## APPENDIX

## A Other related works

Here we will briefly review the history and recent studies that are related to our work. For the historical introductions, we mainly refer to den Boer [2015] as a survey. For bandit approaches, we will review some works that apply bandit algorithms to settle pricing problems. For the structural models, we will introduce different modules based on the review in Chan et al. [2009]. Based on the existing works, we might have a better view of our problem setting and methodology.

## A. 1 History of Pricing

It was the work of Cournot [1897] in 1897 that firstly applied mathematics to analyze the relationship between prices and demands. In that work, the price was denoted as $p$ and the demand was defined as a demand function $F(p)$. Therefore, the revenue could be written as $p F(p)$. This was a straightforward interpretation of the general pricing problem, and the key to solving it was estimations of $F(p)$ regarding different products. Later in 1938, the work Schultz et al. [1938] proposed price-demand measurements on exclusive kinds of products. It is worth mentioning that these problems are "static pricing" ones, because $F$ is totally determined by price $p$ and we only need to insist on the optimal one to maximize our profits.

However, the static settings were qualified by the following two observations: on the one hand, a demand function may not only depends on the static value of $p$, but also be affected by the trend of $p$ 's changing [Evans, 1924, Mazumdar et al., 2005]; on the other hand, even if $F(p)$ is static, $p$ itself might change over time according to other factors such as inventory level [Kincaid and Darling, 1963]. As a result, it is necessary to consider dynamics in both demand and price, which leads to a "dynamic pricing" problem setting.

## A. 2 Dynamic Pricing as Bandits

As is said in Section 2, the pricing problem can be viewed as a stochastic contextual bandits problem [see, e.g., Langford and Zhang, 2007, Agarwal et al., 2014]. Even though we may not know the form of the demand function, we can definitely see feedback of demands, i.e. how many products are sold out, which enables us to learn a better decision-making policy. Therefore, it can be studied in a bandit module. If the demand function is totally agnostic, i.e. the evaluations (the highest prices that customers would accept) come at random or even at adversary over time, then it can be modeled as a Multi-arm bandit (MAB) problem [Whittle, 1980] exactly. In our paper, instead, we focus on selling different products with a great variety of features. This can be characterized as a Contextual bandit (CB) problem [Auer et al., 2002, Langford and Zhang, 2007]. The work Cohen et al. [2020], which applies the "EXP-4" algorithm from Auer et al. [2002], also mentions that "the arms represent prices and the payoffs from the different arms are correlated since the measures of demand evaluated at different price points are correlated random variables". A variety of existing works, including Kleinberg and Leighton [2003], Araman and Caldentey [2009], Chen and Farias [2013], Keskin and Zeevi [2014], Besbes and Zeevi [2015], has been approaching the demand function from a perspective of from either parameterized or non-parameterized bandits.

However, our problem setting is different from a contextual bandits setting in at least two perspectives: feedback and regret. The pricing problem has a specially structured feedback between full information and bandits setting. Specifically, $r_{t}>0$ implies that all policies producing $v<v_{t}$ will end up receiving $r_{t}^{\prime}=v$, and $r_{t}=0$ implies that all policies producing $v>v_{t}$ will end up receiving $r_{t}^{\prime}=0$. However, the missing patterns are confounded with the rewards. Therefore it is non-trivial to leverage this structure to improve the importance sampling approach underlying the algorithm of Agarwal et al. [2014]. We instead consider the natural analog to the linear contextual bandits setting [Chu et al., 2011] ${ }^{9}$ and demonstrate that in this case an exponential improvement in the regret is possible using the additional information from the censored feedback. As for regret, while in contextual bandits it refers to a comparison with the optimal policy, it is here referring to a comparison with the optimal action. In other words, though our approaches (both in EMLP and in ONSP) are finding the true parameter $\theta^{*}$, the regret is defined as the "revenue gap" between the optimal price and our proposed prices. These are actually equivalent in our fully-parametric setting (where we assume a

[^0]linear-valuation-known-noise model), but will differ a lot in partially parametric and totally agnostic settings.

## A. 3 Structural Model

While a totally agnostic model guarantees the most generality, a structural model would help us better understand the mechanism behind the observation of prices and demands. The key to a structural pricing model is the behavior of agents in the market, including customers and/or firms. In other words, the behavior of each side can be described as a decision model. From the perspective of demand (customers), the work Kadiyali et al. [1996] adopts a linear model on laundry detergents market, Iyengar et al. [2007] and Lambrecht et al. [2007] study three-part-tariff pricing problems on wireless and internet services with mixed logit models. Besanko et al. assumed an aggregate logit model on customers in works Besanko et al. [1998] and Besanko et al. [2003] in order to study the competitive behavior of manufacturers in ketchup market. Meanwhile, the supply side is usually assumed to be more strategic, such as Bertrand-Nash behaviors [Kadiyali et al., 1996, Besanko et al., 1998, Draganska and Jain, 2006]. For more details, please see Chan et al. [2009].

## B Proofs

## B. 1 Proof of Lemma 2

Proof. Since $v^{*}=\operatorname{argmax} g(v, u)$, we have:

$$
\begin{aligned}
& \left.\frac{\partial g(v, u)}{\partial v}\right|_{v=v^{*}}=0 \\
\Leftrightarrow & 1-F\left(v^{*}-u\right)-v^{*} \cdot f\left(v^{*}-u\right)=0 \\
\Leftrightarrow & \frac{1-F\left(v^{*}-u\right)}{f\left(v^{*}-u\right)}-\left(v^{*}-u\right)=u
\end{aligned}
$$

Define $\varphi(\omega)=\frac{1-F(\omega)}{f(\omega)}-\omega$, and we take derivatives:

$$
\varphi^{\prime}(\omega)=\frac{-f^{2}(\omega)-(1-F(\omega)) f^{\prime}(w)}{f^{2}(w)}-1=\frac{\mathrm{d}^{2} \log (1-F(\omega))}{\mathrm{d} \omega^{2}} \cdot \frac{(1-F(\omega))^{2}}{(f(\omega))^{2}}-1<-1
$$

where the last equality comes from the strict log-concavity of $(1-F(\omega))$. Therefore, $\varphi(\omega)$ is decreasing and $\varphi(+\infty)=-\infty$. Also, notice $\varphi(-\infty)=+\infty$, we know that for any $u \in \mathbb{R}$, there exists an $\omega$ such that $\varphi(\omega)=u$. For $u \geq 0$, we know that $g(v, u) \geq 0$ for $v \geq 0$ and $g(v, u)<0$ for $v<0$. Therefore, $v^{*} \geq 0$ if $u \geq 0$.

## B. 2 Proofs in Section 5.1

## B.2.1 Proof of Lemma 5

Proof. We again define $\varphi(\omega)=\frac{1-F(\omega)}{f(\omega)}-\omega$ as in Appendix B.1. According to Equation 5, we have:

$$
\begin{align*}
& \left.\frac{\partial g(v, u)}{\partial v}\right|_{v=J(u)}=0 \\
\Rightarrow & 1-F(J(u)-u)-J(u) \cdot f(J(u)-u)=0 \\
\Rightarrow & \varphi(J(u)-u)=u  \tag{19}\\
\Rightarrow & J(u)=u+\varphi^{-1}(u) \\
\Rightarrow & J^{\prime}(u)=1+\frac{1}{\varphi^{\prime}\left(\varphi^{-1}(u)\right)} .
\end{align*}
$$

The last line of Equation 19 is due to the Implicit Function Derivatives Principle. From the result in Appendix B.1, we know that $\varphi^{\prime}(\omega)<-1, \forall \omega \in \mathbb{R}$. Therefore, we have $J^{\prime}(u) \in(0,1), u \in \mathbb{R}$, and hence $J(0) \leq J(u)<u+J(0)$ for $u \geq 0$ and $u+J(0) \leq J(u) \leq J(0)$ for $u \leq 0$. Since $u \in[-B, B]$, we may assume that $v \in[0, B+J(0)]$ without losing generality. In the following part, we will frequently use this range.

Denote $u:=x_{t}^{\top} \theta, u^{*}=x_{t}^{\top} \theta^{*}$. According to Equation 7, we know that:

$$
\begin{aligned}
\operatorname{Reg}_{t}(\theta) & =g\left(J\left(u^{*}\right), u^{*}\right)-g\left(J(u), u^{*}\right) \\
& =-\left.\frac{\partial g\left(v, u^{*}\right)}{\partial v}\right|_{v=J\left(u^{*}\right)}\left(J\left(u^{*}-J(u)\right)\right)+\frac{1}{2}\left(-\left.\frac{\partial^{2} g\left(v, u^{*}\right)}{\partial v^{2}}\right|_{v=\tilde{v}}\right)\left(J\left(u^{*}\right)-J(u)\right)^{2} \\
& \leq 0+\frac{1}{2} \max _{\tilde{v} \in[0, B+J(0)]}\left(-\left.\frac{\partial^{2} g\left(v, u^{*}\right)}{\partial v^{2}}\right|_{v=\tilde{v}}\right) \cdot\left(J\left(u^{*}\right)-J(u)\right)^{2} \\
& =\frac{1}{2} \max _{\tilde{v} \in[0, B+J(0)]}\left(2 f\left(\tilde{v}-u^{*}\right)+\tilde{v} \cdot f^{\prime}\left(\tilde{v}-u^{*}\right)\right) \cdot\left(J\left(u^{*}\right)-J(u)\right)^{2} \\
& \leq \frac{1}{2}\left(2 B_{f}+(B+J(0)) \cdot B_{f^{\prime}}\right)\left(J\left(u^{*}\right)-J(u)\right)^{2} \\
& \leq \frac{1}{2}\left(2 B_{f}+(B+J(0)) \cdot B_{f^{\prime}}\right)\left(u^{*}-u\right)^{2} \\
& =\frac{1}{2}\left(2 B_{f}+(B+J(0)) \cdot B_{f^{\prime}}\right)\left(\theta^{*}-\theta\right)^{\top} x_{t} x_{t}^{\top}\left(\theta^{*}-\theta\right) .
\end{aligned}
$$

Here the first line is from the definition of $g$ and $\operatorname{Reg}(\theta)$, the second line is due to Taylor's Expansion, the third line is from the fact that $J\left(u^{*}\right)$ maximizes $g\left(v, u^{*}\right)$ with respect to $v$, the fourth line is by calculus, the fifth line is from the assumption that $0<f(\omega) \leq B_{f},\left|f^{\prime}(\omega)\right| \leq B_{f^{\prime}}$ and $v \in[0, B+J(0)]$, the sixth line is because of $J^{\prime}(u) \in(0,1), \forall u \in \mathbb{R}$, and the seventh line is from the definition of $u$ and $u^{*}$.

