# Dynamic Pricing in High-dimensions 

Adel Javanmard<br>AJAVANMA@USC.EDU<br>Department of Data Sciences and Operations<br>Marshall School of Business<br>University of Southern California<br>Los Angeles, CA 90089, USA<br>\section*{Hamid Nazerzadeh}<br>NAZERZAD@USC.EDU<br>Department of Data Sciences and Operations<br>Marshall School of Business<br>University of Southern California<br>Los Angeles, CA 90089, USA

Editor: Peter Auer


#### Abstract

We study the pricing problem faced by a firm that sells a large number of products, described via a wide range of features, to customers that arrive over time. Customers independently make purchasing decisions according to a general choice model that includes products features and customers' characteristics, encoded as $d$-dimensional numerical vectors, as well as the price offered. The parameters of the choice model are a priori unknown to the firm, but can be learned as the (binary-valued) sales data accrues over time. The firm's objective is to maximize its revenue. We benchmark the performance using the classic regret minimization framework where the regret is defined as the expected revenue loss against a clairvoyant policy that knows the parameters of the choice model in advance, and always offers the revenue-maximizing price. This setting is motivated in part by the prevalence of online marketplaces that allow for real-time pricing.

We assume a structured choice model, parameters of which depend on $s_{0}$ out of the $d$ product features. Assuming that the market noise distribution is known, we propose a dynamic policy, called Regularized Maximum Likelihood Pricing (RMLP) that leverages the (sparsity) structure of the high-dimensional model and obtains a logarithmic regret in $T$. More specifically, the regret of our algorithm is of $O\left(s_{0} \log d \cdot \log T\right)$. Furthermore, we show that no policy can obtain regret better than $O\left(s_{0}(\log d+\log T)\right)$. In addition, we propose a generalization of our policy to a setting that the market noise distribution is unknown but belongs to a parametrized family of distributions. This policy obtains regret of $O(\sqrt{(\log d) T})$. We further show that no policy can obtain regret better than $\Omega(\sqrt{T})$ in such environments.


Keywords: Revenue Management, Dynamic Pricing, High-dimensional Regression, Maximum Likelihood, Regret, Sparsity

## 1. Introduction

A central challenge in revenue management is determining the optimal pricing policy when there is uncertainty about customers' willingness to pay. Due to its importance, this problem has been studied extensively (Kleinberg and Leighton, 2003; Besbes and Zeevi, 2009;

[^0]Badanidiyuru et al., 2013; Wang et al., 2014; Broder and Rusmevichientong, 2012; Keskin and Zeevi, 2014; den Boer and Zwart, 2014; Cohen et al., 2016). Most of these models are built around the following classic setting: customers arrive over time; the seller posts a price for each customer; if the customer's valuation is above the posted price, a sale occurs and the seller collects a revenue in the amount of the posted price; otherwise, no sale occurs and no revenue is generated. Based on this and the previous feedbacks, the seller updates the posted price. Therefore, the seller is involved in the realm of exploration-exploitation as he needs to choose between learning about the valuations and exploiting what has been learned so far to collect revenue.

In this work, we consider a setting with a large number of products which are defined via a wide range of features. The valuations are given by $v(\theta, x)$ with $x$ being the (observable) feature vectors of products and $\theta_{0}$ representing the customer's characteristics true parameters of the choice model, which is initially unknown to the seller, cf. (Amin et al., 2014; Cohen et al., 2016). An important special case of this setting is the linear model in which

$$
\begin{equation*}
v(\theta, x)=\theta_{0} \cdot x+z \tag{1}
\end{equation*}
$$

where $z$ captures the idiosyncratic noise in valuations.
Our setting is motivated in part by applications in online marketplaces. For instance, a company such as Airbnb recommends prices to hosts based on many features including the space (number of rooms, beds, bathrooms, etc.), amenities (AC, WiFi, washer, parking, etc.), the location (accessibility to public transportation, walk score of the neighborhood, etc.), house rules (pet-friendly, non-smoking, etc.), as well as the prediction of the demand which itself depends on many factors including the date, events in the area, availability and prices of near-by hotels, etc. (Airbnb Documentation, 2015). Therefore, the vector describing each property can have hundreds of features. Another important application comes from online advertising. Online publishers set the (reserve) price of ads based on many features including user's demographic, browsing history, the context of the webpage, the size and location of the ad on the page, etc.

In this work, we propose Regularized Maximum Likelihood Pricing (RMLP) policy for dynamic pricing in high-dimensional environments. As suggested by its name, the policy uses maximum likelihood method to estimate the true parameters of the choice model. In addition, using an ( $\ell_{1}$-norm) regularizer, our policy exploits the structure of the optimal solution; namely, the performance of the RMLP policy significantly improves if the valuations are essentially determined by a small subset of features. More formally, the difference between the revenue obtained by our policy and the benchmark policy that knows in advance the true parameters of the choice model, $\theta_{0}$, is bounded by $O\left(s_{0} \log d \cdot \log T\right)$, where $T, d$, and $s_{0}$ respectively denote the length of the horizon, number of the features, and sparsity (i.e., number of non-zero elements of $\theta_{0}$ ). We show that our results are nearly optimal and no policy can obtain regret better than $O\left(s_{0}(\log d+\log T)\right)$.

In the RMLP policy it is assumed that the noise distribution (distribution of term $z$ in valuation model (1)) is known to the seller. We also consider the setting that the noise distribution is unknown but belongs to a parametrized family of distributions. We propose RMLP-2 policy for such setting and prove that it achieves a regret of $O(\sqrt{(\log d) T})$. We further how that no policy can get regret better than $\Omega(\sqrt{T})$ in this setting.

We point out that our results can be applied to applications where the features dimension is larger than the time horizon of interest. A powerful pricing policy for these applications should obtain regret that scales gracefully with the dimension. Note that in general, little can be learned about the model parameters $\theta_{0}$ if $T<d$, because the number of degrees of freedom $d$ exceeds the number of observations $T$, and therefore, any estimator can be arbitrary erroneous. However, when there is prior knowledge about the structure of unknown parameter $\theta_{0}$, (e.g., sparsity), then accurate estimations are attainable even when $T<d$.

### 1.1. Organization and Contributions

In the remaining part of the introduction, we discuss how our work is positioned with respect to the literature on dynamic pricing and high-dimensional statistics. The rest of the paper presents the following contributions:

- Section 2. We formally present our model and discuss the technical assumptions and the benchmark policy.
- Section 3. We present the RMLP policy that uses a regularized maximum likelihood function to estimate the parameters of the choice model and subsequently sets the best prices based on the obtained estimates. The RMLP policy assumes that the market noise distribution is known and forms the log-likelihood function based on this knowledge.
- Section 4. We analyze the regret of the RMLP policy with respect to a clairvoyant who knows the parameters of the choice model in advance. We show that it achieves a regret of $O\left(s_{0} \log d \cdot \log T\right)$.
- Section 5. We provide a lower bound on the regret of any policy that does not know the parameters of the choice model in advance. Namely, we show that no policy can obtain regret better than $O\left(s_{0}(\log d+\log T)\right)$.
- Section 6. We generalize the RMLP policy to non-linear valuations functions and show that, similar to the case of linear valuations, it achieves a regret of $O\left(s_{0} \log d \cdot \log T\right)$.
- Section 7. We consider a setting where the exact underlying noise distribution is not known but it is assumed to belong to a known class of log-concave distributions. We propose the RMLP-2 policy, so-named, for this setting whose regret is $O(\sqrt{(\log d) T})$. We also argue that no policy can get better than $O(\sqrt{T})$ regret in this setting.
- Section 8. We provide the proofs of our main theorems and results.

We conclude our work along with some future directions in Section 9. Proof of technical lemmas are deferred to the appendices.

### 1.2. Related Work

Our work contributes to the literature on dynamic pricing as well as high dimensional statistics. In the following, we briefly overview the work closest to ours in these contexts.

Dynamic pricing and learning. The literature on dynamic pricing and learning has been growing over the past few years, motivated in part by the advances in big data technology that allow firms to easily collect and utilize information. We briefly discuss some of the recent lines of research in this literature. We refer to den Boer (2015) for an excellent survey on this topic.

Parametric Approach. A natural approach to capture uncertainty about the customers' valuations is to model the uncertainty using a small number of parameters, and then estimate those parameters using classical statistical methods such as maximum likelihood (Broder and Rusmevichientong, 2012; den Boer and Zwart, 2013, 2014; Chen et al., 2015) or least square estimation (Keskin, 2014; Bastani and Bayati, 2016). Our work is similar to this line of work, in that we assume a parametric model for customer's valuations and apply the maximum likelihood method using the randomness of the idiosyncratic noise in valuations. However, the parameter vector $\theta$ is high-dimensional, whose dimension $d$ (that can even exceed the time horizon of interest $T$ ). We use regularized maximum-likelihood in order to promote sparsity structure in the estimated parameter. Further, our pricing policy has an episodic theme which makes the posted prices $p_{t}$ in each episode independent of the idiosyncratic noise in valuations, $z_{t}$, in that episode. This is in contrast to other policies based on maximum-likelihood, such as MLE-GREEDY (Broder and Rusmevichientong, 2012), or greedy iterative least square (GILS) (Keskin, 2014; den Boer and Zwart, 2014; Qiang and Bayati, 2016) that use the entire history of observations to update the estimate for the model parameters at each step.

Bayesian Approach. One of the earliest work on Bayesian parametric approach in this context is by (Rothschild, 1974) who consider a Bayesian framework where the firm can choose from two prices with unknown demand and show that (myopic) Bayesian policies may lead to "incomplete learning." However, carefully designed variations of the myopic policies can (optimally) learn the optimal price (Harrison et al., 2012); see also Keller and Rady (1999); Araman and Caldentey (2009).

Non-Parametric models. An early work in non-parametric setting is by Kleinberg and Leighton (2003). They model the dynamic pricing problem as a multi-armed bandit (MAB) where each arm corresponds to a (discretized) posted price. They propose an $O(\sqrt{T})$ algorithm where $T$ is the length of the horizon. Similar results have been obtained in more general settings (Badanidiyuru et al., 2013; Agrawal and Devanur, 2014) including setting with inventory constraints (Besbes and Zeevi, 2009; Babaioff et al., 2012; Wang et al., 2014).

Feature-based Models. Recent papers on dynamic pricing consider models with features/covariates. Amin et al. (2014), in a model similar to ours, present an algorithm that obtains regret $O\left(T^{2 / 3}\right)$; they also study dynamic incentive compatibility in repeated auctions. Another closely related work to ours is by Cohen et al. (2016). Their model differs from ours in two main aspects: $i$ ) their model is deterministic (no idiosyncratic noise) $i i$ ) the arrivals (of features vectors) is modeled as adversarial. They propose a clever binary-search approach using the Ellipsoid method which obtains regret of $O\left(d^{2} \log (T / d)\right)$. Recently, Golrezaei et al. (2018) studied the problem of dynamic pricing when a group of strategic buyers competes with each other in repeated contextual auctions. To the extent of our knowledge, ours is the first work that highlights the role of structure/sparsity in dynamic pricing.

Qiang and Bayati (2016) study a model where the seller can observe the demand itself, not a binary signal as in our setting. They show that a myopic policy based on least-square estimations can obtain a logarithmic regret. Bastani and Bayati (2016) study a multi-armed bandit setting, with discrete arms, and high-dimensional covariates, generalizing results of Goldenshluger and Zeevi (2013). More precisely, the decision-maker has access to $K$ arms (decisions), where each arm $i$ has an unknown sparse parameter vector $\beta_{i} \in \mathbb{R}^{d}$. Upon pulling arm $i$ at time $t$, the decision-maker gets reward $x_{t} \cdot \beta_{i}+z_{t}$. The authors present an algorithm, using a LASSO estimator, that obtains regret $O\left(K(\log T+\log d)^{2}\right)$ where $K$ denotes the number of arms. By contrast, in our setting decisions (prices) can take continuous values and also the resulting rewards do not follow a similar relationship with the covariates and the decisions.

High-dimensional statistics. There has been a great deal of work on regularized estimator under the high-dimensional scaling; see e.g., Van de Geer (2008). Closer to the spirit of our work is the problem of 1-bit compressed sensing (Plan and Vershynin, 2013; Bhaskar and Javanmard, 2015). In this problem, linear measurements are observed for an unknown parameter of interest but only the sign of these measurements are observed. Note that in our problem, the seller is involved in both the learning task and also the policy design. Specifically, he should decide on the prices, which directly affect collected revenue and also indirectly influence the difficulty of the learning task. The market values are then compared with the posted prices, in contrast to 1-bit compressed sensing where the measurements are compared with zero (sign information). In addition, the pricing problem has an online nature while the 1-bit compressed sensing is mostly studied for offline setting. Finally, note that prices are set based on customer's purchase behavior, and hence introduce dependency among the collected information about the model parameters.

### 1.3. Notations

For a vector $v, \operatorname{supp}(v)$ represents the positions of nonzero entries of $v$. Further, for a vector $v$ and a subset $J, v_{J}$ is the restriction of $v$ to indices in $J$. We write $\|v\|_{p}$ for the standard $\ell_{p}$ norm of a vector $v$, i.e., $\|v\|_{p}=\left(\sum_{i}\left|v_{i}\right|^{p}\right)^{1 / p}$ and $\|v\|_{0}$ for the umber of nonzero entries of $v$. If the subscript $p$ is omitted, it should be deemed as $\ell_{2}$ norm. For two vectors $a, b \in \mathbb{R}^{d}$, the notation $a \cdot b=\sum_{i=1}^{d} a_{i} b_{i}$ represents the standard inner product. For two functions $f(n)$ and $g(n)$, the notation $f(n)=O(g(n))$ means that $f$ is bounded above by $g$ asymptotically, namely, $f(n) \leq C g(n)$ for some fixed positive constant $C>0$. Throughout, $\phi(x)=e^{-x^{2} / 2} / \sqrt{2 \pi}$ is the Gaussian density and $\Phi(x) \equiv \int_{-\infty}^{x} \phi(u) \mathrm{d} u$ is the Gaussian distribution.

## 2. Choice Model

We consider a seller, who has a product for sale in each period $t=1,2, \cdots, T$, where $T$ denotes the length of the horizon and may be unknown the to the seller. Each product is represented by an observable vector of features (covariates) $x_{t} \in \mathcal{X} \subseteq \mathbb{R}^{d}$. Products may vary across periods and we assume that feature vectors $x_{t}$ are sampled independently from a fixed, but a priori unknown, distribution $\mathbb{P}_{X}$, supported on a bounded set $\mathcal{X}$.

The product at time $t$ has a market value $v_{t}=v\left(x_{t}\right)$, which is not observed by the seller and function $v$ is (a priori) unknown. At each period $t$, the seller posts a price $p_{t}$. If $p_{t} \leq v_{t}$, a sale occurs, and the seller collects revenue $p_{t}$. If the price is set higher than the market value, $p_{t}>v_{t}$, no sale occurs and no revenue is obtained. The goal of the seller is to design a pricing policy that maximizes the collected revenue.

We first assume that the market value of a product is a linear function of its covariates, namely

$$
\begin{equation*}
v\left(x_{t}\right)=\theta_{0} \cdot x_{t}+z_{t}, \tag{2}
\end{equation*}
$$

where $a \cdot b$ denotes the inner product of vectors $a$ and $b$. Here, $\left\{z_{t}\right\}_{t \geq 1}$ are idiosyncratic shocks, referred to as noise, which are drawn independently and identically from a distribution with mean zero and cumulative function $F$, with density $f(x)=F^{\prime}(x)$, cf. (Keskin and Zeevi, 2014). The noise can account for the features that are not measured. We generalize our model to non-linear valuation functions in Section 6. Parameter $\theta_{0}$ is a priori unknown to the seller. Therefore, the seller is involved in the realm of exploration-exploitation as he needs to choose between learning $\theta_{0}$ and exploiting what has been learned so far to collect revenue.

Let $y_{t}$ be the response variable that indicates whether a sale has occurred at period $t$ :

$$
y_{t}= \begin{cases}+1 & \text { if } v_{t} \geq p_{t}  \tag{3}\\ -1 & \text { if } v_{t}<p_{t}\end{cases}
$$

Note that the above model can be represented as the following probabilistic model:

$$
y_{t}= \begin{cases}+1 & \text { with probability } 1-F\left(p_{t}-\theta_{0} \cdot x_{t}\right)  \tag{4}\\ -1 & \text { with probability } F\left(p_{t}-\theta_{0} \cdot x_{t}\right)\end{cases}
$$

Our proposed algorithm exploits the structure (sparsity) of the feature space to improve its performance. To this aim, let $s_{0}$ denote the number of nonzero coordinates of $\theta_{0}$, i.e., $s_{0}=\left\|\theta_{0}\right\|_{0}=\sum_{j=1}^{d} \mathbb{I}\left(\theta_{0 j} \neq 0\right)$. We remark that $s_{0}$ is a priori unknown to the seller.

### 2.1. Technical Assumptions

To simplify the presentation, we assume that $\left\|x_{t}\right\|_{\infty} \leq 1$, for all $x_{t} \in \mathcal{X}$, and $\left\|\theta_{0}\right\|_{1} \leq W$ for a known constant $W \geq 1$, where for a vector $u=\left(u_{1}, \ldots, u_{d}\right),\|u\|_{\infty}=\max _{i \in[d]}\left|u_{i}\right|$ denotes the maximum absolute value of its entries and $\|u\|_{1}=\sum_{i=1}^{d}\left|u_{i}\right|$. We denote by $\Omega$ the set of feasible parameters, i.e.,

$$
\begin{equation*}
\Omega=\left\{\theta \in \mathbb{R}^{d+1}:\|\theta\|_{0} \leq s_{0},\|\theta\|_{1} \leq W\right\} . \tag{5}
\end{equation*}
$$

We also make the following assumption on the distribution of noise $F$.
Assumption 1 The function $F(v)$ is strictly increasing. Further, $F(v)$ and $1-F(v)$ are log-concave in $v$.

Log-concavity is a widely-used assumption in the economics literature (Bagnoli and Bergstrom, 2005). Note that if the density $f$ is symmetric and the distribution $F$ is logconcave, then $1-F$ is also log-concave. Assumption 1 is satisfied by several common probability distributions including normal, uniform, Laplace, exponential, and logistic. Note that the cumulative distribution function of all log-concave densities is also log-concave (Boyd and Vandenberghe, 2004).

Our second assumption is on the product feature vectors.
Assumption 2 Product feature vectors are generated independently from a fixed unknown distribution $\mathbb{P}_{X}$ with a bounded support $\mathcal{X} \in \mathbb{R}^{d}$. We denote by $\Sigma=\mathbb{E}\left(x_{t} x_{t}^{\top}\right)$ the second moment matrix of distribution $\mathbb{P}_{X}$, and assume that $\Sigma$ is a positive definite matrix. Namely, all of its eigenvalues are bounded from below by a constant $C_{\min }>0$. We also denote the maximum eigenvalue of $\Sigma$ by $C_{\max }$.

The above assumption holds for many common probability distributions, such as uniform, truncated normal, and in general truncated version of many more distributions. Generally, if $\mathbb{P}_{X}$ is bounded below from zero on an open set around the origin, then it has a positive definite second moment matrix. Let us stress that the seller knows neither the distribution $\mathbb{P}_{X}$, nor its second moment $\Sigma$.

### 2.2. Clairvoyant Policy and Performance Metric

We evaluate the performance of our algorithm using the common notion of regret: the expected revenue loss compared with the optimal pricing policy that knows $\theta_{0}$ in advance (but not the realizations of $\left\{z_{t}\right\}_{t \geq 1}$ ). Let us first characterize this benchmark policy.

Using Eq. (2), the expected revenue from a posted price $p$ is equal to

$$
p \times \mathbb{P}\left(v_{t} \geq p\right)=p\left(1-F\left(p-\theta_{0} \cdot x_{t}\right)\right)
$$

Therefore, using first order conditions, for the optimal posted price, denoted by $p^{*}$, we have

$$
\begin{equation*}
p^{*}\left(x_{t}\right)=\frac{1-F\left(p^{*}-\theta_{0} \cdot x_{t}\right)}{f\left(p^{*}-\theta_{0} \cdot x_{t}\right)} \tag{6}
\end{equation*}
$$

To simplify the presentation, let $p_{t}^{*}=p^{*}\left(x_{t}\right)$ denote the optimal price at time $t$.
We now define $\varphi(v) \equiv v-(1-F(v)) / f(v)$ corresponding to the virtual valuation function commonly used in mechanism design (Myerson, 1981). By Assumption 1, $\varphi$ is injective and hence we can define function $g$ as follows

$$
\begin{equation*}
g(v) \equiv v+\varphi^{-1}(-v) \tag{7}
\end{equation*}
$$

It is easy to verify that $g$ is non-negative. Note that by Eq. (6), for the optimal price we have

$$
\theta_{0} \cdot x_{t}+\varphi\left(p^{*}-\theta_{0} \cdot x_{t}\right)=0
$$

Therefore, by rearranging the terms for the optimal price at time $t$ we have

$$
\begin{equation*}
p_{t}^{*}=g\left(\theta_{0} \cdot x_{t}\right) \tag{8}
\end{equation*}
$$

We can now formally define the regret of a policy. Let $\pi$ be the seller's policy that sets price $p_{t}$ at period $t$, and $p_{t}$ can depend on the history of events up to time $t$. The worst-case regret is defined as:

$$
\begin{equation*}
\operatorname{Regret}_{\pi}(T) \equiv \max _{\substack{\theta_{0} \in \Omega \\ \mathbb{P}_{X} \in Q(\mathcal{X})}} \mathbb{E}\left[\sum_{t=1}^{T}\left(p_{t}^{*} \mathbb{I}\left(v_{t} \geq p_{t}^{*}\right)-p_{t} \mathbb{I}\left(v_{t} \geq p_{t}\right)\right)\right] \tag{9}
\end{equation*}
$$

where the expectation is with respect to the distributions of idiosyncratic noise, $z_{t}$, and $\mathbb{P}_{X}$, the distribution of feature vectors. Also, $Q(\mathcal{X})$ represents the set of probability distributions supported on a bounded set $\mathcal{X}$.

Our algorithm uses the sparsity structure of $\theta_{0}$ and learns the model with order of magnitude less data compared to a structure-ignorant algorithm. In Section 4, we show that our pricing scheme achieves a regret bound of $O\left(s_{0} \log d \cdot \log T\right)$.

Note that the setting with known market noise distribution is equivalent to positing a parametric choice model on the users' behavior. Several parametric choice models have been widely used by practitioners and have been analyzed by researchers. Perhaps the most well-known example is the multinomial logit (MNL) model, which has quite a long history in economic and transportation research (McFadden, 2001), as well as marketing (Hauser and Koppelman, 1979; Gensch and Recker, 1979; Louviere and Woodworth, 1983; Guadagni and Little, 2008). Note that the MNL model corresponds to the case where customers are offered multiple products at each round and the market noise is coming from the Gumbel distribution (Talluri and Van Ryzin, 2006). In the case of single product offered at each round (similar to our setting), this reduces to noise coming from the logistic distribution. Other prominent MNL models include Nested MNL, latent-class MNL, random-parameter mixed MNL, and hierarchical Bayesian MNL (Elshiewy et al., 2017).

In Section 7 we relax the assumption of knowing the market noise distribution.

## 3. A Regularized Maximum Likelihood Pricing (RMLP) Policy

In this section, we present our dynamic pricing policy. Our policy runs in an episodic fashion. Episodes are indexed by $k$ and time periods are indexed by $t$. The length of episode $k$ is denoted by $\tau_{k}$. Throughout episode $k$, we set the prices equal to $p_{t}=g\left(x_{t} \cdot \widehat{\theta}^{k}\right)$ where $\widehat{\theta}^{k}$ denotes the estimate of $\theta_{0}$ which is obtained from the observations $\left\{\left(x_{t}, y_{t}, p_{t}\right)\right\}$ in the previous episode. Note that by Eq. (7), $p_{t}$ is the optimal posted price if $\hat{\theta}^{k}$ was the true underlying parameter of the model.

