

Information-theoretic limitations of data-based price discrimination*

Haitian Xie[†] Ying Zhu[‡] Denis Shishkin[§]

October 24, 2023

Abstract

This paper studies third-degree price discrimination (3PD) based on a random sample of valuation and covariate data, where the covariate is continuous, and the distribution of the data is unknown to the seller. The main results of this paper are twofold. The first set of results is pricing strategy independent and reveals the fundamental information-theoretic limitation of *any* data-based pricing strategy in revenue generation for two cases: 3PD and uniform pricing. The second set of results proposes the K -markets empirical revenue maximization (ERM) strategy and shows that the K -markets ERM and the uniform ERM strategies achieve the optimal rate of convergence in revenue to that generated by their respective true-distribution 3PD and uniform pricing optima. Our numerical results show that the uniform (i.e., 1-market) ERM strategy may generate a larger revenue than the K -markets ERM strategy when the sample size is small enough, and vice versa.

Keywords: price discrimination, information theory, empirical revenue maximization, revenue deficiency, prior-independent mechanism design, optimal rate of convergence, packing sets

MSC classification code: 62B10, 62R07, 91B24, 94A16

JEL classification code: C14, C44, D42, D82

*Haitian Xie and Ying Zhu contributed equally to this work. This paper supersedes a previously circulated draft by Xie and Zhu (<https://arxiv.org/pdf/2204.12723v1.pdf>). The authors would like to thank Dirk Bergemann, Songzi Du, Federico Echenique, Yannai Gonczarowski, Roger Gordon, Nima Haghpahan, Johannes Horner, Jonathan Libgober, Maximilian Schaefer, Joel Sobel, Karl Schlag, Larry Samuelson, J. Miguel Villas-Boas, and Joel Watson for valuable comments and discussions.

[†]Assistant professor at the Department of Business Statistics and Econometrics, Guanghua School of Management, Peking University. The author is grateful to the UC San Diego Department of Economics where this project was developed during his doctoral studies. Email: xht@gsu.pku.edu.cn.

[‡]Assistant professor at the Department of Economics, University of California, San Diego. Email: yiz012@ucsd.edu. The author is grateful to the Society of Hellman Fellows at University of California and the Cowles Foundation at Yale University, and also thanks participants at her seminars.

[§]Assistant professor at the Department of Economics, University of California, San Diego. Email: dshishkin@ucsd.edu.

1 Introduction

In the past few decades, the advances in the theory of mechanism design have been followed by a tremendous interest in its practical applications. At the same time, classic theoretical models typically make strong assumptions about the designer’s knowledge of the environment which may lead the optimal mechanism to be sensitive to the details of the environment (which is sometimes referred to as the Wilson critique).¹

Consider, for example, a classic monopoly pricing problem. Suppose a buyer has a unit demand, a privately-known valuation Y , and a covariate X (which may or may not be observed by the seller) drawn from some joint distribution $F_{Y,X}$, known to the seller. Depending on the environment, the seller may or may not be able to third-degree price discriminate (3PD) between buyers with different values of X . First, suppose that 3PD is infeasible either because the realization of X is not observed by the seller or X corresponds to a protected characteristic such as race or gender, and regulatory constraints prohibit price discrimination. Then, standard results (Riley and Zeckhauser, 1983) imply that it is optimal for the seller to choose a uniform pricing strategy that corresponds to a take-it-or-leave-it offer characterized by a single (posted) price. Similarly, when 3PD is feasible, and the realization of X is observed by the seller, the optimal strategy is to charge the optimal uniform price conditioning on the realization of X . In terms of generating revenue, the classic pricing theory shows that 3PD is at least as good as uniform pricing.

Firms have been using simple 3PD strategies for a long time, e.g., offering discounts to students or seniors. In the Internet era, 3PD strategies have become more sophisticated, e.g., firms might do targeted pricing for online shoppers based on browsing history, location, and even the type of user device or operating system (see, e.g., Hannak, Soeller, Lazer, Mislove and Wilson (2014)). The covariate X can be a single index or score that summarizes the relevant characteristics for pricing and marketing. Hartmann, Nair and Narayanan (2011) provides examples where marketing firms use a one-dimensional continuous score function of customer characteristics, past response histories, and other features of the zip code, and casinos use a one-dimensional continuous score referred to as the average daily win.

The key assumption underlining the classic pricing theory is that the distribution of buyer valuations (and the covariate) is known to the seller. When a seller has only partial information about the distribution, how much revenue can be obtained then?² In this paper, we assume that

¹In some cases, this leads to extreme or unrealistic results as in, e.g., Crémer and McLean (1988).

²We assume that the cost is known to the seller. Therefore, the valuation can be treated as the net valuation, and

the seller has access to a random sample of i.i.d. $\{Y_i, X_i\}_{i=1}^n$ drawn from $F_{Y,X}$, which is unknown to the seller. The main results of this paper are twofold. The first set of results is algorithm (pricing strategy) independent and reveals the fundamental information-theoretic limitation of *any* data-based pricing strategy. Namely, if the seller has to rely on a random sample, then no matter what pricing strategy is used, there is still an inevitable uncertainty about the true distribution, which, in the worst case, precludes the seller from attaining the true-distribution optimal revenue. The second set of results proposes simple data-based pricing strategies that achieve the (inevitable) information-theoretic limitation up to a constant independent of the sample size n . Compared to the existing approaches (discussed at the end of the section), our treatment allows us to study third-degree price discrimination with a continuous covariate.

On the information-theoretic limitation, we establish a lower bound (in terms of n , up to a constant) for the revenue deficiency in *any* data-based pricing strategy relative to the true-distribution optimal strategy in the worst case (by considering the supremum over a class of joint distributions, $F_{Y,X}$, subject to some mild smoothness assumptions). In particular, data-based uniform pricing strategies are algorithms that depend on $\{Y_i\}_{i=1}^n$ only, and the true-distribution optimal strategy corresponds to the optimal uniform pricing strategy derived from F_Y . Similarly, data-based 3PD strategies are algorithms that depend on $\{Y_i, X_i\}_{i=1}^n$, and the true-distribution optimal strategy corresponds to the optimal 3PD strategy derived from $F_{Y,X}$. We show that the minimax revenue deficiency is at least $cn^{-\frac{2}{3}}$ and $cn^{-\frac{1}{2}}$ in the uniform and 3PD cases, respectively, for some constant $c \in (0, \infty)$ independent of n .

On the achievability of the (inevitable) information-theoretic limitation, we consider the optimal posted price that maximizes the revenue based on the empirical CDF of $\{Y_i\}_{i=1}^n$, often referred to as the empirical revenue maximization (ERM),³ for the uniform pricing problem, and propose the K -markets ERM for the 3PD problem. The K -markets ERM strategy divides the covariate space into K equal-length segments (where K grows at a rate of $n^{\frac{1}{4}}$), and the optimal posted price in each segment is calculated for the conditional empirical distribution in that segment. We show that, in terms of the revenue deficiency, the rates $n^{-\frac{2}{3}}$ and $n^{-\frac{1}{2}}$ are achieved by the uniform ERM strategy and by the K -markets ERM strategy in the respective realms of uniform pricing and 3PD.⁴

revenue is the same as profit.

³Even if a seller has access to all data (that allows her to observe the true distribution), she might not be able to compute the optimal price based on all data due to computational/data storage constraints. Therefore, the seller may perform ERM based on random samples.

⁴In the 3PD case where the covariate takes a few values, we can apply our analysis for uniform pricing to each covariate

The two sets of our results reveal several interesting facts. First, for generating revenues in their respective regime (data-based uniform pricing and 3PD), the uniform ERM and the K -markets ERM strategies attain their respective optimal rate of convergence. Somewhat surprisingly, more sophisticated pricing strategies (e.g., with partitioning the covariate space based on observed frequencies) can *at best* improve upon the simple K -markets algorithm in the constant (which is independent of n).

Second, in the 3PD case with a continuous covariate, too much or too little discrimination could undermine revenue generation. Setting the optimal price for each different observed covariate value in the random sample may not “extrapolate” well to unobserved covariate values, so the expected revenue with respect to the entire population can be compromised. On the other hand, too little discrimination underutilizes the information in $\{X_i\}_{i=1}^n$ about $\{Y_i\}_{i=1}^n$ and compromises the expected revenue as well.

When the seller has the access to a sample of i.i.d. $\{Y_i, X_i\}_{i=1}^n$, she can choose the uniform ERM strategy that ignores $\{X_i\}_{i=1}^n$ and makes use of only the valuation data $\{Y_i\}_{i=1}^n$, or the K -markets ERM strategy that makes use of both $\{X_i\}_{i=1}^n$ and $\{Y_i\}_{i=1}^n$. Inherently, the latter is an algorithm trying to learn the $F_{Y,X}$ -optimal pricing function $p(\cdot)$ while the former is an algorithm trying to learn F_Y -optimal (constant) pricing function. As a result, the latter is less demanding in the sample size than the former. On the other hand, in terms of generating revenue, the true-distribution optimal 3PD is at least as good as the true-distribution optimal uniform pricing strategy. This trade-off suggests that the K -markets ERM strategy based on a random sample can be revenue inferior to the uniform ERM strategy if the sample size n is not large enough, and vice versa.

To illustrate the implications of our theoretical results, we also provide two numerical studies. In particular, we calculate the revenues of the K -markets ERM and the uniform ERM strategies based on a real-world data set from eBay auctions and a simulated data set. In both cases, our numerical results illustrate the aforementioned trade-off. When the sample size is small, the uniform ERM strategy generates higher expected revenue than the K -markets ERM strategy. As the sample size grows, the K -market ERM strategy (the uniform ERM strategy) gets closer to the true-distribution optimal 3PD strategy (respectively, the true-distribution optimal uniform pricing strategy). The slower rate of convergence in the revenue from the K -markets ERM strategy (in contrast to the faster rate of convergence in the revenue from the uniform ERM strategy) is

value. This case is much simpler than the problem of our interest, where the covariate is continuous.

dominated by the benefit of price discrimination (based on the underlying $F_{Y,X}$) over uniform pricing (based on F_Y). Consequently, the revenue of the K -markets ERM strategy overtakes that of the uniform ERM strategy when the sample size becomes sufficiently large.

1.1 Related literature

Complexity measures and information theoretic lower bounds. Information theoretic lower bounds and sample complexity are important notions in machine learning. They are two sides of the same coin. Both aim to characterize learnability, i.e., how easy it is to learn an unknown object of interest (in our context, the true-distribution optimal 3PD strategy) from data where the uncertainty arises. Sample complexity derives the rate at which the sample size needs to grow to guarantee a desired learning accuracy. Information theoretic lower bound derives a lower bound on the learning error (in our context, the revenue deficiency) in the worst case for essentially any given sample size. Sample complexity and information theoretic lower bounds are intrinsically tied to the complexity or size of the underlying function class of interest. Vapnik-Chervonenkis (VC) dimensions, shattering dimensions, and metric entropy (such as packing numbers) are popular measures of complexity in machine learning. There have been a number of innovative applications of VC dimensions or shattering dimensions in economic theory and algorithmic economics. Together with the Probably Approximately Correct (PAC) framework, they are used to study the complexity of the classes of demand and utility functions (Beigman and Vohra, 2006; Balcan, Daniely, Mehta, Urner and Vazirani, 2014), k -demand buyer's valuation (Zhang and Conitzer, 2020), theories of choices (Basu and Echenique, 2020), preference functions (Chambers, Echenique and Lambert, 2021, 2023), as well as the resulting learnability from data. VC dimension is useful for deriving sample complexity bounds concerning discrete function sets and finite dimensional vector spaces, and shattering dimension is useful for certain real functions.

From the theory of machine learning, when a class has infinite VC or shattering dimensions, this class is not PAC learnable. For example, a collection of sinusoids have subgraphs with infinite VC dimension. The class of demand functions has infinite shattering dimension (Beigman and Vohra, 2006). Nonetheless, the notion of "learnability" can be generalized using a different type of complexity analysis that gives rise to our information theoretic lower bound in the 3PD problem. This type of analysis is built upon the notion of packing sets, along with tools from information theory. In particular, packing sets and packing numbers are useful for studying classes

with an infinite number of elements; see the discussion about metric entropy in [Wainwright \(2019\)](#).

In our 3PD problem, packing sets are useful as we try to learn an optimal pricing function of the covariate, an infinitely-dimensional parameter. In deriving the lower bound for the deficiency in the *expected* revenue in the 3PD problem, we convert the pricing problem into a multiple classification problem that tries to distinguish between optimal prices in a carefully constructed packing set. Our proof is based on a delicate construction of conditional densities along with the Fano’s inequality and the Gilbert-Varshamov bound from information theory. The key lies in constructing a sufficiently large set of distributions with small enough pairwise Kullback-Leibler divergence such that one cannot distinguish them from the data, but their associated optimal prices have a sufficiently large separation with an unimprovable rate. This rate is the information theoretic limitation, which turns out to be $n^{-1/4}$. The desired set of optimal prices in our proof is a packing set where the separation between the elements is $\Omega(n^{-1/4})$ with respect to the unweighted L_2 norm, and the packing number at that separation is $\Omega(n^{1/4})$. The uniform pricing problem is much easier as it only requires constructing two distributions and simpler techniques.

Prior independent mechanism design. Most of the classic monopoly pricing literature assumes a known distribution of valuations (and covariates).⁵ More recently, some papers (e.g., those surveyed in [Carroll, 2019](#)) studied “prior”-independent mechanism design.⁶ The main focus of that literature is on deriving a robustly optimal mechanism in the absence of both “prior” and data. In particular, [Bergemann and Schlag \(2008, 2011\)](#) derive the minimax-regret uniform pricing strategy in closed form; that is, the strategy that guarantees the smallest deficiency in revenue relative to the known distribution case. Like [Bergemann and Schlag \(2008, 2011\)](#), we study the revenue deficiencies, but in contrast, we assume the availability of data and focus on the (inevitable) information-theoretic limitations of any *data-based* pricing strategies and the achievability of the limitation.

This paper is inspired by the literature that studies approximately optimal “prior”-independent mechanism design, in particular monopoly pricing with a single buyer.⁷ This literature assumes

⁵See also [Segal \(2003\)](#) for a study of optimal multi-unit auctions where the seller has a probabilistic belief about the valuation distribution of the i.i.d. buyers.

⁶Here, “prior” distribution refers to the seller’s prior belief about buyers’ valuations and is often taken to be the true distribution.

⁷There is a less related literature that studies optimal auctions; see, e.g., [Cole and Roughgarden \(2014\)](#); [Dhangwatnotai, Roughgarden and Yan \(2015\)](#); [Fu, Immorlica, Lucier and Strack \(2015\)](#); [Guo, Huang and Zhang \(2019\)](#); [Fu, Haghpahan, Hartline and Kleinberg \(2021\)](#).

that the seller has access to a random sample of i.i.d. $\{Y_i\}_{i=1}^n$ drawn from F_Y and proposes variants of the uniform ERM strategy to derive the revenue guarantee in relation to that from the true-distribution optimal uniform pricing strategy. There are two types of analyses in this literature. The first one focuses on the guarantees for the specific case of $n = 1$ or $n = 2$ (Babaioff, Gonczarowski, Mansour and Moran, 2018; Allouah, Bahamou and Besbes, 2023). The second one (e.g., Huang, Mansour and Roughgarden, 2018) establishes “sample complexity bounds” such that the uniform ERM variants achieve a $(1 - \epsilon)$ fraction guarantee when the sample size grows at a rate depending on ϵ , and also derives the rate at which the sample size needs to grow (as a function of ϵ) for any data-based uniform pricing strategies to obtain a given $(1 - \epsilon)$ fraction guarantee. Allouah, Bahamou and Besbes (2022) involves both types of analyses.

In this paper, we ask the related question, how fast the revenue deficiency decays as a function of n , and provide an answer using information-theoretic lower bounds (independent of algorithms) and upper bounds with respect to specific algorithms in the worst case scenarios.⁸ The main difference with the majority of the data-based literature is that, we study third-degree price discrimination (3PD) with a continuous covariate and compare the performance of data-based 3PD strategies with that of data-based uniform pricing strategies.

To understand why the 3PD problem in our context is more challenging than the uniform pricing problem, note that fundamentally, the latter tries to learn the constant optimal pricing function (a scalar parameter) while the former tries to learn an optimal pricing function of the covariate (an infinitely-dimensional parameter), where the deficiency in the *expected* revenue concerns the entire pricing function at all covariate values. Our framework allows us to tackle several challenging aspects of our 3PD problem, which might be difficult to analyze with the sample complexity toolkit. We describe one example below.