## B.2.2 Proof of Lemma 7

Proof. We take derivatives of $l_{t}(\theta)$, and we get:

$$
\begin{align*}
& l_{t}(\theta)= \mathbb{1}_{t}\left(-\log \left(1-F\left(v_{t}-x_{t}^{\top} \theta\right)\right)\right)+\left(1-\mathbb{1}_{t}\right)\left(-\log \left(F\left(v_{t}-x_{t}^{\top} \theta\right)\right)\right) \\
& \nabla l_{t}(\theta)= \mathbb{1}_{t}\left(-\frac{f\left(v_{t}-x_{t}^{\top} \theta\right)}{1-F\left(v_{t}-x_{t}^{\top} \theta\right)}\right) \cdot x_{t}+\left(1-\mathbb{1}_{t}\right)\left(\frac{f\left(v_{t}-x_{t}^{\top} \theta\right)}{F\left(v_{t}-x_{t}^{\top} \theta\right)}\right) \cdot x_{t} \\
& \nabla^{2} l_{t}(\theta)= \mathbb{1}_{t} \cdot \frac{f\left(v_{t}-x_{t}^{\top} \theta\right)^{2}+f^{\prime}\left(v_{t}-x_{t}^{\top} \theta\right) \cdot\left(1-F\left(v_{t}-x_{t}^{\top} \theta\right)\right)}{\left(1-F\left(v_{t}-x_{t}^{\top} \theta\right)\right)^{2}} \cdot x_{t} x_{t}^{\top} \\
&+\left(1-\mathbb{1}_{t}\right) \cdot \frac{f\left(v_{t}-x_{t}^{\top} \theta\right)^{2}-f^{\prime}\left(v_{t}-x_{t}^{\top} \theta\right) F\left(v_{t}-x_{t}^{\top} \theta\right)}{F\left(v_{t}-x_{t}^{\top} \theta\right)^{2}} \cdot x_{t} x_{t}^{\top} \\
&=\left.\mathbb{1}_{t} \cdot \frac{-\mathrm{d}^{2} \log (1-F(\omega))}{\mathrm{d} \omega^{2}}\right|_{\omega=v_{t}-x_{t}^{\top} \theta} \cdot x_{t} x_{t}^{\top}+\left.\left(1-\mathbb{1}_{t}\right) \frac{-\mathrm{d}^{2} \log (F(\omega))}{\mathrm{d} \omega^{2}}\right|_{\omega=v_{t}-x_{t}^{\top} \theta} \cdot x_{t} x_{t}^{\top} \\
& \succeq \\
& \omega \in[-B, B+J(0)]  \tag{20}\\
&= C_{\text {down }} x_{t} x_{t}^{\top},
\end{align*}
$$

which directly proves the first inequality. For the second inequality, just notice that

$$
\begin{align*}
\nabla l_{t}(\theta) \nabla l_{t}(\theta)^{\top} & =\mathbb{1}_{t}\left(\frac{f\left(v_{t}-x_{t}^{\top} \theta\right)}{1-F\left(v_{t}-x_{t}^{\top} \theta\right)}\right)^{2} x_{t} x_{t}^{\top}+\left(1-\mathbb{1}_{t}\right)\left(\frac{f\left(v_{t}-x_{t}^{\top} \theta\right)}{F\left(v_{t}-x_{t}^{\top} \theta\right)}\right)^{2} x_{t} x_{t}^{\top} \\
& \preceq \sup _{\omega \in[-B, B+J(0)]} \max \left\{\left(\frac{f(\omega)}{F(\omega)}\right)^{2},\left(\frac{f(\omega)}{1-F(\omega)}\right)^{2}\right\} x_{t} x_{t}^{\top}  \tag{21}\\
& =C_{\exp } x_{t} x_{t}^{\top}
\end{align*}
$$

The only thing to point out is that $\frac{f(\omega)}{F(\omega)}$ and $\frac{f(\omega)}{1-F(\omega)}$ are all continuous for $\omega \in[-B, B+J(0)]$, as $F(\omega)$ is strictly increasing and thus $0<F(\omega)<1, \omega \in \mathbb{R}$.

## B.2.3 Proof of Lemma 9

Proof. In the following part, we consider a situation that an epoch of $n \geq 2$ rounds of pricing is conducted, generating $l_{j}(\theta)$ as negative likelihood functions, $j=1,2, \ldots, n$. Define a "leave-one-
out"negative log-likelihood function

$$
\tilde{L}_{i}(\theta)=\frac{1}{n} \sum_{j=1, j \neq i}^{n} l_{j}(\theta),
$$

and let

$$
\tilde{\theta}_{i}:=\underset{\theta}{\arg \min } \tilde{L}_{i}(\theta) .
$$

Based on this definition, we know that $\tilde{\theta}_{i}$ is independent to $l_{i}(\theta)$ given historical data, and that $\tilde{\theta}_{i}$ are identically distributed for all $i=1,2,3, \ldots, n$.
In the following part, we will firstly propose and proof the following inequality:

$$
\begin{equation*}
\frac{1}{n} \sum_{i=1}^{n}\left(l_{i}\left(\tilde{\theta}_{i}\right)-l_{i}(\hat{\theta})\right) \leq \frac{C_{\exp }}{C_{d o w n}} \frac{d}{n}=O\left(\frac{d}{n}\right) \tag{22}
\end{equation*}
$$

where $\hat{\theta}$ is the short-hand notation of $\hat{\theta}_{k}$ as we do not specify the epoch $k$ in this part. We now cite a lemma from Koren and Levy [2015]:

Lemma 13. Let $g_{1}$, $g_{2}$ be 2 convex function defined over a closed and convex domain $\mathcal{K} \subseteq \mathbb{R}^{d}$, and let $x_{1}=\arg \min _{x \in \mathcal{K}} g_{1}(x)$ and $x_{2}=\arg \min _{x \in \mathcal{K}} g_{2}(x)$. Assume $g_{2}$ is locally $\delta$-strongly-convex at $x_{1}$ with respect to a norm $\|\cdot\|$. Then, for $h=g_{2}-g_{1}$ we have

$$
\left\|x_{2}-x_{1}\right\| \leq \frac{2}{\delta}\left\|\nabla h\left(x_{1}\right)\right\|_{*}
$$

Here $\|\cdot\|_{*}$ denotes a dual norm.
The following is a proof of this lemma.
Proof. (of Lemma 13) According to convexity of $g_{2}$, we have:

$$
\begin{equation*}
g_{2}\left(x_{1}\right) \geq g_{2}\left(x_{2}\right)+\nabla g_{2}\left(x_{2}\right)^{\top}\left(x_{1}-x_{2}\right) . \tag{23}
\end{equation*}
$$

According to strong convexity of $g_{2}$ at $x_{1}$, we have:

$$
\begin{equation*}
g_{2}\left(x_{2}\right) \geq g_{2}\left(x_{1}\right)+\nabla g_{2}\left(x_{1}\right)^{\top}\left(x_{2}-x_{1}\right)+\frac{\delta}{2}\left\|x_{2}-x_{1}\right\|^{2} \tag{24}
\end{equation*}
$$

Add Equation (23) and (24), and we have:

$$
\begin{array}{ll} 
& g_{2}\left(x_{1}\right)+g_{2}\left(x_{2}\right) \geq g_{2}\left(x_{2}\right)+g_{2}\left(x_{1}\right)+\left(\nabla g_{2}\left(x_{1}\right)-\nabla g_{2}\left(x_{2}\right)\right)^{\top}\left(x_{2}-x_{1}\right)+\frac{\delta}{2}\left\|x_{2}-x_{1}\right\|^{2} \\
\Leftrightarrow & \left(\nabla g_{2}\left(x_{1}\right)-\nabla g_{2}\left(x_{2}\right)\right)^{\top}\left(x_{1}-x_{2}\right) \geq \frac{\delta}{2}\left\|x_{1}-x_{2}\right\|^{2} \\
\Leftrightarrow & \left(\nabla g_{1}\left(x_{1}\right)+\nabla h\left(x_{1}\right)-\nabla g_{2}\left(x_{2}\right)\right)^{\top}\left(x_{1}-x_{2}\right) \geq \frac{\delta}{2}\left\|x_{1}-x_{2}\right\|^{2} \\
\Leftrightarrow & \nabla h\left(x_{1}\right)^{\top}\left(x_{1}-x_{2}\right) \geq \frac{\delta}{2}\left\|x_{1}-x_{2}\right\|^{2} \\
\Rightarrow \quad & \left\|\nabla h\left(x_{1}\right)\right\|_{*}\left\|x_{1}-x_{2}\right\| \geq \frac{\delta}{2}\left\|x_{1}-x_{2}\right\|^{2} \\
\Rightarrow \quad & \left\|\nabla h\left(x_{1}\right)\right\|_{*} \geq \frac{\delta}{2}\left\|x_{1}-x_{2}\right\| . \tag{25}
\end{array}
$$

The first step is trivial. The second step is a sequence of $g_{2}=g_{1}+h$. The third step is derived by the following 2 first-order optimality conditions: $\nabla g_{1}\left(x_{1}\right)^{\top}\left(x_{1}-x_{2}\right) \leq 0$, and $\nabla g_{2}\left(x_{2}\right)^{\top}\left(x_{2}-x_{1}\right) \leq 0$. The fourth step is derived from Holder's Inequality:

$$
\left\|\nabla h\left(x_{1}\right)\right\|_{*}\left\|x_{1}-x_{2}\right\| \geq \nabla h\left(x_{1}\right)^{\top}\left(x_{1}-x_{2}\right) .
$$

Therefore, the lemma holds.