We estimate $\theta_{0}$ using a regularized maximum-likelihood estimator; see Eq. (10) where the (normalized) negative log-likelihood function for $\theta$ is given by Eq. (11). We note that as a consequence of the log concavity assumption on $F$ and $1-F$, the optimization problem (10) is a convex problem. There is a large toolkit of various optimization methods (e.g., alternating direction method of multipliers (ADMM), fast iterative shrinkage-thresholding algorithm (FISTA), accelerated projected gradient descent, among many others) that can be used to solve this optimization problem. There are also recent developments on distributed solvers for $\ell_{1}$ regularized cost function (Boyd et al., 2011).

Observe that by design, prices posted in the $k$-th episode are independent from the market value noises in this period, i.e., $\left\{z_{t}\right\}_{t=\tau_{k}}^{\tau_{k+1}-1}$. This allows us to estimate $\theta_{0}$ for each

Input: (at time 0) function $g$, regularizations $\lambda_{k}, W$ (bound on $\left\|\theta_{0}\right\|_{1}$ ),
Input: (arrives over time) covariate vectors $\left\{x_{t}\right\}_{t \in \mathbb{N}}$
Output: prices $\left\{p_{t}\right\}_{t \in \mathbb{N}}$
$\tau_{1} \leftarrow 1, p_{1} \leftarrow 0, \widehat{\theta}^{1} \leftarrow 0$
for each episode $k=2,3, \ldots$ do
Set the length of $k$-th episode: $\tau_{k} \leftarrow 2^{k-1}$.
Update the model parameter estimate $\hat{\theta}^{k}$ using the regularized ML estimator obtained
from observations in the previous episode:

$$
\begin{equation*}
\widehat{\theta}^{k}=\underset{\|\theta\|_{1} \leq W}{\arg \min }\left\{\mathcal{L}(\theta)+\lambda_{k}\|\theta\|_{1}\right\}, \tag{10}
\end{equation*}
$$

with

$$
\begin{equation*}
\mathcal{L}(\theta)=-\frac{1}{\tau_{k-1}} \sum_{t=\tau_{k-1}}^{\tau_{k}-1}\left\{\mathbb{I}\left(y_{t}=1\right) \log \left(1-F\left(p_{t}-\theta \cdot x_{t}\right)\right)+\mathbb{I}\left(y_{t}=-1\right) \log \left(F\left(p_{t}-\theta \cdot x_{t}\right)\right)\right\} . \tag{11}
\end{equation*}
$$

5: $\quad$ For each period $t$ during the $k$-th episode, set

$$
\begin{equation*}
p_{t} \leftarrow g\left(\widehat{\theta}^{k} \cdot x_{t}\right) . \tag{12}
\end{equation*}
$$

Algorithm 1: RMLP policy for dynamic pricing
episode separately; see Proposition 10 in Section 8.1. Comparing to policies that use the entire data sale history in making decisions, some remarks are in order:

- Perishability of data: In practical applications, the unknown demand parameters will change over time, raising the concern of perishability of data. Namely, collected data becomes obsolete after a while and cannot be relied on for estimating the model parameters (Keskin and Zeevi, 2016; Javanmard, 2017). Common policies to mitigate this problem (discussed in (Keskin and Zeevi, 2016)) include moving windows and decaying weights which use only recent data to learn the model parameters. By contrast, methods that use the entire historical data, with equal weight, suffers from this problem.
- Simplicity and efficiency: In RMLP policy, estimates of the model parameters are updated only at the first period of each episode ( $\log T$ updates). Further, at each update, the policy uses only the historical data from the previous episode. These two ideas together, not only allow for a neat analysis of the statistical dependency among samples but also decrease the computational cost. Scalability of the pricing policy is indispensable in practical applications as the sales data is collected at an unprecedented rate.
- The impact on regret: By using half of the historical data at each update, our policy loses at most a factor 2 in the total regret. (This becomes clear shortly when we discuss the estimation error rate in terms of number of samples.)
The lengths of episodes in our algorithm increase geometrically ( $\tau_{k}=2^{k-1}$ ), allowing for more accurate estimate of $\theta_{0}$ as the episode index grows. The algorithm terminates at the end of the horizon $(\operatorname{period} T)$, but note that it does not need to know the length of the horizon in advance.

The regularization parameter $\lambda_{k}$ constrains the $\ell_{1}$ norm of the estimator $\widehat{\theta}^{k}$. Selecting the value of $\lambda_{k}$ is of crucial importance as it affects the estimator error. We set it as $\lambda_{k}=O\left(\sqrt{(\log d) / \tau_{k-1}}\right)$. More precisely, define

$$
u_{F} \equiv \sup _{|x| \leq 3 W}\left\{\max \left\{\log ^{\prime} F(x),-\log ^{\prime}(1-F(x))\right\}\right\}
$$

where the derivatives are w.r.t. $x$. By the log-concavity property of $F$ and $1-F$, we have

$$
\begin{equation*}
u_{F}=\max \left\{\log ^{\prime} F(-2 W),-\log ^{\prime}(1-F(2 W))\right\} . \tag{13}
\end{equation*}
$$

Hence, $u_{F}$ captures the steepness of $\log F$. In order to minimize the regret, we run the RMLP policy with

$$
\begin{equation*}
\lambda_{k}=4 u_{F} \sqrt{\frac{\log d}{\tau_{k-1}}} . \tag{14}
\end{equation*}
$$

Note that exploration and exploitation tasks are mixed in our algorithm. In the beginning of each episode, we use what is learned from previous episode to improve the estimation of $\theta_{0}$ and then we exploit this estimate throughout the current episode to incur little regret. Meanwhile, the observations gathered in the current episode are used to update our estimate of $\theta_{0}$ for the next episode.

The knowledge of the market noise distribution is used by the RMLP policy, in estimating the model parameters (constructing the log-likelihood function and choosing the regularization $\lambda_{k}$ ) as well as in setting the prices. (The pricing function $g$ is defined based on the noise distribution $F$.)

In the next section we analyze the regret of the RMLP policy.

## 4. Regret Analysis

The following theorem bounds the regret of our dynamics pricing policy.
Theorem 3 (Regret Upper Bound) Suppose Assumptions 1 and 2 hold and the . Then, the regret of the RMLP policy is of $O\left(\frac{W}{C_{\min }^{2}} s_{0} \log d \cdot \log T\right)$.
Let us stress again that the RMLP policy does not need to know the sparsity $s_{0}$, although its regret scales linearly with $s_{0}$. Further, the RMLP is assumed to know the market noise distribution $F .^{1}$ Below we provide an outline for the proof of Theorem 3 and defer its complete proof to Section 8.1.

[^1]1. In RMLP, the updates in the model parameter estimation only occurs at the beginning of each episode, with using only the samples collected in the previous episode. Therefore, the prices posted in each episode are independent from the market value noises in that episode. This observation also verifies that $\mathcal{L}(\theta)$ given by (11), is indeed the negative log-likelihood of the samples collected in $k$-th episode. Note that this independence is not a mere serendipity, rather it holds because of the specific design of RMLP policy. Using this property, we use tools from high-dimensional statistics to bound the estimation error. To bound the error term $\left\|\theta^{k}-\theta_{0}\right\|_{2}$, we compare the function values $\mathcal{L}\left(\theta^{k}\right)$ and $\mathcal{L}\left(\theta_{0}\right)$. The main challenge here is that $\mathcal{L}(\theta)$ is not strictly convex in $\theta_{.}^{2}$ Hence, there can be, in principle, parameter vectors $\theta_{1}$ and $\theta_{2}$ that are close to each other and nevertheless the values of function $\mathcal{L}$ at these points are far from each other.

To cope with this challenge, we show that a so-called restricted eigenvalue condition holds for the feature products. This notion implies that $\mathcal{L}(\theta)$ is strictly convex on the set of sparse vectors. ${ }^{3}$ Using the restricted eigenvalue condition, we show the following $\ell_{2}$ error for the regularized log-likelihood estimate in the $k$-th episode, $\widehat{\theta}^{k}$, holds true

$$
\left\|\widehat{\theta}^{k}-\theta_{0}\right\|_{2}=O\left(\sqrt{s_{0}} \lambda_{k}\right)=O\left(\sqrt{\frac{s_{0} \log d}{\tau_{k-1}}}\right) .
$$

As expected, the estimate gets more accurate as the episode's length increases; see Section 8.1 for more details.
2. For any $p \geq 0$, denote by $\operatorname{rev}_{t}(p)=p\left(1-F\left(p-x_{t} \cdot \theta_{0}\right)\right)$, the expected revenue under price $p$. We bound regt in terms of $\operatorname{rev}_{t}\left(p_{t}^{*}\right)-\operatorname{rev}_{t}\left(p_{t}\right)$. Since $p_{t}^{*} \in \arg \max \left\{\operatorname{rev}_{t}(p)\right\}$, we have $\operatorname{rev}_{t}^{\prime}\left(p_{t}^{*}\right)=0$, and by Taylor expansion of $\operatorname{rev}_{t}$ around $p_{t}^{*}$, we obtain $\operatorname{rev}_{t}\left(p_{t}^{*}\right)-\operatorname{rev}_{t}\left(p_{t}\right)=$ $O\left(\left(p_{t}^{*}-p_{t}\right)^{2}\right)$.
3. For $t$ in the $k$-th episode, namely $\tau_{k-1} \leq t \leq \tau_{k}-1$, we have

$$
p_{t}^{*}-p_{t}=g\left(\theta_{0} \cdot x_{t}\right)-g\left(\widehat{\theta}^{k} \cdot x_{t}\right) \leq\left|\left(\theta_{0}-\widehat{\theta}^{k}\right) \cdot x_{t}\right|
$$

which follows by showing that $g$ is 1 -Lipschitz. Further, by Assumption 2 (without loss of generality assume $C_{\max }>1$ ), we have

$$
\mathbb{E}\left[\left(\left(\theta_{0}-\widehat{\theta}^{k}\right) \cdot x_{t}\right)^{2}\right]=\mathbb{E}\left[\left(\theta_{0}-\widehat{\theta}^{k}\right)^{\top} \Sigma\left(\theta_{0}-\widehat{\theta}^{k}\right)\right] \leq C_{\max } \mathbb{E}\left[\left\|\widehat{\theta}^{k}-\theta_{0}\right\|_{2}^{2}\right]
$$

where the equality holds because $x_{t}$ is independent of $\widehat{\theta}^{k}$, and the last step follows since the maximum eigenvalue of the second moment matrix $\Sigma=\mathbb{E}\left(x_{t} x_{t}^{\top}\right)$ is at most $C_{\text {max }}>0$.
2. Note that $\nabla_{\theta}^{2} \mathcal{L}=\left(-1 / \tau_{k-1}\right) \sum_{t=\tau k-1}^{\tau_{k}-1}\left(\partial^{2} / \partial_{u_{t}}^{2} \mathcal{L}\right) x_{t} x_{t}^{\top}$, where $u_{t}=p_{t}-\theta \cdot x_{t}-\alpha_{0}$. Therefore, $\nabla_{\theta}^{2} \mathcal{L}$ is a $d \times d$ matrix of rank at most $\tau_{k}-\tau_{k-1}$. Hence, $\mathcal{L}(\theta)$ is strictly convex in $\theta$ only if $\tau_{k}-\tau_{k-1} \geq d$. However, since we are not updating our estimates in the middle of an episode, episodes of length $d$ yield the regret to scale linearly in $d$, which is not desired.
3. It is strictly convex over the set of $s_{0}$ sparse vectors in $d$-dimension if the number of samples is above $c s_{0} \log d$ for a suitable constant $c>0$.

Let reg ${ }_{t}$ be the regret occurred at step $t$. Combining the above bounds (step 2 and 3 ), we arrive at $\mathbb{E}\left[\mathrm{reg}_{t}\right]=O\left(s_{0}(\log d) / \tau_{k-1}\right)$. Therefore, the cumulative expected regret in episode $k$ works out at $O\left(s_{0} \log d\right)$. Since the length of episodes increase geometrically, there are $O(\log T)$ episodes by time $T$. This implies that the total expected regret by time $T$ is $O\left(s_{0} \log d \log T\right)$.

### 4.1. Comparison with the "Common" Regret of Bound $\Omega(\sqrt{T})$

There is an often-seen regret bound $\Omega(\sqrt{T})$ in the literature of online decision making, which can be improved to a logarithmic regret bound if some type of "separability assumption" holds true (Dani et al., 2008; Abbasi-Yadkori et al., 2012). Separability assumption posits that there is a positive constant gap between the rewards of the best and the second best actions. In our framework, the parameter $\theta$ belongs to a continuous set in $\mathbb{R}^{d+1}$ and therefore the separability assumption cannot be enforced because by choosing $\theta$ arbitrary close to $\theta_{0}$, one can obtain suboptimal (but arbitrary close to optimal) reward. However, the regret of our policy scales only logarithmically with $T$. Here, we contrast our logarithmic regret bound with the folklore lower bound $\Omega(\sqrt{T})$, to build further insight on our results.

Uninformative prices and $\Omega(\sqrt{T})$ lower-bound. We focus on (Broder and Rusmevichientong, 2012) which has a close framework to ours in that it considers a dynamic pricing policy from purchasing decisions and presents a pricing policy based on maximum likelihood estimation with regret $O(\sqrt{T})$. Adopting their notation, it is assumed that market values $v_{t}$ are independent and identically distributed random variables coming from a distribution function that belongs to some family parametrized by $\mathbf{z}$. Denote by $d(p ; \mathbf{z})$ the demand curve. This curve determines the probability of a purchase at a given price, i.e., $d(p ; \mathbf{z})=\mathbb{P}_{\mathbf{z}}\left(v_{t} \geq p\right)$. (Broder and Rusmevichientong, 2012) show that the worst-case regret of any pricing policy must be at least $\Omega(\sqrt{T})$ (see Theorem 3.1 therein). The bound is proved by considering a specific family of demand curves $d(p ; \mathbf{z})$, such that all demand curves in this family intersect at a common price. Further, the common price is the optimal price for a specific choice of parameter $\mathbf{z}_{0}$, i.e, $p^{*}\left(\mathbf{z}_{0}\right){ }^{4}$ Therefore, the price $p^{*}\left(\mathbf{z}_{0}\right)$ is "uninformative" since no policy can gain information about the demand parameter $\mathbf{z}$, while pricing $p^{*}\left(\mathbf{z}_{0}\right)$. The idea behind the derived lower bound for the worst-case regret is that for a policy to learn the underlying demand curve fast enough, it must necessarily choose prices that are away from (the uninformative) price $p^{*}\left(\mathbf{z}_{0}\right)$ and this leads to a large regret when the true demand curve is indeed $\mathbf{z}_{0}$.

Intuition behind our results. In contrast to the previous case, for our framework there is no such uninformative price. First, note that the for a choice model with parameters $\theta_{0}=\left(\theta_{0}, \alpha_{0}\right)$, the demand curve at time $t$ is given by

$$
d_{t}\left(p ; \theta_{0}\right)=1-F\left(p-\theta_{0} \cdot x_{t}\right),
$$

For $n \geq 1$, we define the aggregate demand function up to time $n$ as $d_{1}^{n}=\left(d_{1}, d_{2}, \ldots, d_{n}\right)$. In the following, we argue that under our setting, there is no uninformative price. For any

[^2]price $p$ and any $\theta_{1}, \theta_{2}$, we have
\[

$$
\begin{aligned}
\frac{1}{n}\left\|d_{1}^{n}\left(p, \theta_{1}\right)-d_{1}^{n}\left(p, \theta_{2}\right)\right\|_{2}^{2} & =\frac{c^{2}}{n} \sum_{\ell=1}^{n}\left(\left(\theta_{1}-\theta_{2}\right) \cdot x_{t}\right)^{2} \\
& =\frac{c^{2}}{n}\left\|X\left(\theta_{1}-\theta_{2}\right)\right\|_{2}^{2},
\end{aligned}
$$
\]

where $X$ is the matrix with rows $x_{\ell}$, for $1 \leq \ell \leq t$. We also used the fact that $f(z) \geq c>0$ for some constant $c$ because $F$ is strictly increasing by Assumption 1. As we show in Appendix A, for $n \geq c_{0} s_{0} \log d$ (with $c_{0}$ a proper constant), $X$ satisfy a so-called "restricted eigenvalue", by which we have

$$
\begin{equation*}
\frac{1}{n}\left\|X\left(\theta_{1}-\theta_{2}\right)\right\|_{2}^{2} \geq \frac{C_{\min }}{2}\left\|\theta_{1}-\theta_{2}\right\|_{2}^{2} \tag{15}
\end{equation*}
$$

Therefore, for any fixed price $p$, if we vary the demand parameters $\theta_{1}$ to some other value $\theta_{2}$, then the aggregate demand at price $p$ also changes by an amount proportional to $\left\|\theta_{1}-\theta_{2}\right\|_{2}$. Hence, any price in this setting is informative about the model parameters.

To build further insight, let us consider a more general choice model, where the utility of the customer from buying a product with feature vectors $x_{t}$ at price $p$ is given by

$$
\begin{equation*}
u\left(x_{t}\right)=\theta_{0} \cdot x_{t}-\beta_{0} p+z_{t}, \tag{16}
\end{equation*}
$$

where $\theta_{0}, \beta_{0}$ are unknown model parameters and $z_{t}$ is the noise term. The customer buys the product iff $u\left(x_{t}\right) \geq 0$. Note that the model we studied in this paper (see Equation (3)) is a special case when the price sensitivity $\beta_{0}$ is known and hence can be normalized to 1 . We next argue that in case of unknown $\beta_{0}$, the uninformative prices do exist and hence the $\Omega(\sqrt{T})$ is still in place.

To see this, let $\theta_{0, i}=0$ for $2 \leq i \leq d$. Fix arbitrary $a$, and set $\beta_{0}=g(a)-a+\theta_{0,1}$. Also, let $x_{t, 1}=1$, for all $t$. Then, the demand curves will be unaltered over time and are given by

$$
d_{t}(p, \theta)=1-F\left(\beta_{0} p-\theta_{0,1}\right)=1-F\left(\left(g(a)-a+\theta_{0,1}\right) p-\theta_{0,1}\right) .
$$

We next verify that $p^{*}=1$ is the optimal price for the specific choice of $\theta_{0,1}=a$. In this case, the expected revenue form a posted price $p$ is given by

$$
\begin{equation*}
p \times d_{t}(p, \theta)=p(1-F(g(a) p-a)) . \tag{17}
\end{equation*}
$$

By definition of the pricing function $g$, given by (7), it is easy to see that $p^{*}=1$ sets the gradient of the expected revenue to zero.

Next note that all the demand curves, regardless of value $\theta_{0,1}$, intersect at $p^{*}=1$ (they all have the value $1-F(g(a)-a)$ at this price). Therefore, $p^{*}$ is an uninformative price and no policy can gain information about $\theta_{0,1}$ by pricing at $p^{*}$. However, when $\theta_{0,1}=a$, choosing prices that are away from this informative price leads to a large regret. Prices that are close to $p^{*}$ does not have any information gain, and contrasting these two points, it can be shown that the worst case regret is of order $\Omega(\sqrt{T})$. A formal proof follows the same lines of the proof of (Broder and Rusmevichientong, 2012, Theorem 3.1) and is omitted.

Finally, it is worth noting that the rate of learning demand parameter $\theta_{0}$ is chiefly derived by three factors:

- Non-smoothness of distribution function $F$, as it controls the amount of information obtained about $x_{t} \cdot \theta_{0}$ at each $t$.
- The rate by which the feature vectors $x_{t}$ span the parameter space. This is controlled through the minimum eigenvalue of $\Sigma$, i.e., $C_{\min }$. If $C_{\min }$ is small, the randomly generated features are relatively aligned and one requires larger sample size to estimate $\theta_{0}$ within specified accuracy.
- Complexity of $\theta_{0}$. This is captured through the sparsity measure $s_{0}$.

Later in Proposition 10, we see how these factors contribute to the learning rate of $\theta_{0}$.

### 4.2. Role of $C_{\text {min }}$

In establishing our results, we relied on Assumption 2 which requires the population covariance of features to be positive definite. The lower bound on its eigenvalues, denoted by $C_{\min }$, appears in our regret bound as a factor $1 / C_{\min }^{2}$.

As evident from the proof of Proposition 10, Assumption 1 can be replaced by the weaker restricted eigenvalue condition (Bühlmann and van de Geer, 2011; Candes and Tao, 2007), which is a common assumption in high-dimensional statistical learning. While assumption $C_{\min }>0$ allows for a fast learning rate of model parameters and a regret bound $O(\log T)$, RMLP policy can still provably achieve regret $O(\sqrt{T})$, even when $C_{\min }=0$.

Theorem 4 Suppose that product feature vectors are generated independently from a probability distribution $\mathbb{P}_{X}$ with a bounded support $\mathcal{X} \in \mathbb{R}^{d}$. Also, assume that the choice model parameters $\theta_{0} \in \Omega$, where the set $\Omega$ is given by (5). Under Assumption 1, the regret of RMLP policy is of $O\left(W^{3} \sqrt{(\log d) T}\right)$.

Proof of Theorem 4 is given in Section 8.2. We remark that the regret bound in Theorem 4 is independent of the sparsity $s_{0}$ and scales logarithmically with the feature dimension $d$. From the technical point of view, this is due to the fact that in our analysis instead of bounding the estimation errors $\left\|\theta_{0}-\widehat{\theta}^{k}\right\|_{2}$, we directly bound the prediction errors $\left\|X^{(k-1)}\left(\theta_{0}-\widehat{\theta}^{k}\right)\right\|_{2}$, where $X^{(k-1)}$ is the matrix obtained by putting the feature vectors in episode $(k-1)$ as its rows. We then connect the total regret in each episode to the corresponding prediction error. By this analysis, we are able to remove the requirement $C_{\min }>0$ and obtain a regret that is independent of the sparsity $s_{0}$, at the cost of getting $O(\sqrt{T})$ regret bound.

## 5. Lower Bound on Regret

As discussed in Section 2.2, if the true parameter $\theta_{0}$ is known, the optimal policy (in terms of expected revenue) is the one that sets prices as $p_{t}=g\left(x_{t} \cdot \theta_{0}\right)$. Let $\mathcal{H}_{t}=$ $\left\{x_{1}, x_{2}, \ldots, x_{t}, z_{1}, z_{2}, \ldots, z_{t}\right\}$ denote the history set up to time $t$, and recall that $\Omega$ denotes the set of feasible parameters, i.e., $\Omega=\left\{\theta \in \mathbb{R}^{d}:\|\theta\|_{0} \leq s_{0},\|\theta\|_{1} \leq W\right\}$. We consider the following set of policies, П:

$$
\begin{equation*}
\Pi=\left\{\pi: \pi\left(p_{t}\right)=g\left(x_{t} \cdot \theta_{t}\right), \text { for some } \theta_{t} \in \Omega, \text { such that } \theta_{t} \text { is } \mathcal{H}_{t-1} \text {-measurable }\right\} . \tag{18}
\end{equation*}
$$

Here $\pi\left(p_{t}\right)$ denotes the price posted by policy $\pi$ at time $t$.
We provide a lower bound on the achievable regret by any policy in set $\Pi$. Indeed this lower bound applies to an oracle who fully observes the market values after the price is either accepted or rejected. Compared to our setting, where the seller observes only the binary feedbacks (purchase/no purchase), this oracle appears exceedingly powerful at first sight but surprisingly, the derived lower bound matches the regret of our dynamic policy, up to a logarithmic factor.

Theorem 5 Consider linear model (2) where the market values $v\left(x_{t}\right), 1 \leq t \leq T$, are fully observed. We further assume that market value noises are generated as $z_{t} \sim \mathrm{~N}\left(0, \sigma^{2}\right)$. Let $\Pi$ be the set of policies given by (18). Then, there exists constant $C^{\prime}>0$ (depending on $W$ and $\sigma$ ), such that the following holds true for all $T \in \mathbb{N}$.

$$
\begin{equation*}
\min _{\pi \in \Pi} \operatorname{Regret}_{\pi}(T) \geq C^{\prime}\left\{s_{0} \log \left(\frac{T}{s_{0}}\right)+\min \left[\frac{T}{s_{0}}, s_{0} \log \left(\frac{d}{s_{0}}\right)\right]\right\} . \tag{19}
\end{equation*}
$$

In the following we give an outline for the proof of Theorem 5, summarizing its main steps and defer the complete proof to Section 8.3.