Somewhat related, Devanur, Huang and Psomas (2016) studies sample complexity of optimal pricing with “side information”. In their “signals model” (Sections 5.1 and 5.3), there is a covariate (signal) $X \in [0, 1]$, and the seller can condition the data-based reserve price on the covariate. For the single-buyer case (which would correspond to our 3PD problem), they derive upper and lower sample complexity bounds. Importantly, they assume that the true joint distribution $F_{Y,X}$ has the following property: larger values of X are associated with larger values

⁸A large literature studies data-based auctions by focusing on guarantees for revenue deficiencies (instead of fractions), such as how the revenues from the data-based strategies converge in probability to the true-distribution benchmark, e.g., Baliga and Vohra (2003); Goldberg, Hartline, Karlin, Saks and Wright (2006); Gonçalves and Furtado (2020). This line of work does not consider the optimal rates of convergence and learnability issues such as sample complexity or information theoretic lower bounds.

of Y in the sense of first-order stochastic dominance of conditional distributions. In contrast, our 3PD setup imposes no assumptions about the relationship between the covariate X and the valuation Y ; meanwhile, our proposed K -market ERM strategy learns the relationship from the data. Moreover, our K -market ERM strategy attains the optimal rate of convergence in revenue deficiency (as described before), while the upper and lower bounds in [Devanur et al. \(2016\)](#) have different rates, and hence, one cannot conclude whether these rates are optimal or not.

Organization of this paper. The rest of the paper is organized as follows. Section 2 introduces the problem setup. Section 3 presents the information-theoretic lower bounds for data-based pricing strategies. Section 4 proposes the uniform ERM and the K -markets ERM strategies, and shows that they achieve the (inevitable) information-theoretic limitation up to a constant independent of the sample size. Section 5 exhibits the empirical and simulation results. Section 6 concludes the paper and discusses the broader connection of the insights from this paper with literature beyond empirical revenue maximization. Proofs are presented in Appendices A, B and C.

2 Setup

We begin by introducing the problem setup. The seller is selling an item to a buyer. Let $Y \in [0, 1]$ be the valuation (i.e., willingness to pay) of the buyer, and X the covariate (such as a characteristic) associated with the buyer. The joint distribution of (Y, X) is denoted by $F_{Y,X}$. We assume that X is supported on a bounded interval, and without loss of generality, we take the interval to be $[0, 1]$. Given a covariate value, the seller wants to set a price according to a mapping from the covariate to a set of prices. We use \mathcal{D} to denote the set of all pricing functions:

$$\mathcal{D} \equiv \{p: [0, 1] \rightarrow [0, 1], \text{ measurable}\}.$$

For a generic pricing strategy $p \in \mathcal{D}$, the price depends on the covariate value x . This scheme falls in the realm of third-degree price discrimination (3PD). Uniform pricing can be viewed as a special case where the price is the same for all covariate values. We use \mathcal{U} to denote the set of all uniform pricing functions:

$$\mathcal{U} \equiv \{p \in \mathcal{D}: p \text{ is a constant function}\}.$$

To lighten the notation, we express $p \in \mathcal{U}$ as a scalar rather than a function for the uniform pricing problem.

Let $F_{Y|X}$ be the conditional distribution and f_X the marginal density function. Given a price $y \in [0, 1]$ and a covariate value $x \in [0, 1]$, there are $1 - F_{Y|X}(y|x)$ buyers whose valuation is above the price. The revenue generated from these buyers is

$$r(y, x, F_{Y,X}) \equiv (1 - F_{Y|X}(y|x))y, \quad (1)$$

and the *expected* revenue for a pricing function p is

$$R(p, F_{Y,X}) \equiv \int_0^1 r(p(x), x, F_{Y,X}) f_X(x) dx.$$

In various places of the rest of the paper, we will slightly abuse the notation and denote $r(p, x) \equiv r(p(x), x)$ when p is a pricing function and also simply write $r(y, x) = r(y, x, F_{Y,X})$ for brevity when $F_{Y,X}$ is clear from the context. In the special case where the pricing strategy is uniform (i.e., $p \in \mathcal{U}$), the revenue only depends on the marginal distribution F_Y :

$$\begin{aligned} R(p, F_{Y,X}) &= p \int_0^1 (1 - F_{Y|X}(p|x)) f_X(x) dx \\ &= p \int_0^1 \mathbb{P}(Y \geq p | X = x) f_X(x) dx \\ &= p \mathbb{P}(Y \geq p) \\ &= p(1 - F_Y(p)), p \in \mathcal{U}. \end{aligned}$$

The true-distribution optimal 3PD strategy p_D^* is the one that maximizes the revenue:

$$R(p_D^*, F_{Y,X}) = \sup_{p \in \mathcal{D}} \int_0^1 r(p(x), x) f_X(x) dx.$$

In a similar fashion, we denote p_U^* as the true-distribution optimal uniform pricing strategy such that

$$R(p_U^*, F_Y) = R(p_U^*, F_{Y,X}) = \sup_{p \in \mathcal{U}} p(1 - F_Y(p)).$$

Note that p_D^* depends on $F_{Y,X}$ and p_U^* depends on F_Y .

In terms of generating revenue, the classic pricing theory shows that 3PD is at least as good

as uniform pricing when the joint distribution $F_{Y,X}$ is known to the seller. In this case, we can solve analytically for the optimal pricing strategies p_D^* and p_U^* . Since \mathcal{U} is contained in \mathcal{D} , p_D^* must achieve a (weakly) better revenue than p_U^* . Intuitively, when Y is correlated with X , p_D^* utilizes the information in X .

Now suppose that the seller knows neither $F_{Y,X}$ nor F_Y , but instead observes a random sample of $data \equiv \{(Y_i, X_i), 1 \leq i \leq n\}$ drawn from $F_{Y,X}$, or $data_Y \equiv \{Y_i, 1 \leq i \leq n\}$ from F_Y . The following assumption is used throughout this paper.

Assumption 1. *data and data_Y consist of i.i.d. draws from $F_{Y,X}$ and F_Y , respectively.*

The seller wants to construct a pricing strategy based on the sample. That is, let $\check{p}_D(data)$ be a pricing function in \mathcal{D} and $\check{p}_D(x_0; data)$ be its value at a covariate $x_0 \in [0, 1]$. Similarly, let $\check{p}_U(data_Y)$ be a uniform pricing function in \mathcal{U} and $\check{p}_U(x_0; data_Y)$ be its value at a covariate $x_0 \in [0, 1]$. Finally, let $\check{\mathcal{D}}$ be the set of all data-based 3PD functions \check{p}_D and $\check{\mathcal{U}}$ be the set of all data-based uniform pricing functions \check{p}_U .

Notation. For functions $f(n)$ and $g(n)$, we write $f(n) \gtrsim g(n)$ to mean that $f(n) \geq cg(n)$ for some constant $c \in (0, \infty)$. Similarly, we write $f(n) \lesssim g(n)$ to mean that $f(n) \leq c'g(n)$ for some constant $c' \in (0, \infty)$, and write $f(n) \asymp g(n)$ when $f(n) \gtrsim g(n)$ and $f(n) \lesssim g(n)$ hold simultaneously. As a general rule for this paper, the various c and C constants denote positive universal constants that are independent of the sample size n , and may vary from place to place. For functions f and g , the unweighted L_2 norm (L_2 as the short form) $\|f - g\|_2 \equiv \left(\int_0^1 [f(x) - g(x)]^2 dx \right)^{\frac{1}{2}}$.

3 Information-theoretic limitation of data-based pricing

To show the fundamental information-theoretic limitation of *any* data-based pricing strategies by comparing them with the true-distribution optimal pricing strategies, lower bounds are needed. It makes little sense to consider a framework recommending the data-based pricing strategies that are only good for a single distribution. For any *fixed* joint distribution $F_{Y,X}$, there is always a trivial data-based pricing strategy: simply ignore the data and select the optimal pricing scheme given $F_{Y,X}$. For this particular distribution, the revenue deficiency is zero. However, such a pricing strategy may perform poorly under other distributions of (Y, X) . One solution to circumvent this issue is to introduce a class \mathcal{F} of joint distributions and take the supremum over the class \mathcal{F} .

To be specific, we consider the minimax difference in the revenues at a given covariate value x_0 for 3PD,

$$\mathcal{R}_n^D(x_0; \mathcal{F}) \equiv \inf_{\check{p}_D \in \check{\mathcal{D}}} \sup_{F_{Y,X} \in \mathcal{F}} (r(p_D^*, x_0, F_{Y,X}) - \mathbb{E}_{F_{Y,X}} [r(\check{p}_D(data), x_0, F_{Y,X})]),$$

and the minimax difference in the *expected* revenues for 3PD,

$$\mathcal{R}_n^D(\mathcal{F}) \equiv \inf_{\check{p}_D \in \check{\mathcal{D}}} \sup_{F_{Y,X} \in \mathcal{F}} (R(p_D^*, F_{Y,X}) - \mathbb{E}_{F_{Y,X}} [R(\check{p}_D(data), F_{Y,X})]),$$

where the expectation $\mathbb{E}_{F_{Y,X}}$ is taken with respect to $data \sim F_{Y,X}$ and $R(\cdot, \cdot)$ is defined in Section 2.

Similarly, for uniform pricing, we consider

$$\mathcal{R}_n^U(\mathcal{F}^U) \equiv \inf_{\check{p}_U \in \check{\mathcal{U}}} \sup_{F_Y \in \mathcal{F}^U} (R(p_U^*, F_Y) - \mathbb{E}_{F_Y} [R(\check{p}_U(data_Y), F_Y)]),$$

where the expectation \mathbb{E}_{F_Y} is taken with respect to $data_Y \sim F_Y$. In particular, our assumptions on \mathcal{F} and \mathcal{F}^U in the following sections ensure that, for any $F_{Y,X} \in \mathcal{F}$, the marginal distribution F_Y associated with $F_{Y,X}$ belongs to \mathcal{F}^U .

In what follows, we derive a lower bound for $\mathcal{R}_n^D(x_0; \mathcal{F})$, $\mathcal{R}_n^D(\mathcal{F})$ and $\mathcal{R}_n^U(\mathcal{F}^U)$, respectively. These lower bounds are algorithm independent and reveal the fundamental information-theoretic limitation of data-based pricing strategies.

3.1 Price discrimination

3.1.1 The class of joint distributions of interest

Let us introduce the class \mathcal{F} of joint distributions of (Y, X) that we consider in $\mathcal{R}_n^D(x_0; \mathcal{F})$ and $\mathcal{R}_n^D(\mathcal{F})$.

Assumption 2. *The set \mathcal{F} of joint distributions satisfies the following conditions.*

- (i) (Lipschitz continuity) *There exists $C_0 \in (0, \infty)$ such that, for any $y, y', x \in [0, 1]$, the conditional density $f_{Y|X}$ satisfies*

$$|f_{Y|X}(y|x) - f_{Y|X}(y'|x)| \leq C_0 |y - y'|.$$

(ii) (Strong concavity) There exists $C^* > 0$ such that the revenue function $r(y, x) \equiv y(1 - F_{Y|X}(y|x))$ is strictly concave with second-order derivative

$$-2f_{Y|X}(y|x) - y \frac{\partial}{\partial y} f_{Y|X}(y|x) \leq -C^*, \text{ a.e.} \quad (2)$$

(iii) (Interior solution) For each $x \in [0, 1]$, the optimal price is an interior solution; that is, $p_D^*(x; F_{Y,X}) \in (0, 1)$.

(iv) (Differentiability) The conditional distribution function $f_{Y|X}(y|x)$ is continuously differentiable in (x, y) in a neighborhood of the curve $\{(x, p_D^*(x; F_{Y,X})) : x \in [0, 1]\}$.

(v) (Boundedness) The functions

$$\left| 2f_{Y|X}(y|x) + y \frac{\partial}{\partial y} f_{Y|X}(y|x) \right| \quad (3)$$

$$\text{and } \left| \frac{\partial}{\partial x} F_{Y|X}(y|x) + y \frac{\partial}{\partial x} f_{Y|X}(y|x) \right| \quad (4)$$

are bounded from above by $\bar{C} \in (0, \infty)$ a.e.

(vi) (Marginal density) The marginal density f_X is bounded from above by $\bar{C}' \in (0, \infty)$ and bounded away from zero; that is, $f_X \geq \underline{C} > 0$.

Part (i) requires the density function to be sufficiently smooth. The partial derivative $\frac{\partial}{\partial y} f_{Y|X}(y|x)$ is well-defined almost everywhere because $f_{Y|X}$ is Lipschitz continuous and hence absolutely continuous. Under part (ii), the optimal price is well-defined. Part (iii) ensures that the first-order condition holds for the optimal price. Part (iv) ensures that the optimal pricing function $p_D^*(x; F_{Y,X})$ is sufficiently smooth in x . Part (v) requires the partial derivatives of the revenue to be bounded. Part (vi) ensures that the covariate does not take vanishing or dominating values.

3.1.2 Lower bounds

The first theorem presents a lower bound (in terms of n , up to a constant) for the revenue difference at a given covariate value x_0 , between any data-based 3PD strategy and the true-distribution optimal 3PD strategy under the worst-case distribution by taking the supremum over \mathcal{F} .

Theorem 1 (Lower bounds for 3PD, deficiency in pointwise revenue). *For any \mathcal{F} satisfying Assumption 2 with $C^* \in (0, 2)$ in (2), the minimax difference in the revenues at a given covariate*

value x_0 is bounded from below as

$$\mathcal{R}_n^D(x_0; \mathcal{F}) \gtrsim n^{-1/2}, \quad x_0 \in (0, 1),$$

if $x_0 n^{1/4} \geq c'$ and $(1 - x_0) n^{1/4} \geq c''$ for some positive universal constants c' and c'' (independent of n and x_0).

The second theorem presents a lower bound (in terms of n , up to a constant) for the difference in expected revenues between any data-based 3PD strategy and the true-distribution optimal 3PD strategy under the worst-case distribution by taking the supremum over \mathcal{F} .

Theorem 2 (Lower bounds for 3PD, deficiency in expected revenue). *For any \mathcal{F} satisfying Assumption 2 with $C^* \in (0, 2)$ in (2), the minimax difference in the expected revenues is bounded from below as*

$$\mathcal{R}_n^D(\mathcal{F}) \gtrsim n^{-1/2}.$$

Remark. By saying $C^* \in (0, 2)$ in the theorems above, we allow $r(y, x)$ associated with an $f_{Y|X}$ to have a second derivative bounded from above by a number smaller than or equal to -2 . To motivate the use of $C^* \in (0, 2)$, suppose $f_{Y|X} = f_Y$ (that is, the valuation and covariate are independent of each other) and f_Y is the uniform distribution on $[0, 1]$, $U[0, 1]$. In this case, the revenue function equals $R(y) = y(1 - y)$, which is twice-differentiable with second-order derivative $R''(y) = -2$ for any $y \in [0, 1]$. In our proof for the lower bounds, $U[0, 1]$ is used as the benchmark distribution.

Theorems 1 and 2 state that, there is an inevitable deficiency, $\text{constant} \cdot n^{-1/2}$, in the revenue from any data-based 3PD strategy relative to the revenue from the true-distribution optimal 3PD strategy in the worst case by taking the supremum over \mathcal{F} .

3.2 Uniform pricing

3.2.1 The class of marginal distributions of interest

Assumption 3. Let \mathcal{F}^U be the set of marginal distributions such that F_Y satisfies parts (i), (ii), and (v)(3) of Assumption 2 with $f_{Y|X}(y|x)$ replaced by $f_Y(y)$. Moreover, the optimal price is an interior solution; that is, $p_U^*(F_Y) \in (0, 1)$. The distribution function $f_Y(y)$ is continuously differentiable in y in a neighborhood of $p_U^*(F_Y)$.

Remark. By defining \mathcal{F}^U in the way above, note that the marginal distribution associated with any joint distribution satisfying (i), (ii) and (v)(3) of Assumption 2 satisfies the counterpart conditions in Assumption 3.

3.2.2 Lower bounds

We have the following theorem for uniform pricing.

Theorem 3. For any \mathcal{F}^U satisfying the conditions in Assumption 3 with $C^* \in (0, 2)$ in (2), the minimax difference in the revenues is bounded from below as

$$\mathcal{R}_n^U(\mathcal{F}^U) \gtrsim n^{-2/3}.$$

Theorem 3 states that there is an inevitable deficiency of at least $\text{constant} \cdot n^{-2/3}$ in the revenue from any data-based uniform pricing strategy relative to the revenue from the true-distribution optimal uniform pricing strategy by taking the supremum over \mathcal{F}^U .

