)}{2m+1}} \log |I_K| \left(\sqrt{d} + \sqrt{\log \frac{1}{\delta}} \right)^2.$$

By choosing $\delta = 1/|I_k|$, we have

$$\begin{aligned} \mathbb{E}\mathbf{J}_1 &\leq 2\bar{C}_{x,K}^2 |I_K|^{-\frac{2(m-1)}{2m+1}} \log |I_K| \left(\sqrt{d} + \sqrt{\log \frac{1}{\delta}} \right)^2 + 2B^2 \cdot 6\delta \\ &\leq 4\bar{C}_{x,K}^2 |I_K|^{-\frac{2(m-1)}{2m+1}} \log |I_K| \left(d + \log |I_K| \right) + \frac{12B^2}{|I_k|} \end{aligned} \quad (35)$$

Combining (33), (34) and (35), we obtain an upper bound for the expected regret at any time t during episode k :

$$\mathbb{E}R_t \leq \bar{C}_{x,K}^{(1)} |I_K|^{-\frac{2(m-1)}{2m+1}} \log |I_K| \left(d + \log |I_K| \right),$$

where $\bar{C}_{x,K}^{(1)} = \frac{1}{2}(2l_r + Bl'_r) \cdot \left[\frac{4}{c_\phi^2} R_{\mathcal{X}}^2 \left(\frac{10 \max\{\psi_x, 1\} (R_{\mathcal{X}} R_\Theta + B)}{c_{\min}} \right)^2 + 20B^2 + 4C_{x,K}'^2 \right]$. We choose $|I_k| = \lceil (l_k d)^{\frac{2m+1}{4m-1}} \rceil$. The total regret during the k -th episode is

$$\begin{aligned} \text{Regret}_k &= \sum_{t \in I_k} \mathbb{E}R_t + \sum_{t \in I'_k} \mathbb{E}R_t \\ &\leq B|I_k| + l_k \cdot \mathbb{E}R_t \\ &\leq B(l_k d)^{(2m+1)/(4m-1)} + B + l_k \cdot \bar{C}_{x,K}^{(1)} (l_k d)^{-(2m-2)/(4m-1)} \log T(d + \log T) \\ &\leq (2B + \bar{C}_{x,K}^{(1)}) l_k^{\frac{2m+1}{4m-1}} d^{\frac{2m+1}{4m-1}} \log T(1 + \log T/d). \end{aligned}$$

Finally, the total regret defined in (6) can be bounded by

$$\begin{aligned} \text{Regret}_\pi(T) &= \sum_{k=1}^K \text{Regret}_k \leq 2B\sqrt{T} + (2B + \bar{C}_{x,K}^{(1)}) d^{\frac{2m+1}{4m-1}} \log T(1 + \log T/d) \sum_{k=1}^K l_k^{(2m+1)/(4m-1)} \\ &\leq \left[2B + \frac{2l_0^{(2m+1)/(4m-1)} (2B + \bar{C}_{x,K}^{(1)})}{2^{(2m+1)/(4m-1)} - 1} \right] (Td)^{\frac{2m+1}{4m-1}} \log T \left(1 + \frac{\log T}{d} \right). \end{aligned} \quad (36)$$

Here $K = \lceil \log_2 T \rceil$. The proof is then finished by letting $C_{x,K}^* = 2B + \frac{2l_0^{(2m+1)/(4m-1)} (2B + \bar{C}_{x,K}^{(1)})}{2^{(2m+1)/(4m-1)} - 1}$.

D Proof under the strong-mixing feature setting

In this section, we mainly present the proof of Theorem 2. The proof will be decomposed to the following lemmas, and their proof is also attached.

Before stating the lemmas, we introduce the α -mixing condition.

Definition 1. [α -mixing] For a sequence of random variables x_i defined on a probability space $(\Omega, \mathcal{X}, \mathbb{P})$, define

$$\alpha_k = \sup_{l \geq 0} \alpha(\sigma(x_t, t \leq l), \sigma(x_t, t \geq l + k))$$

in which

$$\alpha(\mathcal{A}, \mathcal{B}) = \sup_{A \in \mathcal{A}, B \in \mathcal{B}} \{ |\mathbb{P}(A \cap B) - \mathbb{P}(A)\mathbb{P}(B)| \}$$

From the definition of strong β -mixing, we see that it can infer strong α -mixing conditions. So in this case, our sequence $\mathbf{x}_t$ also follows strong α -mixing conditions, with $\alpha_k \leq e^{-ck}$.

Lemma 6. *[Parametric estimation under dependence] Under Assumption 2 and 6, there exist positive constants c_1 and c_2 (only depend on constants given in Assumptions) such that when $|I_k| \geq \max\{c_1(d+1), c_2 \log^2 |I_k| \log \log |I_k|\}$, for any episode k within the horizon, with probability $1 - 4/|I_k|^2$, we obtain*

$$\|\hat{\boldsymbol{\theta}}_k - \boldsymbol{\theta}_0\|_2 \leq \frac{2}{c_{\min}} \sqrt{\frac{(d+1)(6W_x^2 \log |I_k| + 6W_x \log^2 |I_k| \log \log |I_k|)}{C_w |I_k|}},$$

where $W_x = 2R_{\mathcal{X}}(R_{\mathcal{X}}R_{\Theta} + B)$.

The proof of Lemma 6 can be found in §F.6. Next, we present the following results on estimation error of $F(\cdot)$ and $F'(\cdot)$:

Lemma 7. *Suppose that Assumptions 3, 4, 5, 6 and 7 hold. Then there exist constants $B_{mx,K}, B'_{mx,K}, C_{mx,K}$ only depending on $R_{\mathcal{X}} := \sup_{\mathbf{x} \in \mathcal{X}} \|\mathbf{x}\|_2$ and constants within assumptions, such that as long as*

$$T \geq B_{mx,K}(\log T + 2 \log d)^{\frac{12m-3}{m}} [(d+1) \log(d+1)]^{\frac{4m-1}{m}} / d^2,$$

we have for any $k \geq \lfloor (\log(\sqrt{T} + \ell_0) - \log \ell_0) / \log 2 \rfloor + 2$, and $\delta \in [8 \exp(-|I_k|^{\frac{2m}{2m+1}} / (B'_{mx,K} \log^2 |I_k|)), 1/2]$ with probability at least $1 - 2\delta$,

$$\sup_{u \in I, \boldsymbol{\theta} \in \Theta_k} |\hat{F}_k(u, \boldsymbol{\theta}) - F(u)| \leq C_{mx,K} |I_k|^{-\frac{m}{2m+1}} \log |I_k| \left(\sqrt{(d+1) \log(d+1) \log |I_k|} + \sqrt{2 \log \frac{8}{\delta}} \right). \quad (37)$$

Here $I = [-\delta_z, \delta_z]$ and we choose the bandwidth $b_k = |I_k|^{-\frac{1}{2m+1}}$.

The proof of Lemma 7 can be found in §F.7.

Lemma 8. *Suppose that Assumptions 3, 4, 5, 6 and 7 hold. Then there exist constants $\bar{B}_{mx,K}, \bar{B}'_{mx,K}, \bar{C}_{mx,K}$ that depending only on $R_{\mathcal{X}} := \sup_{\mathbf{x} \in \mathcal{X}} \|\mathbf{x}\|_2$ and the constants within the assumptions such that as long as*

$$T \geq \bar{B}_{mx,K}(\log T + 2 \log d)^{\frac{12m-3}{m}} [(d+1) \log(d+1)]^{\frac{4m-1}{m}} / d^2,$$

for any $k \geq \lfloor (\log(\sqrt{T} + \ell_0) - \log \ell_0) / \log 2 \rfloor + 2$ and $\delta \in [\{8 \exp(-|I_k|^{\frac{2m}{2m+1}} / (\bar{B}'_{mx,K} \log^2 |I_k|)), 1/2]$ we have with probability at least $1 - 4\delta$,

$$\sup_{u \in I, \boldsymbol{\theta} \in \Theta_k} |\hat{F}_k^{(1)}(u, \boldsymbol{\theta}) - F'(u)| \leq \bar{C}_{mx,K} |I_k|^{-\frac{m-1}{2m+1}} \log |I_k| \left(\sqrt{(d+1) \log(d+1) \log |I_k|} + \sqrt{2 \log \frac{8}{\delta}} \right). \quad (38)$$

Here $I = [-\delta_z, \delta_z]$ and we choose the bandwidth $b_k = |I_k|^{-\frac{1}{2m+1}}$.

The proof of this lemma can be found in §F.8.

By combining these two lemmas and following our conclusions from Lemma 4, we are able to achieve the regret bound at the same order with Theorem 1 in Theorem 2.

E Proof under the super smooth noise distribution setting

Proof of Theorem 3 can be followed directly from the proof of Theorem 2 by substituting the Lemma 5 with Lemma 6. Below we'll only present the proof of Lemma 5.

Proof. We only bound $\sup_{u \in I, \theta \in \Theta_k} |\mathbb{E}[f_k(u, \theta)] - f_\theta(u)|$ and $\sup_{u \in I, \theta \in \Theta_k} |\mathbb{E}[f_k^{(1)}(u, \theta)] - f'_\theta(u)|$, since the analysis for $f_k(u, \theta)$ and $h_k(u, \theta)$ are the same. In fact, under the settings of Lemma 5, for any $u \in I, \theta \in \Theta_k$,

$$\begin{aligned} \mathbb{E}[f_k(u, \theta)] - f_\theta(u) &= \int_{\mathbb{R}} \frac{1}{b_k} K\left(\frac{s-u}{b_k}\right) f_\theta(s) ds - f_\theta(u) \\ &= \mathcal{F}\left(\mathcal{F}^{-1}\left(\int_{\mathbb{R}} \frac{1}{b_k} K\left(\frac{s-u}{b_k}\right) f_\theta(s) ds\right) - \mathcal{F}^{-1} \circ f_\theta(u)\right) \\ &= \mathcal{F}\left(\phi_\theta(u) \left[\mathcal{F}^{-1}\left(\frac{1}{b_k} K\left(\frac{-u}{b_k}\right)\right) - 1\right]\right) \\ &= \mathcal{F}(\phi_\theta(u) [\kappa(-b_k u) - 1]). \end{aligned}$$

Here $\mathcal{F}$ is the Fourier transform operator defined by

$$g \rightarrow \mathcal{F} \circ g(u) = \frac{1}{2\pi} \int_{\mathbb{R}} g(x) e^{-iux} dx,$$

and we've utilized the fact that $K = \mathcal{F} \circ \kappa$, $\phi_\theta(u) = \mathcal{F}^{-1} \circ f_\theta$. Since $|\kappa(x)| \leq 1$ for all $x \in \mathbb{R}$ and that $\kappa(x) = 1$ for $|x| \leq c_\kappa$,

$$\begin{aligned} \sup_{u \in I, \theta \in \Theta_k} |\mathbb{E}[f_k(u, \theta)] - f_\theta(u)| &\leq \sup_{u \in I, \theta \in \Theta_k} |\mathcal{F}(\phi_\theta(u) [\kappa(-b_k u) - 1])| \\ &\leq \sup_{\theta \in \Theta_k} \frac{1}{2\pi} \int |\phi_\theta(s)| \cdot |\kappa(-b_k s) - 1| ds \\ &\leq \sup_{\theta \in \Theta_0} \frac{1}{\pi} \int_{|s| > c_\kappa/b_k} |\phi_\theta(s)| ds \\ &\leq \frac{2}{\pi} \int_{s>0} D_\phi e^{-d_\phi(s+c_\kappa/b_k)^\alpha} ds \\ &\leq \frac{2}{\pi} \int_{s>0} D_\phi e^{-d_\phi/2 \cdot [s^\alpha + (c_\kappa/b_k)^\alpha]} ds. \end{aligned}$$

Here, the last inequality is due to the fact that for $x, y \in \mathbb{R}$, $(x+y)^\alpha \geq \min\{2^{\alpha-1}, 1\}(x^\alpha + y^\alpha) \geq \frac{1}{2}(x^\alpha + y^\alpha)$. Thus, by choosing $b_k = c_\kappa(d_\phi/\log|I_k|)^{1/\alpha}$, we obtain that

$$\sup_{u \in I, \theta \in \Theta_k} |\mathbb{E}[f_k(u, \theta)] - f_\theta(u)| \leq C_{\inf}/\sqrt{n},$$

where $C_{\inf} = 2D_\phi/\pi \cdot \int_{s>0} \exp(-d_\phi s^\alpha/2) ds$.

The analysis for $\sup_{u \in I, \theta \in \Theta_k} |\mathbb{E}[f_k^{(1)}(u, \theta)] - f'_\theta(u)|$ is similar as above. In fact, for any $u \in I, \theta \in \Theta_k$,

$$\begin{aligned}
\mathbb{E}[f_k^{(1)}(u, \theta)] - f'_\theta(u) &= - \int_{\mathbb{R}} \frac{1}{b_k^2} K' \left(\frac{s-u}{b_k} \right) f_\theta(s) ds - f'_\theta(u) \\
&= \int_{\mathbb{R}} \frac{1}{b_k} K \left(\frac{s-u}{b_k} \right) f'_\theta(s) ds - f'_\theta(u) \\
&= \mathcal{F} \left(\mathcal{F}^{-1} \left(\int_{\mathbb{R}} \frac{1}{b_k} K \left(\frac{s-u}{b_k} \right) f'_\theta(s) ds \right) - \mathcal{F}^{-1} \circ f'_\theta(u) \right) \\
&= \mathcal{F} \left(\phi_\theta^{(1)}(u) \left[\mathcal{F}^{-1} \left(\frac{1}{b_k} K \left(\frac{-u}{b_k} \right) \right) - 1 \right] \right) \\
&= \mathcal{F}(\phi_\theta^{(1)}(u) [\kappa(-b_k u) - 1]).
\end{aligned}$$

Following the same arguments as above, we deduce that

$$\sup_{u \in I, \theta \in \Theta_k} |\mathbb{E}[f_k^{(1)}(u, \theta)] - f'_\theta(u)| \leq C_{\text{inf}}/\sqrt{n}.$$

□

F Proof of technical lemmas

F.1 Proof of Lemma 1

We only prove (15), since (16) can be proved in the same way.