1. We derive a lower bound for regret in terms of the minimax estimation error. Specifically, for $t \in \mathbb{N}$, let

$$
\begin{equation*}
\operatorname{reg}_{t} \equiv p_{t}^{*} \mathbb{I}\left(v_{t} \geq p_{t}^{*}\right)-p_{t} \mathbb{I}\left(v_{t} \geq p_{t}\right) \tag{20}
\end{equation*}
$$

be the regret at period $t$. We show that

$$
\begin{equation*}
\max _{\theta_{0} \in \Omega} \mathbb{E}\left(\mathrm{reg}_{t}\right) \geq c \max _{\theta_{0} \in \Omega} \mathbb{E}\left\{\min \left(\left\|\theta_{t}-\theta_{0}\right\|_{2}^{2}, C\right)\right\} \tag{21}
\end{equation*}
$$

for some constants $c, C>0$.
2. Let $\theta_{1}^{T}=\left(\theta_{t}\right)_{t=1}^{T}$ and define $\mathrm{d}\left(\theta_{1}^{T}, \theta\right) \equiv \sum_{t=1}^{T} \min \left(\left\|\theta_{t}-\theta\right\|_{2}^{2}, C\right)$. We use a standard argument (Le Cam's method) that relates the minimax $\ell_{2}$ risk, $\min _{\theta_{1}^{T}} \max _{\theta_{0} \in \Omega} \mathbb{E}\left[\mathrm{~d}\left(\theta_{1}^{T}, \theta_{0}\right)\right]$, in terms of the error in multi-way hypothesis problem (Tsybakov, 2008). We first construct a maximal set of points in $\Omega$, such that minimum pairwise distances among them is at least $\delta$. (Such set is usually referred to as a $\delta$-packing in the literature). Here $\delta$ is a free parameter to be determined in the proof. We then use a standard reduction to show that any estimator with small minimax risk should necessarily solve a hypothesis testing problem over the packing set, with small error probability. More specifically, suppose that nature chooses one point from the packing set uniformly at random and conditional on nature's choice of the parameter vector, say $\theta_{0}$, the market value are generated according to $\left\langle x_{t}, \theta_{0}\right\rangle+z_{t}$ with $z_{t} \sim \mathrm{~N}\left(0, \sigma^{2}\right)$. The problem is reduced to lower bounding the error probability in distinguishing $\theta_{0}$ among the candidates in the packing set using the observed market values.
3. We apply Fano's inequality from information theory to lower bound the probability of error (Tsybakov, 2008). The Fano's bound involves the logarithm of the cardinality of the $\delta$-packing set as well as the mutual information between the observations (market
values) and the random parameter vector $\theta_{0}$ chosen uniformly at random from the packing set. Le Cam's method is used to derive minimal risk lower bound for an estimator $\widehat{\theta}$, while here we have a sequence of estimators and need to adjust the Le Cam's method to get the lower bound for $\mathrm{d}\left(\theta_{1}^{T}, \theta_{0}\right)$.

## 6. Nonlinear Valuation Function

In previous sections, we focused exclusively on linear valuation function given by Eq (2). Here, we extend our results and assume that the market valuations are modeled by a nonlinear function that depends on products' features and an independent noise term. Specifically, the market value of a product with feature vector $x_{t}$ is given by

$$
\begin{equation*}
v\left(x_{t}\right)=\psi\left(\theta_{0} \cdot \phi\left(x_{t}\right)+z_{t}\right), \tag{22}
\end{equation*}
$$

where the original features $x_{t}$ are transformed by a feature mapping $\phi: \mathbb{R}^{d} \mapsto \mathbb{R}^{d}$, and function $\psi: \mathbb{R} \mapsto \mathbb{R}$ is a general function that is log-concave and strictly increasing. Important examples of this model include $\log$-log model $\left(\psi(x)=e^{x}, \phi(x)=\ln (x)\right)$, semi-log model $\left(\psi(x)=e^{x}, \phi(x)=x\right)$, and logistic model $\left(\psi(x)=e^{x} /\left(1+e^{x}\right), \psi(x)=x\right)$.

Model (22) allows us to capture correlations and non-linear dependencies on the features. We next state our assumption on the feature mapping $\phi$ and then discuss our dynamic pricing policy and its regret bound for the general setting (22).

Assumption 6 Let $p_{X}$ be an (unknown) distribution from which the original features $x_{t}$ are sampled independently. Suppose that the feature mapping $\phi$ has continuous derivative and denote by $\Sigma_{\phi} \equiv \mathbb{E}\left(\phi(x) \cdot \phi(x)^{\top}\right)$ the second moment matrix of $\phi(x)$ under $\mathbb{P}_{X}$. We assume that there exist constants $C_{\min }$ and $C_{\max }$ such that for every eigenvalue $\sigma$ of $\Sigma_{\phi}$, we have $0<C_{\min } \leq \sigma<C_{\max }<\infty$.

Invoking Assumption 1, $\mathbb{P}_{X}$ has a bounded support $\mathcal{X}$ and since $\phi$ has continuous derivative, it is Lipschitz on $\mathcal{X}$ and hence the image of $\mathcal{X}$ under $\phi$ remains bounded. Therefore, the new features $\phi\left(x_{t}\right)$ are also sampled independently from a bounded set. The condition on $\Sigma_{\phi}$ is analogous to that on $\Sigma$, as required by Assumption 2 for the linear setting.

Based on feature mapping $\phi$, validity of Assumption 6 may depend on all moments of distribution $\mathbb{P}_{X}$. We provide an alternative to this assumption, which only depends on feature mapping $\phi$ and the second moment of $\mathbb{P}_{X}$. In stating the assumption, we use the notation $\mathcal{D}_{\phi}$ to denote the derivative matrix of a feature mapping $\phi$. Precisely, for $\phi=$ $\left(\phi_{1}, \ldots, \phi_{d}\right)$, with $\phi_{i}$ real-valued function defined on $\mathbb{R}^{d}$, we write $\mathcal{D}_{\phi}=\left(\partial \phi_{i} / \partial x_{j}\right)_{1 \leq i \leq j \leq d}$.

Assumption 7 (alternative to Assumption 6) Suppose that feature mapping $\phi$ has continuous derivative and its derivative $\mathcal{D}_{\phi}(x)$ is full-rank for almost all $x$. In addition, there exist constants $C_{\min }$ and $C_{\max }$ such that for every eigenvalue $\sigma$ of covariance $\Sigma$, we have $0<C_{\min } \leq \sigma<C_{\max }<\infty$.

Recall that the noise terms $\left\{z_{t}\right\}_{t \geq 1}$ are drawn independently and identically from a distribution with cumulative function $F$ and density $f(x)$. Let $\lambda(v)=f(v) /(1-F(v))$ be the hazard rate function for distribution $F$. For a log-concave function $\psi$, we define

$$
\begin{equation*}
g_{\psi}^{-1}(v) \equiv v-\lambda^{-1}\left(\frac{\psi^{\prime}(v)}{\psi(v)}\right) . \tag{23}
\end{equation*}
$$

Input: (at time 0) function $g$, regularizations $\lambda_{k}, W$ (bound on $\left\|\theta_{0}\right\|_{1}$ )
Input: (arrives over time) covariate vectors $\left\{\tilde{x}_{t}=\phi\left(x_{t}\right)\right\}_{t \in \mathbb{N}}$
Output: prices $\left\{p_{t}\right\}_{t \in \mathbb{N}}$
$\tau_{1} \leftarrow 1, p_{1} \leftarrow 0, \widehat{\theta}^{1} \leftarrow 0$
for each episode $k=2,3, \ldots$ do
Set the length of $k$-th episode: $\tau_{k} \leftarrow 2^{k-1}$.
Update the model parameter estimate $\widehat{\theta}^{k}$ using the regularized ML estimator obtained
from observations in the previous episode:

$$
\begin{equation*}
\widehat{\theta}^{k}=\underset{\|\theta\|_{1} \leq W}{\arg \min }\left\{\mathcal{L}(\theta)+\lambda_{k}\|\theta\|_{1}\right\}, \tag{24}
\end{equation*}
$$

where $\mathcal{L}(\theta)$ is given by:

$$
\begin{align*}
\mathcal{L}(\theta)=-\frac{1}{\tau_{k-1}} \sum_{t=\tau_{k-1}}^{\tau_{k}-1} & \left\{\mathbb{I}\left(y_{t}=1\right) \log \left(1-F\left(\psi^{-1}\left(p_{t}\right)-\theta \cdot \tilde{x}_{t}\right)\right)\right. \\
& \left.+\mathbb{I}\left(y_{t}=-1\right) \log \left(F\left(\psi^{-1}\left(p_{t}\right)-\theta \cdot \tilde{x}_{t}\right)\right)\right\} \tag{25}
\end{align*}
$$

5: $\quad$ For each period $t$ during the $k$-th episode, set

$$
\begin{equation*}
p_{t} \leftarrow \psi\left(g_{\psi}\left(\widehat{\theta}^{k} \cdot \tilde{x}_{t}\right)\right) . \tag{26}
\end{equation*}
$$

## Algorithm 2: RMLP Policy for dynamic pricing under the nonlinear setting

Note that $\psi^{\prime}(v) / \psi(v)=\log ^{\prime} \psi(v)$ and since $\psi$ is log-concave, this term is decreasing. Further, since $1-F$ is log-concave then its hazard rate $\lambda$ is increasing (See proof of Lemma C.1.) Combining these observations, we have that $-\lambda^{-1}\left(\psi^{\prime}(v) / \psi(v)\right)$ is increasing. Consequently,

- Right-hand side of (23) is strictly increasing and hence, $g_{\psi}^{-1}$ is well-defined.
- We have $\left(g_{\psi}^{-1}\right)^{\prime}(v) \geq 1$, for all $v$. This implies that $0<g_{\psi}^{\prime}(v) \leq 1$, for all $v$.

It is worth noting that for $\psi(v)=v$ (linear model), we have $g_{\psi}=g$, where $g$ is defined by (7). Our pricing policy for the nonlinear model is conceptually similar to the linear setting: The policy runs in an episodic manner. During episode $k$, the prices are set as $p_{t}=\psi\left(g_{\psi}\left(\widehat{\theta}^{k} \cdot x_{t}\right)\right)$, where $\widehat{\theta}^{k}$ denotes the estimate of the true parameter vector $\theta_{0}$, using a regularized maximum-likelihood estimator applied to observations in the previous episode, and $\tilde{x}_{t}=\phi\left(x_{t}\right)$.

We describe our (modified) RMLP policy in Algorithm 2. There a few differences between Algorithm 2 and Algorithm 1: Firstly, the features $x_{t}$ are replaced by $\tilde{x}_{t}=\phi\left(x_{t}\right)$. Secondly, in the regularized estimator, prices $p_{t}$ are replaced by $\psi^{-1}\left(p_{t}\right)$. Thirdly, in the last step of algorithm prices are set as $\psi\left(g_{\psi}\left(\widehat{\theta}^{k} \cdot x_{t}\right)\right)$, with $g_{\psi}$ defined by Equation (23).

Our next theorem bounds the regret of our pricing policy (Algorithm 2).

Theorem 8 Let $\psi$ be log-concave and strictly increasing. Suppose that Assumptions 1 and 6 (or its alternative, Assumption 7) hold. Then, regret of the RMLP policy described as Algorithm 2 is of $O\left(s_{0} \log d \cdot \log T\right)$.

Proof of Theorem 8 is given in Appendix 8.4. Here, we summarize its key ingredients.

1. By increasing property of $\psi$, a sale occurs at period $t$ when $z_{t} \geq \psi^{-1}\left(p_{t}\right)-\theta_{0} \cdot \tilde{x}_{t}$. Hence, the log-likelihood estimator for this setting reads as (25). By virtue of Assumption 6 (or its alternative, Assumption 7) we get a similar estimation error for the regularized estimator to the one in Proposition 10.
2. Similar to our derivation for linear setting, we show that the optimal pricing policy that knows $\theta_{0}$ in advance is given by $p_{t}^{*}=\psi\left(g_{\psi}\left(\theta_{0} \cdot \tilde{x}_{t}\right)\right)$, where $g_{\psi}$ is defined based on Equation (23).
3. The difference between the posted price and the optimal price can be bounded as $p_{t}-p_{t}^{*}=\psi\left(g_{\psi}\left(\widehat{\theta}^{k} \cdot \tilde{x}_{t}\right)\right)-\psi\left(g_{\psi}\left(\theta_{0} \cdot \tilde{x}_{t}\right)\right) \leq L\left|\tilde{x}_{t} \cdot\left(\widehat{\theta}^{k}-\theta_{0}\right)\right|$, for a constant $L>0$. This bound is similar to the corresponding bound for the linear setting, and following the same lines of our regret analysis for that case, we get $R(T)=O\left(s_{0} \log d \cdot \log T\right)$.

## 7. Knowledge of Market Noise Distribution

The proposed RMLP policy has assumed that the market noise distribution $F$ is known to the seller. Knowledge of $F$ has been used both in estimating the model parameters $\theta_{0}$ and in setting the prices $p_{t}$. On the other hand, the benchmark policy is also assumed to have access to model parameters and the distribution $F$. Therefore, the regret bound established in Theorem 3 essentially measures how much the seller loses in revenue due to lack of knowledge of the underlying model parameters. In practice, however, the underlying distribution of valuations is not given and this raises the question of distribution-independent pricing policy.

It is worth mentioning that in some applications, although the underlying distribution of valuations is unknown, it belongs to a known class of distributions. For example, lognormal distributions have proved to be a good fit for the distribution of valuations of advertisers in online advertising markets (Edelman et al., 2007; Lahaie and Pennock, 2007; Xiao et al., 2009; Balseiro et al., 2014). In Section 7.1, we consider a setting where the underlying distribution belongs to a known class of log-concave distributions and propose a policy whose regret is $O(\sqrt{T})$. We also argue that no policy can get a better regret bound.

### 7.1. Unknown Distribution from a Known Class

Suppose that the market noises are generated from a log-concave distribution $F_{m, \sigma}$ (e.g., Log-normal), with unknown mean $m$ and unknown variance $\sigma^{2}$. Without loss of generality, we can assume that $m=0$; otherwise, in the valuation model (2), $m$ can be absorbed in the features as the intercept term. We next explain how the RMLP policy can be adapted to this case.

Input: Pricing function $g$ (corresponding to $F_{0,1}$ ), regularizations $\lambda_{k}, W$
Input: (arrives over time) covariate vectors $\left\{x_{t}\right\}_{t \in \mathbb{N}}$
Output: prices $\left\{p_{t}\right\}_{t \in \mathbb{N}}$
$\tau_{1} \leftarrow 1, p_{1} \leftarrow 0, \widehat{\theta}^{1} \leftarrow 0$
for each episode $k=2,3, \ldots$ do
Set the length of $k$-th episode: $\tau_{k} \leftarrow 2^{k-1}$.
4: Update the model parameter estimate $\widehat{\mu}^{k}$ using the regularized ML estimator:

$$
\begin{equation*}
\left(\widehat{\mu}^{k}, \widehat{\beta}^{k}\right)=\underset{\|(\mu / \beta, \beta)\|_{1} \leq W}{\arg \min }\left\{\mathcal{L}(\mu, \beta)+\lambda_{k}\|\mu\|_{1}\right\}, \tag{28}
\end{equation*}
$$

where $\mathcal{L}(\mu, \beta)$ is given by:

$$
\begin{align*}
\mathcal{L}(\mu, \beta)=-\frac{1}{\tau_{k-1}} \sum_{t=\tau_{k-1}}^{\tau_{k}-1}\{ & \mathbb{I}\left(y_{t}=1\right) \log \left(1-F\left(\beta p_{t}-\mu \cdot x_{t}\right)\right) \\
& \left.+\mathbb{I}\left(y_{t}=-1\right) \log \left(F\left(\beta p_{t}-\mu \cdot x_{t}\right)\right)\right\} . \tag{29}
\end{align*}
$$

5: $\quad$ For each period $t$ during the $k$-th episode, set

$$
\begin{equation*}
p_{t} \leftarrow \frac{1}{\widehat{\beta}^{k}} g\left(\widehat{\mu}^{k} \cdot x_{t}\right) . \tag{30}
\end{equation*}
$$

## Algorithm 3: RMLP-2 policy for dynamic pricing

Define $\beta_{0}=1 / \sigma$ and consider the transformation $\tilde{v}_{t}=\beta_{0} v_{t}, \mu_{0}=\beta_{0} \theta_{0}, \tilde{z}_{t}=\beta_{0} z_{t}$. Then, the valuation model (2) can be written as

$$
\begin{equation*}
\tilde{v}_{t}=x_{t} \cdot \mu_{0}+\tilde{z}_{t}, \tag{27}
\end{equation*}
$$

where $\tilde{z}_{t}$ are drawn from $F_{0,1}$. To lighten the notation, we use the shorthand $F \equiv F_{0,1}$. The response variables $y_{t}$ are then given by $y_{t}=\mathbb{I}\left(\tilde{v}_{t} \geq \beta_{0} p_{t}\right)$.

We propose a variant of RMLP policy, called RMLP-2 for this case. The RMLP2 policy is very similar to RMLP except that the maximum log-likelihood is formed to estimate the parameter vector $\left(\mu_{0}, \beta_{0}\right)$ jointly. Likewise, the regularizations are set as $\lambda_{k}=4 u_{F} \sqrt{(\log d) / \tau_{k-1}}$, where $u_{F}$ is defined based on the distribution $F \equiv F_{0,1}$ as in Equation (13). The policy then offers the optimal prices based on the current parameter estimates at the time. Specifically, for episode $k$, we set $p_{t}=\left(1 / \widehat{\beta}^{k}\right) g\left(\widehat{\mu}^{k} \cdot x_{t}\right)$, where the pricing function $g$, given by (7), is defined based on the distribution $F \equiv F_{0,1}$. A formal description of RMLP-2 is given in Algorithm 3.

Our next result bounds the regret of RMLP-2.
Theorem 9 Consider the valuation model (2), where noises $z_{t}$ are generated from a distribution $F_{m, \sigma}$, with unknown mean $m$ and variance $\sigma^{2}$. Assuming that distribution $F_{m, \sigma}$ satisfies Assumption 1, the regret of RMLP-2 policy is of $O(\sqrt{(\log d) T})$. Further, regret of any pricing policy in this case is $\Omega(\sqrt{T})$.

We refer to Section 8.5 for the proof of Theorem 9. As discussed in the proof, the lower bound $\Omega(\sqrt{T})$ applies to this case due to the existence of non-informative prices; See also Section 4.1.

## 8. Proof of Theorems

### 8.1. Proof of Theorem 3

Following step 1 of the proof outline mentioned in Section 4, we consider the problem of estimating $\theta_{0}$ based on observations from previous episode. Before we proceed, let us emphasize once again that the way RMLP is designed, posted prices at each episode are statistically independent from the market noises in that episode. This can be easily observed because $p_{t}=g\left(x_{t} \cdot \widehat{\theta}^{k}\right)$ for $t$ belonging in the $k$-th episode, and $\widehat{\theta}^{k}$ is estimated based on the samples in the ( $k-1$ )-th episode.

We fix $k \geq 1$ and to lighten the notation, we use the indices $1,2, \ldots, n$ to correspond to periods in the $k$-the episode, i.e., $t=\tau_{k}, \tau_{k}+1, \ldots, \tau_{k+1}-1$.

Using probabilistic model (4), $\theta_{0}$ is estimated by solving a regularized maximum likelihood (ML) optimization problem. The (normalized) negative log-likelihood function for $\theta$ reads as

$$
\begin{equation*}
\mathcal{L}(\theta)=-\frac{1}{n} \sum_{t=1}^{n}\left\{\mathbb{I}\left(y_{t}=1\right) \log \left(1-F\left(p_{t}-\theta \cdot x_{t}\right)\right)+\mathbb{I}\left(y_{t}=-1\right) \log \left(F\left(p_{t}-\theta \cdot x_{t}\right)\right)\right\} . \tag{31}
\end{equation*}
$$

Parameter $\theta$ is estimated as the solution of the following program:

$$
\begin{equation*}
\widehat{\theta}=\underset{\|\theta\|_{1} \leq W}{\arg \min } \mathcal{L}(\theta)+\lambda\|\theta\|_{1} . \tag{32}
\end{equation*}
$$

Define $\ell_{F}$ as follows which corresponds to "flatness" of function $\log F$ :

$$
\begin{equation*}
\ell_{F} \equiv \inf _{|x| \leq 3 W}\left\{\min \left\{-\log ^{\prime \prime} F(x),-\log ^{\prime \prime}(1-F(x))\right\}\right\} \tag{33}
\end{equation*}
$$

By Assumption 1, the log-concavity property of $F$ and $1-F$, we have $\ell_{F}>0$.
The next theorem upper bounds the estimation error of the proposed regularized estimator.

Proposition 10 (Estimation Error) Consider linear model (2) with $\theta_{0} \in \Omega$, under Assumptions 1 and 2. Let $\widehat{\theta}$ be the solution of optimization problem (32) with $\lambda \geq 4 u_{F} \sqrt{(\log d) / n}$. Then, there exist positive constants $c_{0}$ and $C$ such that, for $n \geq c_{0} s_{0} \log (d)$, the following inequality holds with probability at least $1-1 / d-2 e^{-n /\left(c_{0} s_{0}\right)}$ :

$$
\begin{equation*}
\left\|\widehat{\theta}-\theta_{0}\right\|_{2}^{2} \leq \frac{16 s_{0} \lambda^{2}}{\ell_{F}^{2} C_{\min }^{2}} \tag{34}
\end{equation*}
$$

We refer to Appendix A for the proof of Proposition 10.
As we see the $\ell_{2}$ estimation error scales linearly with the sparsity level $s_{0}$. As $s_{0}$ increases, the number of parameters to be estimated becomes larger and this makes the estimation problem harder, leading to worse $\ell_{2}$ bound for a fixed number of samples, n. Further,
choosing $\lambda \sim \sqrt{(\log d) / n}$ (where $\sim$ indicates equality up to a constant factor), our $\ell_{2}$ bound scales logarithmically in the dimension of the demand space, $d$. This allows to deal with high-dimensional applications and obtain a regret that scales logarithmically in $d$. Further, the estimation error shrinks as $\sim 1 / n$; getting more samples with fixed value of $s_{0}$ and $d$ leads to better estimation accuracy. Finally, note that for small values of $\ell_{F}$, the $\log$-likelihood function is very flat and there can be, in principle, vectors $\theta$ of $\log$-likelihood value very close to the optimum and nevertheless far from the optimum. In other words, estimation task becomes harder as $\ell_{F}$ gets smaller and this is clearly reflected in the derived estimation bound.

We next use Proposition 10 to bound the expected estimation error.
Corollary 11 Under assumptions of Proposition 10, the following holds true:

$$
\begin{equation*}
\mathbb{E}\left(\left\|\widehat{\theta}-\theta_{0}\right\|_{2}^{2}\right) \leq \frac{16 s_{0} \lambda^{2}}{\ell_{F}^{2} C_{\min }^{2}}+4 W^{2}\left(\frac{1}{d}+2 e^{-n /\left(c_{0} s_{0}\right)}\right) \tag{35}
\end{equation*}
$$

Proof of Corollary 11 is straightforward and is omitted.
In the next proposition, we improve bound (35) for $n \geq c_{1} d$, for a constant $c_{1}>0$. As we will see, the following result is useful to develop sharper upper bound for regret of RMLP policy.