Remark. We prove Theorem 2 by turning the pricing problem into a multiple classification problem that tries to distinguish between M distributions. These distributions need be sufficiently close to each other but yield significantly different optimal prices. The most intricate part of the proof includes constructing the conditional densities for the M distributions (where M grows with n) and bounding the separation between the optimal prices. The desired set of optimal prices in our proof is a packing set where the separation between elements is $\Omega(n^{-1/4})$ with respect to the unweighted L_2 norm, and the packing number at that separation is $\Omega(n^{1/4})$. In contrast, proofs for Theorems 1 and 3 are relatively straightforward, as they only require constructing two (instead of M) distributions and simpler techniques.

4 Achievability of the information-theoretic limitation

In the previous section, we demonstrate the information-theoretic limitations of any data-based pricing strategies in the realm of 3PD and uniform pricing, respectively. In this section, we propose specific data-based pricing strategies and show that their performance bounds match the previous lower bounds up to constant factors (that are independent of n). This in turn shows that our lower bounds are tight in terms of rates.

4.1 The K -markets ERM strategy for price discrimination

We propose the “ K -markets” ERM strategy for the data-based 3PD problem with a continuous covariate:

1. Divide the individuals into $K(\equiv K_n)$ markets by splitting the covariate space $[0, 1]$ into K equally spaced intervals

$$I_k \equiv [(k-1)/K, k/K), k = 1, \dots, K$$

such that each interval has at least one observation.

2. For each market I_k , based on the empirical distribution of $\{Y_i : X_i \in I_k\}$,

$$\hat{F}_k(\cdot) = \frac{1}{n_k} \sum_{i \in \{j: X_j \in I_k\}} \mathbf{1}\{Y_i \leq \cdot, X_i \in I_k\}$$

where n_k is the cardinality of $\{j : X_j \in I_k\}$, solve for the optimal price \hat{p}_k as follows,

$$\hat{p}_k \equiv \operatorname{argmax}_p p(1 - \hat{F}_k(p)). \quad (5)$$

The resulting pricing function is a piece-wise function

$$\hat{p}(x; \text{data}) = \hat{p}_k, x \in I_k.$$

Theorem 4. Recall \mathcal{F} given in Section 3.1.1. Assume that $K \leq n$.

- (i) At a given covariate value x_0 , the revenue generated by our proposed pricing strategy satisfies

$$\sup_{F_{Y,X} \in \mathcal{F}} (r(p_D^*, x_0) - \mathbb{E}_{F_{Y,X}} [r(\hat{p}_D(\text{data}), x_0)]) \lesssim (K/n)^{2/3} + 1/K^2, \quad x_0 \in I_k;$$

moreover,

$$(K/n)^{2/3} + 1/K^2 \asymp n^{-1/2} \text{ when } K \asymp n^{1/4}.$$

- (ii) The expected revenue generated by our proposed pricing strategy satisfies

$$\sup_{F_{Y,X} \in \mathcal{F}} (R(p_D^*, F_{Y,X}) - \mathbb{E}_{F_{Y,X}} [R(\hat{p}_D, F_{Y,X})]) \lesssim (K/n)^{2/3} + 1/K^2;$$

moreover,

$$(K/n)^{2/3} + 1/K^2 \asymp n^{-1/2} \text{ when } K \asymp n^{1/4}.$$

In the above, the deficiency in revenues comes from two sources. The first part $(K/n)^{2/3}$ is the “variance” component or “estimation” error, which is due to the randomness of the finite sample, leaving $\hat{F}_k(\cdot)$ different from its expectation. The second part $1/K^2$ is the approximation error due to the fact that we set the same price for all covariate values in the market I_k .

Despite the simplicity of the K -markets ERM strategy and in view of Theorems 1 and 2, Theorem 4 shows that the revenue from this algorithm achieves the optimal rate of convergence (as a function of n) to the revenue from the $F_{Y,X}$ -optimal 3PD strategy. In other words, more sophisticated pricing strategies (e.g., with partitioning the covariate space based on observed frequencies) can *at best* improve upon the K -market ERM algorithm in the constant (which is independent of n).

4.2 The uniform (1-market) ERM strategy for uniform pricing

Based on the empirical distribution of $\{Y_i\}_{i=1}^n$

$$\hat{F}(\cdot) = \frac{1}{n} \sum_{i=1}^n \mathbf{1}\{Y_i \leq \cdot\},$$

the uniform ERM strategy simply solves the optimal price \hat{p}_U as follows:

$$\hat{p}_U(\text{data}_Y) \equiv \operatorname{argmax}_{p \in [0,1]} p(1 - \hat{F}(p)). \quad (6)$$

Theorem 5. Recall \mathcal{F}^U given in Section 3.2.1. The revenue generated by our proposed pricing strategy satisfies

$$\sup_{F_Y \in \mathcal{F}^U} (R(p_U^*, F_Y) - \mathbb{E}_{F_Y} [R(\hat{p}_U(\text{data}_Y), F_Y)]) \lesssim n^{-2/3}.$$

Despite the simplicity of the uniform ERM strategy and in view of Theorem 3, Theorem 5 shows that the revenue from this algorithm achieves the optimal rate of convergence (as a function of n) to the revenue from the F_Y -optimal uniform pricing strategy.

4.3 Welfare analysis

From the policy-maker perspective, it is also of interest to study the welfare under the specific pricing strategies in Sections 4.1 and 4.2. In this section, we derive the rate at which the welfare generated by these data-based pricing strategies converges to the welfare generated by their respective true-distribution optimal strategies.

We assume that there is no production cost for the seller, and there is no utility for the seller if the item is not sold. These assumptions are typically imposed in a benchmark model for auction and pricing literature. For any pricing strategy $p \in \mathcal{D}$, its welfare can be written as

$$W(p, F_{Y,X}) \equiv \mathbb{E}_{F_{Y,X}}[Y \mathbf{1}\{Y > p(X)\}].$$

Theorem 6.

(i) Recall \mathcal{F} given in Section 3.1.1. Take $K \asymp n^{-1/4}$ in the “ K -markets” ERM strategy. Then

$$\sup_{F_{Y,X} \in \mathcal{F}} \mathbb{E}_{F_{Y,X}} |W(\hat{p}_D(data), F_{Y,X}) - W(p_D^*, F_{Y,X})| \lesssim n^{-1/4}.$$

(ii) Recall \mathcal{F}^U given in Section 3.2.1. Then

$$\sup_{F_Y \in \mathcal{F}^U} \mathbb{E}_{F_Y} |W(\hat{p}_U(data_Y), F_Y) - W(p_U^*, F_Y)| \lesssim n^{-1/3}.$$

5 Simulation and empirical studies

In this section, we present numerical evidence to illustrate the implications of our theoretical results. Specifically, we calculate the revenues of the pricing strategies proposed in Section 4 using real-world and simulated data. We describe the data in detail below.

Data. For the empirical study, we use an eBay auction data set (Jank and Shmueli, 2010). Because eBay uses a sealed-bid second-price auction format, the bid of each participant can serve as a proxy for an individual valuation of the object. In particular, we use the data on 194 7-day auctions for the new Palm Pilot M515 PDAs.⁹ The data has 3,832 observations at the bid level,

⁹Jank and Shmueli (2010) also provide data on Cartier wristwatches, Swarovski beads, and Xbox game consoles, but each of these data sets may pool various configurations or models of these products categories. Thus, we choose the data on the Palm Pilot M515 to minimize such variations.

and each observation includes an auction id, a bid amount, a bidder id, and a bidder rating. Some bidders appear in the data set several times because either they revised their bid during an auction or participated in several auctions. To be consistent with our assumption of independent sampling, we analyze the data at the bidder level and use the highest bid of each bidder across all auctions she participated in as the one representing her valuation. This leaves 1,203 observations from which we draw samples of various sizes. For Y_i , we use the bid (as described above) of bidder i normalized to $[0, 1]$. For X_i in the 3PD case, we use bidder i 's rating on eBay, which indicates the number of times sellers left feedback after a transaction with i .

For the simulation study, we let the marginal distribution of X be uniform on $[0, 1]$ and the CDF of Y conditional on $X = x$ be

$$F(y|x) = y^{x+1}.$$

Implementation. For each type of data, we calculate (a Monte-Carlo approximation of) the expected revenue generated by the uniform ERM and the K -markets ERM strategies for various sample sizes as follows. First, fix n and K . Then, draw a sample $\{Y_i, X_i\}_{i=1}^n$ and, for each $k = 1, \dots, K$, define the k -th market as

$$\text{market}_k \equiv \{Y_i : X_i \in [(k-1)/K, k/K)\}, \quad n_k \equiv |\text{market}_k|, \quad \hat{F}_k(t) \equiv \frac{|\{Y_i \in \text{market}_k : Y_i \leq t\}|}{n_k}$$

where the function $\lfloor z \rfloor$ denotes the largest integer smaller than or equal to z , and when applied to a set, $|\cdot|$ denotes the cardinality of the set. Because the K -markets ERM strategy will use the data in each market to determine the optimal price, we need to guarantee that each market is non-empty. For this reason, when calculating the revenue corresponding to the K -markets ERM strategy, we use the highest $\tilde{K} \leq K$ such that $n_k \neq 0$ for all k . In what follows, with a slight abuse of notation, we simply replace \tilde{K} with K for simplicity.

Then, the empirical optimal price in the k -th market is given by

$$\hat{p}_k \equiv \operatorname{argmax}_{y \in [0,1]} y(1 - \hat{F}_k(y)) = \operatorname{argmax}_{y \in \text{market}_k} y(1 - \hat{F}_k(y)),$$

where the second equality holds because \hat{F}_k is a step-function. Note that the uniform ERM strategy simply corresponds to the $K = 1$ -market ERM strategy. Finally, we compute the revenue deficiency for the uniform ERM and K -markets ERM strategies (under $K \asymp n^{1/4}$).

Figure 1 plots the expected revenue generated by the K -markets ERM strategy for $K \in \{1, \dots, 5\}$ as a function of the sample size n (with $K = 1$ corresponding to the uniform ERM strategy). To facilitate the exposition, we use a logarithmic scale for the n -axis. For both types of data, one can see that for sufficiently small n , the K -markets revenue is *decreasing* in K . As n grows, the performance of higher K improves faster than that of lower K , and, for sufficiently large n , the K -markets revenue overtakes that with any lower K . This finding can be explained by the bound $(K/n)^{2/3} + 1/K^2$ in Theorem 4(ii), which implies that higher K (more discrimination) approximates the revenue generated by the $F_{Y,X}$ -optimum better but incurs a larger estimation error ("variance"). When the sample size is small, a lower K would be more beneficial and vice versa.

Figure 1 also suggests that the uniform ERM strategy can be revenue-superior to any $K(> 1)$ -markets ERM strategy when n is sufficiently small. Recall from Theorem 4 that the bound $(K/n)^{2/3} + 1/K^2$ is minimized at $K = n^{1/4}$, which gives $n^{-1/2}$, the optimal rate of convergence to the revenue generated by the $F_{Y,X}$ -optimal 3PD strategy. This convergence rate is slower than $n^{-2/3}$, the optimal rate of convergence to the revenue generated by the F_Y -optimal uniform pricing strategy (c.f. Theorem 5). The slower convergence rate of (even) the best K -market ERM strategy can potentially dominate the revenue gain from price discrimination over without discrimination at all for small n .

Figure 2 illustrates the difference in the rates of convergence of the uniform ERM and the K -markets ERM strategies to their respective theoretical benchmarks. In particular, we set $K = \max\{1, \lfloor 2n^{1/4} - 7 \rfloor\}$ for the empirical study and $K = \frac{1}{5} \lfloor n^{1/4} \rfloor$ for the simulation study. As predicted by the rate $n^{-1/2}$ in Theorem 4 and the rate $n^{-2/3}$ in Theorem 5, the revenue from the uniform ERM strategy is converging to the revenue from the F_Y -optimal uniform pricing strategy faster than the K -markets revenue - to the revenue from the $F_{Y,X}$ -optimal 3PD strategy.

6 Conclusion

Data-based price-discrimination presents a clear economic trade-off. More flexible pricing under finer market segmentation comes at a cost of requiring more data in each market segment. In particular, our findings suggest that data-based price discrimination might be revenue-inferior to data-based uniform pricing when the sample size is sufficiently small. This insight may be viewed as being similar in spirit to that of the data-based treatment assignment literature ([Man-](#)

Figure 1: Revenue under uniform and K -markets ERM strategies


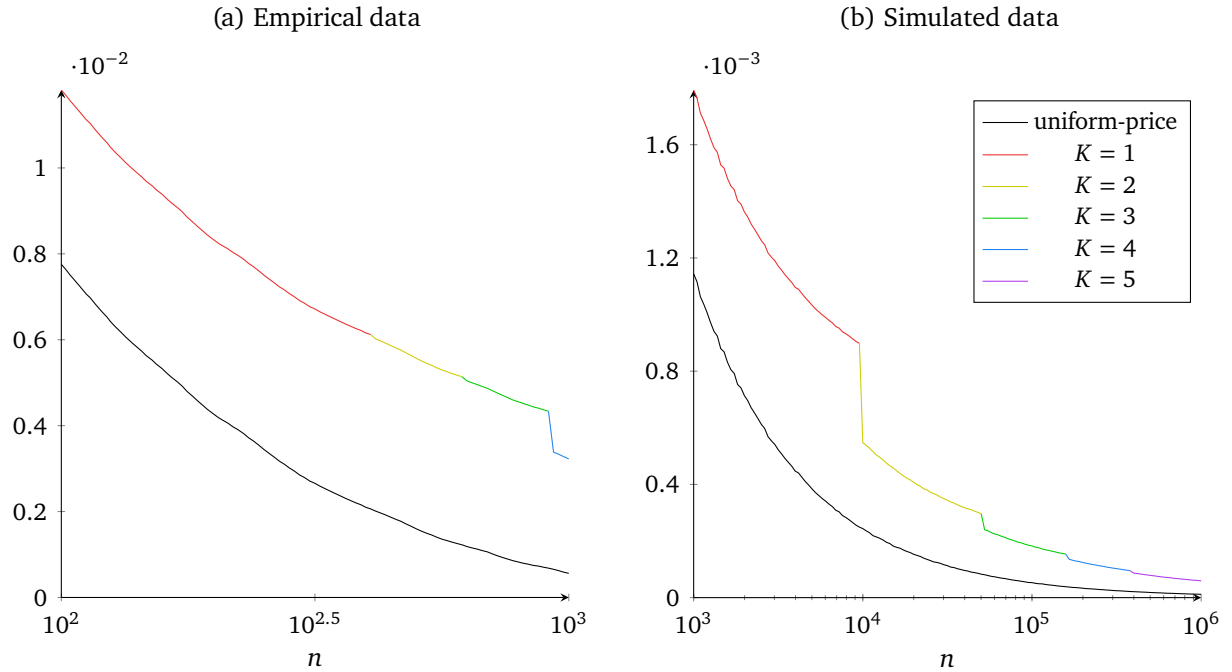


Figure 2: Data-based revenue deficiency under uniform and K -markets ERM strategies (with $K \asymp n^{1/4}$).



ski, 2004; Stoye, 2009, 2012), where the decision-maker chooses from a discrete set of treatment levels. This line of work advocates for treatment rules conditioning on covariates when there is

a substantial amount of data because then, conditioning could be *welfare-regret*-superior. On the other hand, Figure 1 suggests that, in order for conditioning (i.e. price discrimination in our context) to be *revenue*-superior, the sample size needs to be large.

In principle, the machinery that enables the theoretical analyses in this paper can be extended to incorporate multiple covariates. Such a generalization is arduous while carrying similar implications. In particular, we conjecture that with a d -dimensional covariate (such that d is small relative to n), the optimal rate of convergence with respect to the revenue deficiency in the 3PD case is $n^{-\frac{2}{3+d}}$.

A Proofs for lower bounds

A.1 Preliminaries

A.1.1 Preliminary of the proof for Theorem 1

For Theorem 1, we first show that the minimax difference in price at a given covariate value x_0 is bounded from below as follows:

$$\inf_{\check{p}_D \in \check{\mathcal{D}}} \sup_{F_{Y,X} \in \mathcal{F}} \mathbb{E}_{F_{Y,X}} |\check{p}_D(x_0; data) - p_D^*(x_0; F_{Y,X})| \gtrsim n^{-1/4}, x_0 \in (0, 1). \quad (7)$$

Using Taylor expansion type of arguments and condition (2), we can relate the revenue difference to the minimax squared difference in price at x_0 :

$$\begin{aligned} \mathcal{R}_n^D(x_0; \mathcal{F}) &\gtrsim \inf_{\check{p}_D \in \check{\mathcal{D}}} \sup_{F_{Y,X} \in \mathcal{F}} \mathbb{E}_{F_{Y,X}} [|\check{p}_D(x_0; data) - p_D^*(x_0; F_{Y,X})|^2] \\ &\geq \inf_{\check{p}_D \in \check{\mathcal{D}}} \sup_{F_{Y,X} \in \mathcal{F}} \left\{ \mathbb{E}_{F_{Y,X}} [|\check{p}_D(x_0; data) - p_D^*(x_0; F_{Y,X})|] \right\}^2 \end{aligned}$$

where the last line follows from the Jensen's inequality.