Recall that $h_k(u, \theta) = \frac{1}{nb_k} \sum_{t=1}^n K(\frac{w_t(\theta) - u}{b_k}) y_t$, and $\mathbb{E}[y_t | w_t(\theta) = u] = r_\theta(u) = \frac{h_\theta(u)}{f_\theta(u)}$. We have

$$\mathbb{E}h_k(u, \theta) = \frac{1}{b_k} \mathbb{E}K\left(\frac{w_t(\theta) - u}{b_k}\right) y_t = \frac{1}{b_k} \mathbb{E}K\left(\frac{w_t(\theta) - u}{b_k}\right) r(w_t(\theta)).$$

Thus,

$$\begin{aligned}
\mathbb{E}h_k(u, \theta) - h_\theta(u) &= \int \frac{1}{b_k} K\left(\frac{w(\theta) - u}{b_k}\right) r_\theta(w(\theta)) f_\theta(w(\theta)) dw(\theta) - h_\theta(u) \\
&= \int K(s) h_\theta(u + b_k s) ds - h_\theta(u).
\end{aligned} \tag{39}$$

Using Taylor's expansion, $\forall s \in \mathbb{R}$, there exists some $\xi(s, u)$ lying between the points u and $u + b_k s$ such that

$$h_\theta(u + b_k s) = h_\theta(u) + \sum_{i=1}^{m-2} \frac{h_\theta^{(i)}(u)}{i!} (b_k s)^i + \frac{h_\theta^{(m-1)}(\xi(s, u))}{(m-1)!} (b_k s)^{m-1}.$$

Plugging this into (39) gives

$$\begin{aligned}
\mathbb{E}h_k(u, \boldsymbol{\theta}) - h_{\boldsymbol{\theta}}(u) &= \int K(s) \left[h_{\boldsymbol{\theta}}(u) + \sum_{i=1}^{m-2} \frac{h_{\boldsymbol{\theta}}^{(i)}(u)}{i!} (b_k s)^i + \frac{h_{\boldsymbol{\theta}}^{(m-1)}(\xi(s, u))}{(m-1)!} (b_k s)^{m-1} \right] ds - h_{\boldsymbol{\theta}}(u) \\
&= \int K(s) \frac{h_{\boldsymbol{\theta}}^{(m-1)}(\xi(s, u))}{(m-1)!} (b_k s)^{m-1} ds \\
&= \int K(s) \frac{h_{\boldsymbol{\theta}}^{(m-1)}(u)}{(m-1)!} (b_k s)^{m-1} ds + \int K(s) \frac{[h_{\boldsymbol{\theta}}^{(m-1)}(\xi(s, u)) - h_{\boldsymbol{\theta}}^{(m-1)}(u)]}{(m-1)!} (b_k s)^{m-1} ds \\
&= \int K(s) \frac{[h_{\boldsymbol{\theta}}^{(m-1)}(\xi(s, u)) - h_{\boldsymbol{\theta}}^{(m-1)}(u)]}{(m-1)!} (b_k s)^{m-1} ds.
\end{aligned}$$

Thus we have that

$$\begin{aligned}
|\mathbb{E}h_k(u, \boldsymbol{\theta}) - h_{\boldsymbol{\theta}}(u)| &\leq \int |K(s)| \frac{|h_{\boldsymbol{\theta}}^{(m-1)}(\xi(s, u)) - h_{\boldsymbol{\theta}}^{(m-1)}(u)|}{(m-1)!} |b_k s|^{m-1} ds \\
&\leq \int |K(s)| \frac{l_f |b_k s|}{(m-1)!} |b_k s|^{m-1} ds \\
&\leq C_1 b_k^m,
\end{aligned}$$

where $C_1 = l_f \cdot \int |s^m K(s)| ds / (m-1)!$. Moreover, since the inequality holds for any $u \in I$ and $\boldsymbol{\theta} \in \Theta_k$, we finish the proof.

F.2 Proof of Lemma 2

We only prove (17), since (18) can be proved in the same way.

For any $u \in I$, $\boldsymbol{\theta} \in \Theta_k$, denote $Z(u, \boldsymbol{\theta}) := h_k(u, \boldsymbol{\theta}) - \mathbb{E}h_k(u, \boldsymbol{\theta}) = \frac{1}{nb_k} \sum_{t \in I_k} [K(\frac{w_t(\boldsymbol{\theta}) - u}{b_k}) y_t - \mathbb{E}K(\frac{w_t(\boldsymbol{\theta}) - u}{b_k}) y_t]$. Then

$$\sup_{u \in I, \boldsymbol{\theta} \in \Theta_k} |h_k(u, \boldsymbol{\theta}) - \mathbb{E}h_k(u, \boldsymbol{\theta})| = \sup_{u \in I, \boldsymbol{\theta} \in \Theta_k} |Z(u, \boldsymbol{\theta})| = \max \left\{ \sup_{u \in I, \boldsymbol{\theta} \in \Theta_k} Z(u, \boldsymbol{\theta}), \sup_{u \in I, \boldsymbol{\theta} \in \Theta_k} (-Z(u, \boldsymbol{\theta})) \right\}.$$

We can then bound $\sup_{u \in I, \boldsymbol{\theta} \in \Theta_k} |h_k(u, \boldsymbol{\theta}) - \mathbb{E}h_k(u, \boldsymbol{\theta})|$ by upper bounding both $\sup_{u \in I, \boldsymbol{\theta} \in \Theta_k} Z(u, \boldsymbol{\theta})$ and $\sup_{u \in I, \boldsymbol{\theta} \in \Theta_k} (-Z(u, \boldsymbol{\theta}))$. We now give upper bound for $\sup_{u \in I, \boldsymbol{\theta} \in \Theta_k} Z(u, \boldsymbol{\theta})$ with high probability (Bounding $\sup_{u \in I, \boldsymbol{\theta} \in \Theta_k} (-Z(u, \boldsymbol{\theta}))$ is essentially the same).

We use the chaining method to obtain the desired bound. First, we construct a sequence of ε -nets with decreasing scale. Denote the left and right endpoints of the interval I as L_I and R_I respectively. For any $i \in \mathbb{N}^+$, construct set $S_1^{(i)} \subseteq I$ as

$$S_1^{(i)} \triangleq \left\{ L_I + \frac{j}{2^i \sqrt{n}} (R_I - L_I) : j \in \{1, 2, \dots, (2^i - 1) \lfloor \sqrt{n} \rfloor\} \right\}.$$

For any $u \in I$, $i \in \mathbb{N}^+$, let $\pi_1^{(i)}(u) = \arg \min_{s \in S_1^{(i)}} |s - u|$. Moreover, let $\pi_1^{(0)}(u) = u$. Then we can easily verify that $|S_1^{(i)}| \leq 2^i(\sqrt{n} + 1)$, and $\forall u \in I$, $|\pi_i(u) - \pi_{i+1}(u)| \leq \frac{2\delta_z}{2^{i-1}\sqrt{n}}$. At the same time, denote $S_2^{(i)}$ as a $R_k/2^i$ -net with respect to l_2 -distance of Θ_k , where R_k denotes the radius of Θ_k . Similar to $\pi_1^{(i)}$, define $\pi_2^{(i)}(\mathbf{u}) = \arg \min_{s \in S_2^{(i)}} |\mathbf{u} - s|$. By Corollary 4.2.13 in Vershynin [2018], $|S_2^{(i)}| \leq (2^{i+1} + 1)^d$.

Combining the above two nets, we have $S^{(i)} := S_1^{(i)} \times S_2^{(i)}$ is a $2^{-i}\sqrt{4\delta_z^2/n + R_k^2}$ -net of $U_k := I \times \Theta_k$ with cardinality $|S^{(i)}| \leq 2^i(\sqrt{n} + 1) \cdot (2^{i+1} + 1)^d$. In fact, for any $\mathbf{u} := (u, \boldsymbol{\theta}) \in I \times \Theta_k$ with $i \geq 1$, denote $\pi_i(\mathbf{u}) := (\pi_1^{(i)}(u), \pi_2^{(i)}(\boldsymbol{\theta}))$, then $\|\pi_i(\mathbf{u}) - \mathbf{u}\|_2 \leq 2^{-i}\sqrt{4\delta_z^2/n + R_k^2}$.

Now, since $Z(u, \boldsymbol{\theta})$ is continuous a.s., we have for any $M \in \mathbb{N}^+$

$$Z(\mathbf{u}) - Z(\pi_M(\mathbf{u})) = \sum_{i=M}^{\infty} [Z(\pi_{i+1}(\mathbf{u})) - Z(\pi_i(\mathbf{u}))],$$

and thus

$$\sup_{\mathbf{u} \in U_k} Z(\mathbf{u}) \leq \sup_{\mathbf{u} \in U_k} Z(\pi_M(\mathbf{u})) + \sum_{i=M}^{\infty} \sup_{\mathbf{u} \in U_k} [Z(\pi_{i+1}(\mathbf{u})) - Z(\pi_i(\mathbf{u}))] \quad (40)$$

almost surely. Our goal is to choose a suitable M such that both terms on the right hand side of (40) can be controlled in a reasonable manner.

For this reason, Let $M = \lceil \frac{3}{\log 2} \log \frac{1}{b_k} \rceil + 10$. We first upper bound $\sup_{\mathbf{u} \in U_k} Z(\pi_M(\mathbf{u}))$. Note that

$$Z(\mathbf{u}) = \frac{1}{nb_k} \sum_{t \in I_k} A_t(\mathbf{u}),$$

where $A_t(\mathbf{u}) = K(\frac{w_t(\boldsymbol{\theta}) - u}{b_k})Y_t - \mathbb{E}K(\frac{w_t(\boldsymbol{\theta}) - u}{b_k})Y_t$. We have $\mathbb{E}A_t(\mathbf{u}) = 0$ and $|A_t(\mathbf{u})| \leq \bar{K}$ almost surely. Moreover,

$$\begin{aligned} \text{Var}(A_t(\mathbf{u})) &\leq \mathbb{E} \left[K\left(\frac{w_t(\boldsymbol{\theta}) - u}{b_k}\right)Y_t \right]^2 \leq \mathbb{E} \left[K\left(\frac{w_t(\boldsymbol{\theta}) - u}{b_k}\right) \right]^2 \\ &\leq \int K\left(\frac{w_t(\boldsymbol{\theta}) - u}{b_k}\right)^2 f_{\boldsymbol{\theta}}(w_t(\boldsymbol{\theta})) dw_t(\boldsymbol{\theta}) = b_k \int K(s)^2 f_{\boldsymbol{\theta}}(u + b_k s) ds \leq C_4 b_k, \end{aligned}$$

where $C_4 = \max\{\bar{f} \cdot \int K(s)^2 ds, \bar{f} \cdot \int K(s)^2 ds\}$. Thus according to Bernstein's Inequality, for any $\epsilon > 0$,

$$\mathbb{P}(|Z(\mathbf{u})| \geq \epsilon) = \mathbb{P}(|\sum_{t \in I_k} A_t(\mathbf{u})| \geq nb_k \epsilon) \leq 2e^{-\frac{n^2 b_k^2 \epsilon^2}{2C_4 nb_k + \frac{2}{3}\bar{K} nb_k \epsilon}} \leq 2e^{-C_5 \frac{nb_k \epsilon^2}{1 + \epsilon}},$$

where $C_5 = 1/\max\{2C_4, \frac{2}{3}\bar{K}, 1\}$. A union bound then gives

$$\begin{aligned}
\mathbb{P}(\sup_{\mathbf{u} \in \mathcal{U}_k} |Z(\pi_M(\mathbf{u}))| \geq \epsilon) &\leq |S^{(M)}| \cdot \mathbb{P}(|Z(\mathbf{u})| \geq \epsilon) \\
&\leq 2^M (\sqrt{n} + 1) \cdot (2^{M+1} + 1)^d \cdot 2e^{-C_5 \frac{nb_k \epsilon^2}{1+\epsilon}} \\
&\leq \exp \left(4dM \log 2 + \log n - \frac{C_5}{2} nb_k \min\{\epsilon, \epsilon^2\} \right).
\end{aligned}$$

When $\delta \geq 4e^{-nb_k/3}$ and $nb_k \geq \max\{C_b d(\log \frac{1}{b_k} + 1), 3 \log n\}$ for some absolute constant $C_b > 0$, by choosing

$\epsilon = \epsilon(k) = \frac{2}{C_5} \frac{1}{\sqrt{nb_k}} \sqrt{4dM \log 2 + \log n + \log \frac{4}{\delta}}$, we can verify that the last term above is upper bounded by $\frac{\delta}{4}$, and thus we have

$$\mathbb{P} \left(\sup_{\mathbf{u} \in \mathcal{U}_k} |Z(\pi_M(\mathbf{u}))| \geq \epsilon(k) \right) \leq \frac{\delta}{4}. \quad (41)$$

Now we proceed to bound the latter term on the right hand side of (40). For any $\mathbf{u}_1 := (u, \boldsymbol{\theta}_1), \mathbf{u}_2 := (s, \boldsymbol{\theta}_2) \in I \times \Theta_k$, we have

$$Z(\mathbf{u}_1) - Z(\mathbf{u}_2) = Z(u, \boldsymbol{\theta}_1) - Z(s, \boldsymbol{\theta}_2) = \frac{1}{nb_k} \sum_{t \in I_k} B_t(u, \boldsymbol{\theta}_1, s, \boldsymbol{\theta}_2),$$

where

$$B_t(u, \boldsymbol{\theta}_1, s, \boldsymbol{\theta}_2) = y_t \left(K\left(\frac{w_t(\boldsymbol{\theta}_1) - u}{b_k}\right) - K\left(\frac{w_t(\boldsymbol{\theta}_2) - s}{b_k}\right) \right) - \mathbb{E} y_t \left(K\left(\frac{w_t(\boldsymbol{\theta}_1) - u}{b_k}\right) - K\left(\frac{w_t(\boldsymbol{\theta}_2) - s}{b_k}\right) \right).$$

Then $\mathbb{E} B_j(u, \boldsymbol{\theta}_1, s, \boldsymbol{\theta}_2) = 0$, and

$$\begin{aligned}
|Z(\mathbf{u}_1) - Z(\mathbf{u}_2)| &= |B_t(u, \boldsymbol{\theta}_1, s, \boldsymbol{\theta}_2)| \leq 2 \left| y_t \left(K\left(\frac{w_t(\boldsymbol{\theta}_1) - u}{b_k}\right) - K\left(\frac{w_t(\boldsymbol{\theta}_2) - s}{b_k}\right) \right) \right| \\
&\leq \frac{2l_K \sqrt{(\max_{\mathbf{x} \in \mathcal{X}} \|\mathbf{x}\|_2^2 + 1)}}{b_k} \cdot \|\mathbf{u}_1 - \mathbf{u}_2\|_2.
\end{aligned}$$

Using Hoeffding's Inequality, for any $\epsilon > 0$,

$$\mathbb{P}(|\sum_{t \in I_k} B_t(\mathbf{u}_1, \mathbf{u}_2)| \geq \epsilon) \leq 2e^{-\frac{2\epsilon^2}{4l_K^2(R_{\mathcal{X}}^2 + 1)/b_k^2 \cdot n \|\mathbf{u}_1 - \mathbf{u}_2\|_2^2}} = 2e^{-\frac{b_k^2 \epsilon^2}{2l_K^2 n (R_{\mathcal{X}}^2 + 1) \|\mathbf{u}_1 - \mathbf{u}_2\|_2^2}}$$

Therefore,

$$\mathbb{P}(|Z(\mathbf{u}_1) - Z(\mathbf{u}_2)| \geq \epsilon) = \mathbb{P}(|\sum_{t \in I_k} B_t(\mathbf{u}_1, \mathbf{u}_2)| \geq nb_k \epsilon) \leq 2e^{-\frac{nb_k^4 \epsilon^2}{2l_K^2 (R_{\mathcal{X}}^2 + 1) \|\mathbf{u}_1 - \mathbf{u}_2\|_2^2}}.$$