In the following part, we will set up a strongly convex function of $g_{2}$. Denote $H=\sum_{t=1}^{n} x_{t} x_{t}^{\top}$. From Lemma 7, we know that

$$
\nabla^{2} \hat{L}(\theta) \succeq C_{\text {down }} \frac{1}{n} H
$$

Here $\hat{L}(\theta)$ is the short-hand notation of $\hat{L}_{k}(\theta)$ as we do not specify $k$ in this part. Since we do not know if $H$ is invertible, i.e. if a norm can be induced by $H$, we cannot let $g_{2}(\theta)=\hat{L}(\theta)$. Instead, we change the variable as follows:

We first apply singular value decomposition to $H$, i.e. $H=U \Sigma U^{\top}$, where $U \in \mathbb{R}^{d \times r}, U^{\top} U=$ $I_{r}, \Sigma=\operatorname{diag}\left\{\lambda_{1}, \lambda_{2}, \ldots, \lambda_{r}\right\} \succ 0$. After that, we introduce a new variable $\eta:=U^{\top} \theta$. Therefore, we have $\theta=U \eta+V \epsilon$, where $V \in \mathbb{R}^{d \times(d-r)}, V^{\top} V=I_{d-r}, V^{\top} U=0$ is the standard orthogonal bases of the null space of $U$, and $\epsilon \in \mathbb{R}^{(d-r)}$. Similarly, we define $\tilde{\eta}_{i}=U^{\top} \tilde{\theta}_{i}$ and $\hat{\eta}=U^{\top} \hat{\theta}$. According to these, we define the following functions:

$$
\begin{align*}
f_{i}(\eta) & :=l_{i}(\theta)=l_{i}(U \eta+V \epsilon) \\
\tilde{F}_{i}(\eta) & :=\tilde{L}_{i}(\theta)=\tilde{L}_{i}(U \eta+V \epsilon)  \tag{26}\\
\hat{F}(\eta) & :=\hat{L}(\theta)=\hat{L}(U \eta+V \epsilon)
\end{align*}
$$

Now we prove that $\hat{F}(\eta)$ is locally-strongly-convex. Similar to the proof of Lemma 7, we have:

$$
\begin{align*}
\nabla^{2} \hat{F}(\eta) & =\frac{1}{n} \sum_{i=1}^{n} \nabla^{2} f_{i}(\eta) \\
& =\frac{1}{n} \sum_{i=1}^{n} \frac{\partial^{2} l_{i}}{\partial\left(x_{i}^{\top} \theta\right)^{2}}\left(\frac{\partial x_{i}^{\top} \theta}{\partial \eta}\right)\left(\frac{\partial x_{i}^{\top} \theta}{\partial \eta}\right)^{\top} \\
& =\frac{1}{n} \sum_{i=1}^{n} \frac{\partial^{2} l_{i}}{\partial\left(x_{i}^{\top} \theta\right)^{2}}\left(\frac{\partial x_{i}^{\top}(U \eta+V \epsilon)}{\partial \eta}\right)\left(\frac{\partial x_{i}^{\top}(U \eta+V \epsilon)}{\partial \eta}\right)^{\top} \\
& =\frac{1}{n} \sum_{i=1}^{n} \frac{\partial^{2} l_{i}}{\partial\left(x_{i}^{\top} \theta\right)^{2}}\left(U^{\top} x_{i}\right)\left(U^{\top} x_{i}\right)^{\top} \\
& \succeq \frac{1}{n} \sum_{i=1}^{n} C_{\text {down }} U^{\top} x_{i} x_{i}^{\top} U  \tag{27}\\
& =\frac{1}{n} C_{\text {down }} U^{\top}\left(\sum_{i=1}^{n} x_{i} x_{i}^{\top}\right) U^{\top} \\
& =\frac{1}{n} C_{\text {down }} U^{\top} H U \\
& =\frac{1}{n} C_{\text {down }} U^{\top} U \Sigma U^{\top} U \\
& =\frac{1}{n} C_{\text {down }} \Sigma \\
& \succ 0
\end{align*}
$$

That is to say, $\hat{F}(\eta)$ is locally $\frac{C_{d o w n}}{n}$-strongly convex w.r.t $\Sigma$ at $\eta$. Similarly, we can verify that $\tilde{F}_{i}(\eta)$ is convex (not necessarily strongly convex). Therefore, according to Lemma 13 , let $g_{1}(\eta)=$ $\tilde{F}_{i}(\eta), g_{2}(\eta)=\hat{F}(\eta)$, and then $x_{1}=\tilde{\eta}_{i}=U^{\top} \tilde{\theta}_{i}, x_{2}=\hat{\eta}=U^{\top} \hat{\theta}$. Therefore, we have:

$$
\begin{equation*}
\left\|\hat{\eta}-\tilde{\eta}_{i}\right\|_{\Sigma} \leq \frac{1}{C_{\text {down }}}\left\|\nabla f_{i}\left(\tilde{\eta}_{i}\right)\right\|_{\Sigma}^{*} \tag{28}
\end{equation*}
$$

Now let us show the validation of this theorem:

$$
\begin{align*}
& l_{i}\left(\tilde{\theta}_{i}\right)-l_{i}(\hat{\theta})=f_{i}\left(\tilde{\eta}_{i}\right)-f_{i}(\hat{\eta}) \\
& \leq \nabla f_{i}\left(\tilde{\eta}_{i}\right)^{\top}\left(\tilde{\eta}_{i}-\hat{\eta}\right) \\
& \text { convexity }  \tag{29}\\
& \leq\left\|\nabla f_{i}\left(\tilde{\eta}_{i}\right)\right\|_{\Sigma}^{*}\left\|\tilde{\eta}_{i}-\hat{\eta}\right\|_{\Sigma} \\
& \text { Holder inequality } \\
& \underset{\substack{\uparrow \\
\text { Lemma } 13}}{\leq} \frac{1}{C_{\text {down }}}\left(\left\|\nabla f_{i}\left(\tilde{\eta}_{i}\right)\right\|_{\Sigma}^{*}\right)^{2} .
\end{align*}
$$

And thus we have

$$
\begin{align*}
\sum_{i=1}^{n} l_{i}\left(\tilde{\theta}_{i}\right)-l_{i}(\hat{\theta}) & \left.\leq \frac{1}{C_{\text {down }}} \sum_{i=1}^{n}\left\|\nabla f_{i}\left(\tilde{\eta}_{i}\right)\right\|_{\Sigma}^{*}\right)^{2} \\
& \leq \frac{1}{C_{\text {down }}} \sum_{i=1}^{n}\left(\frac{p}{\Phi}\right)_{\max }^{2} x_{i}^{\top} U \Sigma^{-1} U^{\top} x_{i} \\
& =\frac{C_{\exp }}{C_{\text {down }}} C_{\exp } \sum_{i=1}^{n} \operatorname{tr}\left(x_{i}^{\top} U \Sigma^{-1} U^{\top} x_{i}\right) \\
& =\frac{C_{\exp }}{C_{\text {down }}} C_{\exp } \sum_{i=1}^{n} \operatorname{tr}\left(U \Sigma^{-1} U^{\top} x_{i} x_{i}^{\top}\right) \\
& =\frac{C_{\exp }}{C_{\text {down }}} C_{\exp } t r\left(U \Sigma^{-1} U^{\top} \sum_{i=1}^{n} x_{i} x_{i}^{\top}\right) \\
& =\frac{C_{\exp }}{C_{\text {down }}} C_{\exp } t r\left(U \Sigma^{-1} U^{\top} H\right)  \tag{30}\\
& =\frac{C_{\text {exp }}}{C_{\text {down }}} C_{\exp } \operatorname{tr}\left(U \Sigma^{-1} U^{\top} U \Sigma U^{\top}\right) \\
& =\frac{C_{\exp }}{C_{\text {down }}} C_{\exp } t r\left(U U^{\top}\right) \\
& =\frac{C_{\exp }}{C_{\text {down }}} C_{\exp } t r\left(U^{\top} U\right) \\
& =\frac{C_{\exp }}{C_{\text {down }}} C_{\exp } t r\left(I_{r}\right) \\
& =\frac{C_{\exp }}{C_{\text {down }}} C_{\exp } r \\
& \leq \frac{C_{\exp }}{C_{\text {down }}} d .
\end{align*}
$$

Thus the Inequality 22 is proved. After that, we have:

$$
\begin{aligned}
& \mathbb{E}_{h}\left[L\left(\tilde{\theta}_{n}\right)\right]-L\left(\theta^{*}\right) \\
= & \mathbb{E}_{h}\left[L\left(\tilde{\theta}_{n}\right)\right]-\mathbb{E}_{h}\left[\hat{L}\left(\theta^{*}\right)\right] \\
\leq & \mathbb{E}_{h}\left[L\left(\tilde{\theta}_{n}\right)\right]-\mathbb{E}_{h}[\hat{L}(\hat{\theta})] \\
= & \frac{1}{n} \sum_{i=1}^{n} \mathbb{E}_{h}\left[L\left(\tilde{\theta}_{i}\right)\right]-\mathbb{E}_{h}[\hat{L}(\hat{\theta})] \\
= & \frac{1}{n} \sum_{i=1}^{n} \mathbb{E}_{h}\left[l_{i}\left(\tilde{\theta}_{i}\right)\right]-\mathbb{E}_{h}[\hat{L}(\hat{\theta})] \\
= & \frac{1}{n} \sum_{i=1}^{n} \mathbb{E}_{h}\left[l_{i}\left(\tilde{\theta}_{i}\right)\right]-\sum_{i=1}^{n} \mathbb{E}_{h}\left[l_{i}(\hat{\theta})\right] \\
= & \frac{1}{n} \sum_{i=1}^{n} \mathbb{E}_{h}\left[l_{i}\left(\tilde{\theta}_{i}\right)-l_{i}(\hat{\theta})\right] \\
\leq & \frac{C_{\text {exp }}}{C_{\text {down }}} \frac{d}{n} \\
= & O\left(\frac{d}{n}\right)
\end{aligned}
$$

Thus we has proved that $\mathbb{E}_{h}\left[L\left(\tilde{\theta}_{n}\right)\right]-L\left(\theta^{*}\right) \leq \frac{C_{\text {exp }}}{C_{\text {down }}} \cdot \frac{d}{n}$. Notice that $\tilde{\theta}_{n}$ is generated by optimizing the leave-one-out likelihood function $\tilde{L}_{n}(\theta)=\sum_{j=1}^{n-1} l_{j}(\theta)$, which does not contain $l_{n}(\theta)$, and that the expected likelihood function $L(\theta)$ does not depend on any specific result occurring in this round. That is to say, every term of this inequality is not related to the last round $\left(x_{n}, v_{n}, \mathbb{1}_{n}\right)$ at all. In other words, this inequality is still valid if we only conduct this epoch from round 1 to $(n-1)$.