The derivation of the lower bound (7) can be reduced to a binary classification problem. In a binary classification problem, we have two distributions $F_{Y,X}^1, F_{Y,X}^2 \in \mathcal{F}$ whose optimal prices are separated by some number 2ε ; that is,

$$|p_D^*(x_0; F_{Y,X}^{j'}) - p_D^*(x_0; F_{Y,X}^j)| \geq 2\varepsilon, \quad j, j' \in \{1, 2\}. \quad (8)$$

A binary classification rule uses the data to decide whether the true distribution is $F_{Y,X}^1$ or $F_{Y,X}^2$. To relate the binary classification problem to the pricing problem, note that, given any pricing function \check{p}_D , we can use it to distinguish between $F_{Y,X}^1$ and $F_{Y,X}^2$ in the following way. Define the binary classification rule

$$\psi(data) = \arg \min_{j \in \{1, 2\}} |p_D^*(x_0; F_{Y,X}^j) - \check{p}_D(x_0; data)|.$$

We claim that when the underlying distribution is $F_{Y,X}^j$, the decision rule ψ is correct if

$$|p_D^*(x_0; F_{Y,X}^j) - \check{p}_D(x_0; data)| < \varepsilon. \quad (9)$$

To see this, note that by the triangle inequality, (8) and (9) guarantee that

$$\begin{aligned}
& |p_D^*(x_0; F_{Y,X}^{j'}) - \check{p}_D(x_0; data)| \\
& \geq |p_D^*(x_0; F_{Y,X}^{j'}) - p_D^*(x_0; F_{Y,X}^j) - \check{p}_D(x_0; data)| \\
& > 2\varepsilon - \varepsilon = \varepsilon, \text{ where } j' \neq j, j, j' \in \{1, 2\}.
\end{aligned}$$

This implies that

$$\mathbb{P}_{F_{Y,X}^j}(\psi(data) \neq j) \leq \mathbb{P}_{F_{Y,X}^j}(|p_D^*(x_0; F_{Y,X}^j) - \check{p}_D(x_0; data)| \geq \varepsilon), \quad j = 1, 2.$$

Therefore, we can upper bound the average probability of mistakes in the binary classification problem as

$$\begin{aligned}
& \frac{1}{2} \mathbb{P}_{F_{Y,X}^1}(\psi(data) \neq 1) + \frac{1}{2} \mathbb{P}_{F_{Y,X}^2}(\psi(data) \neq 2) \\
& \leq \frac{1}{2} \mathbb{P}_{F_{Y,X}^1}(|p_D^*(x_0; F_{Y,X}^1) - \check{p}_D(x_0; data)| \geq \varepsilon) + \frac{1}{2} \mathbb{P}_{F_{Y,X}^2}(|p_D^*(x_0; F_{Y,X}^2) - \check{p}_D(x_0; data)| \geq \varepsilon) \\
& \leq \sup_{F_{Y,X} \in \mathcal{F}} \mathbb{P}_{F_{Y,X}}(|p_D^*(x_0; F_{Y,X}) - \check{p}_D(x_0; data)| \geq \varepsilon).
\end{aligned}$$

By the Markov inequality, we have

$$\begin{aligned}
& \sup_{F_{Y,X} \in \mathcal{F}} \mathbb{E}|\check{p}_D(x_0; data) - p_D^*(x_0; F_{Y,X})| \\
& \geq \varepsilon \sup_{F_{Y,X} \in \mathcal{F}} \mathbb{P}(|\check{p}_D(x_0; data) - p_D^*(x_0; F_{Y,X})| \geq \varepsilon) \\
& \geq \varepsilon \left(\frac{1}{2} \mathbb{P}_{F_{Y,X}^1}(\psi(data) \neq 1) + \frac{1}{2} \mathbb{P}_{F_{Y,X}^2}(\psi(data) \neq 2) \right).
\end{aligned}$$

Finally, we take the infimum over all pricing strategies on the left-hand side (LHS), and the infimum over the induced set of binary decisions on the right-hand side (RHS). This leads to

$$\begin{aligned}
& \inf_{\check{p}_D \in \check{\mathcal{D}}} \sup_{F_{Y,X} \in \mathcal{F}} \mathbb{E}|\check{p}_D(x_0; data) - p_D^*(x_0; F_{Y,X})| \\
& \geq \varepsilon \inf_{\psi} \left(\frac{1}{2} \mathbb{P}_{F_{Y,X}^1}(\psi(data) \neq 1) + \frac{1}{2} \mathbb{P}_{F_{Y,X}^2}(\psi(data) \neq 2) \right). \tag{10}
\end{aligned}$$

The RHS of the above inequality consists of two parts: (1) ε , related to the separation between two optimal prices, and (2) the average probability of making a mistake in distinguishing the

two distributions. To obtain a meaningful bound, we want to find two distributions $F_{Y,X}^1$ and $F_{Y,X}^2$ that are close to each other (hard to distinguish) but their optimal prices are sufficiently separated. We leave the details of the construction of such distributions to the *proof of Theorem 1* given in Appendix A.2.

A.1.2 Preliminary of the proof for Theorem 2

For Theorem 2, we first show that the minimax (unweighted) L_2 -distance in price is bounded from below as follows:

$$\inf_{\check{p}_D \in \check{\mathcal{D}}} \sup_{F_{Y,X} \in \mathcal{F}} \mathbb{E} \|\check{p}_D(data) - p_D^*(F_{Y,X})\|_2^2 \gtrsim n^{-1/2}$$

where

$$\|\check{p}_D(data) - p_D^*(F_{Y,X})\|_2^2 = \int_0^1 |\check{p}_D(x; data) - p_D^*(x; F_{Y,X})|^2 dx.$$

Using Taylor expansion type of arguments and condition (2), we can relate the difference in the *expected* revenues to the minimax (unweighted) L_2 -distance in price:

$$\mathcal{R}_n^D(\mathcal{F}) \gtrsim \inf_{\check{p}_D \in \check{\mathcal{D}}} \sup_{F_{Y,X} \in \mathcal{F}} \mathbb{E}_{F_{Y,X}} \|\check{p}_D(data) - p_D^*(F_{Y,X})\|_2^2$$

where the expectation $\mathbb{E}_{F_{Y,X}}$ is taken with respect to $data \sim F_{Y,X}$.

The object above concerns the entire pricing function $p_D^*(\cdot; F_{Y,X})$. As a result, bounding the RHS of the above inequality is more complicated than the previous one (7). In particular, we consider a multiple classification problem that tries to distinguish among M distributions, where M is a function of the sample size n . Similar as before, we want the optimal prices of these M distributions to be sufficiently separated. Similar derivations show that the lower bound of the revenue problem can be reduced to that of a multiple classification problem:

$$\inf_{\check{p}_D \in \check{\mathcal{D}}} \sup_{F_{Y,X} \in \mathcal{F}} \mathbb{E}_{F_{Y,X}} \|\check{p}_D(data) - p_D^*(F_{Y,X})\|_2^2 \geq \varepsilon^2 \inf_{\psi} \frac{1}{M} \sum_{j=1}^M \mathbb{P}_{F_{Y,X}^j}(\psi(data) \neq j), \quad (11)$$

where the infimum \inf_{ψ} is taken over the set of all multiple decisions (with M choices). To proceed, we apply the Fano's inequality from information theory (Cover and Thomas, 2005).

Fano's inequality gives a lower bound on the average probability of mistakes:¹⁰

$$\frac{1}{M} \sum_{j=1}^M \mathbb{P}_{F_{Y,X}^j}(\psi(data) \neq j) \geq 1 - \frac{\sum_{j,j'=1}^M \text{KL}(F_{Y,X}^j \| F_{Y,X}^{j'}) / M^2 + \log 2}{\log M}, \quad (12)$$

where $\text{KL}(\cdot \| \cdot)$ denotes the Kullback-Leibler (KL) divergence between two distributions:

$$\text{KL}(F_1 \| F_2) \equiv \int f_1(y, x) \log \frac{f_1(y, x)}{f_2(y, x)} dy dx.$$

To obtain a sharp bound based on the multiple classification problem, we want to find a set of distributions (where the cardinality M of the set is large enough) that are close enough to each other (small enough pairwise KL divergence) but their optimal prices are sufficiently separated. We leave the detailed proof to Appendix A.2. Our proof is based on a delicate construction of conditional densities along with an application of the Gilbert-Varshamov Lemma from coding theory.

A.1.3 Preliminary of the proof for Theorem 3

Relative to the proofs in the case of 3PD, the proofs for the price and revenue lower bounds in uniform pricing are simpler. We first show that the minimax difference in price is bounded from below as follows:

$$\inf_{\check{p}_U \in \check{\mathcal{U}}} \sup_{F_Y \in \mathcal{F}^U} \mathbb{E}_{F_Y} |\check{p}_U(data_Y) - p_U^*| \gtrsim n^{-1/3}. \quad (13)$$

As previously, we can relate the revenue difference to the minimax squared difference in price:

$$\begin{aligned} \mathcal{R}_n^U(\mathcal{F}^U) &\gtrsim \inf_{\check{p}_U \in \check{\mathcal{U}}} \sup_{F_Y \in \mathcal{F}^U} \mathbb{E}_{F_Y} [|\check{p}_U(data_Y) - p_U^*|^2] \\ &\geq \inf_{\check{p}_U \in \check{\mathcal{U}}} \sup_{F_Y \in \mathcal{F}^U} \left\{ \mathbb{E}_{F_Y} [|\check{p}_U(data_Y) - p_U^*|] \right\}^2 \end{aligned}$$

where the last line follows from the Jensen's inequality. The derivation of (13) only requires constructing two distributions, similar to the approach discussed in Section A.1.1.

¹⁰We do not present the Fano's inequality in its standard form as in [Cover and Thomas \(2005\)](#). Instead, we use a version from [Wainwright \(2019\)](#) that is more convenient for our purposes.

A.2 Proofs

Proof of Theorem 1. For Theorem 1, we use Lemma C.4 to prove the lower bound. Define

$$\omega_D(\epsilon) \equiv \sup_{F_1, F_2 \in \mathcal{F}} \left\{ |p_D^*(x_0; F_1) - p_D^*(x_0; F_2)| : H(F_1 \| F_2) \leq \epsilon \right\}.$$

By Lemma C.4, we have

$$\inf_{\check{p}_D \in \check{\mathcal{D}}} \sup_{F_{Y,X} \in \mathcal{F}} \mathbb{E}_{F_{Y,X}} |\check{p}_D(x_0; data) - p_D^*(x_0)| \geq \frac{1}{8} \omega_D \left(1/(2\sqrt{n}) \right).$$

Therefore, we only need to find a lower bound for ω_D . Based on the explanation in Appendix A.1.1, we want to construct two distributions that are hard to distinguish but their optimal prices are well-separated. We start by defining two perturbation functions. Let ϕ_Y be defined as

$$\phi_Y(t) \equiv \begin{cases} t+1, & t \in [-1, 0], \\ -t+1, & t \in [0, 2], \\ t-3, & t \in [2, 3], \\ 0, & \text{otherwise.} \end{cases} \quad (14)$$

Notice that ϕ_Y is Lipschitz continuous on \mathbb{R} . Let ϕ_X be defined as


$$\phi_X(t) \equiv \begin{cases} e^{-(4t-1)^2/(1-(4t-1)^2)}, & t \in (0, 1/2), \\ -e^{-(4t-3)^2/(1-(4t-3)^2)}, & t \in (1/2, 1), \\ 0, & \text{otherwise.} \end{cases}$$

Notice that ϕ_X is infinitely differentiable on \mathbb{R} . We plot the two perturbation functions in Figure A.1.

Now we construct the two distributions. Let $\delta \in (0, 1/4)$ be a small number (that depends on n) to be specified later. Let a be any number in the interval $(0, 4 - 2C^*)$. Define the two conditional density functions of Y given X as

$$\begin{aligned} f_1(y|x) &\equiv 1, \\ f_2(y|x) &\equiv 1 + a\delta\phi_Y\left(\frac{y-1/2}{\delta}\right)\phi_X\left(\frac{x-x_0}{\delta} + 1/4\right). \end{aligned} \quad (15)$$

Figure A.1: Perturbation functions ϕ_Y and ϕ_X .



We let the marginal distribution $f_X(x)$ of X be the uniform distribution on $[0, 1]$. Note that $f_1(y|x)$, $f_2(y|x)$, $f_1(y, x) = f_1(y|x)f_X(x)$, and $f_2(y, x) = f_2(y|x)f_X(x)$ are non-negative everywhere, with integrals over their respective entire spaces all equaling to 1.

The first task is to verify that the two distributions are indeed in the class \mathcal{F}_κ . For $C^* \in (0, 2)$, the first distribution is in \mathcal{F} by Lemma C.2 and the fact that Y is independent of X . Given any $x \in [0, 1]$, we can treat the whole term $a\phi_X((x - x_0)/\delta + 1/4)$ as the coefficient b in Lemma C.3. Then the results of Lemma C.3 applies since $|\phi_X| \leq 1$. In particular, the revenue function at x is twice-differentiable a.e., the absolute value of the second-order partial derivative with respect to y is bounded, and is also bounded from below by C^* . The optimal price is an interior solution and is in the interior of a region on which the revenue function is twice-differentiable. Lastly, the absolute value of the partial derivative of $f_2(y|x)$ with respect to x is bounded. This ensures that the quantity $|\frac{\partial}{\partial x} F_{Y|X}(y|x) + y \frac{\partial}{\partial x} f_{Y|X}(y|x)|$ is bounded.

Next, we want to derive the Hellinger distance between the two joint densities

$$f_1(y, x) = 1,$$

$$f_2(y, x) = 1 + a\delta\phi_Y\left(\frac{y - 1/2}{\delta}\right)\phi_X\left(\frac{x - x_0}{\delta} + 1/4\right).$$

Let $\Psi(t) \equiv \sqrt{1+t}$. Its second-order derivate is bounded when $|t| < 1/2$; that is,

$$\sup_{|t| < 1/2} |\Psi''(t)| < C.$$

We use H to denote the Hellinger distance:

$$H(f_1 \| f_2)^2 \equiv \int_0^1 \left(\sqrt{f_1(y)} - \sqrt{f_2(y)} \right)^2 dy.$$

The Hellinger distance can be bounded as

$$\begin{aligned} H^2(f_1 \| f_2)/2 &= 1 - \int_0^1 \int_0^1 \Psi \left(a\delta \phi_Y \left(\frac{y-1/2}{\delta} \right) \phi_X \left(\frac{x-x_0}{\delta} + 1/4 \right) \right) dy dx \\ &= \int_0^1 \int_0^1 \Psi(0) - \Psi \left(a\delta \phi_Y \left(\frac{y-1/2}{\delta} \right) \phi_X \left(\frac{x-x_0}{\delta} + 1/4 \right) \right) dy dx \\ &\leq -a\Psi'(0)\delta \int_0^1 \int_0^1 \phi_Y \left(\frac{y-1/2}{\delta} \right) \phi_X \left(\frac{x-x_0}{\delta} + 1/4 \right) dy dx \\ &\quad + a^2 C \delta^2 \int_0^1 \int_0^1 \phi_Y^2 \left(\frac{y-1/2}{\delta} \right) \phi_X^2 \left(\frac{x-x_0}{\delta} + 1/4 \right) dy dx, \end{aligned}$$

where we have applied the second-order Taylor expansion to obtain the last inequality. By the change of variables $u = (y-1/2)/\delta$ and $v = (x-x_0)/\delta + 1/4$, for sufficiently small $\delta \in (0, 1/2]$,

$$\int_0^1 \int_0^1 \phi_Y \left(\frac{y-1/2}{\delta} \right) \phi_X \left(\frac{x-x_0}{\delta} + 1/4 \right) dy dx = \delta^2 \int_{-1}^1 \phi_Y(u) du \int_0^1 \phi_X(v) dv = 0, \quad (16)$$

and

$$\int_0^1 \int_0^1 \phi_Y^2 \left(\frac{y-1/2}{\delta} \right) \phi_X^2 \left(\frac{x-x_0}{\delta} + 1/4 \right) dy dx = \delta^2 \int_{-1}^1 \phi_Y^2(u) du \int_0^1 \phi_X^2(v) dv \leq C\delta^2.$$