Recall that $\forall \mathbf{u}, \|\pi_i(\mathbf{u}) - \pi_{i+1}(\mathbf{u})\|_2 \leq 2^{-i} \sqrt{4\delta_z^2/n + R_k^2}$. We use union bound to obtain

$$\begin{aligned} & \mathbb{P}\left(\sup_{\mathbf{u} \in \mathcal{U}_k} |Z(\pi_{i+1}(\mathbf{u})) - Z(\pi_i(\mathbf{u}))| \geq \epsilon\right) \\ & \leq 2^i(\sqrt{n} + 1) \cdot (2^{i+1} + 1)^d \cdot 2e^{-\frac{2^{2i-2}n^2b_k^4\epsilon^2}{2l_K^2(R_{\mathcal{X}}^2+1)(4\delta_z^2+nR_k^2)}}. \end{aligned}$$

Let $\epsilon = \frac{l_K \sqrt{(R_{\mathcal{X}}^2+1)(4\delta_z^2+nR_k^2)}\epsilon_i}{2^{i-1}nb_k^2}$. The above inequality reduces to

$$\begin{aligned} & \mathbb{P}\left(\sup_{\mathbf{u} \in \mathcal{U}_k} |Z(\pi_{i+1}(\mathbf{u})) - Z(\pi_i(\mathbf{u}))| \geq \frac{l_K \sqrt{(R_{\mathcal{X}}^2+1)(4\delta_z^2+nR_k^2)}\epsilon_i}{2^{i-1}nb_k^2}\right) \\ & \leq 2^i(\sqrt{n} + 1) \cdot (2^{i+1} + 1)^d \cdot 2e^{-\frac{\epsilon_i^2}{2}}. \end{aligned} \quad (42)$$

Now we choose $\epsilon_i = \sqrt{2 \log \frac{8}{\delta} + \log n + (2i+4)(d+2) \log 2}$ and define $W^* := \sqrt{(R_{\mathcal{X}}^2+1)(4\delta_z^2+nR_k^2)}$. Notice that

$$\begin{aligned} \sum_{i=M}^{\infty} \frac{l_K W^*}{nb_k^2} \frac{\epsilon_i}{2^{i-1}} & \leq \frac{l_K W^*}{nb_k^2} \sum_{i=M}^{\infty} \frac{\sqrt{2id \log 2} + \sqrt{(4d+8) \log 2 + \log n + 2 \log \frac{8}{\delta}}}{2^{i-1}} \\ & \leq \frac{l_K W^*}{nb_k^2} \left[\sqrt{2d \log 2} \sum_{i=M}^{\infty} \frac{i}{2^{i-1}} + \frac{1}{2^{M-2}} \sqrt{(4d+8) \log 2 + \log n + 2 \log \frac{8}{\delta}} \right] \\ & \leq \frac{l_K W^*}{nb_k^2} \left[\sqrt{2d \log 2} \frac{M+1}{2^{M-2}} + \frac{1}{2^{M-2}} \sqrt{(4d+8) \log 2 + \log n + 2 \log \frac{8}{\delta}} \right] \\ & \leq \frac{l_K W^*}{\sqrt{nb_k}} \cdot \frac{1}{n^{1/2}b_k^{3/2}} \frac{M+2}{2^{M-2}} \left[\sqrt{2 \log \frac{8}{\delta} + \log n + 4\sqrt{d \log 2}} \right] \\ & \leq \frac{l_K W^*}{\sqrt{nb_k}} \left[\sqrt{\frac{2}{n} \log \frac{8}{\delta} + 1} + \frac{6\sqrt{\log 2}}{\sqrt{c_0}} \right] \end{aligned}$$

Here we use the fact that when $B_{x,K} \geq (2c_0)^4$, combining the assumptions in the lemma, we have $n \geq c_0 d$. Combining this fact and a union bound on (42), we get

$$\begin{aligned} & \mathbb{P}\left(\sup_{\mathbf{u} \in \mathcal{U}_k} |Z(\mathbf{u}) - Z(\pi_M(\mathbf{u}))| \geq \frac{l_K W^*}{\sqrt{nb_k}} \left[\sqrt{\frac{2}{n} \log \frac{8}{\delta} + 1} + \frac{6\sqrt{\log 2}}{\sqrt{c_0}} \right]\right) \\ & \leq \mathbb{P}\left(\sup_{\mathbf{u} \in \mathcal{U}_k} |Z(\mathbf{u}) - Z(\pi_M(\mathbf{u}))| \geq \sum_{i=M}^{\infty} \frac{l_K W^*}{nb_k^2} \frac{\epsilon_i}{2^{i-1}}\right) \\ & \leq \mathbb{P}\left(\sum_{i=M}^{\infty} \sup_{\mathbf{u} \in \mathcal{U}_k} |Z(\pi_{i+1}(\mathbf{u})) - Z(\pi_i(\mathbf{u}))| \geq \sum_{i=M}^{\infty} \frac{l_K W^*}{nb_k^2} \frac{\epsilon_i}{2^{i-1}}\right) \end{aligned}$$

$$\begin{aligned}
&\leq \sum_{i=M}^{\infty} \mathbb{P} \left(\sup_{\mathbf{u} \in U_k} |Z(\pi_{i+1}(\mathbf{u})) - Z(\pi_i(\mathbf{u}))| \geq \frac{l_K W^*}{n b_k^2} \frac{\epsilon_i}{2^{i-1}} \right) \\
&\leq \sum_{i=M}^{\infty} 2^i (\sqrt{n} + 1) \cdot (2^{i+1} + 1)^d \cdot 2e^{-\frac{\epsilon_i^2}{2}} \leq \sum_{i=M}^{\infty} \frac{\delta}{4} \cdot \frac{1}{2^{i+1}} \leq \frac{\delta}{4 \cdot 2^M} \leq \frac{\delta}{4}.
\end{aligned} \tag{43}$$

Finally, combining (40), (41) and (43), we obtain that

$$\begin{aligned}
\frac{\delta}{2} &\geq \mathbb{P} \left(\sup_{\mathbf{u} \in U_k} |Z(\pi_M(\mathbf{u}))| \geq \epsilon(k) \right) + \mathbb{P} \left(\sup_{\mathbf{u} \in U_k} |Z(\mathbf{u}) - Z(\pi_M(\mathbf{u}))| \geq \frac{l_K W^*}{\sqrt{n b_k}} \left[\sqrt{\frac{2}{n} \log \frac{8}{\delta}} + 1 + \frac{6\sqrt{\log 2}}{\sqrt{c_0}} \right] \right) \\
&\geq \mathbb{P} \left(\sup_{\mathbf{u} \in U_k} Z(\mathbf{u}) \geq \epsilon(k) + \frac{l_K W^*}{\sqrt{n b_k}} \left[\sqrt{\frac{2}{n} \log \frac{8}{\delta}} + 1 + \frac{6\sqrt{\log 2}}{\sqrt{c_0}} \right] \right) \\
&\geq \mathbb{P} \left(\sup_{\mathbf{u} \in U_k} Z(\mathbf{u}) \geq \frac{4\sqrt{11}/C_5}{\sqrt{n b_k}} \sqrt{d \left(1 + \log \frac{1}{b_k} \right) + \log n + \log \frac{4}{\delta}} + \right. \\
&\quad \left. 16\sqrt{2} \left(1 + \frac{6\sqrt{\log 2}}{c_0} \right) \frac{l_K \sqrt{1 + R_{\mathcal{X}}^2}}{\sqrt{n b_k}} \max \left\{ \delta_z, \frac{\max\{1, \psi_x\}(B + R_{\mathcal{X}} R_{\Theta})}{c_{\min}} \right\} \left(\sqrt{d \log n} + \sqrt{\frac{d \log n}{n} \log \frac{8}{\delta}} \right) \right) \\
&\geq \mathbb{P} \left(\sup_{\mathbf{u} \in U_k} Z(\mathbf{u}) \geq C_x l_K \sqrt{\frac{\log n}{n b_k}} \left(\sqrt{d} + \sqrt{\log 1/\delta} \right) \right).
\end{aligned}$$

Here we let $C_x = 8\sqrt{22}/C_5 + \frac{60(6\sqrt{\log 2} + \sqrt{c_0})}{c_0} \sqrt{1 + R_{\mathcal{X}}^2} \max\{\delta_z, \frac{\max\{1, \psi_x\}(B + R_{\mathcal{X}} R_{\Theta})}{c_{\min}}\}$.

For the same reason, we have that

$$\mathbb{P} \left(\sup_{\mathbf{u} \in U_k} (-Z(\mathbf{u})) \geq C_x l_K \sqrt{\frac{\log n}{n b_k}} \left(\sqrt{d} + \sqrt{\log 1/\delta} \right) \right) \leq \frac{\delta}{2}.$$

Combining the above two inequalities, we finish the proof.

F.3 Proof of Lemma 3

We only prove (23), since (24) can be proved in a similar way. Recall $h_k^{(1)}(u, \boldsymbol{\theta}) = \frac{-1}{n b_k^2} \sum_{t \in I_k} K'(\frac{w_t(\boldsymbol{\theta}) - u}{b_k}) y_t$, we have

$$\mathbb{E} h_k^{(1)}(\boldsymbol{\theta}, u) = \frac{-1}{b_k^2} \mathbb{E} K'(\frac{w_t(\boldsymbol{\theta}) - u}{b_k}) y_u = \frac{-1}{b_k^2} \mathbb{E} K'(\frac{w_t(\boldsymbol{\theta}) - u}{b_k}) r(w_t(\boldsymbol{\theta})).$$

Then

$$\begin{aligned}
\mathbb{E} h_k^{(1)}(u, \boldsymbol{\theta}) - h'_{\boldsymbol{\theta}}(u) &= \int \frac{-1}{b_k^2} K'(\frac{w_t(\boldsymbol{\theta}) - u}{b_k}) h_{\boldsymbol{\theta}}(w_t(\boldsymbol{\theta})) dw_t(\boldsymbol{\theta}) - h'_{\boldsymbol{\theta}}(u) \\
&= \int K(s) h'_{\boldsymbol{\theta}}(u + b_k s) ds - h'_{\boldsymbol{\theta}}(u),
\end{aligned} \tag{44}$$

where (44) follows from integration by parts. By Taylor's expansion, we have

$$h'_\theta(u + b_k s) = h'_\theta(u) + \sum_{i=2}^{m-2} \frac{h_\theta^{(i)}(u)}{(i-1)!} (b_k s)^{i-1} + \frac{h_\theta^{(m-1)}(\xi(s, u))}{(m-1)!} (b_k s)^{m-2}.$$

Similar to our proof procedure of Lemma 1, under Assumption 5, we get

$$\mathbb{E}h_k^{(1)}(u, \theta) - h'_\theta(u) = \int K(s) \frac{h_\theta^{(m-1)}(\xi(s, u)) - h_\theta^{(m-1)}(u)}{(m-2)!} (b_k s)^{m-2} ds.$$

Thus

$$\begin{aligned} |\mathbb{E}h_k^{(1)}(u, \theta) - h'_\theta(u)| &\leq \int |K(s)| \frac{|h_\theta^{(m-1)}(\xi(s, u)) - h_\theta^{(m-1)}(u)|}{(m-2)!} (b_k s)^{m-2} |ds| \\ &\leq |K(s)| \frac{l_f |b_k s|}{(m-2)!} |b_k s|^{m-2} ds \\ &\leq C_{x,K}^{(5)} b_k^{m-1}, \end{aligned} \tag{45}$$

in which $C_{x,K}^{(5)} = \frac{l_f}{(m-2)!} \int |K(s)| s^{m-1} |ds|$. Because (45) holds for any $t \in I$ and $\theta \in \Theta_k$, we have

$$\sup_{u \in I, \theta \in \Theta_k} |\mathbb{E}h_k^{(1)}(u, \theta) - h'_\theta(u)| \leq C_{x,K}^{(5)} b_k^{m-1},$$

which claims inequality 23 of Lemma 3. On the other hand, (24) follows directly from our proof procedure above, so we omit the details.

F.4 Proof of Lemma 4

For any $u \in I$, $\theta \in \Theta_k$, write

$$Z^{(1)}(u, \theta) = h_k^{(1)}(u, \theta) - \mathbb{E}h_k^{(1)}(u, \theta) = \frac{-1}{b_k} \cdot \frac{1}{nb_k} \sum_{t \in I_k} \left[K' \left(\frac{w_t(\theta) - u}{b_k} \right) y_t - \mathbb{E} K' \left(\frac{w_t(\theta) - u}{b_k} \right) y_t \right]$$

Under Assumption 4 and Assumption 5, by following a similar proof procedure with Lemma 2, for $\delta \in [4e^{-nb_k/3}, \frac{1}{2})$, with probability at least $1 - \delta$,

$$\sup_{u \in I, \theta \in \Theta_k} \left| \frac{1}{nb_k} \sum_{t \in I_k} \left[K' \left(\frac{w_t(\theta) - u}{b_k} \right) y_t - \mathbb{E} K' \left(\frac{w_t(\theta) - u}{b_k} \right) y_t \right] \right| \leq C_{x,K}^{(2)} \sqrt{\frac{\log n}{nb_k}} \left(\sqrt{d} + \sqrt{\log 1/\delta} \right),$$

where $C_{x,K}^{(2)} = l_K \left(8\sqrt{22} \max\{2\bar{f} \int K^2 ds, 2\bar{f} \int K'^2 ds, \frac{2}{3} \bar{K}, 1\} + \right.$

$\frac{60(6\sqrt{\log 2} + \sqrt{c_0})}{c_0} \sqrt{1 + R_{\mathcal{X}}^2} \max\{\delta_z, \frac{\max\{1, \psi_x\}(B + R_{\mathcal{X}} R_{\Theta})}{c_{\min}}\}\bigg). \text{ Thus, with probability at least } 1 - \delta,$

$$\sup_{u \in I, \theta \in \Theta_k} |h_k^{(1)}(u, \theta) - \mathbb{E}h_k^{(1)}(u, \theta)| \leq C_{x,K}^{(2)} \sqrt{\frac{\log n}{nb_k^3}} \left(\sqrt{d} + \sqrt{\log 1/\delta} \right),$$

which claims the inequality (25) in Lemma 4. Moreover, (26) also follows directly from our procedure given above. Thus, we claim our conclusion of Lemma 4.

F.5 Proof of Lemma 5

First, we argue that for any $\tilde{\mathbf{x}}_t$,

$$\boldsymbol{\theta}_0^\top \tilde{\mathbf{x}}_t \in [\delta_z + \delta_v, B - \delta_z - \delta_v]. \quad (46)$$

In fact, we have $v_t = \boldsymbol{\theta}_0^\top \tilde{\mathbf{x}}_t + z_t$, where $z_t \in [-\delta_z, \delta_z]$ and that $\boldsymbol{\theta}_0^\top \tilde{\mathbf{x}}_t$ is independent from z_t . Therefore, in order to satisfy the condition $v_t \in [\delta_v, B - \delta_v]$, it ought to be true that $\boldsymbol{\theta}_0^\top \tilde{\mathbf{x}}_t \in [\delta_z + \delta_v, B - \delta_z - \delta_v]$.