Proposition 12 Under assumptions of Proposition 10, there exist constants $c, c_{1}>0$, such that for $n \geq c_{1} d$, the following holds true:

$$
\begin{equation*}
\mathbb{E}\left(\left\|\widehat{\theta}-\theta_{0}\right\|_{2}^{2}\right) \leq \frac{16\left(s_{0}+1\right) \lambda^{2}}{\ell_{F}^{2} C_{\min }^{2}}+4 W^{2} e^{-c n^{2}} \tag{36}
\end{equation*}
$$

Proposition 12 is proved in Appendix B.
We next establish some useful properties of the virtual valuation function $\varphi$ and the price function $g$.

Lemma 13 If $1-F$ is log-concave, then the virtual valuation function $\varphi$ is strictly monotone increasing.

Lemma 14 If $1-F$ is log-concave, then the price function $g$ satisfies $0<g^{\prime}(v)<1$, for all values of $v \in \mathbb{R}$.

Proofs of Lemma 13 and 14 are given in Appendix C. 1 and C.2, respectively.
Given that $\left\|\widehat{\theta}^{k}\right\|_{1} \leq W$ and $\left|x_{t} \cdot \widehat{\theta^{k}}\right| \leq W$ for all $t, k$,

$$
\begin{equation*}
p_{t}=g\left(x_{t} \cdot \widehat{\theta}^{k}\right) \leq 2\left|x_{t} \cdot \widehat{\theta}^{k}\right| \leq 2 W, \tag{37}
\end{equation*}
$$

where in the first inequality we used the fact that $\varphi(v)$ is increasing as per Lemma 13 and hence $g(v)=v+\varphi^{-1}(-v) \leq v+|v| \leq 2|v|$. Similarly, we have $p_{t}^{*} \leq 2 W$ for all $t$.

We are now ready to bound the regret of our policy. For $t \geq 1$, let

$$
\begin{equation*}
\operatorname{reg}_{t} \equiv p_{t}^{*} \mathbb{I}\left(v_{t} \geq p_{t}^{*}\right)-p_{t} \mathbb{I}\left(v_{t} \geq p_{t}\right) \tag{38}
\end{equation*}
$$

be the regret at period $t$. Further, let $\mathcal{H}_{t}=\left\{x_{1}, x_{2}, \ldots, x_{t}, z_{1}, z_{2}, \ldots, z_{t}\right\}$ be the history set, up to time $t$ (more precisely, $\mathcal{H}_{t}$ is the filtration generated by $\left\{x_{1}, x_{2}, \ldots, x_{t}, z_{1}, z_{2}, \ldots, z_{t}\right\}$ ). We also define $\overline{\mathcal{H}}_{t}=\mathcal{H}_{t} \cup\left\{x_{t+1}\right\}$ as the filtration obtained after augmenting by the new feature $x_{t+1}$.

We write

$$
\begin{align*}
\mathbb{E}\left(\operatorname{reg}_{t} \mid \overline{\mathcal{H}}_{t-1}\right) & =\mathbb{E}\left(p_{t}^{*} \mathbb{I}\left(v_{t} \geq p_{t}^{*}\right) \mid \overline{\mathcal{H}}_{t-1}\right)-\mathbb{E}\left(p_{t} \mathbb{I}\left(v_{t} \geq p_{t}\right) \mid \overline{\mathcal{H}}_{t-1}\right) \\
& =p_{t}^{*}\left(1-F\left(p_{t}^{*}-x_{t} \cdot \theta_{0}\right)\right)-p_{t}\left(1-F\left(p_{t}-x_{t} \cdot \theta_{0}\right)\right) . \tag{39}
\end{align*}
$$

$\operatorname{Define}^{\operatorname{rev}_{t}}(p) \equiv p\left(1-F\left(p-x_{t} \cdot \theta_{0}\right)\right)$ as the expected revenue under price $p$. Note that $p_{t}^{*} \in \arg \max \operatorname{rev}_{t}(p)$ and thus $r_{t}^{\prime}\left(p_{t}^{*}\right)=0$. By Taylor expansion,

$$
\begin{equation*}
\operatorname{rev}_{t}\left(p_{t}\right)=\operatorname{rev}_{t}\left(p_{t}^{*}\right)+\frac{1}{2} r_{t}^{\prime \prime}(p)\left(p_{t}-p_{t}^{*}\right)^{2} \tag{40}
\end{equation*}
$$

for some $p$ between $p_{t}$ and $p_{t}^{*}$.
We next show that $\left|r_{t}^{\prime \prime}(p)\right| \leq 2\left(B+W B^{\prime}\right)$, with $B=\max _{v} f(v)$, and $B^{\prime}=\max _{v} f^{\prime}(v)$. To see this, we write

$$
\begin{equation*}
\left|r_{t}^{\prime \prime}(p)\right|=\left|2 f\left(p-x_{t} \cdot \theta_{0}\right)+p f^{\prime}\left(p-x_{t} \cdot \theta_{0}\right)\right| \leq 2 B+2 W B^{\prime} \tag{41}
\end{equation*}
$$

where we use the fact that $p_{t}, p_{t}^{*} \leq 2 W$ and consequently $p \leq 2 W$. We let $C=2 \max \left(B, B^{\prime}\right)$. Then, combining Equations (39), (40), (41), along with 1-Lipschitz property of $g$ gives

$$
\begin{align*}
\mathbb{E}\left(\text { reg }_{t} \mid \overline{\mathcal{H}}_{t-1}\right) & \leq\left(B+W B^{\prime}\right)\left(p_{t}^{*}-p_{t}\right)^{2} \\
& \leq C W\left(p_{t}^{*}-p_{t}\right)^{2} \\
& =C W\left(g\left(\theta_{0} \cdot x_{t}\right)-g\left(\widehat{\theta}^{k} \cdot x_{t}\right)\right)^{2} \\
& \leq C W\left|x_{t} \cdot\left(\theta_{0}-\widehat{\theta}^{k}\right)\right|^{2} . \tag{42}
\end{align*}
$$

Here, in the second inequality we used the fact that $W \geq 1$. Given that $x_{t}$ is independent of $\mathcal{H}_{t-1}$, we have

$$
\begin{equation*}
\mathbb{E}\left(\mathrm{reg}_{t} \mid \mathcal{H}_{t-1}\right) \leq C W\left\langle\widehat{\theta}^{k}-\theta_{0}, \Sigma\left(\widehat{\theta}^{k}-\theta_{0}\right)\right\rangle, \tag{43}
\end{equation*}
$$

where $\Sigma=\mathbb{E}\left(x_{t} x_{t}^{\top}\right)$. Since the maximum eigenvalue of $\Sigma$ is bounded by $C_{\max }$, we obtain

$$
\begin{equation*}
\mathbb{E}\left(\mathrm{reg}_{t}\right)=\mathbb{E}\left(\mathbb{E}\left(\mathrm{reg}_{t} \mid \mathcal{H}_{t-1}\right)\right) \leq C C_{\max } W \mathbb{E}\left(\left\|\widehat{\theta}^{k}-\theta_{0}\right\|_{2}^{2}\right) . \tag{44}
\end{equation*}
$$

Now, since the length of episodes grows exponentially, the number of episodes by period $T$ is logarithmic in $T$. Specifically, $T$ belongs to episode $K=\lfloor\log T\rfloor+1$. Hence,

$$
\begin{equation*}
\operatorname{Regret}(T)=\sum_{k=1}^{K} \operatorname{Regret}(k \text { th Episode }) \tag{45}
\end{equation*}
$$

We bound the total regret over each episode by considering three separate cases:

- $2^{k-2} \leq c_{0} s_{0} \log d:$ Here, $c_{0}$ is the constant in the statement of Proposition 10. In this case, episodes are not large enough to estimate $\theta_{0}$ accurately enough, and thus we use a naive bound on regret. Clearly, by (37), we have $\mathbb{E}\left(\mathrm{reg}_{t}\right) \leq p_{t}^{*} \leq 2 W$. Since the length of $k$ th episode is $2^{k-1} \leq 2 c_{0} s_{0} \log d$, the total regret incurred during episode $k$ is at most $4 c_{0} W s_{0} \log d$.
- $c_{0} s_{0} \log d \leq 2^{k-2} \leq c_{1} d$ : Here, $c_{1}$ is the constant in the statement of Proposition 12. Continuing from Equation (44) and applying Corollary 11 to episode $k$, we obtain

$$
\begin{align*}
\operatorname{Regret}(k \text { th Episode }) & =\sum_{t=\tau_{k}}^{\tau_{k+1}-1} \mathbb{E}\left(\mathrm{reg}_{t}\right) \\
& \leq C C_{\max } W \sum_{t=\tau_{k}}^{\tau_{k+1}-1} \mathbb{E}\left(\left\|\widehat{\theta}^{k}-\theta_{0}\right\|_{2}^{2}\right) \\
& \leq C C_{\max } W\left\{\frac{16 s_{0} \lambda_{k}^{2}}{\ell_{F}^{2} C_{\min }^{2}} \tau_{k}+4 W^{2}\left(\frac{\tau_{k}}{d}+2 \tau_{k} e^{-\tau_{k-1} /\left(c_{0} s_{0}\right)}\right)\right\} \\
& \leq C C_{\max } W\left\{\left(\frac{16 u_{F}}{\ell_{F} C_{\min }}\right)^{2} 2 s_{0} \log d+8 W^{2}\left(2 c_{1}+2 \tau_{k-1} e^{-\frac{\tau_{k-1}}{c_{0} s_{0}}}\right)\right\} \tag{46}
\end{align*}
$$

where in the last step we used $\tau_{k}=2 \tau_{k-1}$ and $\tau_{k}=2^{k-1} \leq 2 c_{1} d$. Therefore, in this case

$$
\begin{equation*}
\operatorname{Regret}(k \text { th Episode }) \leq \frac{C^{\prime} W}{C_{\min }^{2}} s_{0} \log d \tag{47}
\end{equation*}
$$

where $C^{\prime}$ hides various constants in the right-hand side of (46).

- $c_{1} d<2^{k-2}$ : Continuing from Equation (44) and applying Proposition 12 to episode $k$, we obtain

$$
\begin{align*}
\operatorname{Regret}(k \text { th Episode }) & =\sum_{t=\tau_{k}}^{\tau_{k+1}-1} \mathbb{E}\left(\text { reg }_{t}\right) \\
& \leq C C_{\max } W \sum_{t=\tau_{k}}^{\tau_{k+1}-1} \mathbb{E}\left(\left\|\widehat{\theta}^{k}-\theta_{0}\right\|_{2}^{2}\right) \\
& \leq C C_{\max } W\left\{\frac{16\left(s_{0}+1\right) \lambda_{k}^{2}}{\ell_{F}^{2} C_{\min }^{2}} \tau_{k}+4 \tau_{k} W^{2} e^{-c \tau_{k-1}^{2}}\right\} \\
& \leq C C_{\max } W\left\{\left(\frac{16 u_{F}}{\ell_{F} C_{\min }}\right)^{2} 2\left(s_{0}+1\right) \log d+8 W^{2} \tau_{k-1} e^{-c \tau_{k-1}^{2}}\right\}, \tag{48}
\end{align*}
$$

Therefore, in this case

$$
\begin{equation*}
\operatorname{Regret}(k \text { th Episode }) \leq \frac{C^{\prime} W}{C_{\min }^{2}} s_{0} \log d, \tag{49}
\end{equation*}
$$

where $C^{\prime}$ hides various constants in the right-hand side of (48).
Combining the above three cases into Equation (45), we get

$$
\begin{equation*}
\operatorname{Regret}(T) \leq K \frac{C^{\prime} W}{C_{\min }^{2}} s_{0} \log d=O\left(\frac{W}{C_{\min }^{2}} s_{0} \log d \cdot \log T\right) \tag{50}
\end{equation*}
$$

which concludes the proof.

### 8.2. Proof of Theorem 4

By using Equation (43), we have

$$
\begin{equation*}
\mathbb{E}\left(\mathrm{reg}_{t}\right) \leq \frac{C}{2} \mathbb{E}\left(\left\langle\hat{\theta}^{k}-\theta_{0}, \Sigma\left(\widehat{\theta}^{k}-\theta_{0}\right)\right\rangle\right), \tag{51}
\end{equation*}
$$

with $\Sigma=\mathbb{E}\left(x_{t} x_{t}^{\top}\right)$ and $C=2\left(B+W B^{\prime}\right)$, where $B=\max _{v} f(v)$ and $B^{\prime}=\max _{v} f^{\prime}(v)$. Therefore, letting $K=\lfloor\log T\rfloor+1$, we get

$$
\begin{equation*}
\operatorname{Regret}(T)=\sum_{k=1}^{K} \operatorname{Regret}(k \text { th Episode }) \leq \frac{C}{2} \sum_{k=1}^{K} \mathbb{E}\left(\left\langle\hat{\theta}^{k}-\theta_{0}, \Sigma\left(\widehat{\theta}^{k}-\theta_{0}\right)\right\rangle\right) \tau_{k} \tag{52}
\end{equation*}
$$

We next bound the right-hand side of the above bound. Let $X^{(k)} \in \mathbb{R}^{\tau_{k} \times d}$ be the matrix obtained by stacking feature vectors in episode $k$ as rows. By applying bound (91) to samples in episode $(k-1)$, we get that with probability at least $1-1 / d$,

$$
\begin{equation*}
\frac{2 \ell_{F}}{\tau_{k-1}}\left\|X^{(k-1)}\left(\theta_{0}-\widehat{\theta}^{k}\right)\right\|^{2}+2 \lambda_{k}\left\|\widehat{\theta}^{k}\right\|_{1} \leq \lambda_{k}\left\|\hat{\theta}^{k}-\theta_{0}\right\|_{1}+2 \lambda_{k}\left\|\theta_{0}\right\|_{1} \tag{53}
\end{equation*}
$$

Hence,

$$
\begin{equation*}
\frac{2 \ell_{F}}{\tau_{k-1}}\left\|X^{(k-1)}\left(\theta_{0}-\widehat{\theta}^{k}\right)\right\|^{2} \leq \lambda_{k}\left\|\widehat{\theta}^{k}-\theta_{0}\right\|_{1}+2 \lambda_{k}\left\|\theta_{0}\right\|_{1}-2 \lambda_{k}\left\|\hat{\theta}^{k}\right\|_{1} \leq 3 \lambda_{k}\left\|\widehat{\theta}^{k}-\theta_{0}\right\|_{1} \tag{54}
\end{equation*}
$$

For $k \geq 1$, let $S^{(k)} \in \mathbb{R}^{d \times d}$ be the empirical second moment of $X^{(k)}$, and define $E^{(k)}=$ $\Sigma-S^{(k)}$. Then,

$$
\begin{equation*}
\left\langle\widehat{\theta}^{k}-\theta_{0}, \Sigma\left(\widehat{\theta}^{k}-\theta_{0}\right)\right\rangle=\left\langle\widehat{\theta}^{k}-\theta_{0}, S^{(k-1)}\left(\widehat{\theta}^{k}-\theta_{0}\right)\right\rangle+\left\langle\widehat{\theta}^{k}-\theta_{0}, E^{(k-1)}\left(\widehat{\theta}^{k}-\theta_{0}\right)\right\rangle . \tag{55}
\end{equation*}
$$

The first term is bounded using Equation (54) as follows:

$$
\begin{equation*}
\left\langle\widehat{\theta}^{k}-\theta_{0}, S^{(k-1)}\left(\widehat{\theta}^{k}-\theta_{0}\right)\right\rangle=\frac{1}{\tau_{k-1}}\left\|X^{(k-1)}\left(\theta_{0}-\widehat{\theta}^{k}\right)\right\|^{2} \leq \frac{3 \lambda_{k}}{2 \ell_{F}}\left\|\widehat{\theta}^{k}-\theta_{0}\right\|_{1} \leq \frac{3 W}{\ell_{F}} \lambda_{k} \tag{56}
\end{equation*}
$$

with probability at least $1-1 / d$.
The second term can be bounded by virtue of the following lemma, whose proof is deferred to Appendix C. 8

Lemma 15 For any $k \geq 1$ and any vector $v \in \mathbb{R}^{d}$, we have

$$
\left\langle v, E^{(k)} v\right\rangle \leq 3 \sqrt{\frac{\log d}{\tau_{k}}}\|v\|_{1}^{2}
$$

with probability at least $1-8 / d^{2}$.
By Lemma 15, we have

$$
\begin{equation*}
\left\langle\widehat{\theta}^{k}-\theta_{0}, E^{(k-1)}\left(\widehat{\theta}^{k}-\theta_{0}\right)\right\rangle \leq 12 \sqrt{\frac{\log d}{\tau_{k-1}}} W^{2} \tag{57}
\end{equation*}
$$

with probability at least $1-8 / d^{2}$.
Combining Equations (56) and (57), with probability at least $1-9 / d$ we have

$$
\begin{equation*}
\left\langle\widehat{\theta}^{k}-\theta_{0}, \Sigma\left(\widehat{\theta}^{k}-\theta_{0}\right)\right\rangle \leq \frac{3 W}{\ell_{F}} \lambda_{k}+12 \sqrt{\frac{\log d}{\tau_{k-1}}} W^{2}=12\left(\frac{W u_{F}}{\ell_{F}}+W^{2}\right) \sqrt{\frac{\log d}{\tau_{k-1}}} \tag{58}
\end{equation*}
$$

Following a similar argument as in Section 8.1 (see Equation (45) and onwards) we have that the following holds for a suitable constant $C>0$ :

$$
\begin{aligned}
\operatorname{Regret}(T) & \leq C W^{3} \sum_{k=2}^{K} \sqrt{\frac{\log d}{\tau_{k-1}}} \cdot \tau_{k} \\
& \leq \sqrt{2} C W^{3} \sum_{k=2}^{K} \sqrt{(\log d) \tau_{k-1}}=O\left(W^{3} \sqrt{(\log d) T}\right)
\end{aligned}
$$

which completes the proof.

### 8.3. Proof of Theorem 5

The regret benchmark (9) is defined as the maximum gap between a policy and the oracle policy over different $\theta_{0} \in \Omega$ and $p_{X} \in Q(\mathcal{X})$. Without loss of generality, we assume $\mathcal{X}=$ $[-1,1]^{d}$. In order to obtain a lower bound on the regret, it suffices to consider a specific distribution in $Q(\mathcal{X})$. We consider a distribution $p_{X}$ that selects coordinates $x_{i}, 1 \leq i \leq d$, uniformly at random from $\{-1,1\}$ and independent of each other. We further assume $\theta_{0} \in \Omega$, with $\Omega$ given by (5).

Fix an arbitrary policy $\pi$ in family $\Pi$. By definition of policy set $\Pi$, we have $\pi\left(p_{t}\right)=$ $g\left(x_{t} \cdot \theta_{t}\right)$, for some $\theta_{t} \in \Omega$, which is $\mathcal{H}_{t-1}$-measurable. Recalling our notation in the proof of Theorem 3, reg ${ }_{t}$ denotes the regret occurred at step $t$ and by Equations (39), (40), we have

$$
\begin{equation*}
\mathbb{E}\left(\operatorname{reg}_{t} \mid \overline{\mathcal{H}}_{t-1}\right)=\operatorname{rev}_{t}\left(p_{t}^{*}\right)-\operatorname{rev}_{t}\left(p_{t}\right)=-\frac{1}{2} \operatorname{rev}_{t}^{\prime \prime}(p)\left(p_{t}-p_{t}^{*}\right)^{2} \tag{59}
\end{equation*}
$$

for some $p$ between $p_{t}$ and $p_{t}^{*}$.
Our first lemma will be used in lower bounding $\mathbb{E}\left(\right.$ reg $\left._{t} \mid \overline{\mathcal{H}}_{t-1}\right)$.

Lemma 16 There exists a constant $c_{1}>0$ (depending on $W$ and $\sigma$ ) such that, with probability one $e^{5}, r_{t}^{\prime \prime}\left(p_{t}^{*}\right) \leq-c_{1}$, for all $t \geq 1$. Further, there exists constant $\delta>0$ (depending on $W$ and $\sigma$ ) such that $r_{t}^{\prime \prime}(p) \leq-c_{1} / 4$ for $p \in\left[p_{t}^{*}-\delta, p_{t}^{*}+\delta\right]$, with probability one.

Proof of Lemma 16 is given in Appendix C.3.
Continuing from Equation (59), we consider two separate cases:

- $\left|p_{t}-p_{t}^{*}\right| \leq \delta$ : We have $p \in\left[p_{t}^{*}-\delta, p_{t}^{*}+\delta\right]$ and therefore by applying Lemma 16 we obtain

$$
\begin{equation*}
\mathbb{E}\left(\operatorname{reg}_{t} \mid \overline{\mathcal{H}}_{t-1}\right)=\operatorname{rev}_{t}\left(p_{t}^{*}\right)-\operatorname{rev}_{t}\left(p_{t}\right) \geq \frac{c}{8}\left(p_{t}-p_{t}^{*}\right)^{2} . \tag{60}
\end{equation*}
$$

- $\left|p_{t}-p_{t}^{*}\right|>\delta$ : Since function rev ${ }_{t}$ has only one local maximum, namely $p_{t}^{*}$, the function is increasing before $p_{t}^{*}$ and decreasing afterward. Therefore, if $p_{t} \leq p_{t}^{*}-\delta$ then

$$
\begin{equation*}
\operatorname{rev}_{t}\left(p_{t}\right) \leq \operatorname{rev}_{t}\left(p_{t}^{*}-\delta\right)=\operatorname{rev}_{t}\left(p_{t}^{*}\right)+\frac{1}{2} \operatorname{rev}_{t}^{\prime \prime}(p) \delta^{2} \leq \operatorname{rev}_{t}\left(p_{t}^{*}\right)-\frac{c_{1}}{8} \delta^{2}, \tag{61}
\end{equation*}
$$

where $p$ is some point in $\left[p_{t}^{*}-\delta, p_{t}^{*}\right]$ and we applied Lemma 16 in the last step. Similarly, for $p_{t} \geq p_{t}^{*}+\delta$ we obtain

$$
\begin{equation*}
\operatorname{rev}_{t}\left(p_{t}\right) \leq \operatorname{rev}_{t}\left(p_{t}^{*}+\delta\right)=\operatorname{rev}_{t}\left(p_{t}^{*}\right)+\frac{1}{2} \operatorname{rev}_{t}^{\prime \prime}(p) \delta^{2} \leq \operatorname{rev}_{t}\left(p_{t}^{*}\right)-\frac{c_{1}}{8} \delta^{2}, \tag{62}
\end{equation*}
$$

where $p \in\left[p_{t}^{*}-\delta, p_{t}^{*}\right]$ this time. Combining these two inequalities, we get that $\operatorname{rev}_{t}\left(p_{t}^{*}\right)-$ $\operatorname{rev}_{t}\left(p_{t}\right) \geq c_{1} \delta^{2} / 8$, if $\left|p_{t}^{*}-p_{t}\right| \geq \delta$.