Therefore, the Hellinger distance is bounded as

$$H^2(f_1 \| f_2) \lesssim \delta^4.$$

Now we take δ such that $\delta^4 \asymp 1/n$. Note that (16) holds when $\delta > 0$ is small enough such that $\delta \in (0, 1/2]$, $1/4 - x_0/\delta \leq 0$ and $(1 - x_0)/\delta + 1/4 \geq 1$; that is, when $x_0 n^{1/4} \geq c'$ and $(1 - x_0)n^{1/4} \geq c''$ for positive universal constants c' and c'' (independent of n and x_0). This ensures that $H^2(f_1 \| f_2) \lesssim 1/n$. Then from Lemma C.4, we know that

$$\inf_{\check{p}_D \in \check{\mathcal{D}}} \sup_{F_{Y,X} \in \mathcal{F}} \mathbb{E}_{F_{Y,X}} |\check{p}_D(x_0; data) - p_D^*(x_0)| \gtrsim n^{-1/4}, x_0 \in (0, 1).$$

For bounding the revenue, recall that the revenue achieved at the price p and covariate value x_0

is $r(p, x_0) = \max_p p(1 - F_{Y|X}(p|x_0))$. By Lemma C.1, we have

$$r(p_D^*(x_0)) - r(\check{p}_D(x_0; data)) \geq \frac{C^*}{2} |p_D^*(x_0) - \check{p}_D(x_0; data)|^2.$$

As a result, we have

$$\begin{aligned} & \inf_{\check{p}_D \in \check{\mathcal{D}}} \sup_{F_{Y,X} \in \mathcal{F}} \mathbb{E}[r(p_D^*, x_0) - r(\check{p}_D(data), x_0)] \\ & \geq \inf_{\check{p}_D \in \check{\mathcal{D}}} \sup_{F_{Y,X} \in \mathcal{F}} \mathbb{E} \left[\frac{C^*}{2} |p_D^*(x_0) - \check{p}_D(x_0; data)|^2 \right] \\ & \geq \inf_{\check{p}_D \in \check{\mathcal{D}}} \sup_{F_{Y,X} \in \mathcal{F}} \frac{C^*}{2} \left\{ \mathbb{E} [|p_D^*(x_0) - \check{p}_D(x_0; data)|] \right\}^2 \gtrsim n^{-1/2}. \end{aligned}$$

This proves Theorem 1. □

Proof of Theorem 2. To prove Theorem 2, we follow the explanation in Appendix A.1.2 and use the Fano's inequality to bound the probability of mistakes in the multiple classification problem. Before solving the revenue problem, we first study the lower bound for the L_2 -distance of pricing functions. For two pricing functions p_1 and p_2 , we define the (unweighted) L_2 -distance as

$$\|p_1 - p_2\|_2 \equiv \left(\int_0^1 |p_1(x) - p_2(x)|^2 dx \right)^{1/2}.$$

In part (i), we defined the perturbation on the X dimension at a fixed point x_0 . Now we want to define a large set of perturbed distributions. Each of these distributions is perturbed in a small interval on the X dimension. Let $m \geq 8$ be a large number (depending on n) that we specify later. Let $\alpha \in \{0, 1\}^m$ be a vector of length m ; that is,

$$\alpha \equiv (\alpha_1, \dots, \alpha_m), \text{ where } \alpha_j \in \{0, 1\}, j = 1, \dots, m.$$

We construct a set of conditional density functions indexed by α :

$$f_{Y|X}^\alpha(y|x) \equiv 1 + \frac{a}{m} \sum_{j=1}^m \alpha_j \phi_Y(m(y - 1/2)) \phi_X(mx - (j - 1)).$$

The marginal distribution of X is taken to be the uniform distribution on $[0, 1]$, that is, $f_X \equiv \mathbf{1}_{[0,1]}$. We denote the joint distribution by $f_{Y,X}^\alpha \equiv f_{Y|X}^\alpha f_X$.

We briefly describe this construction of the conditional density. The unit interval $[0, 1]$ is divided equally into m subintervals:

$$I_j \equiv [(j-1)/m, j/m], j = 1, \dots, m.$$

For $x \in I_j$, if $\alpha_j = 0$, then the conditional density is 1. If $\alpha_j = 1$, then the conditional density

$$f_{Y|X}^\alpha(y|x) \equiv 1 + \frac{a}{m} \phi_Y(m(y - 1/2)) \phi_X(mx - (j-1)), x \in I_j.$$

By treating $1/m$ as the scalar δ in part (i), we can see that, for m large enough, each $f_{Y,X}^\alpha$ belongs to the set \mathcal{F}_κ .

From the set $\{f_{Y,X}^\alpha : \alpha \in \{0, 1\}^m\}$, we want to pick out a large enough subset of distributions whose optimal price functions are well-separated. For this purpose, we use the Gilbert-Varshamov bound (Lemma 2.9, Chapter 2 [Tsybakov, 2009](#)). The Gilbert-Varshamov bound states that for $m \geq 8$, there exists a subset $\mathcal{A} \subset \{0, 1\}^m$ with cardinality $M \equiv |\mathcal{A}| \geq 2^{m/8}$, and the pairwise rescaled Hamming distance between elements in this set is greater than $1/8$. That is,

$$\frac{1}{m} \sum_{j=1}^m \mathbf{1}\{\alpha_j \neq \alpha'_j\} \geq \frac{1}{8}, \text{ for any } \alpha, \alpha' \in \mathcal{A}.$$

Applying the Gilbert-Varshamov bound, we can show that for $\alpha, \alpha' \in \mathcal{A}$, the optimal pricing functions of $f_{Y,X}^\alpha$ and $f_{Y,X}^{\alpha'}$ are well-separated. Let p_α be the pricing function associated with $f_{Y,X}^\alpha$; that is,

$$p_\alpha(x) \equiv \operatorname{argmax}_{p \in [0,1]} p(1 - F_{Y|X}^\alpha(p|x)),$$

where $F_{Y|X}^\alpha(y|x)$ is the corresponding conditional cumulative distribution function. Note that $\alpha, \alpha' \in \mathcal{A}$ differ in at least $m/8$ positions. This means that $f_{Y|X}^\alpha$ and $f_{Y|X}^{\alpha'}$ differ in $m/8$ intervals. Suppose that I_j is such an interval, where $\alpha_j = 0$ and $\alpha'_j = 1$. We restrict our attention to a subset of this interval:

$$\tilde{I}_j \equiv \left[\frac{1}{6m} + \frac{j-1}{m}, \frac{1}{3m} + \frac{j-1}{m} \right] \subset I_j.$$

When $x \in \tilde{I}_j$, we have

$$mx - (j - 1) \in [1/6, 1/3] \implies \phi_X(mx - (j - 1)) \in [\phi_X(0), \phi_X(1/2)]. \quad (17)$$

By Lemma C.3 (where $b = a\phi_X(mx - (j - 1)) > 0$, $\delta = 1/m$), the choice $a \in (0, 4 - 2\kappa)$, and the fact (17), if we fix $x \in \tilde{I}_j$, then $p_\alpha(x) = 1/2$ while

$$p_{\alpha'}(x) \leq 1/2 - \frac{c}{m}\phi_X(mx - (j - 1)) \leq 1/2 - \frac{c\phi_X(1/6)}{m}, x \in \tilde{I}_j,$$

where $c > 0$ is a universal constant that does not depend on n .¹¹ This implies that

$$|p_\alpha(x) - p_{\alpha'}(x)| \gtrsim \frac{1}{m}, x \in \tilde{I}_j.$$

Therefore, on the interval I_j , the separation between p_α and $p_{\alpha'}$ is lower bounded as

$$\int_{I_j} |p_\alpha(x) - p_{\alpha'}(x)|^2 dx \gtrsim \int_{\tilde{I}_j} 1/m^2 dx = \frac{1}{6m} \times \frac{1}{m^2} \gtrsim 1/m^3.$$

By the Gilbert-Varshamov bound, there are at least $m/8$ such intervals. Therefore, we can lower bound the total separation by

$$\|p_1 - p_2\|_2 \gtrsim (m/8 \times 1/m^3)^{1/2} \gtrsim 1/m.$$

Next, we want to compute the KL divergence between $f_{Y,X}^\alpha$ and $f_{Y,X}^{\alpha'}$. Note that the term $\phi_X(mx - (j - 1))$ is non-zero only when $x \in I_j$. The KL divergence can therefore be treated as a sum of m integrals:

$$\text{KL}(f_{Y,X}^\alpha \| f_{Y,X}^{\alpha'}) = \int_0^1 \int_0^1 f_{Y,X}^\alpha(y, x) \log \frac{f_{Y,X}^\alpha}{f_{Y,X}^{\alpha'}} dy dx = \sum_{j=1}^m E_j,$$

where

$$\begin{aligned} E_j \equiv & \int_{I_j} \int_0^1 \left(1 + \frac{a}{m} \alpha_j \phi_Y(m(y - 1/2)) \phi_X(mx - (j - 1)) \right) \\ & \times \log \frac{1 + \frac{a}{m} \alpha_j \phi_Y(m(y - 1/2)) \phi_X(mx - (j - 1))}{1 + \frac{a}{m} \alpha'_j \phi_Y(m(y - 1/2)) \phi_X(mx - (j - 1))} dy dx. \end{aligned}$$

¹¹For example, c can be equal to $a/8$ according to Lemma C.3.

Notice that when $\alpha_j = \alpha_{j'}$, $E_j = 0$. Therefore, we only need to consider the j 's where $\alpha_j \neq \alpha_{j'}$.

Denote $\Psi_1(t) = -\log(1+t)$ and $\Psi_2(t) = (1+t)\log(1+t)$. Then we can write E_j as

$$E_j = \begin{cases} \int_{I_j} \int_0^1 \Psi_1\left(\frac{a}{m}\phi_Y(m(y-1/2))\phi_X(mx-(j-1))\right) dydx, & \text{if } \alpha_j = 0, \alpha_{j'} = 1, \\ \int_{I_j} \int_0^1 \Psi_2\left(\frac{a}{m}\phi_Y(m(y-1/2))\phi_X(mx-(j-1))\right) dydx, & \text{if } \alpha_j = 1, \alpha_{j'} = 0. \end{cases}$$

By the second-order Taylor expansion at zero, we have

$$\Psi_1(t) = -t + \frac{1}{2(1+t')^2}t^2,$$

for some t' between 0 and t . When $|t| \leq 1/4$,¹² we have

$$\Psi_1(t) \leq -t + Ct^2,$$

for some universal constant $C > 0$. Similarly, we can show that

$$\Psi_2(t) \leq t + Ct^2.$$

Applying these inequalities to E_j , we have

$$\begin{aligned} E_j &\leq \pm \int_{I_j} \int_0^1 \frac{a}{m}\phi_Y(m(y-1/2))\phi_X(mx-(j-1)) dydx \\ &\quad + C \int_{I_j} \int_0^1 \frac{a^2}{m^2}\phi_Y^2(m(y-1/2))\phi_X^2(mx-(j-1)) dydx. \end{aligned}$$

Similar to the derivation in Part (i), we know that the first term on the RHS is zero. For the second term, we can apply change of variables $u = m(y-1/2)$ and $v = mx-(j-1)$ and obtain that

$$\begin{aligned} &\int_{I_j} \int_0^1 \phi_Y^2(m(y-1/2))\phi_X^2(mx-(j-1)) dydx \\ &= \frac{1}{m^2} \int_0^1 \phi_X^2(v) dv \int_{-1}^3 \phi_Y^2(u) du \leq \frac{C'}{m^2} \end{aligned}$$

for some universal constant $C' > 0$. Putting the results together, we know that $E_j \leq \frac{C}{m^4}$

¹²Later we show that m is chosen to be $c_0 n^{1/4}$ where $c_0 > 0$ is a universal constant. As a result, $|t| \leq 1/4$ is guaranteed as long as c_0 is sufficiently large.

for all j . Since there are m intervals, we can bound the KL divergence by

$$\text{KL}(f_{Y,X}^\alpha \| f_{Y,X}^{\alpha'}) = \sum_{j=1}^m E_j \lesssim \frac{1}{m^3}.$$

This is the KL distance for a single observation. For the entire data set with n i.i.d. observations, the KL divergence is upper bounded by Cn/m^3 .

Lastly, we can summarize our results into the Fano inequality presented in Lemma C.5. We have

$$\begin{aligned} \inf_{\check{p}_D \in \check{\mathcal{D}}} \sup_{F_{Y,X} \in \mathcal{F}} \mathbb{E} \|\check{p}_D(data) - p_D^*\|_2^2 &\geq \frac{C_1}{m^2} \left(1 - \frac{C_2 n/m^3 + \log 2}{\log 2^{m/8}} \right) \\ &\geq \frac{C_1}{m^2} \left(1 - \frac{C_2 n/m^3 + \log 2}{C_3 m} \right). \end{aligned}$$

By choosing $m = c_0 n^{1/4}$ for a sufficiently large universal constant $c_0 > 0$, we can make the factor $\left(1 - \frac{C_2 n/m^3 + \log 2}{C_3 m} \right)$ stay above, say, $1/2$. Then we have

$$\inf_{\check{p}_D \in \check{\mathcal{D}}} \sup_{F_{Y,X} \in \mathcal{F}} \mathbb{E} \|\check{p}_D(data) - p_D^*\|_2^2 \gtrsim \frac{1}{m^2} \asymp n^{-1/2}.$$

So far we have derived the lower bound for the L_2 -distance of pricing. Moving onto the revenue problem, recall that the revenue achieved at the price p and covariate value x is $r(p, x) = \max_p p(1 - F_{Y|X}(p|x))$. By Lemma C.1, we have

$$r(p_D^*, x) - r(\check{p}_D(data), x) \geq \frac{C^*}{2} |p_D^*(x) - \check{p}_D(x; data)|^2.$$

Since f_X is bounded away from zero, we have

$$\begin{aligned} &\inf_{\check{p}_D \in \check{\mathcal{D}}} \sup_{F_{Y,X} \in \mathcal{F}} \mathbb{E}[R(p_D^*) - R(\check{p}_D)] \\ &= \inf_{\check{p}_D \in \check{\mathcal{D}}} \sup_{F_{Y,X} \in \mathcal{F}} \mathbb{E} \left[\int_0^1 (r(p_D^*, x) - r(\check{p}_D, x)) f_X(x) dx \right] \\ &\geq \inf_{\check{p}_D \in \check{\mathcal{D}}} \sup_{F_{Y,X} \in \mathcal{F}} \mathbb{E} \left[\frac{C^*}{2} \left(\inf_{x \in [0,1]} f_X(x) \right) \int_0^1 |p_D^*(x) - \check{p}_D(x; data)|^2 dx \right] \gtrsim n^{-1/2}. \end{aligned}$$

□

Proof of Theorem 3. We use Lemma C.4 to prove the lower bound for Theorem 3. Define

$$\omega_U(\epsilon) \equiv \sup_{F_1, F_2 \in \mathcal{F}^U} \{ |p_U^*(F_1) - p_U^*(F_2)| : H(F_1 \| F_2) \leq \epsilon \}.$$

Then by Lemma C.4, we have

$$\inf_{\check{p}_U \in \mathcal{U}} \sup_{F_Y \in \mathcal{F}^U} \mathbb{E}_{F_Y} |\check{p}_U(\text{data}_Y) - p_U^*| \geq \frac{1}{8} \omega_U \left(1/(2\sqrt{n}) \right).$$

Therefore, we only need to find a lower bound for ω_U . The proof proceeds in three steps. In the first step, we construct two distributions and compute the separation between their optimal prices. The second step bounds the Hellinger distance between these two distributions. The third step summarizes.

Step 1. We construct two distribution functions. The first distribution is the uniform distribution on the unit interval $[0, 1]$. We denote this density function as

$$f_1(y) = 1_{[0,1]}(y).$$

The distribution function is $F_1(y) = y$ on the support $[0, 1]$. The revenue function under this distribution is $R_1(p) = p(1 - p)$. The optimal price is

$$p_1 = \operatorname{argmax}_{p \in [0,1]} R_1(p) = \operatorname{argmax}_{p \in [0,1]} p - p^2 = 1/2.$$


The second distribution function is a small twist of the uniform distribution. We use the same perturbation function ϕ_Y defined in (14).