On the other hand,

$$\begin{aligned} \sup_{\tilde{\mathbf{x}}_t \in \mathcal{X}, \theta \in \Theta_0} |\theta^\top \tilde{\mathbf{x}}_t - \boldsymbol{\theta}_0^\top \tilde{\mathbf{x}}_t| &\leq \sup_{\theta \in \Theta_0} \|\theta - \boldsymbol{\theta}_0\| \cdot \sup_{\tilde{\mathbf{x}}_t \in \mathcal{X}} \|\tilde{\mathbf{x}}_t\| \\ &\leq C_\theta T^{-\frac{2m+1}{4(4m-1)}} d^{\frac{m-1}{4m-1}} \sqrt{\log T + 2 \log d} \cdot R_{\mathcal{X}} \\ &\leq \delta_v. \end{aligned} \quad (47)$$

The last inequality is due to the condition on T . The lemma is proved by combining (46) and (47).

F.6 Proof of Lemma 6

The proof of Lemma 6 is similar with our proof of Lemma 1, the major difference between them is that here we assume our covaraites $\tilde{\mathbf{x}}_t, t \geq 0$ follow β -mixing condition instead of of i.i.d. assumption. After following similar proof procedures of (3)-(7), we obtain the same inequality with (7) and we also divide the following proofs into two steps.

Step I: In this step, we prove under β -mixing conditions given in Assumption 6, with high-probability, there exists a constant $c > 0$ such that $\lambda_{\min}(\frac{1}{|I_k|} \sum_{t \in I_k} \tilde{\mathbf{x}}_t \tilde{\mathbf{x}}_t^\top) \geq c$. In order to prove this, we first use the following matrix Bernstein inequality under β -mixing conditions to prove the concentration between $\Sigma_k := \frac{1}{|I_k|} \sum_{t \in I_k} \tilde{\mathbf{x}}_t \tilde{\mathbf{x}}_t^\top$ and $\Sigma := \mathbb{E}[\tilde{\mathbf{x}}_t \tilde{\mathbf{x}}_t^\top]$. Similar to §C.1, here for notational convenience, we also denote $n = |I_k|$ for any $k \geq 1$ respectively.

Lemma 9 (Matrix Bernstein Inequality under Mixing). *We assume $\tilde{\mathbf{x}}_t, t \geq 0$ satisfy Assumption 6, and we also assume there exists a positive constant M_x such that $\|\tilde{\mathbf{x}}_t\|_2 \leq M_x$. Then for any x and integer $n \geq 2$ we have*

$$\mathbb{P}\left(\left\|\sum_{t \in I_k} \tilde{\mathbf{x}}_t \tilde{\mathbf{x}}_t^\top - n\Sigma\right\| \geq nx\right) \leq 2(d+1) \exp\left(-\frac{C_u n^2 x^2}{v^2 n + M_x^4 + nx M_x^2 \log n}\right) \quad (48)$$

where C is a universal constant and

$$v^2 = \sup_{K \in \{1, \dots, n\}} \frac{1}{\text{Card}(K)} \lambda_{\max} \left\{ \mathbb{E} \left[\sum_{i \in K} (\tilde{\mathbf{x}}_i \tilde{\mathbf{x}}_i^\top - \Sigma) \right]^2 \right\}$$

and v^2 is at the order of M_x^4 .

Proof. (48) is a direct consequence of Theorem 1 in Banna et al. [2016], so here we just need to prove the order of v^2 .

$$\begin{aligned} \lambda_{\max} \left\{ \mathbb{E} \left[\sum_{i \in K} (\tilde{\mathbf{x}}_i \tilde{\mathbf{x}}_i^\top - \Sigma) \right]^2 \right\} &= \lambda_{\max} \left\{ \sum_{i, j \in K} \text{Cov}(\tilde{\mathbf{x}}_i \tilde{\mathbf{x}}_i^\top, \tilde{\mathbf{x}}_j \tilde{\mathbf{x}}_j^\top) \right\} \\ &= \lambda_{\max} \left\{ \sum_{i \in K} \text{Var}(\tilde{\mathbf{x}}_i \tilde{\mathbf{x}}_i^\top) + 2 \sum_{j > i, i, j \in K} \text{Cov}(\tilde{\mathbf{x}}_i \tilde{\mathbf{x}}_i^\top, \tilde{\mathbf{x}}_j \tilde{\mathbf{x}}_j^\top) \right\} \end{aligned}$$

Then we get

$$v^2 \leq \max_{i \in K} \lambda_{\max} \left\{ \text{Var}(\tilde{\mathbf{x}}_i \tilde{\mathbf{x}}_i^\top) + 2 \sum_{j > i, i, j \in K} \text{Cov}(\tilde{\mathbf{x}}_i \tilde{\mathbf{x}}_i^\top, \tilde{\mathbf{x}}_j \tilde{\mathbf{x}}_j^\top) \right\}$$

We know $\|\tilde{\mathbf{x}}_i\|_2 \leq M_x$, so we have

$$\lambda_{\max} \{ \text{Var}(\tilde{\mathbf{x}}_i \tilde{\mathbf{x}}_i^\top) \} \leq \|\mathbb{E}[\tilde{\mathbf{x}}_i \tilde{\mathbf{x}}_i^\top \tilde{\mathbf{x}}_i \tilde{\mathbf{x}}_i^\top]\| \leq M_x^4$$

In addition, we obtain

$$\|\text{Cov}(\tilde{\mathbf{x}}_i \tilde{\mathbf{x}}_i^\top, \tilde{\mathbf{x}}_j \tilde{\mathbf{x}}_j^\top)\| = \|\mathbb{E}[\tilde{\mathbf{x}}_i \tilde{\mathbf{x}}_i^\top \tilde{\mathbf{x}}_j \tilde{\mathbf{x}}_j^\top] - \mathbb{E}[\tilde{\mathbf{x}}_i \tilde{\mathbf{x}}_i^\top] \mathbb{E}[\tilde{\mathbf{x}}_j \tilde{\mathbf{x}}_j^\top]\| \quad (49)$$

By Lemma 1.1 (Berbee's Lemma) given in Bosq, we are able to construct a $\tilde{\mathbf{x}}_j^*$ such that the distribution of $\tilde{\mathbf{x}}_j^*$ is the same with $\tilde{\mathbf{x}}_j$ but is independent with $\tilde{\mathbf{x}}_i$. At the same time, we also have $\mathbb{P}(\tilde{\mathbf{x}}_j^* \neq \tilde{\mathbf{x}}_j) = \beta_{j-i}$ according to Berbee's Lemma. We then proceed to bound (49).

$$\begin{aligned} (49) &= \|\mathbb{E}[\tilde{\mathbf{x}}_i \tilde{\mathbf{x}}_i^\top \tilde{\mathbf{x}}_j \tilde{\mathbf{x}}_j^\top] - \mathbb{E}[\tilde{\mathbf{x}}_i \tilde{\mathbf{x}}_i^\top] \mathbb{E}[\tilde{\mathbf{x}}_j^* \tilde{\mathbf{x}}_j^{*\top}]\| \\ &= \|\mathbb{E}[\tilde{\mathbf{x}}_i \tilde{\mathbf{x}}_i^\top (\tilde{\mathbf{x}}_j \tilde{\mathbf{x}}_j^\top - \tilde{\mathbf{x}}_j^* \tilde{\mathbf{x}}_j^{*\top})]\| \\ &\leq \|\mathbb{E}[\tilde{\mathbf{x}}_i \tilde{\mathbf{x}}_i^\top (\tilde{\mathbf{x}}_j \tilde{\mathbf{x}}_j^\top - \tilde{\mathbf{x}}_j^* \tilde{\mathbf{x}}_j^{*\top}) \mid \tilde{\mathbf{x}}_j \neq \tilde{\mathbf{x}}_j^*]\| \beta_{j-i} \leq M_x^4 \beta_{j-i} \end{aligned}$$

Then we obtain that there exists a constant $C_v \geq 1 + \sum_{j > i} \beta_{j-i}$ s.t.

$$v^2 \leq C_v M_x^4,$$

holds, since the term $1 + \sum_{j > i} \beta_{j-i}$ is finite by our Assumption 6 on β_j , $j \geq 0$. Then we conclude our proof of Lemma 9 $\square$

By using conclusions from this Lemma 9, according to Assumption 2 we have $\lambda_{\min}(\Sigma) = c_{\min}$ and $\|\tilde{\mathbf{x}}_t\|_2 \leq M_x := \sqrt{R_{\mathcal{X}}^2 + 1}$, so when

$$n \geq \max\{(12C_v(R_{\mathcal{X}}^2 + 1)^2 \log n + 6(R_{\mathcal{X}}^2 + 1) \log^2 n)/(C_u \min\{c_{\min}^2/4, 1\}), d + 1\},$$

$$\lambda_{\min}(\Sigma_k) \geq c_{\min}/2. \quad (50)$$

holds with probability $1 - 2/n^2$.

Step II: The next step is to prove the upper bound of $\|\nabla_{\theta} L_k(\theta_0)\|_{\infty}$. By definition we know

$$\nabla_{\theta} L_k(\theta_0) = \frac{1}{n} \sum_{t \in I_k} 2(\theta_0^{\top} \tilde{\mathbf{x}}_t - B y_t) \tilde{\mathbf{x}}_t.$$

Since the expression of $\nabla_{\theta} L_k(\theta_0)$ involves both $\tilde{\mathbf{x}}_t$ and y_t , $t \in [n]$, next we show the sequence $(\tilde{\mathbf{x}}_t, y_t), t \geq 0$ satisfy α -mixing condition with $\alpha_k \leq \exp(-ck)$ under Assumption 6.

Lemma 10 (strong α -mixing of both $\tilde{\mathbf{x}}$ and y). *Here we denote $\mathcal{A} = \sigma((\tilde{\mathbf{x}}_t, y_t)_{t \leq l})$ and $\mathcal{B} = \sigma((\tilde{\mathbf{x}}_t, y_t)_{t \leq l+k})$. In addition, we also denote $\mathcal{A}_x = \sigma(\tilde{\mathbf{x}}_t, t \leq l)$ and $\mathcal{B}_x = \sigma(\tilde{\mathbf{x}}_t, t \geq l+k)$. Then under Assumption 6, we have for any $l, k \geq 0$,*

$$\sup_{l \geq 0} \sup_{A \in \mathcal{A}, B \in \mathcal{B}} |\mathbb{P}(A, B) - \mathbb{P}(A) \cdot \mathbb{P}(B)| \leq \alpha_k$$

where the definition of α_k is given in Definition 1.

Proof.

$$\begin{aligned} \sup_{l \geq 0} \sup_{A \in \mathcal{A}, B \in \mathcal{B}} |\mathbb{P}(A, B) - \mathbb{P}(A) \cdot \mathbb{P}(B)| &= \sup_{l \geq 0} \sup_{A \in \mathcal{A}, B \in \mathcal{B}} |\mathbb{E}[\mathbb{I}_{A,B}] - \mathbb{E}[\mathbb{I}_A] \mathbb{E}[\mathbb{I}_B]| \\ &= \sup_{l \geq 0} \sup_{A \in \mathcal{A}, B \in \mathcal{B}} |\mathbb{E}[\mathbb{E}[\mathbb{I}_{A,B} | \mathcal{A}_x, \mathcal{B}_x]] - \mathbb{E}[\mathbb{E}[\mathbb{I}_A | \mathcal{A}_x]] \mathbb{E}[\mathbb{E}[\mathbb{I}_B | \mathcal{B}_x]]| \end{aligned}$$

After conditioning on $\tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_j$, we observe that y_i, y_j are independent with each other, then we get $\mathbb{E}[\mathbb{I}_{A,B} | \mathcal{A}_x, \mathcal{B}_x] = \mathbb{E}[\mathbb{I}_A | \mathcal{A}_x] \cdot \mathbb{E}[\mathbb{I}_B | \mathcal{B}_x]$. Thus, we have for any $k \geq 0$,

$$\begin{aligned} \sup_{l \geq 0} \sup_{A \in \mathcal{A}, B \in \mathcal{B}} |\mathbb{E}[\mathbb{I}_{A,B}] - \mathbb{E}[\mathbb{I}_A] \mathbb{E}[\mathbb{I}_B]| &= \sup_{l \geq 0} \sup_{A \in \mathcal{A}, B \in \mathcal{B}} |\mathbb{E}[\mathbb{E}[\mathbb{I}_A | \mathcal{A}_x] \cdot \mathbb{E}[\mathbb{I}_B | \mathcal{B}_x]] - \mathbb{E}[\mathbb{E}[\mathbb{I}_A | \mathcal{A}_x]] \mathbb{E}[\mathbb{E}[\mathbb{I}_B | \mathcal{B}_x]]| \\ &\leq \alpha_k \|\mathbb{I}_A\|_{\infty} \cdot \|\mathbb{I}_B\|_{\infty} = \alpha_k \end{aligned}$$

The last inequality follows directly from Corollary 1.1 in Bosq, since $\mathbb{E}[\mathbb{I}_A | \mathcal{A}_x]$ lies in $\mathcal{A}_x$ and $\mathbb{E}[\mathbb{I}_B | \mathcal{B}_x]$ lies in $\mathcal{B}_x$. $\square$

By using the same proof given in §C.1, we have $\mathbb{E}[\nabla_{\theta} L_k(\theta_0)] = 0$. In addition, we obtain an upper bound of every entry of $\nabla_{\theta} L_k(\theta_0)$ in a way that there exists a upper bound $W_x = 2R_{\mathcal{X}}(R_{\mathcal{X}}R_{\Theta} + B)$ of $|2(\theta_0^{\top} \tilde{\mathbf{x}}_t - B y_t) \tilde{\mathbf{x}}_{t,i}|$, for every $i \in [d]$. Then using the following vector Bernstein inequality under α -mixing conditions, we obtain an upper bound for $\|\nabla_{\theta} L_k(\theta_0)\|_{\infty}$.

Lemma 11. (Vector Bernstein under α -Mixing Conditions, Theorem 1 in [Merlevède et al. \[2009\]](#)) Let $X_j, j \geq 0$ be a sequence of centered real-valued random variables. Suppose there exists a positive W_x such that $\sup_i \|X_i\|_\infty \leq W_x$, then when $n \geq 4$ and $x \geq 0$, we obtain

$$\mathbb{P}\left(\left|\frac{1}{n} \sum_{i=1}^n X_i\right| \geq x\right) \leq \exp\left(-\frac{C_w n^2 x^2}{nW_x^2 + W_x n x \log n \log \log n}\right)$$

where C_w is a universal constant.