Now let $n=\tau+1$, and then we know that $\tilde{\theta}_{\tau+1}=\hat{\theta}$. Therefore, the theorem holds.

## B. 3 Proof of Lower bound in Section 5.3

Proof. We assume a fixed $u^{*}$ such that $x^{\top} \theta^{*}=u^{*}, \forall x \in \mathbb{D}$. In other words, we are considering a non-context setting. Therefore, we can define a policy as $\Psi:\{0,1\}^{t} \rightarrow \mathbb{R}^{+}, t=1,2, \ldots$ that does not observe $x_{t}$ at all. Before the proof begins, we firstly define a few notations: We denote $\Phi_{\sigma}(\omega)$ and $p_{\sigma}(\omega)$ as the CDF and PDF of Gaussian distribution $\mathcal{N}\left(0, \sigma^{2}\right)$, and the corresponding $J_{\sigma}(u)=\arg \max _{v} v\left(1-\Phi_{\sigma}(v-u)\right)$ as the pricing function.

Since we have proved that $J^{\prime}(u) \in(0,1)$ for $u \in \mathbb{R}$ in Appendix B.2.2, we have the following lemma:

Lemma 14. $u-J_{\sigma}(u)$ monotonically increases as $u \in(0,+\infty), \forall \sigma>0$. Also, we know that $J_{\sigma}(0)>0, \forall \sigma>0$.

Now consider the following cases: $\sigma_{1}=1, \sigma_{2}=1-f(T)$, where $\lim _{T \rightarrow \infty} f(T)=0, f^{\prime}(T)<$ $0,0<f(T)<\frac{1}{2}$. We will later determine the explicit form of $f(T)$.

Suppose $u^{*}$ satisfies $J_{\sigma_{1}}\left(u^{*}\right)=u^{*}$. Solve it and get $u^{*}=\sqrt{\frac{\pi}{2}}$. Therefore, we have $u \in\left(0, u^{*}\right) \Leftrightarrow$ $J_{1}(u)>u$, and $u \in\left(u^{*},+\infty\right) \Leftrightarrow J_{1}(u)<u$. As a result, we have the following lemma.

Lemma 15. For any $\sigma \in\left(\frac{1}{2}, 1\right)$, we have:

$$
\begin{equation*}
J_{\sigma}\left(u^{*}\right) \in\left(0, u^{*}\right) \tag{31}
\end{equation*}
$$

Proof. Firstly, we have:

$$
\begin{aligned}
J_{\sigma}(u) & =\underset{v}{\arg \max } v \Phi_{\sigma}(u-v) \\
& =\underset{v}{\arg \max } v \Phi_{1}\left(\frac{u-v}{\sigma}\right) \\
& =\underset{\omega=\frac{v}{\sigma}}{\arg \max } \sigma \omega \Phi_{1}\left(\frac{u}{\sigma}-\omega\right) \\
& =\sigma \underset{\omega}{\arg \max } \Phi_{1}\left(\frac{u}{\sigma}-\omega\right) \\
& =\sigma J_{1}\left(\frac{u}{\sigma}\right)
\end{aligned}
$$

When $\sigma \in\left(\frac{1}{2}, 1\right)$, we know $\frac{u^{*}}{\sigma}>u^{*}$. Since $J_{1}\left(u^{*}\right)=u^{*}$ and that $u \in\left(u^{*},+\infty\right) \Leftrightarrow J_{1}(u)<u$, we have $\frac{u^{*}}{\sigma}>J_{1}\left(\frac{u^{*}}{\sigma}\right)$. Hence

$$
\begin{equation*}
u^{*}>\sigma J_{1}\left(\frac{u^{*}}{\sigma}\right)=J_{\sigma}\left(u^{*}\right) \tag{32}
\end{equation*}
$$

Therefore, without losing generality, we assume that for the problem parameterized by $\sigma_{2}$, the price $v \in\left(0, u^{*}\right)$. To be specific, suppose $v^{*}(\sigma)=J_{\sigma}\left(u^{*}\right)$. Define $\Psi_{t+1}:[0,1]^{t} \rightarrow\left(0, u^{*}\right)$ as any policy that proposes a price at time $t+1$. Define $\Psi=\left\{\Psi_{1}, \Psi_{2}, \ldots, \Psi_{T-1}, \Psi_{T}\right\}$.
Define the sequence of price as $V=\left\{v_{1}, v_{2}, \ldots, v_{T-1}, v_{T}\right\}$, and the sequence of decisions as $\mathbb{1}=\left\{\mathbb{1}_{1}, \mathbb{1}_{2}, \ldots, \mathbb{1}_{T-1}, \mathbb{1}_{T}\right\}$. Denote $V^{t}=\left\{v_{1}, v_{2}, \ldots, v_{t},\right\}$.
Define the probability (also the likelihood if we change $u^{*}$ to other parameter $u$ ):

$$
\begin{equation*}
Q_{T}^{V, \sigma}(\mathbb{1})=\prod_{t=1}^{T} \Phi_{\sigma}\left(u^{*}-v_{t}\right)^{\mathbb{1}_{t}} \Phi_{\sigma}\left(v_{t}-u^{*}\right)^{1-\mathbb{1}_{t}} \tag{33}
\end{equation*}
$$

Define a random variable $Y_{t} \in\{0,1\}^{t}, Y_{t} \sim Q_{t}^{V^{t}, \sigma}$ and one possible assignment $y_{t}=\left\{\mathbb{1}_{1}, \mathbb{1}_{2}, \ldots, \mathbb{1}_{t-1}, \mathbb{1}_{t}\right\}$. For any price $v$ and any parameter $\sigma$, define the expected reward function as $r(v, \sigma):=v \Phi_{\sigma}\left(u^{*}-v\right)$. Based on this, we can further define the expected regret $\operatorname{Regret}(\sigma, T, \Psi)$ :

$$
\begin{equation*}
\operatorname{Regret}(\sigma, T, \Psi)=\mathbb{E}\left[\sum_{t=1}^{T} r\left(J_{\sigma}\left(u^{*}\right), \sigma\right)-r\left(\Psi_{t}\left(y_{t-1}\right), \sigma\right)\right] \tag{34}
\end{equation*}
$$

Now we have the following properties:
Lemma 16. 1. $r\left(v^{*}(\sigma), \sigma\right)-r(v, \sigma) \geq \frac{1}{60}\left(v^{*}(\sigma)-v\right)^{2}$;
2. $\left|v^{*}(\sigma)-u^{*}\right| \geq \frac{2}{5}|1-\sigma|$;
3. $\left|\Phi_{\sigma}\left(u^{*}-v\right)-\Phi_{1}\left(u^{*}-v\right)\right| \leq\left|u^{*}-v\right| \cdot|\sigma-1|$.

Proof. 1. We have:

$$
\left.\frac{\partial r(v, \sigma)}{\partial v}\right|_{v=v^{*}(\sigma)}=0 \frac{\partial^{2} r(v, \sigma)}{\partial v^{2}}=\frac{1}{\sigma^{2}}\left(v^{2}-u^{*} v-2 \sigma^{2}\right) p_{\sigma}\left(u^{*}-v\right)
$$

Since $v \in\left(0, u^{*}\right)$, we have $\left(v^{2}-u^{*} v-2 \sigma^{2}\right)<-2 \sigma^{2}$. Also, since $\sigma \in(1 / 2,1)$, we have $p_{\sigma}\left(u^{*}-v\right)>\frac{1}{\sqrt{2 \pi}} \cdot e^{-\frac{\left(u^{*}\right)^{2}}{2 \cdot(1 / 2)^{2}}}=\frac{1}{\sqrt{2 \pi} e^{\pi}}>0.017$. Therefore, we have

$$
\frac{\partial^{2} r(v, \sigma)}{\partial v^{2}}<-2 * 0.017<-\frac{1}{30}
$$

As a result, we have:

$$
\begin{align*}
r(v *(\sigma), \sigma)-r(v, \sigma) & =-\left.(v *(\sigma)-v) \frac{\partial r(v, \sigma)}{\partial v}\right|_{v=v^{*}(\sigma)}-\left.\frac{1}{2}(v *(\sigma)-v)^{2} \frac{\partial^{2} r(v, \sigma)}{\partial v^{2}}\right|_{v=\tilde{v}} \\
& =0-\left.\frac{1}{2}(v *(\sigma)-v)^{2} \frac{\partial^{2} r(v, \sigma)}{\partial v^{2}}\right|_{v=\tilde{v}} \\
& \geq \frac{1}{2} \cdot \frac{1}{30}(v *(\sigma)-v)^{2} \tag{35}
\end{align*}
$$

2. According to Equation 32, we know that:

$$
v^{*}(\sigma)=\sigma J_{1}\left(\frac{u^{*}}{\sigma}\right)
$$

For $u \in\left(u^{*},+\infty\right), J_{1}(u)<u$. According to Lemma 14 , we have:

$$
\begin{aligned}
J_{1}^{\prime}(u) & =1+\frac{1}{J_{1}(u)\left(J_{1}(u)-u\right)-2} \\
& >1+\frac{1}{0-2} \\
& =\frac{1}{2}
\end{aligned}
$$

Also, for $u \in\left(u^{*}, \frac{u^{*}}{\sigma}\right)$, we have:

$$
\begin{aligned}
& J_{1}^{\prime}(u)=1-\frac{1}{2+J_{1}(u)\left(u-J_{1}(u)\right)} \\
& \underset{\substack{\uparrow \\
0<J_{1}(u)<u}}{\leq 1-\frac{1}{2+u(u-0)}} \\
& \underset{\substack{\uparrow \\
u<\frac{u^{*}}{\sigma}} \frac{u^{*}}{2}}{\leq} 1-\frac{1}{2+\left(\frac{u^{*}}{2}\right)^{2}} \\
& =1-\frac{1}{2+\frac{\pi}{8}} \\
& <\frac{3}{5} \text {. }
\end{aligned}
$$