Writing the bounds in the two cases together, we get

$$
\mathbb{E}\left(\operatorname{reg}_{t} \mid \overline{\mathcal{H}}_{t-1}\right) \geq \operatorname{rev}_{t}\left(p_{t}^{*}\right)-\operatorname{rev}_{t}\left(p_{t}\right) \geq \begin{cases}\frac{c_{1}}{8}\left(p_{t}-p_{t}^{*}\right)^{2}, & \text { if }\left|p_{t}-p_{t}^{*}\right| \leq \delta,  \tag{63}\\ \frac{c_{1}}{8} \delta^{2}, & \text { if }\left|p_{t}-p_{t}^{*}\right|>\delta\end{cases}
$$

We proceed by relating the lower bound to the error in estimation $\theta_{0}$.

$$
\begin{align*}
\mathbb{E}\left(\text { reg }_{t} \mid \overline{\mathcal{H}}_{t-1}\right) & \geq \frac{c_{1}}{8} \min \left(\left(p_{t}-p_{t}^{*}\right)^{2}, \delta^{2}\right)=\frac{c_{1}}{8} \min \left(\left(g\left(x_{t} \cdot \theta_{t}\right)-g\left(x_{t} \cdot \theta_{0}\right)\right)^{2}, \delta^{2}\right) \\
& \geq \frac{c_{1}}{8} \min \left(c_{2}^{2}\left|x_{t} \cdot\left(\theta_{t}-\theta_{0}\right)\right|^{2}, \delta^{2}\right), \tag{64}
\end{align*}
$$

where we used the fact that by Lemma $14, g^{\prime}(v)>c_{2}$ over the bounded interval $[-W, W]$, for some constant $c_{2}>0$. We recall the definition of history set $\mathcal{H}_{t} \equiv \overline{\mathcal{H}}_{t} \backslash\left\{x_{t+1}\right\}=$ $\left\{x_{1}, x_{2}, \ldots, x_{t}, z_{1}, z_{2}, \ldots, z_{t}\right\}$. Since $\mathcal{H}_{t} \subseteq \overline{\mathcal{H}}_{t}$, by iterated law of expectation, we get

$$
\begin{equation*}
\mathbb{E}\left(\operatorname{reg}_{t} \mid \mathcal{H}_{t-1}\right)=\mathbb{E}\left(\mathbb{E}\left(\operatorname{reg}_{t} \mid \overline{\mathcal{H}}_{t-1}\right) \mid \mathcal{H}_{t-1}\right) \geq \frac{c_{1}}{8} \mathbb{E}\left(\min \left(c_{2}^{2}\left|x_{t} \cdot\left(\theta_{t}-\theta_{0}\right)\right|_{2}^{2}, \delta^{2}\right) \mid \mathcal{H}_{t-1}\right) . \tag{65}
\end{equation*}
$$

Note that $x_{t}$ is independent of $\mathcal{H}_{t-1}$ and $\theta_{t}-\theta_{0}$ is $\mathcal{H}_{t-1}$-measurable.
We use the following lemma to lower bound the right-hand side of (65).
5. The randomness comes from randomness in prices which in turn comes from randomness in features $x_{t}$.

Lemma 17 Let $x \in \mathbb{R}^{d}$ be a random vector such that its coordinates are chosen independently and uniformly at random from $\{-1,1\}$. Further, suppose that $v \in \mathbb{R}^{d}$ and $\delta>0$ are deterministic. Then,

$$
\begin{equation*}
\mathbb{E}\left(\min \left((x \cdot v)^{2}, \delta^{2}\right)\right) \geq 0.1 \min \left(\|v\|_{2}^{2}, \delta^{2}\right) \tag{66}
\end{equation*}
$$

Proof of Lemma 17 is given in Appendix C.4.
Applying Lemma 17 to bound (65), we obtain

$$
\begin{equation*}
\mathbb{E}\left(\operatorname{reg}_{t} \mid \mathcal{H}_{t-1}\right) \geq \frac{c_{1} c_{2}^{2}}{80} \mathbb{E}\left(\min \left(\left\|\theta_{t}-\theta_{0}\right\|_{2}^{2}, \delta^{2} / c_{2}^{2}\right) \mid \mathcal{H}_{t-1}\right) \tag{67}
\end{equation*}
$$

Now, taking expectation from both sides with respect to $\mathcal{H}_{t-1}$, we arrive at

$$
\begin{equation*}
\mathbb{E}\left(\mathrm{reg}_{t}\right) \geq \frac{c_{1} c_{2}^{2}}{80} \mathbb{E}\left(\min \left(\left\|\theta_{t}-\theta_{0}\right\|_{2}^{2}, \delta^{2} / c_{2}^{2}\right)\right) \tag{68}
\end{equation*}
$$

Equation (68) lower bounds the expected regret at each step to the $\ell_{2}$ estimation error.
We continue by establishing a minimax lower bound on $\ell_{2}$-risk of estimation.
Lemma 18 Consider linear model (2), with $\alpha_{0}=0$, and assume that the market values $v\left(x_{t}\right), 1 \leq t \leq T$, are fully observed and the feature vectors are generated according to $p_{X}$, described above. We further assume that the noise in market value is generated as $z_{t} \sim \mathrm{~N}\left(0, \sigma^{2}\right)$. For a sequence of estimators $\theta_{t}$, we let $\theta_{1}^{t}=\left(\theta_{1}, \theta_{2}, \ldots, \theta_{t}\right)$. Then, for any fixed value $C>0$, there exists a nonnegative constant $\widetilde{C}$, depending on $C, \sigma, W$, such that

$$
\begin{equation*}
\min _{\theta_{1}^{T}} \max _{\theta_{0} \in \Omega} \sum_{t=1}^{T} \mathbb{E}\left(\min \left(\left\|\theta_{t}-\theta_{0}\right\|_{2}^{2}, C\right) \mid x_{1}, \ldots, x_{T}\right) \geq \widetilde{C}\left\{s_{0} \log \left(\frac{T}{s_{0}}\right)+\min \left[\frac{T}{s_{0}}, s_{0} \log \left(\frac{d}{s_{0}}\right)\right]\right\} \tag{69}
\end{equation*}
$$

Proof of Lemma 18 is given in Appendix C.5.
We are now ready to lower bound the regret of any policy in $\Pi$.

$$
\begin{align*}
\operatorname{Regret}(T) & \geq \max _{\theta_{0} \in \Omega} \sum_{t=1}^{T} \mathbb{E}\left(\operatorname{reg}_{t}\right) \geq \frac{c_{1} c_{2}^{2}}{80} \sum_{t=1}^{T} \mathbb{E}\left(\min \left(\left\|\theta_{t}-\theta_{0}\right\|_{2}^{2}, \delta^{2} / c_{2}^{2}\right)\right) \\
& \geq \widetilde{C} \frac{c_{1} c_{2}^{2}}{80}\left\{s_{0} \log \left(\frac{T}{s_{0}}\right)+\min \left[\frac{T}{s_{0}}, s_{0} \log \left(\frac{d}{s_{0}}\right)\right]\right\} \tag{70}
\end{align*}
$$

where the last step follows from Lemma 18.

### 8.4. Proof of Theorem 8

Let $\tilde{x}_{t}=\phi\left(x_{t}\right)$ denote the transformed features under the feature-map. Also, let $\tilde{p}_{t}=$ $\psi^{-1}\left(p_{t}\right)$. We first show that Assumption 7 implies Assumption 6, and therefore it suffices to prove the theorem under Assumption 6.

Lemma 19 Suppose that Assumption 1 hold true. Then, Assumption 7 implies Assumption 6 .

Proof of Lemma 19 is given in Appendix C.6.
By Assumption 1, the support of $\mathbb{P}_{X}$ is a bounded set $\mathcal{X}$. Given that $\phi$ has a continuous derivative, it is Lipschitz on the bounded set $\mathcal{X}$ and ergo the image of $\mathcal{X}$ remains bounded under the feature-map $\phi$. Putting differently, features $\tilde{x}_{t}$ are sampled from a bounded set in $\mathbb{R}^{d}$. Without loss of generality, we assume $\left\|\tilde{x}_{t}\right\|_{\infty} \leq 1$. Further, as per Assumption 6 , the covariance of the underlying distribution $\Sigma_{\phi}$ is positive definite with bounded eigenvalues.

On a different note, since $\psi$ is strictly increasing, a sale occurs at period $t$ when $\theta_{0} \cdot \tilde{x}_{t}+$ $z_{t} \geq \psi^{-1}\left(p_{t}\right)=\tilde{p}_{t}$. Therefore the (negative) log-likelihood function for $\theta$ reads as
$\mathcal{L}(\theta)=-\frac{1}{\tau_{k-1}} \sum_{t=\tau_{k-1}}^{\tau_{k}-1}\left\{\mathbb{I}\left(y_{t}=1\right) \log \left(1-F\left(\tilde{p}_{t}-\theta \cdot \tilde{x}_{t}\right)\right)+\mathbb{I}\left(y_{t}=-1\right) \log \left(F\left(\tilde{p}_{t}-\theta \cdot \tilde{x}_{t}\right)\right)\right\}$.
The estimation bound (34) also holds for this setting and the proof goes along the same lines of the proof of Proposition 10, with slight modifications: $(i)$ the features $x_{t}$ and prices $p_{t}$ should be replaced by $\tilde{x}_{t}$ and $\tilde{p}_{t}$. (ii) Quantity $u_{F}$ and $\ell_{F}$ in the statement of Proposition 10 should be set as $M=(1 / 3) g_{\psi}(0)+(2 / 3) W$. This follows from the bounds below

$$
\begin{equation*}
\tilde{p}_{t}=g_{\psi}\left(\tilde{x}_{t} \cdot \widehat{\theta}^{k}\right) \leq g_{\psi}(0)+\left|\tilde{x}_{t} \cdot \widehat{\theta}^{k}\right| \leq g_{\psi}(0)+W, \tag{71}
\end{equation*}
$$

where we used the facts that $g_{\psi}$ is 1-Lipschitz and increasing as explained below Equation (23).

We next characterize the optimal policy when the true parameter $\theta_{0}$ is known. The expected revenue from a posted price $p$ works out at $p\left(1-F\left(\psi^{-1}(p)-\theta_{0} \cdot \tilde{x}_{t}\right)\right)$. Writing this in terms of $\tilde{p}=\psi^{-1}(p)$, the first order condition for the optimal price reads as

$$
\begin{equation*}
\lambda\left(\tilde{p}^{*}-\theta_{0} \cdot \tilde{x}_{t}\right) \equiv \frac{f\left(\tilde{p}^{*}-\theta_{0} \cdot \tilde{x}_{t}\right)}{1-F\left(\tilde{p}^{*}-\theta_{0} \cdot \tilde{x}_{t}\right)}=\frac{\psi^{\prime}\left(\tilde{p}^{*}\right)}{\psi\left(\tilde{p}^{*}\right)}, \tag{72}
\end{equation*}
$$

where $\lambda$ denotes the hazard rate function. Equivalently

$$
\begin{equation*}
\theta_{0} \cdot \tilde{x}_{t}=\tilde{p}^{*}-\lambda^{-1}\left(\frac{\psi^{\prime}\left(\tilde{p}^{*}\right)}{\psi\left(\tilde{p}^{*}\right)}\right) . \tag{73}
\end{equation*}
$$

By definition of function $g_{\psi}$ as per Equation (23), we get $\tilde{p}^{*}=g_{\psi}\left(\theta_{0} \cdot \tilde{x}_{t}\right)$ and thus $p^{*}=\psi\left(g_{\psi}\left(\theta_{0} \cdot \tilde{x}_{t}\right)\right)$.

We are now ready to bound the regret of the algorithm. Similar to Equation (42), we have

$$
\begin{align*}
\mathbb{E}\left(\operatorname{reg}_{t} \mid \overline{\mathcal{H}}_{t-1}\right) \leq \frac{C}{2}\left(p_{t}^{*}-p_{t}\right)^{2} & =\frac{C}{2}\left[\psi\left(g_{\psi}\left(\theta_{0} \cdot \tilde{x}_{t}\right)\right)-\psi\left(g_{\psi}\left(\widehat{\theta}^{k} \cdot \tilde{x}_{t}\right)\right)\right]^{2} \\
& \left.\leq \frac{C}{2} L\left(g_{\psi}\left(\theta_{0} \cdot \tilde{x}_{t}\right)\right)-g_{\psi}\left(\widehat{\theta}^{k} \cdot \tilde{x}_{t}\right)\right)^{2} \leq \frac{L C}{2}\left|\tilde{x}_{t} \cdot\left(\theta_{0}-\widehat{\theta}^{k}\right)\right|^{2} \tag{74}
\end{align*}
$$

where $L \equiv \max _{|v| \leq \psi(M)}\left|\psi^{\prime}(v)\right|$ (since $\psi$ is continuously differentiable, it attains a maximum over a bounded set.) In addition, we used the fact that $g_{\psi}^{\prime}(v) \leq 1$ as explained below Equation (23). The inequalities above then follow from the mean-value theorem.

Given that $\tilde{x}_{t}$ is independent of $\mathcal{H}_{t-1}$, we have

$$
\begin{equation*}
\mathbb{E}\left(\mathrm{reg}_{t} \mid \mathcal{H}_{t-1}\right) \leq \frac{L C}{2}\left\langle\widehat{\theta}^{k}-\theta_{0}, \Sigma_{\phi}\left(\widehat{\theta}^{k}-\theta_{0}\right)\right\rangle, \tag{75}
\end{equation*}
$$

## Dynamic Pricing in High-dimensions

where $\Sigma_{\phi}=\mathbb{E}\left(x_{t} x_{t}^{\boldsymbol{\top}}\right)$. Using Assumption 6,

$$
\begin{equation*}
\mathbb{E}\left(\mathrm{reg}_{t}\right)=\mathbb{E}\left(\mathbb{E}\left(\mathrm{reg}_{t} \mid \mathcal{H}_{t-1}\right)\right) \leq \frac{1}{2} L C C_{\max } \mathbb{E}\left(\left\|\widehat{\theta}^{k}-\theta_{0}\right\|_{2}^{2}\right) . \tag{76}
\end{equation*}
$$

Rest of the proof is similar to proof of Theorem 8.1 (see after Equation (45)).

### 8.5. Proof of Theorem 9

We consider representation (27) of the valuations. Letting $\tilde{x}_{t}=\left(-x_{t} ; p_{t}\right)$ and $\eta=(\mu ; \beta)$ vectors in $\mathbb{R}^{(d+1) \times 1}$, we can write the log-likelihood loss as:

$$
\begin{equation*}
\mathcal{L}(\eta)=-\frac{1}{\tau_{k-1}} \sum_{t=\tau_{k-1}}^{\tau_{k}-1}\left\{\mathbb{I}\left(y_{t}=1\right) \log \left(1-F\left(\tilde{x}_{t} \cdot \eta\right)\right)+\mathbb{I}\left(y_{t}=-1\right) \log \left(F\left(\tilde{x}_{t} \cdot \eta\right)\right)\right\} \tag{77}
\end{equation*}
$$

We let $p^{(k)}=\left(p_{\tau_{k-1}}, \ldots, p_{\tau_{k}-1}\right)^{\top}$ be the vector of prices posted in episode $k$. We also construct feature matrix $X^{(k)}$ by putting features $x_{t}$ that arrive in episode $k$, as its rows. Matrix $\widetilde{X}^{(k)}$ is constructed likewise by stacking augmented features $x_{t}$. Hence, $\widetilde{X}^{(k)}=$ $\left[-X^{(k)}, p^{(k)}\right]$.

The next lemma is a technical lemma that will be used later in bounding the regret. We refer the reader to Appendix C. 9 for its proof.

Lemma 20 Assume that the product features $x_{t} \in \mathbb{R}^{d}$ are generated independently from an unknown distribution with bounded support in $\mathbb{R}^{d}$. Then, with the choice of regularization parameters $\lambda_{k}=4 u_{F} \sqrt{(\log d) / \tau_{k-1}}$ in the RMLP-2 policy, the following inequalities hold true:

$$
\begin{align*}
\frac{1}{\tau_{k-1}}\left\|X^{(k-1)}\left(\mu_{0}-\widehat{\mu}^{k}\right)\right\|^{2} & \leq \frac{4 W}{\ell_{F}} \lambda_{k}  \tag{78}\\
\left(\beta_{0}-\widehat{\beta}^{k}\right)^{2} & \leq \frac{c_{W}}{\ell_{F}} \lambda_{k} \tag{79}
\end{align*}
$$

where $c_{W}>0$ is a constant that depends on $W$.
We next figure out the clairvoyant policy.
Lemma 21 Let $g$ be the pricing function corresponding to distribution $F=F_{0,1}$, given by $g(v)=v+\varphi^{-1}(-v)$, where $\varphi(v)=v-(1-F(v)) / f(v)$ is the virtual valuation function. Then, under model (27), the clairvoyant optimal prices are given by

$$
\begin{equation*}
p_{t}^{*}=\frac{1}{\beta_{0}} g\left(x_{t} \cdot \mu_{0}\right) . \tag{80}
\end{equation*}
$$

Proof of Lemma 21 is given in Appendix C.7.

We next bound the regret at other periods of the episode. Let $\overline{\mathcal{H}}_{t}=\left\{\tilde{x}_{1}, \ldots, \tilde{x}_{t}, \tilde{x}_{t+1}, z_{1}, \ldots, z_{t}\right\}$ be the history set up to time $t$. Similar to (42), we write

$$
\begin{align*}
& \mathbb{E}\left(\operatorname{reg}_{t} \mid \overline{\mathcal{H}}_{t-1}\right) \leq C W\left(p_{t}^{*}-p_{t}\right)^{2} \\
& =C W\left(\frac{1}{\beta_{0}} g\left(x_{t} \cdot \mu_{0}\right)-\frac{1}{\widehat{\beta}^{k}} g\left(x_{t} \cdot \widehat{\mu}^{k}\right)\right)^{2} \\
& \leq \frac{2 C W}{\beta_{0}^{2}}\left(g\left(x_{t} \cdot \mu_{0}\right)-g\left(x_{t} \cdot \widehat{\mu}^{k}\right)\right)^{2}+2 C W\left(\frac{1}{\beta_{0}}-\frac{1}{\widehat{\beta}^{k}}\right)^{2} g\left(x_{t} \cdot \widehat{\mu}^{k}\right)^{2} \\
& \leq \frac{2 C W}{\beta_{0}^{2}}\left|x_{t} \cdot\left(\mu_{0}-\widehat{\mu}^{k}\right)\right|^{2}+\frac{8 C W^{3}}{\beta_{0}^{2}}\left(\widehat{\beta}^{k}-\beta_{0}\right)^{2} . \tag{81}
\end{align*}
$$

In the last step, the first term is bounded by using 1-Lipschitz property of $g$ and the second term is bounded by using the following chain of inequalities:

$$
\frac{1}{\widehat{\beta}^{k}} g\left(x_{t} \cdot \widehat{\mu}^{k}\right) \leq \frac{2}{\widehat{\beta}^{k}}\left|x_{t} \cdot \widehat{\mu}^{k}\right| \leq \frac{2}{\widehat{\beta}^{k}}\left\|x_{t}\right\|_{\infty}\left\|\widehat{\mu}^{k}\right\|_{1} \leq 2 W .
$$

Note that in the first inequality we used the fact that $\varphi(v)$ is increasing for log-concave distributions and hence $g(v)=v+\varphi^{-1}(-v) \leq v+|v| \leq 2|v|$. The last inequality holds because $\left\|\widehat{\mu}^{k} / \widehat{\beta}^{k}\right\| \leq W$ by the optimization constraint in (28).

Recalling our notation $\mathcal{H}_{t}=\overline{\mathcal{H}}_{t} \backslash\left\{\tilde{x}_{t+1}\right\}$ and applying the law of iterated expectations, we have

$$
\begin{align*}
\mathbb{E}\left(\mathrm{reg}_{t} \mid \mathcal{H}_{t-1}\right) & \leq \frac{2 C W}{\beta_{0}^{2}}\left\langle\widehat{\mu}^{k}-\mu_{0}, \Sigma\left(\widehat{\mu}^{k}-\mu_{0}\right)\right\rangle+\frac{8 C W^{3}}{\beta_{0}^{2}}\left(\widehat{\beta}^{k}-\beta_{0}\right)^{2} \\
& \leq C_{1} W^{3}\left(\left\langle\widehat{\mu}^{k}-\mu_{0}, \Sigma\left(\widehat{\mu}^{k}-\mu_{0}\right)\right\rangle+\left(\widehat{\beta}^{k}-\beta_{0}\right)^{2}\right) \tag{82}
\end{align*}
$$

with $C_{1}=8 C W / \beta_{0}^{2}$. Following a similar derivation to (58) we get

$$
\begin{equation*}
\left\langle\widehat{\mu}^{k}-\mu_{0}, \Sigma\left(\hat{\mu}^{k}-\mu_{0}\right)\right\rangle \leq 16\left(\frac{W u_{F}}{\ell_{F}}+W^{2}\right) \sqrt{\frac{\log d}{\tau_{k-1}}} . \tag{83}
\end{equation*}
$$

Using Equations (83) and (79) in Equation (82), we get

$$
\begin{equation*}
\mathbb{E}\left(\mathrm{reg}_{t}\right) \leq \underbrace{C_{1} W^{3}\left[16\left(\frac{W u_{F}}{\ell_{F}}+W^{2}\right)+4 \frac{u_{F} c_{W}}{\ell_{F}}\right]}_{C_{2}} \sqrt{\frac{\log d}{\tau_{k-1}}} \tag{84}
\end{equation*}
$$

Following a similar argument as in Section 8.1 (see Equation (45) and onwards) we have that the following holds for a suitable constant $C>0$ :

$$
\begin{aligned}
\operatorname{Regret}(T) & \leq C_{2} \sum_{k=2}^{K} \sqrt{\frac{\log d}{\tau_{k-1}}} \cdot \tau_{k} \\
& \leq \sqrt{2} C_{2} \sum_{k=2}^{K} \sqrt{(\log d) \tau_{k-1}}=O(\sqrt{(\log d) T})
\end{aligned}
$$

where we used $\tau_{k}=2 \tau_{k-1}$.
For the lower bound $\Omega(\sqrt{T})$, note that under model (27) we can define the (scaled) customer's utility as

$$
\begin{equation*}
\tilde{u}\left(x_{t}\right)=\mu_{0} \cdot x_{t}-\beta_{0} p+\tilde{z}_{t} . \tag{85}
\end{equation*}
$$

Then, a purchase occurs if $\tilde{u}\left(x_{t}\right)>0$. Following our discussion in Section 4.1 (see after Equation (16)), since $\beta_{0}$ is unknown, the uninformative prices do exist and therefore $\Omega(\sqrt{T})$ applies to this case.

## 9. Conclusion

In this work, we leverage tools from statistical learning to design a dynamic pricing policy for a setting wherein the products are described via high-dimensional features. Our policy is computationally efficient and by exploiting the structure of demand parameters, it obtains a regret that scales gracefully with the features dimension and the time horizon. Namely, the regret of our algorithm scales linearly with the sparsity of the optimal solution and logarithmically with the dimension. We also show an $O(\log T)$ dependence of the regret on the length of the horizon. On the flip side, we provide a lower-bound of $O(\log T)$ on the regret of any algorithm that does not know the true parameters of the model in advance.

A natural next step is providing a tight bound on the regret, closing the gap between the derived upper and lower bounds. Another step would be assuming that $\theta^{*}$ is not exactly sparse, but it can be well approximated by a sparse vector, i.e, $\left\|\theta_{0}-\theta_{s_{0}}\right\|_{1} \leq \delta$ for some $s_{0}$-sparse vector $\theta_{s_{0}}$. An interesting question is to figure out how the regret scales with $\delta$.

The choice model that proposed in this work assumes one product arrived at each period, and describes the customer's purchase behavior based on the product features and the posted price. A more general choice model would be the one that assumes multiple products at each period. More specifically, each customer has a "consideration" set which includes products left after the customer has narrowed down her choices based on her own personal screening criteria, and then chooses the product from this set which brings maximum utility. We model the no purchase option as an extra product. This generalization is the focus of a future work.

We also believe the ideas and techniques developed in this work can be be applied to other settings such as personalized pricing where information about the buyers can be used for price differentiation or optimizing reserve prices in online ad auctions. Another application would be assortment optimization and learning consumer choice models both in terms of the role of the structure (Farias et al., 2013; Kallus and Udell, 2016) as well as inventory-constrained environments (Golrezaei et al., 2014; Ferreira et al., 2018; Cheung et al., 2018).

## Acknowledgments

Authors are thankful to Arnoud den Boer and Paat Rusmevichientong for their suggestions that improved this work. A. J. would also like to acknowledge the financial support of the Office of the Provost at the University of Southern California through the Zumberge Fund

Individual Grant Program, and an Outlier Research in Business (iORB) grant from the USC Marshall School of Business (2018). Authors are supported in part by a Google Faculty Research Award (2016).

## Appendix A. Proof of Proposition 10

We start by reviewing the notion of restricted eigenvalue (RE) which is commonplace in high-dimensional statistical estimation.