By leveraging conclusions from Lemma 11, we have

$$\mathbb{P}(\|\nabla_{\theta} L_k(\theta_0)\|_\infty \geq x) \leq 2(d+1) \exp\left(-\frac{C_w n^2 x^2}{nW_x^2 + W_x n x \log n \log \log n}\right).$$

Thus, when $n \geq \max\{(6W_x^2 \log n + 6W_x \log^2 n \log \log n)/C_w, d+1\}$ we obtain, with probability $1 - 2/n^2$, we have

$$\|\nabla_{\theta} L_k(\theta_0)\|_\infty \leq \sqrt{(6W_x^2 \log n + 6W_x \log^2 n \log \log n)/(C_w n)}. \quad (51)$$

Then combining our results given in (7), (50) and (51), with probability $1 - 4/|I_k|^2$ we obtain

$$\|\hat{\theta}_k - \theta_0\|_2 \leq \frac{2}{c_{\min}} \sqrt{\frac{(d+1)(6W_x^2 \log |I_k| + 6W_x \log^2 |I_k| \log \log |I_k|)}{C_w |I_k|}}$$

for any $k \geq 1$.

F.7 Proof of Lemma 7

Proof. Similar with our proof given in §C.2, we suppose $\{w_t(\theta) := p_t - \tilde{\mathbf{x}}_t^\top \theta, y_t\}_{t \in [n]}$ are observations from the stationary distribution $P_{w(\theta), y}$. We assume that the marginal distribution $P_{w(\theta)}$ has density $f_\theta(u)$ and let $r_\theta(u) = \mathbb{E}[y_t | w_t(\theta) = u]$ be the regression function to be estimated by estimator

$$\hat{r}_k(u, \theta) = \frac{h_k(u, \theta)}{f_k(u, \theta)},$$

where

$$h_k(u, \theta) = \frac{1}{nb_k} \sum_{t \in I_k} K\left(\frac{w_t(\theta) - u}{b_k}\right) Y_t, \quad f_k(u, \theta) = \frac{1}{nb_k} \sum_{t \in I_k} K\left(\frac{w_t(\theta) - u}{b_k}\right).$$

Here, $b_k > 0$ is the bandwidth (to be chosen) in episode k , $|I_k|$ is denoted as n for simplicity and $K(\cdot)$ is some kernel function. For the true signal θ_0 , we denote the true regression function as $r_{\theta_0}(u) = \mathbb{E}[y_t | w_t(\theta_0) = u]$. The following proof procedures are similar with that given in §C.2, where their major differences are related to control the biases of $|\mathbb{E}[h_k(u, \theta)] - h_\theta(u)|$ and $|\mathbb{E}[f_k(u, \theta)] - f_\theta(u)|$ given in Lemma 12 and the variances of $h_k(u, \theta)$ and $f_k(u, \theta)$ given in Lemma 13 under strong-mixing settings respectively.

Lemma 12. Under Assumptions 3-5 and 6, with any choice of $b_k \leq 1$, we obtain

$$\sup_{u \in I, \theta \in \Theta_k} |\mathbb{E}h_k(u, \theta) - h_\theta(u)| \leq C_{mx,K}^{(1)} b_k^m$$

$$\sup_{u \in I, \theta \in \Theta_k} |\mathbb{E}f_k(u, \theta) - f_\theta(u)| \leq C_{mx,K}^{(1)} b_k^m$$

where $C_{mx,K} = l_f \frac{\int |s^m K(s)| ds}{(m-1)!}$.

Proof. The proof of Lemma 12 is the same with the proof of Lemma 1. So we omit the details. $\square$

Lemma 13. Under Assumption 3-5 and 6, there exists a constant C'_{17} only depending on constants given in assumptions, such that for $I = [-\delta_z, \delta_z]$, if $b_k \in [1/n, 1]$, $nb_k \geq 4C_{17}'^2 \log^3 n [(d+1) \log(d+1)]$ and $\delta \in [8 \exp(-nb_k/(8C_{17}'^2 \log^2 n)), 1/2]$, the following inequalities hold simultaneously with probability $1 - \delta$:

$$\sup_{u \in I, \theta \in \Theta_k} |h_k(u, \theta) - \mathbb{E}[h_k(u, \theta)]| \leq \frac{C'_{17} \log n}{\sqrt{nb_k}} \left(\sqrt{(d+1) \log(d+1) \log n} + \sqrt{2 \log \frac{8}{\delta}} \right) \quad (52)$$

$$\sup_{u \in I, \theta \in \Theta_k} |f_k(u, \theta) - \mathbb{E}[f_k(u, \theta)]| \leq \frac{C'_{17} \log n}{\sqrt{nb_k}} \left(\sqrt{(d+1) \log(d+1) \log n} + \sqrt{2 \log \frac{8}{\delta}} \right) \quad (53)$$

Proof. We only prove (52), since (53) can be proved in the same way. For any $u \in I$ and $\theta \in \Theta_k$, we denote $Z(u, \theta) := h_k(u, \theta) - \mathbb{E}h_k(u, \theta) = \frac{1}{nb_k} \sum_{t \in I_k} [K(\frac{w_t(\theta) - u}{b_k}) y_t - \mathbb{E}K(\frac{w_t(\theta) - u}{b_k}) y_t]$. Then we have that

$$\sup_{u \in I, \theta \in \Theta_k} |h_k(u, \theta) - \mathbb{E}h_k(u, \theta)| = \sup_{u \in I, \theta \in \Theta_k} |Z(u, \theta)| = \max \left\{ \sup_{u \in I, \theta \in \Theta_k} Z(u, \theta), \sup_{u \in I, \theta \in \Theta_k} (-Z(u, \theta)) \right\}.$$

Similar with our proof procedure of Lemma 2, we then bound $\sup_{u \in I, \theta \in \Theta_k} |h_k(u, \theta) - \mathbb{E}h_k(u, \theta)|$ by upper bounding both $\sup_{u \in I, \theta \in \Theta_k} Z(u, \theta)$ and $\sup_{u \in I, \theta \in \Theta_k} (-Z(u, \theta))$. We next also use chaining method to achieve desired bound. We also construct a sequence of ϵ -nets with decreasing scale.

As a reminder, here we also denote the left and right endpoints of the interval I as L_I and R_I respectively. For any $i \in \mathbb{N}^+$, construct set $S_1^{(i)} \subseteq I$ as

$$S_1^{(i)} \triangleq \left\{ L_I + \frac{j}{2^i \sqrt{n}} (R_I - L_I) : j \in \{1, 2, \dots, (2^i - 1) \lfloor \sqrt{n} \rfloor\} \right\}.$$

For any $u \in I$, $i \in \mathbb{N}^+$, let $\pi_i(u) = \arg \min_{s \in S_1^{(i)}} |s - u|$. Moreover, let $\pi_0(u) = u$. Then we can easily verify that $|S_1^{(i)}| \leq 2^i (\sqrt{n} + 1)$, and that $\forall u \in I$, $|\pi_i(u) - \pi_{i+1}(u)| \leq \frac{2\delta_z}{2^{i-1} \sqrt{n}}$.

As for the ϵ -net of Θ_k , we let S_2^i be a $R_m/(2^i \sqrt{n})$ -net with respect to l_2 -distance of Θ_k , where $R_m = 2/c_{\min} \sqrt{6W_x/C_w}$ (constants are specified in the Lemma 6). By Proposition 4.2.12 in Vershynin [2018], we have $|S_2^{(i)}| \leq (2^{i+1} C(d, n) + 1)^d$, where $C(d, n) = \sqrt{(d+1)(W_x \log n + \log^2 n \log \log n)}$.

Then we have for any $\mathbf{u} := (u, \boldsymbol{\theta}) \in I \times \Theta_k$ with $i \geq 1$, there exist $\pi_i(u) \in S_1^{(i)}$ and $\pi_i(\boldsymbol{\theta}) \in S_2^{(i)}$ such that $\|\pi_i(\mathbf{u}) := (\pi_i(u), \pi_i(\boldsymbol{\theta})) - \mathbf{u}\|_2 \leq \sqrt{4\delta_z^2 + R_m^2}/(2^i \sqrt{n})$. So $S^{(i)} := S_1^{(i)} \times S_2^{(i)}$ is a $\sqrt{4\delta_z^2 + R_m^2}/(2^i \sqrt{n})$ -net of $U_k := I \times \Theta_k$ with size $|S^{(i)}| \leq 2^i(\sqrt{n} + 1) \cdot (2^{i+1}C(d, n) + 1)^d$ and $C(d, n) = \sqrt{(d+1)(W_x \log n + \log^2 n \log \log n)}$.

Because $Z(u, \boldsymbol{\theta})$ is continuous almost surely, we have that for any $M \in \mathbb{N}^+$

$$Z(\mathbf{u}) - Z(\pi_M(\mathbf{u})) = \sum_{i=M}^{\infty} [Z(\pi_{i+1}(\mathbf{u})) - Z(\pi_i(\mathbf{u}))],$$

and thus

$$\sup_{\mathbf{u} \in U_k} Z(\mathbf{u}) \leq \sup_{\mathbf{u} \in U_k} Z(\pi_M(\mathbf{u})) + \sum_{i=M}^{\infty} \sup_{\mathbf{u} \in U_k} [Z(\pi_{i+1}(\mathbf{u})) - Z(\pi_i(\mathbf{u}))] \quad (54)$$

almost surely. If we can choose a M properly then the two terms at the right hand side of (54) can be both well controlled. For this reason, we let $M = \lceil \frac{4}{\log 2} \log \frac{1}{b_n} \rceil$. We then first bound $\sup_{\mathbf{u} \in U_k} Z(\pi_M(\mathbf{u}))$ by using a union bound. By our definition on $Z(\mathbf{u})$, we can write

$$Z(\mathbf{u}) = \frac{1}{nb_k} \sum_{t \in I_k} A_t(\mathbf{u}).$$

in which $A_t(\mathbf{u}) = K(\frac{w_t(\boldsymbol{\theta}) - u}{b_k})y_t - \mathbb{E}K(\frac{w_t(\boldsymbol{\theta}) - u}{b_k})y_t$. Similar with our case in proving Lemma 2, we have that $\mathbb{E}[A_t(\mathbf{u})] = 0$ and $|A_t(\mathbf{u})| \leq \bar{K}$ almost surely. We next prove the bound of variance of $A_t(\mathbf{u})$ and the covariance between $A_j(\mathbf{u})$ and $A_i(\mathbf{u})$ with $j > i$. Following similar procedures with Lemma 2, we first conclude that

$$\text{Var}(A_t(\mathbf{u})) \leq C'_4 b_k,$$

where $C'_4 = C_4 = \max\{\bar{f} \cdot \int K(s)^2 ds, \bar{f} \cdot \int K'(s)^2 ds\}$ is defined in the same way with our proof of Lemma 2. We next control the covariance of $A_j(\mathbf{u})$ and $A_i(\mathbf{u})$ with $j > i$.

$$\begin{aligned} \text{Cov}(A_j(\mathbf{u}), A_i(\mathbf{u})) &= \mathbb{E}\left[K\left(\frac{w_j(\boldsymbol{\theta}) - u}{b_k}\right)y_j K\left(\frac{w_i(\boldsymbol{\theta}) - u}{b_k}\right)y_i\right] - \mathbb{E}\left[K\left(\frac{w_j(\boldsymbol{\theta}) - u}{b_k}\right)y_j\right] \mathbb{E}\left[K\left(\frac{w_i(\boldsymbol{\theta}) - u}{b_k}\right)y_i\right] \\ &= \mathbb{E}\left[K\left(\frac{w_j(\boldsymbol{\theta}) - u}{b_k}\right)K\left(\frac{w_i(\boldsymbol{\theta}) - u}{b_k}\right)\mathbb{E}[y_j y_i | w_j(\boldsymbol{\theta}), w_i(\boldsymbol{\theta})]\right] \\ &\quad - \mathbb{E}\left[K\left(\frac{w_j(\boldsymbol{\theta}) - u}{b_k}\right)y_j\right] \mathbb{E}\left[K\left(\frac{w_i(\boldsymbol{\theta}) - u}{b_k}\right)y_i\right] \end{aligned}$$

For simplicity, for any $\boldsymbol{\theta} \in \Theta_0$, we define $r(u_i, u_j) := \mathbb{E}[y_i y_j | w_j(\boldsymbol{\theta}) = u_j, w_i(\boldsymbol{\theta}) = u_i]$ and $r(u_j) = \mathbb{E}[y_j | w_j(\boldsymbol{\theta}) = u_j]$. Then after some simple calculation, we further obtain

$$\text{Cov}(A_j(\mathbf{u}), A_i(\mathbf{u})) = \int \int K\left(\frac{w_j(\boldsymbol{\theta}) - u}{b_k}\right)K\left(\frac{w_i(\boldsymbol{\theta}) - u}{b_k}\right)r(w_i(\boldsymbol{\theta}), w_j(\boldsymbol{\theta}))f(w_i(\boldsymbol{\theta}), w_j(\boldsymbol{\theta}))dw_i(\boldsymbol{\theta})dw_j(\boldsymbol{\theta})$$

$$\begin{aligned}
& - \int \int K\left(\frac{w_j(\boldsymbol{\theta}) - u}{b_k}\right) K\left(\frac{w_i(\boldsymbol{\theta}) - u}{b_k}\right) r(w_i(\boldsymbol{\theta})) r(w_j(\boldsymbol{\theta})) f(w_i(\boldsymbol{\theta})) f(w_j(\boldsymbol{\theta})) dw_i(\boldsymbol{\theta}) dw_j(\boldsymbol{\theta}) \\
& = b_k^2 \int \int K(s_1) K(s_2) [r(b_k s_1 + u, b_k s_2 + u) f(b_k s_1 + u, b_k s_2 + u) \\
& \quad - r(b_k s_1 + u) r(b_k s_2 + u) f(b_k s_1 + u) f(b_k s_2 + u)] ds_1 ds_2
\end{aligned}$$

We next prove that $h(u_i, u_i) := r(u_i, u_j) f(u_i, u_j)$ stays close to $h(u_i) h(u_j) := r(u_i) f(u_i) r(u_j) f(u_j)$ for all (u_i, u_j) in the following Lemma 14.