We apply a small perturbation to the uniform density. Let $\delta > 0$ be a small number (that depends on n) specified later. Let $a \in (0, 4 - 2C^*)$. The formula of the density f_2 is given by

$$f_2(y) \equiv 1 + a\delta\phi_Y\left(\frac{y - 1/2}{\delta}\right) = \begin{cases} 1, & \text{if } y \in [0, 1/2 - \delta), \\ ay + 1 - \frac{a}{2} + a\delta, & \text{if } y \in [1/2 - \delta, 1/2), \\ -ay + 1 + \frac{a}{2} + a\delta, & \text{if } y \in [1/2, 1/2 + 2\delta), \\ ay + 1 - \frac{a}{2} - 3a\delta, & \text{if } y \in [1/2 + 2\delta, 1/2 + 3\delta), \\ 1, & \text{if } y \in [1/2 + 3\delta, 1]. \end{cases}$$

We compare the two densities f_1 and f_2 in the following graph.

Figure A.2: Density functions f_1 and f_2 .



Denote the optimal price under f_2 by p_2 . By Lemma C.3(ii), we have

$$|p_2 - p_1| \geq a\delta/8$$

when δ is sufficiently small.

Step 2. We want to bound the Hellinger distance $H(F_1||F_2)$. Define the function $\Psi(t) = \sqrt{1+t}$. Its second-order derivative is bounded when $|t| < 1/2$; that is,

$$\sup_{|t| < 1/2} |\Psi''(t)| \leq \frac{\sqrt{2}}{2}.$$

Since $f_1(y) = 1$, we have

$$\begin{aligned} H(F_1||F_2)^2/2 &= 1 - \int_0^1 \Psi\left(a\delta\phi_Y\left(\frac{y-1/2}{\delta}\right)\right) dy \\ &= \int_0^1 \Psi(0) - \Psi\left(a\delta\phi_Y\left(\frac{y-1/2}{\delta}\right)\right) dy. \end{aligned}$$

By the second-order Taylor expansion, we have

$$\begin{aligned} &\Psi(0) - \Psi\left(a\delta\phi_Y\left(\frac{y-1/2}{\delta}\right)\right) \\ &\leq -\Psi'(0)a\delta\phi_Y\left(\frac{y-1/2}{\delta}\right) + \frac{\sqrt{2}}{4}a^2\delta^2\phi_Y^2\left(\frac{y-1/2}{\delta}\right). \end{aligned}$$

By the construction of ϕ_Y , we have

$$\int_0^1 \phi_Y \left(\frac{y - 1/2}{\delta} \right) dy = 0.$$

By the change of variables $u = (y - 1/2)/\delta$, we have

$$\int_0^1 \phi_Y^2 \left(\frac{y - 1/2}{\delta} \right) dy = \delta \int_{\mathbb{R}} \phi_Y^2(u) du \leq 4\delta \int_{-1}^0 (x + 1)^2 dx = \frac{4}{3}\delta.$$

Combining these results together, we obtain a bound on the Hellinger distance

$$H(F_1 \| F_2)^2 \leq \frac{2\sqrt{2}}{3} a^2 \delta^3.$$

Step 3. By setting $\delta = c'_0 (3/8\sqrt{2})^{1/3} a^{-2/3} n^{-1/3}$ for $c'_0 \in (0, 1)$, we can ensure that $H(F_1 \| F_2) \leq$

$1/(2\sqrt{n})$. Previously, we assumed that $a\delta \leq 1/2$ for the second-order Taylor expansion. This is true if c'_0 is chosen to be sufficiently small. In this case, the separation between p_1 and p_2 is lower bounded as below:

$$|p_1 - p_2| \geq a\delta/8 = \frac{c'_0}{16} \left(\frac{3}{\sqrt{2}} \right)^{1/3} \left(\frac{a}{n} \right)^{1/3}.$$

By Lemma C.4, we have

$$\inf_{\check{p}_U \in \check{\mathcal{U}}} \sup_{F_Y \in \mathcal{F}^U} \mathbb{E} |\check{p}_U(\text{data}_Y) - p_U^*| \geq \frac{c'_0}{16} \left(\frac{3}{\sqrt{2}} \right)^{1/3} \left(\frac{a}{n} \right)^{1/3}.$$

Lastly, we want to lower bound the revenue. By Lemma C.1, we have

$$\begin{aligned} \mathcal{R}_n^U(\mathcal{F}^U) &= \inf_{\check{p}_U \in \check{\mathcal{U}}} \sup_{F_Y \in \mathcal{F}^U} \mathbb{E} |R(\check{p}_U(\text{data}_Y), F_Y) - R(p_U^*, F_Y)| \\ &\geq \inf_{\check{p}_U \in \check{\mathcal{U}}} \sup_{F_Y \in \mathcal{F}^U} \mathbb{E} \left[\frac{C^*}{2} |\check{p}_U(\text{data}_Y) - p_U^*|^2 \right] \\ &\geq \inf_{\check{p}_U \in \check{\mathcal{U}}} \sup_{F_Y \in \mathcal{F}^U} \frac{C^*}{2} \{ \mathbb{E} [|\check{p}_U(\text{data}_Y) - p_U^*|] \}^2 \\ &\gtrsim \left(\frac{1}{n} \right)^{2/3}. \end{aligned}$$

□

B Proofs for upper bounds

To facilitate the presentation, we first give the proof for Theorem 5.

Proof of Theorem 5. For simplicity, we omit writing $F_{Y,X}$ and *data* in the proof. Denote $\kappa' \equiv \inf_{p \in [0,1]} |R''(p)|/2 > 0$. By Taylor expansion, for any p ,

$$R(p_U^*) - R(p) \geq \kappa' (p - p_U^*)^2.$$

Denote $\hat{R}(p) \equiv p(1 - \hat{F}(p))$. Combining the inequality above with the basic inequality (i.e., $\hat{R}(\hat{p}_U) \geq \hat{R}(p_U^*)$), we have

$$\kappa' (\hat{p}_U - p_U^*)^2 \leq R(p_U^*) - R(\hat{p}_U) \leq R(p_U^*) - \hat{R}(p_U^*) - (R(\hat{p}_U) - \hat{R}(\hat{p}_U)). \quad (18)$$

For $\delta \in (0, p_U^*]$, define

$$\mathcal{G}_\delta \equiv \{y \mapsto p \mathbf{1}\{y \geq p\} - p_U^* \mathbf{1}\{y \geq p_U^*\} : p \in [p_U^* - \delta, p_U^* + \delta]\}$$

and

$$G_\delta(y) \equiv \begin{cases} 0, & \text{if } y < p_U^* - \delta, \\ p_U^*, & \text{if } p_U^* - \delta \leq y \leq p_U^* + \delta, \\ \delta, & \text{if } y > p_U^* + \delta. \end{cases}$$

Then G_δ is an envelope function of the class \mathcal{G}_δ . The $L_2(P)$ -norm of G_δ is bounded by

$$\|G_\delta\|_{L_2(P)} = ((p_U^*)^2 \mathbb{P}(Y \in [p_U^* - \delta, p_U^* + \delta]) + \delta^2 \mathbb{P}(Y > p_U^* + \delta))^{1/2} \leq C\sqrt{\delta}.$$

Since \mathcal{G}_δ is a VC-subgraph class, we have

$$\mathbb{E} \sup_{g \in \mathcal{G}_\delta} \left| \frac{1}{n} \sum_{i=1}^n g(Y_i) - \mathbb{E} g(Y_i) \right| \leq C\sqrt{\delta/n}. \quad (19)$$

We derive the convergence rate of $\hat{p} - p^*$ via a peeling argument. Consider the following decom-

position

$$\mathbb{P}\left(n^{1/3}|\hat{p}_U - p_U^*| > M\right) = \sum_{j=M+1}^{\infty} \mathbb{P}\left(n^{1/3}|\hat{p}_U - p_U^*| \in (j-1, j]\right).$$

For any $j \geq M+1$, we have

$$\begin{aligned} & \{|\hat{p}_U - p_U^*| \in ((j-1)n^{-1/3}, jn^{-1/3}]\} \\ &= \{|\hat{p}_U - p_U^*| > (j-1)n^{-1/3}, |\hat{p}_U - p_U^*| \leq jn^{-1/3}\} \\ &\subset \left\{R(p_U^*) - \hat{R}(p_U^*) - (R(\hat{p}_U) - \hat{R}(\hat{p}_U)) \geq \kappa'(j-1)^2 n^{-2/3}, |\hat{p}_U - p_U^*| \leq jn^{-1/3}\right\} \\ &\subset \left\{\Delta_{j,n} \geq \kappa'(j-1)^2 n^{-2/3}\right\}, \end{aligned}$$

where the third line follows from (18), and $\Delta_{j,n}$ in the last line is defined as

$$\Delta_{j,n} \equiv \sup_{g \in \mathcal{G}_{jn^{-1/3}}} \left| \frac{1}{n} \sum_{i=1}^n g(Y_i) - \mathbb{E}g(Y_i) \right|.$$

Therefore,

$$\mathbb{P}\left(|\hat{p}_U - p_U^*| \in ((j-1)n^{-1/3}, jn^{-1/3}]\right) \leq \mathbb{P}\left(\Delta_{j,n} \geq \kappa'(j-1)^2 n^{-2/3}\right).$$

To bound the probability on the RHS of the above inequality, we use the concentration inequality given by Theorem 7.3 in [Bousquet \(2003\)](#), which is a version of [Talagrand's \(1996\)](#) inequality. The concentration inequality states that for all $t > 0$,

$$\mathbb{P}\left(\Delta_{j,n} \geq \mathbb{E}\Delta_{j,n} + \sqrt{2t(\sigma^2 + 2\mathbb{E}\Delta_{j,n})/n} + t/(3n)\right) \leq \exp(-ct),$$

for some universal constant $c > 0$, where

$$\sigma^2 \equiv \sup_{g \in \mathcal{G}_{jn^{-1/3}}} \mathbb{E}g(Y_1)^2 \leq \|\mathcal{G}_{jn^{-1/3}}\|_{L_2}^2 \leq Cjn^{-1/3}.$$

From (19), we have

$$\mathbb{E}\Delta_{j,n} \leq C\sqrt{jn^{-1/3}/n} = C\sqrt{j}n^{-2/3}.$$

By setting $t = \kappa' j^2$, we have

$$\begin{aligned} & \mathbb{E}\Delta_{j,n} + \sqrt{2t(\sigma^2 + 2\mathbb{E}\Delta_{j,n})/n + t/(3n)} \\ & \leq C\sqrt{j}n^{-2/3} + \sqrt{2\kappa' j^2(Cjn^{-1/3} + 2C\sqrt{j}n^{-2/3})/n + \kappa' j^2/(3n)} \\ & \leq C' j^{3/2} n^{-2/3} \leq C^*(j-1)^2 n^{-2/3}, \end{aligned}$$

when j is large enough. Then we have

$$\mathbb{P}\left(\Delta_{j,n} \geq C^*(j-1)^2 n^{-2/3}\right) \leq \mathbb{P}\left(\Delta_{j,n} \geq Cjn^{-2/3}\right) \leq \exp(-c\kappa' j^2), \text{ for } j \text{ large.}$$

To summarize, we have shown that

$$\mathbb{P}\left(n^{1/3}|\hat{p}_U - p_U^*| > M\right) \leq \sum_{j=M+1}^{\infty} \exp(-C_1 j^2) \leq C_3 \exp(-C_2 M^2).$$

By integrating the tail probability, we have

$$\mathbb{E}|\hat{p}_U - p_U^*|^s \lesssim n^{-s/3}.$$

For revenue, we use the second-order Taylor expansion and obtain that

$$\mathbb{E}[R(p_U^*) - R(\hat{p}_U)] \leq \sup_p |R''(p)| \mathbb{E}(\hat{p}_U - p_U^*)^2 \lesssim n^{-2/3}.$$

□

Proof of Theorem 4. We introduce some notations. Let $\tilde{R}_k(p)$ denote the revenue collected from the k th market by charging price p ; that is,

$$\begin{aligned} \tilde{R}_k(p) & \equiv p\mathbb{P}(Y > p, X \in I_k) \\ & = p \int_p^1 \int_{I_k} f_{Y|X}(y|x) f_X(x) dx dy. \end{aligned}$$

Denote $\tilde{p}_k \equiv \operatorname{argmax}_{p \in [0,1]} \tilde{R}_k(p)$ as the maximizer of \tilde{R}_k . The first- and second-order derivatives

of \tilde{R}_k are respectively

$$\begin{aligned}\tilde{R}'_k(p) &= \int_p^1 \int_{I_k} f_{Y|X}(y|x) f_X(x) dx dy - p \int_{I_k} f_{Y|X}(p|x) f_X(x) dx, \\ \tilde{R}''_k(p) &= \int_{I_k} \left(-2f_{Y|X}(p|x) - p \frac{\partial}{\partial y} f_{Y|X}(p|x) \right) f_X(x) dx.\end{aligned}$$

By the Lipschitz continuity assumption, the second-order derivative $\tilde{R}''_k(p)$ exists for almost all $p \in [0, 1]$. Recall that

$$-2f_{Y|X}(p|x) - p \frac{\partial}{\partial y} f_{Y|X}(p|x) \leq -C^*,$$

and f_X is bounded away from zero. Denote $2\kappa'' \equiv C^* \inf_{x \in [0,1]} f_X(x)$. Then

$$\tilde{R}''_k(p) \leq -2\kappa'' \int_{I_k} dx = -2\kappa''/K$$

for almost all $p \in [0, 1]$. By Lemma C.1, we have

$$\tilde{R}_k(\tilde{p}_k) - \tilde{R}_k(p) = |\tilde{R}_k(\tilde{p}_k) - \tilde{R}_k(p)| \geq \frac{\kappa''}{K} (\tilde{p}_k - p)^2, p \in [0, 1]. \quad (20)$$

Note that \tilde{p}_k is not the true optimal price under $F_{Y,X}$. We need to relate it to the true optimal price. Let $k(x_0)$ be such that $x_0 \in I_k$. Then by the triangle inequality, we can decompose the pricing difference into estimation error and approximation error:

$$\begin{aligned}\mathbb{E}|\hat{p}_D(x_0; \text{data}) - p_D^*(x_0)| &= \mathbb{E}|\hat{p}_{k(x_0)} - p_D^*(x_0)| \\ &\leq \underbrace{\mathbb{E}|\hat{p}_{k(x_0)} - \tilde{p}_{k(x_0)}|}_{\text{Estimation error}} + \underbrace{|\tilde{p}_{k(x_0)} - p_D^*(x_0)|}_{\text{Approximation error}}.\end{aligned} \quad (21)$$

Estimation error. Denote \hat{R}_k as the empirical counterpart of \tilde{R}_k ; that is,

$$\hat{R}_k(p) \equiv \frac{p}{n_k} \sum_{i \in \{j: X_j \in I_k\}} \mathbf{1}\{Y_i > p, X_i \in I_k\}.$$

Recall that \hat{p}_k is the maximizer of \hat{R}_k . The basic inequality (i.e., $\hat{R}_k(\hat{p}_k) \geq \hat{R}_k(\tilde{p}_k)$) gives that

$$\begin{aligned}\tilde{R}_k(\tilde{p}_k) - \tilde{R}_k(\hat{p}) &= \tilde{R}_k(\tilde{p}_k) - \tilde{R}_k(\hat{p}) - \hat{R}_k(\tilde{p}_k) + \hat{R}_k(\tilde{p}_k) \\ &\leq \tilde{R}_k(\tilde{p}_k) - \tilde{R}_k(\hat{p}) - \hat{R}_k(\tilde{p}_k) + \hat{R}_k(\hat{p}_k).\end{aligned}\tag{22}$$

Combining (20) and (22) yields

$$\frac{\kappa''}{K}(\tilde{p}_k - \hat{p}_k)^2 \leq \tilde{R}_k(\tilde{p}_k) - \hat{R}_k(\tilde{p}_k) - (\tilde{R}_k(\hat{p}) - \hat{R}_k(\hat{p}_k)).\tag{23}$$

In each I_k ($k = 1, \dots, K$), the optimal price is the same. Based on the proof for Theorem 5, we obtain

$$\mathbb{E}|\hat{p}_k - \tilde{p}_k|^s \lesssim (1/n_k)^{s/3} \lesssim (K/n)^{s/3}$$

where the second “ \lesssim ” follows from the assumption that $f_X \geq \underline{C} > 0$.