Definition 22 For a given matrix $A \in \mathbb{R}^{d \times d}$ and some integer $s$ such that $1 \leq s_{0} \leq d$ and a positive number $c$, we say that Restricted Eigenvalue ( $R E$ ) condition is met if

$$
\kappa^{2}\left(A, s_{0}, c\right) \equiv \min _{\substack{J \subseteq[p] \\|J| \leq s_{0}}} \min _{\substack{v \neq 0 \\\left\|v_{J} c\right\|_{1} \leq c\left\|v_{J}\right\|_{1}}} \frac{v^{\top} A v}{\left\|v_{J}\right\|_{2}^{2}}>0
$$

It is shown in (Bühlmann and van de Geer, 2011) and (Rudelson and Zhou, 2013) that when two matrices $A_{0}, A_{1}$ are close to each other (in the maximum element-wise norm) compared to sparsity $s_{0}$, the RE condition for $A_{0}$ implies the RE condition for $A_{1}$. This is particularly useful when $A_{0}$ is a population second moment matrix and $A_{1}$ is a corresponding empirical second moment matrix. To apply this result to our case, let $X \in \mathbb{R}^{n \times d}$ be the feature matrix with rows $x_{t}$, corresponding to $n$ products. Let $\Sigma=\mathbb{E}\left(x_{t} x_{t}^{\boldsymbol{\top}}\right)$ be the empirical second moment. By Assumption 2 , we have $\Sigma \succeq C_{\min } \mathrm{I}$ and therefore, $\Sigma$ satisfies RE condition with $\kappa^{2}\left(\widetilde{\Sigma}, s_{0}, 3\right) \geq C_{\text {min }}$. By using the following result, we conclude that $\widehat{\Sigma}=\left(X^{\top} X\right) / n$ also satisfies RE condition with $\kappa^{2}\left(\widehat{\Sigma}, s_{0}, 3\right) \geq C_{\text {min }} / 2$.

Proposition 23 Let $\widehat{\Sigma}=\left(X^{\top} X\right) / n$ and let $S=\operatorname{supp}\left(\theta_{0}\right)$ be the support of $\theta_{0}$. Under Assumption 2, $\widehat{\Sigma}$ satisfies the restricted eigenvalue condition with constant $\kappa\left(\widehat{\Sigma}, s_{0}, 3\right) \geq$ $\sqrt{C_{\min } / 2}$, with probability $1-e^{-2 n /\left(c_{0} s_{0}\right)}$ and $c_{0}=768 / C_{\min }^{2}$, provided that $n \geq c_{0} s_{0} \log d$,

Proposition 23 follows from the results established in (Bühlmann and van de Geer, 2011) and (Rudelson and Zhou, 2013). We outline the main steps of its proof in Appendix A. 1 for the reader's convenience.

By the second-order Taylor's theorem, expanding around $\theta_{0}$ we have

$$
\begin{equation*}
\mathcal{L}\left(\theta_{0}\right)-\mathcal{L}(\widehat{\theta})=-\left\langle\nabla \mathcal{L}\left(\theta_{0}\right), \widehat{\theta}-\theta_{0}\right\rangle-\frac{1}{2}\left\langle\widehat{\theta}-\theta_{0}, \nabla^{2} \mathcal{L}(\tilde{\theta})\left(\widehat{\theta}-\theta_{0}\right)\right\rangle, \tag{86}
\end{equation*}
$$

for some $\tilde{\theta}$ on the line segment between $\theta_{0}$ and $\widehat{\theta}$. Invoking (31), we have

$$
\begin{equation*}
\nabla \mathcal{L}(\theta)=\frac{1}{n} \sum_{t=1}^{n} \xi_{t}(\theta) x_{t}, \quad \nabla^{2} \mathcal{L}(\theta)=\frac{1}{n} \sum_{t=1}^{n} \eta_{t}(\theta) x_{t} x_{t}^{\top} \tag{87}
\end{equation*}
$$

where $\nabla$ and $\nabla^{2}$ represents the gradient and the hessian w.r.t $\theta$. Further,

$$
\begin{gathered}
\xi_{t}(\theta)=-\frac{f\left(u_{t}(\theta)\right)}{F\left(u_{t}(\theta)\right)} \mathbb{I}\left(y_{t}=-1\right)+\frac{f\left(u_{t}(\theta)\right)}{1-F\left(u_{t}(\theta)\right)} \mathbb{I}\left(y_{t}=+1\right) \\
=-\log ^{\prime} F\left(u_{t}(\theta)\right) \mathbb{I}\left(y_{t}=-1\right)-\log ^{\prime}\left(1-F\left(u_{t}(\theta)\right)\right) \mathbb{I}\left(y_{t}=+1\right) . \\
\eta_{t}(\theta)=\left(\frac{f\left(u_{t}(\theta)\right)^{2}}{F\left(u_{t}(\theta)\right)^{2}}-\frac{f^{\prime}\left(u_{t}(\theta)\right)}{F\left(u_{t}(\theta)\right)}\right) \mathbb{I}\left(y_{t}=-1\right)+\left(\frac{f\left(u_{t}(\theta)\right)^{2}}{\left(1-F\left(u_{t}(\theta)\right)\right)^{2}}+\frac{f^{\prime}\left(u_{t}(\theta)\right)}{1-F\left(u_{t}(\theta)\right)}\right) \mathbb{I}\left(y_{t}=+1\right) \\
=-\log ^{\prime \prime} F\left(u_{t}(\theta)\right) \mathbb{I}\left(y_{t}=-1\right)-\log ^{\prime \prime}\left(1-F\left(u_{t}(\theta)\right)\right) \mathbb{I}\left(y_{t}=+1\right),
\end{gathered}
$$

where $u_{t}(\theta)=p_{t}-\left\langle x_{t}, \theta\right\rangle$, and $\log ^{\prime} F(x)$ and $\log ^{\prime \prime} F(x)$ represent first and second derivative w.r.t $x$, respectively.

By Equation (37), we have

$$
\left|u_{t}\left(\theta_{0}\right)\right| \leq\left|p_{t}\right|+\left\|x_{t}\right\|_{\infty}\left\|\theta_{0}\right\|_{1} \leq 3 W .
$$

Further, recall that the sequences $\left\{p_{t}\right\}_{t=1}^{n}$ and $\left\{x_{t}\right\}_{t=1}^{n}$ are independent of $\left\{z_{t}\right\}_{t=1}^{n}$. Therefore, $\left\{u_{t}\left(\theta_{0}\right)\right\}_{t=1}^{T}$ and $\left\{z_{t}\left(\theta_{0}\right)\right\}_{t=1}^{T}$ are independent and by (4), we have $\mathbb{E}\left[\xi_{t}\left(\theta_{0}\right)\right]=\mathbb{E}\left[\mathbb{E}\left[\xi_{t}\left(\theta_{0}\right) \mid u_{t}\left(\theta_{0}\right)\right]\right]=$ 0 . Further, by definition of $u_{F}$, cf. Equation (13), we have $\left|\xi_{t}\left(\theta_{0}\right)\right| \leq u_{F}$.

We next introduce the set

$$
\begin{equation*}
\mathcal{F} \equiv\left\{\left\|\nabla \mathcal{L}\left(\theta_{0}\right)\right\|_{\infty} \leq 2 u_{F} \sqrt{\frac{\log d}{n}}\right\} \tag{88}
\end{equation*}
$$

By applying Azuma-Hoeffding inequality followed by union bounding over $d$ coordinates of feature vectors, we obtain $\mathbb{P}(\mathcal{F}) \geq 1-1 / d$.

On the other note, $\left\|\theta_{0}\right\|_{1},\|\widehat{\theta}\|_{1} \leq W$ and hence $\|\tilde{\theta}\|_{1} \leq W$. This implies that $\left|u_{t}(\tilde{\theta})\right| \leq$ $3 W$. Therefore, by definition of $\ell_{F}$, cf. Equation (33), we have $\eta_{t}(\tilde{\theta}) \geq \ell_{F}$. Recalling Equation (87), we get $\nabla^{2} \mathcal{L}(\tilde{\theta}) \succeq \ell_{F}\left(\widetilde{X}^{\top} \widetilde{X} / n\right)$.

By optimality of $\widehat{\theta}$, we write

$$
\begin{equation*}
\mathcal{L}(\widehat{\theta})+\lambda\|\widehat{\theta}\|_{1} \leq \mathcal{L}\left(\theta_{0}\right)+\lambda\left\|\theta_{0}\right\|_{1}, \tag{89}
\end{equation*}
$$

and by rearranging the terms and using (86), we arrive at

$$
\begin{equation*}
\frac{\ell_{F}}{n}\left\|X\left(\theta_{0}-\widehat{\theta}\right)\right\|^{2}+\lambda\|\widehat{\theta}\|_{1} \leq\left\|\nabla \mathcal{L}\left(\theta_{0}\right)\right\|_{\infty}\left\|\widehat{\theta}-\theta_{0}\right\|_{1}+\lambda\left\|\theta_{0}\right\|_{1} \tag{90}
\end{equation*}
$$

Form now on, the analysis is exactly similar to the oracle inequality for Lasso estimator. We bring the analysis here for the reader's convenience.

Choosing $\lambda \geq 4 u_{F} \sqrt{(\log d) / n}$, we have on $\mathcal{F}$

$$
\begin{equation*}
\frac{2 \ell_{F}}{n}\left\|X\left(\theta_{0}-\widehat{\theta}\right)\right\|^{2}+2 \lambda\|\widehat{\theta}\|_{1} \leq \lambda\left\|\widehat{\theta}-\theta_{0}\right\|_{1}+2 \lambda\left\|\theta_{0}\right\|_{1} . \tag{91}
\end{equation*}
$$

Let $S=\operatorname{supp}\left(\theta_{0}\right)$. On the left-hand side using triangle inequality, we have

$$
\|\widehat{\theta}\|_{1}=\left\|\widehat{\theta}_{S}\right\|_{1}+\left\|\widehat{\theta}_{S^{c}}\right\|_{1} \geq\left\|\widehat{\theta}_{S}\right\|_{1}-\left\|\widehat{\theta}_{S}-\theta_{0, S}\right\|_{1}+\left\|\widehat{\theta}_{S^{c}}\right\|_{1}
$$

On the right-hand side, we have

$$
\left\|\widehat{\theta}-\theta_{0}\right\|_{1}=\left\|\widehat{\theta}_{S}-\theta_{0, S}\right\|_{1}+\left\|\widehat{\theta}_{S^{c}}\right\|_{1} .
$$

Using these two inequalities in (91), we get

$$
\begin{equation*}
\frac{2 \ell_{F}}{n}\left\|X\left(\theta_{0}-\widehat{\theta}\right)\right\|^{2}+\lambda\left\|\widehat{\theta}_{S^{c}}\right\|_{1} \leq 3 \lambda\left\|\widehat{\theta}_{S}-\theta_{0, S}\right\|_{1} \tag{92}
\end{equation*}
$$

We next write

$$
\begin{aligned}
\frac{2 \ell_{F}}{n}\left\|X\left(\theta_{0}-\widehat{\theta}\right)\right\|^{2}+\lambda\left\|\widehat{\theta}-\theta_{0}\right\|_{1} & =\frac{2 \ell_{F}}{n}\left\|X\left(\theta_{0}-\widehat{\theta}\right)\right\|^{2}+\lambda\left\|\widehat{\theta}_{S}-\theta_{0, S}\right\|_{1}+\lambda\left\|\widehat{\theta}_{S^{c}}\right\|_{1} \\
& \stackrel{(a)}{\leq} 4 \lambda\left\|\widehat{\theta}_{S}-\theta_{0, S}\right\|_{1} \stackrel{(b)}{\leq} 4 \lambda \sqrt{s_{0}}\left\|\widehat{\theta}_{S}-\theta_{0, S}\right\|_{2} \\
& \stackrel{(c)}{\leq} \frac{4 \lambda \sqrt{2 s_{0}}}{\sqrt{n C_{\min }}}\left\|X\left(\widehat{\theta}-\theta_{0}\right)\right\|_{2} \\
& \stackrel{(d)}{\leq} \frac{\ell_{F}}{n}\left\|X\left(\widehat{\theta}-\theta_{0}\right)\right\|_{2}^{2}+\frac{8 \lambda^{2} s_{0}}{\ell_{F} C_{\min }}
\end{aligned}
$$

where (a) follows from Equation (92); (b) holds by Cauchy-Shwarz inequality; (c) follows form the RE condition, which holds for $\widehat{\Sigma}=\left(X^{\top} X\right) / n$ as stated by Proposition 23, with $\kappa\left(\widehat{\Sigma}, s_{0}, 3\right) \geq \sqrt{C_{\min } / 2}$, and recalling the inequality $\left\|\widehat{\theta}_{S^{c}}-\theta_{0, S^{c}}\right\|_{1}=\left\|\widehat{\theta}_{S^{c}}\right\|_{1} \leq 3\left\|\widehat{\theta}_{S}-\theta_{0, S}\right\|$ as per Equation (92); Finally ( $d$ ) follows from the inequality $2 \sqrt{a b} \leq a^{2}+b^{2}$. Rearranging the terms, we obtain

$$
\begin{equation*}
\frac{\ell_{F}}{n}\left\|X\left(\theta_{0}-\widehat{\theta}\right)\right\|^{2}+\lambda\left\|\widehat{\theta}-\theta_{0}\right\|_{1} \leq \frac{8 \lambda^{2} s_{0}}{\ell_{F} C_{\min }} \tag{93}
\end{equation*}
$$

Applying the RE condition again to the l.h.s of (93), we get

$$
\begin{equation*}
C_{\min } \frac{\ell_{F}}{2}\left\|\theta_{0}-\widehat{\theta}\right\|_{2}^{2} \leq \frac{\ell_{F}}{n}\left\|X\left(\theta_{0}-\widehat{\theta}\right)\right\|^{2} \leq \frac{8 \lambda^{2} s_{0}}{\ell_{F} C_{\min }} \tag{94}
\end{equation*}
$$

and therefore,

$$
\begin{equation*}
\left\|\theta_{0}-\widehat{\theta}\right\|_{2}^{2} \leq \frac{16 s_{0} \lambda^{2}}{\ell_{F}^{2} C_{\min }^{2}} \tag{95}
\end{equation*}
$$

The result follows.

## A.1. Proof of Proposition 23

The proof follows by combining two lemmas from (Bühlmann and van de Geer, 2011).
We show the desired result holds for a more general case, namely for $X$ with subgaussian entries. Before stating the proof, we recall a few definitions and notations.

Definition $24 A$ random variable $\nu$ is subgaussian if there exist constants $L, \sigma_{0}$ such that

$$
\mathbb{E}\left(e^{\frac{\nu^{2}}{L^{2}}}\right) \leq \frac{\sigma_{0}^{2}}{L^{2}}+1
$$

Note that bounded random variables are subgaussian. Specifically, if $|\nu| \leq \nu_{\max }$, then $\nu$ is subgaussian with $L=\nu_{\text {max }}$ and $\sigma_{0}=\nu_{\max } \sqrt{e-1}$.

For a matrix $A$, we let $\|A\|_{\infty}$ denote its (element wise) maximum norm, i.e., $\|A\|_{\infty}=$ $\max _{i, j}\left|A_{i j}\right|$. The next lemma shows that if two matrices are close enough in maximum norm and if the compatibility condition holds for one of them then it would also hold for the other one.

Lemma 25 Suppose that the restricted eigenvalue ( $R E$ ) condition holds for $\Sigma_{0}$ with constant $\kappa\left(\Sigma_{0}, s_{0}, 3\right)>0$. If

$$
\left\|\Sigma_{0}-\Sigma_{1}\right\|_{\infty} \leq \frac{\kappa^{2}\left(\Sigma_{0}, s_{0}, 3\right)}{32 s_{0}}
$$

then the RE condition holds for $\Sigma_{1}$ with constant $\kappa\left(\Sigma_{1}, s_{0}, 3\right) \geq \kappa\left(\Sigma_{0}, s_{0}, 3\right) / \sqrt{2}$.
Proof Proof of Lemma 25 We refer to Problem 6.10 of (Bühlmann and van de Geer, 2011).

Lemma 26 Consider $X \in \mathbb{R}^{n \times p}$ with i.i.d. rows generated from a distribution with covariance $\Sigma \in \mathbb{R}^{d \times d}$. Let $\widehat{\Sigma}=X^{\top} X / n$ be the corresponding empirical second moment. Further, suppose that the entries of $X$ are uniformly subgaussian with parameters $L, \sigma_{0}$. If $n \geq c_{0} L s_{0} \log d$ with $c_{0}=768 L / \kappa^{2}\left(\Sigma, s_{0}, 3\right)$, then

$$
\begin{equation*}
\mathbb{P}\left[\|\widehat{\Sigma}-\Sigma\|_{\infty} \geq \frac{\kappa^{2}\left(\Sigma, s_{0}, 3\right)}{384 s_{0}}\left(2+\frac{7 \sigma_{0}}{L}\right)\right] \leq e^{-\frac{2 n}{c s_{0}}} \tag{96}
\end{equation*}
$$

Proof Proof of Lemma 26 The result follows readily from Problem 14.3 on page 535 of (Bühlmann and van de Geer, 2011).

Next we note that $\Sigma$ satisfies the restricted eigenvalue condition with constant $\kappa^{2}\left(\Sigma, s_{0}, 3\right) \geq$ $C_{\text {min }}$ because of Assumption 2. Further, since $\left\|x_{t}\right\|_{\infty} \leq 1$, we can apply the result of Lemma 26 with $L=1, \sigma_{0}=\sqrt{e-1}$. Proposition 23 then follows from Lemma 25 .

## Appendix B. Proof of Proposition 12

Define the event $\mathcal{B}_{n}$ as follows:

$$
\begin{equation*}
\mathcal{B}_{n} \equiv\left\{X \in \mathbb{R}^{n \times d}: \sigma_{\min }\left(X^{\top} X / n\right)>C_{\min } / 2\right\} \tag{97}
\end{equation*}
$$

Using concentration bounds on the spectrum of random matrices with subgaussian rows (see (Vershynin, 2010, Equation (5.26))), there exist constants $c, c_{1}>0$ such that for $n>$ $c_{1} d$, we have $\mathbb{P}\left(\mathcal{B}_{n}\right) \geq 1-e^{-c n^{2}}$.

For $\gamma>0$, we define the event $\mathcal{F}_{\gamma}=\left\{\left\|\nabla \mathcal{L}\left(\theta_{0}\right)\right\|_{\infty} \leq \gamma\right\}$. Using characterization (87), and by applying Azuma-Hoeffding inequality (similar to our argument after Equation (88)), we obtain

$$
\begin{equation*}
\mathbb{P}\left(\mathcal{F}_{\gamma}\right) \geq 1-d \exp \left(-\frac{n \gamma^{2}}{2 u_{F}^{2}}\right) \tag{98}
\end{equation*}
$$

We also let $\mathcal{E}_{1, n} \equiv \mathcal{B}_{n} \cap \mathcal{F}_{\lambda / 2}, \mathcal{E}_{2, n} \equiv \mathcal{B}_{n} \cap \mathcal{F}_{\lambda / 2}^{c}$. To lighten the notation, we use the shorthand $D \equiv\left\|\widehat{\theta}-\theta_{0}\right\|_{2}^{2}$.

We then have

$$
\begin{equation*}
\mathbb{E}(D)=\mathbb{E}\left(D \cdot \mathbb{I}\left(\mathcal{B}_{n}^{c}\right)\right)+\mathbb{E}\left(D \cdot \mathbb{I}\left(\mathcal{E}_{1, n}\right)\right)+\mathbb{E}\left(D \cdot \mathbb{I}\left(\mathcal{E}_{2, n}\right)\right) \tag{99}
\end{equation*}
$$

We treat each of the terms on the right-hand side separately.

- Term 1: We have

$$
\begin{equation*}
\mathbb{E}\left(D \cdot \mathbb{I}\left(\mathcal{B}_{n}^{c}\right)\right) \leq 4 W^{2} \mathbb{P}\left(\mathcal{B}_{n}^{c}\right) \leq 4 W^{2} e^{-c n^{2}} \tag{100}
\end{equation*}
$$

- Term 2: Similar to proof of Proposition 10 , on $\mathcal{E}_{1, n}$, we have $D \leq 16 s_{0} \lambda^{2} /\left(\ell_{F}^{2} C_{\text {min }}^{2}\right)$. Hence,

$$
\begin{equation*}
\mathbb{E}\left(D \cdot \mathbb{I}\left(\mathcal{E}_{1, n}\right)\right) \leq \frac{16 s_{0} \lambda^{2}}{\ell_{F}^{2} C_{\min }^{2}} \tag{101}
\end{equation*}
$$

- Term 3: To bound term 3, we first prove the following lemma.

Lemma 27 On event $\mathcal{E}_{n}(\gamma) \equiv \mathcal{B}_{n} \cap \mathcal{F}_{\gamma}$, with $\gamma>\lambda / 2$, we have

$$
D \leq\left(\frac{36}{C_{\min }^{2} \ell_{F}^{2}}\right) \gamma^{2} d
$$

Lemma 27 is proved in Section B.1.
We next bound term 3 as follows. Let $L=9 \lambda^{2} d /\left(C_{\min }^{2} \ell_{F}^{2}\right)$.

$$
\begin{align*}
\mathbb{E}\left(D \cdot \mathbb{I}\left(\mathcal{E}_{2, n}\right)\right) & =\int_{0}^{\infty} \mathbb{P}\left(D \cdot \mathbb{I}\left(\mathcal{E}_{2, n}\right)>\alpha\right) \mathrm{d} \alpha \\
& =L \int_{0}^{\infty} \mathbb{P}\left(D \cdot \mathbb{I}\left(\mathcal{E}_{2, n}\right)>L c\right) \mathrm{d} c \\
& =L \int_{0}^{1} \mathbb{P}\left(D \cdot \mathbb{I}\left(\mathcal{E}_{2, n}\right)>L c\right) \mathrm{d} c+L \int_{1}^{\infty} \mathbb{P}\left(D \cdot \mathbb{I}\left(\mathcal{E}_{2, n}\right)>L c\right) \mathrm{d} c . \tag{102}
\end{align*}
$$

For the first term on the right-hand side we write

$$
\begin{equation*}
L \int_{0}^{1} \mathbb{P}\left(D \cdot \mathbb{I}\left(\mathcal{E}_{2, n}\right)>L c\right) \mathrm{d} c \leq L \int_{0}^{1} \mathbb{P}\left(\mathcal{E}_{2, n}\right) \leq L \mathbb{P}\left(\mathcal{F}_{\lambda / 2}^{c}\right) \leq \frac{L}{d} \tag{103}
\end{equation*}
$$

where the last step holds from Equation (98) with $\gamma=\lambda / 2$.
We next upper bound the second term. For arbitrary fixed $c>1$, let $\gamma=\sqrt{L c /(d \kappa)}$, with $\kappa=36 /\left(\ell_{F}^{2} C_{\text {min }}^{2}\right)$. It is easy to verify that $\gamma=\sqrt{c} \lambda / 2>\lambda / 2$. Further, by virtue of Lemma 27 , on $\mathcal{E}(\gamma)$ we have $D \leq \kappa \gamma^{2} d=L c$. Hence,

$$
\begin{equation*}
\mathbb{P}\left(D \cdot \mathbb{I}\left(\mathcal{E}_{2, n}\right)>L c\right) \leq \mathbb{P}\left(\mathcal{E}(\gamma)^{c} \cap \mathcal{E}_{2, n}\right) \leq \mathbb{P}\left(\mathcal{F}_{\gamma}^{c} \cap \mathcal{B}_{n}\right) \leq \mathbb{P}\left(\mathcal{F}_{\gamma}^{c}\right) \tag{104}
\end{equation*}
$$

Further, applying Equation (98) and plugging for $\gamma$, we obtain

$$
\begin{equation*}
\mathbb{P}\left(\mathcal{F}_{\gamma}^{c}\right) \leq d \exp \left(-\frac{n L c}{2 u_{F}^{2} \kappa d}\right)=d \exp \left(-\frac{n \lambda^{2}}{8 u_{F}^{2}}\right) \leq d^{1-2 c} . \tag{105}
\end{equation*}
$$

Here, the second step follows from definition of $L$ and the last step holds because $\lambda \geq$ $4 u_{F} \sqrt{(\log d) / n}$.