Lemma 14. *Under Assumptions given in Lemma 13. We let $g^*(u_i, u_j) := h(u_i, u_j) - h(u_i) h(u_j)$, if we further assume $g^*(u_i, u_j)$ is Lipschitz continuous w.r.t. (u_i, u_j) with Lipschitz constant l , then we have*

$$\sup_{u_i, u_j} |g^*(u_i, u_j)| \leq (1/4 + \sqrt{2}l) \beta_{j-i}^{1/3}$$

Proof. For any x we define

$$B(x, \epsilon) := \{x' : \|x' - x\| \leq \epsilon\}, \quad \epsilon > 0, x \in \mathbb{R}$$

First, we prove $|\mathbb{E}[y_i y_j \mathbb{I}_{\{w_i(\boldsymbol{\theta}) \in B(x, \epsilon), w_j(\boldsymbol{\theta}) \in B(y, \epsilon)\}}] - \mathbb{E}[y_i \mathbb{I}_{\{w_i(\boldsymbol{\theta}) \in B(x, \epsilon)\}}] \mathbb{E}[y_j \mathbb{I}_{\{w_j(\boldsymbol{\theta}) \in B(y, \epsilon)\}}]| \leq \beta_{j-i}$. We have

$$\begin{aligned}
& |\mathbb{E}[y_i y_j \mathbb{I}_{\{w_i(\boldsymbol{\theta}) \in B(x, \epsilon), w_j(\boldsymbol{\theta}) \in B(y, \epsilon)\}}] - \mathbb{E}[y_i \mathbb{I}_{\{w_i(\boldsymbol{\theta}) \in B(x, \epsilon)\}}] \mathbb{E}[y_j \mathbb{I}_{\{w_j(\boldsymbol{\theta}) \in B(y, \epsilon)\}}]| \\
& = |\mathbb{E}[\mathbb{I}_{\{w_i(\boldsymbol{\theta}) \in B(x, \epsilon), w_j(\boldsymbol{\theta}) \in B(y, \epsilon)\}}] \mathbb{E}[y_i y_j | \tilde{\mathbf{x}}_i, \tilde{\mathbf{x}}_j, p_i, p_j] \\
& \quad - \mathbb{E}[\mathbb{I}_{\{w_i(\boldsymbol{\theta}) \in B(x, \epsilon)\}}] \mathbb{E}[y_i | \tilde{\mathbf{x}}_i, p_i] \mathbb{E}[\mathbb{I}_{\{w_j(\boldsymbol{\theta}) \in B(y, \epsilon)\}}] \mathbb{E}[y_j | \tilde{\mathbf{x}}_j, p_j]| \\
& = |\mathbb{E}[\mathbb{E}[y_i \mathbb{I}_{\{w_i(\boldsymbol{\theta}) \in B(x, \epsilon)\}} | \tilde{\mathbf{x}}_i, p_i] \mathbb{E}[y_j \mathbb{I}_{\{w_j(\boldsymbol{\theta}) \in B(y, \epsilon)\}} | \tilde{\mathbf{x}}_j, p_j] \\
& \quad - \mathbb{E}[\mathbb{E}[y_i \mathbb{I}_{\{w_i(\boldsymbol{\theta}) \in B(x, \epsilon)\}} | \tilde{\mathbf{x}}_i, p_i] \mathbb{E}[\mathbb{E}[y_j \mathbb{I}_{\{w_j(\boldsymbol{\theta}) \in B(y, \epsilon)\}} | \tilde{\mathbf{x}}_j, p_j]]|
\end{aligned}$$

As $p_i, i \in |I_k|, k \geq 0$ are independent, so the σ -algebra generated by the joint distribution of $\tilde{\mathbf{x}}_i, p_i$ still follows strong- β and $-\alpha$ conditions given in our Assumption 6. Moreover, we have $\mathbb{E}[y_i \mathbb{I}_{\{w_i(\boldsymbol{\theta}) \in B(x, \epsilon)\}} | \tilde{\mathbf{x}}_i, p_i]$ lies in $\sigma(\tilde{\mathbf{x}}_i, p_i)$ and $\mathbb{E}[y_j \mathbb{I}_{\{w_j(\boldsymbol{\theta}) \in B(y, \epsilon)\}} | \tilde{\mathbf{x}}_j, p_j]$ lies in $\sigma(\tilde{\mathbf{x}}_j, p_j)$ with $j > i$. So we are able to obtain the upper bound:

$$|\mathbb{E}[y_i y_j \mathbb{I}_{\{w_i(\boldsymbol{\theta}) \in B(x, \epsilon), w_j(\boldsymbol{\theta}) \in B(y, \epsilon)\}}] - \mathbb{E}[y_i \mathbb{I}_{\{w_i(\boldsymbol{\theta}) \in B(x, \epsilon)\}}] \mathbb{E}[y_j \mathbb{I}_{\{w_j(\boldsymbol{\theta}) \in B(y, \epsilon)\}}]| \leq \beta_{j-i} \quad (55)$$

by using Corollary 1.1 in Bosq.

Next, we get an upper bound of $\sup_{(u_i, u_j)} |g^*(u_i, u_j)|$. From (55) and our definition on g^* , we obtain

$$\beta_{j-i} \geq \left| \int_{B(x, \epsilon) \times B(y, \epsilon)} g^*(u_i, u_j) du_i du_j \right| := \mathcal{I}$$

Then by the mean value property we have $\mathcal{I} = 4\epsilon^2|g^*(x', y')|$ for some $(x', y') \in B(x, \epsilon) \times B(y, \epsilon)$. Moreover, as we assume g is Lipschitz, then we get

$$|g^*(x, y)| \leq |g^*(x', y')| + \sqrt{2}l\epsilon$$

Hence, we finally achieve

$$|g^*(x, y)| \leq \beta_{j-i}/(4\epsilon^2) + \sqrt{2}l\epsilon.$$

for any fixed (x, y) . As this inequality holds for all $\epsilon > 0$, we choose $\epsilon = \beta_{j-i}^{1/3}$ and we conclude the proof of our Lemma 14. $\square$

By our conclusion from Lemma 14, we are able to find a constant C'_5 such that $|\sum_{j>i} \text{Cov}(A_j(\mathbf{u}), A_i(\mathbf{u}))| \leq C'_5 b_n$ holds according to our assumptions on β_{j-i} , $j > i$, where we set $C'_5 = (1/4 + \sqrt{2}l) \sum_{j>0} \beta_j^{1/3}$. Next we introduce the following Bernstein inequality under strong-mixing conditions, in order to achieve an upper bound of $Z(\mathbf{u})$.

Lemma 15. [Theorem 2 in Merlevède et al. [2009]] Under conditions of Lemma 13, for all $n \geq 2$, we have

$$\mathbb{P}(|Z(\mathbf{u})| \geq nb_k x) = \mathbb{P}(|\sum_{j \in I_k} A_j(\mathbf{u})| \geq nb_k x) \leq 2 \exp\left(-\frac{C_b b_k^2 n^2 x^2}{v^2 n + \bar{K}^2 + nb_k x \log^2 n}\right)$$

Here

$$v^2 = \sup_{i>0} (\text{Var}(A_i(\mathbf{u})) + 2 \sum_{j>i} |\text{Cov}(A_i(\mathbf{u}), A_j(\mathbf{u}))|),$$

C_b is a pure constant and $\bar{K}$ is defined as the upper bound of $|A_j(\mathbf{u})|$ with any $j \in [n]$.

By our conclusions from Lemma 14 and Lemma 15, we conclude there exists a constant $C'_6 = (C'_4 + 2C'_5)$ such that $v^2 \leq C'_6 b_n$, so we obtain

$$\begin{aligned} \mathbb{P}(|Z(\mathbf{u})| \geq x) &\leq 2 \exp\left(-\frac{C_b b_k^2 n^2 x^2}{C'_6 n b_k + \bar{K}^2 + n b_k x \log^2 n}\right) \\ &\leq 2 \exp\left(-\frac{C_b n b_k x^2}{(C'_6 + \bar{K}^2 + \log^2 n)(1+x)}\right) \end{aligned}$$

The last inequality follows from our assumption that $b_k \geq 1/n = 1/|I_k|$ for any $k \geq 1$ in given Lemma 13. Further, we set $C'_7 = C_b/(2C'_6 + 2\bar{K}^2 + 2)$. Then we take the union bound over U_k , which gives

$$\begin{aligned} \mathbb{P}(\sup_{\mathbf{u} \in U_k} |Z(\pi_M(\mathbf{u}))| \geq x) &\leq |S^{(M)}| \cdot \mathbb{P}(|Z(\mathbf{u})| \geq x) \\ &\leq 2 \cdot 2^M (\sqrt{n} + 1) \cdot (2^{M+1} \sqrt{d} + 1)^d \cdot e^{-\frac{C'_7 n b_k}{\log^2 n} \min\{x, x^2\}} \end{aligned}$$

$$\leq 2e^{(d+1)M \log 2 + \log(\sqrt{n}+1) + d \log(2C(n,d)+2) - \frac{C'_7 n b_k}{\log^2 n} \min\{x, x^2\}}.$$

Since we define $M = \lceil \frac{4}{\log 2} \log \frac{1}{b_k} \rceil$, then we choose

$$x(n, d) := \frac{\log n}{\sqrt{n b_k}} \sqrt{\left[(d+1)4 \log \frac{1}{b_k} + 2(d+1) \log 2 + \log(\sqrt{n}+1) + d \log(2C(n, d) + 2) + \log \frac{8}{\delta} \right] / C'_7}, \quad (56)$$

where $C(n, d) = \sqrt{(d+1)(W_x \log n + \log^2 n \log \log n)}$. We then have

$$\mathbb{P}(\sup_{\mathbf{u} \in U_k} |Z(\pi_M(\mathbf{u}))| \geq x(n, d)) \leq \frac{\delta}{4}.$$

when $\delta > 8 \exp(-n b_k / (C'_7 \log^2 n))$ and $n b_k \geq 2 \log^2 n [(d+1)4 \log \frac{1}{b_k} + 2(d+1) \log 2 + \log(\sqrt{n}+1) + d \log(2C(n, d) + 2)] / C'_7$ (because under such conditions, we have $x(n, d) \leq 1$). Now, we proceed to bound the later term at the right hand side of (54). Similar with our cases stated in the proof of Lemma 2, for any $\mathbf{u}_1 := (u, \boldsymbol{\theta}_1)$, $\mathbf{u}_2 := (s, \boldsymbol{\theta}_2) \in I \times \Theta_k$, we have that

$$Z(\mathbf{u}_1) - Z(\mathbf{u}_2) = Z(u, \boldsymbol{\theta}_1) - Z(s, \boldsymbol{\theta}_2) = \frac{1}{n b_k} \sum_{t \in I_k} B_t(u, \boldsymbol{\theta}_1, s, \boldsymbol{\theta}_2),$$

where

$$B_t(u, \boldsymbol{\theta}_1, s, \boldsymbol{\theta}_2) = y_t \left(K\left(\frac{w_t(\boldsymbol{\theta}_1) - u}{b_k}\right) - K\left(\frac{w_t(\boldsymbol{\theta}_2) - s}{b_k}\right) \right) - \mathbb{E} y_t \left(K\left(\frac{w_t(\boldsymbol{\theta}_1) - t}{b_k}\right) - K\left(\frac{w_t(\boldsymbol{\theta}_2) - s}{b_k}\right) \right).$$

We have $\mathbb{E} B_t(u, \boldsymbol{\theta}_1, s, \boldsymbol{\theta}_2) = 0$, and that

$$\begin{aligned} |Z(\mathbf{u}_1) - Z(\mathbf{u}_2)| &= |B_t(u, \boldsymbol{\theta}_1, s, \boldsymbol{\theta}_2)| \leq 2 \left| y_j \left(K\left(\frac{w_t(\boldsymbol{\theta}_1) - u}{b_k}\right) - K\left(\frac{w_t(\boldsymbol{\theta}_2) - s}{b_k}\right) \right) \right| \\ &\leq \frac{2l_K \sqrt{1 + \max_{\mathbf{x} \in \mathcal{X}} \|\mathbf{x}\|_2^2 + 1}}{b_n} \cdot \|\mathbf{u}_1 - \mathbf{u}_2\|_2 := \frac{C^*}{b_n} \|\mathbf{u}_1 - \mathbf{u}_2\|_2. \end{aligned}$$

The last inequality follows from the Lipschitz property of $K(\cdot)$ and for simplicity we use C^* to denote the constant $2l_K \sqrt{\max_{\mathbf{x} \in \mathcal{X}} \|\mathbf{x}\|_2^2 + 2} = 2l_K \sqrt{R_{\mathcal{X}}^2 + 2}$. Then according to the Bernstein inequality given in Lemma 11, we have

$$\mathbb{P}(|\sum_{t=1}^n B_t(\mathbf{u}_1, \mathbf{u}_2)| \geq n b_k x) \leq 2 \exp \left(- \frac{C_w n^2 b_k^2 x^2}{n \frac{C^{*2} \|\mathbf{u}_1 - \mathbf{u}_2\|_2^2}{b_k^2} + n b_k x \frac{C^* \|\mathbf{u}_1 - \mathbf{u}_2\|_2}{b_k} \log^2 n} \right).$$

Recall that $\forall \mathbf{u} \in U_n$, we have $\|\pi_i(\mathbf{u}) - \pi_{i+1}(\mathbf{u})\|_2 \leq \frac{\sqrt{4\delta_z^2 + R_m^2}}{2^{i-1} \sqrt{n}}$. We then use the union bound to get

$$\begin{aligned} & \mathbb{P}\left(\sup_{\mathbf{u} \in U_k} |Z(\pi_{i+1}(\mathbf{u})) - Z(\pi_i(\mathbf{u}))| \geq x\right) \\ & \leq 2^{2i+2}(\sqrt{n} + 1)^2(2^{i+2}C(n, d) + 1)^{2d} \cdot 2e^{-\frac{C'_8 2^{i-1} n^{3/2} b_k^4 x^2}{(\frac{4\delta_z^2 + R_m^2}{2^{i-1}\sqrt{n}} + b_k^2 \sqrt{4\delta_z^2 + R_m^2} \log^2 n)(1+x)}} \end{aligned}$$

in which $C'_8 = C_w / \max\{C^{*2}, C^*\}$. We let $x = \frac{\sqrt{(4\delta_z^2 + R_m^2)/(2^{i-1}\sqrt{n}) + b_k^2 \sqrt{4\delta_z^2 + R_m^2} \log^2 n}}{2^{(i-1)/2} n^{3/4} b_k^2} \cdot \epsilon_i$. Then we have

$$\mathbb{P}\left(\sup_{\mathbf{u} \in U_k} |Z(\pi_{i+1}(\mathbf{u})) - Z(\pi_i(\mathbf{u}))| \geq \sqrt{(4\delta_z^2 + R_m^2)/(2^{i-1}\sqrt{n}) + b_k^2 \sqrt{4\delta_z^2 + R_m^2} \log^2 n} / (2^{(i-1)/2} n^{3/4} b_k^2) \cdot \epsilon_i\right) \quad (57)$$

$$\leq 2^{2i+2}(\sqrt{n} + 1)^2(2^{i+2}C(n, d) + 1)^{2d} \cdot 2e^{-\frac{C'_8 \epsilon_i^2}{1 + \frac{\sqrt{(4\delta_z^2 + R_m^2)/(2^{i-1}\sqrt{n}) + b_k^2 \sqrt{4\delta_z^2 + R_m^2} \log^2 n}}{2^{(i-1)/2} n^{3/4} b_k^2} \cdot \epsilon_i}} \quad (58)$$

We observe that if we could choose ϵ_i such that

$$\frac{\sqrt{(4\delta_z^2 + R_m^2)/(2^{i-1}\sqrt{n}) + b_k^2 \sqrt{4\delta_z^2 + R_m^2} \log^2 n}}{2^{(i-1)/2} n^{3/4} b_k^2} \cdot \epsilon_i < 1,$$

holds, then the right hand side of (58) satisfies

$$(58) \leq 2^{2i+2}(\sqrt{n} + 1)^2(2^{i+2}C(n, d) + 1)^{2d} \cdot 2e^{-\frac{C'_8 \epsilon_i^2}{2}}. \quad (59)$$