Approximation error. The second term $|\tilde{p}_{k(x_0)} - p_D^*(x_0)|$ in (21) is deterministic and can be controlled by using the smoothness conditions. By definition, $p_D^*(x_0)$ satisfies the first-order condition

$$0 = \frac{\partial}{\partial p} r(p_D^*(x), x).$$

By the differentiability condition of \mathcal{F} , $\frac{\partial}{\partial p} r(p, x)$ is continuously differentiable in (p, x) in a neighborhood of $(p_D^*(x), x)$. By the strong concavity, $\frac{\partial^2}{\partial p^2} r(p_D^*(x), x)$ is non-zero. Then by the implicit function theorem, the function $p_D^*(x)$ is well-defined (uniquely determined by the first-order condition) and is differentiable. Its derivative is given as follows:

$$\frac{d}{dx} p_D^*(x) = -\frac{\frac{\partial^2}{\partial p \partial x} r(p_D^*(x), x)}{\frac{\partial^2}{\partial p^2} r(p_D^*(x), x)}.$$

By the strong concavity, the absolute value of $\frac{\partial^2}{\partial p^2} r(p, x)$ is bounded away from zero; also, the function $\left| \frac{\partial}{\partial x} F_{Y|X}(y|x) + y \frac{\partial}{\partial x} f_{Y|X}(y|x) \right|$ is bounded above by \bar{C} . This implies that $p_D^*(x)$ is Lipschitz continuous on $[0, 1]$. We use L_1 to denote the Lipschitz constant. By applying Taylor

expansion to the first-order condition of \tilde{p}_k , we have

$$\begin{aligned} 0 &= \int_{I_k} \frac{\partial}{\partial p} r(\tilde{p}_k, x) f_X(x) dx = \int_{I_k} \underbrace{\frac{\partial}{\partial p} r(p_D^*(x), x)}_{=0} f_X(x) dx \\ &\quad + \int_{I_k} \frac{\partial^2}{\partial p^2} r(\bar{p}(x), x) (\tilde{p}_k - p_D^*(x)) f_X(x) dx, \end{aligned}$$

for some $\bar{p}(x)$ between \tilde{p}_k and $p_D^*(x)$. Rearranging terms shows that \tilde{p}_k is a weighted average of $p_D^*(x), x \in I_k$; that is,

$$\tilde{p}_k = \frac{\int_{I_k} \frac{\partial^2}{\partial p^2} r(\bar{p}(x), x) p_D^*(x) f_X(x) dx}{\int_{I_k} \frac{\partial^2}{\partial p^2} r(\bar{p}(x), x) f_X(x) dx}.$$

This implies that there exists some $x^* \in I_k$ such that $\tilde{p}_k = p_D^*(x^*)$. Since $p_D^*(x)$ is Lipschitz continuous, given x_0 , we have

$$|\tilde{p}_k - p_D^*(x_0)|^s = |p_D^*(x^*) - p_D^*(x_0)|^s \leq L_1^s / K^s, \text{ for any } s \geq 1.$$

Therefore, we obtain the following upper bound

$$\mathbb{E} |\hat{p}_D(x_0; \text{data}) - p_D^*(x_0)|^2 \lesssim (K/n)^{2/3} + 1/K^2.$$

By choosing $K \asymp n^{-1/4}$, the above bound becomes $n^{-1/4}$. Then Lemma C.1(iii) gives that

$$\mathbb{E} [r(p_D^*, x_0) - r(\hat{p}_D(\text{data}), x_0)] \lesssim \mathbb{E} [|\hat{p}_D(x_0; \text{data}) - p_D^*(x_0)|^2] \lesssim (K/n)^{2/3} + 1/K^2.$$

This proves part (i) of the theorem.

For part (ii), we want to bound the expected revenue difference. Consider the following decomposition:

$$\begin{aligned} &R(p_D^*, F_{Y,X}) - R(\hat{p}_D(\text{data}), F_{Y,X}) \\ &\leq R(p_D^*, F_{Y,X}) - R(\tilde{p}, F_{Y,X}) + |R(\tilde{p}, F_{Y,X}) - R(\hat{p}_D(\text{data}), F_{Y,X})|. \end{aligned}$$

The first term on the RHS is deterministic and can be bounded by using Lemma C.1 (iii) as follows:

$$\begin{aligned}
& |R(p_D^*, F_{Y,X}) - R(\tilde{p}, F_{Y,X})| \\
& \leq \int_0^1 |r(p_D^*(x), x) - r(\tilde{p}(x), x)| f_X(x) dx \\
& = \sum_{k=1}^K \int_{I_k} |r(p_D^*(x), x) - r(\tilde{p}_k, x)| f_X(x) dx \\
& \leq \sum_{k=1}^K \int_{I_k} \frac{1}{2} \left| 2f_{Y|X}(y|x) + y \frac{\partial}{\partial y} f_{Y|X}(y|x) \right| \sup_{y,x} |p_D^*(x) - \tilde{p}_k|^2 f_X(x) dx \\
& \lesssim 1/K^2.
\end{aligned}$$

where we have used the first-order condition of p_D^* . For the second term, we have

$$R(\tilde{p}, F_{Y,X}) - R(\hat{p}_D(data), F_{Y,X}) = \sum_{k=1}^K \tilde{R}_k(\tilde{p}_k) - \tilde{R}_k(\hat{p}_k).$$

This is because both $\hat{p}(data)$ and \tilde{p} are constant within each I_k . Their revenues on I_k are reduced to \tilde{R}_k . Note that for every k , $\tilde{R}'_k(\tilde{p}_k) = 0$, and

$$\begin{aligned}
|\tilde{R}''_k(p)| & \leq \int_{I_k} \left| 2f_{Y|X}(y|x) + y \frac{\partial}{\partial y} f_{Y|X}(y|x) \right| f_X(x) dx \\
& \leq \frac{1}{K} \sup_{y,x} \left(\left| 2f_{Y|X}(y|x) + y \frac{\partial}{\partial y} f_{Y|X}(y|x) \right| f_X(x) \right).
\end{aligned}$$

Then Lemma C.1 (iii) gives that

$$\tilde{R}_k(\tilde{p}_k) - \tilde{R}_k(\hat{p}_k) \lesssim 1/K(\tilde{p}_k - \hat{p}_k)^2.$$

Hence, we have

$$\mathbb{E}|R(\tilde{p}, F_{Y,X}) - R(\hat{p}_D(data), F_{Y,X})| \leq \sum_{k=1}^K \mathbb{E}|\tilde{R}_k(\tilde{p}_k) - \tilde{R}_k(\hat{p}_k)| \lesssim (K/n)^{2/3}.$$

To summarize, we have shown that

$$R(p_D^*, F_{Y,X}) - R(\hat{p}_D(data), F_{Y,X}) \lesssim (K/n)^{2/3} + 1/K^2.$$

By choosing $K \asymp n^{-1/4}$, the above bound becomes $n^{-1/2}$. This proves part (ii) of the theorem. \square

Proof of Theorem 6. For part (i), notice that the welfare can be written as a double integral

$$W(p, F_{Y,X}) = \int_0^1 \int_0^{p(x)} y f_{Y|X}(y|x) dy f_X(x) dx.$$

The function $y f_{Y|X}(y|x)$ is nonnegative and bounded for $y, x \in [0, 1]$. Then by the integral mean value theorem, we have

$$\begin{aligned} & \mathbb{E}|W(\hat{p}_D(data), F_{Y,X}) - W(p_D^*, F_{Y,X})| \\ &= \mathbb{E} \left| \int_0^1 \int_{p_D^*(x)}^{\hat{p}_D(x;data)} y f_{Y|X}(y|x) dy f_X(x) dx \right| \\ &\leq \sup_{y,x} |y f_{Y|X}(y|x)| \mathbb{E} \int_0^1 |\hat{p}_D(x;data) - p_D^*(x)| f_X(x) dx. \end{aligned}$$

The integral on the last line can be decomposed based on the K markets:

$$\begin{aligned} \mathbb{E} \int_0^1 |\hat{p}_D(x;data) - p_D^*(x)| f_X(x) dx &\leq \sum_{k=1}^K \int_{I_k} [\mathbb{E}|\hat{p}_D(x;data) - \tilde{p}_k| + |\tilde{p}_k - p_D^*(x)|] f_X(x) dx \\ &= \sum_{k=1}^K \mathbb{E}|\hat{p}_k - \tilde{p}_k|/K + \sum_{k=1}^K \int_{I_k} |\tilde{p}_k - p_D^*(x)| f_X(x) dx \\ &\lesssim (K/n)^{1/3} + 1/K \asymp n^{-1/4}, \end{aligned}$$

where the last line follows from the proof of Theorem 4.

For part (ii), since p_U^* is a scalar, the welfare can be simplified to

$$W(p_U^*, F_Y) = \int_0^{p_U^*} y f_Y(y) dy.$$

Then we have

$$\begin{aligned} \mathbb{E}|W(\hat{p}_U(data_Y), F_Y) - W(p_U^*, F_Y)| &= \mathbb{E} \left| \int_{p_U^*}^{\hat{p}_U(data_Y)} y f_Y(y) dy \right| \\ &\leq \sup_y |y f_Y(y)| \mathbb{E}|\hat{p}_U(data_Y) - p_U^*| \\ &\lesssim n^{-1/3}, \end{aligned}$$

where we have used Theorem 5 along with the fact that $y f_Y(y)$ is nonnegative and bounded for $y \in [0, 1]$. \square

C Auxiliary Lemmas

Lemma C.1. *Let f be a function on $[0, 1]$. Assume that f is differentiable and its derivative f' is Lipschitz continuous. Let z^* be a point in $[0, 1]$ such that $f'(z^*) = 0$.*

- (i) *The derivative f' is a.e. differentiable on $[0, 1]$.*
- (ii) *Assume that there exists $\kappa_1 > 0$ such that $f''(z) \leq -\kappa_1$ for almost all $z \in [0, 1]$. Then, for any $z \in [0, 1]$, we have*

$$|f(z) - f(z^*)| \geq \frac{\kappa_1}{2}(z - z^*)^2.$$

- (iii) *Assume that there exists $\kappa_2 > 0$ such that $|f''(z)| \leq \kappa_2$ for almost all $z \in [0, 1]$. Then, for any $z \in [0, 1]$, we have*

$$|f(z) - f(z^*)| \leq \frac{\kappa_2}{2}(z - z^*)^2.$$

Proof of Lemma C.1. For part (i), notice that a Lipschitz continuous function is absolutely continuous. By Theorem 3.35 in Chapter 3 of [Folland \(1999\)](#), we know that f' is differentiable a.e. with

$$f'(z_1) - f'(z_2) = \int_{z_2}^{z_1} f''(z) dz.$$

For part (ii), we can apply the fundamental theorem of calculus twice and obtain that

$$\begin{aligned} f(z) - f(z^*) &= \int_{z^*}^z f'(\tilde{z}) d\tilde{z} \\ &= \int_{z^*}^z (f'(z_1) - f'(z^*)) dz_1 \\ &= \int_{z^*}^z \int_{z^*}^{z_1} f''(z_2) dz_2 dz_1 \\ &\leq -\kappa_1 \int_{z^*}^z \int_{z^*}^{z_1} dz_2 dz_1, \end{aligned}$$

where in the second line we have used the assumption that $f'(z^*) = 0$, and in the last line we have used the assumption that $f''(z) \leq -\kappa_1$ for almost all $z \in [0, 1]$. The double integral in the

last line is equal to

$$\int_{z^*}^z \int_{z^*}^{z_1} dz_2 dz_1 = \int_{z^*}^z (z_1 - z^*) dz_1 = \frac{(z - z^*)^2}{2}.$$

Therefore, we have

$$|f(z) - f(z^*)| \geq \frac{\kappa_1}{2}(z - z^*)^2.$$

Part (iii) can be proved analogously. \square

Lemma C.2. *For the uniform distribution on $[0, 1]$, the revenue function $R(y) = y(1 - y)$. The revenue function is twice-differentiable with second-order derivative $R''(y) = -2, y \in [0, 1]$. The optimal price is $1/2$.*

Proof of Lemma C.2. The proof is straightforward. \square

Lemma C.3. *Recall the perturbation function ϕ_Y defined in (14). Consider the following density function*

$$f(y) \equiv 1 + b\delta\phi_Y\left(\frac{y - 1/2}{\delta}\right) = \begin{cases} 1, & \text{if } y \in [0, 1/2 - \delta), \\ by + 1 - \frac{b}{2} + b\delta, & \text{if } y \in [1/2 - \delta, 1/2), \\ -by + 1 + \frac{b}{2} + b\delta, & \text{if } y \in [1/2, 1/2 + 2\delta), \\ by + 1 - \frac{b}{2} - 3b\delta, & \text{if } y \in [1/2 + 2\delta, 1/2 + 3\delta), \\ 1, & \text{if } y \in [1/2 + 3\delta, 1], \\ 0, & \text{otherwise.} \end{cases}$$


Denote F as the corresponding cumulative distribution function, $R(y) \equiv y(1 - F(y))$ the revenue function, and $p^* \equiv \operatorname{argmax}_{y \in [0, 1]} R(y)$ the optimal price. If $C^* \in (0, 2)$, $|b| < 4 - 2C^*$, and $\delta > 0$ is sufficiently small, then the following statements hold.

- (i) *The density f is Lipschitz continuous.*
- (ii) *The revenue function is twice-differentiable a.e. The second-order derivative is bounded a.e. and satisfies that*

$$-2f(y) - yf'(y) \geq -C^* \text{ for almost all } y.$$

(iii) For $b > 0$, the optimal price $p^* \in (1/2 - \delta, 1/2 - b\delta/8)$. For $b < 0$, the optimal price $p^* \in (1/2 - b\delta/8, 1/2 + 2\delta)$. For $b = 0$, the optimal price $p^* = 1/2$. In particular, p^* is always an interior solution, and f is always differentiable in a neighborhood of p^* .

Figure C.1: Perturbation function and perturbed density.



Proof of Lemma C.3. For reference, we plot here the perturbation function ϕ_Y and the perturbed density f . Part (i) is straightforward. The density f is piecewise linear and hence Lipschitz continuous with Lipschitz constant b . To verify the strong concavity in part (ii), note that the corresponding revenue function R is continuously differentiable and twice-differentiable a.e. on the support $[0, 1]$. Its second-order derivative

$$R''(y) = -2f(y) - yf'(y) = \begin{cases} -2, & \text{if } y \in [0, 1/2 - \delta], \\ -3by - 2 + b - 2b\delta, & \text{if } y \in [1/2 - \delta, 1/2], \\ 3by - 2 - b - 2b\delta, & \text{if } y \in [1/2, 1/2 + 2\delta], \\ -3by - 2 + b + 6b\delta, & \text{if } y \in [1/2 + 2\delta, 1/2 + 3\delta], \\ -2, & \text{if } y \in [1/2 + 3\delta, 1]. \end{cases}$$

We can see that R'' is piecewise linear and hence bounded a.e. We further show that R'' is bounded away from zero by κ . On the intervals $[0, 1/2 - \delta]$ and $[1/2 + 3\delta, 1]$, we have $R''(y) = -2 < -C^*$. We check the remaining three intervals one by one. On the interval $[1/2 - \delta, 1/2]$,

the condition $|b| < 4 - 2C^*$ ensures that

$$b \geq 0 \implies R''(y) \leq R''(1/2 - \delta) = -b/2 - 2 + b\delta \leq -C^*,$$

$$b < 0 \implies R''(y) \leq R''(1/2) = -b/2 - 2 - 2b\delta \leq -C^*,$$

when δ is sufficiently small. On the interval $[1/2, 1/2 + 2\delta]$, we have

$$b \geq 0 \implies R''(y) \leq R''(1/2 + 2\delta) = b/2 - 2 + 4b\delta \leq -C^*,$$

$$b < 0 \implies R''(y) \leq R''(1/2) = b/2 - 2 - 2b\delta \leq -C^*,$$

when δ is sufficiently small. On the interval $[1/2 + 2\delta, 1/2 + 3\delta]$, we have

$$b \geq 0 \implies R''(y) \leq R''(1/2 + 2\delta) = -b/2 - 2 < -C^*,$$

$$b < 0 \implies R''(y) \leq R''(1/2 + 3\delta) = -b/2 - 2 - 3b\delta < -C^*,$$

To summarize, we have shown that $R''(y) \leq -C^*$ a.e. on $[0, 1]$ provided that $\delta > 0$ is sufficiently small.

For part (iii), we first consider the case $b > 0$. We only need to consider the interval $[1/2 - \delta, 1/2]$. The reason will become clear later. The cumulative distribution function

$$F(y) = \frac{b}{2}y^2 + \left(1 - \frac{b}{2} + b\delta\right)\delta y + \frac{b}{2}(1/2 - \delta)^2, y \in [1/2 - \delta, 1/2].$$

The revenue function

$$R(y) = -\frac{b}{2}y^3 - \left(1 - \frac{b}{2} + b\delta\right)y^2 + \left(1 - \frac{b}{2}(1/2 - \delta)^2\right)y, y \in [1/2 - \delta, 1/2].$$

The marginal revenue

$$R'(y) = -\frac{3b}{2}y^2 - (2 - b + 2b\delta)y + 1 - \frac{b}{2}(1/2 - \delta)^2, y \in [1/2 - \delta, 1/2].$$

We evaluate the marginal revenue at two points $1/2 - \delta$ and $1/2 - \frac{b\delta}{8}$. When $y = 1/2 - \delta$, the marginal revenue

$$R'(1/2 - \delta) = \delta > 0.$$

When $y = 1/2 - b\delta/8$, the marginal revenue

$$R' \left(1/2 - \frac{b\delta}{8} \right) \approx \frac{b(b-4)}{16} \delta < 0,$$

where we have omitted higher order terms involving δ^2 . Therefore, $R' \left(1/2 - \frac{b\delta}{8} \right)$ is negative for sufficiently small δ . Since the marginal revenue R' is strictly decreasing on the entire domain $[0, 1]$, we know that the only zero of R' (which is the optimal price p^*) is within the region $(1/2 - \delta, 1/2 - \frac{b\delta}{8})$. Within this region, the revenue is twice-differentiable everywhere.