Combining Equations (104) and (105), we have

$$
\begin{equation*}
L \int_{1}^{\infty} \mathbb{P}\left(D \cdot \mathbb{I}\left(\mathcal{E}_{2, n}\right)>L c\right) \mathrm{d} c \leq L \int_{1}^{\infty} d^{1-2 c} \mathrm{~d} c \leq \frac{L}{2 d \log d} \tag{106}
\end{equation*}
$$

Using bounds (103) and (106) in Equation (102), we obtain

$$
\begin{equation*}
\mathbb{E}\left(D \cdot \mathbb{I}\left(\mathcal{E}_{2, n}\right)\right) \leq \frac{L}{d}\left(1+\frac{1}{2 \log d}\right) \leq \frac{3 L}{2 d}<\frac{16 \lambda^{2}}{\ell_{F}^{2} C_{\min }^{2}} \tag{107}
\end{equation*}
$$

The result follows by putting the upper bounds on the three terms together.

## B.1. Proof of Lemma 27

We start by rewriting Equation (90), which follows from optimality of $\widehat{\theta}$ and log-concave property of the loss function:

$$
\begin{equation*}
\frac{\ell_{F}}{n}\left\|X\left(\theta_{0}-\widehat{\theta}\right)\right\|^{2}+\lambda\|\widehat{\theta}\|_{1} \leq\left\|\nabla \mathcal{L}\left(\theta_{0}\right)\right\|_{\infty}\left\|\widehat{\theta}-\theta_{0}\right\|_{1}+\lambda\left\|\theta_{0}\right\|_{1} . \tag{108}
\end{equation*}
$$

On event $\mathcal{E}_{n}(\gamma)$, Equation (108) implies that

$$
\begin{equation*}
\frac{1}{2} C_{\min } \ell_{F}\left\|\theta_{0}-\widehat{\theta}\right\|_{2}^{2}+\lambda\|\widehat{\theta}\|_{1} \leq \gamma\left\|\widehat{\theta}-\theta_{0}\right\|_{1}+\lambda\left\|\theta_{0}\right\|_{1} \tag{109}
\end{equation*}
$$

Using the assumption $\gamma>\lambda / 2$ and our shorthand $D \equiv\left\|\widehat{\theta}-\theta_{0}\right\|_{2}^{2}$, we get

$$
\begin{aligned}
\frac{1}{2} C_{\min } \ell_{F}\left\|\widehat{\theta}-\theta_{0}\right\|_{2}^{2} & \leq \gamma\left\|\widehat{\theta}-\theta_{0}\right\|_{1}+\lambda\left\|\theta_{0}\right\|_{1}-\lambda\|\widehat{\theta}\|_{1} \\
& \leq(\gamma+\lambda)\left\|\widehat{\theta}-\theta_{0}\right\|_{1} \\
& \leq 3 \gamma\left\|\widehat{\theta}-\theta_{0}\right\|_{1} \leq 3 \gamma \sqrt{d}\left\|\widehat{\theta}-\theta_{0}\right\|_{2}
\end{aligned}
$$

Writing the above bound in terms of our shorthand $D \equiv\left\|\widehat{\theta}-\theta_{0}\right\|_{2}^{2}$, we obtain the desired result.

## Appendix C. Proof of Technical Lemmas

## C.1. Proof of Lemma 13

We write the virtual valuation function as $\varphi(v)=v-1 / \lambda(v)$ where $\lambda(v)=\frac{f(v)}{1-F(v)}=$ $-\log ^{\prime}(1-F(v))$ is the hazard rate function. Since $1-F$ is log-concave, the hazard function $\lambda(v)$ is increasing which implies that $\varphi$ is strictly increasing. Indeed, by this argument $\varphi^{\prime}(v)>1$.

## C.2. Proof of Lemma 14

Recalling the definition $g(v)=v+\varphi^{-1}(-v)$, we have $g^{\prime}(v)=1-1 / \varphi^{\prime}\left(\varphi^{-1}(-v)\right)$. Since $\varphi$ is strictly increasing by Lemma 13, we have $g^{\prime}(v)<1$. The claim $g^{\prime}(v)>0$ follows if we show $\varphi^{\prime}\left(\varphi^{-1}(-v)\right)>1$. For this we refer to the proof of Lemma 13, where we showed that $\varphi^{\prime}(v)>1$ for all $v$.

## C.3. Proof of Lemma 16

Let $\phi(v)$ and $\Phi(v)$ respectively denote the density and the distribution function of standard normal variable. Function $h_{t}$ and its derivatives read as

$$
\begin{align*}
\operatorname{rev}_{t}(p) & =p\left(1-\Phi\left(\frac{p-x_{t} \cdot \theta_{0}}{\sigma}\right)\right)  \tag{110}\\
\operatorname{rev}_{t}^{\prime}(p) & =1-\Phi\left(\frac{p-x_{t} \cdot \theta_{0}}{\sigma}\right)-\frac{p}{\sigma} \phi\left(\frac{p-x_{t} \cdot \theta_{0}}{\sigma}\right)  \tag{111}\\
\operatorname{rev}_{t}^{\prime \prime}(p) & =\frac{1}{\sigma}\left[\frac{p}{\sigma}\left(\frac{p-x_{t} \cdot \theta_{0}}{\sigma}\right)-2\right] \phi\left(\frac{p-x_{t} \cdot \theta_{0}}{\sigma}\right) \tag{112}
\end{align*}
$$

Define $\xi \equiv p_{t}^{*}-x_{t} \cdot \theta_{0}=g\left(x_{t} \cdot \theta_{0}\right)-x_{t} \cdot \theta_{0}$. Writing $\operatorname{rev}_{t}^{\prime \prime}\left(p_{t}^{*}\right)$ in term of $\xi$, we obtain

$$
\begin{equation*}
r_{t}^{\prime \prime}\left(p_{t}^{*}\right)=\frac{1}{\sigma}\left[\frac{\xi}{\sigma}\left(\frac{1-\Phi(\xi / \sigma)}{\phi(\xi / \sigma)}\right)-2\right] \phi(\xi / \sigma) \tag{113}
\end{equation*}
$$

By tail bound inequality for Gaussian distribution $1-\Phi(\xi / \sigma) \leq(\sigma / \xi) \phi(\xi / \sigma)$ for $\xi \geq 0$. Therefore,

$$
\begin{equation*}
\frac{\xi}{\sigma}\left(\frac{1-\Phi(\xi / \sigma)}{\phi(\xi / \sigma)}\right)-2 \leq-1 \tag{114}
\end{equation*}
$$

and the same bound obviously holds for $\xi<0$.
By definition of function $g,|\xi| \leq 3 W$ with $\varphi$ being the virtual valuation fusion corresponding to the Gaussian distribution. Hence, $\phi(\xi / \sigma) \geq \phi(3 W / \sigma)$. Putting this together with (114), we get $\operatorname{rev}_{t}^{\prime \prime}\left(p_{t}^{*}\right) \leq-c_{1}$ with $c_{1}=(1 / \sigma) \phi(3 W / \sigma)$.

For the second part of the Lemma statement, set $\delta=\min \left\{3 W, \sigma^{2} /(18 W), \sigma^{2} \phi(3 W / \sigma)\right\}$. For $p \in\left[p_{t}^{*}-\delta, p_{t}^{*}+\delta\right]$, we have

$$
\begin{equation*}
\left|\frac{p}{\sigma}\left(\frac{p-x_{t} \cdot \theta_{0}}{\sigma}\right)-\frac{p_{t}^{*}}{\sigma}\left(\frac{p_{t}^{*}-x_{t} \cdot \theta_{0}}{\sigma}\right)\right| \leq \frac{1}{\sigma^{2}}\left|p-p_{t}^{*}\right| \cdot\left|p+p_{t}^{*}-x_{t} \cdot \theta_{0}\right| \leq \frac{1}{\sigma^{2}} \delta(5 W+\delta) \leq \frac{1}{2} \tag{115}
\end{equation*}
$$

Using Equation (114) we get $p\left(p-x_{t} \cdot \theta_{0}\right) / \sigma^{2} \leq-1 / 2$. Further,

$$
\begin{equation*}
\left|\phi\left(\frac{p-x_{t} \cdot \theta_{0}}{\sigma}\right)-\phi\left(\frac{p_{t}^{*}-x_{t} \cdot \theta_{0}}{\sigma}\right)\right| \leq \frac{\left|p-p_{t}^{*}\right|}{2 \sigma^{2}} \leq \frac{\delta}{2 \sigma^{2}} \leq \frac{1}{2} \phi(3 W / \sigma) \tag{116}
\end{equation*}
$$

Therefore,

$$
\begin{equation*}
\phi\left(\frac{p-x_{t} \cdot \theta_{0}}{\sigma}\right) \geq \phi(\xi / \sigma)-\frac{1}{2} \phi(3 W / \sigma) \geq \frac{1}{2} \phi(3 W / \sigma) . \tag{117}
\end{equation*}
$$

Combining (115), (117) we obtain

$$
\begin{equation*}
\operatorname{rev}_{t}^{\prime \prime}(p)=\frac{1}{\sigma}\left[\frac{p}{\sigma}\left(\frac{p-x_{t} \cdot \theta_{0}}{\sigma}\right)-2\right] \phi\left(\frac{p-x_{t} \cdot \theta_{0}}{\sigma}\right) \leq-\frac{1}{4 \sigma} \phi(3 W / \sigma)=-\frac{c_{1}}{4} \tag{118}
\end{equation*}
$$

The result follows.

## C.4. Proof of Lemma 17

Let $Z=x \cdot v$ and $\tilde{Z}=Z /\|v\|_{2}$. Note that $\operatorname{Var}(\tilde{Z})=1$. Write the expectation in terms of the tail probability

$$
\begin{equation*}
\mathbb{E}\left(\min \left(Z^{2}, \delta^{2}\right)\right)=\int_{0}^{\delta^{2}} \mathbb{P}\left(Z^{2} \geq t\right) \mathrm{d} t=\int_{0}^{\delta^{2}} \mathbb{P}\left(|\tilde{Z}| \geq \frac{\sqrt{t}}{\|v\|_{2}}\right) \mathrm{d} t \tag{119}
\end{equation*}
$$

We consider two cases:

- $\delta \leq\|v\|_{2}$ : The right-hand side in (119) can be lower bounded as

$$
\begin{equation*}
\int_{0}^{\delta^{2}} \mathbb{P}\left(|\tilde{Z}| \geq \frac{\sqrt{t}}{\|v\|_{2}}\right) \mathrm{d} t \geq \int_{0}^{\delta^{2}} \mathbb{P}\left(|\tilde{Z}| \geq \frac{\sqrt{t}}{\delta}\right) \mathrm{d} t=2 \delta^{2} \int_{0}^{1} t \mathbb{P}(|\tilde{Z}| \geq t) \mathrm{d} t \tag{120}
\end{equation*}
$$

In the sequel, we provide two separate lower bounds for the right-hand side.
Let $\xi \equiv \mathbb{P}(|\tilde{Z}| \geq 1)$. We have

$$
\begin{equation*}
\int_{0}^{1} t \mathbb{P}(|\tilde{Z}| \geq t) \mathrm{d} t \geq \int_{0}^{1} t \xi \mathrm{~d} t \geq \frac{\xi}{2} \tag{121}
\end{equation*}
$$

We proceed to obtain another bound which utilizes the fact $\operatorname{Var}(\tilde{Z})=1$.

$$
\begin{align*}
\int_{0}^{1} t \mathbb{P}(|\tilde{Z}| \geq t) \mathrm{d} t & =\int_{0}^{\infty} t \mathbb{P}(|\tilde{Z}| \geq t) \mathrm{d} t-\int_{1}^{\infty} t \mathbb{P}(|\tilde{Z}| \geq t) \mathrm{d} t \\
& =\frac{1}{2} \operatorname{Var}(\tilde{Z})-\int_{1}^{\infty} t \mathbb{P}(|\tilde{Z}| \geq t) \mathrm{d} t \tag{122}
\end{align*}
$$

For $t \geq 1$, we have $\mathbb{P}(|\tilde{Z}| \geq t) \leq \xi$. Further, by applying Chernoff bound, we get

$$
\begin{aligned}
\mathbb{P}(|\tilde{Z}| \geq t) & =2 \mathbb{P}(\tilde{Z} \geq t)=2 \mathbb{P}\left(e^{\lambda \tilde{Z}} \geq e^{\lambda t}\right) \leq e^{-\lambda t} \mathbb{E}\left(e^{\lambda \tilde{Z}}\right) \\
& =2 e^{-\lambda t} \prod_{i=1}^{d}\left(\frac{e^{\lambda \frac{v_{i}}{\|v\|_{2}}}+e^{-\lambda \frac{v_{i}}{\|v\|_{2}}}}{2}\right) \leq 2 e^{-\lambda t} \prod_{i=1}^{d} e^{\lambda^{2} \frac{v_{i}^{2}}{2\|v\|_{2}^{2}}}=2 e^{\frac{\lambda^{2}}{2}-\lambda t}
\end{aligned}
$$

Setting $\lambda=t$ leads to $\mathbb{P}(|\tilde{Z}| \geq t) \leq 2 e^{-\frac{t^{2}}{2}}$. Combining these bounds into (121), we obtain

$$
\begin{equation*}
\int_{0}^{1} t \mathbb{P}(|\tilde{Z}| \geq t) \mathrm{d} t \geq \frac{1}{2}-\int_{1}^{\infty} t \min \left(2 e^{-\frac{t^{2}}{2}}, \xi\right) \mathrm{d} t=\frac{1-\xi}{2}-\xi \log \left(\frac{2}{\xi}\right) \tag{123}
\end{equation*}
$$

We summarize bounds (121) and (123) as in

$$
\begin{equation*}
\int_{0}^{1} t \mathbb{P}(|\tilde{Z}| \geq t) \mathrm{d} t \geq \min _{\xi \in[0,1]} \max \left(\xi, \frac{1-\xi}{2}-\xi \log \left(\frac{2}{\xi}\right)\right)>0.05 . \tag{124}
\end{equation*}
$$

Turning back to Equation (120), in this case we have

$$
\begin{equation*}
\int_{0}^{\delta^{2}} \mathbb{P}\left(|\tilde{Z}| \geq \frac{\sqrt{t}}{\|v\|_{2}}\right) \mathrm{d} t \geq 0.1 \delta^{2} \tag{125}
\end{equation*}
$$

- $\delta \geq\|v\|_{2}$ : Similar to the previous case, the right-hand side in (119) can be lower bounded as

$$
\begin{equation*}
\int_{0}^{\delta^{2}} \mathbb{P}\left(|\tilde{Z}| \geq \frac{\sqrt{t}}{\|v\|_{2}}\right) \mathrm{d} t \geq \int_{0}^{\|v\|_{2}^{2}} \mathbb{P}\left(|\tilde{Z}| \geq \frac{\sqrt{t}}{\|v\|_{2}}\right) \mathrm{d} t . \geq 0.1\|v\|_{2}^{2} \tag{126}
\end{equation*}
$$

The above two cases can be summarized as $\mathbb{E}\left(\min \left(Z^{2}, \delta^{2}\right)\right) \geq 0.1 \min \left(\|v\|_{2}^{2}, \delta^{2}\right)$.

## C.5. Proof of Lemma 18

We use a standard argument that relates minimax $\ell_{2}$-risk in terms of the error in multi-way hypothesis testing problem; See e.g. (Yang and Barron, 1999; Yu, 1997). Let $\left\{\tilde{\theta}_{1}, \ldots, \tilde{\theta}_{m}\right\}$ be a $\delta$-packing of set $\Omega$, meaning that their pairwise distances are all at least $\delta$. Parameter $\delta$ is free for now and its value will be determined later in the proof. We further let $P_{j}$ denote the induced probability on market values $\left(v\left(x_{1}\right), \ldots, v\left(x_{T}\right)\right)$, conditional on $\left(x_{1}, \ldots, x_{T}\right)$ and for $\theta_{0}=\tilde{\theta}_{j}$. In other words, in defining distributions $P_{j}$ we treat feature vectors fixed. Let $\nu$ be random variable uniformly distributed on the hypothesis set $\{1,2, \ldots, m\}$ which indicates the index of the true parameter, i.e, $\nu=j$ means $\theta_{0}=\tilde{\theta}_{j}$.

Define $\mathrm{d}\left(\theta_{1}^{T}, \theta\right) \equiv \sum_{t=1}^{T} \min \left(\left\|\theta_{t}-\theta\right\|_{2}^{2}, C\right)$ and let $\mu$ be the value of $j$ for which $\mathrm{d}\left(\theta_{1}^{T}, \tilde{\theta}_{j}\right)$ is a minimum. Suppose that $\delta$ is chosen such that $\delta^{2} \leq C$. If $\mathrm{d}\left(\theta_{1}^{T}, \tilde{\theta}_{j}\right)<\delta^{2} T / 4$ then $\mu=j$, because assuming otherwise, we have $\mu=j^{\prime} \neq j$, and by triangle inequality

$$
\begin{align*}
\min \left(\left\|\tilde{\theta}_{j^{\prime}}-\tilde{\theta}_{j}\right\|_{2}^{2}, C\right) & \leq \min \left(2\left\|\theta_{t}-\tilde{\theta}_{j}\right\|_{2}^{2}+2\left\|\theta_{t}-\tilde{\theta}_{j^{\prime}}\right\|_{2}^{2}, C\right) \\
& \leq \min \left(2\left\|\theta_{t}-\tilde{\theta}_{j}\right\|_{2}^{2}, C\right)+\min \left(2\left\|\theta_{t}-\tilde{\theta}_{j^{\prime}}\right\|_{2}^{2}, C\right) \tag{127}
\end{align*}
$$

for all $t$, where we used the inequality $\min (a+b, c) \leq \min (a, c)+\min (b, c)$ for $a, b, c \geq 0$. Summing over $t=1,2, \ldots, T$, we get

$$
T \min \left(\left\|\tilde{\theta}_{j^{\prime}}-\tilde{\theta}_{j}\right\|_{2}^{2}, C\right) \leq 2 \mathrm{~d}\left(\theta_{1}^{T}, \tilde{\theta}_{j}\right)+2 \mathrm{~d}\left(\theta_{1}^{T}, \tilde{\theta}_{j^{\prime}}\right) \leq 4 \mathrm{~d}\left(\theta_{1}^{T}, \tilde{\theta}_{j}\right)<\delta^{2} T,
$$

where we used the assumption $\mu=j^{\prime}$. But this is a contradiction because $\left\|\tilde{\theta}_{j^{\prime}}-\tilde{\theta}_{j}\right\|_{2} \geq \delta$ (they form a $\delta$-packing of $\Omega$ ) and $\delta^{2} \leq C$.

Using Markov inequality, we can write

$$
\begin{align*}
\max _{j} \mathbb{E}_{P_{j}} \mathrm{~d}\left(\theta_{1}^{T}, \tilde{\theta}_{j}\right) & \geq \frac{\delta^{2} T}{4} \max _{j} \mathbb{P}\left(\left.\mathrm{~d}\left(\theta_{1}^{T}, \tilde{\theta}_{j}\right) \geq \frac{\delta^{2} T}{4} \right\rvert\, \nu=j\right) \\
& \geq \frac{\delta^{2} T}{4 m} \sum_{j=1}^{m} \mathbb{P}(\mu \neq j \mid \nu=j)=\frac{\delta^{2} T}{4} \mathbb{P}(\mu \neq \nu) . \tag{128}
\end{align*}
$$

We use Fano's inequality to lower bound the error probability on the right-hand side. We first construct a $\delta$-packing of $\Omega$ similar to the one proposed in (Raskutti et al., 2011, proof of Theorem 1).

Let $s=s_{0} / 2 \leq d / 2$ and define

$$
\mathcal{A}=\left\{q \in\{-1,0,1\}^{d}:\|q\|_{0}=s\right\} .
$$

As proved in (Raskutti et al., 2011, Lemma 5), there exists a subset $\tilde{\mathcal{A}} \subseteq \mathcal{A}$ of cardinality $|\tilde{\mathcal{A}}| \geq \exp \left(\frac{s}{2} \log \frac{d-s / 2}{s}\right)$ such that the Hamming distance between any two elements in $\tilde{\mathcal{A}}$
is at least $s / 2$. Next, consider the set $\sqrt{\frac{2}{s}} \delta \tilde{\mathcal{A}}$ for some $\delta \leq W / \sqrt{2 s}$ whose exact value to be determined later. Then, for $q$ in this set, $\|q\|_{1}=\sqrt{2 s} \delta \leq W$ and hence $\sqrt{\frac{2}{s}} \delta \tilde{\mathcal{A}} \subseteq \Omega$. Further, for $q, q^{\prime} \in \sqrt{\frac{2}{s}} \delta \tilde{\mathcal{A}}$, we have the following bounds:

$$
\begin{align*}
& \left\|q-q^{\prime}\right\|_{2}^{2} \geq \delta^{2}  \tag{129}\\
& \left\|q-q^{\prime}\right\|_{2}^{2} \leq 8 \delta^{2} \tag{130}
\end{align*}
$$

By (129), the set $\sqrt{\frac{2}{s}} \delta \tilde{\mathcal{A}}$ forms a $\delta$-packing for $\Omega$ with size $|\tilde{\mathcal{A}}|$.
We now turn back to bound (128). Left-hand side can be lower bounded using Fano's inequality. We omit the details here as it is a standard argument and instead we refer to (Raskutti et al., 2011, proof of Theorem 1) for details. Using Fano's inequality and bound (130), we get

$$
\begin{equation*}
\mathbb{P}(\mu \neq \nu)=1-\frac{\frac{8 T}{2 \sigma^{2}} \delta^{2}+\log (2)}{\frac{s}{2} \log \left(\frac{d-s / 2}{s}\right)} . \tag{131}
\end{equation*}
$$

Choosing $\delta^{2} \leq \delta_{1}^{2} \equiv \frac{\sigma^{2} s}{32 T} \log \left(\frac{d-s / 2}{s}\right)$, we obtain $\mathbb{P}(\mu \neq \nu) \geq 1 / 4$. Therefore, setting $\delta^{2}=$ $\min \left(\frac{W^{2}}{2 s}, \delta_{1}^{2}, C\right)$ and combining with bound (128), we conclude that

$$
\begin{equation*}
\min _{\theta_{1}^{T}} \max _{\theta_{0} \in \Omega} \mathbb{E}\left(\mathrm{~d}\left(\theta_{1}^{T}, \theta_{0}\right)\right) \geq \frac{\delta^{2} T}{16}=\frac{1}{16} \min \left\{\frac{W^{2} T}{2 s}, \frac{\sigma^{2} s}{32} \log \left(\frac{d-s / 2}{s}\right), C T\right\} . \tag{132}
\end{equation*}
$$

Now since $s=s_{0} / 2 \leq d / 2$, we have $\log ((d-s / 2) / s) \geq c \log (d / s)$ with some constant $c>0$. Therefore, by using Equation (132) and substituting for $s=s_{0} / 2$, we obtain

$$
\begin{equation*}
\min _{\theta_{1}^{T}} \max _{\theta_{0} \in \Omega} \mathbb{E}\left(\mathrm{~d}\left(\theta_{1}^{T}, \theta_{0}\right)\right) \geq L_{1}, \tag{133}
\end{equation*}
$$

with

$$
\begin{equation*}
L_{1} \equiv \frac{1}{16} \min \left\{\frac{W^{2} T}{s_{0}}, \frac{c \sigma^{2} s_{0}}{64} \log \left(\frac{d}{s_{0}}\right), C T\right\} . \tag{134}
\end{equation*}
$$

We next derive another separate lower bound for minimax risk, by assuming that an oracle gives us the true support of $\theta_{0}$. In this case, the least square estimator, applied to the observed features restricted to the true support $S$, achieves the optimal minimax $\ell_{2}$ rate. This implies that $\left\|\theta_{t}-\theta_{0}\right\|_{2}^{2} \geq c \sigma^{2} s_{0} / t$, for $t \geq s_{0}$ and a constant $c>0$. Therefore,

$$
\begin{equation*}
\min _{\theta_{1}^{T}} \max _{\theta_{0} \in \Omega} \mathbb{E}\left(\mathrm{~d}\left(\theta_{1}^{T}, \theta_{0}\right)\right) \geq \sum_{t=1}^{T} \min \left(c \sigma^{2} \frac{s_{0}}{t}, C\right) \tag{135}
\end{equation*}
$$

from which we obtain

$$
\begin{equation*}
\min _{\theta_{1}^{T}} \max _{\theta_{0} \in \Omega} \mathbb{E}\left(\mathrm{~d}\left(\theta_{1}^{T}, \theta_{0}\right)\right) \geq L_{2} \equiv c^{\prime} s_{0} \log \left(T / s_{0}\right), \tag{136}
\end{equation*}
$$

for some constant $c^{\prime}>0$, depending on $\sigma$ and $C$.
Combining bounds in (134) and (136), we have

$$
\begin{align*}
\min _{\theta_{1}^{T}} \max _{\theta_{0} \in \Omega} \mathbb{E}\left(\mathrm{~d}\left(\theta_{1}^{T}, \theta_{0}\right)\right) & \geq \frac{1}{2}\left(L_{1}+L_{2}\right) \\
& \geq \widetilde{C}\left\{s_{0} \log \left(\frac{T}{s_{0}}\right)+\min \left[\frac{T}{s_{0}}, s_{0} \log \left(\frac{d}{s_{0}}\right)\right]\right\} \tag{137}
\end{align*}
$$

for a constant $\widetilde{C}$ that depends on $C, \sigma, W$. The proof is complete.