Now we choose $\epsilon_i = \sqrt{[(4d + 6)(i + 1) \log 2 + 4 \log(\sqrt{n} + 1) + 4d \log(2C(n, d) + 2) + 2 \log(8/\delta)]/C'_8}$. Then we have

$$\begin{aligned} & \frac{\sqrt{(4\delta_z^2 + R_m^2)/(2^{i-1}\sqrt{n}) + b_k^2 \sqrt{4\delta_z^2 + R_m^2} \log^2 n}}{2^{(i-1)/2} n^{3/4} b_k^2} \cdot \epsilon_i \\ & \leq \frac{1}{2^{(i-1)/2} n^{3/4} b_k} \left[\frac{\sqrt{4\delta_z^2 + R_m^2}}{2^{(i-1)/2} b_k n^{1/4}} + (4\delta_z^2 + R_m^2)^{1/4} \log n \right] \cdot \epsilon_i. \end{aligned}$$

Here we only consider $i \geq M = \lceil \frac{4}{\log 2} \log \frac{1}{b_k} \rceil$, and we have $2^{M/4} \cdot b_k = 1$. In addition, we also get $\max_i(i + 1)/2^{(i-2)/2} \leq 3$. Hence, we have

$$\frac{1}{2^{(i-1)/2} n^{3/4} b_k} \left[\frac{\sqrt{4\delta_z^2 + R_m^2}}{2^{(i-1)/2} b_k n^{1/4}} + (4\delta_z^2 + R_m^2)^{1/4} \log n \right] \cdot \epsilon_i < 1,$$

if $\delta \geq 8 \exp(-C'_8 n^{3/2}/(16(4\delta_z^2 + R_m^2) \log^2 n))$ and $n \geq \{8(4\delta_z^2 + R_m^2) \log^2 n \cdot [(12d + 18) \log 2 + 4 \log(\sqrt{n} + 1) + 4d \log(2C(n, d) + 2)]/C'_8\}^{2/3}$. Then after plugging our setting of ϵ_i into (59), we obtain

$$\begin{aligned} \mathbb{P}\left(\sup_{\mathbf{u} \in U_k} |Z(\pi_{i+1}(\mathbf{u})) - Z(\pi_i(\mathbf{u}))| \geq \sqrt{(4\delta_z^2 + R_m^2)/(2^{i-1}\sqrt{n}) + b_k^2 \sqrt{4\delta_z^2 + R_m^2} \log^2 n / (2^{(i-1)/2} n^{3/4} b_k^2)} \cdot \epsilon_i\right) \\ \leq \frac{1}{2^{i+1}} \cdot \frac{\delta}{4}. \end{aligned}$$

And we notice

$$\begin{aligned} \sum_{i=M}^{\infty} \frac{\sqrt{(4\delta_z^2 + R_m^2)/(2^{i-1}\sqrt{n}) + b_k^2 \sqrt{4\delta_z^2 + R_m^2} \log^2 n}}{2^{(i-1)/2} n^{3/4} b_k^2} \cdot \epsilon_i \\ \leq \sum_{i=M}^{\infty} \frac{\sqrt{4\delta_z^2 + R_m^2}}{2^{i-1} n b_k^2} \cdot \epsilon_i + \frac{\sqrt{(4\delta_z^2 + R_m^2) \log^2 n}}{2^{(i-1)/2} n^{3/4} b_k^2} \cdot \epsilon_i := \mathbf{I} + \mathbf{II}. \end{aligned}$$

For term $\mathbf{I}$, we have

$$\begin{aligned} \mathbf{I} &= \sum_{i=M}^{\infty} \frac{\sqrt{4\delta_z^2 + R_m^2}}{2^{i-1} n b_k^2} \cdot \sqrt{[(4d + 6)(i + 1) \log 2 + 4 \log(\sqrt{n} + 1) + 4d \log(2C(n, d) + 2) + 2 \log(8/\delta)]/C'_8} \\ &\leq \frac{\sqrt{(4\delta_z^2 + R_m^2)/C'_8}}{n b_k^2} \left[\sqrt{(4d + 6) \log 2} \sum_{i=M}^{\infty} \frac{i + 1}{2^{i-1}} + \frac{\sqrt{4 \log(\sqrt{n} + 1) + 4d \log(2C(n, d) + 2) + 2 \log(8/\delta)}}{2^{M-2}} \right] \\ &\leq \frac{\sqrt{(4\delta_z^2 + R_m^2)/C'_8}}{n b_k^2} \frac{2M}{2^{M-2}} \left[\sqrt{(4d + 6) \log 2} + \sqrt{4 \log(\sqrt{n} + 1)} + \sqrt{4d \log(2C(n, d) + 2)} + \sqrt{2 \log(8/\delta)} \right] \\ &\leq \frac{\sqrt{(4\delta_z^2 + R_m^2)/C'_8}}{n} \frac{8M}{2^{M/2}} \frac{1}{2^{M/2} b_k^2} \left[\sqrt{(4d + 6) \log 2} + \sqrt{4 \log(\sqrt{n} + 1)} \right. \\ &\quad \left. + \sqrt{4d \log(2C(n, d) + 2)} + \sqrt{2 \log(8/\delta)} \right] \\ &\leq \frac{C'_9}{n} \left[\sqrt{(4d + 6) \log 2} + \sqrt{4 \log(\sqrt{n} + 1)} + \sqrt{4d \log(2C(n, d) + 2)} + \sqrt{2 \log(8/\delta)} \right], \end{aligned}$$

in which C'_9 is a pure constant such that $C'_9 = \sqrt{(4\delta_z^2 + R_m^2)/C'_8} \cdot \max_i (8i/2^{i/2}) = 16\sqrt{(4\delta_z^2 + R_m^2)/C'_8}$ and $C(n, d) \leq \sqrt{(d + 1)(W_x \log n + \log^3 n)}$. Then we obtain

$$\begin{aligned} \sqrt{4d \log(2C(n, d) + 2)} &\leq \sqrt{4d \log \left(4\sqrt{(d + 1)(W_x \log n + \log^3 n)} \right)} \\ &\leq \sqrt{4d \log \left(4\sqrt{2} \sqrt{(d + 1) \max\{1, W_x\} \log^3 n} \right)} \end{aligned}$$

$$\leq \sqrt{4d \log(4\sqrt{2})} + \sqrt{2d \log(\max\{W_x, 1\}(d+1))} + \sqrt{6d \log n}. \quad (60)$$

Next, we are able to find a pure constant $C'_{10} = 6\sqrt{6}$ such that $\sqrt{(4d+6) \log 2} + \sqrt{4 \log(\sqrt{n}+1)} + \sqrt{4d \log(2C(n, d) + 2)} \leq 6\sqrt{6} \sqrt{(d+1) \log(\max\{W_x, 1\}(d+1)) \log n}$ as long as $n \geq 3$ according to (60). Thus, we finally achieve

$$\mathbf{I} \leq \frac{C'_{11}}{n} \left(\sqrt{(d+1) \log(\max\{W_x, 1\}(d+1)) \log n} + \sqrt{2 \log(8/\delta)} \right),$$

where $C'_{11} = C'_{10} \cdot C'_9$. For term **II**, we obtain

$$\begin{aligned} \mathbf{II} &= \sum_{i=M}^{\infty} \frac{\sqrt{(4\delta_z^2 + R_m^2) \log^2 n / C'_8}}{2^{(i-1)/2} n^{3/4} b_k} [\sqrt{(4d+6)(i+1) \log 2} \\ &\quad + \sqrt{4 \log(\sqrt{n}+1)} + \sqrt{4d \log(2C(n, d) + 2)} + \sqrt{2 \log(8/\delta)}] \\ &\leq \frac{\sqrt{(4\delta_z^2 + R_m^2) / C'_8} \log n}{n^{3/4} b_k} \left[\sqrt{(4d+6) \log 2} \sum_{i=M}^{\infty} \frac{i+1}{2^{(i-1)/2}} \right. \\ &\quad \left. + \frac{\sqrt{4 \log(\sqrt{n}+1)} + 4d \log(2C(n, d) + 2) + 2 \log(8/\delta)}{2^{(M-2)/2}} \right] \\ &\leq \frac{\sqrt{(4\delta_z^2 + R_m^2) / C'_8} \log n}{n^{3/4}} \frac{8\sqrt{2}M}{2^{M/4}} \frac{1}{2^{M/4} b_k} \left[\sqrt{(4d+6) \log 2} + \sqrt{4 \log(\sqrt{n}+1)} \right. \\ &\quad \left. + \sqrt{4d \log(2C(n, d) + 2)} + \sqrt{2 \log(8/\delta)} \right]. \end{aligned}$$

We are also able to find a pure constant C'_{12} such that $C'_{12} = \sqrt{(4\delta_z^2 + R_m^2) / C'_8} \max_i (8\sqrt{2}i/2^{i/4}) = 24\sqrt{(4\delta_z^2 + R_m^2) / C'_8}$ and $C'_{13} = C'_{10} \cdot C'_{12}$. Then we obtain

$$\mathbf{II} \leq \frac{C'_{13} \log n}{n^{3/4}} \left(\sqrt{(d+1) \log(\max\{W_x, 1\}(d+1)) \log n} + \sqrt{2 \log 8/\delta} \right).$$

After combining our inequalities of **I** and **II**, we obtain a union bound:

$$\begin{aligned} \mathbb{P} \left(\sup_{\mathbf{u} \in U_k} |Z(\mathbf{u}) - Z(\pi_M(\mathbf{u}))| \geq x_2(n, d) : = \frac{C'_{14} \log n}{n^{3/4}} \left(\sqrt{(d+1) \log(\max\{W_x, 1\}(d+1)) \log n} \right. \right. \\ \left. \left. + \sqrt{2 \log(8/\delta)} \right) \right) \\ \leq \sum_{i=M}^{\infty} \frac{1}{2^{i+1}} \frac{\delta}{4} \leq \frac{\delta}{4}, \end{aligned}$$

in which we choose $C'_{14} = 2 \max\{C'_{11}, C'_{13}\}$. Then we get

$$\mathbb{P} \left(\sup_{\mathbf{u} \in U_k} Z(\mathbf{u}) \geq x(n, d) + x_2(n, d) \right) \leq \frac{\delta}{4} + \frac{\delta}{4} = \frac{\delta}{2}.$$

where the expression of $x(n, d)$ is given in (56). As a reminder, we have

$$x(n, d) := \frac{\log n}{\sqrt{nb_k}} \sqrt{\left[(d+1)4 \log \frac{1}{b_k} + 2(d+1) \log 2 + \log(\sqrt{n}+1) + d \log(2C(n, d) + 2) + \log \frac{8}{\delta}\right] / C'_7}.$$

We obtain there exist a universal constant $C'_{15} = 8/\sqrt{C'_7}$ such that

$$x(n, d) \leq \frac{C'_{15} \log n}{\sqrt{nb_k}} \left(\sqrt{(d+1) \log(\max\{W_x, 1\}(d+1)) \log n} + \sqrt{2 \log \frac{8}{\delta}} \right).$$

Then we finally achieve

$$\mathbb{P} \left(\sup_{\mathbf{u} \in U_k} Z(\mathbf{u}) \geq \frac{C'_{16} \log n}{\sqrt{nb_k}} \left(\sqrt{(d+1) \log(\max\{W_x, 1\}(d+1)) \log n} + \sqrt{2 \log \frac{8}{\delta}} \right) \right) = \frac{\delta}{2},$$

where we let $C'_{16} = 2 \max\{C'_{14}, C'_{15}\}$ and $C'_{17} = C'_{16} \log(\max\{W_x, e\})$. Thus, $nb_k \geq 4C'^2_{17} \log^3 n [(d+1) \log(d+1)]$ and $\delta \geq 8 \exp(-nb_k/(8C'^2_{17} \log^2 n))$ becomes a sufficient condition to make $x(n, d) + x_2(n, d)$ be smaller than 1. Following similar procedure, we are able to prove the same inequality for f_n , so we conclude our proof of Lemma 13. $\square$

The remaining part of Lemma 7 only involves getting a uniform upper bound for $|r_\theta(u) - r_{\theta_0}(u)|$ and thus $|\hat{r}_k(u, \theta) - r_{\theta_0}(u)|$ for any $\theta \in \Theta_k$ and $u \in I$. Similar with the corresponding proof of Lemma 2, we have

$$\sup_{u \in I, \theta \in \Theta_k} |r_\theta(u) - r_{\theta_0}(u)| \leq l_r R_{\mathcal{X}} \cdot \frac{2}{c_{\min}} \sqrt{\frac{(d+1)(6W_x^2 \log n + 6W_x \log^2 n \log \log n)}{C_w n}}.$$

Finally, by setting $b_k = n^{-1/(2m+1)}$ and combining our results obtained in Lemma 12 and Lemma 13, we conclude our results for Lemma 7. In addition, our way of deriving constants $B_{mx,K}$, $B'_{mx,K}$ and $C_{mx,K}$ is similar with that in Lemma 7, so we omit the details here. $\square$

F.8 Proof of Lemma 8 and Theorem 2

The proof of Lemma 8 and Theorem 2 are straight forward by combining the proof of Lemma 3 and Lemma 7, so we omit the details here.

G Additional Plots

In this section, we directly plot $\text{reg}(T)$ for all the settings discussed in the main paper. From Figure 9 - Figure 11, we see that the blue solid lines depicted in every figure are close to the other two lines that depict regrets with either known θ_0 or $g(\cdot)$ in Algorithm 1. This fact reflects the robustness of our estimators on θ_0 and $g(\cdot)$ in every episode.




Figure 9: From left to right, we plot empirical regret $\text{reg}(T)$ against $T^{(2m+1)/(4m-1)} \log^2 T$ with $m \in [2, 4, 6]$ in the setting with i.i.d. covariates with independent entries. Solid blue, green, red lines, represent the mean regret collected by implementing the Algorithm 1 for 30 times with unknown $g(\cdot)$, θ_0 , unknown $g(\cdot)$ but known θ_0 and known $g(\cdot)$ but unknown θ_0 in the exploitation phase respectively. Light color areas around those solid lines depict the standard error of our estimation of $\text{reg}(T)$.




Figure 10: From left to right, we plot empirical regret $\text{reg}(T)$ against $T^{(2m+1)/(4m-1)} \log^2 T$ with $m \in [2, 4, 6]$ in the setting with i.i.d. covariates but dependent entries. The rest caption is the same as in Figure 9.




Figure 11: From left to right, we plot empirical regret $\text{reg}(T)$ against $T^{(2m+1)/(4m-1)} \log_{10}^4 T$ with $m \in [2, 4, 6]$ in the setting with strong-mixing covariates. The rest caption is the same as in Figure 9.