Therefore, we have:

$$
\begin{aligned}
J_{1}\left(u^{*}\right)-J_{\sigma}\left(u^{*}\right) & =J_{1}\left(u^{*}\right)-\sigma J_{1}\left(\frac{u^{*}}{\sigma}\right) \\
& =J_{1}\left(u^{*}\right)-\sigma J_{1}\left(u^{*}\right)+\sigma\left(J_{1}\left(u^{*}\right)-J_{1}\left(\frac{u^{*}}{\sigma}\right)\right) \\
& =J_{1}\left(u^{*}\right)(1-\sigma)-\sigma\left(J_{1}\left(\frac{u^{*}}{\sigma}\right)-J_{1}\left(u^{*}\right)\right) \\
& >u^{*}(1-\sigma)-\sigma \cdot \frac{3}{5}\left(\frac{u^{*}}{\sigma}-u^{*}\right) \\
& =u^{*}(1-\sigma)-\frac{3}{5} \sigma\left(\frac{1}{\sigma}-1\right) u^{*} \\
& =u^{*}(1-\sigma)\left(1-\frac{3}{5}\right) \\
& >\frac{2}{5}(1-\sigma)
\end{aligned}
$$

3. This is because:

$$
\begin{aligned}
&\left|\Phi_{\sigma}\left(u^{*}-v\right)-\Phi_{1}\left(u^{*}-v\right)\right|=\left|\Phi_{\sigma}\left(u^{*}-v\right)-\Phi_{\sigma}\left(\sigma\left(u^{*}-v\right)\right)\right| \\
& \leq \max \left|p_{\sigma}\right| \cdot\left|\left(u^{*}-v\right)-\sigma\left(u^{*}-v\right)\right| \\
& \leq \frac{1}{\sqrt{2 \pi} \sigma} \cdot\left|\left(u^{*}-v\right)-\sigma\left(u^{*}-v\right)\right| \\
& \leq(1-\sigma)\left|u^{*}-v\right| . \\
& \substack{\uparrow>\frac{1}{2}}
\end{aligned}
$$

In the following part, we will propose two theorems, which balance the cost of learning and that of uncertainty. This part is mostly similar to [BR12] Section 3, but we adopt a different family of demand curves here.

Theorem 17 (Learning is costly). Let $\sigma \in(1 / 2,1)$ and $v_{t} \in\left(0, u^{*}\right)$, and we have:

$$
\begin{equation*}
\mathcal{K}\left(Q^{V, 1} ; Q^{V, \sigma}\right)<9900(1-\sigma)^{2} \operatorname{Regret}(1, T, \Psi) \tag{36}
\end{equation*}
$$

Here $v_{t}=\Psi\left(y_{t-1}\right), t=1,2, \ldots, T$.
Proof. First of all, we cite the following lemma that would facilitate the proof.
Lemma 18 (Corollary 3.1 in Taneja and Kumar, 2004). Suppose $B_{1}$ and $B_{2}$ are distributions of Bernoulli random variables with parameters $q_{1}$ and $q_{2}$, respectively, with $q_{1}, q_{2} \in(0,1)$. Then,

$$
\mathcal{K}\left(B_{1} ; B_{2}\right) \leq \frac{\left(q_{1}-q_{2}\right)^{2}}{q_{2}\left(1-q_{2}\right)}
$$

According to the definition of KL-divergence, we have:

$$
\mathcal{K}\left(Q_{T}^{V, 1} ; Q_{T}^{V, \sigma}\right)=\sum_{s=1}^{T} \mathcal{K}\left(Q_{s}^{V^{s}, 1} ; Q_{s}^{V^{s}, \sigma} \mid Y_{s-1}\right)
$$

For each term of the RHS, we have:

$$
\begin{aligned}
& \mathcal{K}\left(Q_{s}^{V^{s}, 1}, Q_{s}^{V^{s}, \sigma} \mid Y_{s-1}\right) \\
& =\sum_{y_{s} \in\{0,1\}^{s}} Q_{s}^{V^{s}, 1}\left(y_{s}\right) \log \left(\frac{Q_{s}^{V^{s}, 1}\left(\mathbb{1}_{s} \mid y_{s-1}\right)}{Q_{s}^{V^{s}, \sigma}\left(\mathbb{1}_{s} \mid y_{s-1}\right)}\right) \\
& \underset{\substack{\uparrow \\
\text { split } y_{s} \text { as } y_{s-1} \text { and ind } d_{s}}}{=} \sum_{y_{s-1} \in\{0,1\}^{s-1}} Q_{s-1}^{V^{s-1}, 1}\left(y_{s-1}\right) \cdot \sum_{\mathbb{1}_{s} \in\{0,1\}} Q_{s}^{V^{s}, 1}\left(\mathbb{1}_{s} \mid y_{s-1}\right) \log \left(\frac{Q^{V^{s}, 1}\left(\mathbb{1}_{s} \mid y_{s-1}\right)}{Q_{s}^{V^{s}, \sigma}\left(\mathbb{1}_{s} \mid y_{s-1}\right)}\right) \\
& =\quad \sum_{y_{s-1} \in\{0,1\}^{s-1}} Q_{s-1}^{V^{s-1}, 1}\left(y_{s-1}\right) \mathcal{K}\left(Q_{s}^{V^{s}, 1}\left(\cdot \mid y_{s-1}\right), Q_{s}^{V^{s}, \sigma}\left(\cdot \mid y_{s-1}\right)\right) \\
& \underset{\substack{\uparrow \\
\text { Lemma } 18}}{\leq} \sum_{y_{s-1} \in\{0,1\}^{s-1}} Q_{s-1}^{V^{s-1}, 1}\left(y_{s-1}\right) \frac{\left(\Phi_{1}\left(u^{*}-v_{s}\right)-\Phi_{\sigma}\left(u^{*}-v_{s}\right)\right)^{2}}{\Phi_{\sigma}\left(u^{*}-v_{s}\right)\left(1-\Phi_{\sigma}\left(u^{*}-v_{s}\right)\right)} \\
& =\frac{1}{\Phi_{\sigma}\left(u^{*}-v_{s}\right)\left(1-\Phi_{\sigma}\left(u^{*}-v_{s}\right)\right)} \sum_{y_{s-1} \in\{0,1\}^{s-1}} Q_{s-1}^{V^{s-1}, 1}\left(y_{s-1}\right)\left(\Phi_{1}\left(u^{*}-v_{s}\right)-\Phi_{\sigma}\left(u^{*}-v_{s}\right)\right)^{2} \\
& \underset{\substack{\uparrow \\
(* *)}}{\leq} \quad 165 \cdot \sum_{y_{s-1} \in\{0,1\}^{s-1}} Q_{s-1}^{V_{s-1}^{s-1}, 1}\left(y_{s-1}\right)\left(\Phi_{1}\left(u^{*}-v_{s}\right)-\Phi_{\sigma}\left(u^{*}-v_{s}\right)\right)^{2} \\
& \underset{\substack{\uparrow \\
\text { Lemma } 6 \text { Property } 3}}{\leq} 165 \cdot \sum_{y_{s-1} \in\{0,1\}^{s-1}} Q_{s-1}^{V_{s-1}^{s-1}\left(y_{s-1}\right)\left(u^{*}-v_{s}\right)^{2}(1-\sigma)^{2}} \\
& =165(1-\sigma)^{2} \mathbb{E}_{Y_{s-1}}\left[\left(u^{*}-v_{s}\right)^{2}\right] .
\end{aligned}
$$

Here inequality $\left({ }^{* *}\right)$ above is proved as follows: since $v_{s} \in\left(0, u^{*}\right)$ as is assumed, we have:

$$
\begin{aligned}
\frac{1}{2}<\Phi_{\sigma}\left(u^{*}-v_{s}\right) & <\Phi_{\sigma}\left(u^{*}\right) \\
& =\sigma \cdot \Phi_{1}\left(\frac{u^{*}}{\sigma}\right) \\
& \leq 1 \cdot \Phi_{1}\left(\frac{\sqrt{\frac{\pi}{2}}}{\frac{1}{2}}\right) \\
& \leq \Phi_{1}(\sqrt{2 \pi}) \\
& \leq 0.9939
\end{aligned}
$$

As a result, we have $\frac{1}{\Phi_{\sigma}\left(u^{*}-v_{s}\right)\left(1-\Phi_{\sigma}\left(u^{*}-v_{s}\right)\right)} \leq \frac{1}{0.9939 \times 0.0061}=164.7988 \leq 165$. Therefore, by summing up all $s$, we have:

$$
\begin{aligned}
& \mathcal{K}\left(Q_{T}^{V, 1} ; Q_{T}^{V, \sigma}\right)=\sum_{s=1}^{T} \mathcal{K}\left(Q_{s}^{V^{s}, 1} ; Q_{s}^{V^{s}, \sigma} \mid Y_{s-1}\right) \\
& \leq 165(1-\sigma)^{2} \sum_{s=1}^{T} \mathbb{E}_{Y_{s-1}}\left[\left(u^{*}-v_{s}\right)^{2}\right] \\
& \leq 165 \times 60 \cdot(1-\sigma)^{2} \sum_{s=1}^{T}\left(r\left(u^{*}, 1\right)-r\left(v_{s}, 1\right)\right) \\
& \text { Lemma 16 Property 1 } \\
&=9900(1-\sigma)^{2} \operatorname{Regret}(1, T, \Psi) \\
& \text { definition of regret and } v_{s}=\Psi\left(y_{s-1}\right) .
\end{aligned}
$$

which concludes the proof.

Theorem 19 (Uncertainty is costly). Let $\sigma \leq 1-T^{-\frac{1}{4}}$, and we have:

$$
\begin{equation*}
\operatorname{Regret}(1, T, \Psi)+\operatorname{Regret}(\sigma, T, \Psi) \geq \frac{1}{24000} \cdot \sqrt{T} \cdot e^{-\mathcal{K}\left(Q^{V, 1} ; Q^{V, \sigma}\right)} \tag{37}
\end{equation*}
$$

Here $v_{t}=\Psi\left(y_{t-1}\right), t=1,2, \ldots, T$.