Next, we consider the case $b < 0$. In this case, we only need to study the region $[1/2, 1/2 + 2\delta]$. The cumulative distribution function

$$F(y) = -\frac{b}{2}y^2 + \left(1 + \frac{b}{2} + b\delta\right)y + \frac{b}{2}\delta^2 - \frac{b}{2}\delta - \frac{b}{8}, y \in [1/2, 1/2 + 2\delta].$$

The revenue function

$$R(y) = y(1 - F(y)) = \frac{b}{2}y^3 - \left(1 + \frac{b}{2} + b\delta\right)y^2 + \left(1 + \frac{b}{8} - \frac{b}{2}\delta^2 + \frac{b}{2}\delta\right)y, y \in [1/2, 1/2 + 2\delta].$$

The marginal revenue

$$R'(y) = \frac{3b}{2}y^2 - (2 + b + 2b\delta)y + \left(1 + \frac{b}{8} - \frac{b}{2}\delta^2 + \frac{b}{2}\delta\right), y \in [1/2, 1/2 + 2\delta].$$

We evaluate the marginal revenue at two points $1/2 + \delta$ and $1/2 - \frac{b\delta}{8}$. When $y = 1/2 + \delta$, the marginal revenue

$$R'(1/2 + \delta) \approx -2\delta < 0,$$

where we have omitted higher order terms involving δ^2 . When $y = 1/2 - b\delta/8$, the marginal revenue

$$R' \left(1/2 - \frac{b\delta}{8} \right) \approx \frac{b(b+4)}{16} \delta > 0,$$

where we have omitted higher order terms involving δ^2 . Since the marginal revenue R' is strictly decreasing on the entire domain $[0, 1]$, we know that the only zero of R' (which is the optimal price p^*) is within the region $(1/2 - \frac{b\delta}{8}, 1/2 + \delta)$. Within this region, the revenue is twice-

differentiable everywhere.

Lastly, when $b = 0$, the density function is constant, and Lemma C.2 shows that the optimal price is $1/2$. Therefore, regardless of the sign of b , the optimal price is always an interior solution, and is in the interior of a region on which the revenue function is twice-differentiable. \square

Lemma C.4. Take $x_0 \in [0, 1]$. Recall the following definition of $\omega_D(\epsilon)$ and $\omega_U(\epsilon)$:

$$\begin{aligned}\omega_D(\epsilon) &\equiv \sup_{F_1, F_2 \in \mathcal{F}} \left\{ |p_D^*(x_0; F_1) - p_D^*(x_0; F_2)| : H(F_1 \| F_2) \leq \epsilon \right\}, \\ \omega_U(\epsilon) &\equiv \sup_{F_1, F_2 \in \mathcal{F}^U} \left\{ |p_U^*(F_1) - p_U^*(F_2)| : H(F_1 \| F_2) \leq \epsilon \right\}.\end{aligned}$$

Then

$$\begin{aligned}\inf_{\check{p}_D \in \check{\mathcal{D}}} \sup_{F_{Y,X} \in \mathcal{F}} \mathbb{E}_{F_{Y,X}} |\check{p}_D(x_0; \text{data}) - p_D^*(x_0; F_{Y,X})| &\geq \frac{1}{8} \omega_D \left(1/(2\sqrt{n}) \right), \\ \inf_{\check{p}_U \in \check{\mathcal{U}}} \sup_{F_Y \in \mathcal{F}^U} \mathbb{E}_{F_Y} |\check{p}_U(\text{data}_Y) - p_U^*(F_Y)| &\geq \frac{1}{8} \omega_U \left(1/(2\sqrt{n}) \right).\end{aligned}$$

Proof of Lemma C.4. By treating $p_D^*(x_0; \cdot)$ and $p_U^*(\cdot)$ as functionals, the desired results directly follow from Corollary 15.6 (Le Cam for functionals) in Chapter 15 of [Wainwright \(2019\)](#). \square

Lemma C.5. Let $\{F_{Y,X}^j : 1 \leq j \leq M\} \subset \mathcal{F}$ be such that

$$\|p_D^*(F_{Y,X}^j) - p_D^*(F_{Y,X}^{j'})\|_2 \geq 2\delta, j \neq j'.$$

Then we have

$$\inf_{\check{p}_D \in \check{\mathcal{D}}} \sup_{F_{Y,X} \in \mathcal{F}} \mathbb{E} \|\check{p}_D(\text{data}) - p_D^*(F_{Y,X})\|_2^2 \geq \delta^2 \left(1 - \frac{\sum_{j,j'=1}^M KL(F_{Y,X}^j \| F_{Y,X}^{j'}) / M^2 + \log 2}{\log M} \right)$$

Proof of Lemma C.5. The result follows from Proposition 15.12 (the Fano's inequality) and inequality (15.34) (convexity of the KL divergence) in Chapter 15 of [Wainwright \(2019\)](#), where Φ is taken to be the square function, ρ the L_2 -distance, and θ the functional p_D^* . \square

Lemma C.6. Consider the following function class:

$$\{(y, x) \mapsto (p\mathbf{1}\{y \geq p\} - \tilde{p}\mathbf{1}\{y \geq \tilde{p}\})\mathbf{1}\{x \in [k/K, (k+1)/K]\} : p \in [0, 1]\}.$$

For any $\tilde{p} \in [0, 1]$, $K \geq 1$, and $0 \leq k \leq K-1$, the above class is a VC-subgraph with VC-dimension no greater than 2.

Proof of Lemma C.6. By Lemma 2.6.22 in Chapter 2 of [van der Vaart and Wellner \(1996\)](#), the class

$$\{(y, x) \mapsto p \mathbf{1}\{y \geq p\} : p \in [0, 1]\}$$

is a VC-subgraph with VC-dimension no greater than 2.¹³ The function $(y, x) \mapsto \tilde{p} \mathbf{1}\{y \geq \tilde{p}\}$ is a fixed function that does not depend on the index p . By the proof Lemma 2.6.18(v) in [van der Vaart and Wellner \(1996\)](#), the class

$$\{(y, x) \mapsto p \mathbf{1}\{y \geq p\} - \tilde{p} \mathbf{1}\{y \geq \tilde{p}\} : p \in [0, 1]\}$$

is a VC-subgraph with VC-dimension no greater than 2. Lastly, we multiply each function in the class by an indicator $\mathbf{1}\{x \in [k/K, (k+1)/K]\}$. This does not increase the VC-dimension. \square

Lemma C.7. Let Z_1, \dots, Z_n be an i.i.d. sequence of random variables from distribution P . Let \mathcal{G} be a class of VC-subgraph functions with VC-dimension ν and envelope function G . Assume that $\|G\|_{L_2(P)} < \infty$. Then we have

$$\mathbb{E} \sup_{g \in \mathcal{G}} \left| \frac{1}{n} \sum_{i=1}^n g(Z_i) - \mathbb{E} g(Z_i) \right| \leq 8\sqrt{2} \frac{\|G\|_{L_2(P)}}{\sqrt{n}} (\log(2C) + \log(\nu) + (\log(16) + 3)\nu),$$

for some universal constant C , where the $L_2(P)$ norm $\|f - g\|_{L_2(P)} \equiv \left(\int_X [f(x) - g(x)]^2 \mathbb{P}(dx) \right)^{\frac{1}{2}}$.

Proof of Lemma C.7. This is a well-known result in the literature. We include it here for completeness. Let $N(\mathcal{G}, L_2(Q), \tau)$ denote the covering number of $(\mathcal{G}, L_2(Q))$. By Remark 3.5.5 in Chapter 3 of [Giné and Nickl \(2015\)](#), we know that

$$\mathbb{E} \sup_{g \in \mathcal{G}} \left| \frac{1}{n} \sum_{i=1}^n g(X_i) - \mathbb{E} g(X_i) \right| \leq 8\sqrt{2} \frac{\|G\|_{L_2(P)}}{\sqrt{n}} \int_0^1 \sup_Q \sqrt{\log 2N(\mathcal{G}, L_2(Q), \tau) \|G\|_{L_2(Q)}} d\tau,$$

¹³In the original statement of the lemma, the VC dimension is no greater than 3. This is because the definition of VC dimension in [van der Vaart and Wellner \(1996\)](#) is the smallest number n for which no set of n points is shattered. The definition we use in this paper is the largest number n that some set of n points is shattered.

where the supremum is taken over all discrete probabilities with a finite number of atoms. By Theorem 2.6.7 in Chapter 2 of [van der Vaart and Wellner \(1996\)](#), we know that for any probability measure Q ,

$$N(\mathcal{G}, L_2(Q), \tau \|G\|_{L_2(Q)}) \leq C\nu(16e)^\nu (1/\tau)^{2\nu},$$

for some universal constant C . Therefore,

$$\int_0^1 \sup_Q \sqrt{\log 2N(\mathcal{G}, L_2(Q), \tau \|G\|_{L_2(Q)})} d\tau \leq \log(2C) + \log(\nu) + (\log(16) + 3)\nu$$

Then the desired result follows. □

References

- Allouah, Amine, Achraf Bahamou, and Omar Besbes**, “Pricing with samples,” *Operations Research*, 2022, 70 (2), 1088–1104.
- , —, and —, “Optimal pricing with a single point,” *Management Science*, 2023.
- Babaioff, Moshe, Yannai A Gonczarowski, Yishay Mansour, and Shay Moran**, “Are two (samples) really better than one?,” in “Proceedings of the 2018 ACM Conference on Economics and Computation” 2018, pp. 175–175.
- Balcan, Maria-Florina, Amit Daniely, Ruta Mehta, Ruth Urner, and Vijay V Vazirani**, “Learning economic parameters from revealed preferences,” in “Web and Internet Economics: 10th International Conference, WINE 2014, Beijing, China, December 14-17, 2014. Proceedings 10” Springer 2014, pp. 338–353.
- Baliga, Sandeep and Rakesh Vohra**, “Market research and market design,” *Advances in theoretical Economics*, 2003, 3 (1).
- Basu, Pathikrit and Federico Echenique**, “On the falsifiability and learnability of decision theories,” *Theoretical Economics*, 2020, 15 (4), 1279–1305.
- Beigman, Eyal and Rakesh Vohra**, “Learning from revealed preference,” in “Proceedings of the 7th ACM Conference on Electronic Commerce” 2006, pp. 36–42.

- Bergemann, Dirk and Karl H. Schlag**, “Pricing without Priors,” *Journal of the European Economic Association*, 2008, 6 (2-3), 560–569.
- and —, “Robust Monopoly Pricing,” *Journal of Economic Theory*, 2011, 146 (6), 2527–2543.
- Bousquet, Olivier**, “Concentration inequalities for sub-additive functions using the entropy method,” in “Stochastic inequalities and applications,” Springer, 2003, pp. 213–247.
- Carroll, Gabriel**, “Robustness in Mechanism Design and Contracting,” *Annual Review of Economics*, 2019, 11 (1), 139–166.
- Chambers, Christopher P, Federico Echenique, and Nicolas S Lambert**, “Recovering preferences from finite data,” *Econometrica*, 2021, 89 (4), 1633–1664.
- , —, and —, “Recovering utility,” *arXiv preprint arXiv:2301.11492*, 2023.
- Cole, Richard and Tim Roughgarden**, “The sample complexity of revenue maximization,” in “Proceedings of the forty-sixth annual ACM symposium on Theory of computing” 2014, pp. 243–252.
- Cover, Thomas M. and Joy A. Thomas**, *Elements of Information Theory*, John Wiley & Sons, Ltd, 2005.
- Cr  mer, Jacques and Richard P McLean**, “Full extraction of the surplus in Bayesian and dominant strategy auctions,” *Econometrica: Journal of the Econometric Society*, 1988, pp. 1247–1257.
- Devanur, Nikhil R, Zhiyi Huang, and Christos-Alexandros Psomas**, “The sample complexity of auctions with side information,” in “Proceedings of the forty-eighth annual ACM symposium on Theory of Computing” 2016, pp. 426–439.
- Dhangwatnotai, Peerapong, Tim Roughgarden, and Qiqi Yan**, “Revenue maximization with a single sample,” *Games and Economic Behavior*, 2015, 91, 318–333.
- Folland, Gerald B**, *Real analysis: modern techniques and their applications*, 2 ed., Vol. 40, New York, NY: John Wiley & Sons, 1999.

- Fu, Hu, Nicole Immorlica, Brendan Lucier, and Philipp Strack**, “Randomization beats second price as a prior-independent auction,” in “Proceedings of the sixteenth ACM conference on economics and computation” 2015, pp. 323–323.
- , **Nima Haghpanah, Jason Hartline, and Robert Kleinberg**, “Full surplus extraction from samples,” *Journal of Economic Theory*, 2021, 193, 105230.
- Giné, Evarist and Richard Nickl**, *Mathematical Foundations of Infinite-Dimensional Statistical Models* Cambridge Series in Statistical and Probabilistic Mathematics, Cambridge University Press, 2015.
- Goldberg, Andrew V, Jason D Hartline, Anna R Karlin, Michael Saks, and Andrew Wright**, “Competitive auctions,” *Games and Economic Behavior*, 2006, 55 (2), 242–269.
- Gonçalves, Duarte and Bruno Furtado**, “Statistical mechanism design: robust pricing and reliable projections,” 2020.
- Guo, Chenghao, Zhiyi Huang, and Xinzhi Zhang**, “Settling the sample complexity of single-parameter revenue maximization,” in “Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing” 2019, pp. 662–673.
- Hannak, Aniko, Gary Soeller, David Lazer, Alan Mislove, and Christo Wilson**, “Measuring Price Discrimination and Steering on E-commerce Web Sites,” in “Proceedings of the Internet Measurement Conference (IMC 2014)” Vancouver, Canada November 2014.
- Hartmann, Wesley, Harikesh S Nair, and Sridhar Narayanan**, “Identifying causal marketing mix effects using a regression discontinuity design,” *Marketing Science*, 2011, 30 (6), 1079–1097.
- Huang, Zhiyi, Yishay Mansour, and Tim Roughgarden**, “Making the most of your samples,” *SIAM Journal on Computing*, 2018, 47 (3), 651–674.
- Jank, Wolfgang and Galit Shmueli**, *Modeling online auctions*, John Wiley & Sons, 2010.
- Manski, Charles F.**, “Statistical Treatment Rules for Heterogeneous Populations,” *Econometrica*, 2004, 72 (4), 1221–1246.

- Riley, John and Richard Zeckhauser**, “Optimal selling strategies: When to haggle, when to hold firm,” *The Quarterly Journal of Economics*, 1983, 98 (2), 267–289.
- Segal, Ilya**, “Optimal pricing mechanisms with unknown demand,” *American Economic Review*, 2003, 93 (3), 509–529.
- Stoye, Jörg**, “Minimax regret treatment choice with finite samples,” *Journal of Econometrics*, 2009, 151 (1), 70–81.
- , “Minimax regret treatment choice with covariates or with limited validity of experiments,” *Journal of Econometrics*, 2012, 166 (1), 138–156. Annals Issue on “Identification and Decisions”, in Honor of Chuck Manski’s 60th Birthday.
- Talagrand, Michel**, “New concentration inequalities in product spaces,” *Inventiones mathematicae*, 1996, 126 (3), 505–563.
- Tsybakov, Alexandre B.**, *Introduction to Nonparametric Estimation* Springer Series in Statistics, 1 ed., Springer, New York, NY, 2009.
- van der Vaart, Aad W and Jon A Wellner**, *Weak Convergence and Empirical Processes*, New York, NY: Springer, 1996.
- Wainwright, Martin J**, *High-dimensional statistics: A non-asymptotic viewpoint*, Vol. 48, Cambridge University Press, 2019.
- Zhang, Hanrui and Vincent Conitzer**, “Learning the Valuations of a k -demand Agent,” in “International Conference on Machine Learning” PMLR 2020, pp. 11066–11075.