H Regret bounds when $F(\cdot)$ is Lipschitz

All our main results require bounded second derivatives of F . This allows the pricing strategy $p_t = \hat{\phi}_k^{-1}(-\mathbf{x}_t^\top \hat{\boldsymbol{\theta}}) + \mathbf{x}_t^\top \hat{\boldsymbol{\theta}}$ to achieve low regret if the revenue function has bounded second derivative. When $F(\cdot)$ is only Lipschitz continuous, the above method is no longer applicable. Fortunately, we can directly define the offered price based on the substitution of $\hat{\boldsymbol{\theta}}$ and $\hat{F}$ into (5). We summarize these in

the following Algorithm 3.

Algorithm 3: Feature based dynamic pricing with unknown noise distribution when $F(\cdot)$ is ℓ -Lipschitz

- 1: **Input:** Upper bound of market value ($\{v_t\}_{t \geq 1}$): $B > 0$, minimum episode length: ℓ_0 , degree of smoothness: $m = 0$.
- 2: **Initialization:** $p_1 = 0$, $\hat{\theta}_1 = 0$.
- 3: **for** each episode $k = 1, 2, \dots$, **do**
- 4: Set length of the k -th episode $\ell_k = 2^{k-1}\ell_0$; Length of the exploration phase $a_k = \lceil (\ell_k d)^{\frac{3}{4}} \rceil$.
- 5: **Exploration Phase** ($t \in I_k := \{\ell_k, \dots, \ell_k + a_k - 1\}$):
- 6: Offer price $p_t \sim \text{Unif}(0, B)$.
- 7: **Updating Estimates (at the end of the exploration phase with data $\{(\tilde{\mathbf{x}}_t, y_t)\}_{t \in I_k}$):**
- 8: Update estimate of θ_0 by $\hat{\theta}_k = \hat{\theta}_k(\{(\tilde{\mathbf{x}}_t, y_t)\}_{t \in I_k})$;

$$\hat{\theta}_k = \underset{\theta}{\operatorname{argmin}} L_k(\theta) := \frac{1}{|I_k|} \sum_{t \in I_k} (By_t - \theta^\top \mathbf{x}_t)^2 \quad (61)$$

- 9: Update estimates of F , by $F_k(u, \hat{\theta}_k) = F_k(u; \hat{\theta}_k, \{(\tilde{\mathbf{x}}_t, y_t, p_t)\}_{t \in I_k})$ given by (14).
- 10: **Exploitation Phase** ($t \in I'_k := \{\ell_k + a_k, \dots, \ell_{k+1} - 1\}$):
- 11: Offer p_t as

$$p_t = \operatorname{argmax}_{p \geq 0} \{p(1 - \hat{F}_k(p - \mathbf{x}_t^\top \hat{\theta}_k))\} \quad (62)$$

12: **end for**

Theorem 1. Let Assumptions 1, 3, 4 and 5 hold. Then there exist constants C (depending only on the absolute constants within the assumptions) such that for all T satisfying $T \geq Cd$, the regret of Algorithm 3 over time T is no more than $C_{x,K}^*(Td)^{\frac{3}{4}} \log T(1 + \log T/d)$.

Proof. We write

$$\mathbb{E}[R_t | \bar{\mathcal{H}}_{t-1}] = p_t^*(1 - F(p_t^* - \mathbf{x}_t^\top \theta_0)) - p_t(1 - F(p_t - \mathbf{x}_t^\top \theta_0)) \quad (63)$$

$$= \operatorname{rev}_t(p_t^*, \theta_0, F) - \operatorname{rev}_t(p_t, \theta_0, F). \quad (64)$$

The last inequality follows from our definition of (4). When $t \in I'_k$ (the k -th exploitation phase) we can then further expand (64) into

$$(64) = \operatorname{rev}_t(p_t^*, \theta_0, F) - \operatorname{rev}_t(p_t^*, \hat{\theta}_k, F) \quad (65)$$

$$+ \operatorname{rev}_t(p_t^*, \hat{\theta}_k, F) - \operatorname{rev}_t(p_t^*, \hat{\theta}_k, \hat{F}_k) \quad (66)$$

$$+ \operatorname{rev}_t(p_t^*, \hat{\theta}_k, \hat{F}_k) - \operatorname{rev}_t(p_t, \hat{\theta}_k, \hat{F}_k) \quad (67)$$

$$+ \operatorname{rev}_t(p_t, \hat{\theta}_k, \hat{F}_k) - \operatorname{rev}_t(p_t, \hat{\theta}_k, F) \quad (68)$$

$$+ \text{rev}_t(p_t, \hat{\theta}_k, F) - \text{rev}_t(p_t, \theta_0, F). \quad (69)$$

For (67), by our definition on p_t in (8), we have

$$(67) = \text{rev}_t(p_t^*, \hat{\theta}_k, \hat{F}_k) - \text{rev}_t(p_t, \hat{\theta}_k, \hat{F}_k) \leq 0 \quad (70)$$

For terms (65) and (69), we can control both of them by difference between $\hat{\theta}_k$ and θ_0 in a sense that

$$(65), (69) \lesssim |\langle x_t, \hat{\theta}_k - \theta_0 \rangle| \lesssim \frac{1}{\sqrt{a_k}} \quad (71)$$

holds with high probability by our Lemma 4.1, since we assume F is Lipschitz continuous. Recall $a_k = |I_k|$, which is the length of the k -th exploration phase

For the rest two parts (66) and (68), as we are able to control $\hat{F}(x)$ to $F(x)$ uniformly with rate $a_k^{-1/3}$ using data in the exploration phase. we can then bound $\mathbb{E}[R_t]$ by

$$\mathbb{E}[R_t] = \mathbb{E}[\mathbb{E}[R_t | \mathcal{H}_{t-1}]] \lesssim \frac{1}{a_k^{1/3}} \quad (72)$$

Then for the regret in k -th episode we can bound it as

$$\text{Regret}_k = \sum_{t \in I_k} (\text{rev}_t^* - \text{rev}_t) + \sum_{t \in E_k \setminus I_k} (\text{rev}_t^* - \text{rev}_t) \quad (73)$$

$$\leq Ba_k + a_k/a_k^{1/3} = a_k^{3/4} + a_k/a_k^{1/4} = \mathcal{O}(a_k^{3/4}) \quad (74)$$

let $K = \lceil \log_2 T \rceil + 1$, we have our total regret can be bounded by

$$\text{Regret}_\pi(T) = \sum_{k=1}^K 2^{3(k-1)/4} = \mathcal{O}(T^{3/4}). \quad (75)$$

□

I A Data Driven Way to Determine m

As mentioned in Remark 12, we are able to adopt the cross-validation method [Hall and Racine, 2015] to determine the order of smoothness m using data from the previous exploration phase. In the below, we briefly introduce how the order of smoothness can be determined in local polynomial regression in the context of nonparametric regression.

Given training data $\{x_i, y_i\}_{i=1}^n$ and we assume they are generated following model

$$Y = g^*(X) + \epsilon,$$

with $\mathbb{E}[\epsilon | X] = 0$. Define

$$\text{CV}(h, m) = \frac{1}{n} \sum_{i=1}^n (Y_i - \hat{g}_{-i}(X_i))^2,$$

where $\hat{g}_{-i}(\cdot)$ is fitted using all samples except the i -th pair (x_i, y_i) . Here we use bandwidth h and m -th order local polynomial to fit the regression function. According to Theorem 3.2 given in Hall and Racine [2015], optimizing $\text{CV}(h, m)$ is equivalent to optimizing over (h, m) with respect to the averaged summed squared errors defined in (76) up to some small order terms.

$$\frac{1}{n} \sum_{i=1}^n (\hat{g}(x_i) - g^*(x_i))^2. \tag{76}$$

Thus, this method is a valid way to determine the order of smoothness. We summarize the combined

procedure in Algorithm 4.

Algorithm 4: Feature based dynamic pricing with unknown m

- 1: **Input:** Upper bound of market value ($\{v_t\}_{t \geq 1}$): $B > 0$, minimum episode length: ℓ_0 .
- 2: **Initialization:** $p_1 = 0$, $\hat{\theta}_1 = 0$.
- 3: **for** each episode $k = 1, 2, \dots$, **do**
- 4: If $k \geq 2$, use $\{\mathbf{x}_t^\top \hat{\theta}_{k-1}, p_t, y_t\}_{t \in I_{k-1}}$ and Algorithm 4 to determine m . If $k = 1$, set $\hat{m} = 2$.
- 5: Set length of the k -th episode $\ell_k = 2^{k-1} \ell_0$; Length of the exploration phase $a_k = \lceil (\ell_k d)^{\frac{2\hat{m}+1}{4\hat{m}-1}} \rceil$.
- 6: **Exploration Phase** ($t \in I_k := \{\ell_k, \dots, \ell_k + a_k - 1\}$):
- 7: Offer price $p_t \sim \text{Unif}(0, B)$.
- 8: **Updating Estimates (at the end of the exploration phase with data $\{(\tilde{\mathbf{x}}_t, y_t)\}_{t \in I_k}$):**
- 9: Update estimate of θ_0 by $\hat{\theta}_k = \hat{\theta}_k(\{(\tilde{\mathbf{x}}_t, y_t)\}_{t \in I_k})$;

$$\hat{\theta}_k = \underset{\theta}{\operatorname{argmin}} L_k(\theta) := \frac{1}{|I_k|} \sum_{t \in I_k} (B y_t - \theta^\top \tilde{\mathbf{x}}_t)^2. \quad (77)$$

- 10: If $\hat{m} \geq 1$, update estimates of F, F' by $F_k(u, \hat{\theta}_k) = F_k(u; \hat{\theta}_k, \{(\tilde{\mathbf{x}}_t, y_t, p_t)\}_{t \in I_k}, \hat{h}_k)$, $F_k^{(1)}(u, \hat{\theta}_k) = F_k^{(1)}(u, \hat{\theta}_k, \{(\tilde{\mathbf{x}}_t, y_t, p_t)\}_{t \in I_k}, \hat{h}_k)$. The detailed formulas are given by (14) and (16).
- 11: Update estimate of ϕ by $\hat{\phi}_k(u) = u - \frac{1 - \hat{F}_k(u)}{\hat{F}^{(1)}(u)}$ and estimate of g by $\hat{g}_k(u) = u + \hat{\phi}_k^{-1}(-u)$.
 If $\hat{m} = 0$, update estimates of F , by $F_k(u, \hat{\theta}_k) = F_k(u; \hat{\theta}_k, \{(\tilde{\mathbf{x}}_t, y_t, p_t)\}_{t \in I_k})$, The detailed formulas are given by (14).
- 12: **Exploitation Phase** ($t \in I'_k := \{\ell_k + a_k, \dots, \ell_{k+1} - 1\}$):
- 13: If $\hat{m} \geq 1$, offer p_t as

$$p_t = \min\{\max\{\hat{g}_k(\tilde{\mathbf{x}}_t^\top \hat{\theta}_k), 0\}, B\}. \quad (78)$$

If $\hat{m} = 0$, offer p_t as

$$p_t = \operatorname{argmax}_{p \geq 0} \{p(1 - \hat{F}_k(p - \mathbf{x}_t^\top \hat{\theta}_k))\}.$$

14: **end for**

Algorithm 5: Selection of m .

- 1: **Input:** Data $\{\mathbf{x}_t^\top \hat{\theta}_{t-1}, p_t, y_t\}_{t \in I_{k-1}}$
- 2: **For** $(m, h) \in \mathcal{M} \times \mathcal{H}$, **compute:**

$$(\hat{m}, \hat{h}) = \underset{(m, h)}{\operatorname{argmin}} L(m, h) = \frac{1}{|I_{k-1}|} \sum_{i=1}^{|I_{k-1}|} (Y_i - \hat{g}_{-i}^{(m, h)}(X_i))^2$$

- 3: **Output:** $\hat{m}$
-

References

- M. Banna, F. Merlevède, and P. Youssef. Bernstein-type inequality for a class of dependent random matrices. *Random Matrices: Theory and Applications*, 05(02):1650006, 2016.
- D. Bosq. *Nonparametric Statistics for Stochastic Processes: Estimation and Prediction*, volume 110. Springer Science & Business Media.
- J. Broder and P. Rusmevichientong. Dynamic pricing under a general parametric choice model. *Operations Research*, 60(4):965–980, 2012.
- R. J. Carroll, J. Fan, I. Gijbels, and M. P. Wand. Generalized partially linear single-index models. *Journal of the American Statistical Association*, 92(438):477–489, 1997. doi: 10.1080/01621459.1997.10474001.
- R. Eubank and P. L. Speckman. Confidence bands in nonparametric regression. *Journal of the American Statistical Association*, 88(424):1287–1301, 1993.
- J. Fan and I. Gijbels. *Local polynomial modelling and its applications*. Chapman and Hall, 1996.
- J. Fan and R. Li. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American statistical Association*, 96(456):1348–1360, 2001.
- J. Fan, M. Farman, and I. Gijbels. Local maximum likelihood estimation and inference. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 60(3):591–608, 1998.
- P. Hall. *The bootstrap and Edgeworth expansion*. Springer-Verlag New York, 1992.
- P. Hall and J. Horowitz. A simple bootstrap method for constructing nonparametric confidence bands for functions. *The Annals of Statistics*, 41(4):1892 – 1921, 2013.
- P. G. Hall and J. S. Racine. Infinite order cross-validated local polynomial regression. *Journal of Econometrics*, 185(2):510–525, 2015. ISSN 0304-4076.
- J. L. Horowitz. *The Bootstrap*, volume 5 of *Handbook of Econometrics*. Elsevier, 2001.
- A. Javanmard and H. Nazerzadeh. Dynamic pricing in high-dimensions. *The Journal of Machine Learning Research*, 20(1):315–363, 2019.
- F. Merlevède, M. Peligrad, and E. Rio. *Bernstein inequality and moderate deviations under strong mixing conditions*, volume Volume 5 of *Collections*. Institute of Mathematical Statistics, Beachwood, Ohio, USA, 2009. doi: 10.1214/09-IMSCOLL518.
- M. H. Neumann and J. Polzehl. Simultaneous bootstrap confidence bands in nonparametric regression. *Journal of Nonparametric Statistics*, 9(4):307–333, 1998.

- C. Stone. Optimal rates of convergence for nonparametric estimators. *The Annals of Statistics*, 8(6): 1348–1360, 1980.
- C. J. Stone. Optimal Global Rates of Convergence for Nonparametric Regression. *The Annals of Statistics*, 10(4):1040 – 1053, 1982. doi: 10.1214/aos/1176345969.
- A. B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer Publishing Company, Incorporated, 2008. ISBN 0387790519.
- R. Vershynin. *Introduction to the non-asymptotic analysis of random matrices*. Cambridge University Press, 2012.
- R. Vershynin. *High-Dimensional Probability: An Introduction with Applications in Data Science*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2018.
- Y. Wang, X. Chen, X. Chang, and D. Ge. Uncertainty quantification for demand prediction in contextual dynamic pricing. *Production and Operations Management*, 30, 12 2020. doi: 10.1111/poms.13337.
- C.-H. Zhang. Nearly unbiased variable selection under minimax concave penalty. *The Annals of Statistics*, 38(2):894– 942, 2010.
- P. Zhao and B. Yu. On model selection consistency of lasso. *The Journal of Machine Learning Research*, 7:2541–2563, 2006.