Proof. First of all, we cite a lemma that would facilitate our proof:

Lemma 20. Let $Q_{0}$ and $Q_{1}$ be two probability distributions on a finite space $\mathcal{Y}$; with $Q_{0}(y), Q_{1}(y)>$ $0, \forall y \in \mathcal{Y}$. Then for any function $F: \mathcal{Y} \rightarrow\{0,1\}$,

$$
Q_{0}\{F=1\}+Q_{1}\{J=0\} \geq \frac{1}{2} e^{-\mathcal{K}\left(Q_{0} ; Q_{1}\right)}
$$

where $\mathcal{K}\left(Q_{0} ; Q_{1}\right)$ denotes the $K L$-divergence of $Q_{0}$ and $Q_{1}$.

Define two intervals of prices:

$$
C_{1}=\left\{v:\left|u^{*}\right| \leq \frac{1}{10 T^{\frac{1}{4}}}\right\} \text { and } C_{2}=\left\{v:\left|J_{\sigma}\left(u^{*}\right)-v\right| \leq \frac{1}{10 T^{\frac{1}{4}}}\right\}
$$

Note that $C_{1}$ and $C_{2}$ are disjoint, since $\left|u^{*}-J_{\sigma}\left(u^{*}\right)\right| \geq \frac{2}{5}|1-\sigma|=\frac{2}{5 T^{1 / 2}}$ according to Lemma 16 Property 2. Also, for $v \in\left(0, u^{*}\right) \backslash C_{2}$, the regret is large according to Lemma 16 Property 1 , because:

$$
r\left(v^{*}(\sigma), \sigma\right)-r(v, \sigma) \geq \frac{1}{60}\left(v-v^{*}(\sigma)\right)^{2} \geq \frac{1}{6000 T^{\frac{1}{2}}}
$$

Then, we have:

$$
\begin{aligned}
& \operatorname{Regret}(1, T, \Psi)+\operatorname{Regret}(\sigma, T, \Psi) \\
& \geq \sum_{t=1}^{T-1} \mathbb{E}_{1}\left[r\left(u^{*}, 1\right)-r\left(v_{t+1}, 1\right)\right]+\mathbb{E}_{\sigma}\left[r\left(J_{\sigma}\left(u^{*}\right), \sigma\right)-r\left(v_{t+1}, \sigma\right)\right] \\
& \geq \frac{1}{6000 \sqrt{T}} \sum_{t=1}^{T-1} \mathbb{P}_{1}\left[v_{t+1} \notin C_{1}\right]+\mathbb{P}_{\sigma}\left[v_{t+1} \notin\left\{C_{2}\right\}\right] \\
& \underset{\substack{\uparrow \\
t+1 \\
=\mathbb{1}\left[v_{t+1} \in C_{2}\right.}}{\geq} \frac{1}{6000 \sqrt{T}} \sum^{t=1} \mathbb{P}_{1}\left[F_{t+1}=1\right]+\mathbb{P}_{\sigma}\left[F_{t+1}=0\right] \\
& \underset{\substack{\uparrow \\
\uparrow \operatorname{man} \\
20}}{\geq} \frac{1}{6000 \sqrt{T}} \sum_{t=1}^{T-1} \frac{1}{2} e^{-\mathcal{K}\left(Q_{t}^{V^{t}, 1} ; Q_{t}^{V^{t}, \sigma}\right)} \\
& \text { Lemma } 20 \\
& \geq \frac{1}{6000 \sqrt{T}} \frac{T-1}{2} e^{-\mathcal{K}\left(Q_{T}^{V, 1} ; Q_{T}^{V, \sigma}\right)} \\
& \mathcal{K}\left(Q_{t}^{V^{t}, 1} ; Q_{t}^{V^{t}, \sigma}\right) \text { not decreasing } \\
& \geq \frac{1}{24000} \sqrt{T} e^{-\mathcal{K}\left(Q_{T}^{V, 1} ; Q_{T}^{V, \sigma}\right)} .
\end{aligned}
$$

According to Theorem 17 and Theorem 19, we can then prove Theorem 12. Let $\sigma=1-T^{-\frac{1}{4}}$

$$
\begin{aligned}
& 2(\operatorname{Regret}(1, T, \Psi)+\operatorname{Regret}(\sigma, T, \Psi)) \\
& \geq \operatorname{Regret}(1, T, \Psi)+(\operatorname{Regret}(1, T, \Psi)+\operatorname{Regret}(\sigma, T, \Psi)) \\
& \geq \frac{1}{9900 T^{-1 / 2}} \mathcal{K}\left(Q^{V, 1} ; Q^{V, \sigma}\right)+\frac{1}{24000} \cdot \sqrt{T} \cdot e^{-\mathcal{K}\left(Q^{V, 1} ; Q^{V, \sigma}\right)} \\
& \geq \frac{1}{24000} \sqrt{T}\left(\mathcal{K}\left(Q^{V, 1} ; Q^{V, \sigma}\right)+e^{-\mathcal{K}\left(Q^{V, 1} ; Q^{V, \sigma}\right)}\right) \\
& \geq \frac{1}{24000} \sqrt{T} . \\
& \text { The fact } e^{x} \geq x+1, \forall x \in \mathbb{R}
\end{aligned}
$$

Thus Theorem 12 is proved valid.

## C More Discussions

## C. 1 Dependence on $B$ and Noise Variance

Here we use a concrete example to analyze the coefficients of regret bounds. Again, we assume that $N_{t} \sim \mathcal{N}\left(0, \sigma^{2}\right)$. Notice that both $C_{s}$ and $C_{a}$ have a component of $\frac{C_{\text {exp }}}{C_{\text {down }}}$. In order to analyze $\frac{C_{\text {exp }}}{C_{\text {down }}}$, we define a hazard function denoted as $\lambda(\omega)$ with $\omega \in \mathbb{R}$ :

$$
\begin{equation*}
\lambda(\omega):=\frac{p_{1}(\omega)}{1-\Phi_{1}(\omega)}=\frac{p_{1}(-\omega)}{\Phi_{1}(-\omega)} \tag{38}
\end{equation*}
$$

where $\Phi_{1}$ and $p_{1}$ are the CDF and PDF of standard Gaussian distribution. The concept of hazard function comes from the area of survival analysis. From Equation 11 and 13, we plug in Equation 38 and get:

$$
\begin{align*}
C_{\text {down }} & \geq \inf _{\omega \in\left[-\frac{B}{\sigma}, \frac{B}{\sigma}\right]}\left\{\frac{1}{\sigma^{2}} \lambda(-\omega)^{2}+\omega \cdot \lambda(-\omega)\right\} \\
C_{\exp } & \leq \sup _{\omega \in\left[-\frac{B}{\sigma}, \frac{B}{\sigma}\right]}\left\{\frac{1}{\sigma^{2}} \lambda(-\omega)^{2}\right\} . \tag{39}
\end{align*}
$$

In Lemma 21, we will prove that $\lambda(\omega)$ is exponentially small as $\omega \rightarrow+\infty$, and is asymptotically close to $-\omega$ as $\omega \rightarrow-\infty$. Therefore, $C_{d o w n}$ is exponentially small and $C_{\text {exp }}$ is quadratically large with respect to $B / \sigma$. Although we assume that $B$ and $\sigma$ are constant, we should be alert that the scale of $B / \sigma$ can be very large as $\sigma$ goes to zero, i.e. as the noise is "insignificant". In practice (especially when $T$ is finite), this may cause extremely large regret at the beginning. A "Shallow Pricing" method introduced by Cohen et al. [2020] (as well as other domain-cutting methods in contextual searching) may serve as a good pre-process as it frequently conducts bisections to cut the feasible region of $\theta^{*}$ with high probability. According to Theorem 3 in Cohen et al. [2020], their Shallow Pricing algorithm will bisect the parameter set for at most logarithmic times to ensure that $\frac{B}{\sigma}$ has been small enough (i.e. upper-bounded by $O($ poly $\log (T))$ ). However, this does not necessarily means that we can use a $O(\log T)$-time pre-process to achieve the same effect, since they run the algorithm throughout the session while we only take it as a pre-process. Intuitively, at least under the adversarial feature assumption, we cannot totally rely on a few features occurring at the beginning (as they might be misleading) to cut the parameter set once and for all. A mixture approach of Shallow Pricing and EMLP/ONSP might work, as the algorithm can detect whether current $\frac{B}{\sigma}$ is larger than a threshold of bisection. However, this requires new regret analysis as the operations parameter domain are changing over time. Therefore, we claim in Section 7 that the regret bound is still open if $\sigma=\Theta\left(T^{-\alpha}\right)$ for $\alpha \in(0,1)$.

Lemma 21 (Properties of $\lambda(\omega)$ ). For $\lambda(\omega):=\frac{p_{1}(\omega)}{1-\Phi_{1}(\omega)}$, we have:

$$
\begin{array}{ll}
1, & \frac{d}{d \omega} \lambda(\omega)>0 . \\
2, & \lim _{\omega \rightarrow-\infty} \omega^{k} \lambda(\omega)=0, \forall k>0 . \\
3, & \lim _{\omega \rightarrow+\infty} \lambda(\omega)-\omega=0 . \\
4, & \lim _{\omega \rightarrow+\infty} \omega(\lambda(\omega)-\omega)=1 .
\end{array}
$$

Proof. We prove the Lemma 21 sequentially:

1. We have:

$$
\begin{align*}
\lambda^{\prime}(\omega) & =\frac{p_{1}^{2}(-\omega)-p_{1}^{\prime}(-\omega) \Phi_{1}(-\omega)}{\Phi_{1}(-\omega)^{2}} \\
& =\frac{p_{1}^{2}(-\omega)-\omega p_{1}(-\omega) \Phi_{1}(-\omega)}{\Phi_{1}(-\omega)^{2}}  \tag{40}\\
& =\frac{p_{1}(-\omega)\left(p_{1}(-\omega)-\omega \Phi_{1}(-\omega)\right)}{\Phi_{1}(-\omega)^{2}} .
\end{align*}
$$

Therefore, it is equivalent to prove that $p_{1}(-\omega)-\omega \Phi_{1}(-\omega)>0$.