## C.6. Proof of Lemma 19

We recall the notion of 0-property established by (Ponomarev, 1987).
Definition 28 A continuous function has the 0-property, if the pre-image of any set of probability zero is a set of probability zero.

As proved in (Ponomarev, 1987, Theorem 1), if a function $\phi: \mathcal{X} \subseteq \mathbb{R}^{d} \mapsto \mathbb{R}^{d}$ is continuously differentiable, then it satisfies 0-property if and only if its derivative $\mathcal{D} \phi$ is full rank for almost all $x \in \mathcal{X}$. Therefore, we need to show that under Assumption 1, if $\phi$ has 0-property, then Assumption 6 holds true.

Supposing otherwise, there exists a nonzero $v \in \mathbb{R}^{d}$ such that $v^{\top} \Sigma_{\phi} v=0$. Therefore, $\mathbb{E}\left((z \cdot \phi(x))^{2}\right)=0$ which implies that $z \cdot \phi(x)=0$, almost surely. Define $S \equiv\left\{z \in \mathbb{R}^{d}\right.$ : $z \cdot \phi(x)=0\}$. Space $S$ is $(d-1)$-dimensional and all the points in $\phi(\mathcal{X})$ belong to $S$ almost surely, i.e., $\mathbb{P}\left(\phi(\mathcal{X}) \cap S^{c}\right)=0$. However, since $\Sigma$ is positive definite (with all of its eigenvalues target than $C_{\min }$, by Assumption 1), $\mathbb{P}_{X}(S)=0$. Combining these observations, $\mathbb{P}(\phi(\mathcal{X})) \leq \mathbb{P}(S)+\mathbb{P}\left(\phi(\mathcal{X}) \cap S^{c}\right)=0$. Since $\phi$ has the 0 -property, this implies that $\mathbb{P}_{X}(\mathcal{X})=0$, which is a contradiction because $\mathcal{X}$ is the support of $\mathbb{P}_{X}$ and thus $\mathbb{P}_{X}(\mathcal{X})=1$. The result follows.

## C.7. Proof of Lemma 21

Under model (27), a purchase occurs at time $t$ with the posted price price $p$ if $\tilde{v}_{t} \geq \beta_{0} p$. This is equivalent to $\tilde{z}_{t} \geq \beta_{0} p-x_{t} \cdot \mu_{0}$. Therefore, the expected revenue from a posted price $p$ is given by

$$
\begin{equation*}
p \times \mathbb{P}\left(\tilde{v}_{t} \geq \beta_{0} p\right)=p\left(1-F\left(\beta_{0} p-\mu_{0} \cdot x_{t}\right)\right) . \tag{138}
\end{equation*}
$$

By setting the first order conditions, the optimal price $p_{t}^{*}$ is given by the solution of the following equation:

$$
\begin{equation*}
\beta_{0} p_{t}^{*}=\frac{1-F\left(\beta_{0} p_{t}^{*}-\mu_{0} \cdot x_{t}\right)}{f\left(\beta_{0} p_{t}^{*}-\mu_{0} \cdot x_{t}\right)} \tag{139}
\end{equation*}
$$

It is straightforward to verify that the solution $p_{t}^{*}$ of the above equation is given by $p_{t}^{*}=\left(1 / \beta_{0}\right) g\left(\mu_{0} \cdot x_{t}\right)$.

## C.8. Proof of Lemma 15

Recall the notation $\left\|E^{(k)}\right\|_{\infty}=\max _{i, j}\left|E_{i j}^{(k)}\right|$. Note that

$$
\begin{align*}
\left\langle v, E^{(k)} v\right\rangle & =\sum_{i, j=1}^{d+1}\left|E_{i j}^{(k)}\left\|v_{i}\right\| v_{j}\right| \leq\left\|E^{(k)}\right\|_{\infty} \sum_{i, j=1}^{d+1}\left|v_{i}\right|\left|v_{j}\right| \\
& \leq\left\|E^{(k)}\right\|_{\infty}\left(\sum_{i=1}^{d}\left|v_{i}\right|\right)^{2}=\left\|E^{(k)}\right\|_{\infty}\|v\|_{1}^{2} . \tag{140}
\end{align*}
$$

Therefore, we only need to bound $\left\|E^{(k)}\right\|_{\infty}$. Fix $1 \leq i, j \leq d+1$. We then have

$$
\begin{equation*}
E_{i j}^{(k)}=\Sigma_{i j}^{(k)}-S_{i j}^{(k)}=\frac{1}{\tau_{k}} \sum_{\ell=1}^{\tau_{k}}\left\{\mathbb{E}\left(X_{\ell i} X_{\ell j}\right)-X_{\ell i} X_{\ell j}\right\} . \tag{141}
\end{equation*}
$$

Let $u_{\ell}^{(i j)}=X_{\ell i} X_{\ell j}$. Then, $\left|u_{\ell}^{(i j)}\right| \leq 1$ because $\left\|x_{\ell}\right\|_{\infty} \leq 1$. By applying Hoeffding's inequality,

$$
\begin{equation*}
\mathbb{P}\left(\left|E_{i j}^{(k)}\right| \geq 3 \sqrt{\frac{\log d}{\tau_{k}}}\right) \leq \frac{2}{d^{4}} . \tag{142}
\end{equation*}
$$

Therefore, by union bonding over all indices $1 \leq i, j \leq d+1$, we obtain that $\left\|E^{(k)}\right\|_{\infty} \leq$ $3 \sqrt{(\log d) / \tau_{k}}$, with probability at least $1-8 / d^{2}$. The claim follows from this result along with (140).

## C.9. Proof of Lemma 20

Define parameter vectors $\eta_{0,0}=\left(\mu_{0} ; \beta_{0}\right), \eta_{0, \wedge}=\left(\mu_{0} ; \widehat{\beta}^{k}\right), \eta_{\wedge, 0}=\left(\widehat{\mu}^{k} ; \beta_{0}\right)$, and $\eta_{\wedge, \wedge}=$ $\left(\widehat{\mu}^{k} ; \widehat{\beta}^{k}\right)$ all in $\mathbb{R}^{d+1}$. By optimality of $\eta_{\wedge, \wedge}$, we have $\nabla_{\mu} \mathcal{L}\left(\widehat{\mu}^{k}\right)+\lambda_{k} \nu^{k}=0$ and $\nabla_{\beta} \mathcal{L}\left(\widehat{\beta}^{k}\right)+$ $\lambda_{k} \rho^{k}=0$, where $\nu^{k}$ is a subgradient of $\ell_{1}$ norm at $\widehat{\mu}^{k}$ and $\rho^{k}$ is a subgradient of the absolute value function at $\widehat{\beta}^{k}$. Therefore,

$$
\begin{align*}
& \nabla \mathcal{L}\left(\eta_{0, \wedge}\right)=\left(\nabla_{\mu} \mathcal{L}\left(\mu_{0}\right) ; \nabla_{\beta} \mathcal{L}\left(\widehat{\beta}^{k}\right)\right)=\left(\nabla_{\mu} \mathcal{L}\left(\mu_{0}\right) ;-\lambda_{k} \rho^{k}\right),  \tag{143}\\
& \nabla \mathcal{L}\left(\eta_{\wedge, 0}\right)=\left(\nabla_{\mu} \mathcal{L}\left(\widehat{\mu}^{k}\right) ; \nabla_{\beta} \mathcal{L}\left(\beta_{0}\right)\right)=\left(-\lambda_{k} \nu^{k} ; \nabla_{\mu} \mathcal{L}\left(\beta_{0}\right)\right) .
\end{align*}
$$

By a similar argument following (88), we have $\left\|\nabla \mathcal{L}\left(\eta_{0,0}\right)\right\|_{\infty} \leq \lambda_{k} / 2$, with probability at least $1-1 / d$. Hence, invoking relations (143), we get

$$
\begin{equation*}
\left\|\nabla \mathcal{L}\left(\eta_{0, \wedge}\right)\right\|_{\infty} \leq \lambda_{k}, \quad\left\|\nabla \mathcal{L}\left(\eta_{\wedge, 0}\right)\right\|_{\infty} \leq \lambda_{k} \tag{144}
\end{equation*}
$$

By optimality of $\eta_{\wedge, \wedge}$, we have

$$
\mathcal{L}\left(\eta_{\wedge, \wedge}\right)+\lambda_{k}\left\|\eta_{\wedge, \wedge}\right\|_{1} \leq \mathcal{L}\left(\eta_{0, \wedge}\right)+\lambda_{k}\left\|\eta_{0, \wedge}\right\|_{1} .
$$

Following a similar argument to (91), we have

$$
\frac{2 \ell_{F}}{\tau_{k-1}}\left\|\widetilde{X}^{(k-1)}\left(\eta_{0, \wedge}-\eta_{\wedge, \wedge}\right)\right\|^{2}+2 \lambda_{k}\left\|\eta_{\wedge, \wedge}\right\|_{1} \leq 2 \lambda_{k}\left\|\eta_{\wedge, \wedge}-\eta_{0, \wedge}\right\|_{1}+2 \lambda_{k}\left\|\eta_{0, \wedge}\right\|_{1} .
$$

Applying triangle inequality and rearranging the terms, this yields to the following inequality:

$$
\begin{equation*}
\frac{2 \ell_{F}}{\tau_{k-1}}\left\|\widetilde{X}^{(k-1)}\left(\eta_{0, \wedge}-\eta_{\wedge, \wedge}\right)\right\|^{2} \leq 4 \lambda_{k}\left\|\eta_{\wedge, \wedge}-\eta_{0, \wedge}\right\|_{1} \leq 8 W \lambda_{k} \tag{145}
\end{equation*}
$$

By substituting for $\eta_{0, \wedge}$ and $\eta_{\wedge, \wedge}$, we get Equation (79).
Likewise, we can write (145) with $\eta_{\wedge, 0}$ instead of $\eta_{0, \wedge}$ :

$$
\begin{equation*}
\frac{2 \ell_{F}}{\tau_{k-1}}\left\|\widetilde{X}^{(k-1)}\left(\eta_{\wedge, 0}-\eta_{\wedge, \wedge}\right)\right\|^{2} \leq 8 W \lambda_{k} \tag{146}
\end{equation*}
$$

By substituting for $\eta_{\wedge, 0}$ and $\eta_{0,0}$ we obtain

$$
\begin{equation*}
\frac{2 \ell_{F}}{\tau_{k-1}}\left(\beta_{0}-\widehat{\beta}^{k}\right)^{2}\left\|p^{(k-1)}\right\|^{2}=\frac{2 \ell_{F}}{\tau_{k-1}}\left\|\tilde{X}^{(k-1)}\left(\eta_{\wedge, 0}-\eta_{\wedge, \wedge}\right)\right\|^{2} \leq 8 W \lambda_{k} \tag{147}
\end{equation*}
$$

where $p^{(k)}=\left(p_{\tau_{k-1}}, \ldots, p_{\tau_{k}-1}\right)^{\top}$ is the vector of prices offered in episode $(k-1)$. Hence, in order to prove (79), it suffices to show that $p_{t} \geq \tilde{c}_{W}$ for all $t$. To see this, note that since $g(v)>0$ it attains its minimum over any bounded interval. Therefore, there exists a constant $c_{W}^{\prime}>0$, such that $g(v)>c_{W}^{\prime}$ for $v \in[-W, W]$. Further, by the constraint in the optimization (28) we have $\left|\widehat{\beta}^{(k-1)}\right| \leq W$. Combining these two facts, we get

$$
\begin{equation*}
p_{t}=\frac{1}{\widehat{\beta}^{(k-1)}} g\left(x_{t} \cdot \widehat{\mu}^{(k-1)}\right) \geq \frac{c_{W}^{\prime}}{W} \equiv \tilde{c}_{W} \tag{148}
\end{equation*}
$$

The result follows from (147) and setting $c_{W}=4 W / \tilde{c}_{W}^{2}$.

## References

Yasin Abbasi-Yadkori, David Pal, and Csaba Szepesvari. Online-to-confidence-set conversions and application to sparse stochastic bandits. In AISTATS, pages 1-9, 2012.

Shipra Agrawal and Nikhil R. Devanur. Bandits with concave rewards and convex knapsacks. In Proceedings of the Fifteenth ACM Conference on Economics and Computation, EC '14, pages 989-1006, 2014.

Airbnb Documentation. Smart pricing: Set prices based on demand. https://www. airbnb. com/help/article/1168/smart-pricing--set-prices-based-on-demand, 2015.

Kareem Amin, Afshin Rostamizadeh, and Umar Syed. Repeated contextual auctions with strategic buyers. In Advances in Neural Information Processing Systems, pages 622-630, 2014.

Victor F Araman and René Caldentey. Dynamic pricing for nonperishable products with demand learning. Operations research, 57(5):1169-1188, 2009.

Moshe Babaioff, Shaddin Dughmi, Robert Kleinberg, and Aleksandrs Slivkins. Dynamic pricing with limited supply. In Proceedings of the 13th ACM Conference on Electronic Commerce, EC '12, pages 74-91, 2012. ISBN 978-1-4503-1415-2.

Ashwinkumar Badanidiyuru, Robert Kleinberg, and Aleksandrs Slivkins. Bandits with knapsacks. In Foundations of Computer Science (FOCS), 2013 IEEE 54th Annual Symposium on, pages 207-216. IEEE, 2013.

Mark Bagnoli and Ted Bergstrom. Log-concave probability and its applications. Economic theory, 26(2):445-469, 2005.

Santiago R Balseiro, Jon Feldman, Vahab Mirrokni, and S Muthukrishnan. Yield optimization of display advertising with ad exchange. Management Science, 60(12):2886-2907, 2014.

Hamsa Bastani and Mohsen Bayati. Online decision-making with high-dimensional covariates. Working Paper, 2016.

Omar Besbes and Assaf Zeevi. Dynamic pricing without knowing the demand function: risk bounds and near-optimal algorithms. Operations Research, 57:1407-1420, 2009.

Sonia A Bhaskar and Adel Javanmard. 1-bit matrix completion under exact low-rank constraint. In Information Sciences and Systems (CISS), 201549 th Annual Conference on, pages 1-6. IEEE, 2015.

Stephen Boyd and Lieven Vandenberghe. Convex optimization. Cambridge university press, 2004.

Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. Foundations and Trends® in Machine Learning, 3(1):1-122, 2011.

Josef Broder and Paat Rusmevichientong. Dynamic pricing under a general parametric choice model. Operations Research, 60(4):965-980, 2012.

Peter Bühlmann and Sara van de Geer. Statistics for high-dimensional data. SpringerVerlag, 2011.

Emmanuel Candes and Terence Tao. The dantzig selector: Statistical estimation when p is much larger than n. The Annals of Statistics, pages 2313-2351, 2007.

Xi Chen, Zachary Owen, Clark Pixton, and David Simchi-Levi. A statistical learning approach to personalization in revenue management. Working Paper, 2015.

Wang Chi Cheung, Will Ma, David Simchi-Levi, and Xinshang Wang. Inventory balancing with online learning. Working Paper, 2018.

Maxime C Cohen, Ilan Lobel, and Renato Paes Leme. Feature-based dynamic pricing. ACM Conference on Economics and Computation, 2016.

Varsha Dani, Thomas P Hayes, and Sham M Kakade. Stochastic linear optimization under bandit feedback. In COLT, pages 355-366, 2008.
A. V. den Boer and A. P. Zwart. Mean square convergence rates for maximum(quasi) likelihood estimation. Stochastic systems, 4:1-29, 2014. ISSN 1946-5238.

Arnoud V den Boer. Dynamic pricing and learning: historical origins, current research, and new directions. Surveys in operations research and management science, 20(1):1-18, 2015.

Arnoud V den Boer and Bert Zwart. Simultaneously learning and optimizing using controlled variance pricing. Management Science, 60(3):770-783, 2013.

Benjamin Edelman, Michael Ostrovsky, and Michael Schwarz. Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. The American economic review, 97(1):242-259, 2007.

Ossama Elshiewy, Daniel Guhl, and Yasemin Boztug. Multinomial logit models in marketing-from fundamentals to state-of-the-art. Marketing ZFP, 39(3):32-49, 2017.

Vivek F Farias, Srikanth Jagabathula, and Devavrat Shah. A nonparametric approach to modeling choice with limited data. Management Science, 59(2):305-322, 2013.

Kris Johnson Ferreira, David Simchi-Levi, and He Wang. Online network revenue management using thompson sampling. Operations research, 2018.

Dennis H Gensch and Wilfred W Recker. The multinomial, multiattribute logit choice model. Journal of Marketing Research, pages 124-132, 1979.

Alexander Goldenshluger and Assaf Zeevi. A linear response bandit problem. Stochastic Systems, 3(1):230-261, 2013.

Negin Golrezaei, Hamid Nazerzadeh, and Paat Rusmevichientong. Real-time optimization of personalized assortments. Management Science, 60(6):1532-1551, 2014.

Negin Golrezaei, Adel Javanmard, and Vahab Mirrokni. Dynamic incentive-aware learning: Robust pricing in contextual auctions. http://dx.doi.org/10.2139/ssrn.3144034, 2018.

Peter M Guadagni and John DC Little. A logit model of brand choice calibrated on scanner data. Marketing Science, 27(1):29-48, 2008.

J Michael Harrison, Bora Keskin, and Assaf Zeevi. Bayesian dynamic pricing policies: Learning and earning under a binary prior distribution. Management Science, 58(3): 570-586, 2012.

John R Hauser and Frank S Koppelman. Alternative perceptual mapping techniques: Relative accuracy and usefulness. Journal of marketing Research, pages 495-506, 1979.

Adel Javanmard. Perishability of data: Dynamic pricing under varying-coefficient models. Journal of Machine Learning Research, 18(53):1-31, 2017. URL http://jmlr.org/ papers/v18/17-061.html.

Nathan Kallus and Madeleine Udell. Dynamic assortment personalization in high dimensions. Working Paper, 2016.

Godfrey Keller and Sven Rady. Optimal experimentation in a changing environment. The review of economic studies, 66(3):475-507, 1999.

Bora Keskin. Optimal dynamic pricing with demand model uncertainty: A squared-coefficient-of-variation rule for learning and earning. Working Paper, 2014.

Bora Keskin and Assaf Zeevi. Dynamic pricing with an unknown demand model: Asymptotically optimal semi-myopic policies. Operations Research, 62(5):1142-1167, 2014.

N Bora Keskin and Assaf Zeevi. Chasing demand: Learning and earning in a changing environment. Mathematics of Operations Research, 2016.

Robert Kleinberg and Tom Leighton. The value of knowing a demand curve: Bounds on regret for online posted-price auctions. In Proceedings of 44 th Annual IEEE Symposium on Foundations of Computer Science, pages 594-605. IEEE, 2003.

Sébastien Lahaie and David M Pennock. Revenue analysis of a family of ranking rules for keyword auctions. In Proceedings of the 8th ACM conference on Electronic commerce, pages 50-56. ACM, 2007.

Jordan J Louviere and George Woodworth. Design and analysis of simulated consumer choice or allocation experiments: an approach based on aggregate data. Journal of marketing research, pages 350-367, 1983.

Daniel McFadden. Economic choices. American economic review, 91(3):351-378, 2001.
Roger B. Myerson. Optimal auction design. Mathematics of Operations Research, 6(1): 58-73, 1981.

Yaniv Plan and Roman Vershynin. One-bit compressed sensing by linear programming. Communications on Pure and Applied Mathematics, 66(8):1275-1297, 2013.

Stanislav P Ponomarev. Submersions and preimages of sets of measure zero. Siberian Mathematical Journal, 28(1):153-163, 1987.

Sheng Qiang and Mohsen Bayati. Dynamic pricing with demand covariates. Working Paper, 2016.

Garvesh Raskutti, Martin J Wainwright, and Bin Yu. Minimax rates of estimation for highdimensional linear regression over-balls. IEEE Transactions on Information Theory, 57 (10):6976-6994, 2011.

Michael Rothschild. A two-armed bandit theory of market pricing. Journal of Economic Theory, 9(2):185-202, 1974.

Mark Rudelson and Shuheng Zhou. Reconstruction from anisotropic random measurements. IEEE Trans. on Inform. Theory, 59(6):3434-3447, 2013.

Kalyan T Talluri and Garrett J Van Ryzin. The theory and practice of revenue management, volume 68. Springer Science \& Business Media, 2006.
A.B. Tsybakov. Introduction to Nonparametric Estimation. Springer Series in Statistics. Springer New York, 2008. ISBN 9780387790527. URL https://books.google.com/ books?id=mwB8rUBsbqoC.

Sara A Van de Geer. High-dimensional generalized linear models and the lasso. The Annals of Statistics, pages 614-645, 2008.

Roman Vershynin. Introduction to the non-asymptotic analysis of random matrices. arXiv preprint arXiv:1011.3027, 2010.

Zizhuo Wang, Shiming Deng, and Yinyu Ye. Close the gaps: A learning-while-doing algorithm for single-product revenue management problems. Operations Research, 62(2): 318-331, 2014.

Baichun Xiao, Wei Yang, and Jun Li. Optimal reserve price for the generalized second-price auction in sponsored search advertising. Journal of Electronic Commerce Research, 10 (3):114, 2009.

Yuhong Yang and Andrew Barron. Information-theoretic determination of minimax rates of convergence. Annals of Statistics, pages 1564-1599, 1999.

Bin Yu. Assouad, fano and le, cam. In Research Papers in Probability and Statistics: Festschrift in Honor of Lucien Le Cam, pages 423-435. Springer-Verlag, 1997.


[^0]:    (C)2019 Adel Javanmard and Hamid Nazerzadeh.

    License: CC-BY 4.0, see https://creativecommons.org/licenses/by/4.0/. Attribution requirements are provided at http://jmlr.org/papers/v20/17-357.html

[^1]:    1. It is useful to compare Theorem 3 with the result of Chen et al. (2015), [Theorem 5.4], which also uses a maximum likelihood estimator without exploiting sparsity structure and achieves regret $O\left(\frac{d}{C_{\text {min }}} \sqrt{\log (T d) T}\right)$.
[^2]:    4. Specifically, they consider $d(p ; \mathbf{z})=0.5+\mathbf{z}-\mathbf{z} p$. Hence $d(1 ; \mathbf{z})=1$, for all $\mathbf{z}$ and it is shown that $p^{*}\left(\mathbf{z}_{0}\right)=1$ for $\mathbf{z}_{0}=0.5$.