Suppose $f(\omega)=p_{1}(\omega)+\omega \Phi_{1}(\omega)$. We now take its derivatives as follows:

$$
\begin{align*}
f^{\prime}(\omega) & =p_{1}^{\prime}(\omega)+\left(\Phi_{1}(\omega)+\omega \cdot p_{1}(\omega)\right) \\
& =(-\omega) p_{1}(\omega)+\Phi_{1}(\omega)+\omega \cdot p_{1}(\omega) \\
& =\Phi_{1}(\omega)  \tag{41}\\
& >0
\end{align*}
$$

Therefore, we know that $f(\omega)$ monotonically increases in $\mathbb{R}$. Additionally, since we have:

$$
\begin{align*}
& \lim _{\omega \rightarrow-\infty} f(\omega) \\
= & \lim _{\omega \rightarrow-\infty} p_{1}(\omega)+\lim _{\omega \rightarrow-\infty} \omega \Phi_{1}(\omega) \\
= & 0+\lim _{\omega \rightarrow-\infty} \frac{1}{\sigma^{2}} \cdot \frac{\Phi_{1}(\omega)}{1 / \omega} \\
= & \lim _{\omega \rightarrow-\infty} \cdot \frac{p_{1}(\omega)}{-1 / \omega^{2}}  \tag{42}\\
= & \lim _{\omega \rightarrow-\infty} \cdot\left(-\frac{1}{\sqrt{2 \pi}} \cdot \frac{\omega^{2}}{\exp \left\{\frac{\omega^{2}}{2}\right\}}\right) \\
= & 0
\end{align*}
$$

Therefore, we know that $f(\omega)>0, \forall \omega \in \mathbb{R}$, and as a result, $\lambda^{\prime}(\omega)>0$.
2. We have:

$$
\begin{aligned}
& \lim _{\omega \rightarrow-\infty} \omega^{k} \lambda(\omega) \\
&= \lim _{\omega \rightarrow-\infty} \omega^{k} \frac{p_{1}(-\omega)}{\Phi_{1}(-\omega)} \\
&= \frac{\omega \rightarrow-\infty}{\lim _{\omega \rightarrow-\infty} \omega^{k} p_{1}(-\omega)} \Phi_{1}(-\omega) \\
&= \frac{\omega \rightarrow-\infty}{} \omega^{k}\left(\frac{1}{\sqrt{2 \pi}} \exp \left\{-\frac{\omega^{2}}{2}\right\}\right) \\
& 1 \\
&= \frac{0}{1} \\
&= 0 .
\end{aligned}
$$

3. We only need to prove that

$$
\lim _{\omega \rightarrow+\infty} \lambda(\omega)-\omega=0
$$

Actually, we have:

$$
\begin{align*}
& \lim _{\omega \rightarrow+\infty} \lambda(\omega)-\omega \\
&= \lim _{\omega \rightarrow+\infty} \frac{p_{1}(-\omega)-\omega \Phi_{1}(-\omega)}{\Phi_{1}(-\omega)} \\
&= \lim _{\omega \rightarrow-\infty} \frac{p_{1}(\omega)+\omega \Phi_{1}(\omega)}{\Phi_{1}(\omega)} \\
& \stackrel{\text { LHospital's rule }}{\stackrel{\perp}{=}} \lim _{\omega \rightarrow-\infty} \frac{(-\omega) p_{1}(\omega)+\Phi_{1}(\omega)+\omega p_{1}(\omega)}{p_{1}(\omega)}  \tag{44}\\
&= \lim _{\omega \rightarrow-\infty} \frac{\Phi_{1}(\omega)}{p_{1}(\omega)} \\
&= \lim _{\omega \rightarrow-\infty} \frac{p_{1}(\omega)}{(-\omega) p_{1}(\omega)} \\
&= 0
\end{align*}
$$

4. 

$$
\begin{align*}
& \lim _{\omega \rightarrow+\infty} \omega(\lambda(\omega)-\omega) \\
&= \lim _{\omega \rightarrow+\infty} \frac{\omega\left(p_{1}(-\omega)-\omega \Phi_{1}(-\omega)\right)}{\Phi_{1}(-\omega)} \\
&= \lim _{\omega \rightarrow-\infty} \frac{-\omega p_{1}(\omega)-\omega^{2} \Phi_{1}(\omega)}{\Phi_{1}(\omega)} \\
& \stackrel{\rightharpoonup}{\text { LHospital's rule }} \stackrel{=}{\omega} \lim _{\omega \rightarrow-\infty} \frac{-p_{1}(\omega)-\omega(-\omega) p_{1}(\omega)-\omega^{2} p_{1}(\omega)-2 \omega \cdot \Phi_{1}(\omega)}{p_{1}(\omega)}  \tag{45}\\
&=-1-2 \lim _{\omega \rightarrow-\infty} \frac{\omega \Phi_{1}(\omega)}{p_{1}(\omega)} \\
&=-1+2 \lim _{\omega \rightarrow+\infty} \frac{1}{\frac{\lambda(\omega)}{\omega}} \\
&=-1+2 \\
&= 1 .
\end{align*}
$$

Thus the lemma holds.

## C. 2 Algorithmic Design

## C.2.1 Probit and Logistic Regressions

A probit/logit model is described as follows: a Boolean random variable $Y$ satisfies the following probabilistic distribution: $\mathbb{P}[Y=1 \mid X]=F\left(X^{\top} \beta\right)$, where $X \in \mathbb{R}$ is a random vector, $\beta \in \mathbb{R}$ is a parameter, and $F$ is the cumulative distribution function (CDF) of a (standard) Gaussian/logistic distribution. In our problem, we may treat $\mathbb{1}_{t}$ as $Y,\left[x_{t}{ }^{\top}, v_{t}\right]^{\top}$ as $X$ and $\left[\theta^{* \top},-1\right]^{\top}$ as $\beta$, which exactly fits this model if we assume the noise as Gaussian or logistic. Therefore, $\hat{\theta}_{k}=\arg \min _{\theta} \hat{L}_{k}(\theta)$ can be solved via the highly efficient implementation of generalized linear models, e.g., GLMnet, rather than resorting to generic tools for convex programming. As a heuristic, we could leverage the vast body of statistical work on probit or logit models and adopt a fully Bayesian approach that jointly estimates $\theta$ and hyper-parameters of $F$. This would make the algorithm more practical by eliminating the need to choose the hyper-parameters when running this algorithm.

## C.2.2 Advantages of EMLP over ONSP.

For the stochastic setting, we specifically propose EMLP even though ONSP also works. This is because EMLP only "switch" the pricing policy $\hat{\theta}$ for $\log T$ times. This makes it appealing in many applications (especially for brick-and-mortar sales) where the number of policy updates is a bottleneck. In fact, the iterations within one epoch can be carried out entirely in parallel.

## C.2.3 Agnostic Dynamic Pricing: Explorations versus Exploitation

At the moment, the proposed algorithm relies on the assumption of a linear valuation function (see Appendix C. 3 for more discussion on problem modeling). It will be interesting to investigate the settings of model-misspecified cases and the full agnostic settings. The key would be to exploit the structural feedback in model-free policy-evaluation methods such as importance sampling. The main reason why we do not explore lies in the noisy model: essentially we are implicitly exploring a higher (permitted) price using the naturally occurring noise in the data. In comparison, there is another problem setting named "adversarial irrationality" where some of the customers will valuate the product adaptively and adversarially ${ }^{10}$. Existing work Krishnamurthy et al. [2021] adopts this setting and shows a linear regret dependence on the number of irrational customers, but they consider a different loss function (See Related Works Section).

[^1]
## C. 3 Problem Modeling

## C.3.1 Noise Distributions

In this work, we have made four assumptions on the noise distribution: strict log-concavity, $2^{\text {nd }}-$ order smooth, known, and i.i.d.. Here we explain each of them specifically.

- The assumption of knowing the exact $F$ is critical to the regret bound: If we have this knowledge, then we achieve $O(\log T)$ even with adversarial features; otherwise, an $\Omega(\sqrt{T})$ regret is unavoidable even with stochastic features.
- The strictly log-concave distribution family includes Gaussian and logistic distributions as two common noises. In comparison, Javanmard and Nazerzadeh [2019] assumes log-concavity that further covers Laplacian, exponential and uniform distributions. Javanmard and Nazerzadeh [2019] also considers the cases when (1) the noise distribution is unknown but log-concave, and (2) the noise distribution is zero-mean and bounded by support of $[-\delta, \delta]$. For case (1), they propose an algorithm with regret $O(\sqrt{T})$ and meanwhile prove the same lower bound. For case (2), they propose an algorithm with linear regret.
- The assumption that $F$ is $2^{\text {nd }}$-order smooth is also assumed by Javanmard and Nazerzadeh [2019] by taking derivatives $f^{\prime}(v)$ and applying its upper bound in the proof. Therefore, we are still unaware of the regret bound if the noise distribution is discrete, where a lower bound of $\Omega(\sqrt{T})$ can be directly applied from Kleinberg and Leighton [2003].
- We even assume that the noise is identically distributed. However, the noise would vary among different people. The same problem happens on the parameter $\theta^{*}$ : can we assume different people sharing the same evaluation parameter? We may interpret it in the following two ways, but there are still flaws: (1) the "customer" can be the public, i.e. their performance is quite stable in general; or (2) the customer can be the same one over the whole time series. However, the former explanation cannot match the assumption that we just sell one product at each time, and the latter one would definitely undermine the independent assumption of the noise: people would do "human learning" and might gradually reduce their noise of making decisions. To this extent, it is closer to the fact if we assume noises as martingales. This assumption has been stated in Qiang and Bayati [2016].


## C.3.2 Linear Valuations on Features

There exist many products whose prices are not linearly dependent on features. One famous instance is a diamond: a kilogram of diamond powder is very cheap because it can be produced artificially, but a single 5-carat (or 1 gram) diamond might cost more than $\$ 100,000$. This is because of an intrinsic non-linear property of diamond: large ones are rare and cannot be (at least easily) compound from smaller ones. Another example lies in electricity pricing [Joskow and Wolfram, 2012], where the more you consume, the higher unit price you suffer. On the contrary, commodities tend to be cheaper than retail prices. These are both consequences of marginal costs: a large volume consuming of electricity may cause extra maintenance and increase the cost, and a large amount of purchasing would release the storage and thus reduce their costs. In a word, our problem setting might not be suitable for those large-enough features, and thus an upper bound of $x^{\top} \theta$ becomes a necessity.

## C. 4 Ex Ante v.s. Ex Post Regrets

In this work, we considered the ex ante regret $R e g_{e a}=\sum_{t=1}^{T} \max _{\theta} \mathbb{E}\left[v_{t}^{\theta} \cdot \mathbb{1}\left(v_{t}^{\theta} \leq w_{t}\right)\right]-\mathbb{E}\left[v_{t} \cdot \mathbb{1}\left(v_{t} \leq\right.\right.$ $\left.\left.w_{t}\right)\right]$, where $v_{t}^{\theta}=J\left(x_{t}^{\top} \theta\right)$ is the greedy price with parameter $\theta$ and $w_{t}=x_{t}^{\top} \theta^{*}+N_{t}$ is the realized random valuation. The ex post definition of the cumulative regret, i.e., $R e g_{e p}=$ $\max _{\theta} \sum_{t=1}^{T} v_{t}^{\theta} \mathbb{1}\left(v_{t}^{\theta} \leq w_{t}\right)-v_{t} \mathbb{1}\left(v_{t} \leq w_{t}\right)$ makes sense, too. Note that we can decompose $\mathbb{E}\left[\operatorname{Reg}_{e p}\right]=R e g_{e a}+\mathbb{E}\left[\max _{\theta} \sum_{t=1}^{T} v_{t}^{\theta} \mathbb{1}\left(v_{t}^{\theta} \leq w_{t}\right)-\sum_{t=1}^{T} v_{t}^{\theta^{*}} \mathbb{1}\left(v_{t}^{\theta^{*}} \leq w_{t}\right)\right]$. While it might be the case that the second term is $\Omega(\sqrt{d T})$ as the reviewer pointed out, it is a constant independent of the algorithm. For this reason, we believe using $R e g_{e a}$ is without loss of generality, and it reveals more nuanced performance differences of different algorithms.

For an ex post dynamic regret, i.e., $R e g_{d}=\sum_{t=1}^{T} w_{t}-v_{t} \cdot \mathbb{1}\left(v_{t} \leq w_{t}\right)$, it is argued in Cohen et al. [2020] that any policy must suffer an expected regret of $\Omega(T)$ (even if $\theta^{*}$ is known). We may also present a good example lies in $N_{t} \sim \mathcal{N}(0,1), x_{t}^{\top} \theta^{*}=\sqrt{\frac{\pi}{2}}$ where the optimal price is $\sqrt{\frac{\pi}{2}}$
as well but the probability of acceptance is only $1 / 2$, and this leads to a constant per-step regret of $\frac{1}{2} \sqrt{\frac{\pi}{2}}$.

## C. 5 Ethic Issues

A field of study lies in "personalized dynamic pricing" [Aydin and Ziya, 2009, Chen and Gallego, 2021], where a firm makes use of information of individual customers and sets a unique price for each of them. This has been frequently applied in airline pricing [Krämer et al., 2018]. However, this causes first-order pricing discrimination. Even though this "discrimination" is not necessarily immoral, it must be embarrassing if we are witted proposing the same product with different prices towards different customers. For example, if we know the coming customer is rich enough and is not as sensitive towards a price (e.g., he/she has a variance larger than other customers), then we are probably raising the price without being too risky. Or if the customer is used to purchase goods from ours, then he or she might have a higher expectation on our products (e.g., he/she has a $\theta=a \theta^{*}, a>1$ ), and we might take advantage and propose a higher price than others. These cases would not happen in an auction-based situation (such as a live sale), but might frequently happen in a more secret place, for instance, a customized travel plan.

## References

A. Agarwal, D. Hsu, S. Kale, J. Langford, L. Li, and R. Schapire. Taming the monster: A fast and simple algorithm for contextual bandits. In International Conference on Machine Learning (ICML-14), pages 1638-1646, 2014.
J. H. Aldrich, F. D. Nelson, and E. S. Adler. Linear Probability, Logit, and Probit Models. Number 45. Sage, 1984.
K. Amin, A. Rostamizadeh, and U. Syed. Repeated contextual auctions with strategic buyers. In Advances in Neural Information Processing Systems (NIPS-14), pages 622-630, 2014.
V. F. Araman and R. Caldentey. Dynamic pricing for nonperishable products with demand learning. Operations Research, 57(5):1169-1188, 2009.
P. Auer, N. Cesa-Bianchi, Y. Freund, and R. E. Schapire. The nonstochastic multiarmed bandit problem. SIAM Journal on Computing, 32(1):48-77, 2002.
G. Aydin and S. Ziya. Personalized dynamic pricing of limited inventories. Operations Research, 57 (6):1523-1531, 2009.
D. Besanko, S. Gupta, and D. Jain. Logit demand estimation under competitive pricing behavior: An equilibrium framework. Management Science, 44(11-part-1):1533-1547, 1998.
D. Besanko, J.-P. Dubé, and S. Gupta. Competitive price discrimination strategies in a vertical channel using aggregate retail data. Management Science, 49(9):1121-1138, 2003.
O. Besbes and A. Zeevi. On the (surprising) sufficiency of linear models for dynamic pricing with demand learning. Management Science, 61(4):723-739, 2015.
J. Broder and P. Rusmevichientong. Dynamic pricing under a general parametric choice model. Operations Research, 60(4):965-980, 2012.
T. Chan, V. Kadiyali, and P. Xiao. Structural models of pricing. Handbook of pricing research in marketing, pages 108-131, 2009.
N. Chen and G. Gallego. A primal-dual learning algorithm for personalized dynamic pricing with an inventory constraint. Mathematics of Operations Research, 2021.
Y. Chen and V. F. Farias. Simple policies for dynamic pricing with imperfect forecasts. Operations Research, 61(3):612-624, 2013.
W. Chu, L. Li, L. Reyzin, and R. Schapire. Contextual bandits with linear payoff functions. In International Conference on Artificial Intelligence and Statistics (AISTATS-11), pages 208-214, 2011.
M. C. Cohen, I. Lobel, and R. Paes Leme. Feature-based dynamic pricing. Management Science, 66 (11):4921-4943, 2020.
A. A. Cournot. Researches into the Mathematical Principles of the Theory of Wealth. Macmillan, 1897.
A. V. den Boer. Dynamic pricing and learning: historical origins, current research, and new directions. Surveys in Operations Research and Management Science, 20(1):1-18, 2015.
M. Draganska and D. C. Jain. Consumer preferences and product-line pricing strategies: An empirical analysis. Marketing science, 25(2):164-174, 2006.
G. C. Evans. The dynamics of monopoly. The American Mathematical Monthly, 31(2):77-83, 1924.
E. Hazan. Introduction to online convex optimization. Foundations and Trends in Optimization, 2 (3-4):157-325, 2016.
R. Iyengar, A. Ansari, and S. Gupta. A model of consumer learning for service quality and usage. Journal of Marketing Research, 44(4):529-544, 2007.
A. Javanmard and H. Nazerzadeh. Dynamic pricing in high-dimensions. The Journal of Machine Learning Research, 20(1):315-363, 2019.
P. L. Joskow and C. D. Wolfram. Dynamic pricing of electricity. American Economic Review, 102 (3):381-85, 2012.
V. Kadiyali, N. J. Vilcassim, and P. K. Chintagunta. Empirical analysis of competitive product line pricing decisions: Lead, follow, or move together? Journal of Business, pages 459-487, 1996.
N. B. Keskin and A. Zeevi. Dynamic pricing with an unknown demand model: Asymptotically optimal semi-myopic policies. Operations Research, 62(5):1142-1167, 2014.
W. Kincaid and D. Darling. An inventory pricing problem. Journal of Mathematical Analysis and Applications, 7:183-208, 1963.
R. Kleinberg and T. Leighton. The value of knowing a demand curve: Bounds on regret for online posted-price auctions. In IEEE Symposium on Foundations of Computer Science (FOCS-03), pages 594-605. IEEE, 2003.
T. Koren and K. Levy. Fast rates for exp-concave empirical risk minimization. In Advances in Neural Information Processing Systems (NIPS-15), pages 1477-1485, 2015.
A. Krämer, M. Friesen, and T. Shelton. Are airline passengers ready for personalized dynamic pricing? a study of german consumers. Journal of Revenue and Pricing Management, 17(2): 115-120, 2018.
A. Krishnamurthy, T. Lykouris, C. Podimata, and R. Schapire. Contextual search in the presence of irrational agents. In Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing (STOC-21), pages 910-918, 2021.
A. Lambrecht, K. Seim, and B. Skiera. Does uncertainty matter? consumer behavior under three-part tariffs. Marketing Science, 26(5):698-710, 2007.
J. Langford and T. Zhang. The epoch-greedy algorithm for contextual multi-armed bandits. In Advances in Neural Information Processing Systems (NIPS-07), pages 817-824, 2007.
R. P. Leme and J. Schneider. Contextual search via intrinsic volumes. In 2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS-18), pages 268-282. IEEE, 2018.
A. Liu, R. P. Leme, and J. Schneider. Optimal contextual pricing and extensions. In Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA-21), pages 1059-1078. SIAM, 2021.
I. Lobel, R. P. Leme, and A. Vladu. Multidimensional binary search for contextual decision-making. Operations Research, 66(5):1346-1361, 2018.
T. Mazumdar, S. P. Raj, and I. Sinha. Reference price research: Review and propositions. Journal of Marketing, 69(4):84-102, 2005.
J. Mourtada. Exact minimax risk for linear least squares, and the lower tail of sample covariance matrices. arXiv preprint arXiv:1912.10754, 2019.
K. P. Murphy. Machine Learning: a Probabilistic Perspective. MIT press, 2012.
S. Qiang and M. Bayati. Dynamic pricing with demand covariates. arXiv preprint arXiv:1604.07463, 2016.
H. Schultz et al. Theory and Measurement of Demand. The University of Chicago Press, 1938.
C.-H. Wang, Z. Wang, W. W. Sun, and G. Cheng. Online regularization for high-dimensional dynamic pricing algorithms. arXiv preprint arXiv:2007.02470, 2020.
P. Whittle. Multi-armed bandits and the gittins index. Journal of the Royal Statistical Society: Series B (Methodological), 42(2):143-149, 1980.
R. E. Wright. Logistic regression. 1995.


[^0]:    ${ }^{9}$ But do notice that our expected reward above is not linear, even if the valuation function is.

[^1]:    ${ }^{10}$ An adaptive adversary may take actions adversarially in respond to the environmental changes. In comparison, what we allow for the "adversarial features" is actually chosen by an oblivious adversary before the interactions start.

